Mobile App for Teaching & Learning Statistics

Bernhard Klingenberg
Director, Masters in Applied Data Science
New College Florida

eCots 2026
Electronic Conference on Teaching Statistics

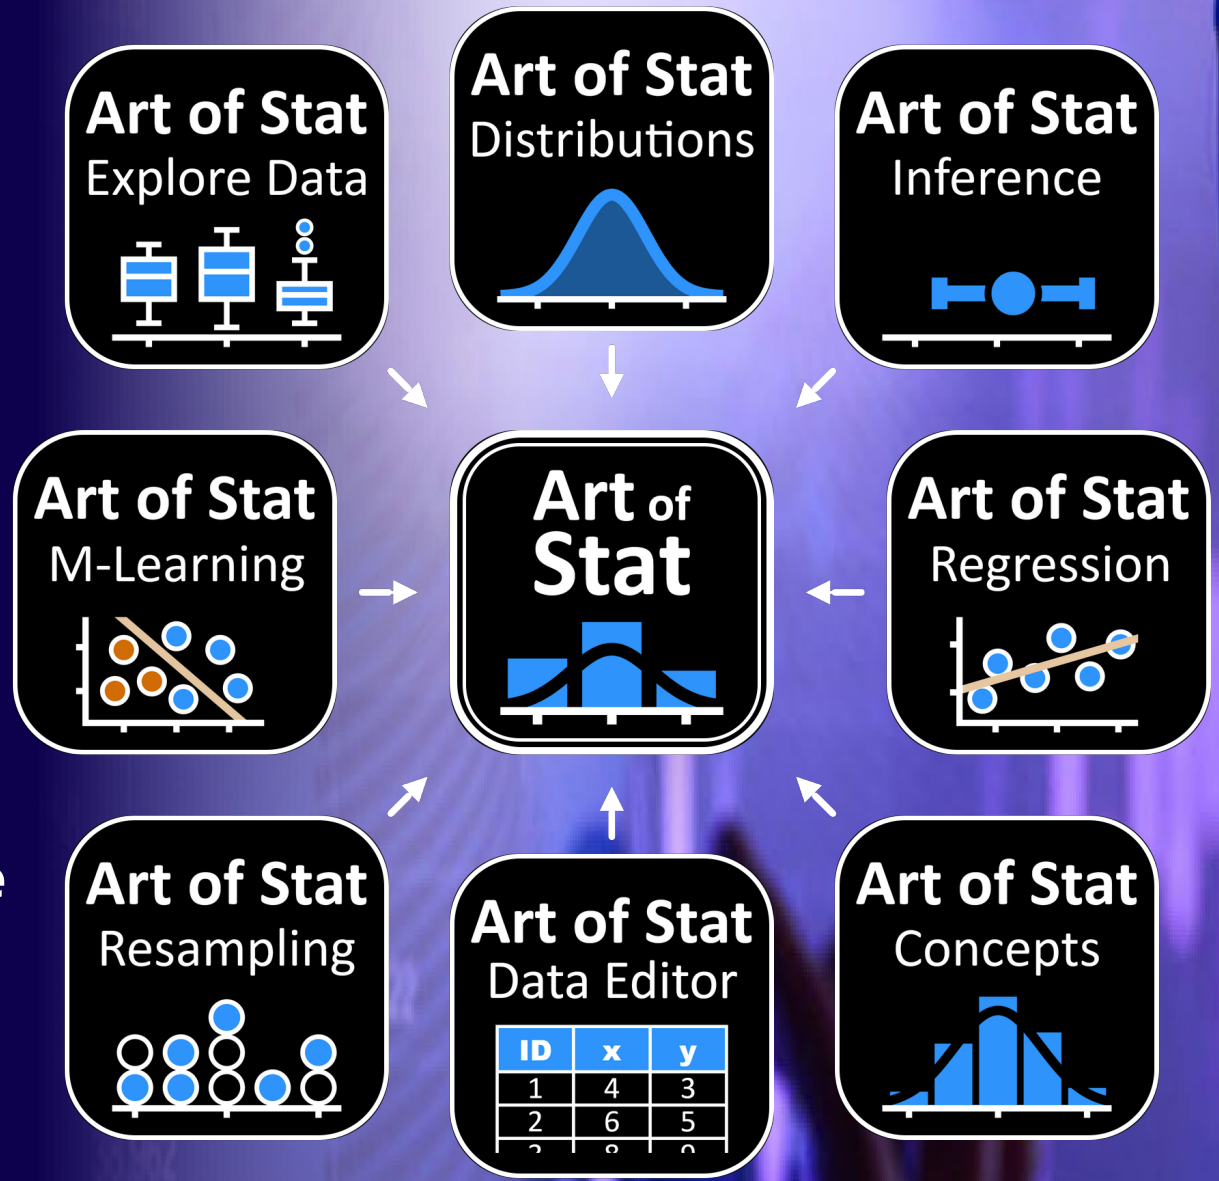

Introduction

- Ph.D. in Statistics (@UF in 2004)
- Worked at Williams College for 18 years
- Research in Categorical Data Analysis, ML for Solar Flare identification
- Co-author of: “Statistics, the Art & Science of Learning from Data” (with Alan Agresti, Chris Franklin, Pearson 2021)
- Web- and Mobile Apps: “Art of Stat”

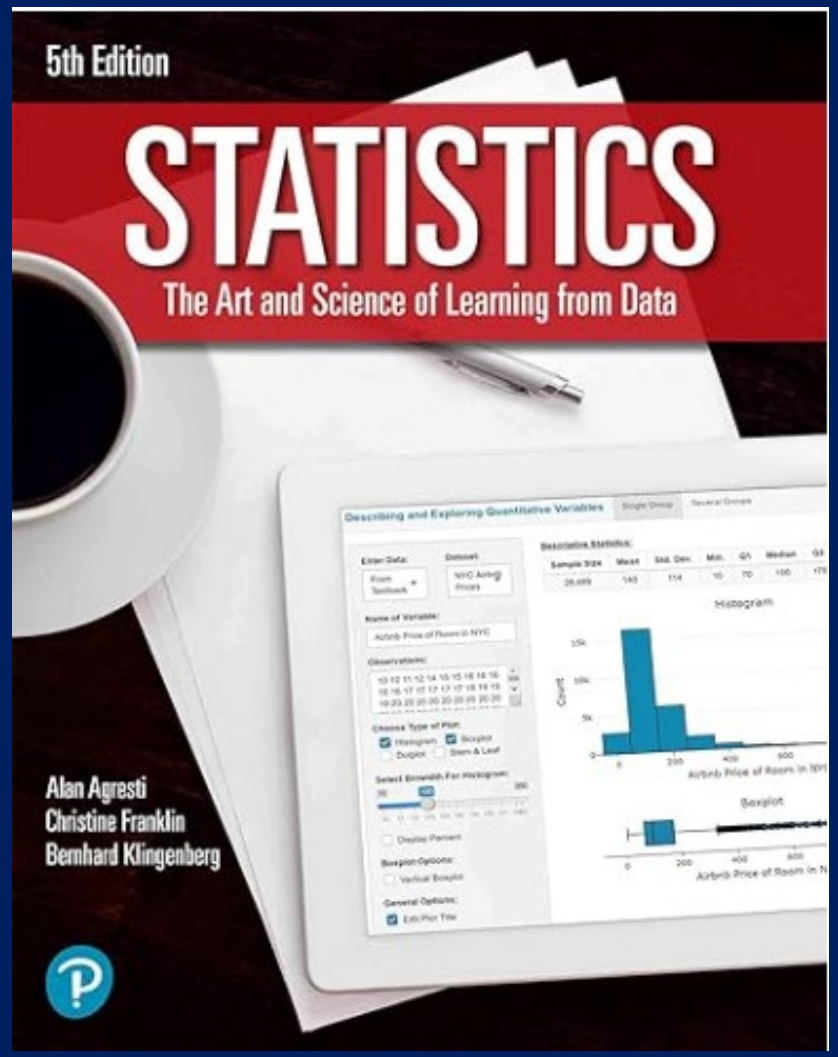

Art of Stat Distributions

Art of Stat Explore Data

Art of Stat Inference

Art of Stat Regression

Art of Stat Concepts

Art of Stat Resampling

Art of Stat M-Learning

Download on the App Store

GET IT ON Google Play

The All-in-One Art of Stat Mobile App

Art of Stat

Download on the App Store

GET IT ON Google Play

Contains 8 Modules

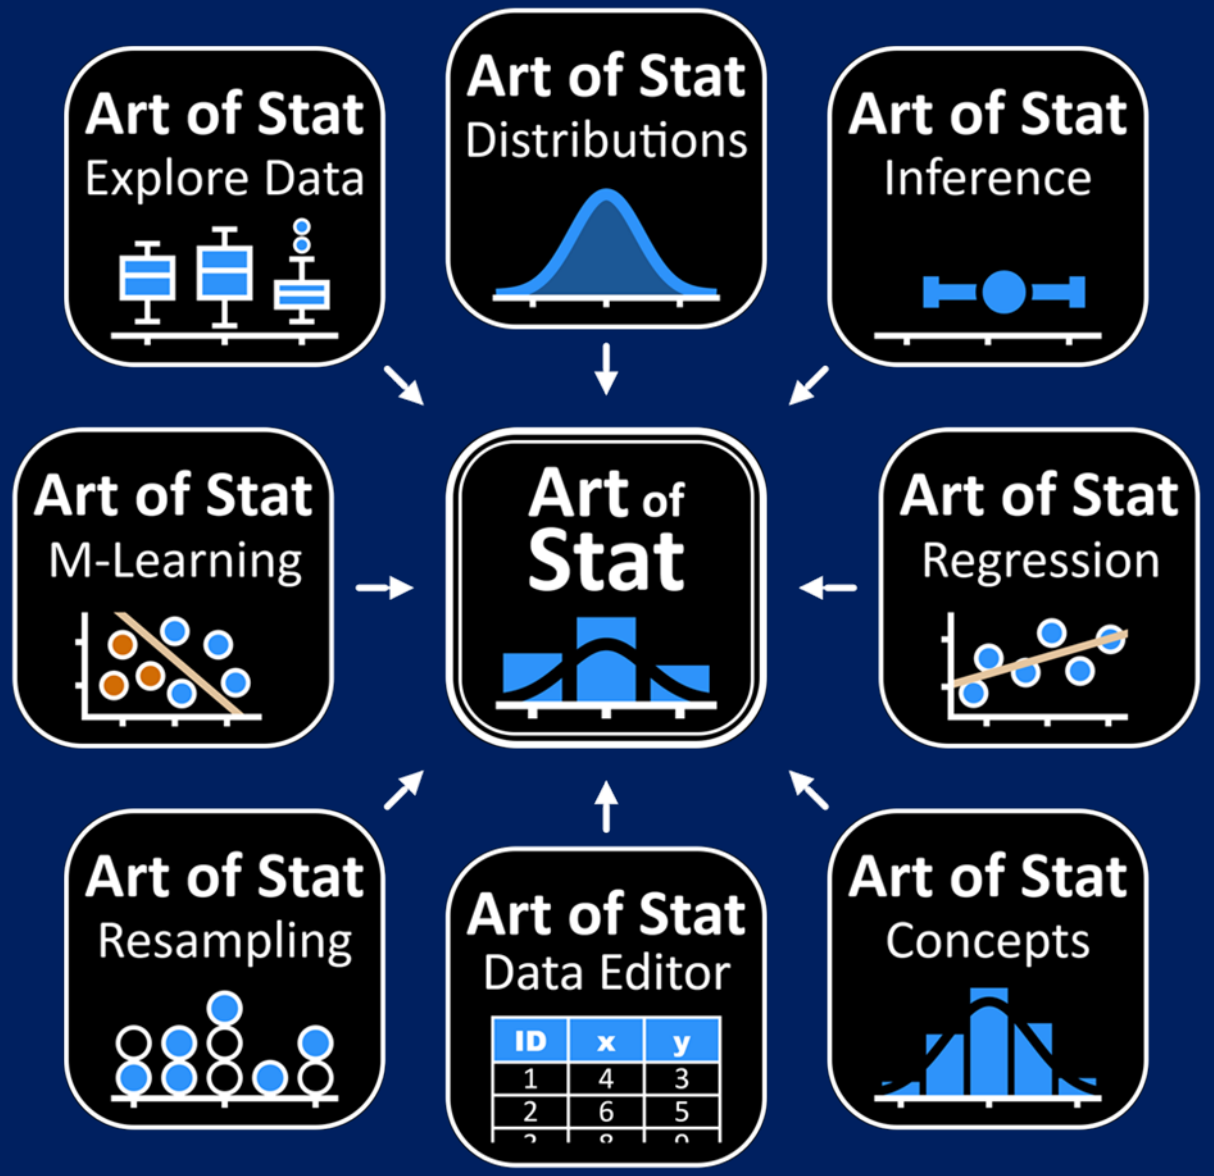

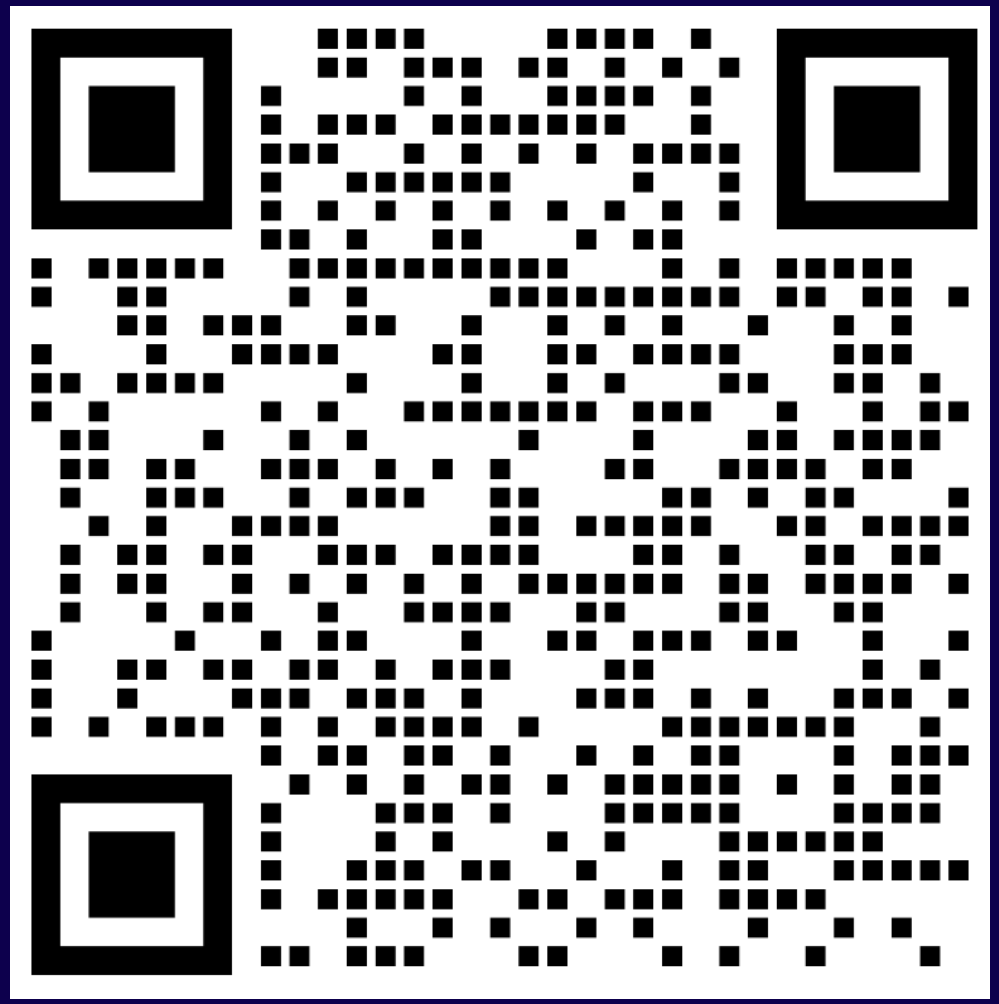

Download on the App Store

GET IT ON
Google Play

Why Teaching with Technology (Apps)?

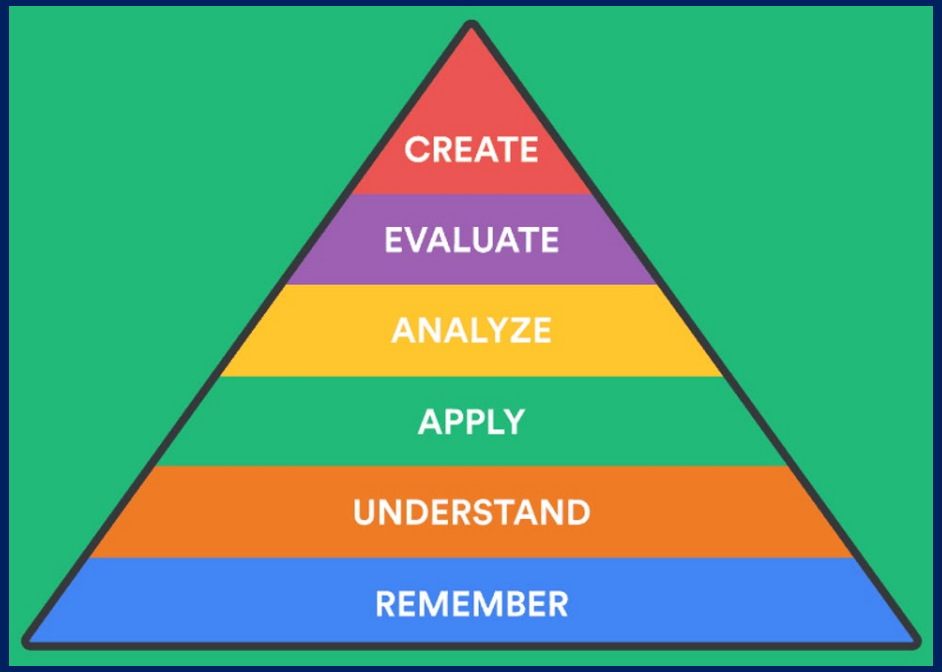

Blooms Taxonomy of effective learning:
Move student from simple recall toward increasingly sophisticated forms of thinking

Methods & Concepts in Stats…

- are understood when you can "see" and interact with them.
- become memorable when you can apply them in an activity.
- stay relevant when you can analyze and create insights and knowledge using your own data.

App Technology

- Connect Phone/Tablet to Screen
- Share Screenshots, Join Zoom Calls
- Screen Recording (Videos)
- Upload your own data (.csv, .xls)
- Available 24/7, no lab needed
- Work Off-Line (Airplane Mode)

Me in a Zoom call sharing my screen with the class

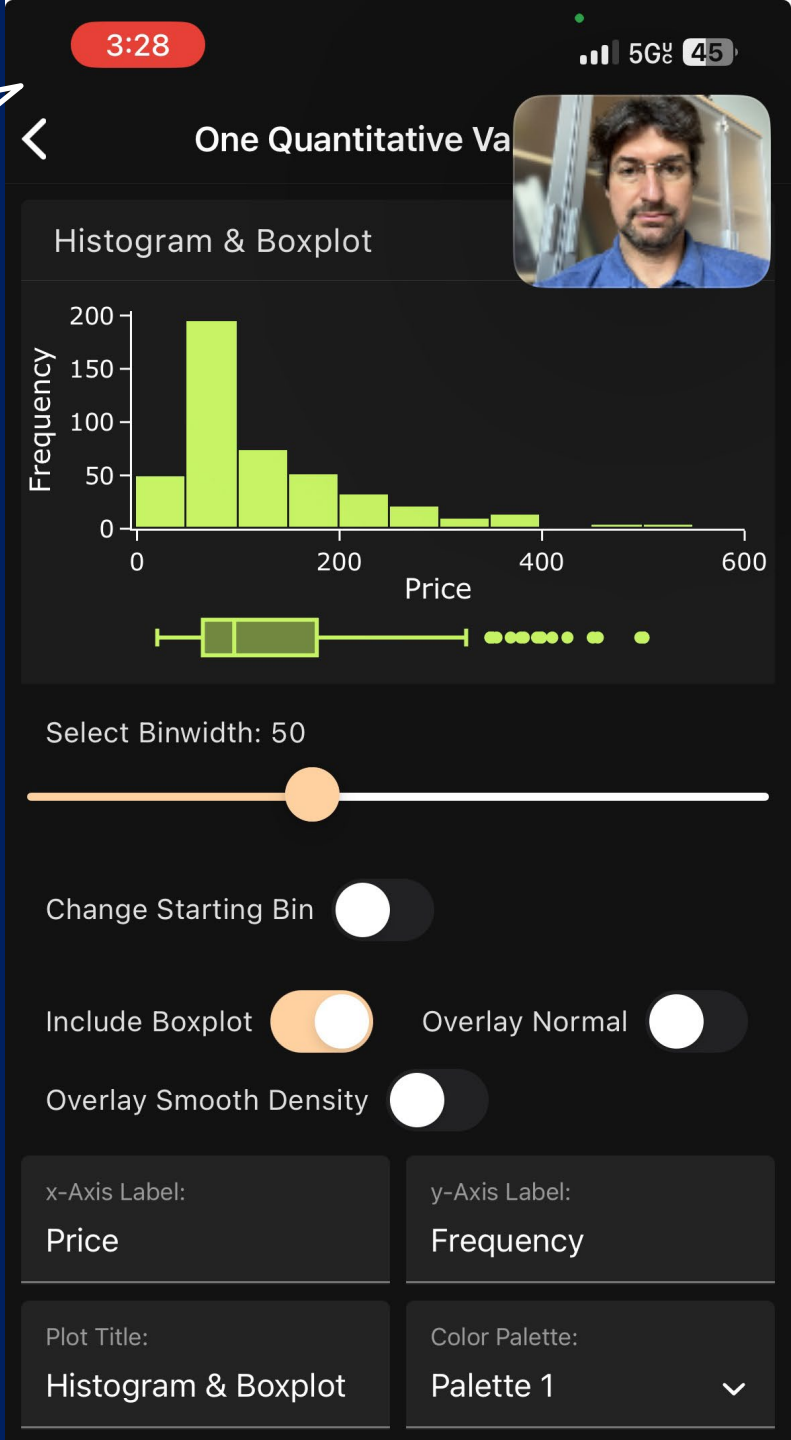

Overview of Modules

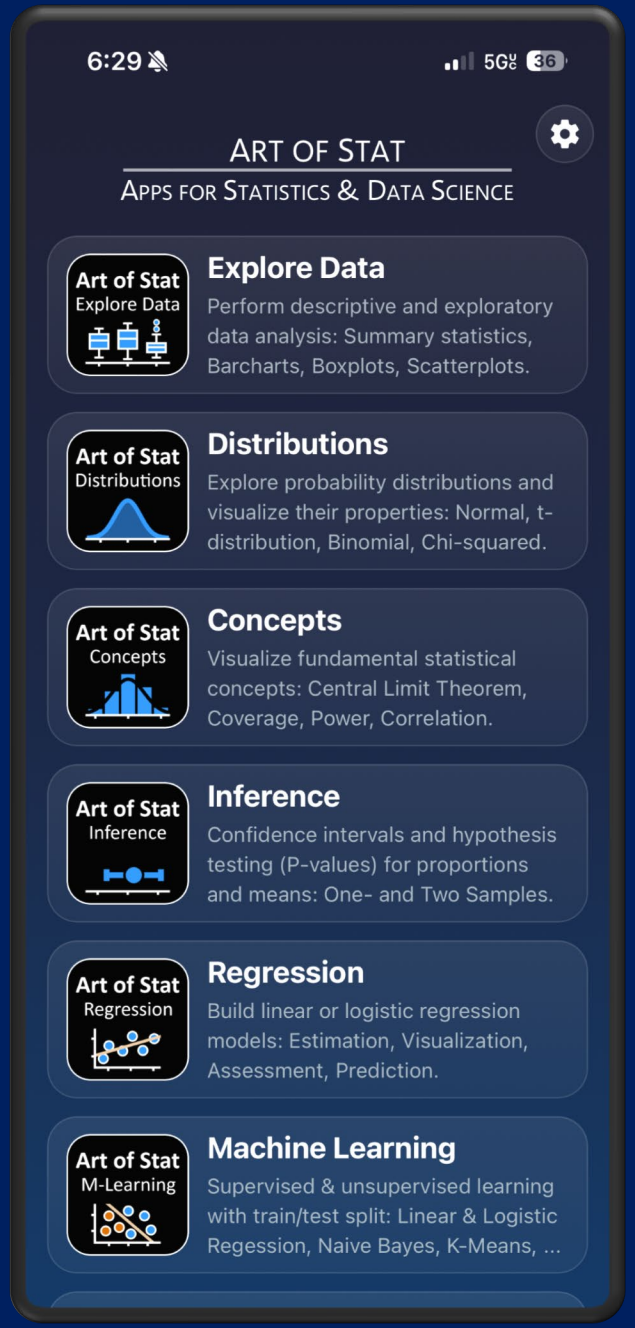

Art of Stat
Explore Data

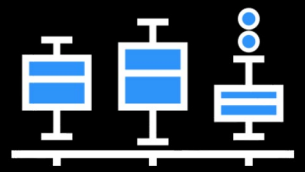

Descriptive Statistics, Plots for Categorical and Quantitative Variables, Scatterplots

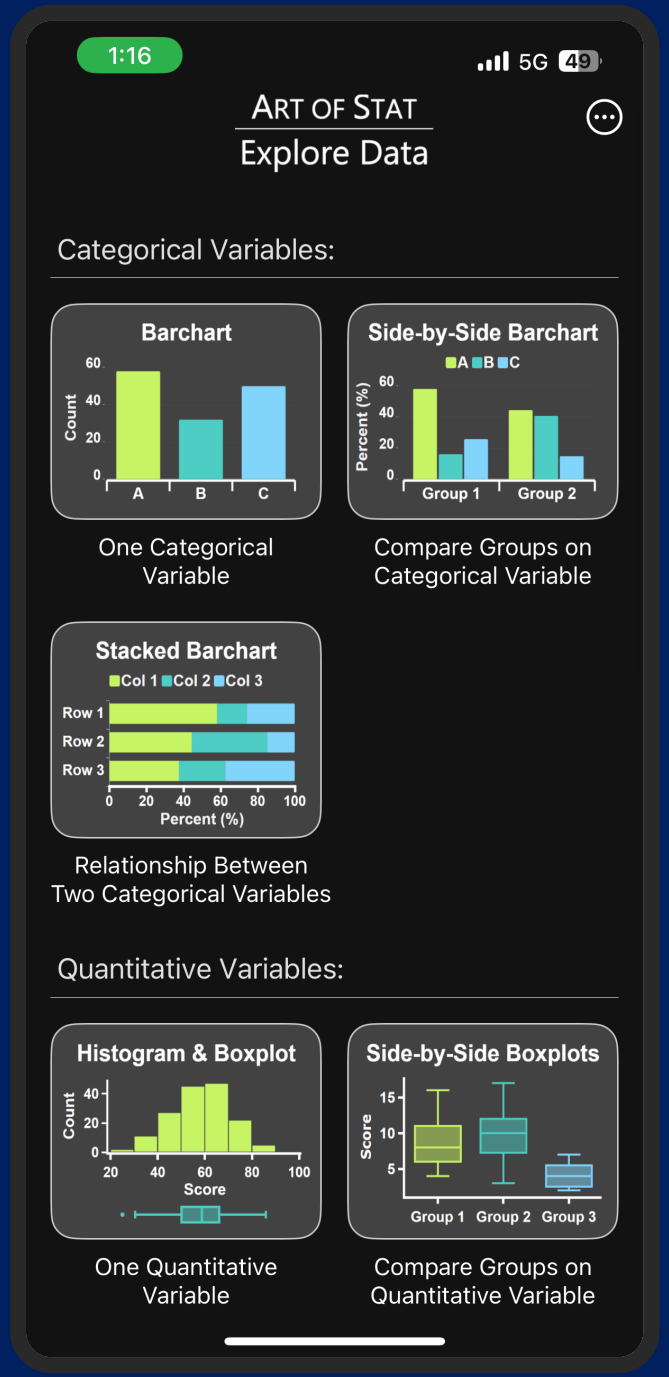

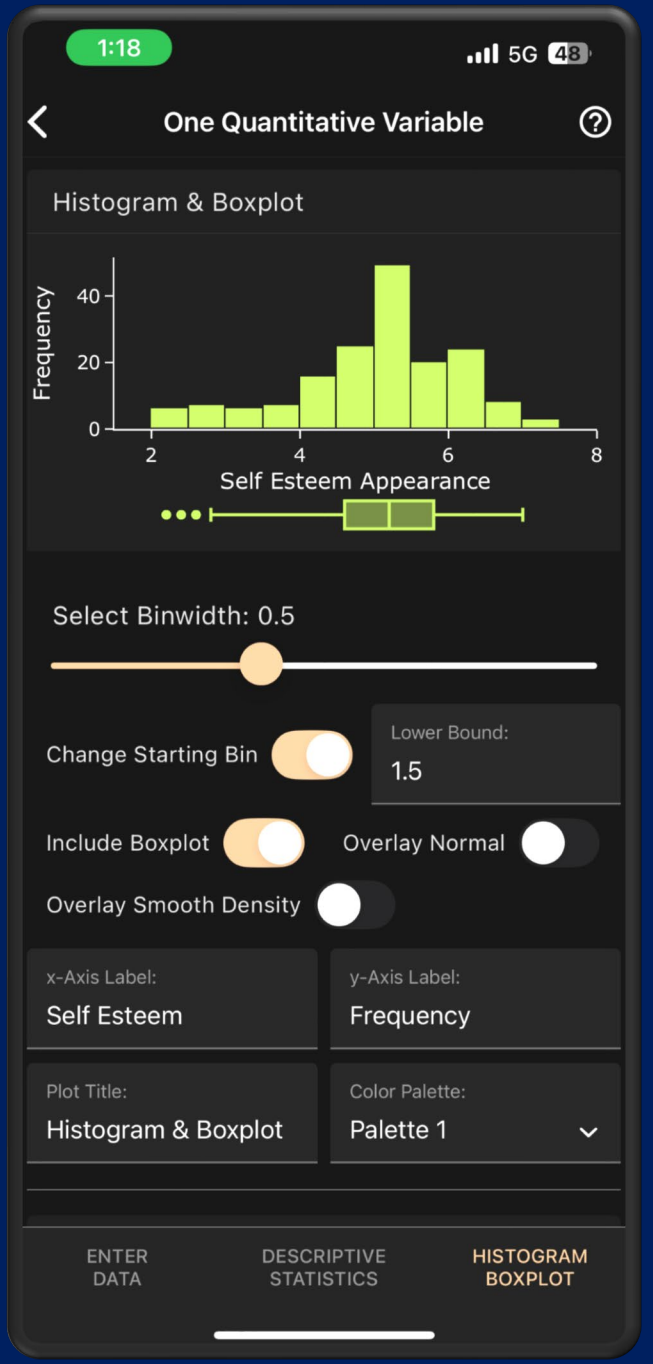

Art of Stat Distributions

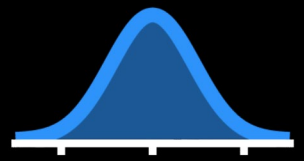

Explore & Visualize Discrete and Continuous Probability Distributions

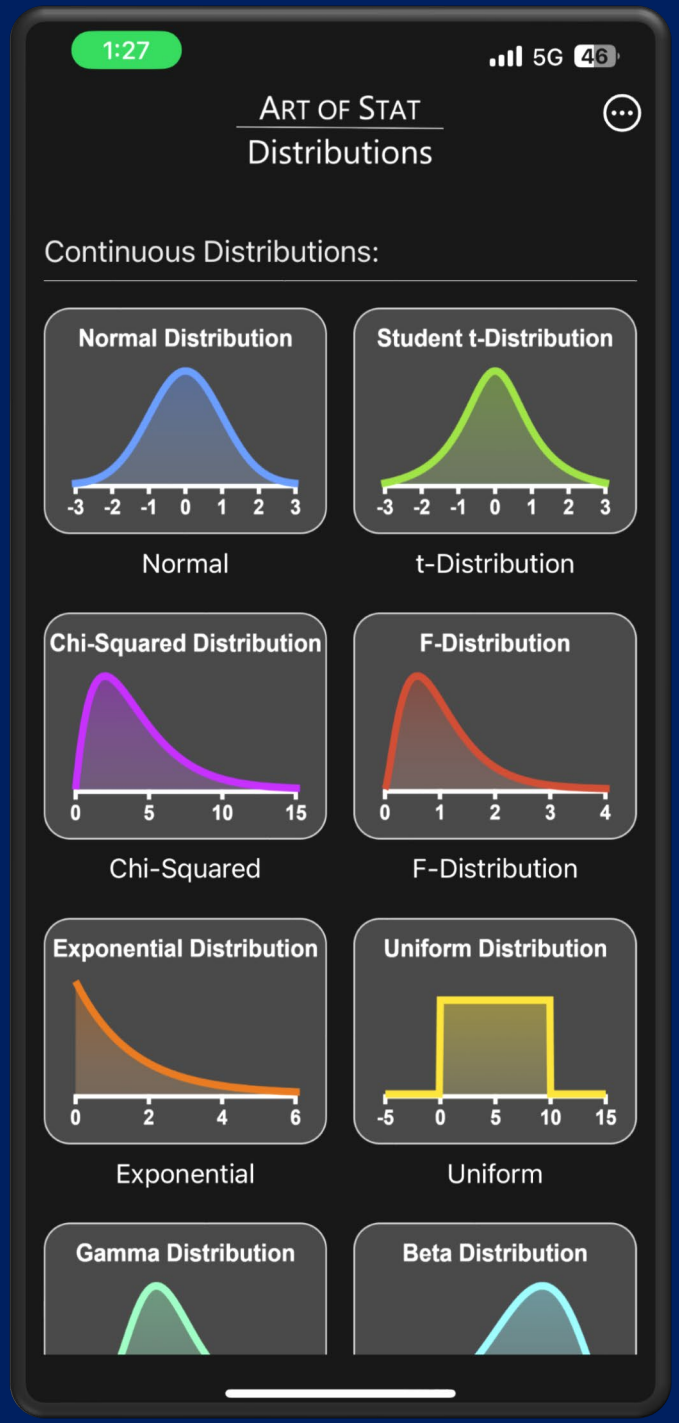

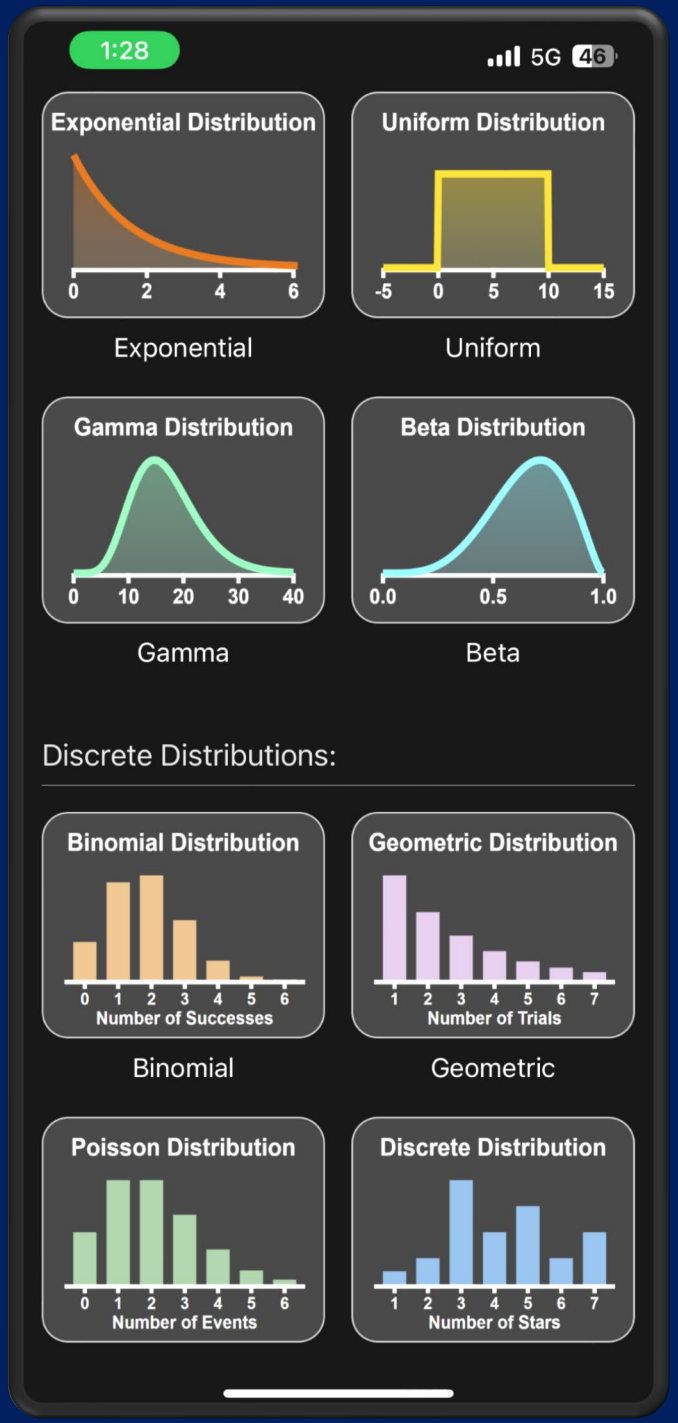

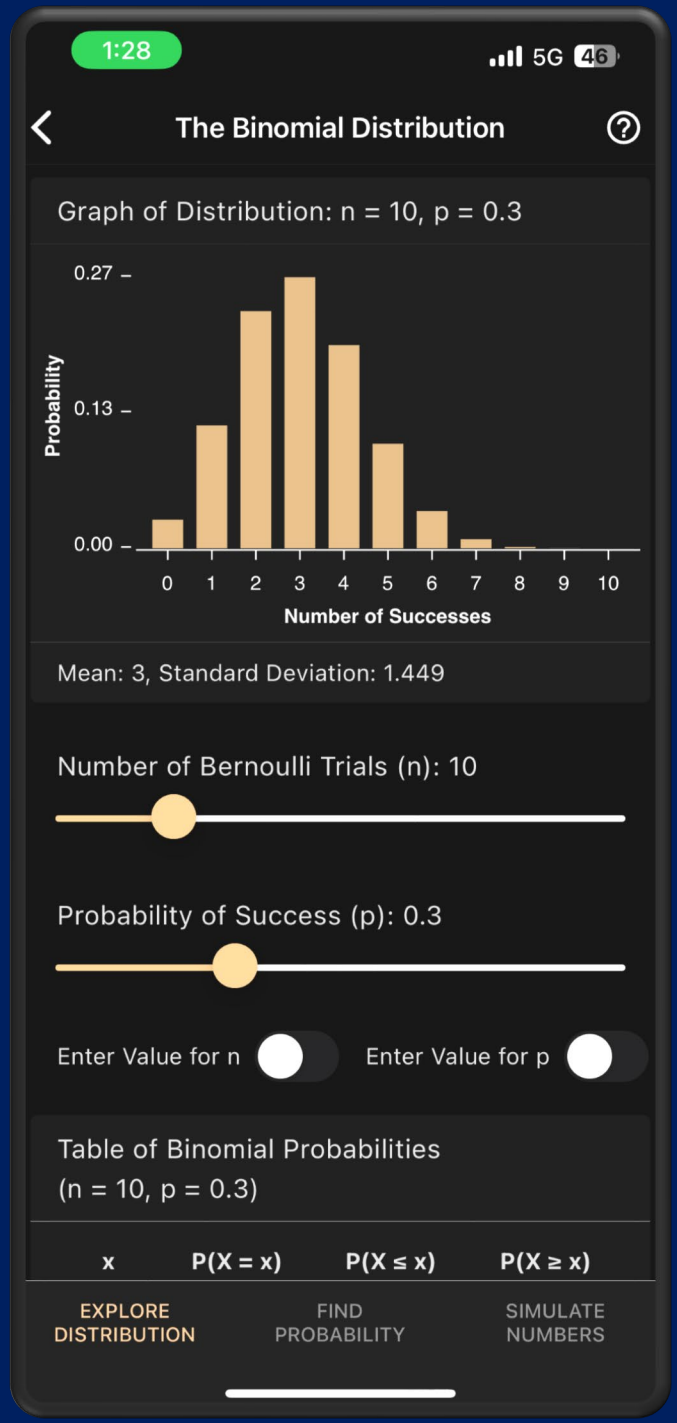

Art of Stat
Inference

Confidence Intervals, Hypothesis Tests for one, two & paired samples. Inference for Linear Regression

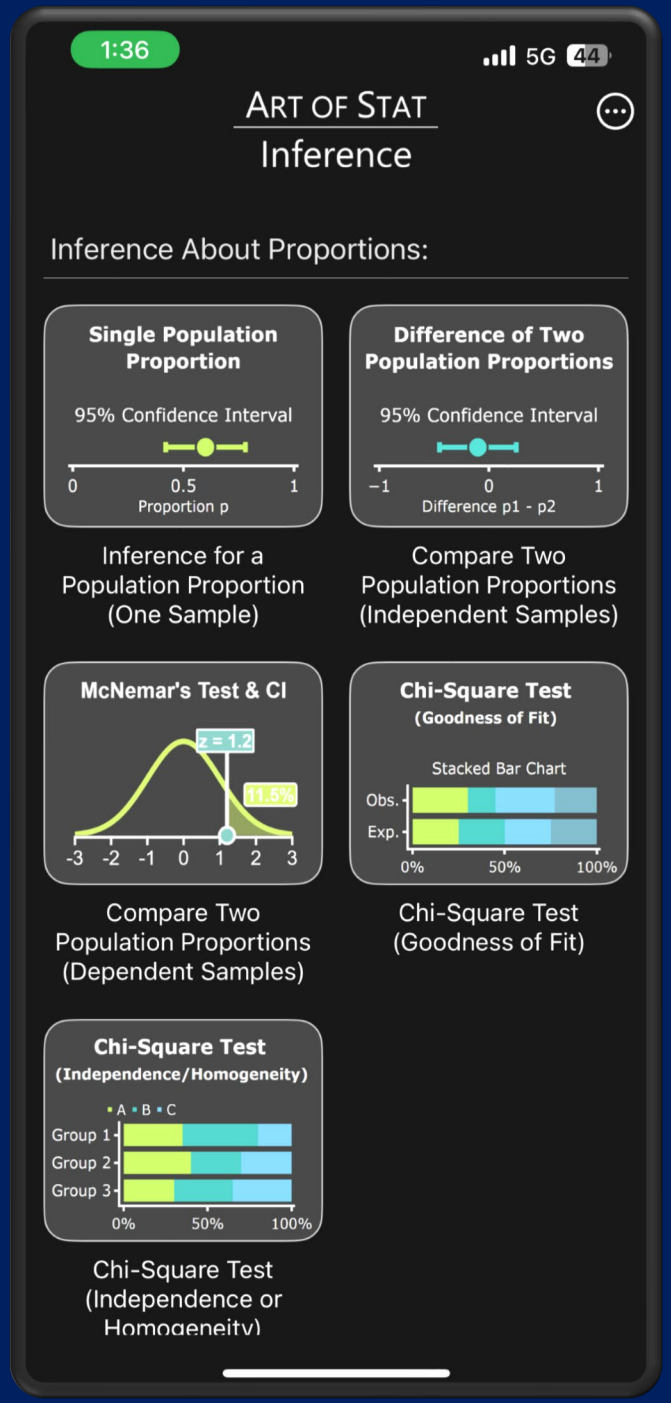

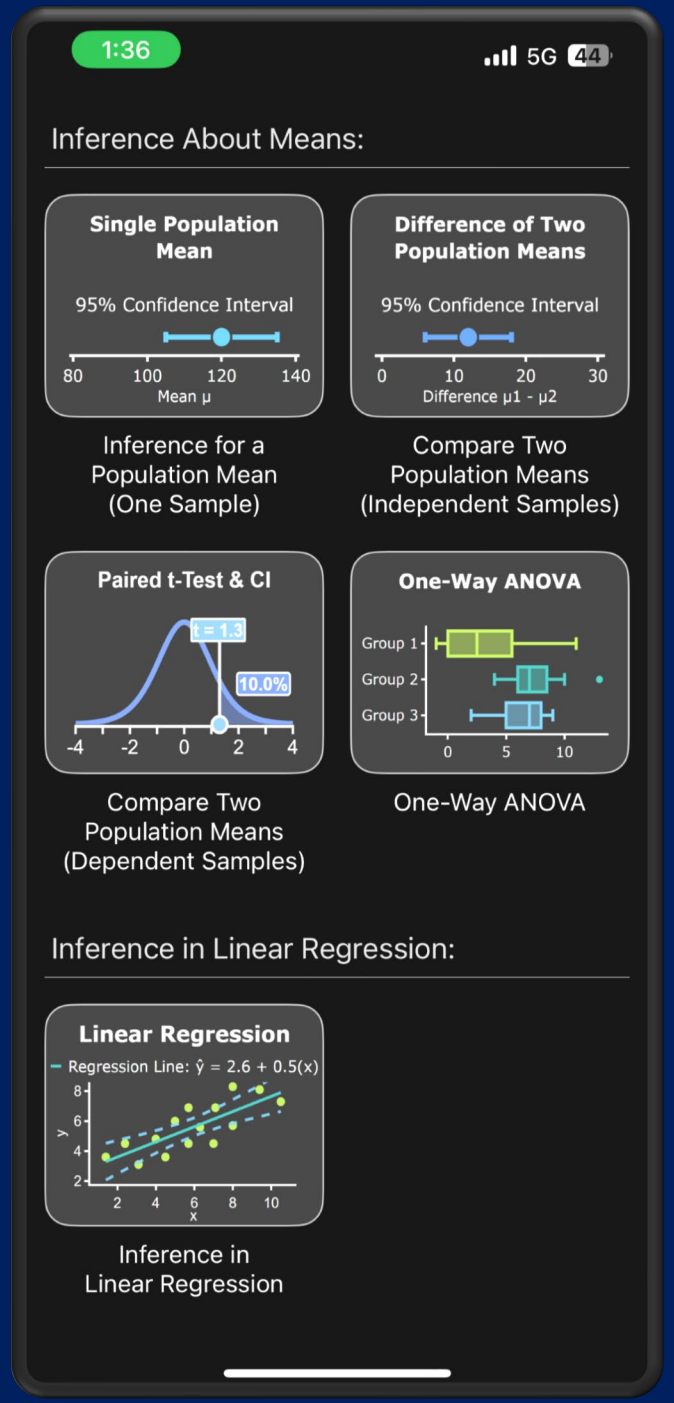

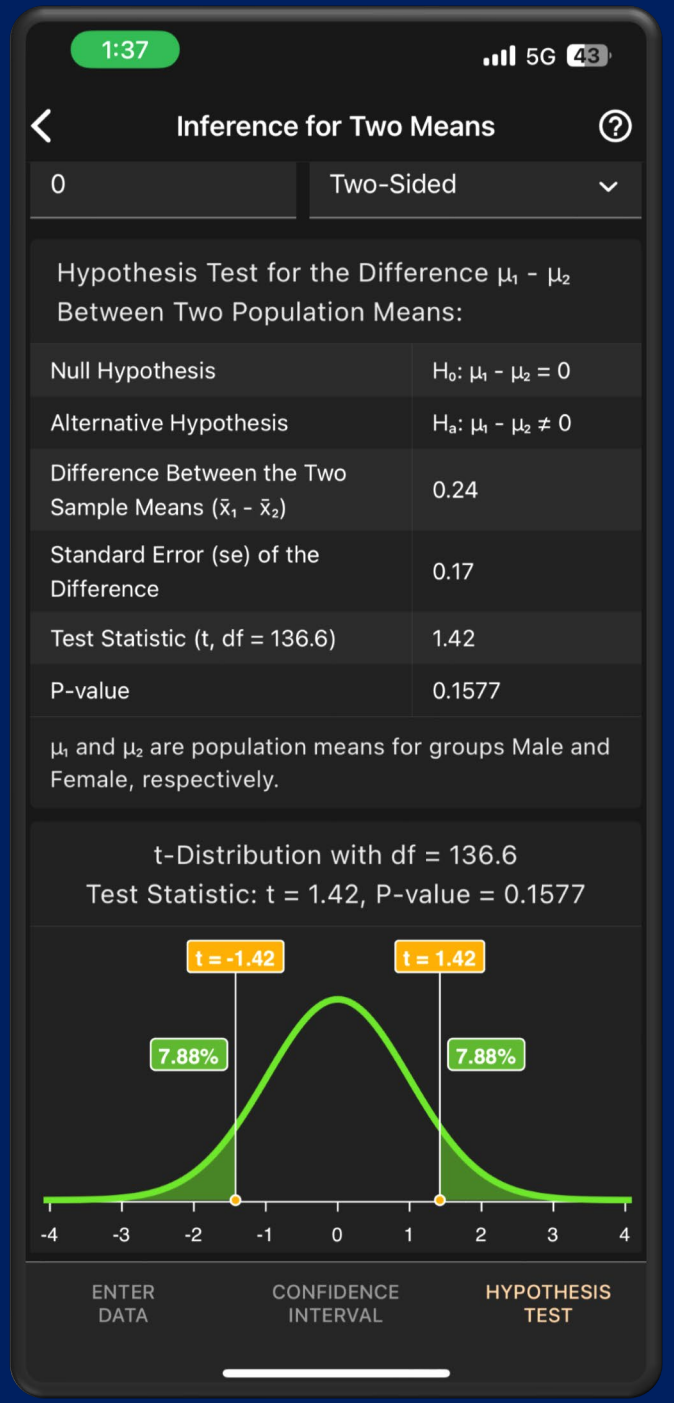

Art of Stat
Regression

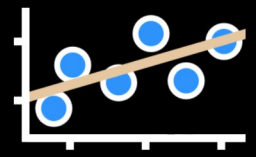

Simple Linear, Poisson, Exponential & Logistic Regression.

Multiple Linear & Logistic Regression

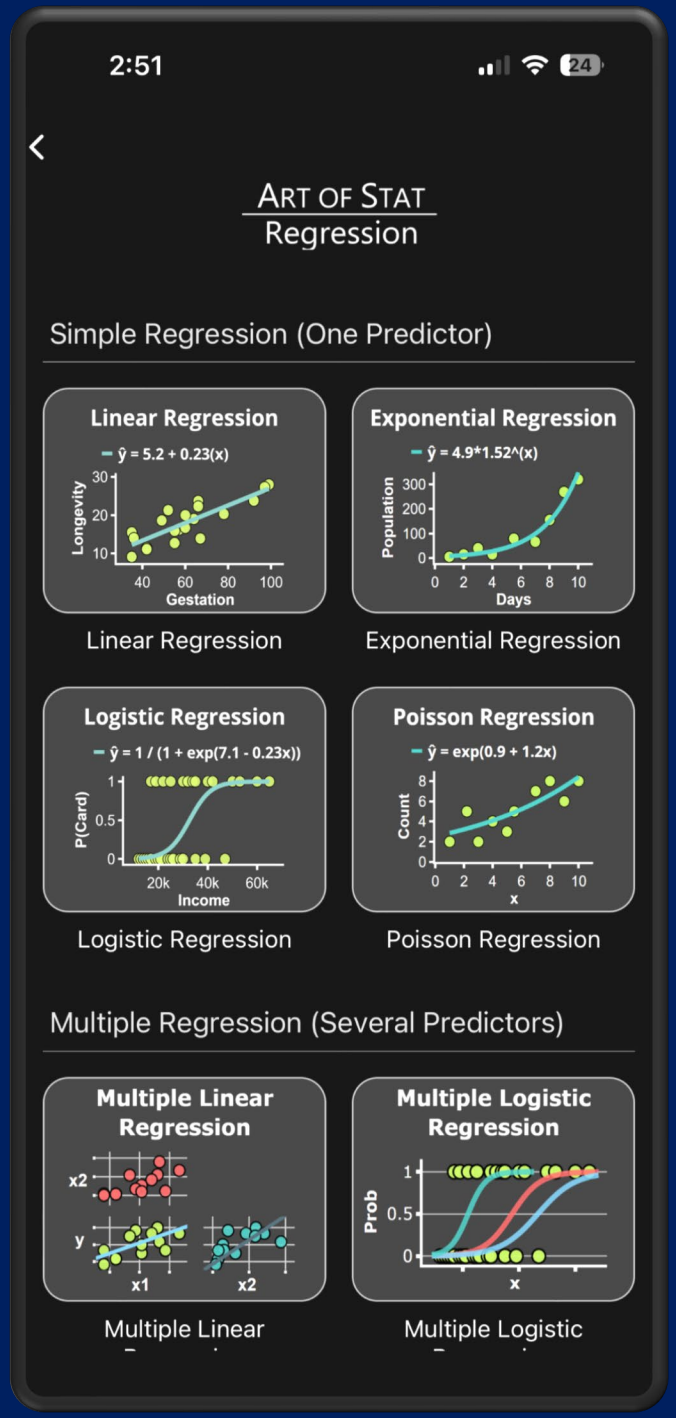

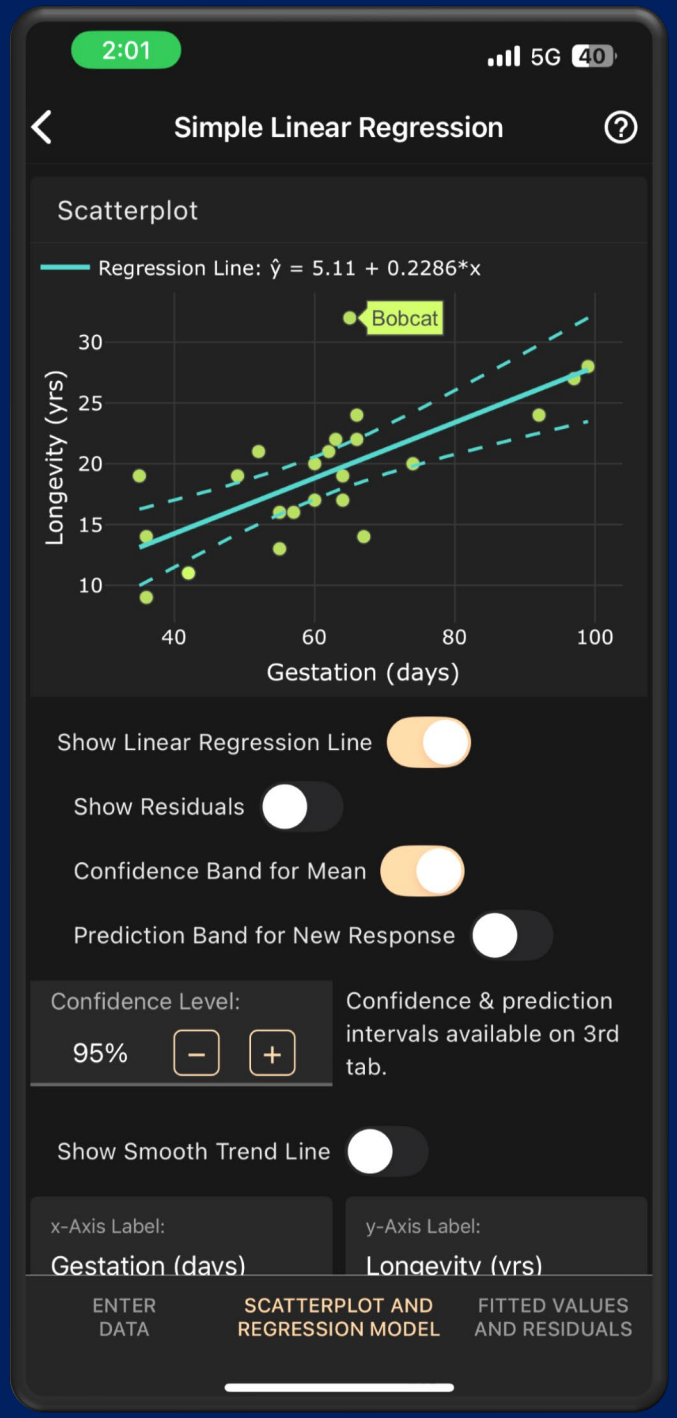

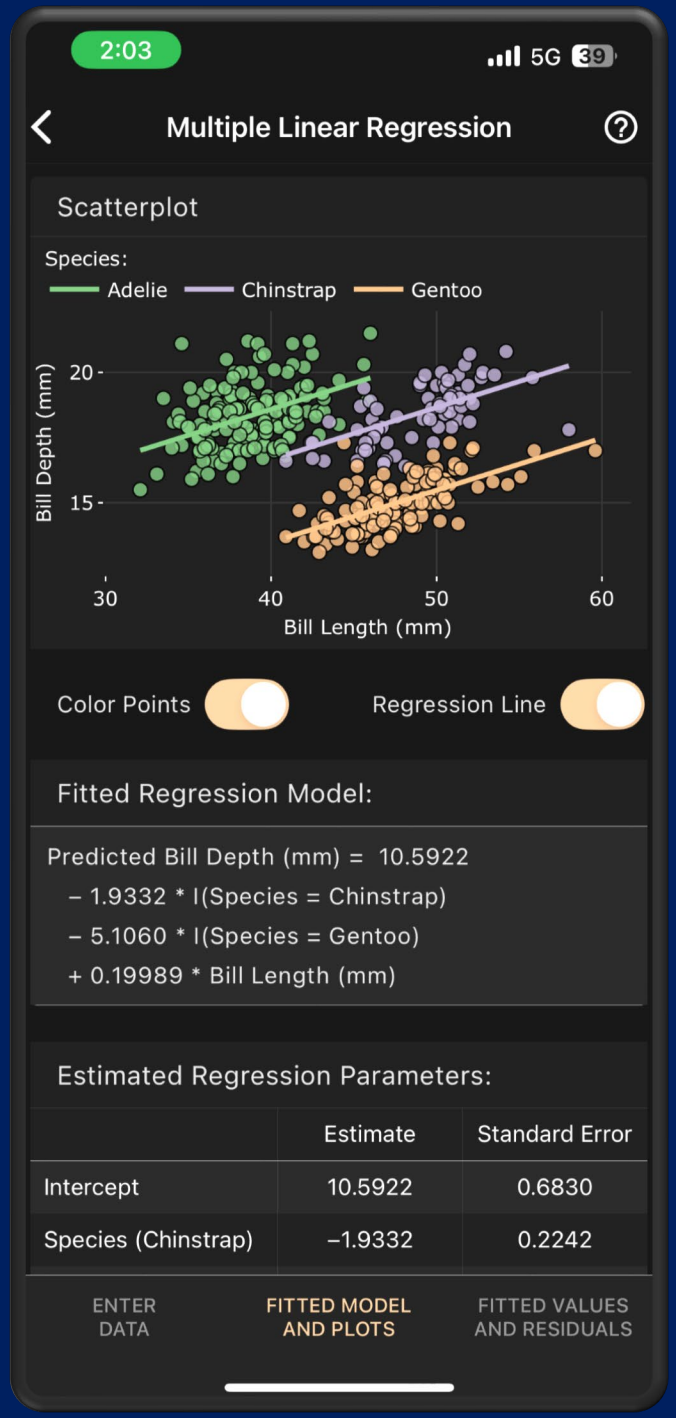

Art of Stat Concepts

Central Limit Theorem, Correlation, Regression, Coverage & Power

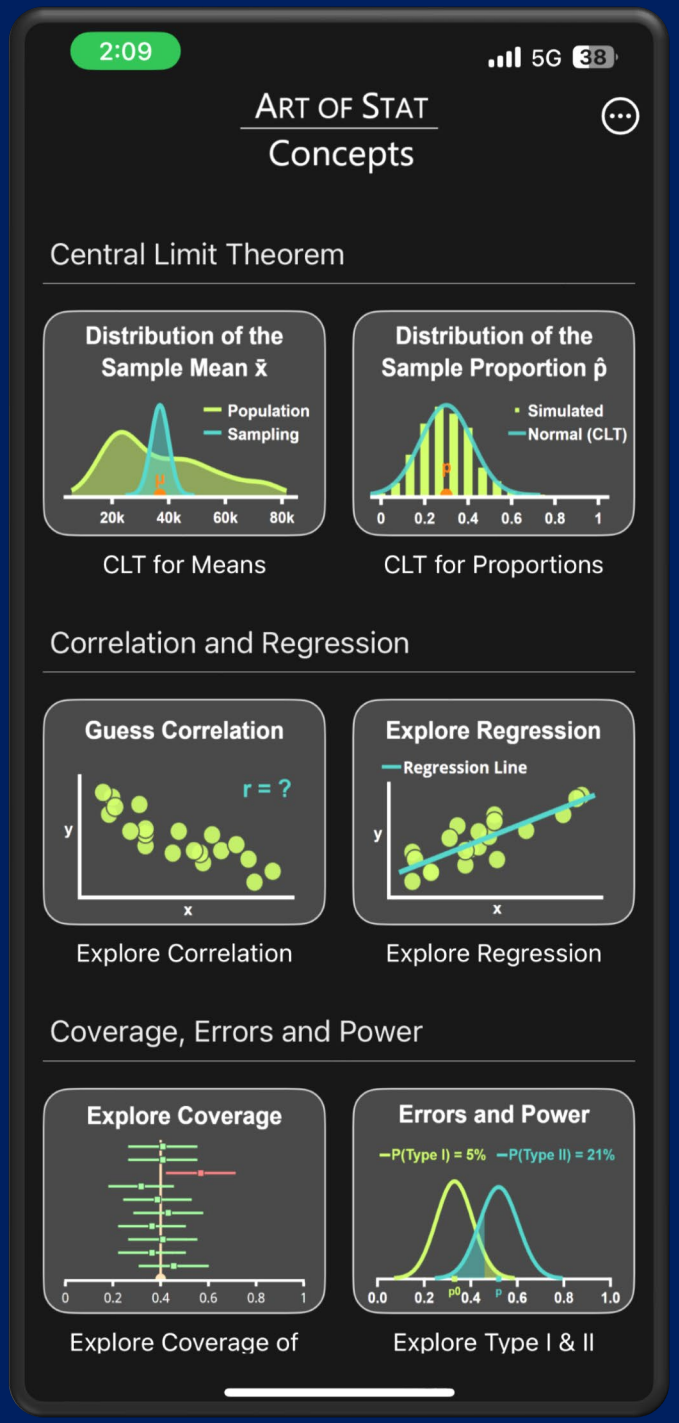

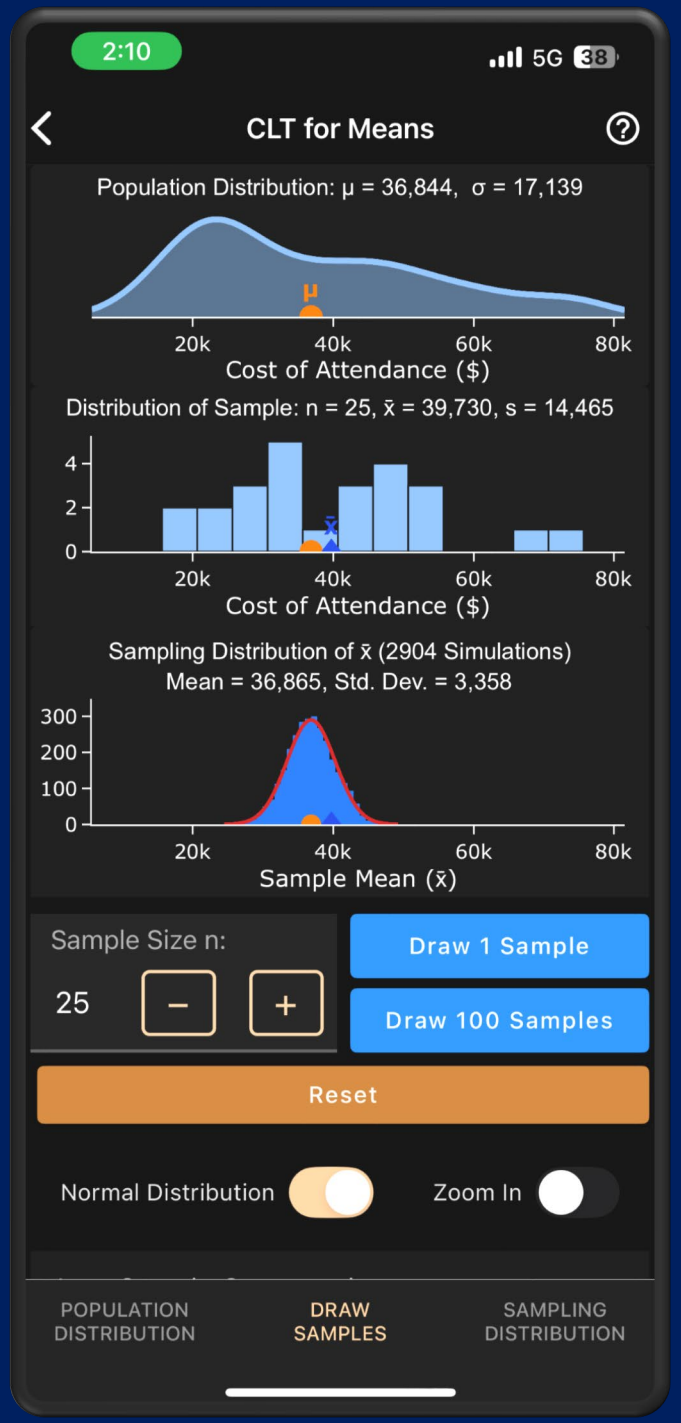

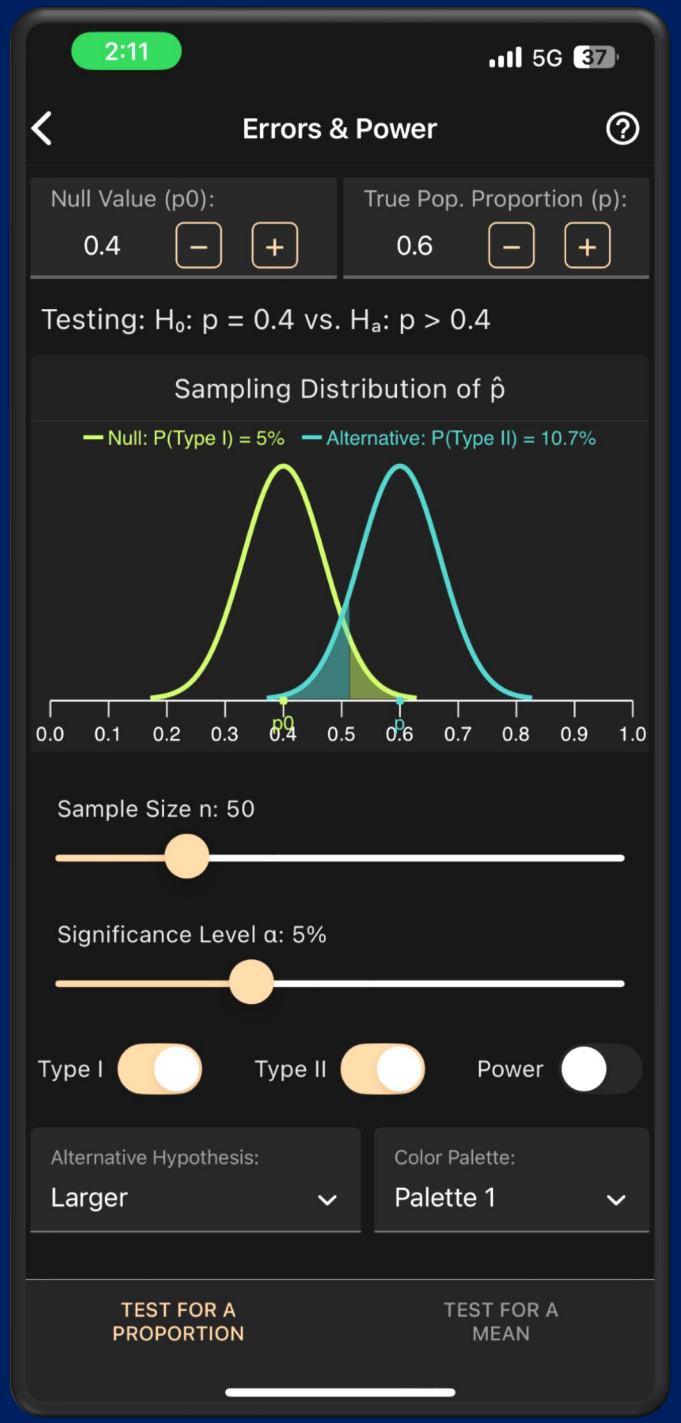

Art of Stat
Resampling

Bootstrap Confidence Intervals. Permutation Tests. For One & Two Samples. Chi-Squared Test.

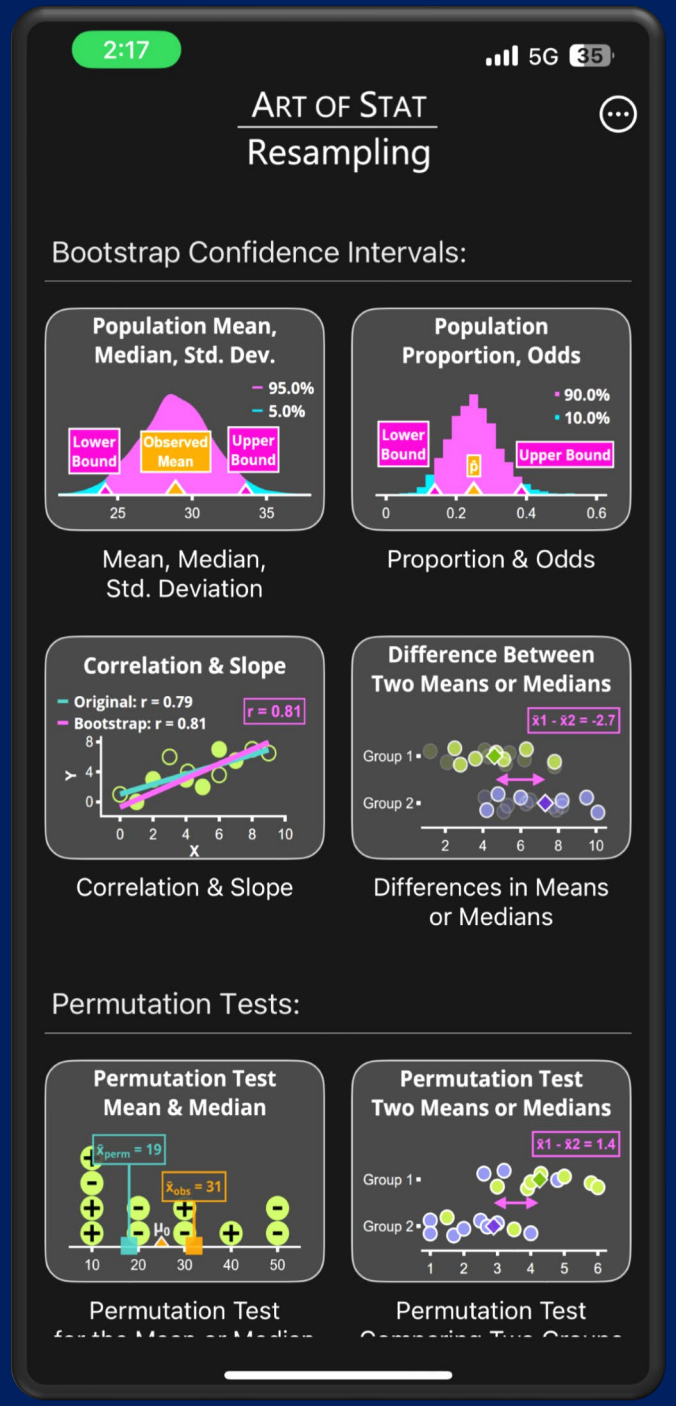

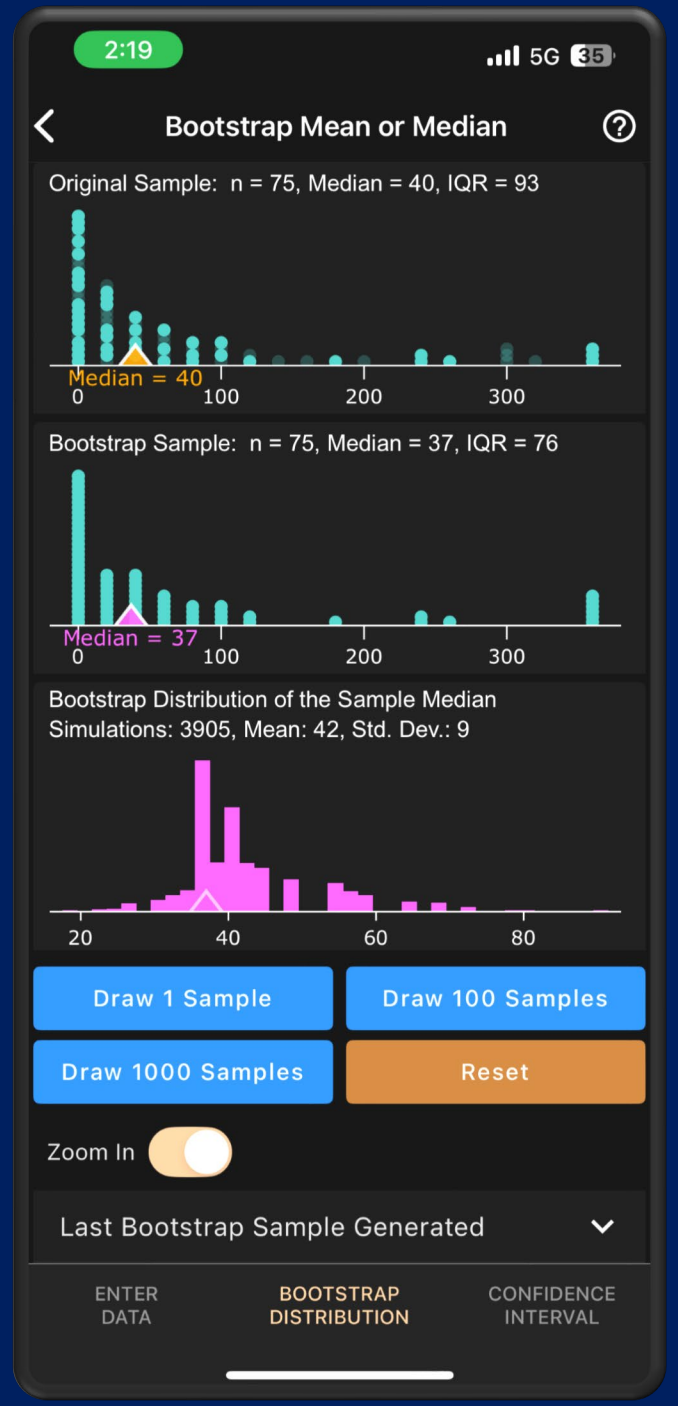

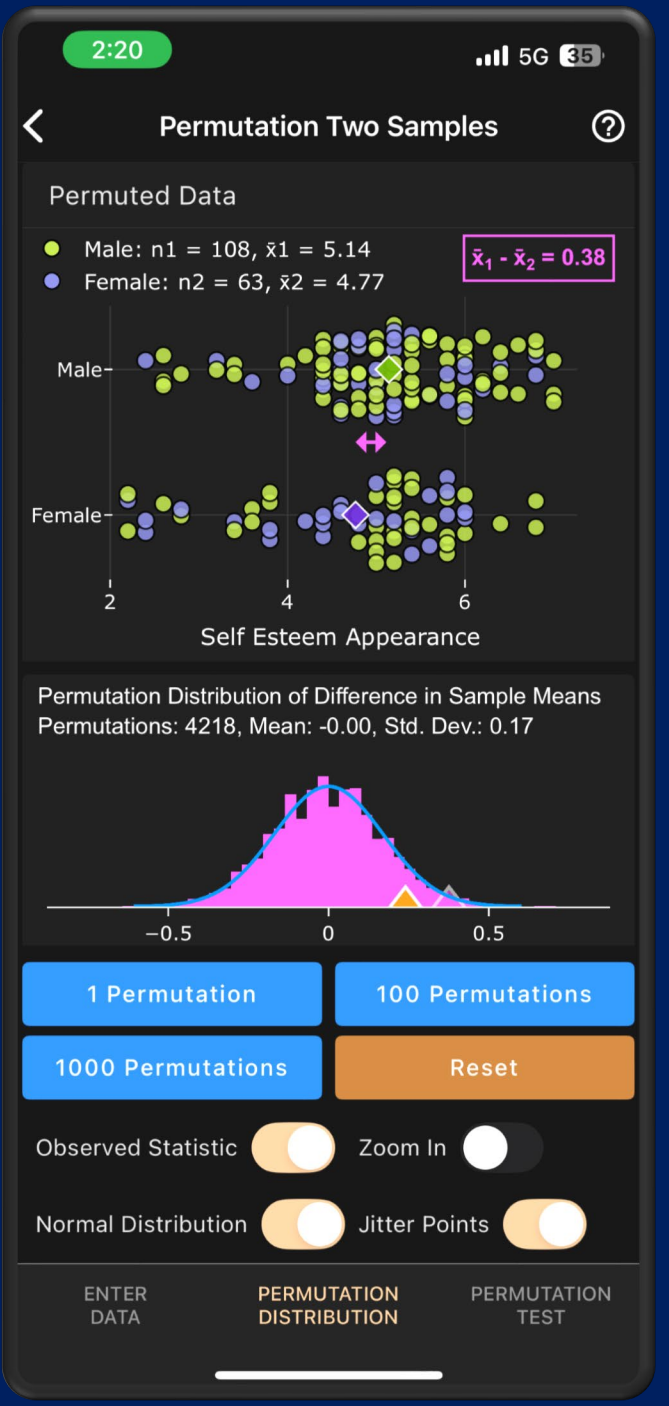

Art of Stat
M-Learning

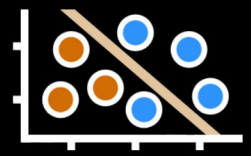

Classical Methods for Supervised & Unsupervised Learning

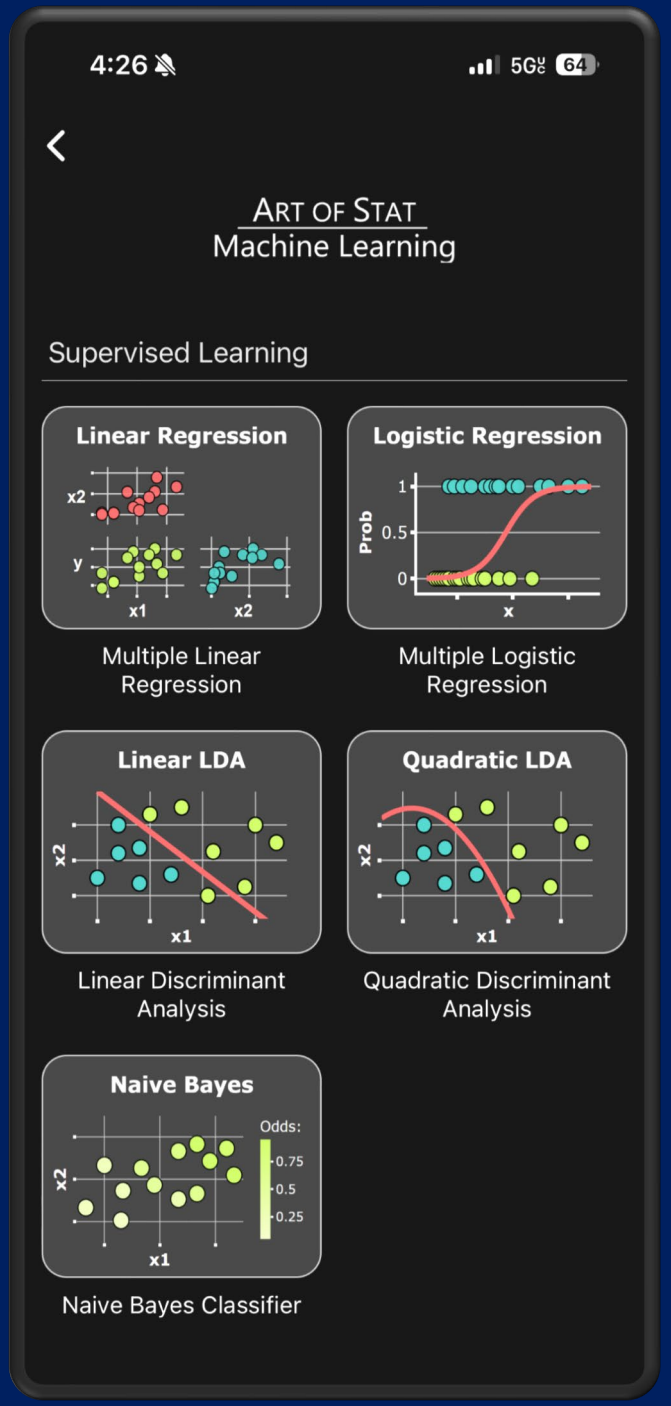

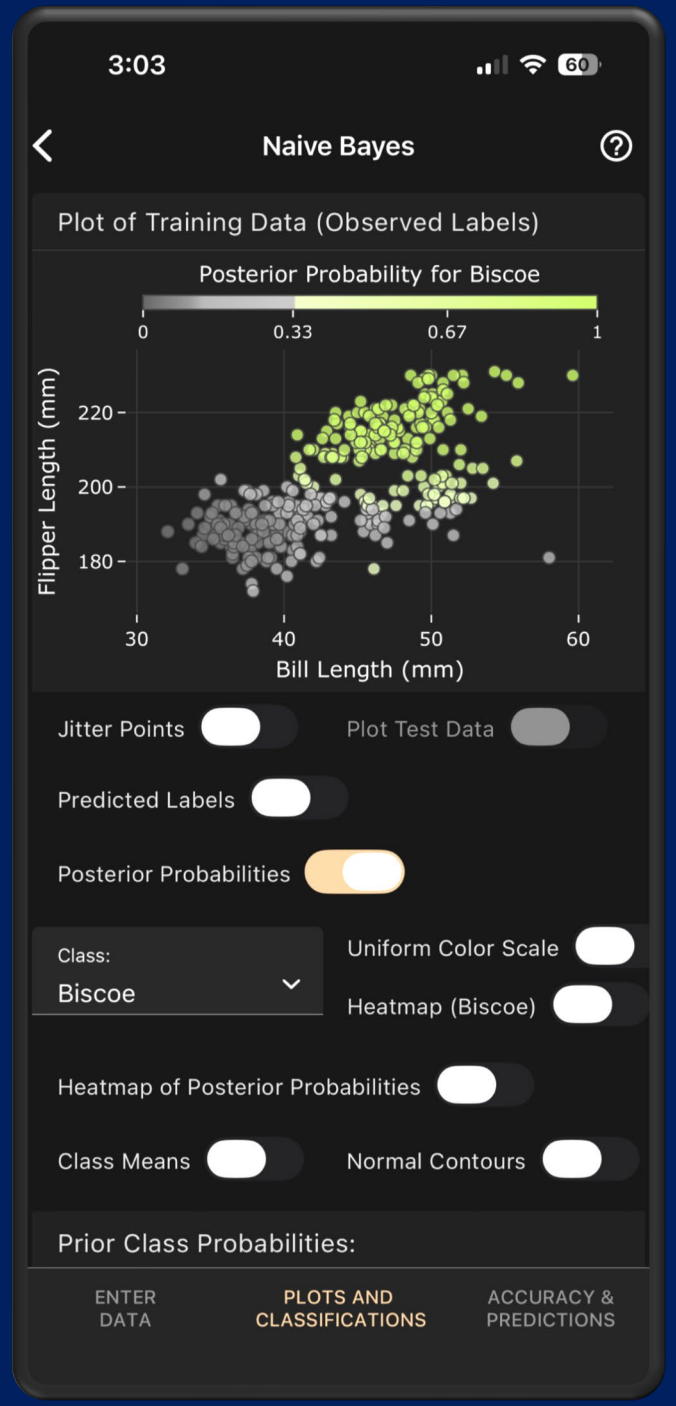

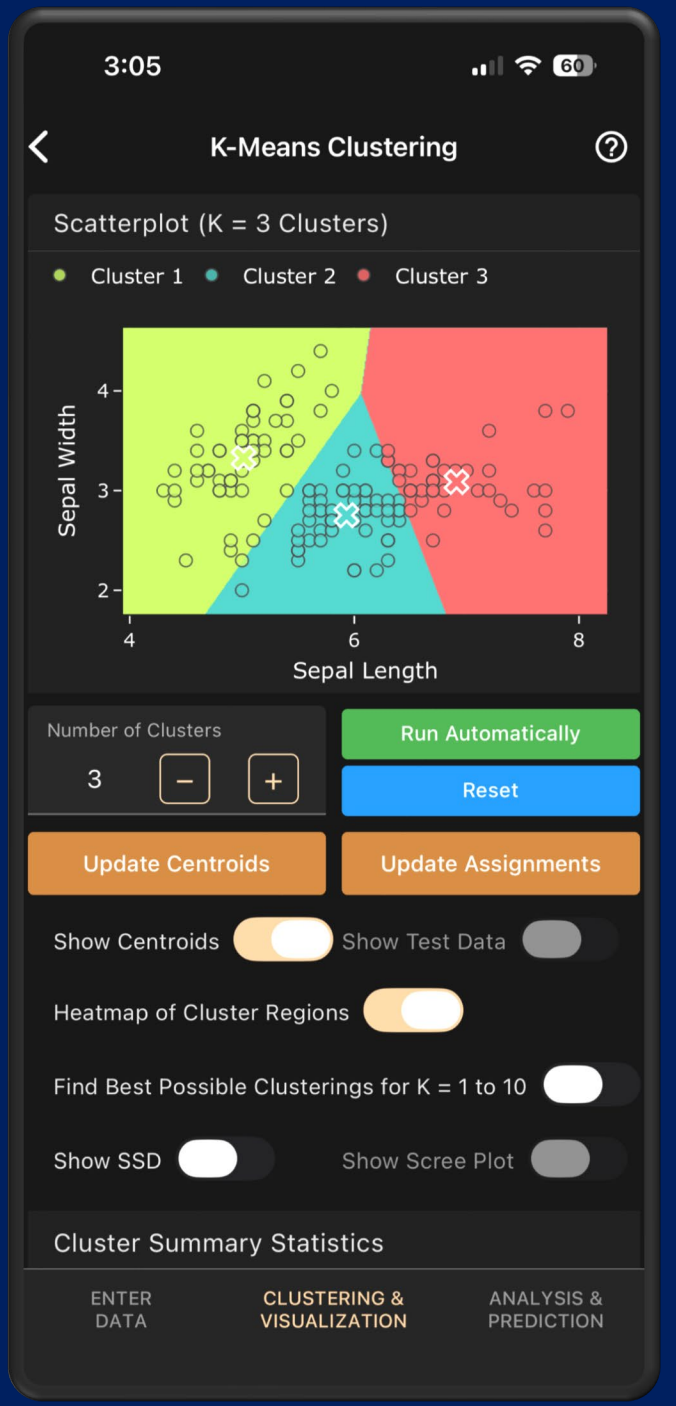

Is there anything you want to see? (Go Live)

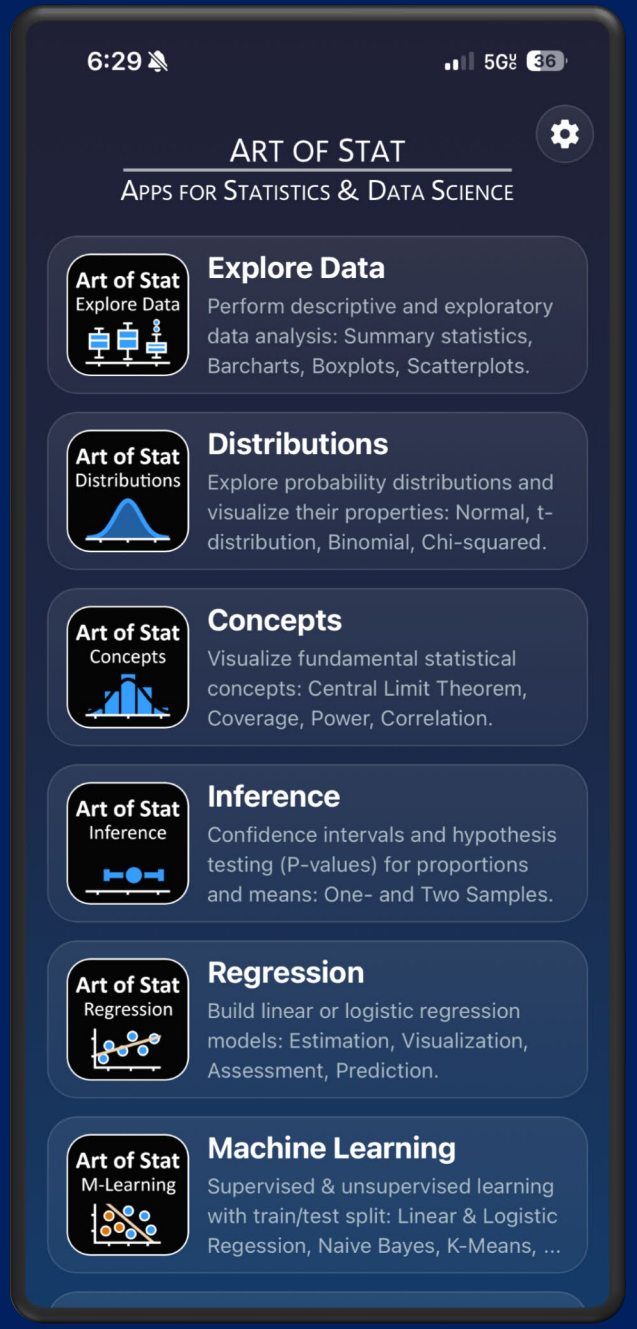

Run-Through of the Naïve Bayes App

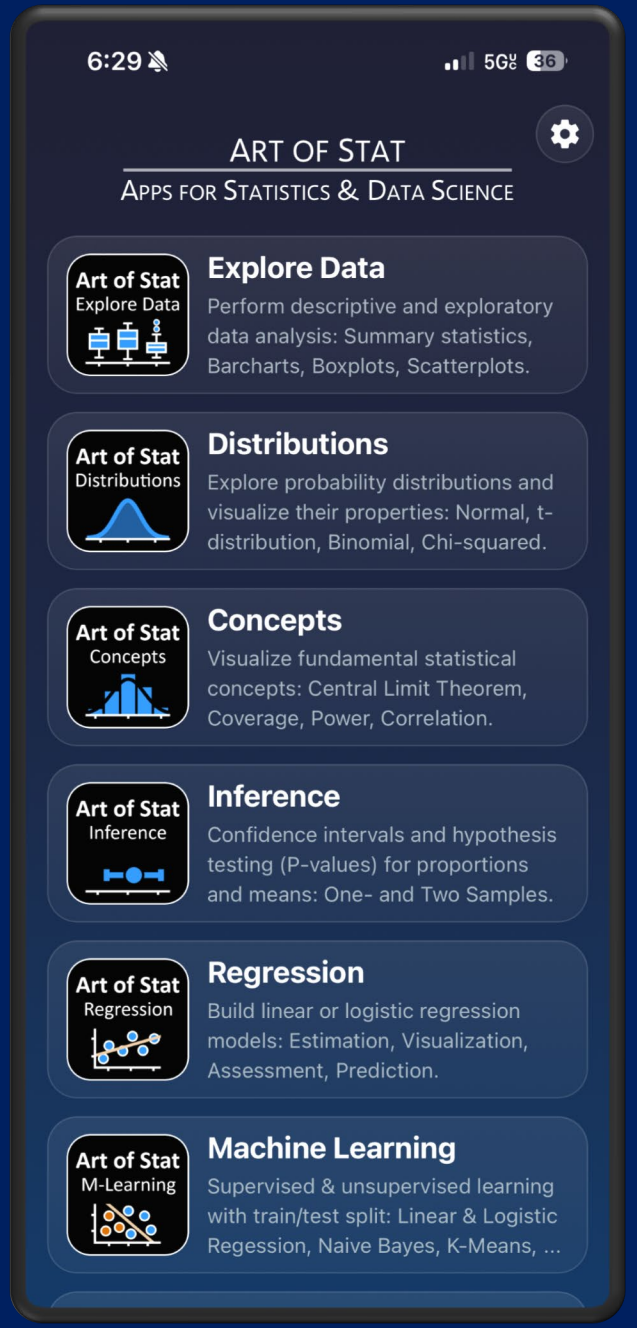

Art of Stat: Machine Learning

- **Supervised Learning**
  - Regression (Linear, Logistic)
  - Discriminant Analysis
  - Naïve Bayes
  - K-Nearest Neighbors
  - Trees, Random Forests

- **Unsupervised Learning**
  - K-Means Clustering

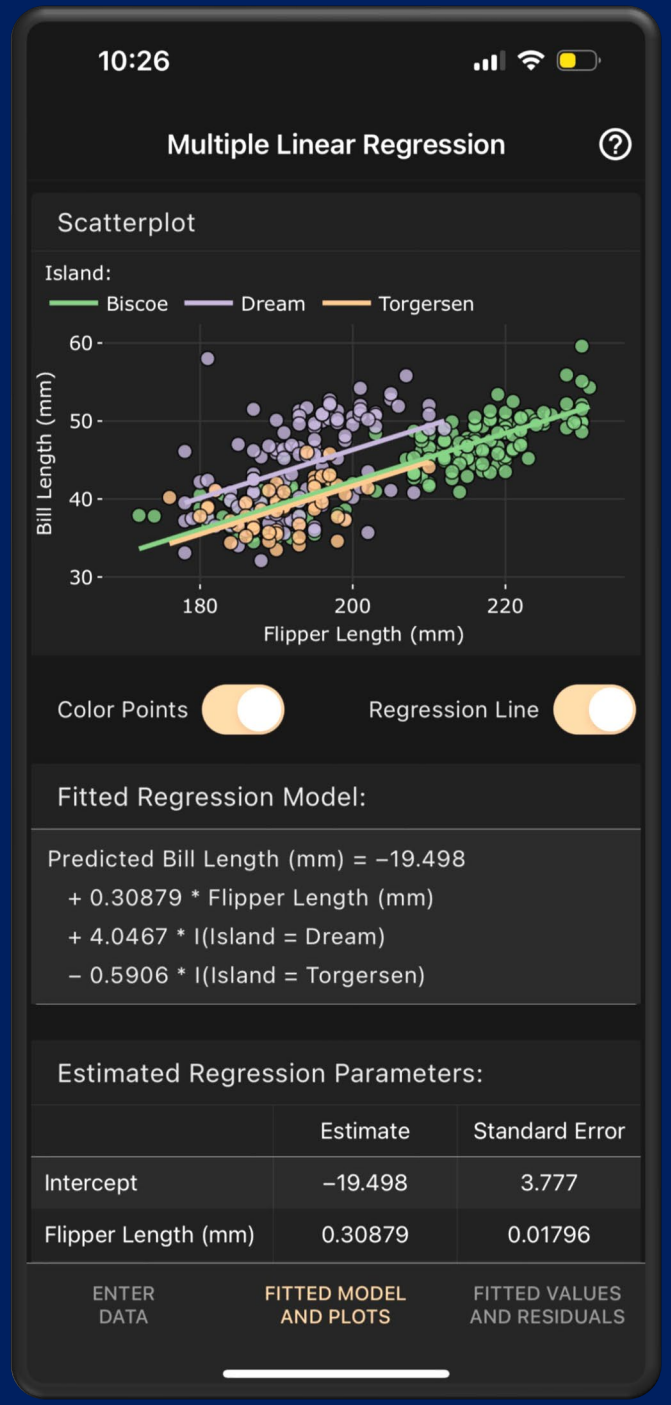

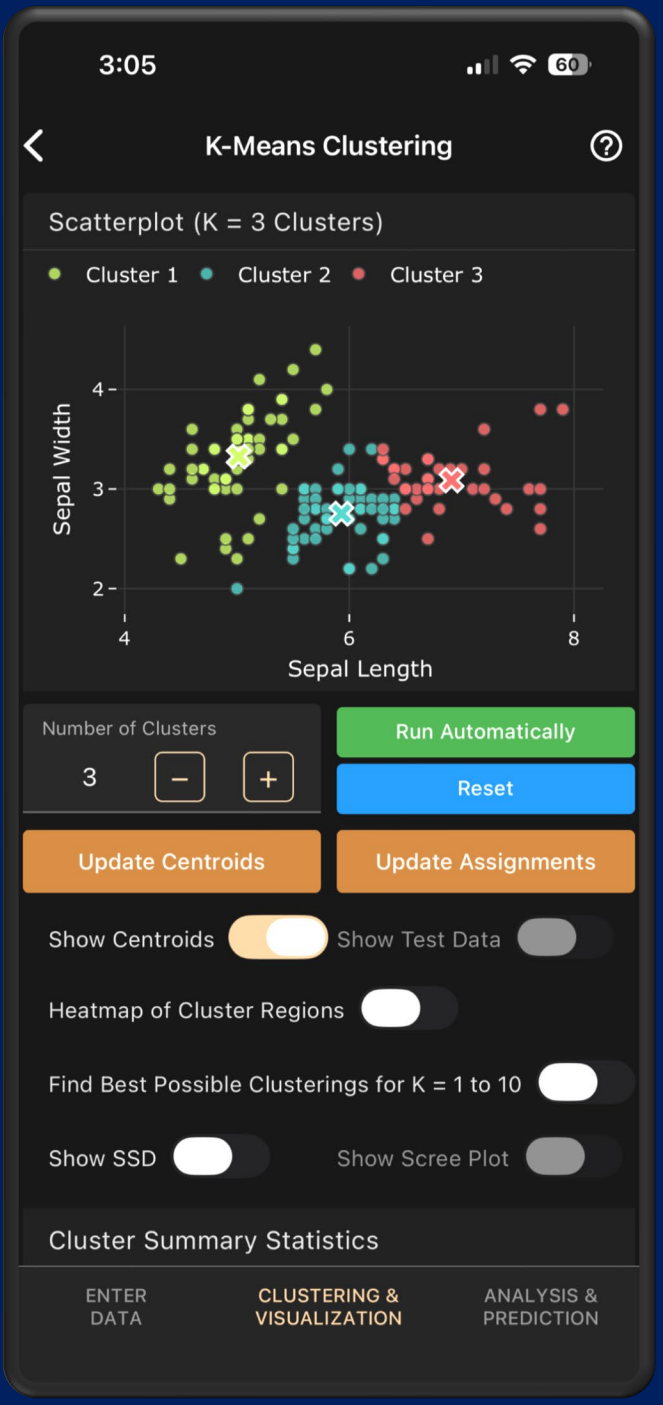

6:29 5G 36

ART OF STAT

APPS FOR STATISTICS & DATA SCIENCE

Explore Data
Perform descriptive and exploratory data analysis: Summary statistics, Barcharts, Boxplots, Scatterplots.

Distributions
Explore probability distributions and visualize their properties: Normal, t-distribution, Binomial, Chi-squared.

Concepts
Visualize fundamental statistical concepts: Central Limit Theorem, Coverage, Power, Correlation.

Inference
Confidence intervals and hypothesis testing (P-values) for proportions and means: One- and Two Samples.

Regression
Build linear or logistic regression models: Estimation, Visualization, Assessment, Prediction.

Machine Learning
Supervised & unsupervised learning with train/test split: Linear & Logistic Regession, Naive Bayes, K-Means, ...

4:26 5G 64

ART OF STAT
Machine Learning

Supervised Learning

Linear Regression
Multiple Linear Regression

Logistic Regression
Multiple Logistic Regression

Linear LDA
Linear Discriminant Analysis

Quadratic LDA
Quadratic Discriminant Analysis

Naive Bayes
Naive Bayes Classifier

5:10 5G 61

Naive Bayes

Data Source:
Open Example Datasets

Choose from several example datasets. Then, select the response (=target) variable which provides the labels and the explanatory variables (=features) you want to include in the naive bayes classification.

Select Dataset:
Palmer Penguins

Penguin	Species	Island	Bill Length (mm)
1	Adelie	Torgersen	39.1
2	Adelie	Torgersen	39.5
3	Adelie	Torgersen	40.3
4	Adelie	Torgersen	36.7
5	Adelie	Torgersen	39.3

Showing all 342 rows, and all 8 columns (swipe left).

Response Variable (Class Labels):
Species

Two Features Selected:
Bill Length (mm) | Flipper Length (mm)

Split into Train and Test Data

ENTER DATA | PLOTS AND CLASSIFICATIONS | ACCURACY & PREDICTIONS

5:12 5G 61

Naive Bayes

Plot of Training Data (Observed Labels)

Species: Adelie Chinstrap Gentoo

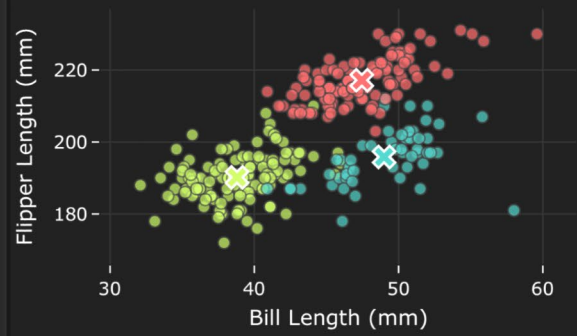

Jitter Points
Plot Test Data
Predicted Labels
Posterior Probabilities
Heatmap of Posterior Probabilities
Class Means
Normal Contours

Prior Class Probabilities:

Species	Prior Probabilities
Adelie	43.8%
Chinstrap	19.3%

ENTER DATA | PLOTS AND CLASSIFICATIONS | ACCURACY & PREDICTIONS

Naïve Bayes

- Classification method, trying to predict class membership based on some features (supervised learning)
- Uses Bayes Theorem and "naïve" independence assumption about joint distribution of features x1 and x2
- Finds Posterior Probability for observation falling in class k:

$$P(Class\ k \mid x1, x2) \propto P(x1, x2 \mid Class\ k) * P(Class\ k)$$
$$= P(x1 \mid Class\ k) * P(x2 \mid Class\ k) * P(Class\ k)$$

- Need prior class probabilities $P(Class\ k)$ and distribution of feature in a given class $P(x1 \mid Class\ k)$
- Predicted class is the one with largest posterior probability

Example: Naïve Bayes

Tab 1:

- Load Data
- Select Target
- Select Predictors
- Train/Test Split
- Descriptive Stats

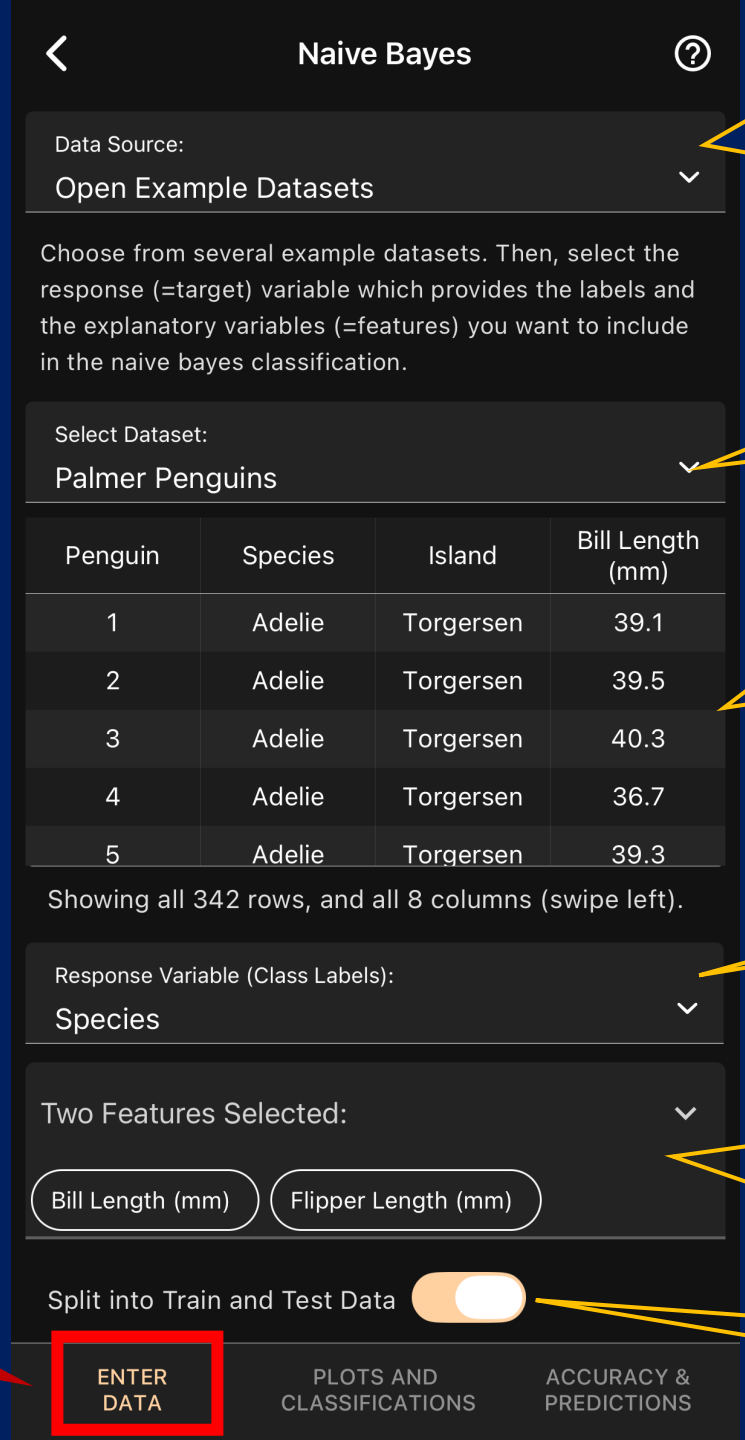

Pick a Data Source:
- Pre- Implemented Examples
- Upload your own CSV/Excel
- Data Editor

Pick the particular dataset

Scroll through the dataset

Pick target variable

Select Features:
- continuous
- categorical

Split into testing & training set

Tab 1

Example: Naïve Bayes

Tab 1:

- Load Data
- Select Target
- Select Predictors
- **Train/Test Split**
- Descriptive Stats

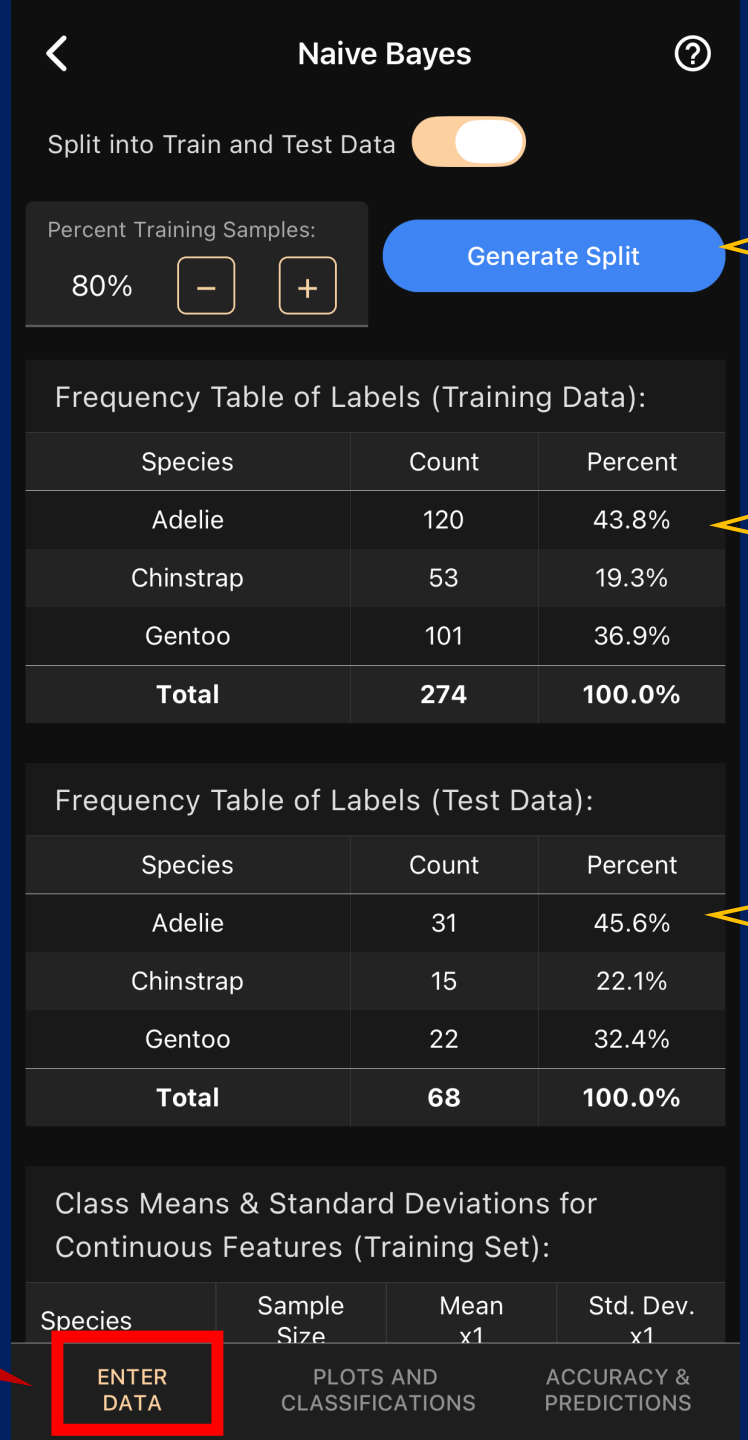

Generate 80/20 split by pressing button

See distribution of labels in training set

See distribution of labels in test set

Tab 1

Example: Naïve Bayes

Tab 1:

- Load Data
- Select Target
- Select Predictors
- Train/Test Split
- **Descriptive Stats**

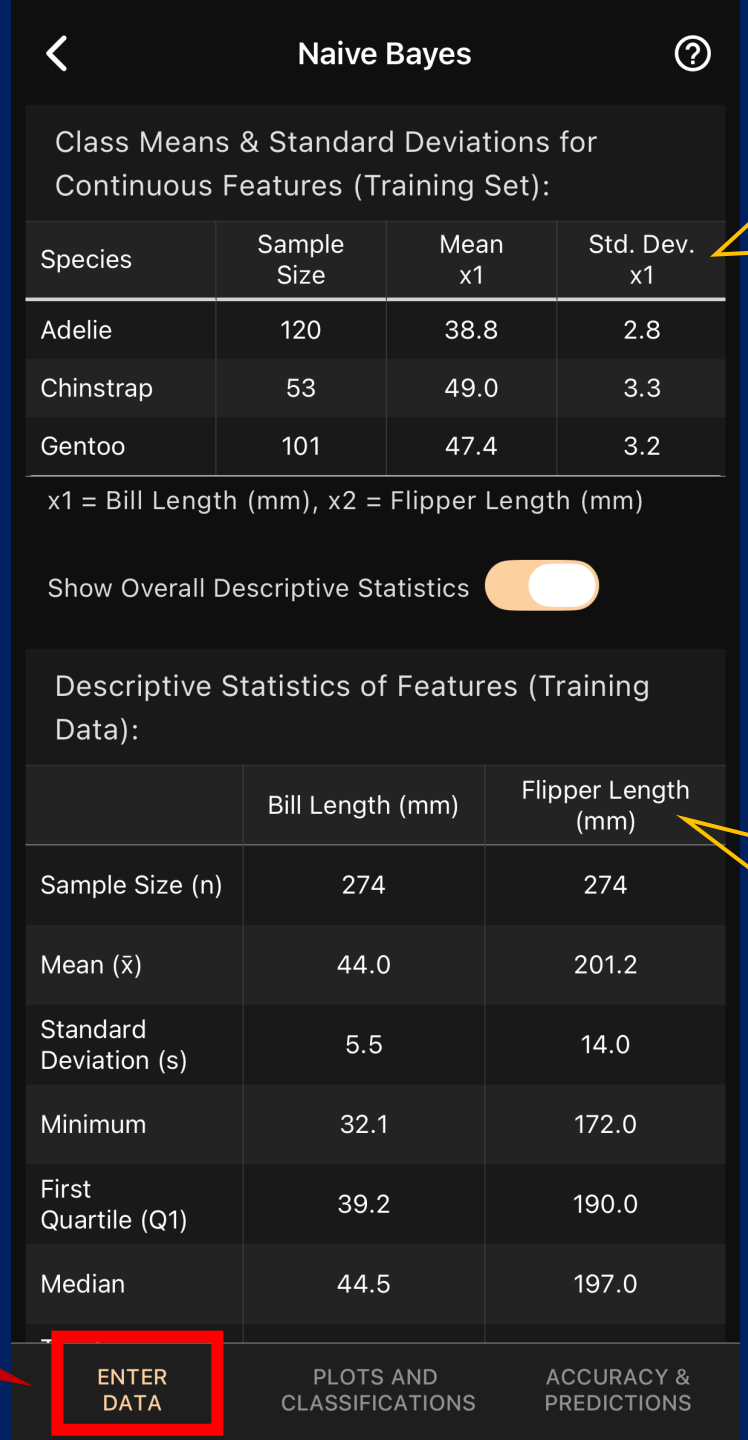

Tab 1

Mean and standard deviation for each continuous predictor, stratified by label.

If we had categorical predictors (e.g., Sex or Island), app would show distribution of categories for each label.

See descriptive stats for each continuous predictor (not stratified)

Example: Naïve Bayes

Tab 2:

- Scatterplot
- Class Means
- Posterior Odds
- Predictions
- Heatmap
- Contours

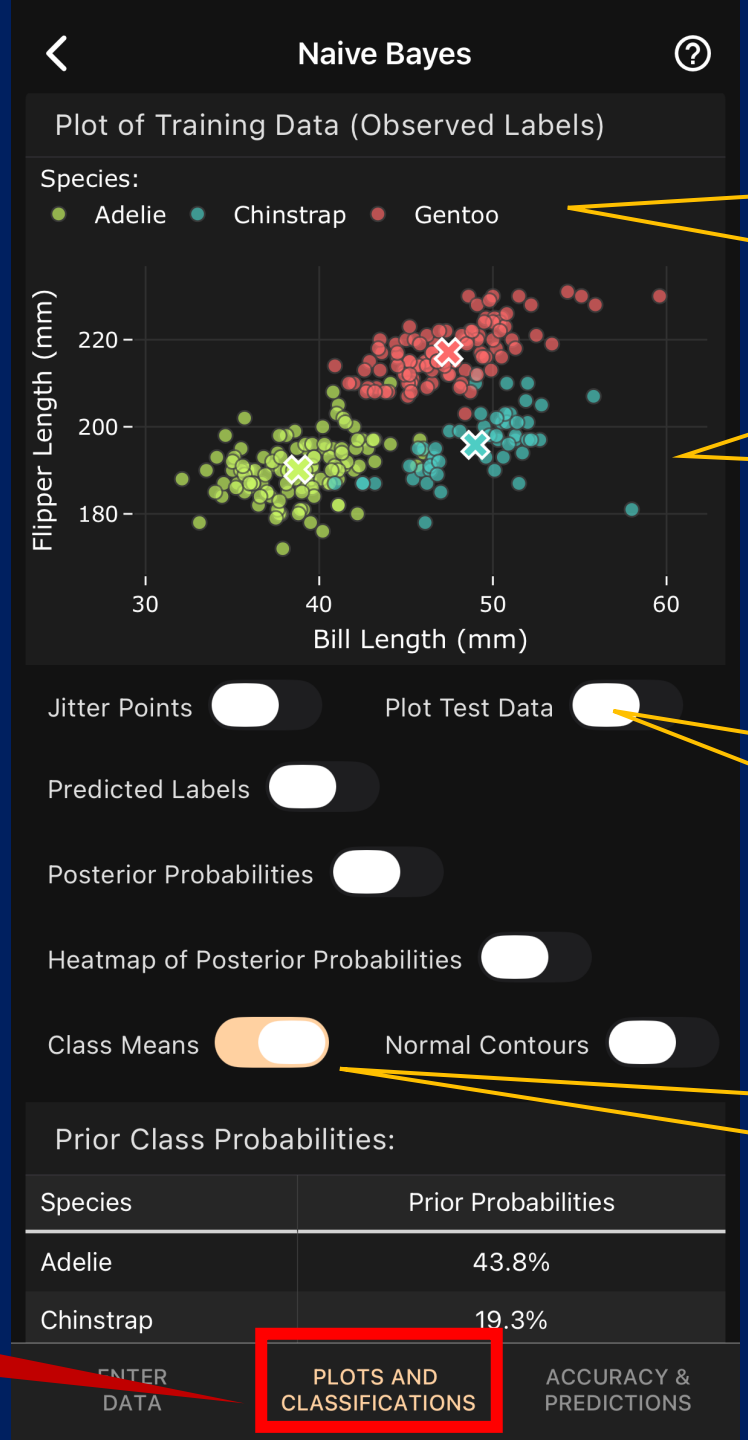

Different target labels are represented by different colors (plot interactive)

Scatterplot with the two features on the x - and y-axis

Request to plot test data instead of training data

Show the mean for each target label

Tab 2

Example: Naïve Bayes

Tab 2:

- Scatterplot
- Class Means
- **Posterior Odds**
- Predictions
- Heatmap
- Contours

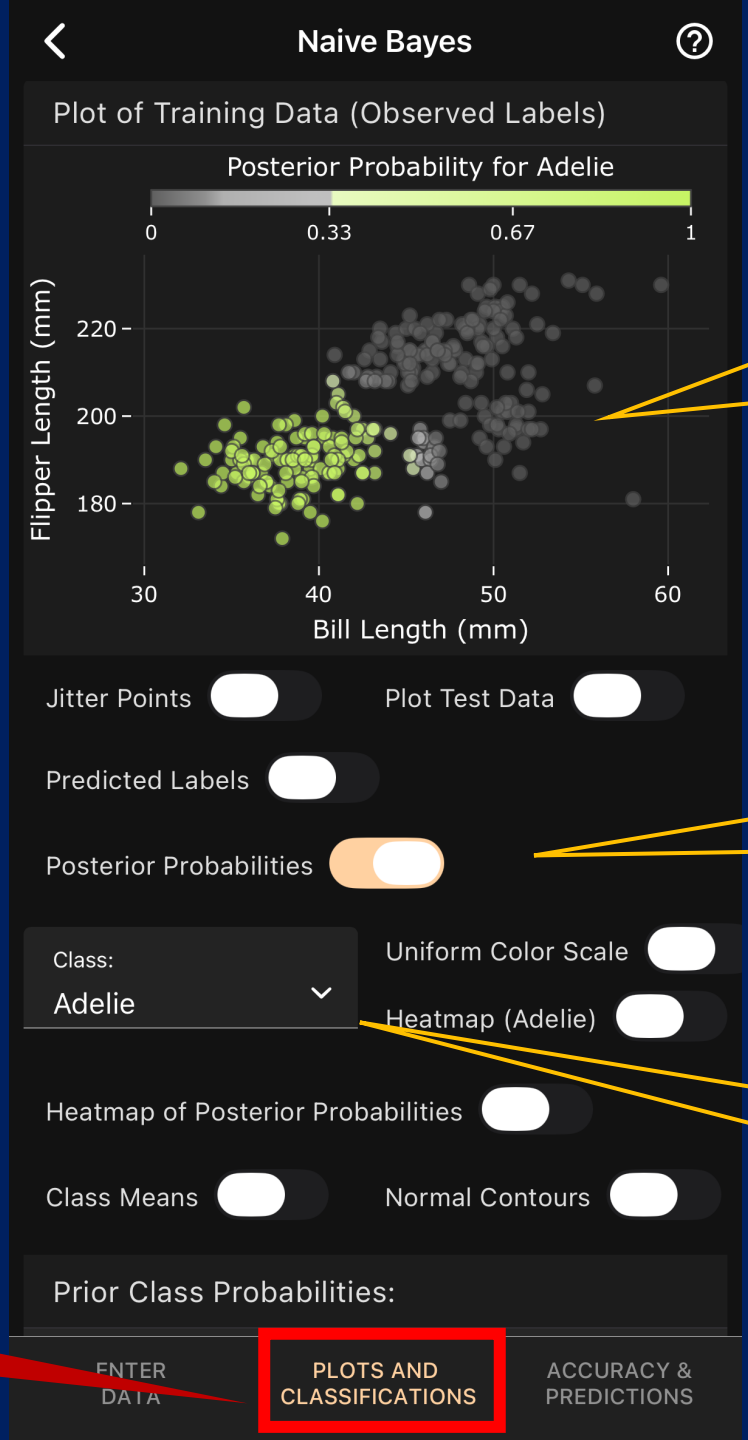

Dots colored by their posterior probability of belonging to species “Adelie”

Request to show posterior odds of class membership

Select which class you want to show posterior odds for

Tab 2

Example: Naïve Bayes

Tab 2:

- Scatterplot
- Class Means
- Posterior Odds
- **Predictions**
- Heatmap
- Contours

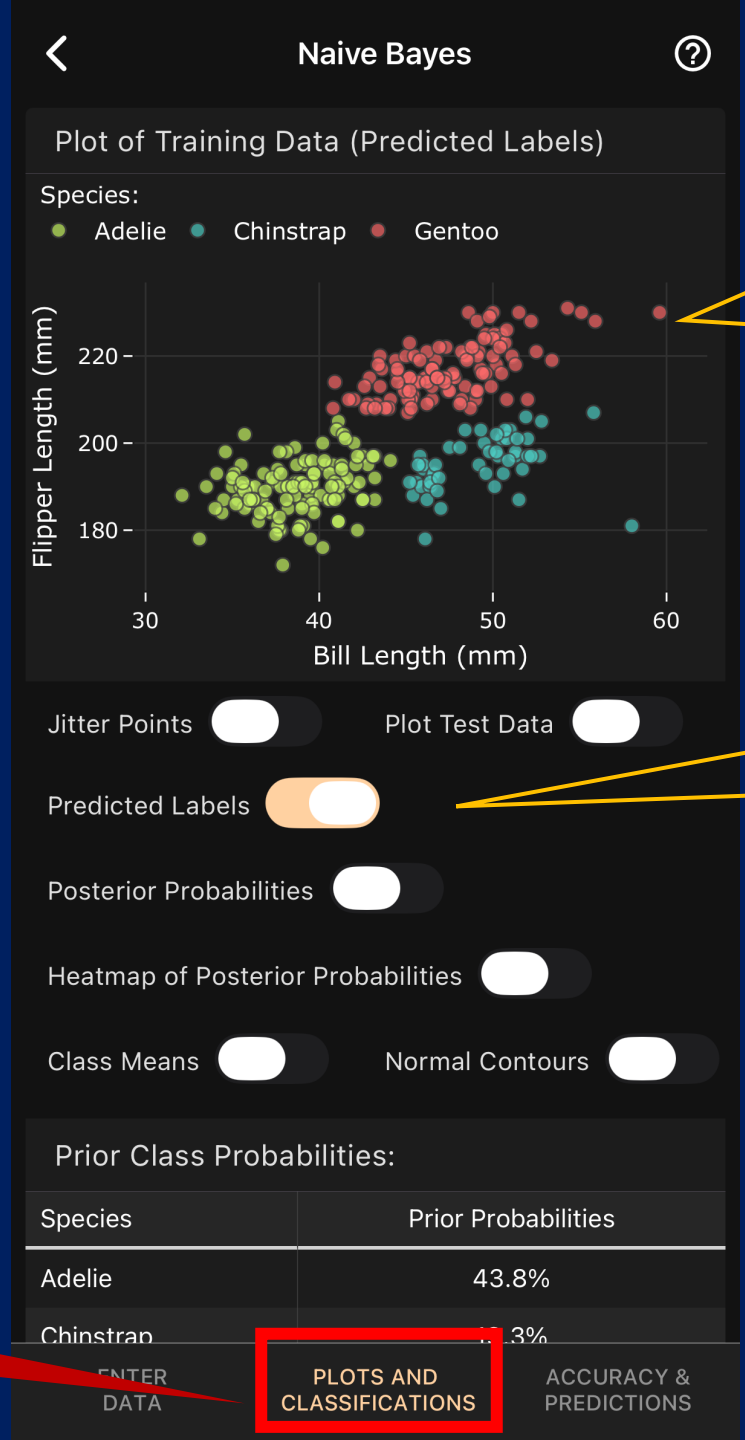

Scatterplot with points colored by predicted label (training set)

Switch to show the predicted label, based on maximum posterior odds

Tab 2

Example: Naïve Bayes

Tab 2:

- Scatterplot
- Class Means
- Posterior Odds
- Predictions
- **Heatmap**
- Contours

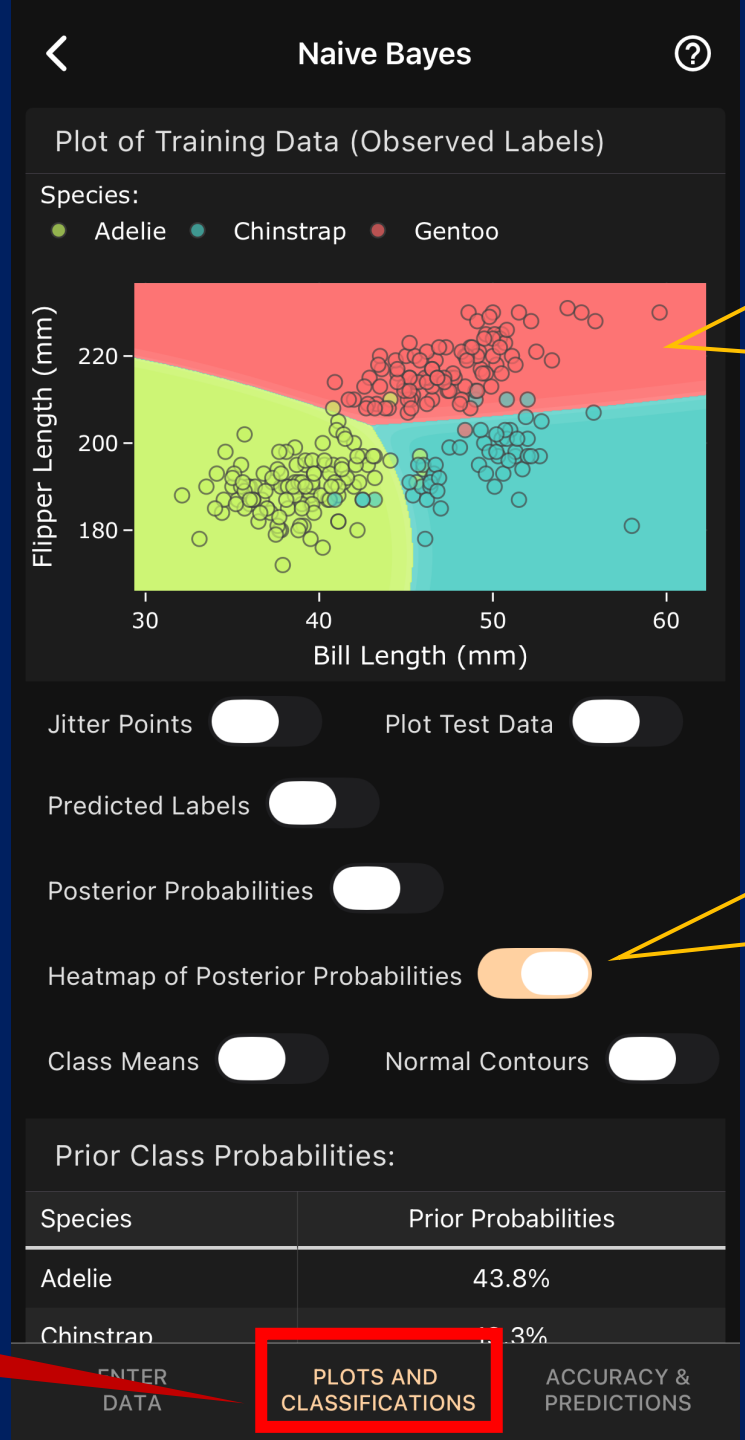

Tab 2

Scatterplot with points colored by original label but area colored by predicted label. (Training set)

Switch to show heatmap of posterior odds. (Works when there are only continuous features.)

Example: Naïve Bayes

Tab 3:

- Dataset with predictions, posterior probs and odds (not shown)
- **Find Prediction**
- Confusion Matrix
- Performance

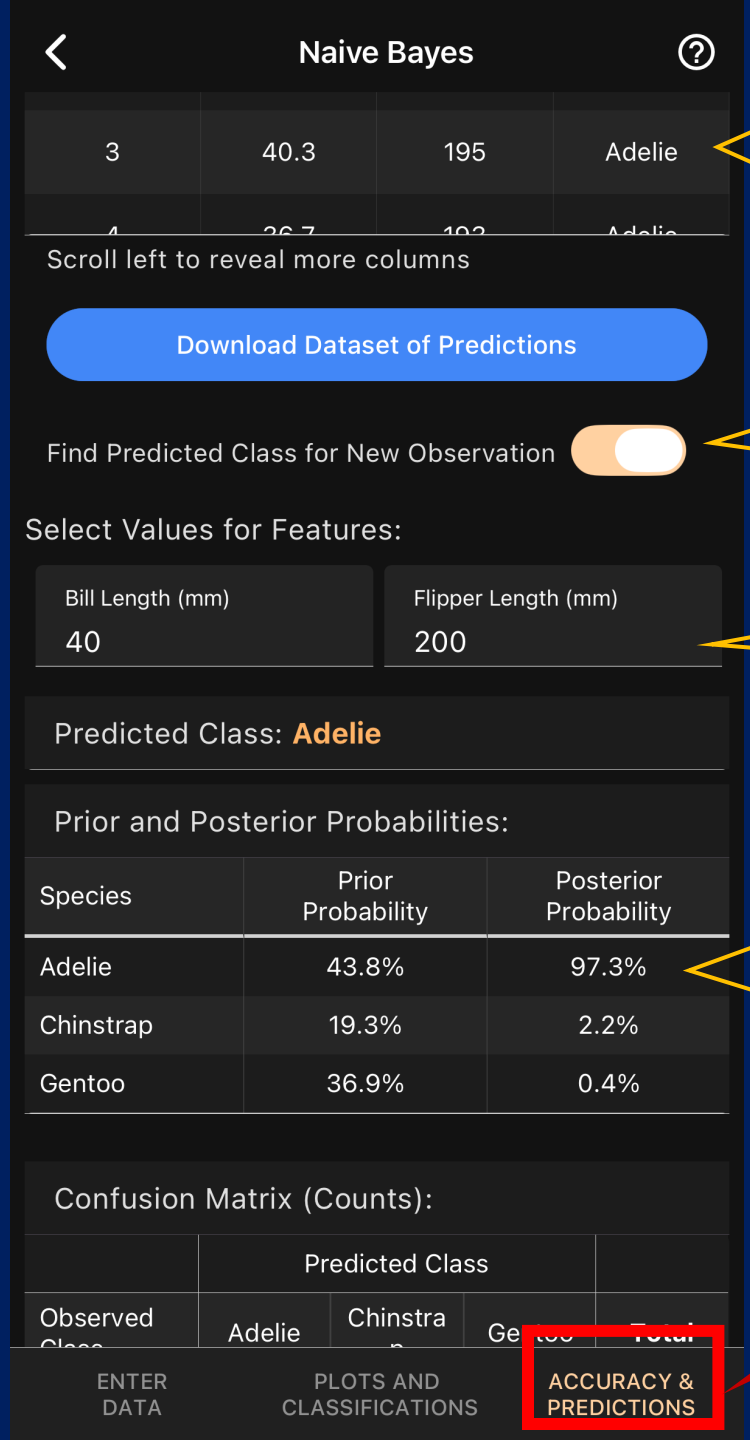

Dataset (train and test) with posterior probabilities and odds (only shown partially in this screenshot)

Ask for class prediction of **new observation**

Enter values for features

Table showing, for each class, the posterior probabilities for each class, based on the supplied feature values

Tab 3

Example: Naïve Bayes

Tab 3:

- Dataset with predictions, posterior probs . and odds
- Find Prediction
- **Confusion Matrix**
- Performance

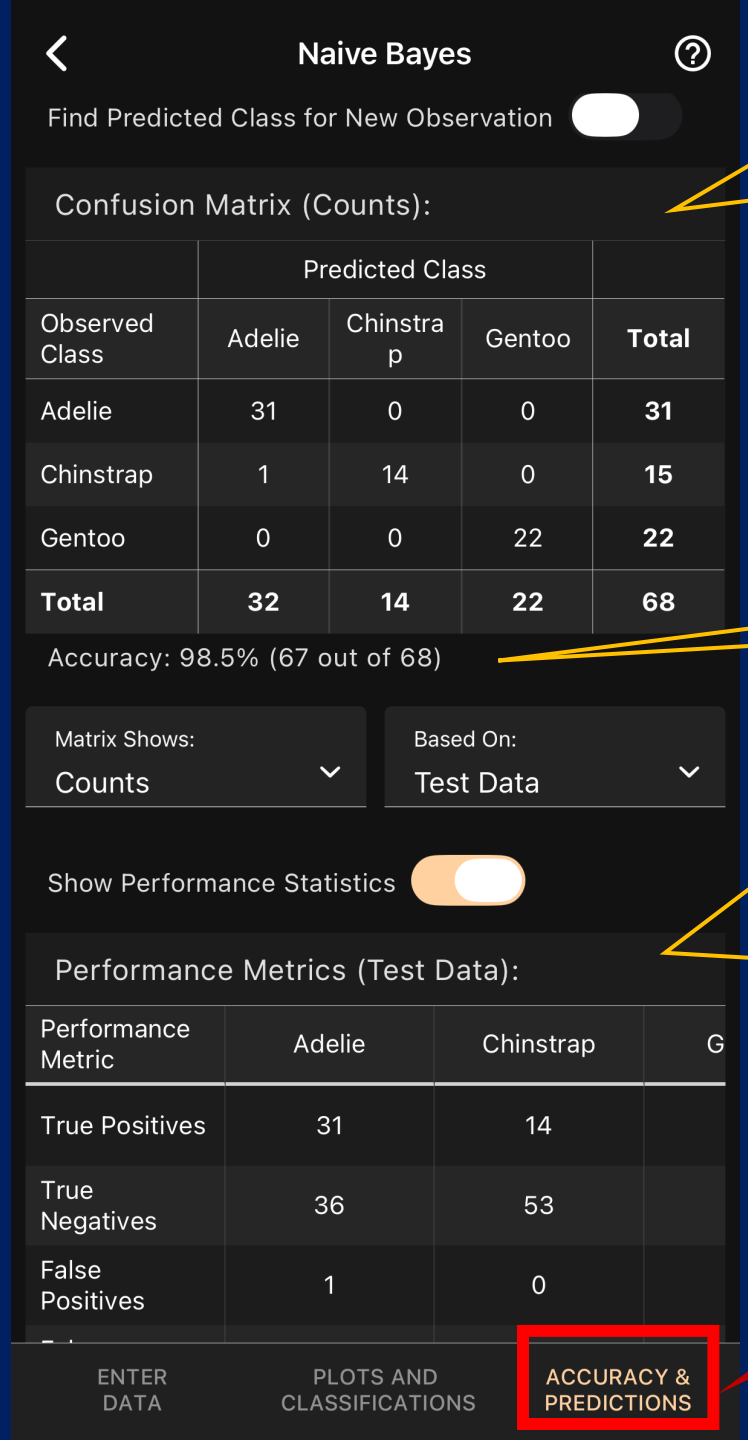

3x3 Confusion Matrix comparing observed and predicted labels

Overall test accuracy

Table showing, for each class, true positives, true negatives, false positives and false negatives, sensitivity, specificity, precision, recall and F1 score.

Tab 3

Example: Naïve Bayes

Tab 3:

- Dataset with predictions, posterior probs . and odds
- Find Prediction
- Confusion Matrix
- **Performance**

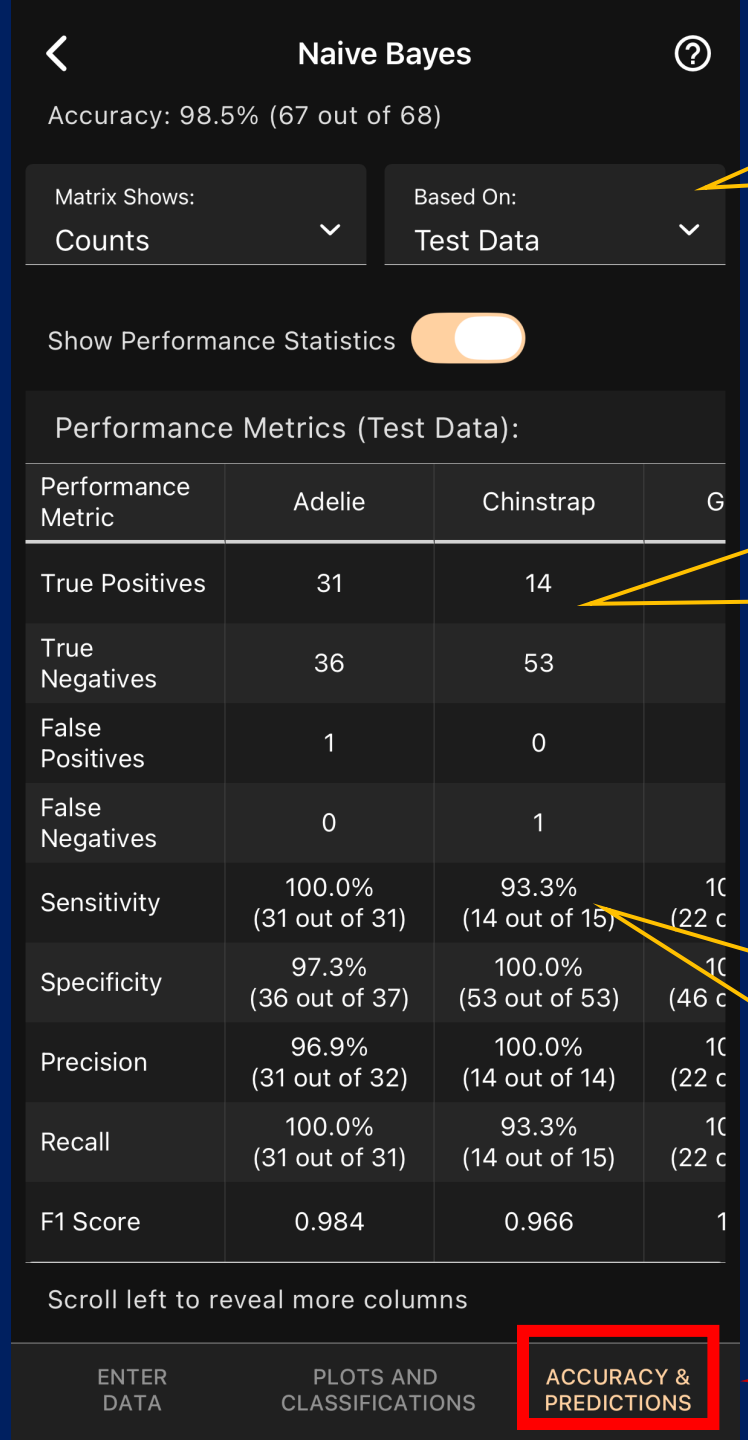

Requested performance statistics for test data

Table showing, for each class, true positives, true negatives, false positives and false negatives.

Also showing resulting sensitivity, specificity, precision, recall and F1 score.

Third tab

Run-Through of the K-Means Clustering App (Go Live)

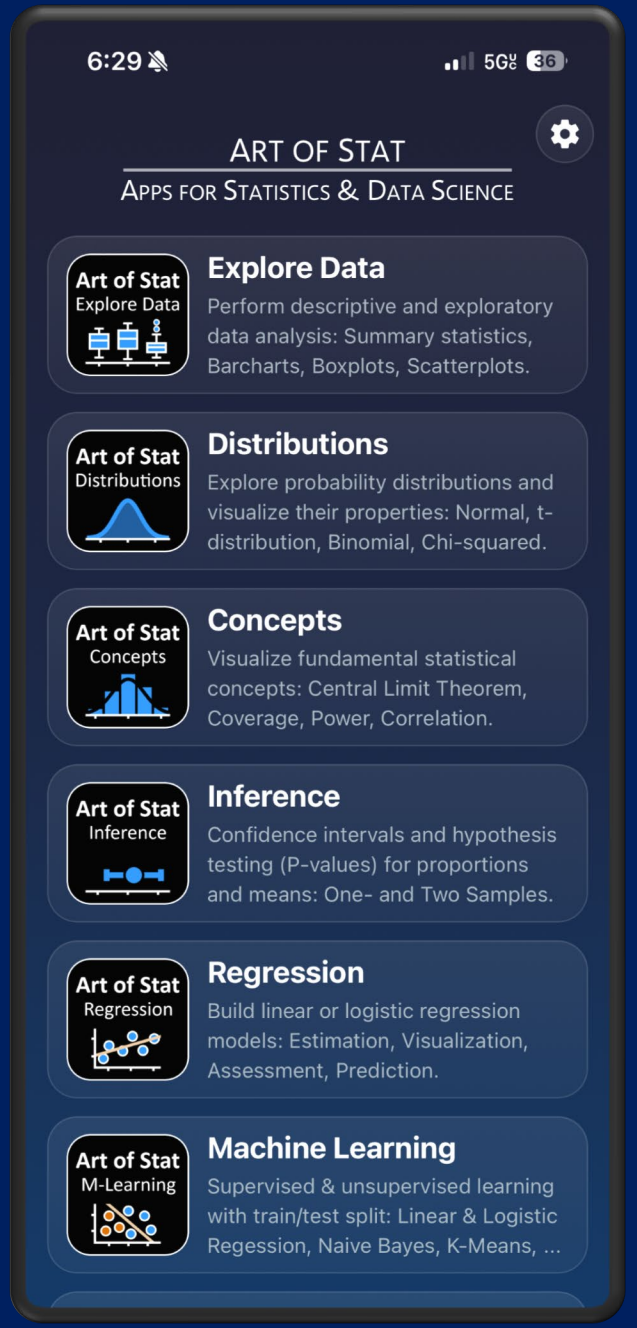

Example: K-Means Clustering

Tab 2:

- Scatterplot
- Select K
- Carry out clustering
- Heatmap
- Screeplot
- Cluster Summaries

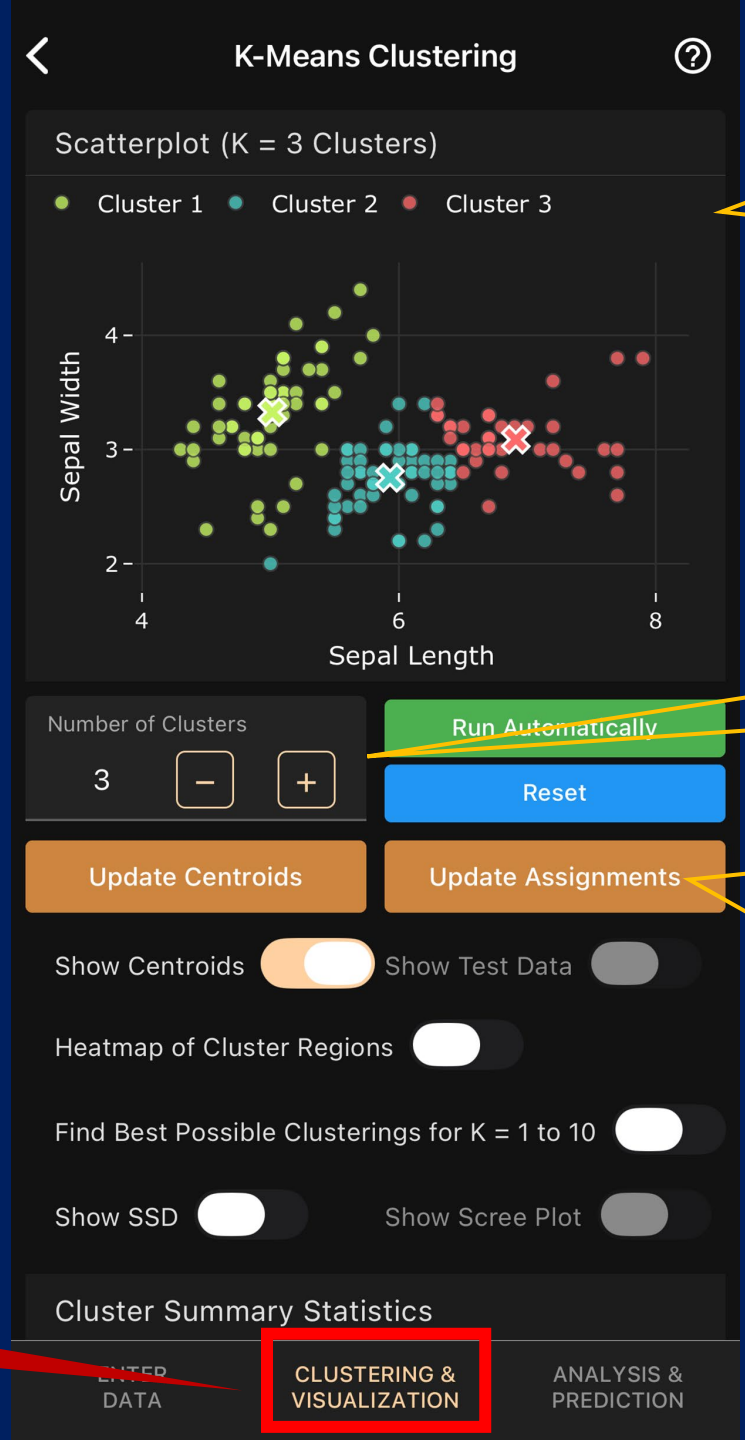

Scatterplot of two features to be clustered on

Select the number of clusters (K)

Step through clustering step - by - step, iterating between finding centroids and then updating assignment.

Tab 2

Run-Through of the Concepts App (Go Live)

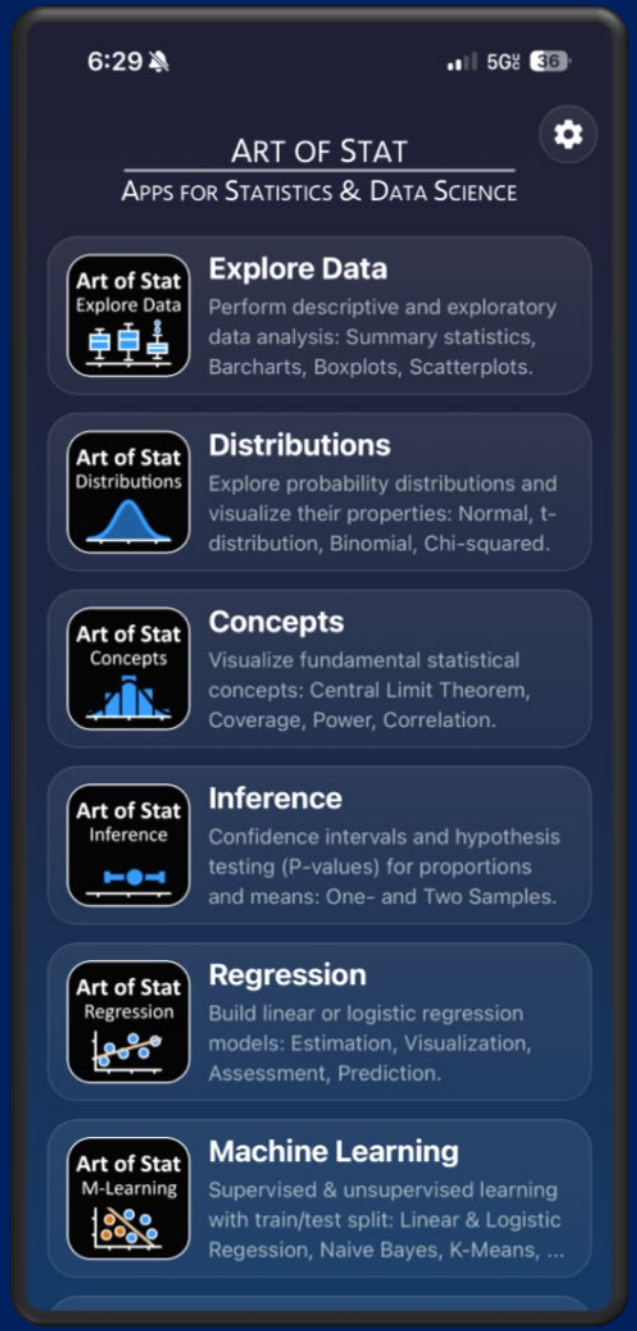

Art of Stat: Concepts

Central Limit Theorem for Means

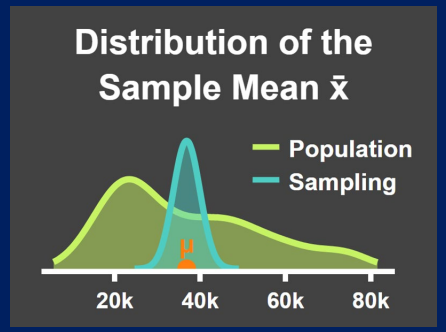

Screens:
- Population Distribution
- Draw Samples
- Sampling Distribution

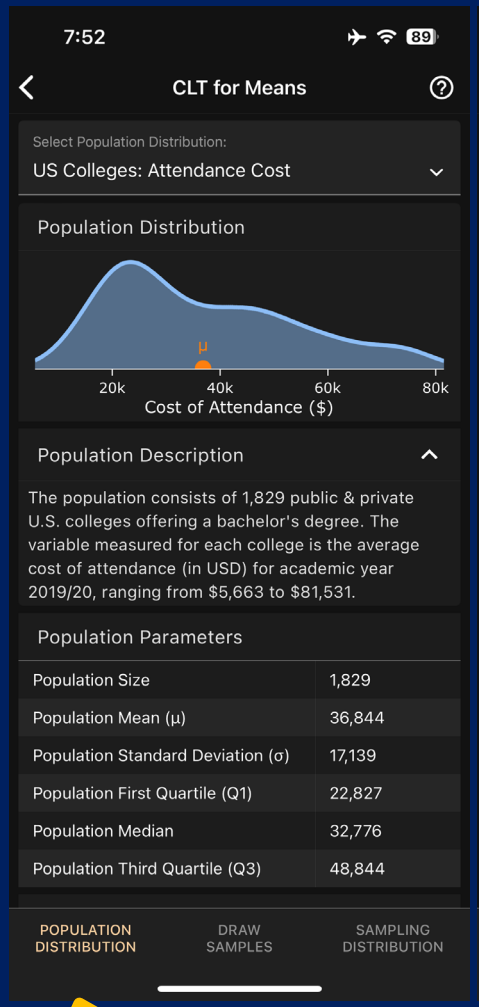

Explore real population distribution, and its parameters (mean, standard deviation).

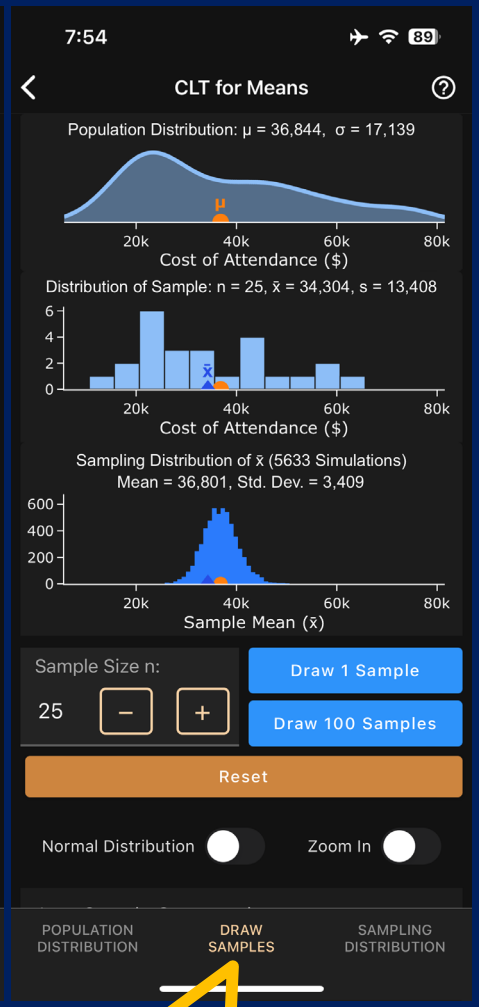

Draw samples of a given size from this population distribution, and keep track of the sample means

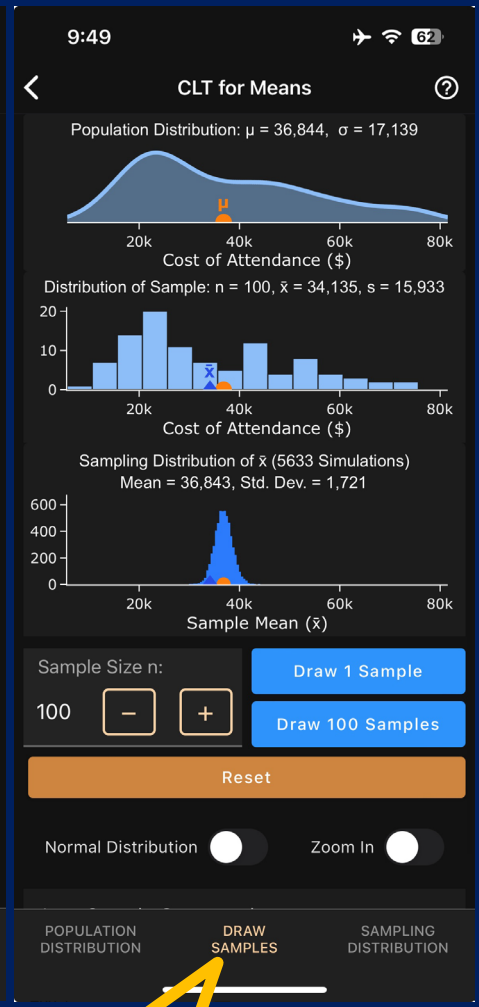

See what happens as the sample size increases

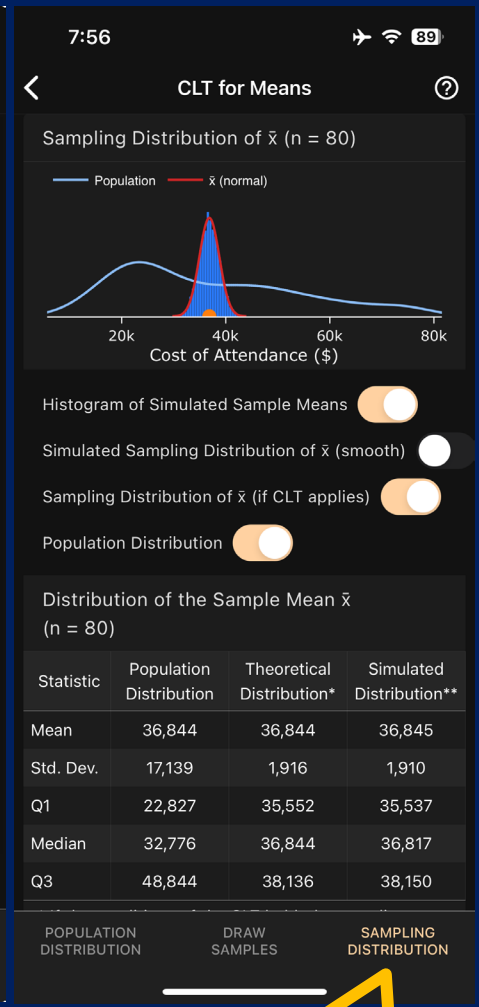

Show (simulated) sampling distribution, compare to population distribution

Run-Through of the Multiple Linear Regression App (Go Live)

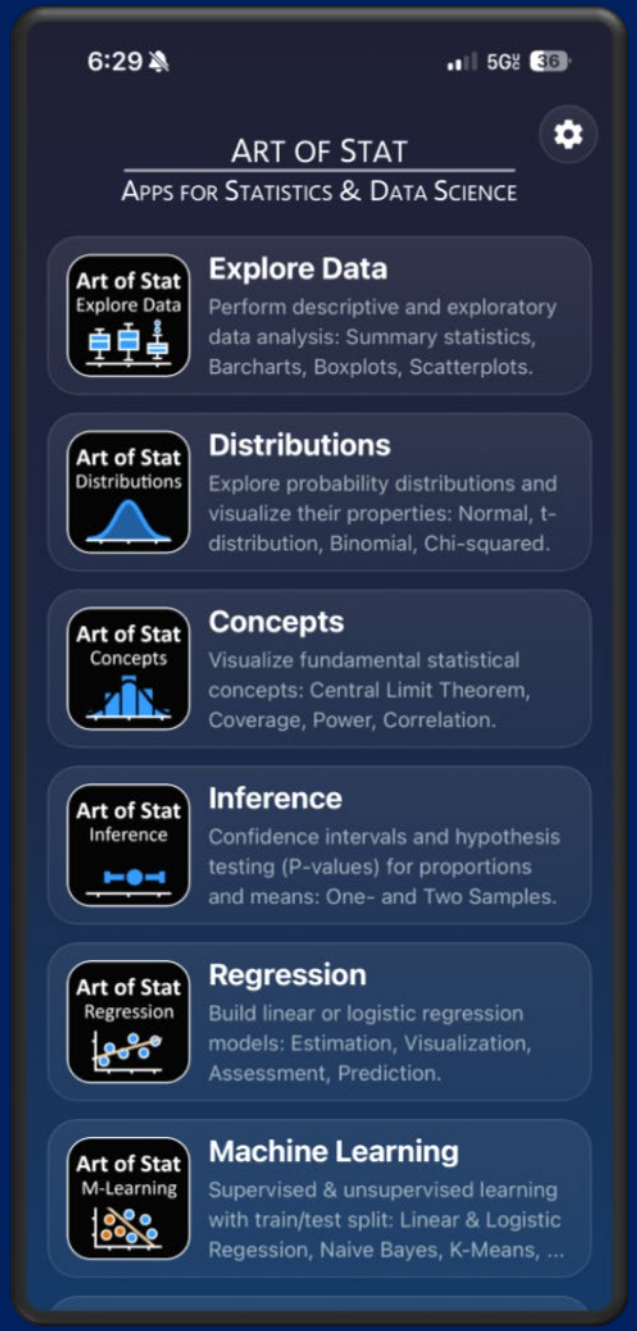

Art of Stat: Regression

Multiple Linear Regression

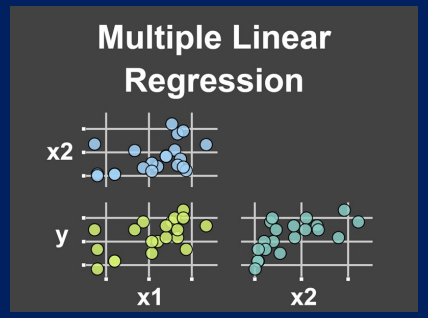

Screens:
- Enter Data
- Fitted Model & Plots
- Fitted Values, Residuals

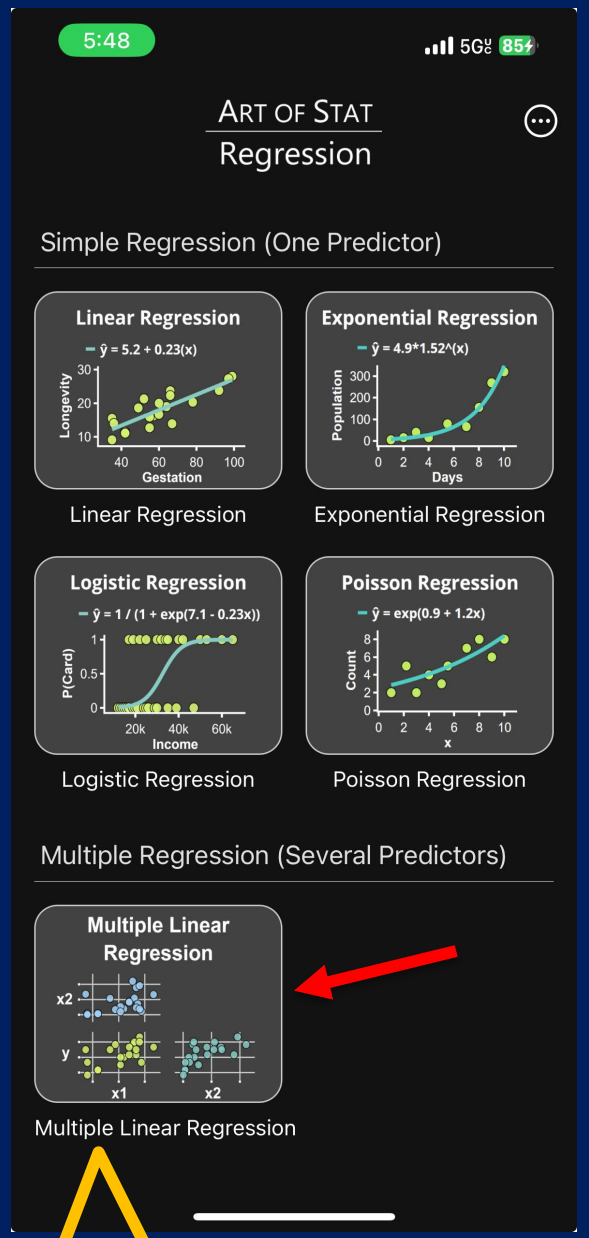

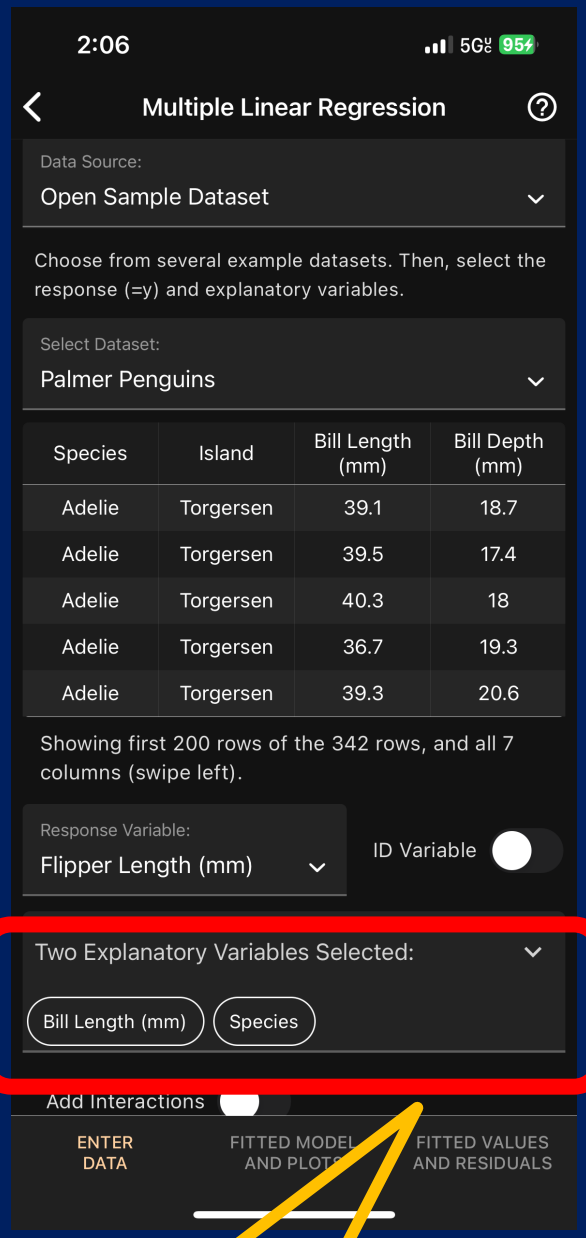

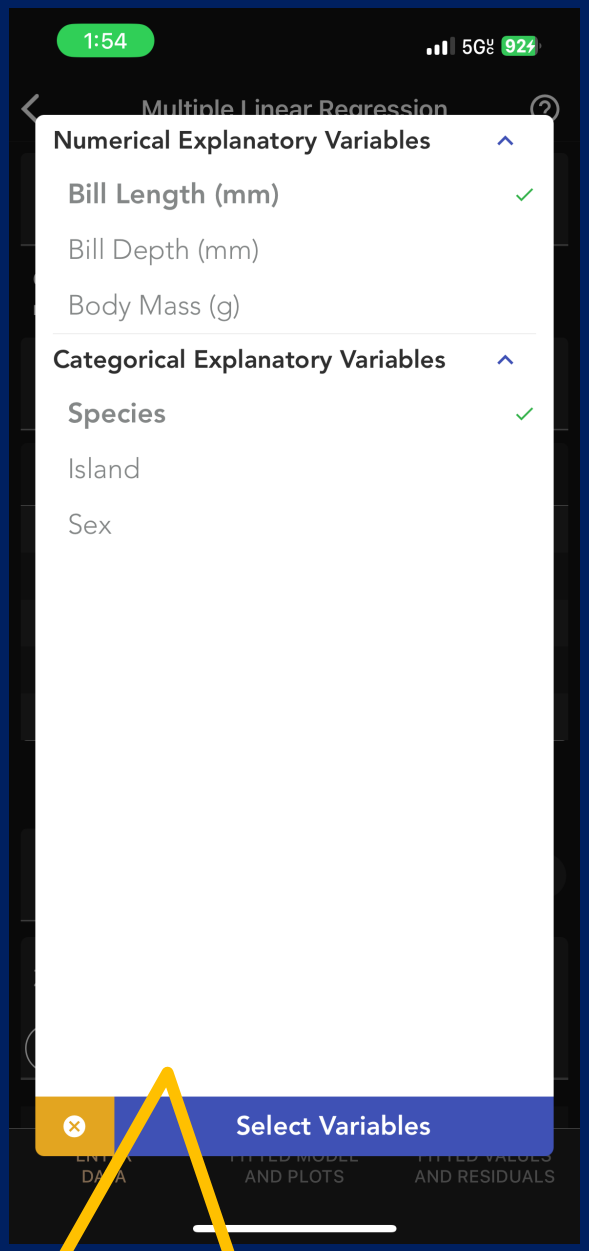

Select Multiple Linear Regression

Select response & explanatory variables

Click on variables you want to select, continuous or categorical

Art of Stat: Regression

Multiple Linear Regression

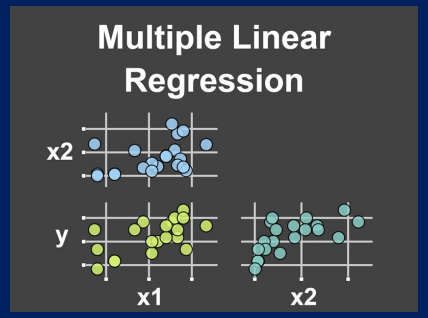

Screens:
- Enter Data
- Fitted Model & Plots
- Fitted Values, Residuals

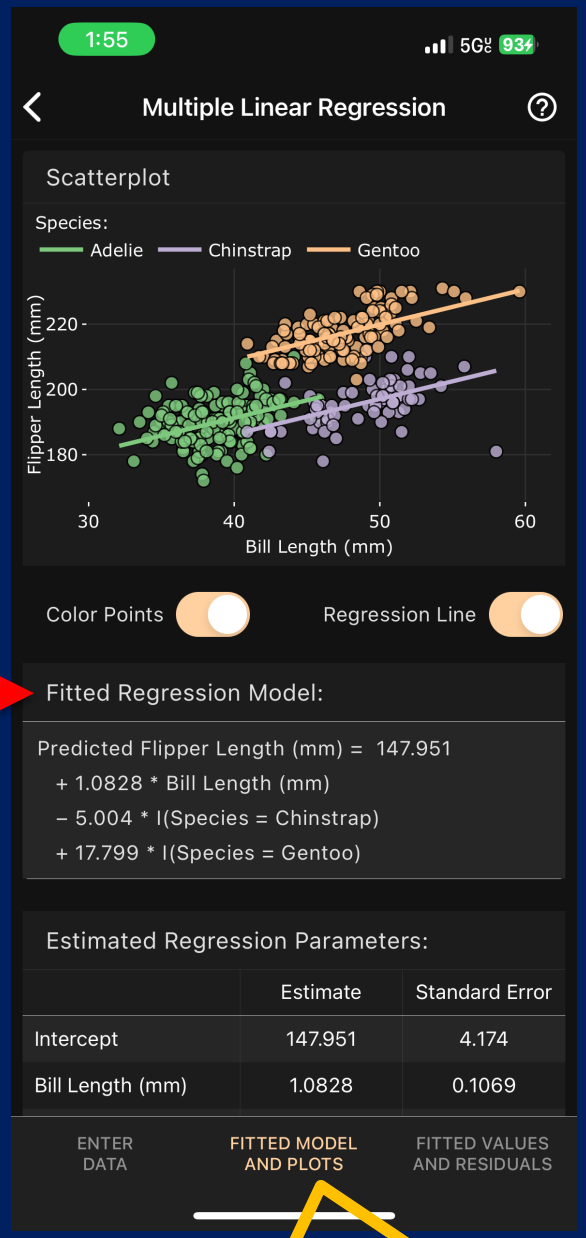

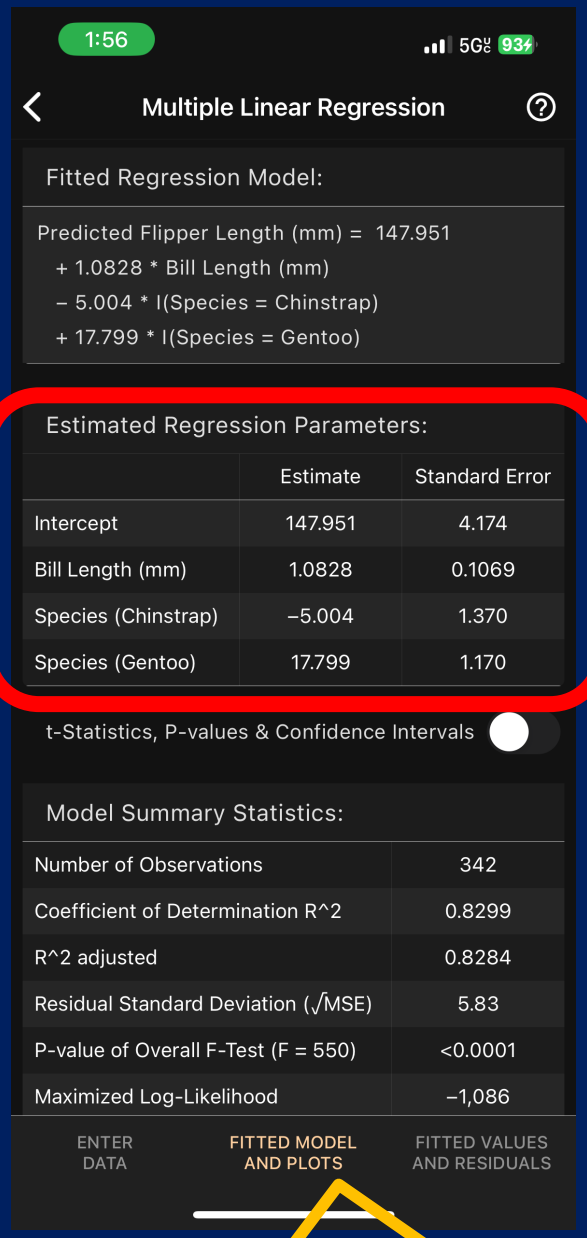

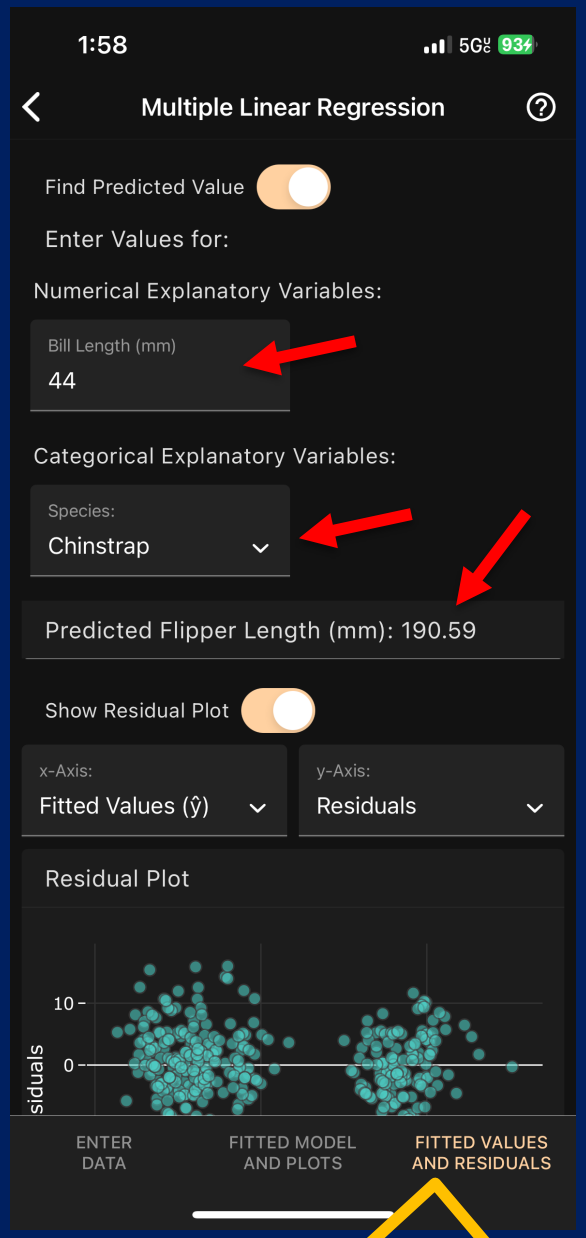

Three separate lines, and equation for fitted model

Table of regression coefficient, plus std. errors and P-values

Prediction: drop-downs for categorical predictors

Art of Stat: Regression

Multiple Linear Regression

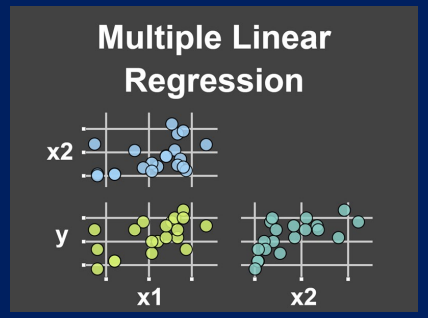

Screens:
- Enter Data
- Fitted Model & Plots
- Fitted Values, Residuals

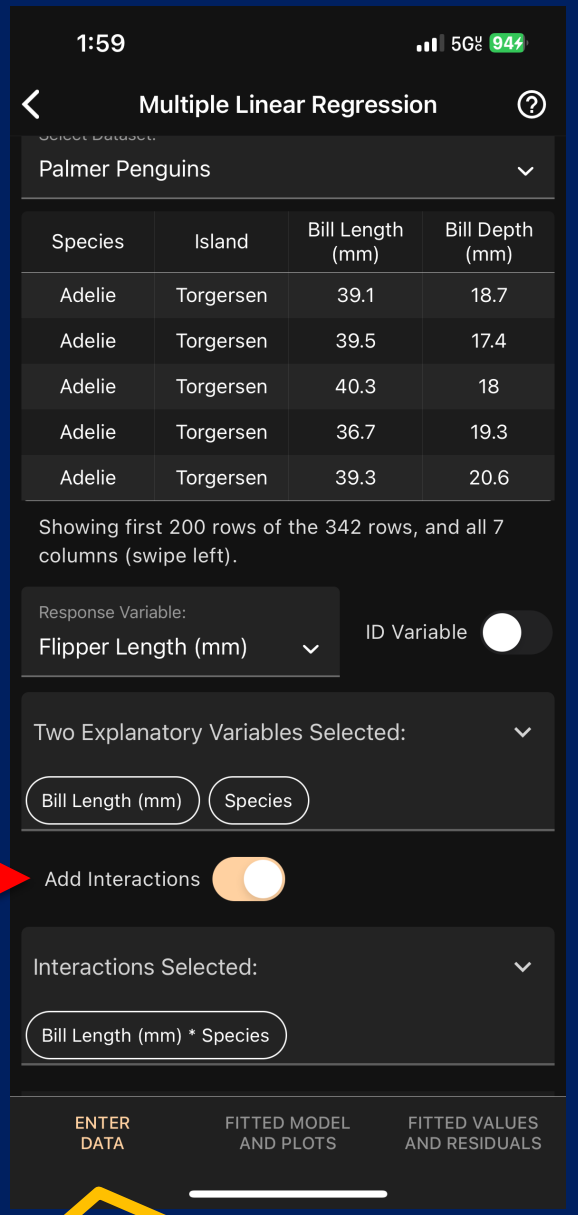

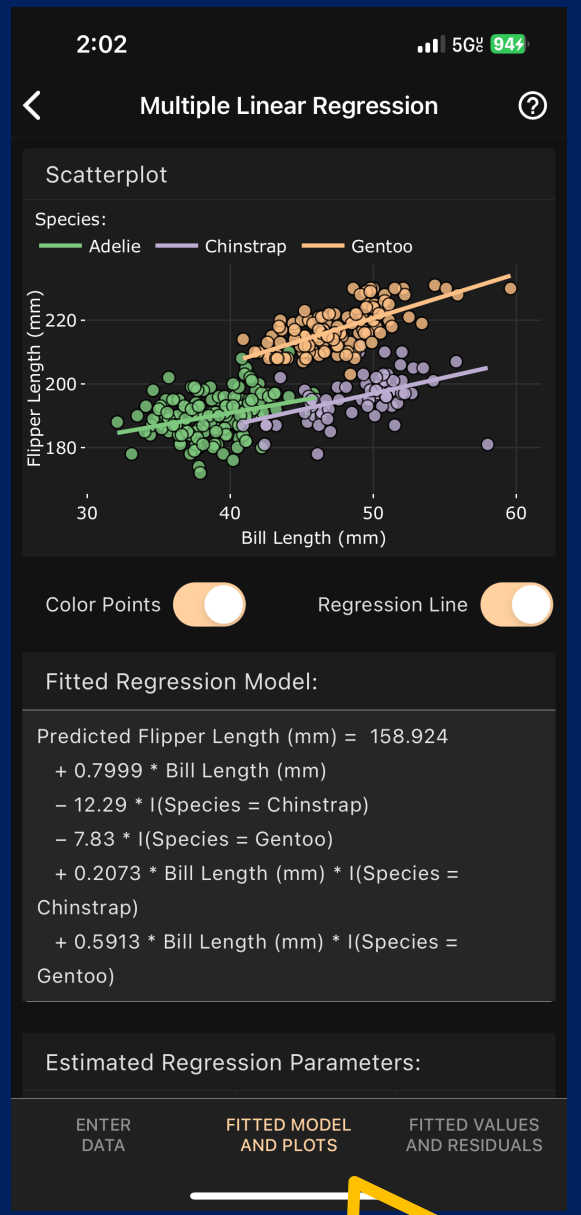

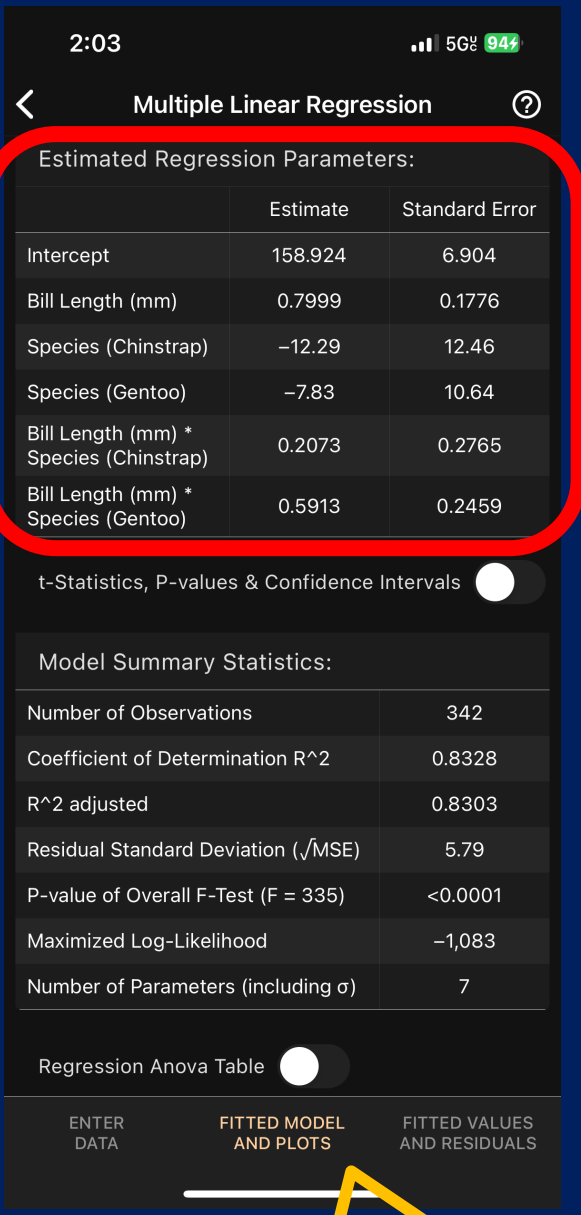

Multivariate Thinking also means understanding **Interactions**

Now we have different slopes: different effects

Regression coefficient table easy to read

Homework Assignment Examples

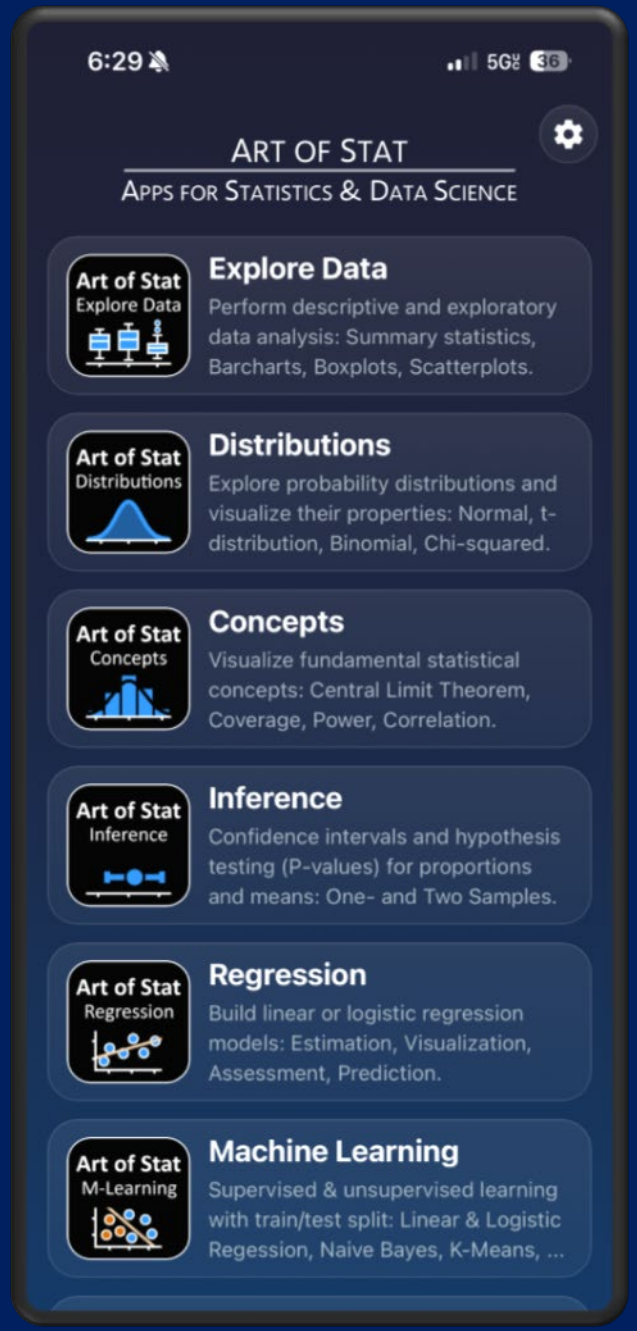

STAT 161: HW2

Reading: Read Chapter 2, Sections 2.1 to 2.4.

1. Eyewitness testimony

Read this story, which ran on NPR in January 2016. The original article can be found at this link. Under Supporting Information, the article gives a link to the data used for this article.

a. Download these data and save the data (as a .csv file) to your computer. (Nothing to turn in.)
b. Open the data set in Excel, Google Sheets or some other spreadsheet program to take an initial look at the data. Clean the data as you see fit. For instance, it is a very good idea to delete the first row and any special characters! Save the cleaned data as a .csv file. (Nothing to turn in.)
c. Choose one categorical variable from this dataset. (Nothing to turn in.)
d. **Email** the cleaned csv file to yourself (then download it from your email to your device) **or** put the cleaned csv file on your **Google Drive/iCloud**. Open the Art of Stat app, go to Explore Data and from there open an appropriate procedure to describe the distribution of the variable. What to turn in:
    i. A description of the distribution of the variable, including an appropriate graph, table and a short paragraph. Use screenshots from the app to include in your description.
e. Now, turn to the internet or, better, generative AI and ask "*How to open a .csv file in Stata*". I tried this with chatGPT and got very useful answers I pasted below, but I rather have you try on

PLOS ONE

RESEARCH ARTICLE

The Novel New Jersey Eyewitness Instruction Induces Skepticism but Not Sensitivity

Athan P. Papailiou[1☯]*, David V. Yokum[1,2☯], Christopher T. Robertson[1]

1 James E. Rogers College of Law, University of Arizona, Tucson, Arizona, United States of America,
2 Department of Psychology, University of Arizona, Tucson, Arizona, United States of America

☯ These authors contributed equally to this work.
* athan@email.arizona.edu

Abstract

In recent decades, social scientists have shown that the reliability of eyewitness identifications is much worse than laypersons tend to believe. Although courts have only recently begun to react to this evidence, the New Jersey judiciary has reformed its jury instructions to notify jurors about the frailties of human memory, the potential for lineup administrators to nudge witnesses towards suspects that they police have already identified, and the advantages of alternative lineup procedures, including blinding of the administrator. This experiment tested the efficacy of New Jersey's jury instruction. In a 2×2 between-subjects design, mock jurors ($N = 335$) watched a 35-minute murder trial, wherein identification quality was either "weak" or "strong" and either the New Jersey or a "standard" instruction was delivered. Jurors were more than twice as likely to convict when the standard instruction was used (OR = 2.55; 95% CI = 1.37–4.89, $p < 0.001$). The New Jersey instruction, however, did *not* improve juror's ability to discern quality; rather, jurors receiving those instructions indiscriminatingly discounted "weak" and "strong" testimony in equal measure.

Citation: Papailiou AP, Yokum DV, Robertson CT (2015) The Novel New Jersey Eyewitness Instruction Induces Skepticism but Not Sensitivity. PLoS ONE 10 (12): e0142695. doi:10.1371/journal.pone.0142695

Editor: Francesco Di Russo, University of Rome,

One Student Solution:

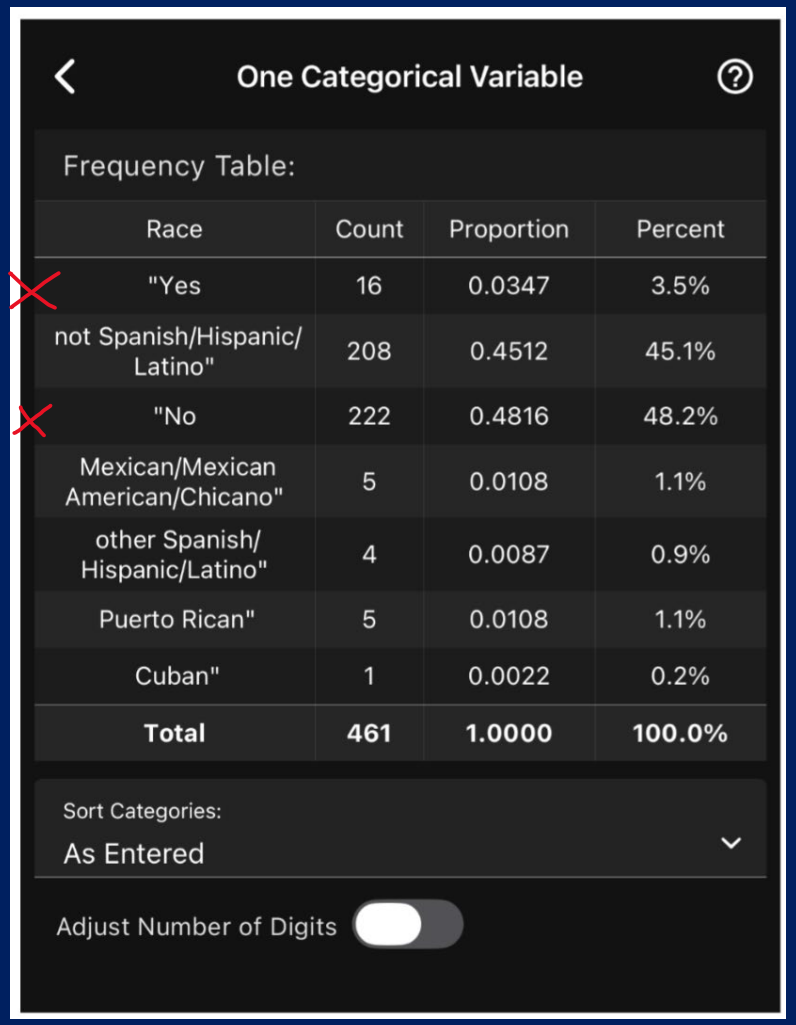

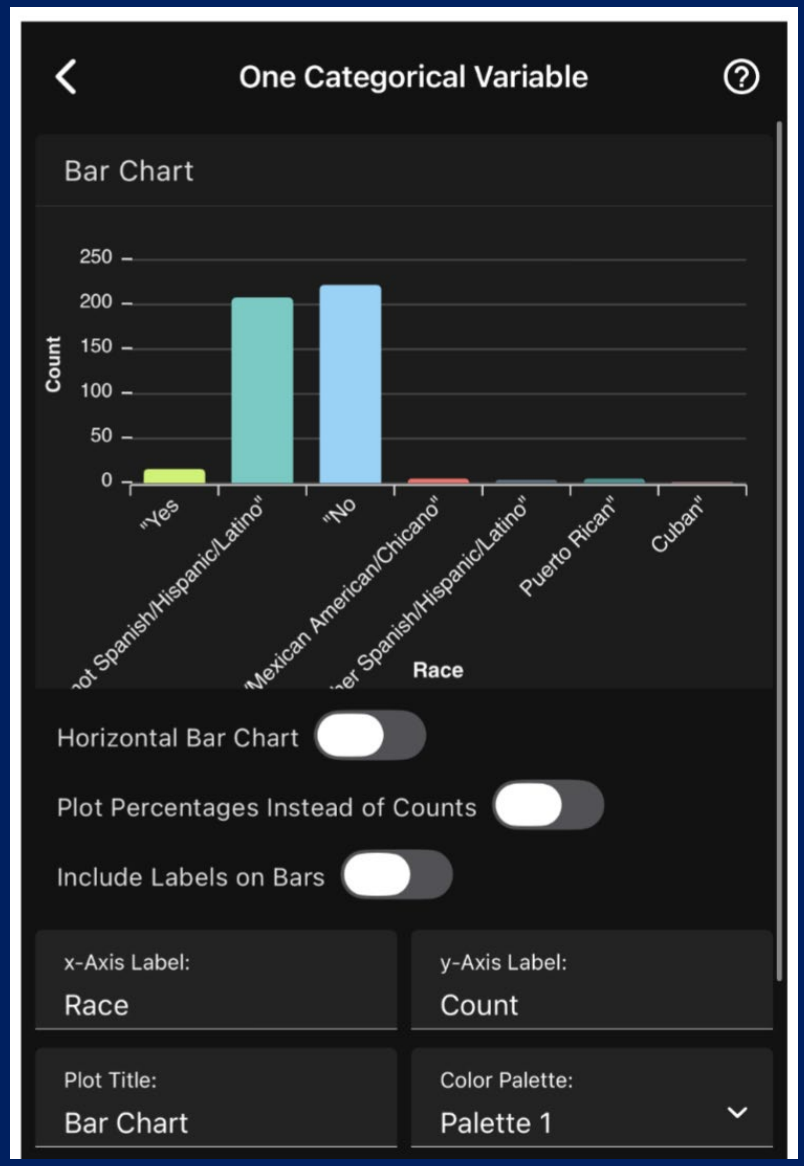

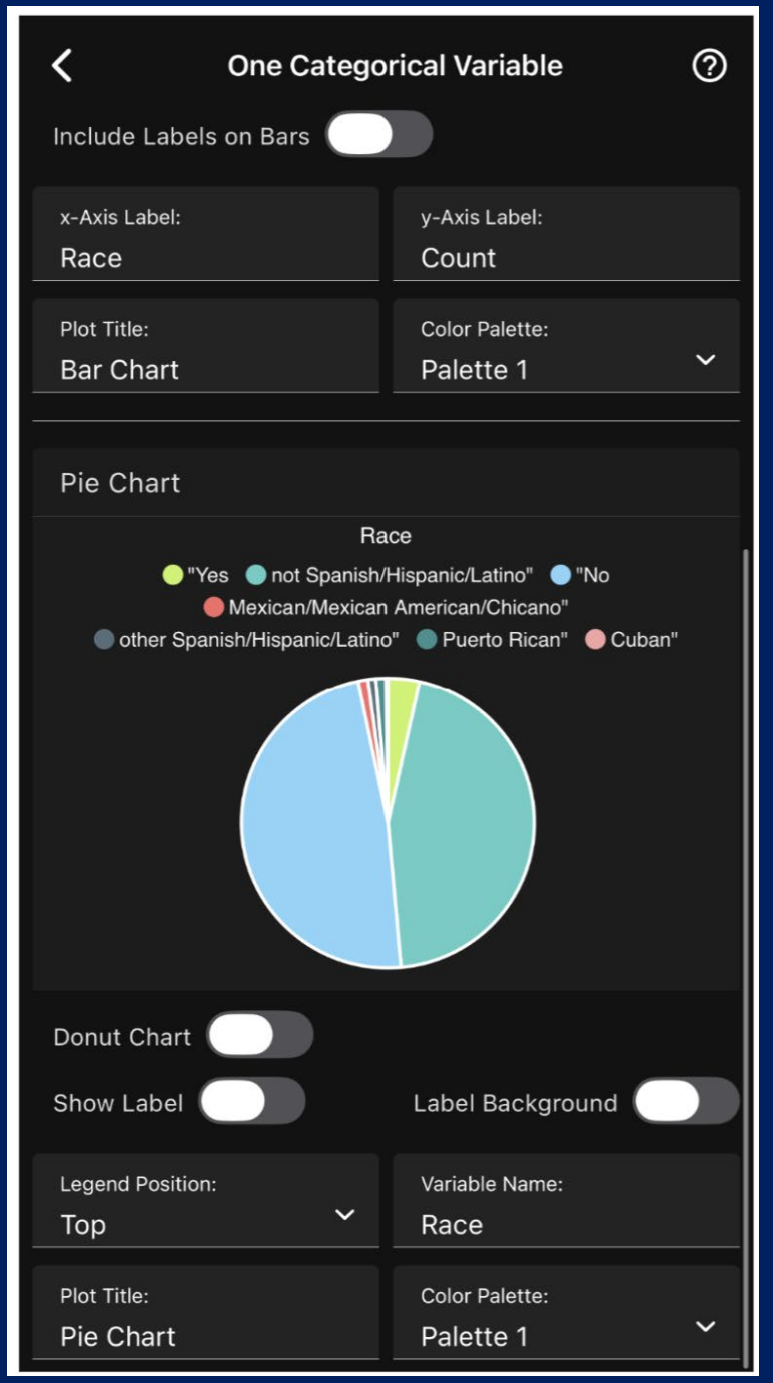

Another Student Solution:

Frequency Table:

Variable	Count	Proportion	Percent
White	296	0.7363	73.6%
Black	51	0.1269	12.7%
Asian	28	0.0697	7.0%
American Indian	6	0.0149	1.5%
Native Hawaiian	3	0.0075	0.7%
Other	18	0.0448	4.5%
Total	**402**	**1.0000**	**100.0%**

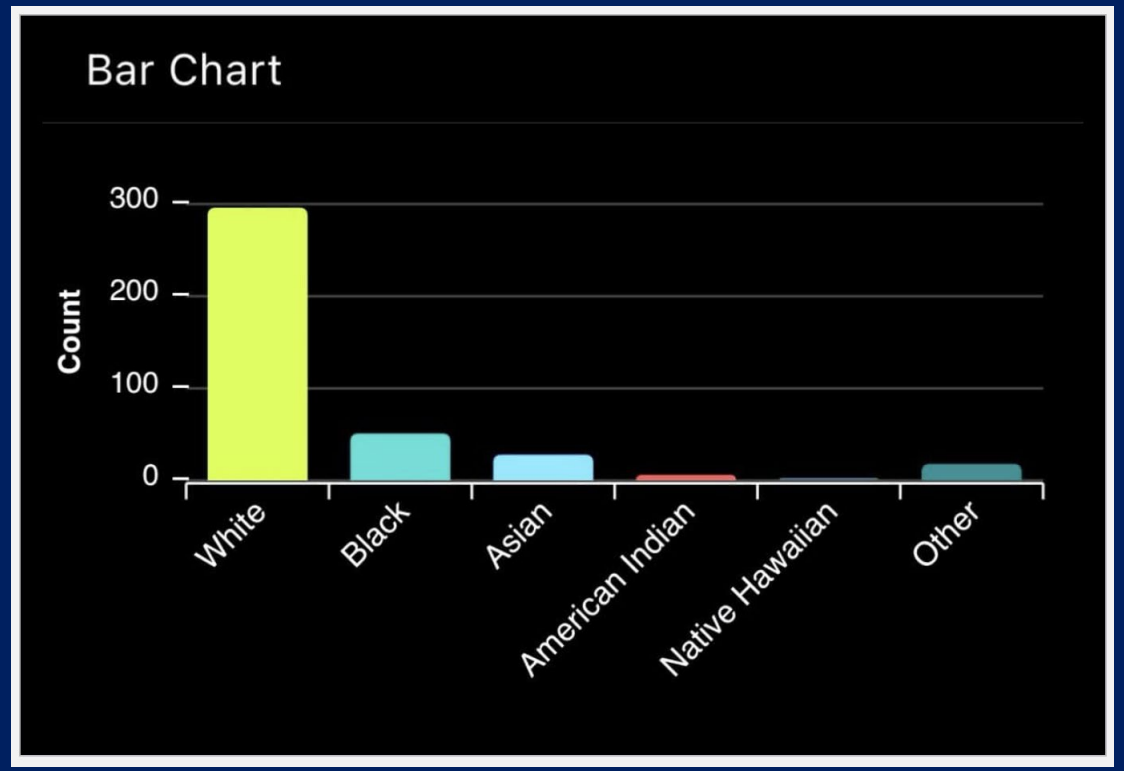

Thank you!

Please download, try and give me feedback on the Art of Stat app!

Slides available at ArtofStat.com