**Monte Carlo Simulations to Develop Guidelines for the Use of Single-Case Designs:
A Tutorial**

Marc J. Lanovaz^{1,2}

¹École de psychoéducation, Université de Montréal

²Centre de recherche de l'Institut universitaire en santé mentale de Montréal

Author Note

This study was supported in part by a salary award from the Fonds de recherche du Québec – Santé (#269462) to the first author.

Correspondence concerning this article should be addressed to Marc J. Lanovaz, École de psychoéducation, Université de Montréal, C.P. 6128, succursale Centre-Ville, Montreal, QC, Canada, H3C 3J7. Email: marc.lanovaz@umontreal.ca, Phone: 1 514-343-6111 #81774

Abstract

Despite being a cornerstone of the science of behavior analysis, researchers and practitioners often rely on tradition and consensus-based guidelines, rather than empirical evidence, to make decisions about single-case designs. One approach to develop empirically-based guidelines is to use Monte Carlo simulations for validation, but behavior analysts are not necessarily trained to apply this type of methodology. Therefore, the purpose of our technical article is to walk the reader through conducting Monte Carlo simulations to examine the accuracy, Type I error rate, and power of a visual aid for AB graphs using R Code. Additionally, the tutorial provides code to replicate the procedures with single-case experimental designs as well as with the Python programming language. Overall, a broader adoption of Monte Carlo simulations to validate guidelines should lead to an improvement in how researchers and practitioners use single-case designs.

Keywords: Monte Carlo, Python, R code, n-of-1 trial, single-case experimental designs, validity

Monte Carlo Simulations to Develop Guidelines for the Use of Single-Case Designs: A Tutorial

From its early beginnings, single-case experimental designs have been a cornerstone in the development of the science of behavior analysis (Baer et al., 1968; Sidman, 1960). To this day, the *Journal of Applied Behavior Analysis* (JABA) remains the journal that publishes the most studies that rely on single-case designs (Rader et al., 2021; Shadish & Sullivan, 2010; Smith, 2012). Despite its heavy reliance on single-case designs to reach data-driven conclusions, JABA has published surprisingly few articles on the validity of single-case designs and their associated methods of analysis. The two notable exceptions involve the visual inspection of single-case designs (e.g., Ford et al., 2020; Kahng et al., 2010; Wolfe et al., 2018) and the dual-criteria methods of analysis (e.g., Fisher et al., 2003; Falligant, McNulty et al., 2020), both of which have been the subject of multiple studies published in the journal.

One popular area of research in single-case designs involves examining correspondence between multiple methods of analysis on published or clinical data (Ninci et al., 2015). For example, Ford et al. (2020) have examined interrater agreement between visual analysts on graphs published in JABA whereas Wolfe et al. (2018) have extracted graphs from multiple journals to monitor correspondence between visual analysis and the conservative dual-criteria (CDC) method. Both studies have found agreement to be high between raters and methods (i.e., above 80%). That said, the main issue with these types of studies is that examining correspondence, or agreement, is not the same as assessing validity. Two raters or two different methods may agree, but still be wrong. Correspondence is especially concerning when using visual inspection as high agreement may be the result of the experts and trainees reading similar

journals, learning the same rules, or contacting similar contingencies when submitting research for publication.

The problem with using nonsimulated data is that researchers and clinicians can never know with certainty whether a graph shows a *true* effect or not, which is central to calculating Type I error rate and power. In single-case designs, Type I error rate involves the probability of concluding that a graph shows an effect when, in reality, the graph contains no true effect (i.e., false positive). On the other hand, power consists of the probability of detecting an effect in a graph when the graph contains a true effect. When developing guidelines to analyze single-case designs, the goal involves minimizing Type I error rate (5% or less) while maximizing power (80% or more). Given that nonsimulated data preclude the identification of Type I error rate and power, researchers have used Monte Carlo simulations to generate novel graphs (e.g., Fisher et al., 2003; Falligant, Kranak et al., 2020; Lanovaz et al., 2020). At its broadest, Monte Carlo simulation involves using the properties of known distributions to sample observations and then analyzing their properties (Mooney, 1997). In single-case design, Monte Carlo simulations can generate graphs with and without added effects, which are known to the researcher. Because the researcher knows with certainty which graphs include true effects (and which do not), they may compute the Type I error rate and power for the method of analysis.

Although Monte Carlo simulations can be an efficient method to conduct research on single-case designs, one potential limiting factor is that coding and statistics are not typically tasks associated with training in behavior analysis (see Behavior Analyst Certification Board, n.d.). Some researchers have used Microsoft Excel to conduct this type of research (e.g., Fisher et al., 2003; Falligant, Kranak et al., 2020), but this approach is inefficient with large datasets. An alternative involves coding the simulations and analyses using R and Python, two of the most

popular statistics and computing coding languages in science (Baker, 2017). However, behavior analysts are probably unfamiliar with these tools. Therefore, the purpose of our technical article is to teach behavior analysts how to conduct Monte Carlo simulations using common programming languages to encourage more research on guidelines for single-case designs.

Monte Carlo Simulations

To improve the flow of the manuscript, the technical article will walk the reader through the procedures using R code. R code has the advantage of being a common programming language to teach statistics to graduate students and is widely used by academics (Robinson, 2017). To encourage replication, the same code and procedures using Python are available in our repository. Similarly, our walkthrough will focus on simulating and analyzing AB graphs, but code for generating and graphing multiple baseline designs, reversal designs, and alternating-treatment designs are also available in our repository. Table 1 describes all files available in our online repository at: <https://github.com/labrl/monte-carlo-simulations-for-single-case-designs>.

Preliminary Steps

To complete this tutorial, you will need to first download and install R code at: <https://www.r-project.org/>. Once R code (version 4.1.1) has been installed, we strongly recommend that you download and install an integrated development environment, RStudio Desktop (version 14.4.1717), at: <https://www.rstudio.com/products/rstudio/download/>. Figure 1 shows the layout of the integrated development environment for RStudio (with code), which you should open at this point. The upper left panel contains the area where you will write and save your code. The lower left panel represents the console, which will print error messages and the results. The upper right panel shows the variables, objects, and functions from the global environment that R is working with. Finally, the lower right panel may display the working

directory, plots or the help page according to the tab selected by the user. The working directory is where you download and save files related to the project. To create a new project, click on “File” and select “New Project”. In “New Project,” select “New Directory” and name it, “MonteCarlo”. Finally, browse your computer to select where you want to save this directory. Note that you should remember this directory as you will need to download a file in that location later. Whenever you open RStudio, make sure that you also open the “MonteCarlo” project so that your working directory is set correctly.

Step 1 – Generating a Time Series

To create an AB data series, the first step involves generating a time series (i.e., a series of temporally dependent points). A basic time series consists of ordered autocorrelated points. In single-case designs, autocorrelation considers the association between a point and points that precede it. In other words, the autocorrelation represents how the frequency of a behavior in one session is related to the frequency of the behavior in prior sessions. To keep it simple, we will focus on the lag 1 autocorrelation, which involves the correlation between a point and the point that immediately precedes it. The mean lag 1 autocorrelation in single-case research is above 0. Shadish and Sullivan (2010) reported a value around .2 whereas a more recent study supports a larger value of .4 (Barnard-Brak et al., 2021). The basic formula to generate a time series with a lag 1 autocorrelation is as follows:

$$y_t = x_t + c$$

$$\text{where } x_t = ax_{t-1} + \varepsilon_t$$

In the previous equations, y_t represents the value of point y at session t , which is based on x_t plus a constant of c . The constant is designed to avoid negative graphs: We will set it at 10 as part of the technical article. The variable, x_t , involves multiplying the autocorrelation, a , with the prior

data point (before the addition of the constant), x_{t-1} , and adding an error term, ε_t . Prior studies have typically used an error term with a mean of 0 and standard deviation of 1 (Falligant, Kranak et al., 2020; Fisher et al., 2003; Lanovaz et al., 2020).

Now, we need to translate these equations to R code. In R, users can create functions so that they only need to write the code once to repeat a series of operations. In the upper left panel of RStudio, write the following code¹:

```
1 set.seed(48151)
2 create_time_series <- function(n, a, ct){
3   time_series = c()
4   point1 = rnorm(1)
5   time_series = c(time_series, point1)
6   for (i in 2:n){
7     point = a*time_series[i-1] + rnorm(1)
8     time_series = c(time_series, point)
9   }
10  time_series = time_series + ct
11  return(time_series)
12 }
```

The first line sets the random seed. Given that our code involves randomness, setting the random seed will allow you to obtain the exact same results as those reported in the current article when comparing outputs. Line 2 creates a function called `create_time_series`, which will create a time series with n points, an autocorrelation of a , and a constant of ct . Note that user-defined functions must begin and end with curly brackets in R (see the end of line 2 and line 12). Line 3 to 11 represent the operations conducted by the function. The third line creates an empty time series², the fourth line creates the first point (i.e., `rnorm(1)` instructs R to randomly pick one point from a normal distribution with mean of 0 and a standard deviation of 1), and the fifth line adds

¹ R code is case sensitive. If you wish to copy and paste the code, we recommend that you do it from the `functions.R` (steps 1 to 5) and `MonteCarlo.R` (steps 6 to 10) files in our repository to keep the formatting and indentations. You should avoid copying the line numbers.

² In R, you can create ordered combinations of values by using the built-in function `c()`. For example, `c()` produces an empty vector whereas `c(1,2,3)` creates a vector containing the values one to three in order.

the first point to the time series. This first point does not have autocorrelation because no point precedes it. The sixth line produces a loop to create all subsequent points (from point 2 to point n). Loops also begin and end with curly brackets (see the end of line 6 and line 9). During the loop, line 7 runs $ax_{t-1} + \varepsilon_t$ to create an autocorrelated point and line 8 adds the point to the time series. This process is repeated to create all points. Finally, line 10 adds the constant to all points and line 11 will return the time series when we call the function (i.e., when we write `create_time_series(n, a, ct)`).

To load your function, click anywhere in the code of the function and press on Ctrl+Alt+F on a Windows computer or Cmd+Option+F on an Apple computer. You know that your function has been loaded if you see it appear in your global environment (i.e., you should see `create_time_series` appear in the upper right panel). To test the function, write the following in the console³ in the lower left panel: `time_series = create_time_series(15, 0.2, 10)`. This code creates one time series (referred to as `time_series`) with 15 points, an autocorrelation of .2, and a constant of 10. To check if it worked, you may look for the time series in your global environment (upper right panel) or type the following code in your console: `time_series`. When you press the Enter key, a series of 15 values beginning with 10.59 and ending with 10.22 should appear in your console. We have now created a function to produce a time series with autocorrelation.⁴

Step 2 – Creating an AB Data Series

³ Whenever we ask you to write code in the console (lower left panel) instead of your script file (upper left panel), we will display this code directly in the text (rather than in code blocks) to minimize errors.

⁴ To teach the logic underlying Monte Carlo simulation, we showed you how to produce a time series by creating a function. That said, you can create an autocorrelated time series with n points and a lag 1 autocorrelation of a using a single line of code by applying a built-in function in R:
`time_series = arima.sim(list(order = c(1,0,0), ar = a), sd=1, n = n) + ct`

The next step involves creating a second function (below the `create_time_series` function) that produces an AB data series with a specific number of points in Phase A, a specific number of points in Phase B, and an added effect (if applicable) in Phase B. For the effect, we will use the standardized mean difference (SMD), which represents the number of standard deviations above the baseline mean that the behavior has changed. In single-case research, we typically expect SMDs of 3 or more, but testing lower values is important to examine the robustness of the analyses (Lanovaz & Hranchuk, 2021). Graphs with no added effects ($SMD = 0$) test for Type I error rate whereas graphs with an effect ($SMD > 0$) allows us to test for power. On the line below the end of your first function, write the following code⁵:

```
1  create_AB_data <- function(time_series, nb_pointsA, nb_pointsB,
    smd) {
2    total_points = nb_pointsA + nb_pointsB
3    PhaseA = time_series[1:nb_pointsA]
4    PhaseB = time_series[(nb_pointsA+1):total_points] + smd
5    labels = c(rep('A', nb_pointsA), rep('B', nb_pointsB))
6    AB_data = list(labels, c(PhaseA, PhaseB))
7    return(AB_data)
8  }
```

The `create_AB_data` function (line 1) takes as input the time series that we previously created (`time_series`), the number of points in Phase A (`nb_pointsA`), the number of points in Phase B (`nb_pointsB`), and an SMD value (`smd`). Line 2 computes the total number of points in the AB data series and line 3 extract the points from the time series that should be added to Phase A. You may notice that we used brackets to extract the data from the time series : `time_series[1:nb_pointsA]` tells R to take the points of the time series from 1 to number of points in Phase A, inclusively. If the number of points in Phase A was 5: `time_series[1:nb_pointsA]` would extract points 1, 2, 3, 4 and 5 from the time series and include them in Phase A. Line 4

⁵ A line without a number indicates that it is a continuation of the prior line.

extracts the points for Phase B and adds the SMD value to these points only. Line 5 produces vector with the labels for each phase (e.g., [“A”, “A”, “A”, ... “B”, “B”, “B”]). The rep built-in function repeats each label for the nb_pointsA for Phase A and nb_pointsB for Phase B. Finally, the sixth line put the labels and the data values in a list, which is returned on line 7. To load the function, press Ctrl+Alt+F which will then appear in your environment as create_AB_data. Type the following code in your console to test the function: `AB_data = create_AB_data(time_series, 5, 10, 3)`. This code creates AB data for a graph with 5 points in Phase A, 10 points in Phase B, and a SMD of 3. If you type `AB_data` in your console, you should see a list of labels (starting with “A” and ending with “B”) and a list of numbers appears (starting with 10.59 and ending with 13.22).

Step 3 – Adding Trend (Optional)

In some cases, researchers may want to add a trend to a graph. For example, Kahng et al. (2010) added a 30-degree trend to their data by pivoting around the center point. To add trend, we can use the following function:

```
1 add_trend <- function(AB_data, tr){
2   middle_point = median(1:length(AB_data[[2]]))
3   for (i in 1:length(AB_data[[2]])){
4     distance = i - middle_point
5     AB_data[[2]][i] = AB_data[[2]][i] + distance*tan(tr*pi/180)
6   }
7   return (AB_data)
8 }
```

In this case, add_trend is a function that takes the AB data (AB_data) and a trend in degrees (tr) as input. Line 2 identifies the middle point around which the others will pivot. In our code, AB_data[[2]] refers to the values of the graphs whereas AB_data[[1]] refers to the labels. For a specific value or label, we need to add an index. For example, we would type AB_data[[2]][3] to retrieve the third value in the AB data series. The length built-in function provides the number of

items in the list (i.e., number of data series). Line 3 applies a function to pivot each point by using a loop. Next, the function computes the distance between the point and the middle point (line 4) and pivots the point by multiplying the distance by the tangent of the trend in radians ($\text{tr} \cdot \pi / 180$). In simpler terms, the program uses trigonometry and algebra to add a trend. During the pivoting, points closer to the middle point vary less than those at the extremes. Points on the left side of the graph decrease in value whereas those on the right side increase. To add a 30 degree trend to the data, load your function with Ctrl+Alt+F and then write the following code in the console: `AB_data = add_trend(AB_data, 30)`. When typing `AB_data` again in the console, you should see that your point values now start at 6.55 and end at 17.26.

Step 4 – Graphing a Single AB Data Series

Researchers may have an interest in examining the validity of visual inspection. The purpose of the technical article does not include teaching graphing with R code, which is already described elsewhere (Bulté & Onghena, 2012), but we do provide an exemplar to support researchers. To graph the AB data, we adapted the code proposed Bulté and Onghena (2012) to work with our data organization format. You can copy and paste the following function in your code editor:

```
1 ABgraph <- function (AB_data){
2   xlab = "Measurement Times"
3   ylab = "Behavior"
4   x <- 1:length(AB_data[[2]])
5   plot(x, AB_data[[2]], xlab = xlab, ylab = ylab, ylim = c(0,
6     max(AB_data[[2]])*1.2), xlim = c(0, length(AB_data[[2]])),
7     pch = 16, axes = FALSE)
8   lines(x[AB_data[[1]] == "A"], AB_data[[2]][AB_data[[1]] == "A"])
9   lines(x[AB_data[[1]] == "B"], AB_data[[2]][AB_data[[1]] == "B"])
10  lines(c(sum(AB_data[[1]] == "A") + 0.5, sum(AB_data[[1]] == "A")
11    + 0.5), c(0, max(AB_data[[2]])*1.2), lty = 2)
12  axis(1, tck=-0.01)
13  axis(2, tck=-0.01, label = NA)
14  mtext('Phase A', side = 3, at = (sum(AB_data[[1]] == "A") + 1)/2)
15  mtext('Phase B', side = 3, at = (sum(AB_data[[1]] == "A") +
16    (sum(AB_data[[1]] == "B") + 1)/2))
```

```
13 }
```

We will not explain each line above, but a commented version of this code is available in the repository. Instead, you can copy and paste it, and then run it by pressing Ctrl+Alt+F or the Apple equivalent. Then, you can graph the data series by writing the following instructions in your console: `ABgraph(AB_data)`. You should see an AB graph appear either in a new window or in the “Plots” tab of the lower right panel. If you initially set the `random.seed` to the recommended value, your graph should look like the one presented in Figure 2.

Step 5 – Applying Structured Criteria to One AB Data Series

Our final function aims to automate data analysis. As the CDC method is arguably the best-known and most widely applied structured aid to analyze single-case graphs amongst behavior analysts, our example will focus on this method. The CDC involves tracing a continuation of the mean line and of the slope line from Phase A unto Phase B and increasing their heights by 0.25 standard deviations (Fisher et al., 2003). Then, the experimenter must count the number of points in Phase B that fall above both lines and compare it to a table of values from Fisher et al. (2003). If the count is equal to or greater than the threshold value, the graph is categorized as showing an effect. Otherwise, the graph is labeled as showing no effect.

The function to apply the CDC method is the following:

```
1 Fisheretal = c(NA, NA, 3, 4, 5, 6, 6, 7, 8, 8, 9, 9, 10, 11, 12,
2   12, 12, 13, 13, 13, 14, 14, 15)
3 CDC_method <- function(AB_data){
4   idx_A = which(AB_data[[1]]=='A')
5   idx_B = which(AB_data[[1]]=='B')
6   mean_line = mean(AB_data[[2]][idx_A]) +
7     0.25*sd(AB_data[[2]][idx_A])
8   lr_A = lm(AB_data[[2]][idx_A] ~ idx_A)
9   trend_line = idx_B*lr_A$coefficients[2]+lr_A$coefficients[1] +
10    0.25*sd(AB_data[[2]][idx_A])
11   count= sum(AB_data[[2]][idx_B] > trend_line & AB_data[[2]][idx_B]
12     > mean_line)
13   if (count >= Fisheretal[length(idx_B)]){
```

```
10     return (1)
11   } else {
12     return (0)
13   }
14 }
```

The first line imports the threshold values from Fisher et al. (2003). Line 2 creates a function, `CDC_method`, that takes our `AB_data` as input. Lines 3 and 4 identify the indices for the values of Phases A and B. Specifically, the `which` built-in function returns the indices for points that should be associated with Phase A and those that should be associated with Phase B. The fifth line computes the mean for Phase A and increases it by 0.25 standard deviation. Line 6 computes the slope and the intercept for Phase A using the `lm` built-in function from R. The code then uses these data to project a trend line on Phase B that has been increased by 0.25 standard deviation (line 7). Line 8 counts the number of points of Phase B that falls above both lines. Finally, line 9 and 10 return a value of 1 (effect) if the count is equal to or greater than the cutoff from Fisher et al. (2003). If not, the function returns a 0, indicating no effect (lines 11 and 12).

To examine the results, load the function and then run the following in the console: `CDC_method(AB_data)`. A value of 1 should appear in the console, indicating a change according to the CDC. Now, you should save your file with all the functions as `functions.R` (click on “File”, select “Save As”, and write “functions” in File name). At this point, you should note that Step 5 should not be limited to the CDC method and can include any guideline or rule that is developed by a research team (e.g., Vannest & Manolov, 2019). You may write a function for any guideline that you want to test. The important issue to bear in mind is that the output should remain 1 for an effect and 0 for no effect. In this manner, you can apply all the other functions and code presented in this technical article.

Step 6 – Producing Multiple AB Data Series

If applied to create a single graph, the previous procedures would seem highly effortful when compared to simply using a spreadsheet software. Monte Carlo simulation involves creating hundreds, or even thousands, of data series with minimal effort. The current step involves creating multiple AB data series by using our previous functions. First, open a new R file by clicking on “File” and selecting “New File”. To use the functions that we have previously created, the first line should import them with the following code:

```
1 source('functions.R')
```

To run the code, select the line and press on Ctrl+Enter on Windows or Cmd+Enter on Apple. If you get an error message, make sure that your functions.R file is named correctly and is in the project (working) directory. Next, we will set the values of each characteristic to test in our analyses:

```
1 nb_pointsA_list = c(3,5)
2 nb_pointsB_list = c(5,10)
3 a_list = c(0, 0.2, 0.4)
4 tr_list = c(0, 15, 30)
5 smd_list = c(0, 0, 0, 1, 2, 3)
6 ct = 10
```

The values represent each characteristic to use when generating AB data series. The code indicates that (1) the number of points in Phase A should be 3 or 5, (2) the number of points in Phase B should be 5 or 10, (3) the autocorrelation value should be 0, 0.2., or 0.4, (4) the trend should be 0 degrees, 15 degrees, or 30 degrees, (5) the SMD should be 0, 1, 2, or 3, and (6) the constant should be 10. Note that line 5 contains three 0s, which is designed to ensure that we have the same number of graphs showing no effect ($SMD = 0$) as those showing an effect ($SMD < 0$). You may vary these values according to the hypotheses to test.

Now, we will produce a list containing data series with each possible combination of the prior characteristics. To do so, write and run the following code:

```

1  all_AB_data = c()
2  true_values = c()
3  trend_values = c()
4  for (nb_pointsA in nb_pointsA_list){
5    for (nb_pointsB in nb_pointsB_list){
6      for (a in a_list){
7        for (tr in tr_list){
8          for (smd in smd_list){
9            time_series = create_time_series(nb_pointsA+nb_pointsB,
10              a, ct)
11            AB_data = create_AB_data(time_series, nb_pointsA,
12              nb_pointsB, smd)
13            AB_data = add_trend(AB_data, tr)
14            all_AB_data = c(all_AB_data, AB_data)
15            if (smd == 0){
16              true_values = c(true_values, 0)
17            } else {
18              true_values = c(true_values, 1)
19            }
20            trend_values = c(trend_values, tr)
21          }
22        }
23      }
24    }
25  }

```

This block of code will produce one list and two vectors. The list, `all_AB_data` (line 1) will contain 216 data series, one for each combination of characteristics ($2 \times 2 \times 3 \times 3 \times 6 = 216$). The list has 432 items because the odd numbered items represent the labels whereas the subsequent pair numbered items represent the values. The vector, `true_values` (line 2), indicates whether each data series shows a change (value = 1) or not (value = 0). The vector, `trend_values` (line 3), includes the trend in degrees for each graph

Lines 4 through 8 run loops to combine every possible combination of characteristics. The ninth line produces the time series, the tenth line splits the time series in Phases A and B while adding the SMD value to Phase B, the eleventh line adds the trend, and the twelfth line

includes the series in the complete dataset. Lines 13 to 16 adds the true values to the `true_values` vector, which will allow us to examine Type I error rate and power at the step 9. The final line with `text` compiles the trend for each graph, which will be used in step 10 to examine its effects.

Step 7 – Graphing Multiple AB Data Series

In the previous step, we produced 216 data series, which we can now readily graph with our `ABgraph` function. To facilitate analyses, we will collate all the graphs in a single .pdf file by running the following lines of code:

```
1 pdf("ABgraphs.pdf", paper="USr", height = 300, width=400)
2 for (i in seq(1,length(all_AB_data), 2)){
3   ABgraph(all_AB_data[i:(i+1)])
4 }
5 dev.off()
```

The first line opens a .pdf file (`ABgraphs.pdf`) in which to save each graph with specified dimensions. Line 2 runs a loop through each data series (remember, each graph is represented by a pair of objects) and line 3 graphs the data using our function written in step 4. Line 5 saves and closes the .pdf file. When you run the above code (select the code and then press `Ctrl+Return` or `Cmd+Return`), a new file should appear in your working directory. This file, `ABgraphs.pdf`, contains the 216 graphs that we have produced. Each graph matches a true value in the `true_values` vector. For example, the first graph in the .pdf file corresponds to the first value in the `true_values` vector.

Step 8 – Analyzing Multiple Graphs

If applying structured criteria, you should run your analysis function as instructed below:

```
1 cdc_results = c()
2 for (i in seq(1,length(all_AB_data), 2)){
3   cdc_results = c(cdc_results, CDC_method(all_AB_data[i:(i+1)]))
4 }
```


Line 1 creates a vector to hold the results, line 2 runs through all the data series, and line 3 applies the CDC method. The result is a vector, `cdc_results`, containing 1s and 0s. Similarly to the `true_values` vectors, the first item in the `cdc_results` vector corresponds to the first graph in the .pdf file, the second item corresponds to the second graph, and so on.

If the purpose of the study is to examine the validity of visual inspection, a blind visual rater could categorize each graph as showing a change (value = 1) or no change (value = 0). In this case, we recommend that the visual rater use a Microsoft Excel, Google Sheets, or Apple Numbers spreadsheet to enter their binary values. All data should be entered in the first column and each row should represent a different graph. For example, if the graph on page 5 of the .pdf file shows no effect according to the visual rater, they should enter a value of 0 in the first column and fifth row of the spreadsheet. At the end, the rater should save the spreadsheet as a .csv file (see `Expert_data.csv` for an exemplar).

Step 9 – Computing Accuracy, Type I Error Rate, and Power

The subsequent step involves measuring accuracy, Type I error, and power for the CDC method with the following code:

```
1 accuracy_CDC = sum(cdc_results == true_values)/length(true_values)
2 idx_0 = which(true_values == 0)
3 typeI_error_CDC = sum(cdc_results[idx_0])/length(idx_0)
4 idx_1 = which(true_values == 1)
5 power_CDC = sum(cdc_results[idx_1])/length(idx_1)
```

Line 1 shows the accuracy computation, which involves calculating the number of times that the CDC method matches the true values and dividing it by the total number of data series (i.e., 216). To see the result, run the code and write `accuracy_CDC` in the console; the value of .648 should appear. Line 2 identifies the graphs that include no effect and line 3 computes Type I error rate by dividing the number of graphs with no true change that show a change according to the CDC

by the total number of graphs showing no true change. If you type `typeI_error_CDC` in the console, you should see a value of .231. The last two lines identify the graphs that contain a true change (line 4) and computes power by dividing the number of graphs with a true change that show a change according to the CDC by the total number of graphs showing a true change (line 5). To obtain the power, you can enter `power_CDC` in the console, which will print a value of .527. In our example, both Type I error rate and power seem unacceptable as they are higher than .05 and lower than .80, respectively. These results were expected as prior research has found that the CDC struggles with autocorrelated and trended graphs (Fisher et al., 2003; Lanovaz & Hranchuk, 2021).

To show you how to apply the same analyses to visual inspection, we created a .csv file with simulated values (`Expert_data.csv`). The code is nearly identical to the one that we used for the CDC, except that line 1 imports the data from the .csv file:

```
1 expert_data = read.csv('Expert_data.csv', header = FALSE)[1]
2 accuracy_VI = sum(expert_data[1] == true_values)
  /length(true_values)
3 typeI_error_VI = sum(expert_data[idx_0,1])/length(idx_0)
4 power_VI = sum(expert_data[idx_1,1])/length(idx_1)
```

In this case, the `expert_data` vector replaces the `cdc_results`. The `read.csv` function in the first line comes in handy when importing data that were collected beforehand. As indicated earlier, you may produce .csv files using Microsoft Excel, Google Sheets, or Apple Numbers. We will let the reader print the accuracy, Type I error rate and power on their own for practice. Table 2 contains the results for both the analyses. It is important to repeat here that the values for the expert data were simulated to show you how to import them. However, these data should not be used to draw any conclusions regarding the validity of visual inspection.

Step 10 – Examining the Effect of Manipulating a Specific Characteristic

Researchers may show interest in conducting more fine-grained analyses of the effects of graph characteristics. In our example, we will examine the effects of trend on the CDC method as few studies have examined its effects on Type I error rate and power (Manolov & Vannest, 2019). The following code will produce a vector containing the Type I error rate for different values of trend:

```
1 error_by_trend = c()
2 for (tr in tr_list){
3   idx = intersect(which(trend_values == tr), which(true_values ==
4     0))
5   error_by_trend = c(error_by_trend,
6     sum(cdc_results[idx])/length(idx))
7 }
8 names(error_by_trend) = tr_list
```

The first line initializes a vector in which the Type I error rates will be saved. The next line repeats the analysis for each value of trend. In our case, the analysis involved three values of trend: 0, 15, and 30 degrees. The third line identifies which data series includes the programmed trend and no true effect. The intersect built-in function select data series that meet both conditions. The fourth line repeats the computation for Type I error rate from step 9. To facilitate analysis, line 8 associates each Type I error rate to its trend. Once you have run the code, you should write, `error_by_trend`, in the console. The console will display the following Type I error rates: .138 (0 degrees), .167 (15 degrees), and .388 (30 degrees). This analysis clearly shows that the high level of Type I error rate observed in the previous step may have been the product of the graphs with 30-degree trends.

The next step involves examining the power of the CDC according to trend:

```
1 power_by_trend = c()
2 for (tr in tr_list){
3   idx = intersect(which(trend_values == tr), which(true_values ==
4     1))
5   power_by_trend = c(power_by_trend,
6     sum(cdc_results[idx])/length(idx))
7 }
```

```
5 }  
6 names(power_by_trend) = tr_list
```

This block of code is similar to the one preceding it with two exceptions: the result vector is named `power_by_trend` and we only kept data series showing a true effect (i.e., `true_changes == 1` instead of `true_changes == 0`). To examine the effects of trend on power, select the code, press on Ctrl+Enter to run it, and enter `power_by_trend` in the console. Power increases as the trend increases: .444 (0 degrees), .527 (15 degrees), and .611 (30 degrees). Taken together, the patterns observed for Type I error rate and power underline a common problem when developing guidelines to analyze single-case data. Increases in power tend to be accompanied by increases in Type I error rate and vice versa (see Lanovaz & Hranchuk, 2021, for example with visual raters). As the aim of designing guidelines for single-case designs is to minimize Type I error rate while simultaneously increasing power, experimenters must carefully balance the two when developing novel methodology.

Some Study Design Considerations

When designing a study using Monte Carlo simulations, one of the first questions that comes to mind is: How many data series or graphs do we need to generate? The simple answer is the more the better. As you increase the number graphs, the margin of error on your measures decreases. For example, having only 200 graphs (100 graphs showing an effect and 100 graphs showing no effect) may produce an absolute error as high as $\pm .10$ whereas increasing to 1,000 graphs in total reduces this maximum error to $\pm .04$. If you want to analyze the effects of individual characteristics (e.g., trend, autocorrelation), the number of graphs needs to be even larger. Producing 10,000 graphs per characteristic type reduces the error to $\pm .01$. Beyond 10,000 graphs per type, the reductions in error are minimal and arguably inconsequential when it comes to applying these guidelines to the real-world. For visual inspection, researchers are limited by

the number of graphs that a rater can analyze. To our knowledge, the study that had the most graphs analyzed by visual raters included 1,024 different graphs (Lanovaz & Hranchuk, 2021).

This issue raises an interesting question for behavior analysts: Should we have a few raters analyze many graphs or more raters analyzing fewer graphs? Although many researchers have adopted the latter approach (e.g., Ford et al., 2020; Kahng et al., 2010; Wolfe et al., 2018), we would argue here that having fewer raters with more graphs produces results more consistent with the dimensions of applied behavior analysis. First, collapsing the data of multiple raters together does not provide valuable information about individual responding. When analyzing graphs in research or practice, visual raters do not typically calculate the mean of multiple raters to reach a conclusion. As the behavior of the individual is central to our science, we contend that we should examine the behavior of individual raters. Having more graphs with fewer raters allows for the examination of within-participant patterns, a hallmark of behavior analysis. Furthermore, having few exemplars of each combination of characteristics increases the risk of relying exclusively on outliers during the analysis, which may produce inconsistent results across studies. Having more graphs should balance the exemplars and outweigh potential outliers that may occur.

In the current article, we focused on how to conduct Monte Carlo simulations with AB graphs. In contrast, most researchers use single-case *experimental* designs in their published studies. Thus, extending and replicating this line of work with multiple baseline, reversal, and alternating-designs seems paramount. To support researchers, we provide the code to generate the data and graphs for experimental designs in our repository. Nevertheless, researchers will run into a question when developing their analyses: Should each panel or phase change be analyzed separately, or should graphs be analyzed as a whole? At this time, the research literature does not

provide a clear response to this question. Some studies have adopted an approach where the components of the design (e.g., phase change, panels) were analyzed individually (e.g., Ferron et al., 2017; Wolfe et al., 2016) while others have suggested analyzing the data of single-case designs as a whole (e.g., Manolov, 2019; Wolfe et al., 2019). The Monte Carlo methodology proposed in the current study could provide an answer to this conundrum by comparing the results of analysis of individual components versus whole graphs. Behavior analysts need more empirical evidence to support either approach.

A final consideration is what type of analysis should we test using Monte Carlo simulations. To keep it simple, the report focused on the CDC method, which has been the topic of multiple studies in JABA. Should behavior analysts pursue their efforts in improving visual inspection (e.g., Ford et al., 2020; Lane et al., 2021; Wolfe et al., 2019; Wolfe & Slocum, 2015) or should they focus on developing novel methodologies (e.g., Lanovaz & Hranchuk, 2021; Manolov & Vannest, 2019)? The best approach to analyzing single-case data probably remains unknown, but it is important to rely on data to make informed decisions. As most decision-making guidelines can now be computerized and integrated into easy-to-use graphing apps, the difference in effort of conducting visual inspection in comparison to applying structured aids is becoming minimal to nonexistent. Behavior analysts should rely on empirical evidence, rather than tradition, to decide how to analyze graphs. To this end, we advocate for more comparison studies (e.g., visual inspection vs. structured methods) to better inform researchers and practitioners.

Conclusion

Our tutorial provides step-by-step instructions to implementing Monte Carlo simulations for single-case designs in R while providing the necessary code to replicate them in Python. As

single-case methodology is central to both research and practice, conducting more studies using simulated data to examine guidelines appears essential. This tutorial may serve as a first step for researchers to conduct studies more thoroughly examining the properties and analysis of single-case designs. An additional benefit of this technical article is that it may serve as an introduction to R code and Python for some behavior analysts. Given that these two languages are amongst the most widely used for scientific computing (Baker, 2017), an introduction to these programming languages may lead some researchers to adopt them in their work and to develop potentially innovative studies, which they could not conduct using traditional software (e.g., Microsoft Excel, Google Sheets). In sum, our tutorial teaches skills that researchers can use to contribute to the advancement of the science of behavior analysis.

References

- Baer, D. M., Wolf, M. M., & Risley, T. R. (1968). Some current dimensions of applied behavior analysis. *Journal of Applied Behavior Analysis, 1*, 91-97.
<https://doi.org/10.1901/jaba.1968.1-91>
- Baker, M. (2017). Scientific computing: Code alert. *Nature, 541*(7638), 563-565.
<https://doi.org/10.1038/nj7638-563a>
- Barnard-Brak, L., Watkins, L., & Richman, D. M. (2021). Autocorrelation and estimates of treatment effect size for single-case experimental design data. *Behavioral Interventions*. Advance online publication. <https://doi.org/10.1002/bin.1783>
- Behavior Analyst Certification Board (n.d.). *Task lists for BACB certifications*.
<https://www.bacb.com/task-lists/>
- Bulté, I., & Onghena, P. (2013). The single-case data analysis package: Analysing single-case experiments with R software. *Journal of Modern Applied Statistical Methods, 12*(2), 28.
<https://doi.org/10.22237/jmasm/1383280020>
- Falligant, J. M., Kranak, M. P., Schmidt, J. D., & Rooker, G. W. (2020). Correspondence between fail-safe k and dual-criteria methods: Analysis of data series stability. *Perspectives on Behavior Science, 43*, 303-319. <https://doi.org/10.1007/s40614-020-00255-x>
- Falligant, J. M., McNulty, M. K., Hausman, N. L., & Rooker, G. W. (2020). Using dual-criteria methods to supplement visual inspection: Replication and extension. *Journal of Applied Behavior Analysis, 53*(3), 1789-1798. <https://doi.org/10.1002/jaba.665>

- Ferron, J. M., Joo, S. H., & Levin, J. R. (2017). A Monte Carlo evaluation of masked visual analysis in response-guided versus fixed-criteria multiple-baseline designs. *Journal of Applied Behavior Analysis*, 50(4), 701-716. <https://doi.org/10.1002/jaba.41>
- Fisher, W. W., Kelley, M. E., & Lomas, J. E. (2003). Visual aids and structured criteria for improving visual inspection and interpretation of single-case designs. *Journal of Applied Behavior Analysis*, 36(3), 387-406. <https://doi.org/10.1901/jaba.2003.36-387>
- Ford, A. L., Rudolph, B. N., Pennington, B., & Byiers, B. J. (2020). An exploration of the interrater agreement of visual analysis with and without context. *Journal of Applied Behavior Analysis*, 53(1), 572-583. <https://doi.org/10.1002/jaba.560>
- Kahng, S. W., Chung, K. M., Gutshall, K., Pitts, S. C., Kao, J., & Girolami, K. (2010). Consistent visual analyses of intrasubject data. *Journal of Applied Behavior Analysis*, 43(1), 35-45. <https://doi.org/10.1901/jaba.2010.43-35>
- Lane, J. D., Shepley, C., & Spriggs, A. D. (2021). Issues and improvements in the visual analysis of AB single-case graphs by pre-service professionals. *Remedial and Special Education*, 42(4), 235-247. <https://doi.org/10.1177/0741932519873120>
- Lanovaz, M. J., Giannakakos, A. R., & Destras, O. (2020). Machine learning to analyze single-case data: A proof of concept. *Perspectives on Behavior Science*, 43(1), 21-38. <https://doi.org/10.1007/s40614-020-00244-0>
- Lanovaz, M. J., & Hranchuk, K. (2021). Machine learning to analyze single-case to analyze single-case graphs: A comparison to visual analysis. *Journal of Applied Behavior Analysis*. Advance online publication. <https://doi.org/10.1002/jaba.863>

Manolov, R., & Vannest, K. J. (2019). A visual aid and objective rule encompassing the data features of visual analysis. *Behavior Modification*. Advance online publication.

<https://doi.org/10.1177/0145445519854323>.

Manolov, R. (2019). A simulation study on two analytical techniques for alternating treatments designs. *Behavior Modification*, 43(4), 544-563.

<https://doi.org/10.1177/0145445518777875>

Mooney, C. Z. (1997). *Monte Carlo simulation*. Sage.

Ninci, J., Vannest, K. J., Willson, V., & Zhang, N. (2015). Interrater agreement between visual analysts of single-case data: A meta-analysis. *Behavior Modification*, 39(4), 510-541.

<https://doi.org/10.1177/0145445515581327>

Rader, A. E., Rader, K. A., Katz, J. S., & Leaf, J. B. (2021). The progression of experimental design and data analysis in applied behavior analysis. *European Journal of Behavior Analysis*. Advance online publication. <https://doi.org/10.1080/15021149.2021.1932199>

Robinson, D. (2017, October 10). The impressive growth of R. *The Overflow*

<https://stackoverflow.blog/2017/10/10/impressive-growth-r/>

Shadish, W. R., & Sullivan, K. J. (2011). Characteristics of single-case designs used to assess intervention effects in 2008. *Behavior Research Methods*, 43(4), 971-980.

<https://doi.org/10.3758/s13428-011-0111-y>

Sidman, M. (1960). *Tactics of scientific research: Evaluating experimental data in psychology*. Basic Books.

Smith, J. D. (2012). Single-case experimental designs: A systematic review of published research and current standards. *Psychological Methods*, 17(4), 510-550.

<https://doi.org/10.1037/a0029312>

- Wolfe, K., Barton, E. E., & Meadan, H. (2019). Systematic protocols for the visual analysis of single-case research data. *Behavior Analysis in Practice, 12*(2), 491-502.
- Wolfe, K., Seaman, M. A., & Drasgow, E. (2016). Interrater agreement on the visual analysis of individual tiers and functional relations in multiple baseline designs. *Behavior Modification, 40*(6), 852-873. <https://doi.org/10.1177/0145445516644699>
- Wolfe, K., Seaman, M. A., Drasgow, E., & Sherlock, P. (2018). An evaluation of the agreement between the conservative dual-criterion method and expert visual analysis. *Journal of Applied Behavior Analysis, 51*(2), 345-351. <https://doi.org/10.1002/jaba.453>
- Wolfe, K., & Slocum, T. A. (2015). A comparison of two approaches to training visual analysis of AB graphs. *Journal of Applied Behavior Analysis, 48*(2), 472-477.
<https://doi.org/10.1002/jaba.212>

Table 1*Description of Each File in the Repository*

File name	Description
functions.R	Contains all the functions from Steps 1 to 5 to create a time series, to produce an AB data series, to add trend, to graph AB data and to apply the CDC method.
functions_commented.R	Same as function.R, except that it includes comments to explain the role of each line of code.
MonteCarlo.R	Contains all the code for Steps 6 to 10. This code uses the previous functions (i.e., functions.R) to conduct Monte Carlo simulations and analyses.
MonteCarlo_commented.R	Same as MonteCarlo.R, except that it includes comments to explain the role of each line of code.
ABABdata.R	Commented file that contains functions to generate and graph ABAB data series.
MBdata.R	Commented file that contains functions to generate and graph multiple baseline data series.
ATdata.R	Commented file that contains functions to generate and graph alternating-treatment series.
Expert_data.csv	Simulated data for a visual rater

Note. The equivalent files for Python share the same name, but end with the .py (instead of .R). Available at <https://github.com/labrl/monte-carlo-simulations-for-single-case-designs>

Table 2*Accuracy, Type I Error Rate and Power Obtained in Step 9 of the Analyses*

	Measure		
	Accuracy	Type I Error Rate	Power
Conservative Dual-Criteria	.648	.231	.527
Expert data (simulated)	.777	.259	.814

Figure 1*Screenshot from RStudio*

Figure 2

Copy of the Graph Produced in Step 4