

‘You Draw It’: Implementation of Visually Fitted Trends with r2d3

EMILY A. ROBINSON^{1,*}, REKA HOWARD², AND SUSAN VANDERPLAS²

¹*Department of Statistics, CalPoly-San Luis Obispo, USA*

²*Department of Statistics, University of Nebraska-Lincoln, USA*

Abstract

How do statistical regression results compare to intuitive, visually fitted results? Fitting lines by eye through a set of points has been explored since the 20th century. Common methods of fitting trends by eye involve maneuvering a string, black thread, or ruler until the fit is suitable, then drawing the line through the set of points. In 2015, the New York Times introduced an interactive feature, called ‘You Draw It,’ where readers are asked to input their own assumptions about various metrics and compare how these assumptions relate to reality. This research is intended to implement ‘You Draw It’, adapted from the New York Times, as a way to measure the patterns we see in data. In this paper, we describe the adaptation of an old tool for graphical testing and evaluation, eye-fitting, for use in modern web-applications suitable for testing statistical graphics. We present an empirical evaluation of this testing method for linear regression, and briefly discuss an extension of this method to non-linear applications.

1 Introduction

We all use statistical graphics, but how do we know that the graphics we use are communicating effectively? Through experimentation, graphical testing methods allow researchers to conduct studies designed to understand how we perceive graphics and perform graphical tasks such as differentiation, prediction, estimation, and extrapolation. Each of these levels of interaction with a graph require a different method of engagement with and use of the information presented in a chart. In this paper, we describe the adaptation of an old tool for graphical testing and evaluation, eye-fitting, for use in modern web-applications suitable for testing statistical graphics. We present an empirical evaluation of this testing method for linear regression, and briefly discuss an extension of this method to non-linear applications.

One of the most common charts created is a scatterplot of points over time; these charts show up regularly in news articles and in scientific publications alike. These charts rely on our ability to identify and detect trends in data. Our visual system is naturally built to look for structure and identify patterns, including patterns and trends over time; many times we do not even notice this process happening subconsciously. As shown in Figure 1, a viewer engaging with a plot of weekly average gas prices over time may perform several cognitive operations. First, they scan the plot and assess points to determine if there are any outliers or otherwise remarkable points. Then, they may fit a rough mental smooth/trend line to the points to summarize the useful information and remove variability. Finally, an interested and engaged viewer may pull in additional contextual information from long-term memory, seeking to explain variation in

*Corresponding author. Email: erobin17@calpoly.edu.

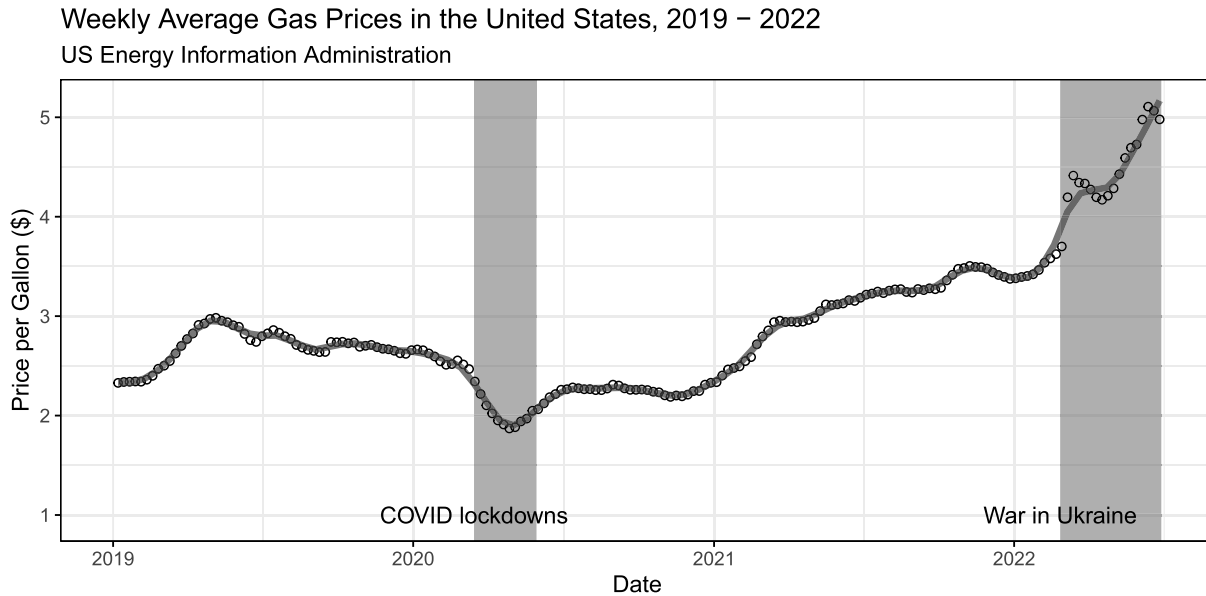


Figure 1: Demonstration of a viewer’s engagement with a scatterplot showing the weekly average gas prices in the United States, from 2019–June 2022. The viewer might perform several cognitive operations such as assessing points to determine outliers, fitting a rough mental smooth/trend line (overlying solid gray line), and using additional contextual information from long term memory to explain the variation (shaded rectangular gray regions annotated in black text) (US Energy Information Administration, 2022).

the mental ‘trend’ with supplemental information, such as COVID lock downs which reduced gasoline demand and the beginning of the war in Ukraine, which wreaked havoc on the supply and demand for global energy sources.

Initial studies in the 20th century explored the use of fitting lines by eye through a set of points (Finney 1951; Mosteller et al. 1981; Unwin and Wills 1988). Common methods of fitting trends by eye involved maneuvering a string, black thread, or ruler until the fit was suitable, then drawing the line through the set of points. Results from these early studies provided ground-work for visually judging slopes and selecting accurate starting values for common iterative calculations. Recently, Ciccione and Dehaene (2021) conducted a comprehensive set of studies investigating human ability to detect trends in graphical representations from a psychophysical approach. This set of studies asked participants to judge trends, estimate slopes, and conduct extrapolation by using a track-pad to adjust the tilt of a line on the screen. Results indicated the slopes participants reported were always in excess of the ideal slopes, both in the positive and in the negative direction, and those biases increased with noise and with number of points. This supported the results found in Mosteller et al. (1981) and suggested that participants might use Deming regression (Deming, 1943), which minimizes the Euclidean distance of points from the line, when fitting a line to a noisy scatterplot.

While psychologists and statisticians have been using eye-fitting techniques to assess our innate perceptual statistical modeling abilities, news organizations have used similar techniques in order to draw readers in and demonstrate the difference between readers’ assumptions and reality. In 2015, the New York Times (NYT) introduced an interactive feature, called ‘You Draw It’ (Aisch et al., 2015), where readers input their own assumptions about various metrics of news

interest and compare these assumptions to reality. The NYT team used Data Driven Documents (D3), a JavaScript library that allows readers to interact with a chart directly by drawing a line on their computer screen with a mouse. Despite the somewhat different purpose behind this feature, the D3 driven method used by the NYT is wonderfully intuitive, and does not require the assumption of linearity, making it much more adaptable to testing how viewers perceive and predict when presented with non-linear data.

We set out to implement ‘You Draw It’, adapted from the NYT feature, as a way to experimentally assess the patterns we see in data. Here, we provide technical details of the software development, utilizing interactive graphics in R (R Core Team, 2022). We then share results from our study which validates ‘You Draw It’ as a method for graphical testing on linear trends and apply an appropriate data analysis method to the participant data. We also briefly demonstrate the use of the ‘You Draw It’ method and analysis on non-linear data.

2 Development

2.1 ‘You Draw It’ v0

The NYT uses D3, Data Driven Documents (Bostock et al., 2011), to create many of the interactive data graphics that users rely on, from the famous “election needle” in 2016 to their COVID dashboards. Naturally, this same framework was used to create the ‘You Draw It’ feature. The NYT code used to create the initial D3 plot and ‘You Draw It’ interactive feature includes multiple working parts. First, the relationship between coordinate mappings and pixels sets up the chart and the initial data is plotted. Next, the `d3.drag` function adds user interactivity and observes the movement of the user’s mouse. Additionally, the `forEach` loop finds the closest point in the coordinate system to the user’s mouse. Finally, the code detects if all points are filled by the user and updates the status of the task.

When working with a new programming language, setting up a new development environment can be challenging. As R users, we found the software overhead required to develop straight JavaScript code to be intimidating (Wattenberger, 2019), but luckily, the R community anticipated this barrier. In the next section, we will see how the `r2d3` package (Strayer et al., 2020) in R can be used to render the JavaScript code and display D3 visuals in familiar R HTML formats such as Rmarkdown and Shiny.

2.2 Generating D3 Plots in R

We leveraged the `r2d3` package to take data randomly generated in R and create a D3 plot that could be modified with the ‘You Draw It’ script. This step allowed us to more easily connect the data generation process with the resulting D3 code without having to generate new code manually each time we re-generated the data. While not strictly necessary for the validation test of the ‘You Draw It’ method compared to past linear regression methods presented in this paper, this abstraction makes it much easier to test arbitrary or model-generated data which is unique to each participant.

We conducted all data simulation and processing in R and output two data sets – *point data* and *line data* – containing (x, y) coordinates corresponding to either a simulated point or fitted value predicted by a statistical model respectively. Then, the `r2d3` package converts the data sets in R to JavaScript Object Notation (JSON) to be interpreted by JavaScript code included with the `r2d3` package. Parameters for aesthetic design choices are defined in a list of options

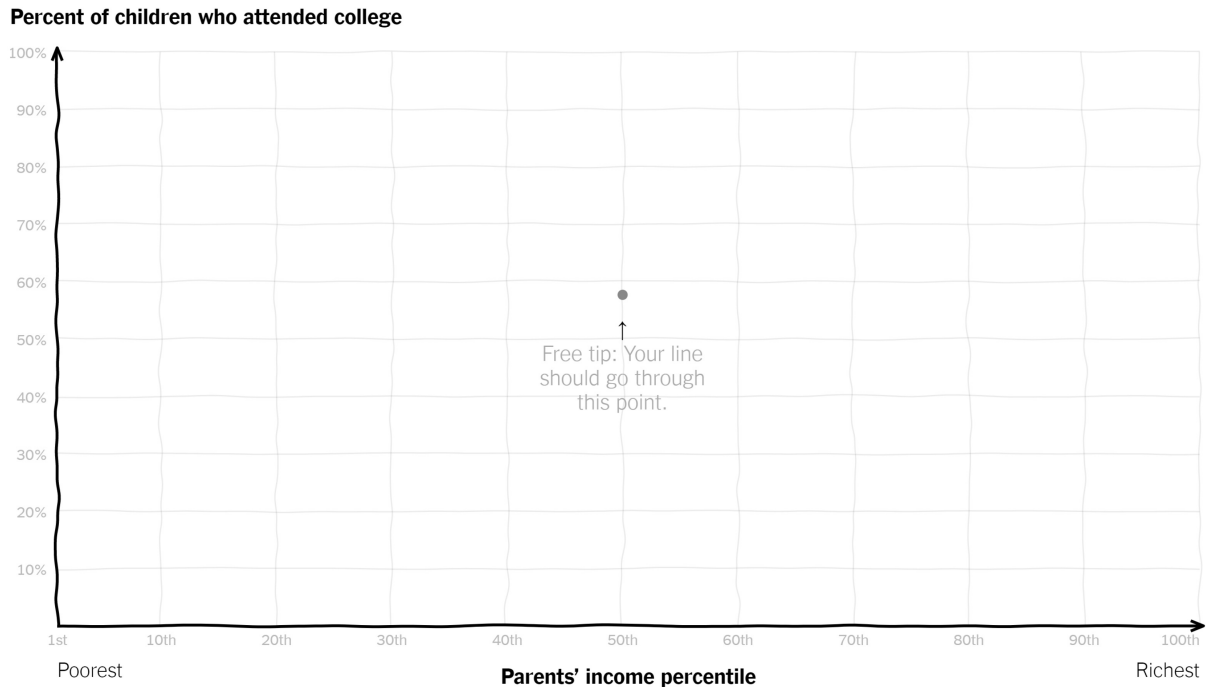


Figure 2: Screenshot of the ‘You Draw It’ application originally developed by the New York Times.

in the `r2d3` function call; `r2d3` passes these to the generated JavaScript code. For instance, we can specify the buffer space allowed for the x and y axes to avoid users anchoring their lines to the axes limits.

2.3 Adding ‘You Draw It’ Functionality

In order to make use of the original ‘You Draw It’ code for perceptual experiments, we first needed to accommodate an additional layer. The original NYT features asked participants to draw on a blank coordinate grid, as shown in Figure 2.

When testing perception of graphics, however, we want to pre-populate the chart with additional information – points, and in some cases, portions of a trend line. This requires that we modify the original JavaScript source code to accommodate these additional elements, which we generate using `r2d3`, as described above.

As previously discussed, we modified `r2d3` code to draw the initial plot and data. Drag events in `D3.js` are utilized to observe and react to user input, as in the original ‘You Draw It’ script from the NYT.¹ We adapted initial code from Pearce (2019) and Furic (2017).

One constraint we inherited from the NYT code is that users can only draw one-to-one functions; because the code works with drag events, each point in x can correspond to only one point in y , at least as the code is currently written. While this is not particularly problematic for our applications to date, it might limit applications of this type of user interaction when

¹The full source code we developed for the ‘You Draw It’ task is available at <https://github.com/earobinson95/presentations/blob/master/can-you-draw-it/www/you-draw-it.js>

assessing user drawn confidence bands, ribbons, and other situations where two or more vertical points are required.

2.4 Visual Cues

During testing of our modifications to the ‘You Draw It’ script, we discovered that frequently the JavaScript code would “skip” from point to point out of sequence. This resulted in a jagged line (visually) with missing values in the underlying stored array of data; as a result, the for-loop controlling the user-drawn line never exited and the user’s data was not recorded. While it would have been possible to fix this using some sort of interpolation algorithm, we did not want to compromise user results by introducing interpolation artifacts, so we opted instead for a visual cue that fixed the problem by modifying the human behavior.

As shown in Figure 3, the user-filled portion of the plot is represented with a yellow rectangle, which adapts to the user’s input and spans any missing values. When the box disappears, the array is filled in and the user can submit their response.

One challenge introduced by adding this visual cue is that D3 uses Scalable Vector Graphics (SVG) to render elements. Adding an additional layer meant that we had to ensure that not only the grids, lines, and points of the default scatterplot and trend line rendered in the correct order, but that the yellow box rendered behind all of these points as well so that no information was masked; and that this layer order updated in real time.

2.5 Connecting to Shiny

Shiny Messages are used to communicate the user interaction between the JavaScript code and the R environment. The initial plot is rendered using Shiny’s `RenderD3` and `d3Output` functions; subsequent user interactions are controlled via the JavaScript code. Once the user is finished modifying the plot, they can submit their response, so long as they have filled in all of the points. This is enforced using Shiny’s message-passing interface and a JavaScript hook that notifies Shiny when the array of user-drawn points is completely filled in. Once the user is done drawing the line, we save the results of the drawn line to a SQLite database, using Shiny to pass data from JavaScript to R. The connections between each portion of the app are shown in Figure 4.

One additional challenge when integrating the D3 graphics into the Shiny application is that by default, most Shiny applications use a reactive framework that adjusts to the user’s browser size. When working with D3, however, this can be problematic: most D3 parameters are specified at the pixel level. The discontinuity between the assumptions of Shiny and D3 meant that we had to fix the size of the Shiny element, and then perform a check to ensure that the user’s screen was sufficiently big to render the plot. While in an ideal world, users could participate using cell phones, tablets, and traditional laptop/desktop computers, functionally our checks limited users’ ability to participate using smaller screen sizes found in cell phones and some tablets. Additionally, after some feedback by laptop users during pilot testing, we included an additional requirement that participants have a computer mouse available; the results from touch-pad users were qualitatively different (more jagged) in ways suggesting that the recorded data did not adequately describe users’ perceptions.

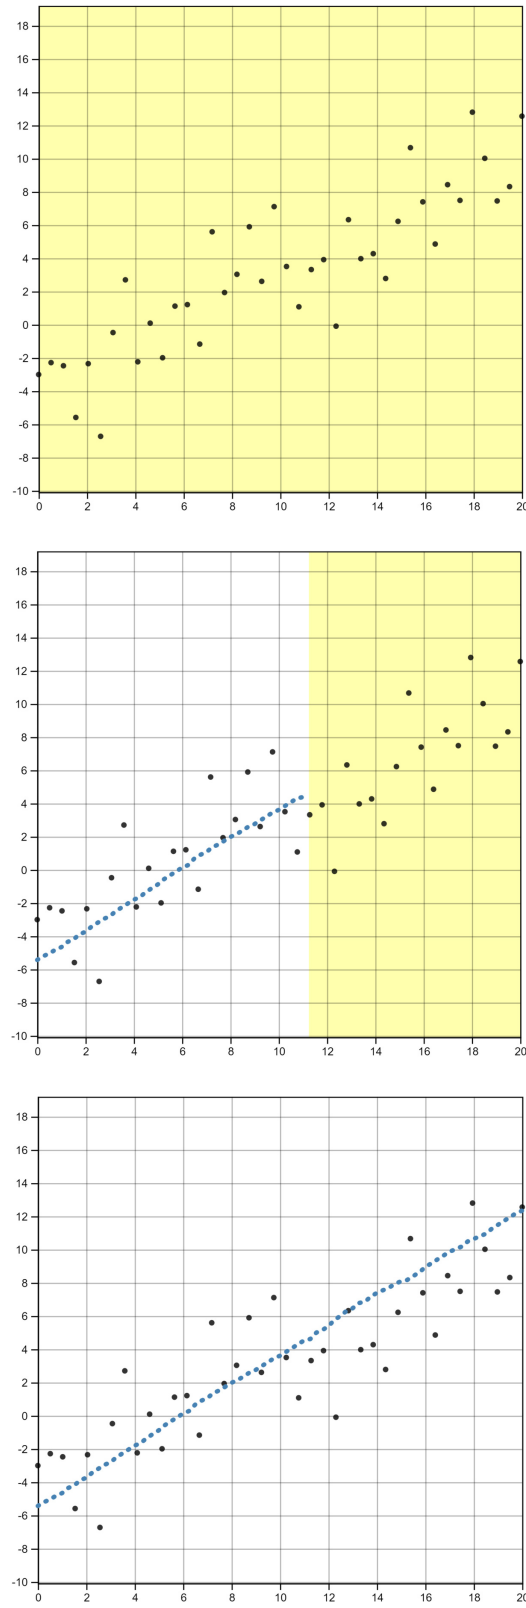


Figure 3: ‘You Draw It’ task plot as shown to user. **top:** Initial state, with instructions, *Use your mouse to fill in the trend in the yellow box region.* **middle:** User view during task completion. **bottom:** Finished state.

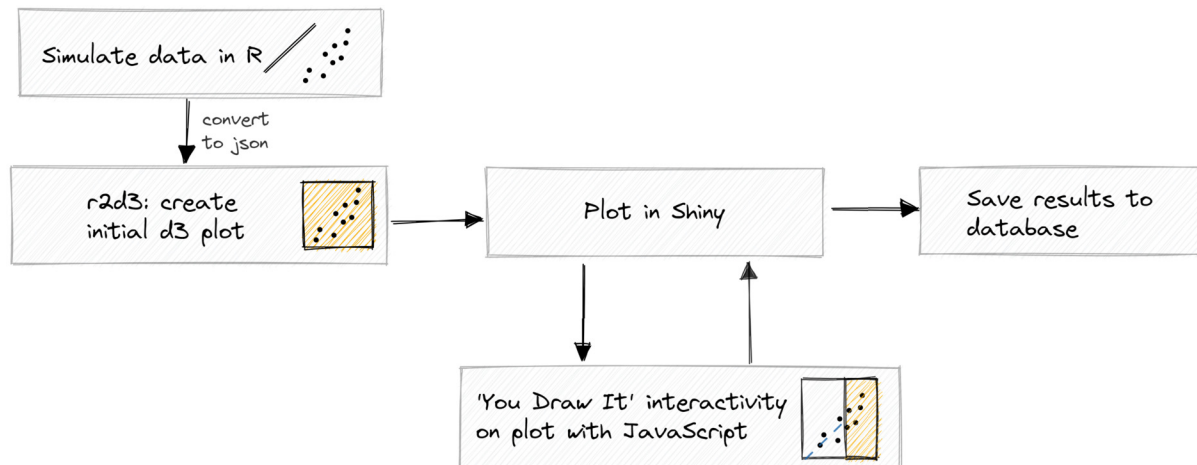


Figure 4: Sketch of underlying code for ‘You Draw It’, illustrating the data simulation conducted in R, the initial setup of the visual stimuli with D3 source code, along with the iterative process between the user interaction and plotting in Shiny. Once the user is done drawing the line, we save the results of the drawn line to a database for analysis.

3 Application

3.1 Validation Study

We conducted a study in order to validate ‘You Draw It’ as a method for graphical testing, comparing results to the less technological method utilized in Mosteller et al. (1981). The original study asked 153 graduate students and postdoctoral researchers in an Introductory Biostatistics course to fit lines by eye to a set of four points (S, F, V, N), each set with differing slope and variance properties, using an 8.5 x 11 inch transparency with a straight line etched across the middle. Researchers conducted the study with a Latin square experimental design to test for a practice effect on the task. Qualitative analysis suggested that participants tended to fit the slope of the first principal component (minimizes Euclidean distance) as opposed to the slope from the ordinary least squares regression (minimizes vertical distance) Figure 5 and found no effect of order.

In our study, we replicated Mosteller et al. (1981) using the ‘You Draw It’ method and simulated data with parameter coefficients selected to reflect those from the four data sets in the original study. Data were simulated based on a linear model with additive errors:

$$y_i = \beta_0 + \beta_1 x_i + e_i \quad (1)$$

with $e_i \sim N(0, \sigma^2)$.

For each participant, four unique data sets were randomly and independently generated from the underlying parameters at the beginning of the study and mapped to a scatterplot graphic. When participants started the study, they were first asked to complete two practice plots – accompanied by instructions and a .gif demonstrating the task – in order to train them in the skills associated with executing the task. The four ‘You Draw It’ task plots associated with the selected parameters followed the practice plots in a random order for each individual; these plots were interspersed with plots from a different experiment.

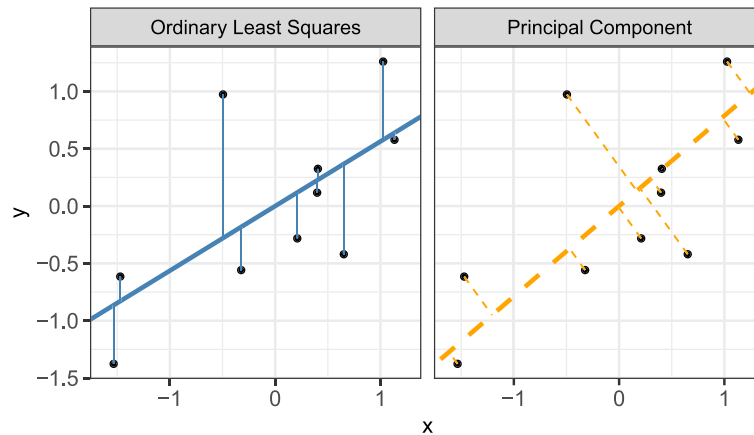


Figure 5: Comparison between an OLS regression equation which minimizes the vertical distance of points from the line and a regression equation with a slope calculated by the first principal component which minimizes the smallest distance of points from the line.

Participants were recruited via Prolific, a crowd sourcing platform for researchers to recruit participants for their online studies, in March 2022; while this study does utilize a convenience sample, as this is primarily a perceptual task, previous results have found few differences between expert and non-expert participants in this context (VanderPlas and Hofmann, 2015). The study was conducted and distributed via a Shiny application² and participants completed the tasks on their own computer, with a computer mouse, and in an environment of their choosing. The data from this study were collected to validate this method of graphical testing, with the hopes of providing a new tool to assess graphical perception interactively.

3.2 Data Analysis and Results

As a part of the validation study, we collected data for each participant drawn trend line, allowing us to conduct formal analyses which compared the intuitive visually fit lines to statistical regression results – the original study lacked these formal comparisons, providing only summaries of averages, variances, and actual (least squares) values for the slope and intercept of each data set. The feedback data, stored in a SQLite database, contained the simulated data points $\{ (x_{ijk}, y_{ijk}) \}$ – the predicted values from the statistical regression models – $\hat{y}_{ijk,OLS}$, and $\hat{y}_{ijk,PCA}$ –, and the predicted values from the user drawn line – $y_{ijk,drawn}$ – for parameter choice $i = 1, 2, 3, 4$, $j = 1, \dots, N_{\text{participant}}$, and x_{ijk} value $k = 1, \dots, 4x_{\text{max}} + 1$. Figure 6 displays an example of all three fitted trend lines for parameter choice F.

A unique data set was simulated independently for each participant, therefore, comparisons of vertical residuals between the user drawn line and the OLS fitted values ($e_{ijk,OLS} = y_{ijk,drawn} - \hat{y}_{ijk,OLS}$) and PCA fitted values ($e_{ijk,PCA} = y_{ijk,drawn} - \hat{y}_{ijk,PCA}$) were used to assess the accuracy of participant drawn lines. We use mixed model methods to analyze residual trends and capture the variability due to participant differences.

²The R Shiny applet used to conduct the study can be accessed at <https://shiny.srvanderplas.com/perception-of-statistical-graphics/>

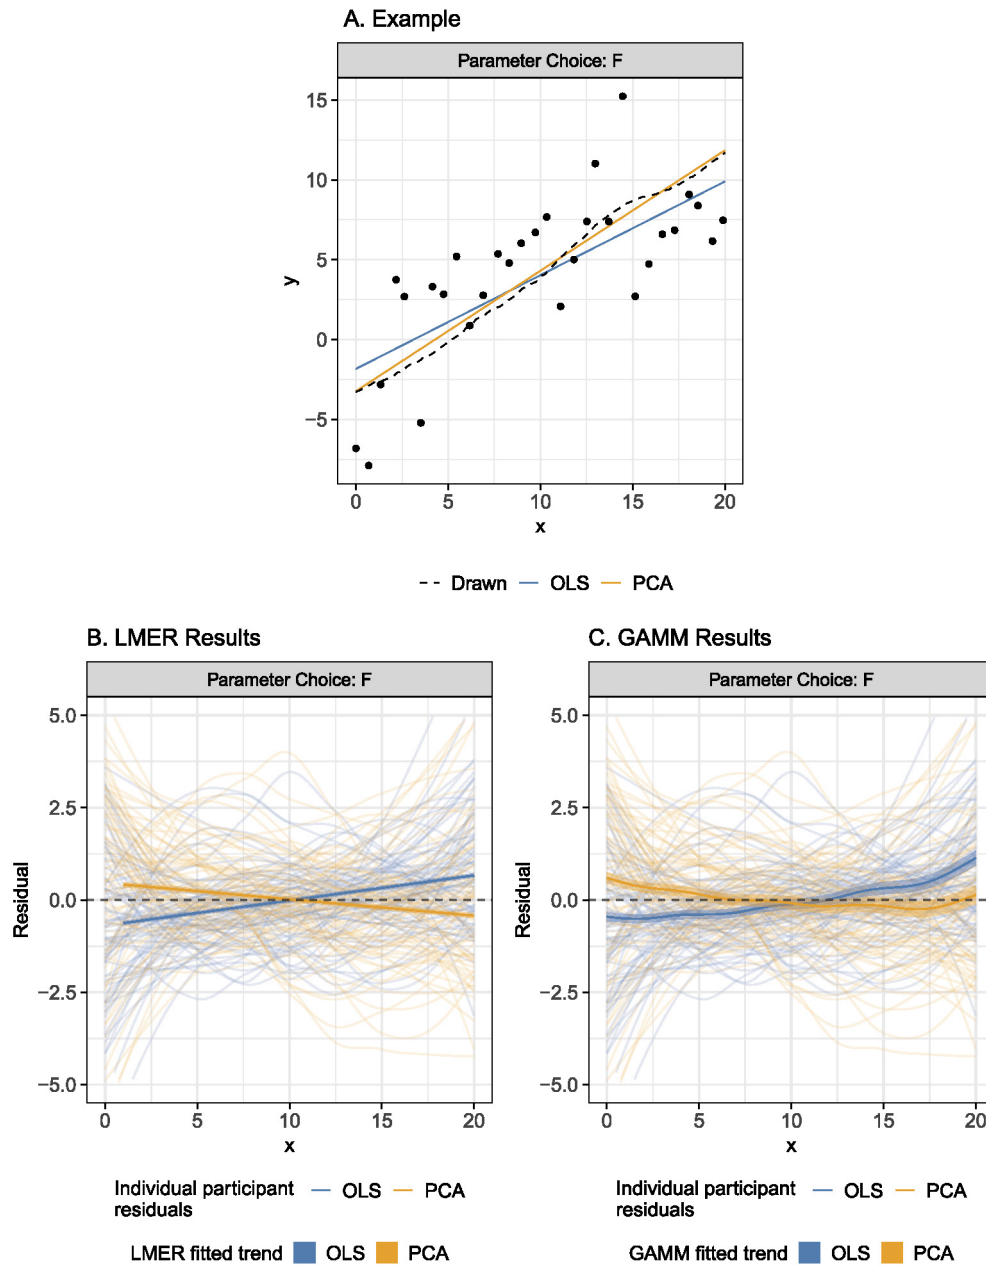


Figure 6: Visual analysis results for parameter choice set F. **A.** Example of validation feedback data where three trend lines show the OLS fitted, PCA fitted, and participant drawn values overlaid on the simulated data points. **B & C.** Estimated trends of residuals (vertical deviation of participant drawn points from both the OLS (blue) and PCA (orange) fitted points) as fit by the (B) linear mixed model and (C) generalized additive mixed model. A random sample of 75 participants was selected to display the individual participant residuals behind the overall trend.

We first fit a linear mixed model (LMM) to the OLS and PCA residuals separately using the `lmer` function in the `lme4` package (Bates et al., 2015), constraining the fit to a linear trend.

Parameter choice, x , and the interaction between x and parameter choice were treated as fixed effects with a random participant effect accounting for variation due to participant. The LMM equation for each fit (OLS and PCA) residuals is given by:

$$y_{ijk,drawn} - \hat{y}_{ijk,fit} = e_{ijk,fit} = [\gamma_0 + \alpha_i] + [\gamma_1 x_{ijk} + \gamma_{2i} x_{ijk}] + p_j + \epsilon_{ijk} \quad (2)$$

where

- $y_{ijk,drawn}$ is the drawn y-value for the i th parameter choice, j th participant, and k th increment of x -value
- $\hat{y}_{ijk,fit}$ is the fitted y-value for the i th parameter choice, j th participant, and k th increment of x -value corresponding to either the OLS or PCA fit
- $e_{ijk,fit}$ is the residual between the drawn and fitted y-values for the i th parameter choice, j th participant, and k th increment of x -value corresponding to either the OLS or PCA fit
- γ_0 is the overall intercept
- α_i is the effect of the i th parameter choice (F, S, V, N) on the intercept
- γ_1 is the overall slope for x
- γ_{2i} is the effect of the parameter choice on the slope
- x_{ijk} is the x -value for the i th parameter choice, j th participant, and k th increment
- $p_j \sim N(0, \sigma_{\text{participant}}^2)$ is the random error due to the j th participant's characteristics
- $\epsilon_{ijk} \sim N(0, \sigma^2)$ is the residual error.

Eliminating the linear trend constraint, we can extend the method and analysis beyond ordinary least squares by using generalized additive mixed models (GAMM) to estimate smoothing splines and allow for a more flexible residual trend. The `bam` function in the `mgcv` (Wood, 2017) package was used to fit a GAMM separately to the OLS and PCA. Parameter choice was treated as a fixed effect with no estimated intercept and a separate smoothing spline for x was estimated for each parameter choice. A random participant effect accounting for variation due to participant and a random spline for each participant accounted for variation in spline for each participant. The GAMM equation for each fit (OLS and PCA) residuals is given by:

$$y_{ijk,drawn} - \hat{y}_{ijk,fit} = e_{ijk,fit} = \alpha_i + s_i(x_{ijk}) + p_j + s_j(x_{ijk}) \quad (3)$$

where

- $y_{ijk,drawn}$ is the drawn y-value for the i th parameter choice, j th participant, and k th increment of x -value
- $\hat{y}_{ijk,fit}$ is the fitted y-value for the i th parameter choice, j th participant, and k th increment of x -value corresponding to either the OLS or PCA fit
- $e_{ijk,fit}$ is the residual between the drawn and fitted y-values for the i th parameter choice, j th participant, and k th increment of x -value corresponding to either the OLS or PCA fit
- α_i is the intercept for the parameter choice i
- s_i is the smoothing spline for the i th parameter choice
- x_{ijk} is the x -value for the i th parameter choice, j th participant, and k th increment
- $p_j \sim N(0, \sigma_{\text{participant}}^2)$ is the error due to participant variation
- s_j is the random smoothing spline for each participant.

This method allows us to quantitatively compare participant lines to the OLS and principal component regression lines, which is an improvement over the qualitative approach in Mosteller et al. (1981). While the statistical models were fit to all four parameter choices, Figure 6 displays the estimated trends of residuals (vertical deviation of participant drawn points from both the

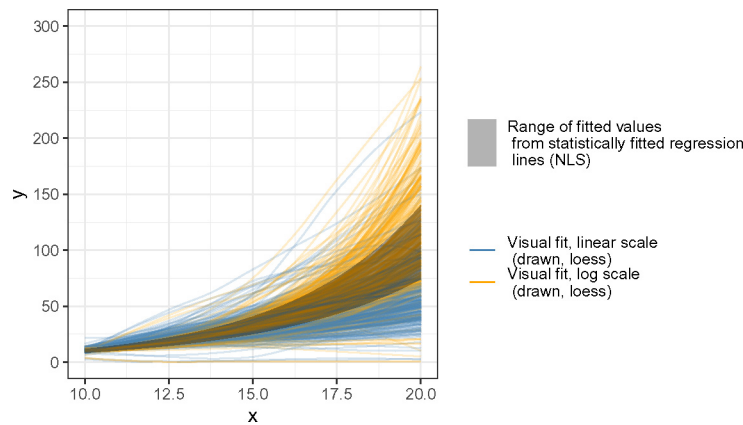


Figure 7: Spaghetti plot of results from a study which asked participants to forecast trends of exponentially increasing data. Participants drawn lines on the linear scale are shown in blue and the log scale are shown in orange. Variability in the statistically fitted regression lines occurred due to a unique data set being simulated for each individual; the gray band shows the range fitted values from the statistically fitted regression lines.

OLS and PCA fitted points) as modeled by a LMM and GAMM respectively for parameter choice set F. A random sample of 75 participants was selected to display individual participant residuals behind the overall residual trend. Examining the plots, the estimated trends of PCA residuals (orange) appear to align more parallel and closer to the $y = 0$ horizontal (dashed) line than the OLS residuals (blue). In particular, this trend was more prominent in parameter choices with large variances. Results from our study were consistent with those found in the original study; when shown points following a linear trend, participants tended to fit the slope of the first principal component. Our study established ‘You Draw It’ as a method for graphical testing and reinforced the differences between intuitive visual model fitting and statistical model fitting, providing information about human perception as it relates to the use of statistical graphics.

3.3 Extension to Non-linear Data

The analysis method presented above extends nicely beyond the linear regressions tested in Mosteller et al. (1981). Here we briefly demonstrate this with an example from a study examining participants ability to make forecasts for exponentially increasing data on log and linear scales. Along with analyzing the feedback data with the GAMM method for flexibility due to non-linear data, we used spaghetti plots to conduct a visual analysis of participant forecasts compared to the non-linear least squares statistical model and make comparisons between two chart design features Figure 7. The combination of the GAMM analysis and visual display demonstrates the strength of the ‘You Draw It’ method for testing statistical graphics.

4 Discussion

The ‘You Draw It’ graphical task presented in this work is an example of an interactive graphic. Here we discuss the use of ‘You Draw It’ as a tool for interactive reading, interactive testing, and interactive teaching as a way to increase user experience and user engagement with data. The New York Times first introduced the ‘You Draw It’ feature as a tool for interactive reading.

The inclusion of interactive graphics in news and media have been used to inform readers in efforts to reduce misconceptions and aid in storytelling (Geidner et al. 2015; George-Palilonis and Spillman, 2011; 2013). In our work, we demonstrated the use of the ‘You Draw It’ method as a tool for interactive testing of graphics. We collected and compared data from user drawn lines to statistical results in order to understand perceptual biases. While this work has many possible directions for future work, there is a strong application to extend the ‘You Draw It’ tool for interactive teaching. The line of best fit is introduced in secondary schools as a student’s first encounter with the concept of statistical association between multiple variables (Casey and Nagle, 2016). Casey (2015) and Casey (2016) conducted task based interviews to examine student’s conceptions of the line of best fit prior to introducing the topic. The implications of using interactive teaching tools for the line of best fit can help students make connections between the shift in slope and the location of the best fit line (Nagle et al., 2017). Casey and Wasserman (2015) conducted similar task based interviews with pre-service mathematics teachers in order to understand the effects regarding teacher preparation for the teaching of this topic. This demonstrates the use of the ‘You Draw It’ interactive tool for training purposes both within a classroom setting and for those in the workforce.

5 Conclusion

The data presented in this paper are part of a much broader experiment on the perception of log scales; while this paper focuses primarily on the computational implementation of the ‘You Draw It’ method and its adaptation to testing graphics, we are currently finishing the analysis and description of the broader experiment, which uses this ‘You Draw It’ method to assess user prediction of exponential trends.

By introducing an interactive method for assessing eye-fit data summaries, along with an analysis method which allows us to test competing hypotheses to determine whether user responses are similar to particular statistical models, we have provided another tool for experimentally testing statistical charts. In the future, we hope to create an R package to more easily facilitate these types of user experiments, making this technique available to other researchers who may not be willing to tinker with JavaScript directly.

As with any method for testing statistical graphics, there are a host of additional studies which would be useful to understand how users react to the ‘You Draw It’ method and what parameters are most important for the researcher to control. One avenue of future exploration is to investigate the effect of axis limits on user anchoring. We know that the perceptual experience is heavily affected by anchoring to e.g. axis breaks, and we would expect that a similar anchoring effect might be present based on the limits of the plot and the amount of space in x and y provided for the user to draw.

Supplementary Material

- **‘You Draw It’ Demonstration Applet:** The shiny app used to demonstrate the ‘You Draw It’ method can be accessed at <https://emily-robinson.shinyapps.io/can-you-draw-it/>.
- **Participant Data (Linear):** De-identified participant data collected in the linear study and used for analyses are available to be downloaded from GitHub at <https://github.com/earobinson95/sdss-2022-you-draw-it-manuscript/raw/master/data/eyefitting-model-data.csv>.
- **Participant Data (Non-linear):** De-identified participant data collected in the non-linear

study and used for analyses are available to be downloaded from GitHub at <https://github.com/earobinson95/sdss-2022-you-draw-it-manuscript/raw/master/data/youdrawit-model-data.csv>.

- **Data Analysis Code:** The code used to replicate the linear and non-linear study analysis in this paper can be found at <https://github.com/earobinson95/sdss-2022-you-draw-it-manuscript/blob/master/analysis/data-analysis.Rmd>.

References

- Aisch G, Cox A, Quealy K (2015). You draw it: How family income predicts children’s college chances. *The New York Times*, 28.
- Bates D, Mächler M, Bolker B, Walker S (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1): 1–48. <https://doi.org/10.18637/jss.v067.i01>.
- Bostock M, Ogievetsky V, Heer J (2011). D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12): 2301–9.
- Casey SA (2015). Examining student conceptions of covariation: A focus on the line of best fit. *Journal of Statistics Education*, 23(1).
- Casey SA (2016). Finding what fits. *Mathematics Teaching in the Middle School*, 21(8): 484–91.
- Casey SA, Nagle C (2016). Students’ use of slope conceptualizations when reasoning about the line of best fit. *Educational Studies in Mathematics*, 92(2): 163–77.
- Casey SA, Wasserman NH (2015). Teachers’ knowledge about informal line of best fit. *Statistics Education Research Journal*, 14(1): 8–35.
- Ciccione L, Dehaene S (2021). Can humans perform mental regression on a graph? Accuracy and bias in the perception of scatterplots. *Cognitive Psychology*, 128: 101406.
- Deming EW (1943). *Statistical Adjustment of Data*.
- Finney DJ (1951). Subjective judgment in statistical analysis: An experimental study. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2): 284–97.
- Furic B (2017). “You-Draw-It-Graph: The User Can Draw His Own Graph Line and Compare with the Reality.” GitHub. <https://github.com/BenoitFuric/you-draw-it-graph>.
- Geidner N, Pjesivac I, Imre I, Coman I, Yuran D (2015). The role of interactive graphics in reducing misperceptions in the electorate. *Visual Communication Quarterly*, 22(3): 133–45.
- George-Palilonis J, Spillman M (2011). Interactive graphics development: A framework for studying innovative visual story forms. *Visual Communication Quarterly*, 18(3): 167–77.
- George-Palilonis J, Spillman M (2013). Storytelling with interactive graphics: An analysis of editors’ attitudes and practices. *Visual Communication Quarterly*, 20(1): 20–27.
- Mosteller F, Siegel AF, Trapido E, Youtz C (1981). Eye fitting straight lines. *The American Statistician*, 35(3): 150–52.
- Nagle C, Casey S, Moore-Russo D (2017). Slope and line of best fit: A transfer of knowledge case study. *School Science and Mathematics*, 117(1–2): 13–26.
- Pearce A (2019). “You-Draw-It.” Observable. <https://bl.ocks.org/1wheel/07d9040c3422dac16bd5be741433ff1e>.
- R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- Strayer N, Luraschi J, Allaire JJ (2020). R2d3: Interface to ‘D3’ Visualizations. <https://CRAN.R-project.org/package=r2d3>.
- Unwin AR, Wills G (1988). Eyeballing time series. In: *Proceedings of the 1988 ASA Statistical Computing Section*, 263–68.

- US Energy Information Administration (2022). Weekly U.S. All Grades All Formulations Retail Gasoline Prices (Dollars Per Gallon). *US Energy Information Administration Independent Statistics & Analysis*. https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=EMM_EPM0_PTE_NUS_DPG&f=W.
- VanderPlas S, Hofmann H (2015). Spatial reasoning and data displays. *IEEE Transactions on Visualization and Computer Graphics*, 22(1): 459–68.
- Wattenberger A (2019). Fullstack D3 and Data Visualization: Build Beautiful Data Visualizations and Dashboards with D3. Fullstack.io.
- Wood SN (2017). *Generalized Additive Models: An Introduction with R*. 2nd ed. Chapman; Hall/CRC.