ACCP Critical Care PRN: Guide to Performing Basic Statistical Testing using R

INTRODUCTION

Given the increasing focus on evidence-based medicine, knowledge of the appropriate use of statistical tests and appropriate interpretation of results is increasingly necessary. Statistics education might be provided during formal pharmacy education, during pharmacy residency, or through certifications or advanced degrees in the area. Specific to statistics education during PGY1 pharmacy residencies, a recent study found that 25% of respondents' institutions provided no assistance with statistical training and approximately 60% of residency program directors (RPDs) had low confidence in their resident's ability to perform statistical analysis. Pharmacy board certification for any specialty, including critical care requires knowledge of research design, including the appropriate use of statistics and the interpretation of statistical analysis. Therefore, an area of improvement is developing appropriate educational materials for pharmacists to aid in understanding of and performance of statistical analysis.

A charge was created in 2017 for the ACCP Critical Care PRN Research Committee to identify what resources residents and RPDs might find helpful in performing research projects. This resulted in a survey distributed in late 2017 to PGY1 and PGY2 residents, along with PGY2 Critical Care RPDs as a needs assessment. In analyzing the results of this survey, we identified that a statistical guide was one of the most sought resources. A statistical guide could allow for consistency in statistical analysis methodology, with a goal to enhance the quality of published research performed by pharmacy residents.

Several statistical analysis programs exist that can be purchased (e.g., SAS, SPSS, Stata). The statistical analysis programs available to residents will vary by each training site and availability will be dependent on institutional resources. For this reason the content of this guide focuses on the use of R, an open-source and freely available statistical analysis program that is increasingly used in published literature. The guide will focus on a sample data set with step-by-step instructions on which test to perform and an explanation of why. Topics in this guide include: 1. Data management; 2. How to prepare data for analysis; 3. Which statistical test to perform; 4. How to perform the test; 5. How to interpret the findings of the test.

DISCLAIMER

The purpose of this guide is only to serve as guidance for what statistical tests to use and how to perform them using statistics programs. We have created this guide with the intent of disseminating general information regarding statistics and acknowledge there are many caveats to statistics that may not have been included in this guide. Therefore, we recommend using this as general guidance and consulting biostatisticians or peers who are more familiar with your data, when necessary.

The data used in this guide are practice data and should not be used as actual data for any study, as much of this data was created to illustrate certain aspects of statistical testing.

ACKNOWLEDGEMENTS

This guide was created by members of the Critical Care PRN Research Committee with advanced education in statistical analysis in 2018. We would like to acknowledge the following members for their assistance in creating this resource: Melissa Thompson Bastin, PharmD, BCPS; Brittany Bissell, PharmD, BCCCP; Benjamin Hohlfelder, PharmD, BCPS; Joshua DeMott, PharmD, MSc, BCPS, BCCCP; Kendall Gross, PharmD, BCPS, BCCCP; Zachary Smith, PharmD, BCPS, BCCCP; Adrian Wong, PharmD, MPH, BCPS, BCCCP (chair).

R Commander

A more user-friendly interface using R is available, called R Commander. R Commander does not rely as heavily on coding to perform statistical analyses. One of the main differences is that R Commander does not have the ability to perform a Fisher's exact test. Additionally, R Studio is much more flexible as to what it can do. However, R Commander should address most basic analyses that are performed for residency projects. A guide to basics of R Commander is available in a different guide.

Installing R Studio

R Studio is a freely available program which makes standard R (also freely available) easier to use. The assessments for this course can be done in standard R or R Studio, but R Studio is recommended. R Studio can be installed here: https://www.rstudio.com/products/rstudio/download/.
If you do not already have standard R installed on your computer, you must first install it by clicking on link in the red box below in Figure 1 (https://cran.r-project.org/mirrors.html). If you already have standard R, you can click on the appropriate version of R Studio from within the green box depending on your computer's operating system.

Figure 1. Downloading R Studio

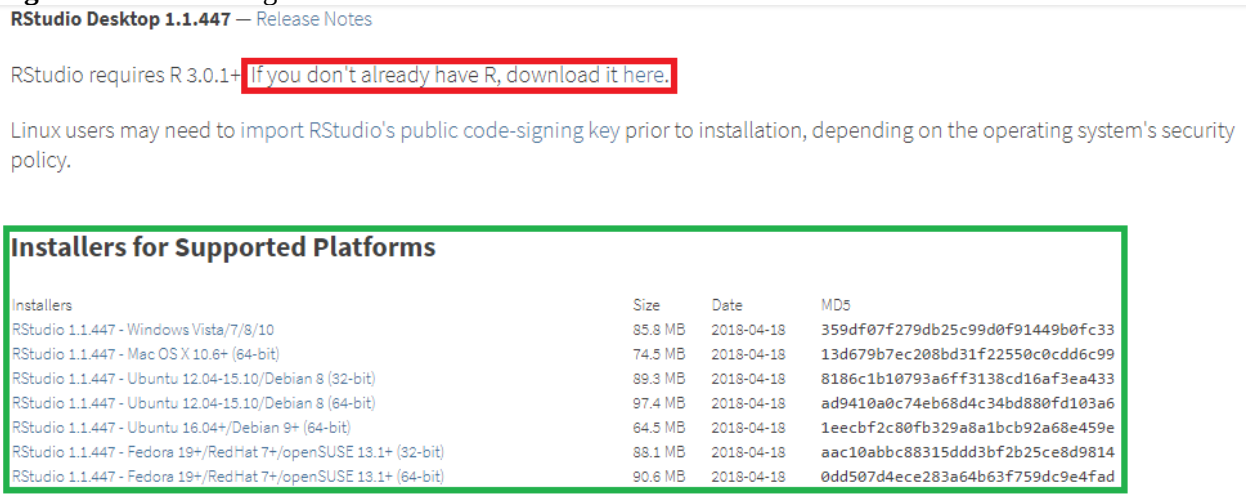

Once you have standard R installed (or if you already had it before starting this course), choose the version of R Studio corresponding to your computer's operating system (listed in the green box in the first image), and follow the installation instructions.

For a brief introduction to the layout of R Studio, please see the following 5 minute YouTube video: https://www.youtube.com/watch?v=5YmcEYTSN7k

There are numerous free online resources for those who would like to learn more. In addition to YouTube, Stack Overflow and online resources from UCLA are particularly helpful (https://stats.idre.ucla.edu/r/).

If you do not see all 4 panes when you open R Studio, click the View dropdown menu, select "Panes", and then choose "Show All Panes". The shortcut for this is Ctrl+Alt+Shift+0 in Windows.

Using R Studio

The screen below in Figure 2 is what you see when you open R Studio. We will now walk through the analysis of a dataset.
Some notes:

- R is case-sensitive so make sure to mirror the case in the variables that you have created in your data collection sheet
- Inputting a ? before a command (e.g., ?wilcox.test) will open up a window that explains the command including information such as what variables can be substituted and what they signify

Figure 2. Screenshot of R Studio opening screen

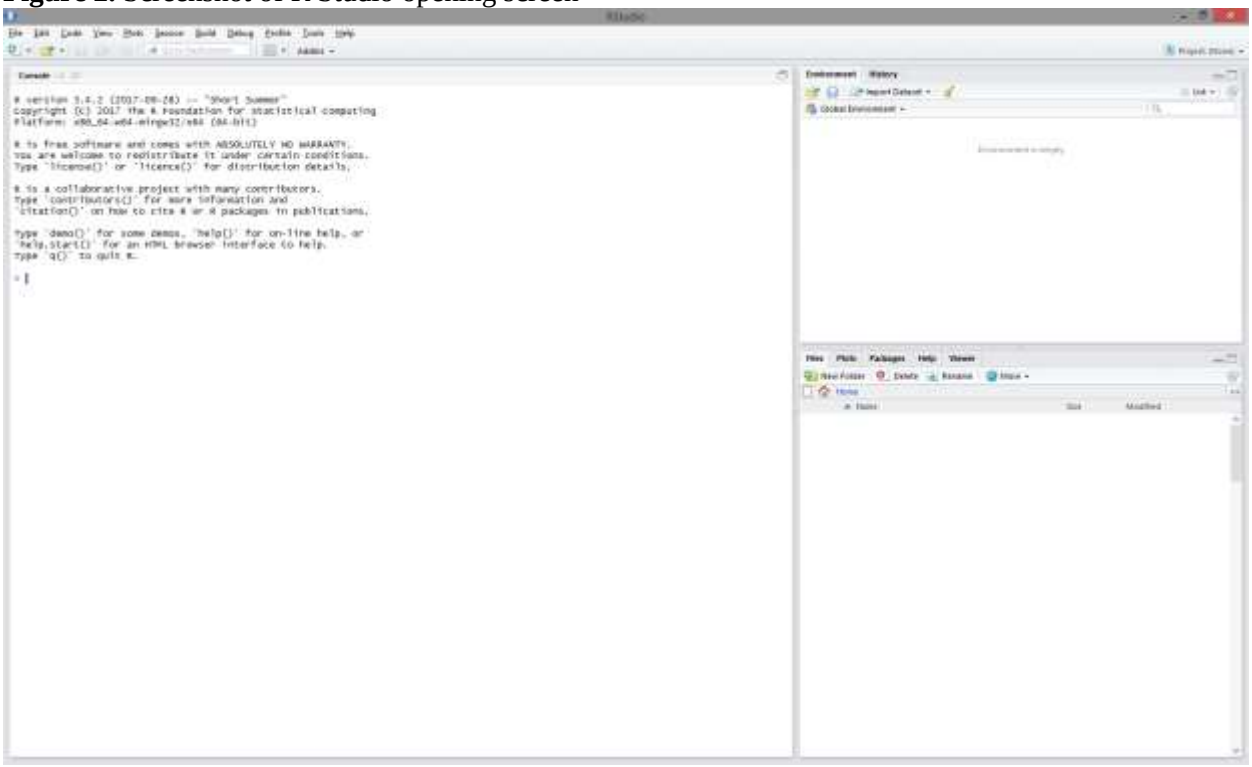

Importing the Data

There are multiple ways to import your data into R Studio. The easiest method is to "Import Dataset", which is accomplished by using the File tab in Figure 3.

Figure 3. Importing a data set into R Studio

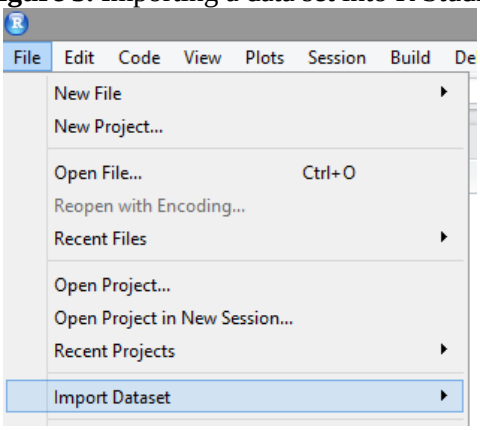

We are interested in importing the dataset "AWSSample" which is saved as an .xlsx file. This dataset includes data evaluating the effect of dexmedetomidine and ketamine on the standard treatment of

patients with alcohol withdrawal syndrome (AWS). Therefore, we will use the option “From Excel…” to import the data as seen below in Figure 4.

Figure 4. Selecting a dataset to input into R Studio – Step 1

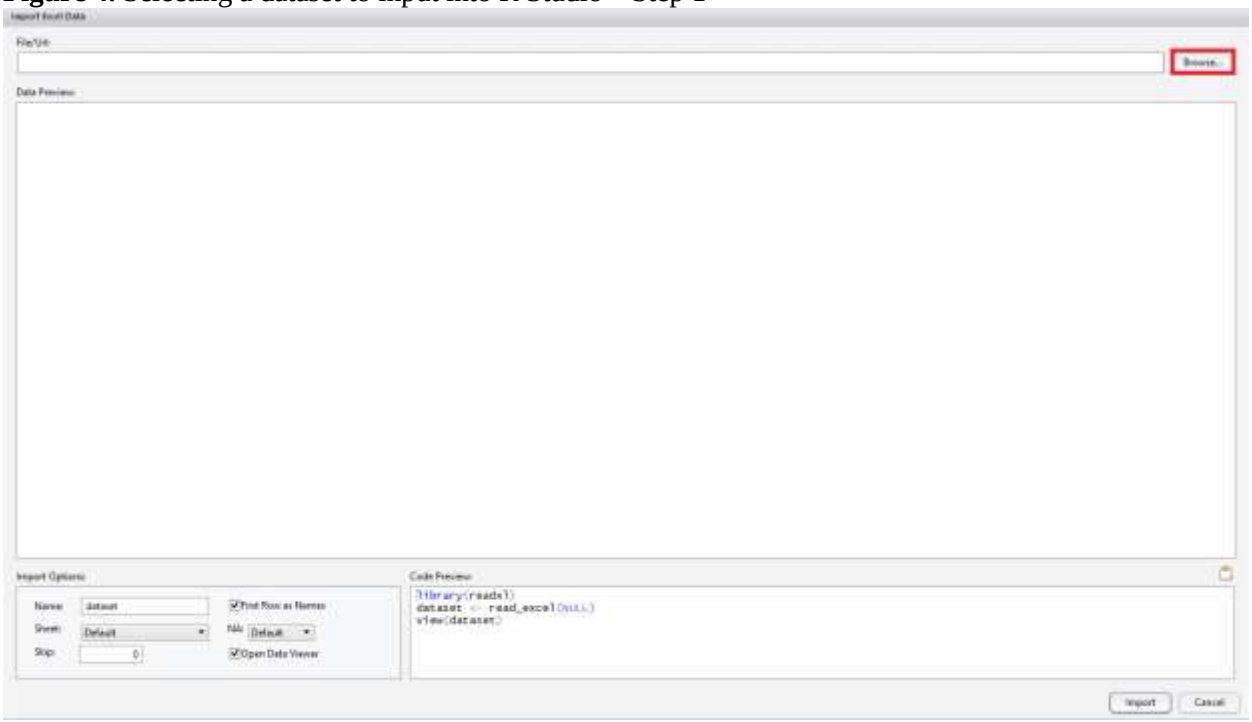

This is the next screen you should encounter. Here, you will select “Browse” to look for the “AWSSample.xlsx” file, as shown here in Figure 5:

Figure 5. Selecting a dataset to input into R Studio – Step 2

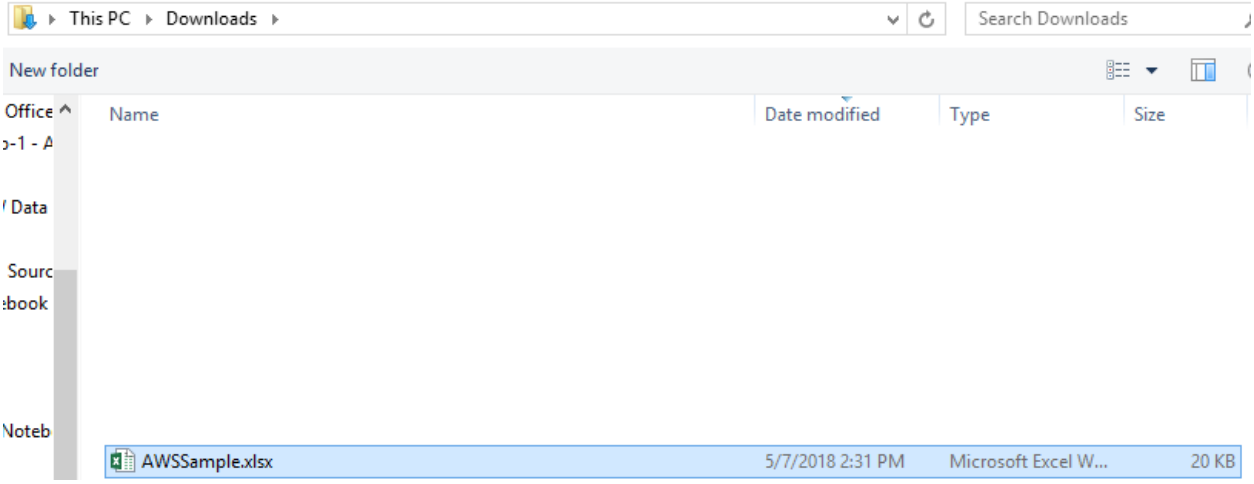

Once selected, you will now see a preview of the dataset (as shown below in Figure 6).

Figure 6. Data set imported into R Studio

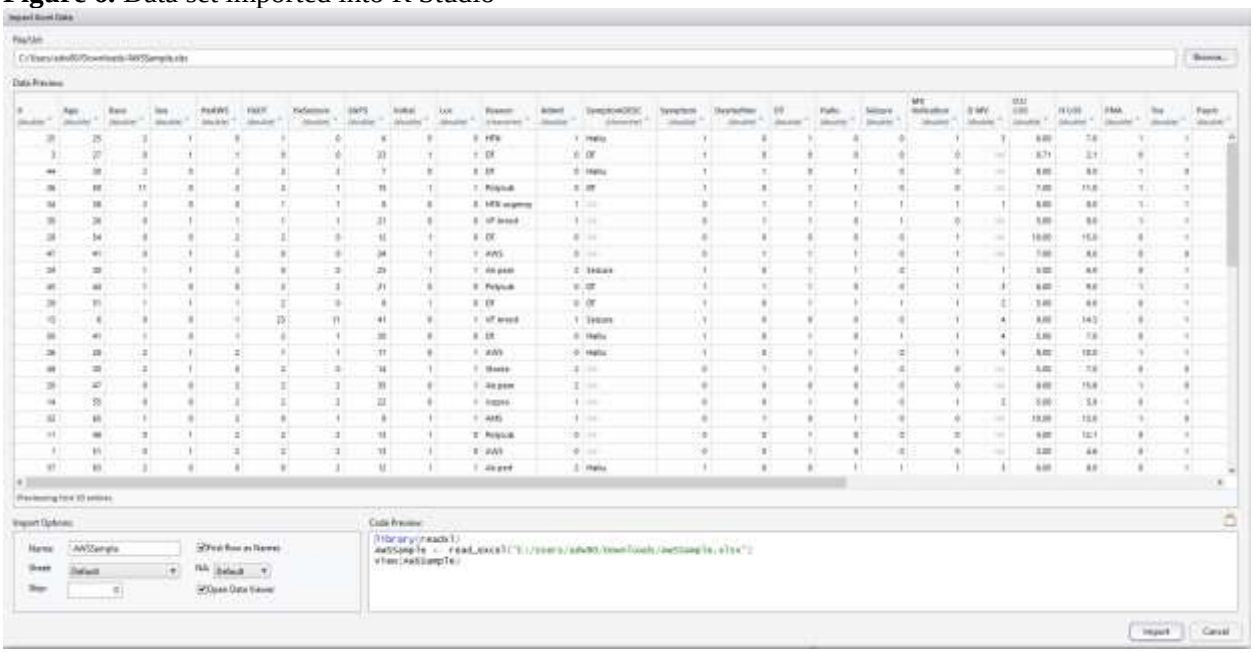

Here, you are able to change the type of data (character, date, numeric) associated with each column (if desired) by clicking on the down arrow by each column in Figure 7. Data types are important as statistical analyses may only be compatible with certain types. You can also choose to "skip" certain parts of your dataset.

Figure 7. Selecting data type for variable column

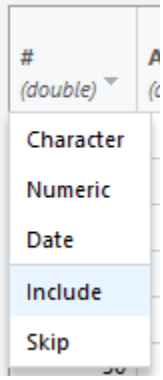

Click "Import" to load this dataset for analysis.

Cleaning the Data

Your screen should now look as below in Figure 8. In the upper-left corner, we see our dataset "AWSSample". The upper-right corner indicates that there are 50 records (i.e., rows) within the dataset, with 35 variables (i.e., columns) associated with these records. The bottom-left corner indicates the code that has been run associated with your analysis. The bottom-right corner has multiple tabs available. "Plots" is where graphics appear while "Packages" are downloadable programs that can help with your analysis.

Figure 8. Display of R Studio once data points are formatted correctly

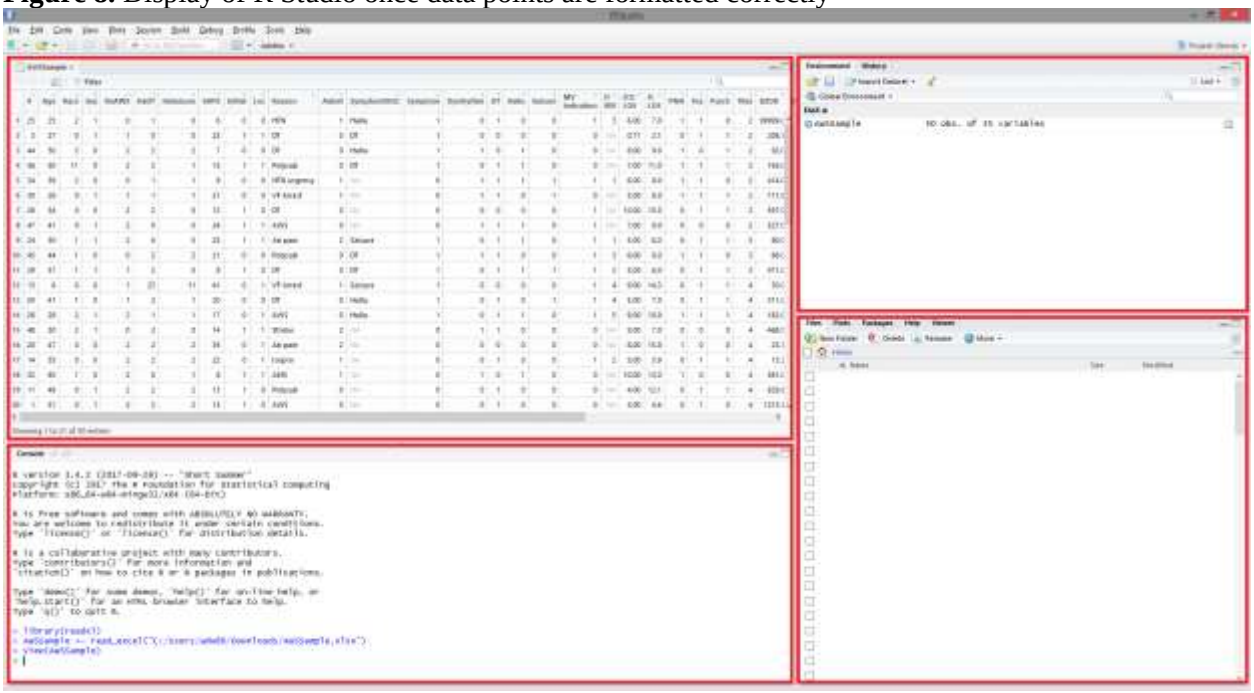

To begin analyzing the data, we will want to know what each of our columns indicates. Table 1 is a description of the data, as well as how this data was coded. This table is known as a data dictionary. It is extremely important to understand what your data means and this key should be one of the first things you create while making a data collection form.

Table 1. Data dictionary

Variable	Description	0	1	2
Age	Patient age on admission (y)			
Race	Patient race	Caucasian	African American	Other
Sex	Patient gender	Female	Male	
HxAWS	Past medical history of AWS (n)	No	Yes	Unknown
HxDT	Past medical history of delirium tremens (n)	No	Yes	Unknown
HxSeizure	Past medical history of seizure (n)	No	Yes	Unknown
SAPS	Simplified Acute Physiology Score			
Initial	Initial presentation of AWS at admission (n)	No	Yes	
Loc	Unit patient was first admitted to	Floor	ICU	
Reason	Reason for patient admission			
Admit	Reason for patient admission, categorized	AWS	Medical	Surgical
SymptomDESC	AWS patient had during admission			
Symptom	AWS symptoms (n)	No	Yes	

Dysrhythm	Dysrhythmia during admission (n)	No	Yes	
DT	Delirium tremens during admission (n)	No	Yes	
Hallu	Hallucinations during admission (n)	No	Yes	
Seizure	Seizure during admission (n)	No	Yes	
MV Indication	Indication for MV (n)	Not AWS	AWS	
D MV	Days of mechanical ventilation			
ICU LOS	ICU length of stay (days)			
H LOS	Hospital length of stay (days)			
PNA	Pneumonia during admission (n)	No	Yes	
Tox	Toxicology consult during admission (n)	No	Yes	
Psych	Psychiatric consult during admission (n)	No	Yes	
TRes	Time to resolution of AWS (days)			
BZDB	Diazepam equivalents prior to study therapy (mg)			
Therapy	Study therapy	Dexmedetomidine	Ketamine	
Time	Time to start of study therapy (hours)			
24B	Diazepam equivalents 12-24 hr prior to study therapy start (mg)			
12B	Diazepam equivalents 0-12 hr prior to study therapy start (mg)			
12A	Diazepam equivalents 12-24 hr after study therapy start (mg)			
24A	Diazepam equivalents 0-12 hr after study therapy start (mg)			
12C	Change in diazepam equivalents from 12 hr before study therapy to 12 hr after study therapy (mg)			
24C	Change in diazepam equivalents from 24 hr before study therapy to 24 hr after study therapy (mg)			

To ensure that our data is correct, we should evaluate our data for correctness. This is also known as "data cleaning". This will be accomplished by running code to evaluate the data. Choose the option to start a new "R Script", as shown below in Figure 9.

Figure 9. Evaluating data for correctness in R Studio

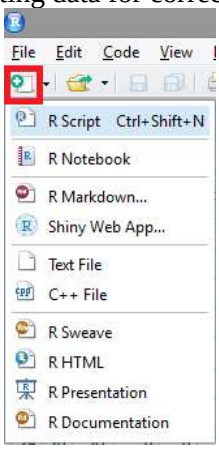

We will first start with the variable "Age", which should only have values ≥18, as only adult patients should have been included in this study. Therefore, any values that are <18 or seem to not be an actual age of a patient, should be evaluated for correction. A simple way to do this would be to use the window of our data view to sort by "Age" as below in Figure 10.

Figure 10. Sorting data in R Studio by data contained within a single column

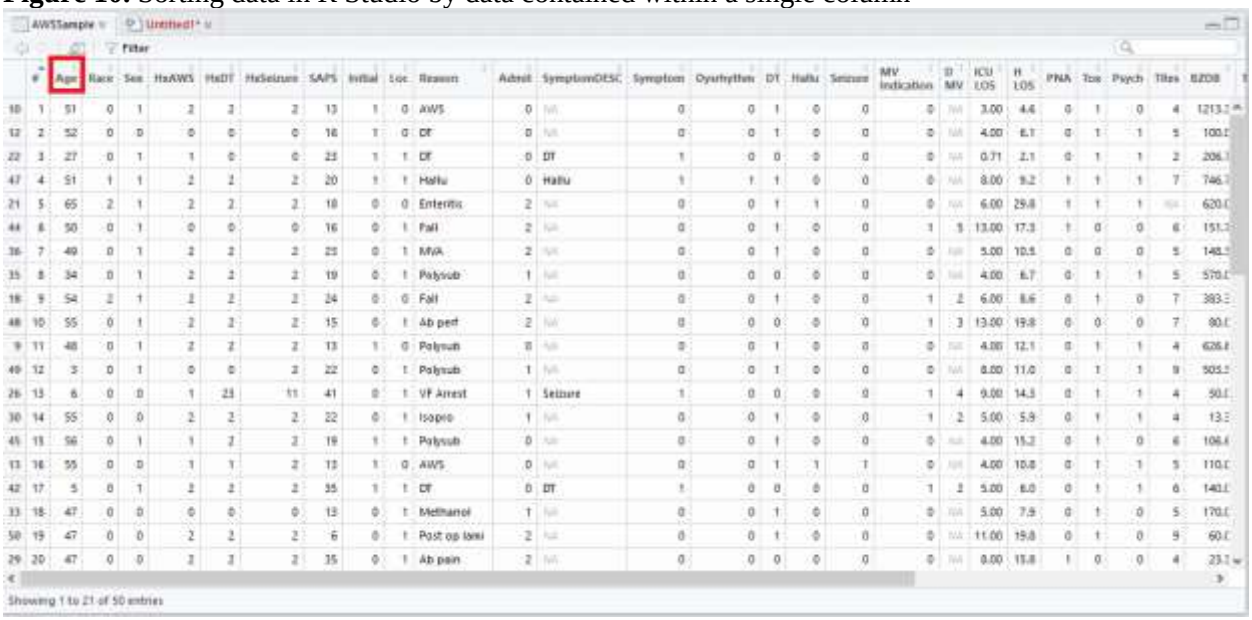

By clicking on the arrow to the right of the variable, we can sort by ascending or descending order as seen in Figure 11.

Figure 11. Sorting by ascending or descending order in R Studio

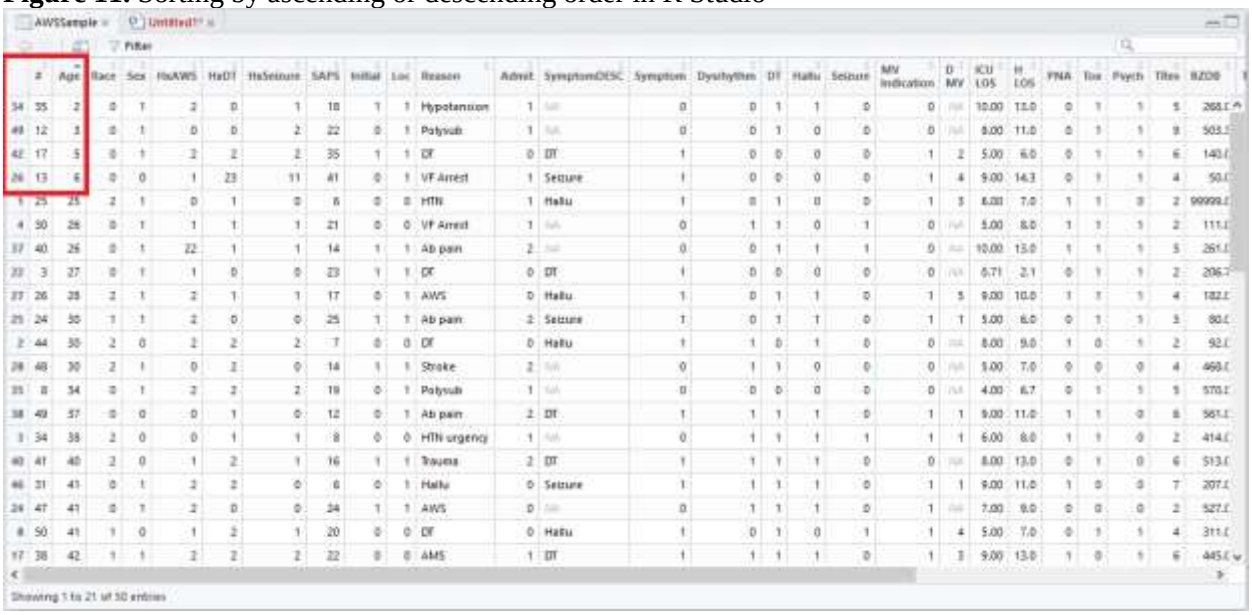

We identify that records 35, 12, 17 and 13 have ages <18. These can be fixed in your original dataset or it can be edited in R Studio itself by installing the "editData" package. To install this, simply input the following into the R script tab (Untitled1 screen in upper left hand corner in Figure 11):

install.packages("editData")

Once installed, type the following (i.e., file name) in the R script tab:

AWSSample

Now, just highlight AWSSample by dragging your mouse over it and then clicking on the Addins tab and choosing "editData" as in Figure 12:

Figure 12. Editing dataset in R Studio

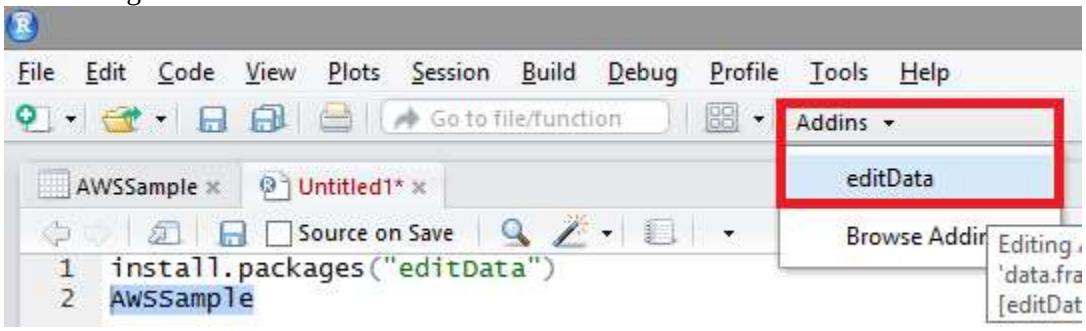

The following screen then pops up as depicted in Figure 13:

Figure 13. Initial screen upon selecting editData

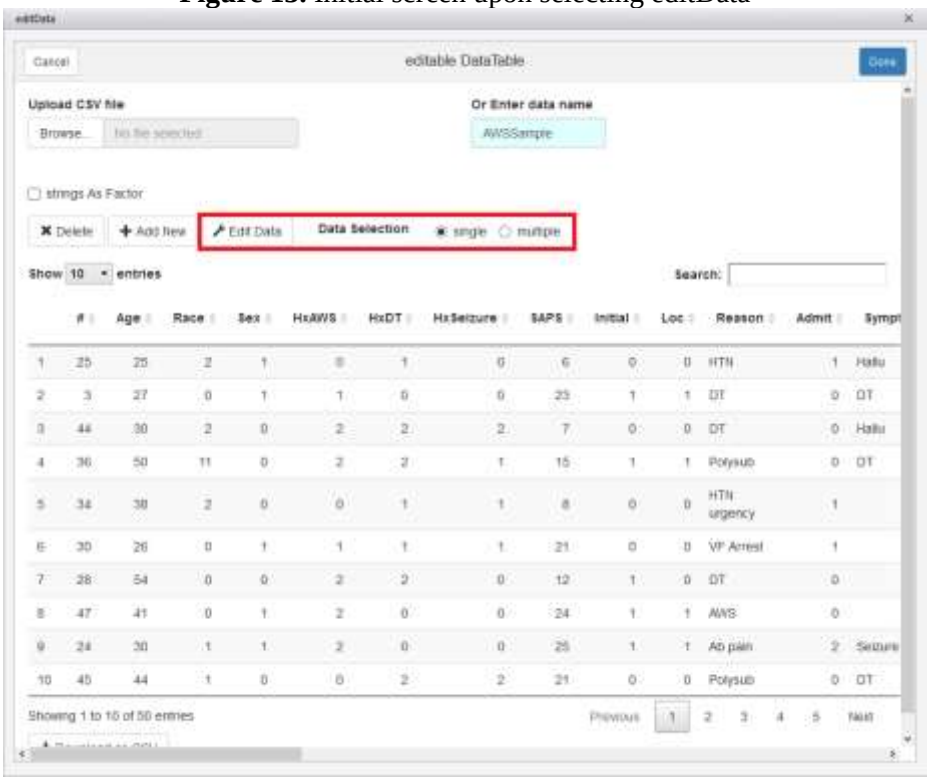

You can “Edit Data” in single or multiple rows by selecting the options within the red box above in Figure 13. We are interested in the records mentioned above, which correspond to the rows identified below in Figure 14:

Figure 14. Selecting data points to edit in R Studio

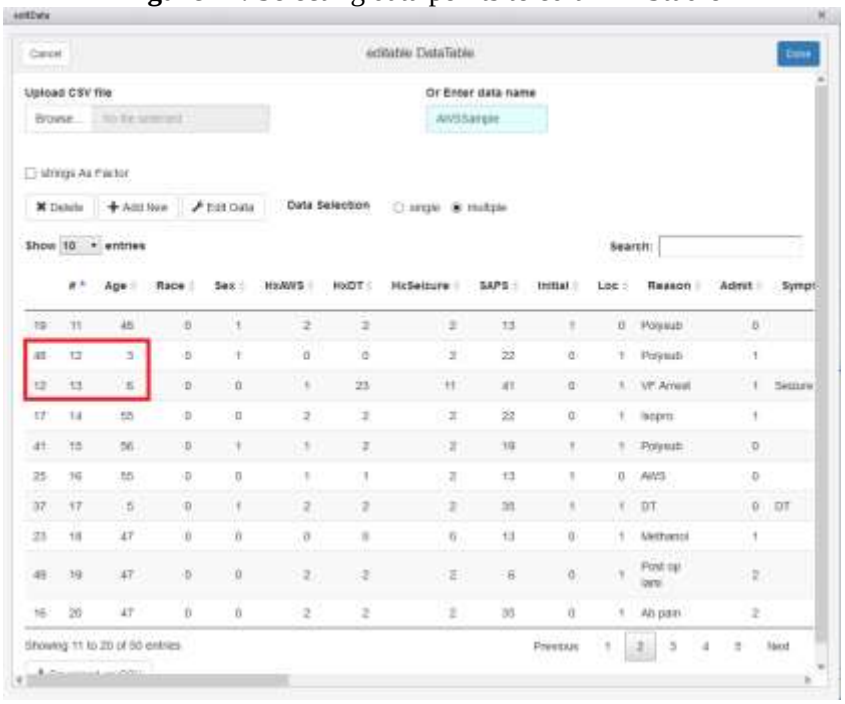

For record #13, which is row #12 in this dataset, we can then edit the data (by clicking and filling out cells) as necessary on the next screen (Figure 15).

Figure 15. Editing single data points in R Studio

Once the data editing is complete, scroll to the bottom of this screen depicted in Figure 15 to update this record.

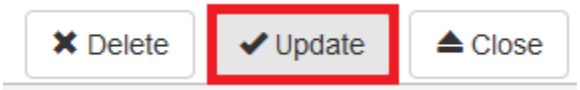

Another method that can quickly identify if any data was input incorrectly is by using the command "table". In using this to evaluate our "Race" variable:

```
table(AWSSample$Race)
```

Format: table(filename$variable)

This notation indicates that you are using the AWSSample data, given the variable Race. This is the output of this command:

```
 0  1  2 11 22
29  9 10  1  1
```

You can see that there are 29 records with a value of 0 (i.e., Caucasian), nine with a value of 1 (i.e., African American), and 10 with a value of 2 (i.e., Other). You can see that there are two records that have incorrect data (one with a value of 11, one with a value of 22).

By going through all the data variables in a systematic manner and correcting the data as necessary, we now have a dataset to evaluate (AWSFinal). The corrections that should be made are the following:

Variable	**Comments**
Sex	#21: Input as 11, should be 1
HxAWS	#40: Input as 22, should be 2
HxDT	#13: Input as 23, should be 2
HxSeizure	#13: Input as 11, should be 1
TRes	#5: Missing data, should be 3
BZDB	#25: Input as 99999, should be 99

You can then save your updated sheet by doing the following:

```
install.packages("xlsx")
library(xlsx)
write.xlsx(AWSSample, 'C:/Users/m0338687/Documents/AWSUpdated.xlsx')
```

Format: write.xlsx(data, directory and file name for new data)

Analyzing the Data

We will now be using the "AWSFinal" dataset for our analysis. To perform any type of data analysis mentioned in this primer the user will need to type the commands given in the boxes. First, we will target how to characterize our two study groups. The most time efficient way to do this is to install the "psych" package:

```
install.packages("psych", dependencies=TRUE)
```

Including the "dependencies" ensures that all aspects of this specific package are included. This package allows us to easily obtain pertinent descriptive statistics. Input the following:

```
library(psych)
describe(AWSFinal)
```

Format: describe(filename)

The output of this command is below in Figure 16, with a description of the terms in Table 2.

Figure 16. Output from command for data set

	vars	n	mean	sd	median	trimmed	mad	min	max	range	skew	kurtosis	se
#	1	50	25.50	14.58	25.50	25.50	18.53	1.00	50.0	49.00	0.00	-1.27	2.06
Age	2	50	46.32	12.26	48.50	46.33	11.12	25.00	69.0	44.00	-0.15	-0.87	1.73
Race	3	50	0.64	0.83	0.00	0.55	0.00	0.00	2.0	2.00	0.73	-1.18	0.12
Sex	4	50	0.56	0.50	1.00	0.57	0.00	0.00	1.0	1.00	-0.23	-1.98	0.07
HxAWS	5	50	1.32	0.84	2.00	1.40	0.00	0.00	2.0	2.00	-0.64	-1.32	0.12
HxDT	6	50	1.28	0.86	2.00	1.35	0.00	0.00	2.0	2.00	-0.55	-1.44	0.12
HxSeizure	7	50	1.20	0.86	1.00	1.25	1.48	0.00	2.0	2.00	-0.38	-1.56	0.12
SAPS	8	50	16.88	7.42	16.00	16.27	5.93	6.00	41.0	35.00	0.96	1.31	1.05
Initial	9	50	0.44	0.50	0.00	0.42	0.00	0.00	1.0	1.00	0.23	-1.98	0.07
Loc	10	50	0.58	0.50	1.00	0.60	0.00	0.00	1.0	1.00	-0.31	-1.94	0.07
Reason*	11	50	NaN	NA	NA	NaN	NA	Inf	-Inf	-Inf	NA	NA	NA
Admit	12	50	0.98	0.87	1.00	0.98	1.48	0.00	2.0	2.00	0.04	-1.70	0.12
SymptomDESC*	13	18	NaN	NA	NA	NaN	NA	Inf	-Inf	-Inf	NA	NA	NA
Symptom	14	50	0.36	0.48	0.00	0.32	0.00	0.00	1.0	1.00	0.57	-1.71	0.07
Dysrhythm	15	50	0.10	0.30	0.00	0.00	0.00	0.00	1.0	1.00	2.59	4.79	0.04
DT	16	50	0.18	0.39	0.00	0.10	0.00	0.00	1.0	1.00	1.62	0.63	0.05
Hallu	17	50	0.10	0.30	0.00	0.00	0.00	0.00	1.0	1.00	2.59	4.79	0.04
Seizure	18	50	0.12	0.33	0.00	0.02	0.00	0.00	1.0	1.00	2.27	3.21	0.05
MV Indication	19	50	0.40	0.49	0.00	0.38	0.00	0.00	1.0	1.00	0.40	-1.88	0.07
D MV	20	32	3.00	1.57	3.00	2.92	1.48	1.00	7.0	6.00	0.49	-0.61	0.28
ICU LOS	21	50	6.87	2.54	6.00	6.78	2.97	0.71	13.0	12.29	0.30	-0.21	0.36
H LOS	22	50	11.14	4.95	10.30	10.57	4.00	2.10	29.8	27.70	1.30	2.60	0.70
PNA	23	50	0.42	0.50	0.00	0.40	0.00	0.00	1.0	1.00	0.31	-1.94	0.07
Tox	24	50	0.76	0.43	1.00	0.82	0.00	0.00	1.0	1.00	-1.18	-0.62	0.06
Psych	25	50	0.54	0.50	1.00	0.55	0.00	0.00	1.0	1.00	-0.16	-2.01	0.07
TRes	26	50	4.80	1.84	5.00	4.75	1.48	2.00	9.0	7.00	0.15	-0.56	0.26
BZDB	27	50	327.24	266.36	254.50	297.45	258.71	13.30	1213.3	1200.00	1.08	1.04	37.67
Therapy	28	50	0.50	0.51	0.50	0.50	0.74	0.00	1.0	1.00	0.00	-2.04	0.07
Time	29	50	27.66	13.32	27.50	28.40	17.79	0.00	47.0	47.00	-0.32	-1.01	1.88
24B	30	50	240.34	195.51	177.00	218.95	188.29	13.30	1000.0	986.70	1.29	2.40	27.65
12B	31	50	212.83	159.75	190.00	195.54	171.24	13.30	810.0	796.70	1.19	1.93	22.59
12A	32	50	128.00	71.44	127.50	126.40	69.68	0.00	366.7	366.70	0.48	1.06	10.10
24A	33	50	154.59	110.82	137.50	140.36	74.13	0.00	560.0	560.00	1.47	2.68	15.67
12C	34	50	-84.83	173.34	-56.65	-71.46	114.90	-810.00	306.7	1116.70	-1.40	4.50	24.51
24C	35	50	-85.75	209.14	-25.00	-70.57	160.86	-1000.00	340.0	1340.00	-1.56	5.25	29.58

Table 2. Terminology from Figure 16 data

Value	Description
N	Number of records of variable included
Mean	Mean of the variable
SD	Standard deviation of the variable
Median	Median of the variable
Trimmed	Trimmed mean of the variable
MAD	Median absolute deviation
Min	Minimum value of the variable
Max	Maximum value of the variable
Range	Range of the variable
Skew	Skew of the variable
Kurtosis	Kurtosis of the variable
SE	Standard error of the variable

There are some variables that are not pertinent here due to the data type. These would be the categorical variables (e.g., race, sex, Hx). The above command is most helpful for continuous variables. To only include certain variables (vars 2 [Age] and vars 8 [SAPS]) in the describe command:

```
describe(AWSFinal[,c(2,8)])
```

Format: describe(filename[,c(var#)])
(where var# is the number associated with the vars in the above results)

```
      vars  n  mean    sd median trimmed   mad min max range  skew kurtosis   se
Age      1 50 46.32 12.26   48.5   46.33 11.12  25  69    44 -0.15    -0.87 1.73
SAPS     2 50 16.88  7.42   16.0   16.27  5.93   6  41    35  0.96     1.31 1.05
```

Additionally, you may have noticed in the initial exercise that some patients had data missing, which was accounted for in this analysis. If you only wanted to include records with fully complete data, you could use this command:

```
describe(AWSFinal, na.rm=FALSE)
```

Format: describe(filename, na.rm=FALSE)
(where na.rm means "not available, remove")

Since we are interested in comparing the two groups (Therapy = 0,1), it would be best to obtain descriptive statistics by the type of "Therapy":

```
describe.by(AWSFinal,AWSFinal$Therapy, mat=TRUE)
```

Format: describe.by(filename, filename$variable, mat=TRUE)
(where mat means "matrix", instead of a list)

If mat=TRUE is not included, the two groups are evaluated separately. Including mat=TRUE makes the statistics easier to read. Our goal is to fill out this table (Table 3) for baseline characteristics to allow the reader to understand potential differences between these two groups.

Table 3. Patient demographics by treatment

Variable	Dexmedetomidine	Ketamine	p-value
Age			
Race			
Sex			
HxAWS			
HxDT			
HxSeizure			
SAPS			
Initial			
Location			
Admit			

The following is the table that may be presented in a manuscript (Table 4):

Table 4. Patient demographics by treatment (for manuscript)

Variable	Dexmedetomidine	Ketamine	p-value
Age, y			
Race, n (%)			
Caucasian			
African American			
Other			
Male, n (%)			
History of AWS, n (%)			
History of DT, n (%)			
History of AWS Seizure, n (%)			
SAPS			

Initial AWS, n (%)			
Initial admit to ICU, n (%)			
Admission type, n (%)			
AWS			
Medical			
Surgical			

AWS = alcohol withdrawal syndrome; DT = delirium tremens; ICU = intensive care unit; SAPS = Simplified Acute Physiology Score

Next we will fill out the table based on the analysis we have already completed, which is dependent on the type of distribution of the data (parametric, non-parametric). Using the "describe" command, we are able to obtain four potential ways to determine the type of distribution (Table 5):

Table 5. Methods to evaluate data distribution

Method	Comments
Evaluation of mean and SD	- A parametric distribution is more likely to have a smaller SD (due to larger sample size) - If the SD is large relative to the mean, it is unlikely to be a parametric distribution
Evaluation of mean compared to median	- The mean and median should be relatively close to each other in value in a parametric distribution
Evaluation of skewness	- Measures asymmetry of the distribution - Values close to 0 indicate a parametric distribution - A value <0 indicates a "left skew" - A value >0 indicates a "right skew"
Evaluation of kurtosis	- Measures "tailedness" of the distribution - Values close to 0 indicate a parametric distribution - A value <0 indicates the tails of the distribution are longer than normal - A value >0 indicates the tails of the distribution are shorter than normal

Based on the small sample size (n=25 in each group), one would expect the distributions to be non-parametric. One additional way to evaluate for the distribution is a histogram (as seen in Figure 17).

```
hist (AWSSample$Race)
```

Figure 17. Histogram of subject race

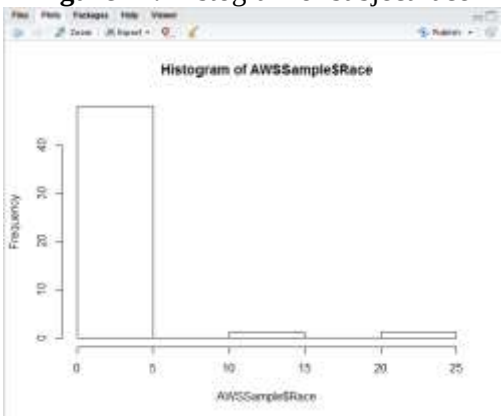

Based on the values for Age and SAPS, median (IQR) will be used to describe these populations. To only choose the two continuous variables (Age and SAPS) and their IQR, use the following:

```
describe.by(AWSFinal[,c(2,8)],AWSFinal$Therapy, quant=c(0.25,0.75), mat=TRUE)
```

Format: describe.by(filename, filename$variable, mat=TRUE)
(where mat means "matrix", instead of a list)

```
      item group1 vars  n  mean        sd median  trimmed     mad min max range       skew   kurtosis       se Q0.25 Q0.75
Age1     1      0    1 25 44.92 13.972950     42 44.52381 17.7912  25  69    44  0.2972681 -1.1452216 2.794590    37    54
Age2     2      1    1 25 47.72 10.378343     51 48.38095  5.9304  25  65    40 -0.8918903 -0.2010008 2.075669    47    54
SAPS1    3      0    2 25 14.20  5.283622     14 14.09524  5.9304   6  24    18  0.1591310 -1.2869244 1.056724    11    18
SAPS2    4      1    2 25 19.56  8.332067     19 19.09524  5.9304   6  41    35  0.7597306  0.3554228 1.666413    13    23
```

Table 4a. Patient demographics by treatment (for manuscript)

Variable	Dexmedetomidine	Ketamine	p-value
Age, y	42 (37, 54)	51 (47, 54)	
Race, n (%)			
Caucasian	9 (36.0)	20 (80.0)	
African American	8 (32.0)	2 (8.0)	
Other	8 (32.0)	3 (12.0)	
Male, n (%)	13 (52.0)	16 (64.0)	
History of AWS, n (%)	6 (24.0)	4 (16.0)	
History of DT, n (%)	8 (32.0)	2 (8.0)	
History of AWS Seizure, n (%)	11 (44.0)	1 (4.0)	
SAPS	14 (11, 18)	19 (13, 23)	
Initial AWS, n (%)	13 (52.0)	9 (36.0)	
Initial admit to ICU, n (%)	14 (56.0)	15 (60.0)	
Admission type, n (%)			
AWS	11 (44.0)	8 (32.0)	
Medical	5 (20.0)	8 (32.0)	
Surgical	9 (36.0)	9 (36.0)	

Values in median (IQR) unless otherwise specified
AWS = alcohol withdrawal syndrome; DT = delirium tremens; ICU = intensive care unit; SAPS = Simplified Acute Physiology Score

To obtain the differences between these two groups on select variables, we will use the "chisq.test", "fisher.test", and "wilcox.test" commands.

To determine the differences by % Male between the dexmedetomidine and ketamine groups, we can use the following:

```
chisq.test(AWSFinal$Therapy, AWSFinal$Sex)
```

Format: chisq.test(filename$group, filenamel$var)

```
        Pearson's Chi-squared test with Yates' continuity correction

data:  AWSFinal$Therapy and AWSFinal$Sex
X-squared = 0.081169, df = 1, p-value = 0.7757
```

The output above indicates there is a p-value of 0.776. This is the probability (assuming the null hypothesis is true [i.e., there is no difference in the proportion of males between the dexmedetomidine and ketamine groups]) of observing a more extreme test statistic in the direction of the alternative hypothesis than the one observed.

We can then do this for the rest of the categorical variables (Table 4b):

Table 4b. Patient demographics by treatment (for manuscript)

Variable	Dexmedetomidine	Ketamine	p-value
Age, y	42 (37, 54)	51 (47, 54)	
Race, n (%)			
Caucasian	9 (36.0)	20 (80.0)	0.007
African American	8 (32.0)	2 (8.0)	
Other	8 (32.0)	3 (12.0)	
Male, n (%)	13 (52.0)	16 (64.0)	0.776
History of AWS, n (%)*	6 (24.0)	4 (16.0)	0.521
History of DT, n (%)*	8 (32.0)	2 (8.0)	0.100
History of AWS Seizure, n (%)*	11 (44.0)	1 (4.0)	<0.001
SAPS	14 (11, 18)	19 (13, 23)	
Initial AWS, n (%)	13 (52.0)	9 (36.0)	0.393
Initial admit to ICU, n (%)	14 (56.0)	15 (60.0)	1.000
Admission type, n (%)			
AWS	11 (44.0)	8 (32.0)	0.558
Medical	5 (20.0)	8 (32.0)	
Surgical	9 (36.0)	9 (36.0)	

Values in median (IQR) unless otherwise specified
* Includes "No" and "Unknown" as one group for statistical analysis
AWS = alcohol withdrawal syndrome; DT = delirium tremens; ICU = intensive care unit; SAPS = Simplified Acute Physiology Score

One thing to note is that since there were more than two levels of the variable for Race and Admit, we only have the p-value for the aggregate comparison (i.e., we do not know the differences by each level of the variable).

Additionally, we may want to consider using the "fisher.test" command for some of the comparisons we have already done (depending on your thoughts on chi-square vs. Fisher's exact test; more details in "Appropriate Statistical Testing" guide). The most significant ones would be the ones with a cell value of <5 (e.g., HxAWS, HxDT, HxSeizure).The command is similar to the one that we have already used for the chisq.test:

```
fisher.test(AWSFinal$Therapy, AWSFinal$HxAWS)
```

Format: fisher.test(filename$group, filenamel$var)

You can see from the output below that the p-value is relatively similar between these two tests:

```
> chisq.test(AWSFinal$Therapy, AWSFinal$HxAWS)

        Pearson's Chi-squared test

data:  AWSFinal$Therapy and AWSFinal$HxAWS
X-squared = 1.3048, df = 2, p-value = 0.5208

> fisher.test(AWSFinal$Therapy, AWSFinal$HxAWS)

        Fisher's Exact Test for Count Data

data:  AWSFinal$Therapy and AWSFinal$HxAWS
p-value = 0.5874
alternative hypothesis: two.sided
```

Another way to quickly evaluate a chi-square or Fisher's test is to input the numbers in each cell of the contingency table:

x <- matrix(c(13, 16, 12, 9), byrow=TRUE, 2, 2)	Creates 2x2 table (byrow=TRUE fills out matrix by rows; "2, 2" indicates 2x2 table)
colnames(x) <- c("Dex", "Ket")	Adds column names
rownames(x) <- c("M", "F")	Adds row names
prop.table(x, 2)	Demonstrates % in each cell (two decimal points)
chisq.test(x)	Chi-square test of "x"
fisher.test(x)	Fisher's test of "x"

The updated table now looks like this (using Fisher's test) (Table 4c):

Table 4c. Patient demographics by treatment (for manuscript)

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value
Age, y	42 (37, 54)	51 (47, 54)	
Race, n (%)			
Caucasian	9 (36.0)	20 (80.0)	0.007
African American	8 (32.0)	2 (8.0)	
Other	8 (32.0)	3 (12.0)	
Male, n (%)	13 (52.0)	16 (64.0)	0.776
History of AWS, n (%)	6 (24.0)	4 (16.0)	0.587
History of DT, n (%)	8 (32.0)	2 (8.0)	0.124
History of AWS Seizure, n (%)	11 (44.0)	1 (4.0)	<0.001
SAPS	14 (11, 18)	19 (13, 23)	
Initial AWS, n (%)	13 (52.0)	9 (36.0)	0.393
Initial admit to ICU, n (%)	14 (56.0)	15 (60.0)	1.000
Admission type, n (%)			
AWS	11 (44.0)	8 (32.0)	0.558
Medical	5 (20.0)	8 (32.0)	
Surgical	9 (36.0)	9 (36.0)	

Values in median (IQR) unless otherwise specified

***** Includes "No" and "Unknown" as one group for statistical analysis

AWS = alcohol withdrawal syndrome; DT = delirium tremens; ICU = intensive care unit; SAPS = Simplified Acute Physiology Score

To evaluate the differences between the two therapies for the continuous variables, we can use the "wilcox.test" command, which is the Mann-Whitney U test. This test was chosen because the two groups are independent (if the two values were dependent, then a Wilcoxon Rank Sum test would be used [will be covered shortly]). This is most appropriate given that the distribution of our data is nonparametric:

```
wilcox.test(AWSFinal$Age ~ AWSFinal$Therapy, paired=FALSE)
```

Format: wilcox.test(filename$var ~ filenamel$group)
(where var is the continuous variable of interest and group is the variable with two levels of interest; paired=FALSE indicates that we want the Mann-Whitney U test and not the Wilcoxon Rank Sum test [paired=TRUE])

You may note the error message that pops up here:

```
        Wilcoxon rank sum test with continuity correction

data:  AWSFinal$Age by AWSFinal$Therapy
W = 250.5, p-value = 0.2323
alternative hypothesis: true location shift is not equal to 0

Warning message:
In wilcox.test.default(x = c(41, 30, 41, 65, 63, 40, 58, 38, 43,  :
  cannot compute exact p-value with ties
```

This error message is because there were the same exact ages for nine records (which is a significant portion of our study sample) in each of the two treatments. Methods to troubleshoot (e.g., jittering) will not be described in this guide.

This allows us to fill-out the remainder of the table (Table 4d).

Table 4d. Patient demographics by treatment (for manuscript)

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value
Age, y	42 (37, 54)	51 (47, 54)	0.232
Race, n (%)			
Caucasian	9 (36.0)	20 (80.0)	0.007
African American	8 (32.0)	2 (8.0)	
Other	8 (32.0)	3 (12.0)	
Male, n (%)	13 (52.0)	16 (64.0)	0.567
History of AWS, n (%)	6 (24.0)	4 (16.0)	0.587
History of DT, n (%)	8 (32.0)	2 (8.0)	0.124
History of AWS Seizure, n (%)	11 (44.0)	1 (4.0)	<0.001
SAPS	14 (11, 18)	19 (13, 23)	0.012
Initial AWS, n (%)	13 (52.0)	9 (36.0)	0.393
Initial admit to ICU, n (%)	14 (56.0)	15 (60.0)	1.000
Admission type, n (%)			
AWS	11 (44.0)	8 (32.0)	0.558
Medical	5 (20.0)	8 (32.0)	
Surgical	9 (36.0)	9 (36.0)	

Values in median (IQR) unless otherwise specified
* Includes “No” and “Unknown” as one group for statistical analysis
AWS = alcohol withdrawal syndrome; DT = delirium tremens; ICU = intensive care unit; SAPS = Simplified Acute Physiology Score

We have now completed the evaluation of the patient demographics for this study. We will now evaluate the outcomes of therapy.

The following tables illustrate potential ways we could present our data. The first table illustrates symptoms associated with AWS therapy (Table 5), while the second table illustrates the patient outcomes (Table 6).

Table 5. Patient characteristics for AWS

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value
Mechanical ventilation for AWS, n (%)			
AWS Symptom, n (%)			
Delirium tremens			
Dysrhythmia			
Hallucination			
Seizure			
Time to resolution of AWS, d			

AWS = alcohol withdrawal syndrome

Table 6. Patient outcomes associated with AWS

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value
Toxicology consultation, n (%)			
Psychiatry consultation, n (%)			
Total BZD dose prior to therapy, mg			
12 hour change in BZD requirements, mg			
24 hour change in BZD requirements, mg			
Length of mechanical ventilation, d			
Pneumonia during admission, n (%)			
ICU length of stay, d			
Hospital length of stay, d			

BZD = benzodiazepine; ICU = intensive care unit

We can fill out the majority of these with the commands we have already used (Table 5a and 6a).

Table 5a. Patient characteristics for AWS

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value
Mechanical ventilation for AWS, n (%)	12 (48.0)	8 (32.0)	0.387
AWS Symptom, n (%)	11 (44.0)	7 (28.0)	0.377
Delirium tremens	4 (36.4)	1 (14.3)	0.138
Dysrhythmia	7 (63.6)	2 (28.6)	0.348
Hallucination	1 (9.1)	4 (57.1)	0.348
Seizure	4 (36.4)	2 (28.6)	0.667
Time to resolution of AWS, h	29 (25.0, 41.0)	27 (14.0, 39.0)	0.122

AWS = alcohol withdrawal syndrome

Table 6a. Patient outcomes associated with AWS

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value

Toxicology consultation, n (%)	17 (68.0)	21 (84.0)	0.321
Psychiatry consultation, n (%)	14 (56.0)	13 (52.0)	1.000
Total BZD dose prior to therapy, mg	311.0 (165.0, 513.0)	151.7 (99.0, 513.3)	0.352
12 hour change in BZD requirements, mg	-97.0 (-202.0, +1.0)	-13.03 (-100.0, +21.7)	
24 hour change in BZD requirements, mg	-112.0 (-227.0, +23.0)	-13.3 (-86.7, +15.6)	
Length of mechanical ventilation, d	3.0 (2.0, 4.0)	2.0 (2.0, 4.0)	0.714
Pneumonia during admission, n (%)	16 (64.0)	5 (20.0)	0.004
ICU length of stay, d	8.0 (6.0, 9.0)	5.0 (4.0, 8.0)	0.019
Hospital length of stay, d	10.0 (9.0, 13.0)	10.5 (6.7, 15.8)	0.899

BZD = benzodiazepine; ICU = intensive care unit

You will note that two of the variables do not have a p-value associated with them. This is because the data is paired (i.e., dependent) as the change in BZD requirements is dependent on the amount the patient required before therapy and the amount the patient required after therapy. Therefore, a Wilcoxon Rank Sum test would be more appropriate:

```
wilcox.test(AWSFinal$`12C` ~ AWSFinal$Therapy, paired=TRUE)
```

Format: wilcox.test(filename$var ~ filenamel$group)

```
> wilcox.test(AWSFinal$`12C` ~ AWSFinal$Therapy, paired=FALSE)

        Wilcoxon rank sum test with continuity correction

data:  AWSFinal$`12C` by AWSFinal$Therapy
W = 221, p-value = 0.07743
alternative hypothesis: true location shift is not equal to 0

Warning message:
In wilcox.test.default(x = c(50, -248, -345, -77, -115, -75, 6,  :
  cannot compute exact p-value with ties
> wilcox.test(AWSFinal$`12C` ~ AWSFinal$Therapy, paired=TRUE)

        Wilcoxon signed rank test

data:  AWSFinal$`12C` by AWSFinal$Therapy
V = 98, p-value = 0.08514
alternative hypothesis: true location shift is not equal to 0
```

You can see that there is a slight difference in the p-value between the two tests (Table 6b).

Table 6b. Patient outcomes associated with AWS

Variable	Dexmedetomidine (n=25)	Ketamine (n=25)	p-value
Toxicology consultation, n (%)	17 (68.0)	21 (84.0)	0.321
Psychiatry consultation, n (%)	14 (56.0)	13 (52.0)	1.000

Total BZD dose prior to therapy, mg	311.0 (165.0, 513.0)	151.7 (99.0, 513.3)	0.352
12 hour change in BZD requirements, mg	-97.0 (-202.0, +1.0)	-13.03 (-100.0, +21.7)	0.085
24 hour change in BZD requirements, mg	-112.0 (-227.0, +23.0)	-13.3 (-86.7, +15.6)	0.149
Length of mechanical ventilation, d	3.0 (2.0, 4.0)	2.0 (2.0, 4.0)	0.714
Pneumonia during admission, n (%)	16 (64.0)	5 (20.0)	0.004
ICU length of stay, d	8.0 (6.0, 9.0)	5.0 (4.0, 8.0)	0.019
Hospital length of stay, d	10.0 (9.0, 13.0)	10.5 (6.7, 15.8)	0.899

BZD = benzodiazepine; ICU = intensive care unit

Summary

Data analysis can be a tricky task but with adequate resources, it is not as difficult as you may first imagine. Some final words/advice:

- Seek help early!
- Ensure that your data is clean and valid prior to performing statistical analysis
- Perform descriptive statistics first to visually inspect data for distribution, kurtosis, outliers, etc.
- Determine the appropriate statistical test to perform (in another section of this guide)
- Ensure assumptions of tests aren't violated prior to evaluation