CS-BIGS 1(2): 130-138 http://www.bentley.edu/cdbigs/vol1-2/carver.pdf
© 2007 CS-BIGS

Tales of Huffman: An Exercise in Dealing with Messy Data

Robert H. Carver
Stonehill College, Easton MA 02357, USA

Statistics education reformers have for years called for the use of real data in teaching introductory statistics (Ballman, 1997; Garfield et al., 2004; Hogg, 1991). Instructors now have ready access to cases, textbook problems and other exercises with accompanying well-documented sets of real or realistic data. On-line portals and data libraries provide a huge array of real data sets keyed variously to substantive topics and statistical techniques suitable for introductory students.

The vast majority of these real datasets tend to have already been cleaned up by their preparers. As enriching as these resources are, relatively few of them offer students first-hand experience with the essential messiness of "real" real data. There is a good case to be made that data cleaning and preparation belong in introductory courses (Burger & Leopold, 2001). Certainly, problems of missing, dirty, and incomplete data are important topics within the field (Hoyle, 1971; Rubin, 1976; Wagner, 2002).

Using field data from the Wright Brothers' 1904 experiments, this case leads introductory or intermediate students through a process of data preparation, illustrating five common steps in data preparation and cleaning: standardizing the format of data records, deciding how to treat ambiguously recorded data, conversion of measurements to a single standard unit, detecting and resolving issues with outliers, and imputation of missing data.

Keywords: *Statistical Reasoning, Missing Data, Messy Data, Wright Brothers*

1. Background

The Wright Brothers' first successful powered flight in December 1903 was a singular moment in human history. Few technological developments are more famous or have had a more profound effect on the modern world than the birth of air travel. However, it is less widely known that the Wrights devoted several years of research and development prior to that first twelve-second flight, and that they continued to refine their design in the years following. It was another two full years before they truly had a commercially viable prototype. The setting for this case is their 1904 efforts to develop reliable methods to control sustained flight. It was one thing to build a plane that could lift off the ground and fly in a straight line. It was quite another to build one that could turn, adapt to shifting currents, and land safely. To understand the context of their 1904 experiments, it is necessary to begin the story a few years earlier.

Their successful 1903 flyer represented the consolidation and application of principles and ideas published by others over a period starting in the late eighteenth century with the work of Sir George Cayley in England (Hallion, 2003a). Throughout their experiments and research, the Wrights adopted a thoroughly systematic and scientific approach, carefully digesting previously published work and maintaining thorough lab and field notes on their own investigations. Their 1904 dataset has come down to us in its original form and is the basis for this case.

2. Designing the Wright Flyer

The Wrights had a lifelong fascination with flying machines and had been particularly interested in the possibility of human flight since the death of German engineer Otto Lilienthal in a glider accident in 1896. Lilienthal had experimented with sixteen different glider designs between 1891 and 1896, completing roughly 2000 glides. Lilienthal meticulously recorded his experiences with his designs and provided important mathematical groundwork upon which Wilbur Wright built (Crouch, 1989).

In 1899, Wilbur wrote to the Smithsonian Institution in Washington DC requesting a comprehensive bibliography on the subject of human powered flight. He wrote "I have been interested in the problem of mechanical and human flight ever since as a boy I constructed a number of bats[1] of various sizes after the style of Cayley's and Pénaud's machines" (Hallion, 2003b).

As a result of Wilbur's review of the earlier literature he came to understand that the problem of flight required the solution to three general problems, described as follows in a lecture he delivered to a group of engineers in Chicago in 1901. A practical flying "machine would require wings that would lift it into the air; a power plant to move it forward with sufficient speed so that the air flowing over the wings would generate that lift; and a means of controlling the machine in the air. It seemed obvious that the basic solutions to the first two problems had already been achieved…" (Crouch, 1989).

Unlike others working on the problem of flight, the Wright brothers "understood that since flight involved so many variables, such as the shape of the wing, the shape of the rudder, the strength of the materials etc. that they had to find a way to test each design component as much as possible in isolation without changing the others" ("Wright history", 2001). They methodically focused on each major variable in sequence, eventually re-integrating their results into a successful flying machine. Much like their contemporary Thomas Edison, they viewed experimental failures as instructive and invaluable.

2.1. 1901-1902: Wind, Lift, and wing design

Wilbur's review of the literature directed his choice of an optimal location to field-test his models and prototypes. He studied the weather data and settled on the North Carolina coast as a promising region that offered steady winds. In August 1900, after a busy summer at their bicycle shop in Dayton Ohio, he wrote to the communities of Myrtle Beach and Kitty Hawk, inquiring about the local conditions and the prospects of finding a suitable location for his investigations. He never heard back from Myrtle Beach (Tobin, 2003).

The Wrights experimented with gliders during the summer of 1901, focusing on the shape of the wings and on accurately calculating lift, relying on data and equations published by Lilienthal. Their experiments led them to a more accurate estimate of an important constant in the formula for lift, the coefficient of air pressure, originally estimated by Sir John Smeaton more than one hundred years earlier (Adams, 2003). With the primitive instrumentation available, it was difficult to make accurate measurements of pressure, but the design of a practical wing was highly sensitive to Smeaton's constant (Crouch, 1989; Jakab & Young, 2000). Ultimately, with a revised constant they had a more reliable estimate of lift and had their wing design and a calculation of the power requirements for an engine.

2.2. 1902-1903: Power and Flight

Working with a local machinist in Dayton, by December 1902 they had built a small 179-pound, 8-horsepower engine. With additional research and design, by spring they had a functional engine capable of generating up to 12 horsepower. This exceeded their initial estimates and would, assuming the correctness of their calculations, provide enough power to lift their 1903 flyer design (Hobbs, 1971).

In the summer and fall of 1903, their testing continued at Kitty Hawk NC, culminating in their famous success on December 17, 1903. The work leading to their first flight was marked by their characteristic attention to detail, data gathering, persistence, and caution. In sharp contrast to our current era of high-visibility product announcements, the Wrights studiously shunned publicity for their achievement, instead devoting the next two years to the refinement of their airplane, concentrating their efforts on "improving longitudinal stability and cir-

[1] By "bats" he was referring to the nocturnal flying mammal rather than baseball or cricket equipment. Today kids make paper airplanes but of course when Wilbur was a child there were no airplanes!

cular flight" (Hallion, 2003a). Simultaneously, they decided to elevate their flying-machine activity from a leisure pursuit to a business proposition, and secured a flight field closer to home, at Huffman Prairie near Dayton.

2.3. 1904-1905: Huffman Prarie and Controlled Flight

Huffman Prairie afforded them the ability to work near their home base, but the wind conditions were far less amenable than at Kitty Hawk. Their 1904 efforts at Huffman brought a mixture of success and failure for the Wrights, but they persisted in their work on controlling their aircraft and improving the plane's performance in windless and other challenging conditions. They developed new modes for takeoff, including a catapult that would sling the plane into flight "much in the same way a modern jet fighter is propelled from the deck of an aircraft carrier" (Adams, 2003, "Huffman prairie—1904. Born of dreams—inspired by freedom", 2003).

The 1904 engine was larger than the one on the 1903 vehicle, weighing 240 pounds and developing nearly 16 horsepower (Jakab & Young, 2000). This would provide greater lift and enable the plane to take off in relatively little wind. This engine design "was their sole power source during all their flying 1904 and 1905…progressing from a short straightaway flight of 59 seconds to a flight controllable in all directions with the duration limited only by the fuel supply" (Hobbs, 1971).

As previously, they mounted recording devices on the aircraft, under control of the pilot, to record both the distance traveled and elapsed time of each flight. That year they completed 105 flights, recording their flight data for about 70% of them. Flights were erratic, varying in duration from a few seconds to over five minutes. Many ended with crash landings, often damaging the airplane in some way. They averted serious personal injury, but just barely in some instances (Crouch, 1989, "Huffman prairie—1904. Born of dreams—inspired by freedom", 2003).

Crouch (1989) writes: "By the end of the following year, [t]hey had done it. The 1905 Wright airplane was one of the most extraordinary machines in the history of technology. Capable of rising into the air, flying for an extended period under the complete control of the operator, and landing safely, it was the world's first practical airplane. Nine years of trial and error, discouragement and hope, risk to life and limb, and brilliant engineering effort had culminated in the air over this Ohio cow pasture".

3. The Dataset

The dataset for this case comes from the Huffman Prairie flight trials during the summer and fall of 1904 now available to us on-line from the U.S. Centennial of Flight Commission ("Flight log: Huffman prairie 1904", 2003). As noted above, the Wrights maintained careful records of their work. A cursory review of the log (see Table 1) reveals many problems typical in real datasets: missing data, inconsistent metrics, variation in the precision of measurements and so on. In addition to measurement data, we also find notes about damage to the planes, visitors on site, and other commentary.

Table 1. Excerpt from Huffman Prairie Flight Log, 1904

DATE	FLIGHT	PILOT	TIME	DISTANCE	ALTITUDE	REMARKS
26-May		OW		ca. 25 ft.	ca. 6-8 ft.	Several pine spars were broken in landing, requiring postponement of further flights until repairs were made. The start of this flight was witnessed by Bishop Milton Wright, John G. Feight, Mr. and Mrs. Frank B. Hale, Mrs. William Werthner, a press reporter, and several others. Rain and lack of
10-Jun		WW		ca. 60 ft.		Machine struck ground because of faulty steering. Machine was damaged, preventing further trials until repairs were made.
21-Jun						
	1			ca. 100 ft.		
	2			ca. 100 ft.		
	3	OW		225 ft.		
23-Jun						
	1			264 ft		
	2					Tail damaged in landing.
25-Jun						Machine struck ground while turning at full speed and was damaged, preventing further trials until repairs were made.
July						Two trials made.
2-Aug						No systematic flight records were kept for 1904 by the Wrights before this
	1	WW		160 ft.		date. The first flight on this date is recorded in Wilbur Wright's Diary E as
	2	WW		370 ft		End bow broken when machine landed near fence.
4-Aug						
	1		10-2/5 sec. (down track)	195 ft		

In all, 91 flights were documented with some data or remarks recorded for 88 of them. We can summarize the key characteristics of the raw dataset as shown in Table 2.

Table 2. Properties of Raw Dataset

Variable or property of dataset	N of cases
Total number of flights documented	91
Some flight data reported	88
Pilot identified	78
Time (flight duration)	68
Distance (ground covered)	80
Altitude	1
Both time and distance recorded	68
Distance measured in feet	26
Distance measured in meters	53
Other distance metric	2
Distance estimated (e.g. "ca. 25 ft.")	4

4. Cleaning the Data

One major goal of their 1904 work was to reduce variation in the behavior of the aircraft so that they could successfully control the plane for extended flights. To enable them to achieve that goal, the Wrights needed metrics that would reliably monitor their performance. They measured what they could, and we can look at those measurements to track their progression towards their goal. Before doing any analysis, we'll need to address some of the defects in the data values themselves. Most of the data analysis in this article will be performed using Minitab, but we'll clean and edit the data in Excel. Other software tools can work as well.

4.1. Consistency in Record Format

There are different data types; many observations are missing; rarely is there altitude data; as noted above, distances are measured in three different scales with varying degrees of precision. Small problems abound and further analysis demands decisions about how to deal with them. We can speculate about the reasons for the deficiencies in the data, but we have no definitive historical record to account for all of them.

In addition, the layout of the data in the worksheet presents problems. We know that statistical software 'expects' a dataset to have one column per variable and one row per unique observation. As we inspect the original flight log, we find that there are rows with no data entries except for the date of the flight. In a large dataset, it would be necessary or advisable to automate as many of the data cleaning steps as possible. In this relatively small dataset, we can use a combination of manual methods and Excel functions to convert the data to useful format. As a first step, it will be wise to re-save the raw data into a new workbook (or worksheet), preserving the original file as a backup. Additionally, unless we intend to convert the remarks (Column G) into codes for later analysis, we can remove that column from the dataset.

We should next eliminate blank rows and be sure that each row follows a consistent format. For many dates, the Wrights recorded multiple flights per day. In the raw flight log file, there is a header row just containing the date for such days, and each flight is numbered in the subsequent rows. We can copy the dates downward in Excel, being sure not to increment the dates but only copy them.

After inserting a date into every row of the worksheet, we then sort the worksheet by Column B. The empty rows will "drop" to the bottom of the worksheet. We delete those rows, leaving a title row plus 88 data-filled rows, and resort by date and flight number to restore the original sequence of the data. Now we have a consistent rectangular array of flight data.

4.2. Standardizing Units of Measurement: Flight Distance

Having standardized the layout of the data rows, we can turn our attention to the inconsistent data values starting with the distance variable. We can identify several distinct problems to resolve. As recorded in the flight log and transcribed into Excel, this comes to us as a text (alpha-numeric) variable, because most of the values include unit labels like "ft" or "m." So, we'll need to separate the numerical and text portions of the 80 recorded values. In addition, we need to standardize the measurements either as feet or meters. Because more of the values are initially presented in meters, let's choose that unit of measurement and convert feet to meters (1 foot = 0.3048 meters).

There are other problems: some distances are listed as "ca. 25 ft", standing for circa (around, approximately) 25 feet. That is not as precise as "264 feet," but we need to decide how to list these approximate values. One flight's distance was "150 to 200" feet. More challenging still, we have one flight whose length was described as "almost 4 rounds of the field."

For the sake of this exercise we will treat approximate measurements as precise and interpolate to the center of ranges, so that 150 to 200 feet becomes 175 feet. We rely on a secondary reference to convert one circuit of Huffman Prairie to 1300 meters ("Huffman prairie—1904. Born of dreams—inspired by freedom", 2003). Four rounds of the field would be 5,200 meters; as an initial equivalent of "almost four rounds" let us use 5,000 meters.

4.3. Looking for Data Errors

Having rendered all distances in meters, we should next explore the data with a simple time series plot of flight distances (Figure 1), and with basic descriptive statistics (Table 3). Gaps in the lines connecting points represent missing observations. I have added a reference line in the graph indicating when the Wrights abandoned the launching rail and switched to their catapult system. At this stage in the data cleaning we want to see whether the converted values are credible or whether there are any signs of data errors. When cleaning data it is good practice to examine outlying values that might indicate erroneously recorded data.

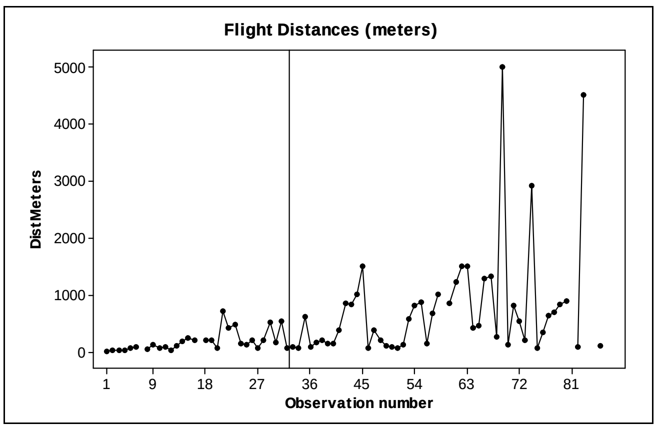

Figure 1. Flight Distances, 1904

Table 3. Summary Statistics for Distance in meters

Variable	N	N*	Mean	SE mean	SD
DistMeters	80	8	537.4	93.8	838.8
	Min	Q1	Median	Q3	Max
DistMeters	7.6	86.0	205.0	701.3	5000.0

We find that that we are still missing eight distance observations, and that the distribution of flight distances is strongly skewed to the right. The minimum flight distance was just 7.6 meters (about 25 feet) which, though short, is credible because similar values occur frequently. We find three exceptionally long flights, but they are borne out by reference to the flight log.

In terms of general patterns, there are several noteworthy features in the time-series plot. The first eighteen flights were consistently short in length, the next fourteen or so were more variable and slightly longer, the catapult-launched flights tended to be longer and more variable still, and the final dozen or so flights of the season varied very widely. The Remarks column of the original dataset can shed some light on assignable sources of variation in some instances.

4.4. Dealing with Implausible and Missing Data: Flight Durations

With the distances expressed in a common unit of measurement, we can move on to the issue of missing data. Of the 91 logged flights, we have recorded distances for 80 but only have flight durations recorded 68 times. Times are recorded in two different units of measurement (minutes and seconds) and with differing degrees of precision. Initially, we'll convert all time measurements to seconds. The results of that conversion are shown in Figure 2 and Table 4. We find patterns similar to the converted data for flight distances, though for this variable we are missing 20 of 88 measurements, or nearly 23% of the possible observations. Shortly we'll consider a possible remedy for this large number of missing values. Notice that the flight log contained no entries for duration for the first group of flights, and that there are many more gaps in this time-series plot than in the corresponding plot for distances (Figure 1).

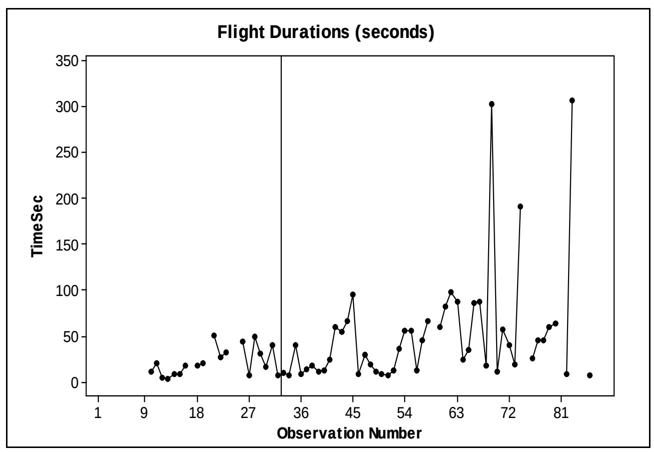

Figure 2. Flight Durations in Sequence

Table 4. Summary Statistics for Flight Durations

Variable	N	N*	Mean	SE mean	SD
TimeSec	68	20	43.43	6.78	55.92
	Min	Q1	Median	Q3	Max
TimeSec	2.75	11.10	25.60	55.75	308.00

We are aware that there are extreme values for both the distance and duration, and have concluded that these values are at least consistent with the Wright Brothers' journals. But it is still possible that the original data was written down or posted on the internet incorrectly. It may therefore be instructive to ask whether the distance

Table 5. Outliers in Scatterplot

Obs	Date/Flight	TimeSec	DistMeters	Standardized Residual	Comments
26	22 Aug. #2	44	193.5	–3.99	160’ track used; 1st recorded flight after damage to flyer a week earlier.
28	22 Aug. #4	49	192.0	–4.65	Switched to 195 ft. launching rail.
69	9 Nov #2	304	5000.0	3.24	5000 was our approximation of “almost four rounds of the field”
83	1 Dec #3	308	4515.0	–2.90	Flight log lists duration as “5 min 8 sec. (?)”. Log expresses doubt about the duration.

and duration data are internally consistent—that is, to ask if the times are plausible in the context of the distances traversed and vice versa.

One way to pursue this question is to construct a scatterplot of times versus distances. One such graph is shown in Figure 3. This scatter plot also displays a fitted regression line, indicating a strong relationship between the two variables. The estimated regression equation is:

$$\text{DistMeters} = -60.52 + 15.6921\ \text{TimeSec} \qquad (1)$$

In this graph we also see four values that deviate from the linear pattern. We can identify the four particular values either by inspection of the graph and dataset or by running a regression and reviewing the observations with unusually large standardized residuals. Table 5 lists the four values, identifying the dates of the flights and noting comments based on the original field notes and our own data cleaning thus far. The two outlying points at 44 and 49 seconds lie well below the regression line, but there are no assignable causes that are obvious from the field notes or from our data conversion.

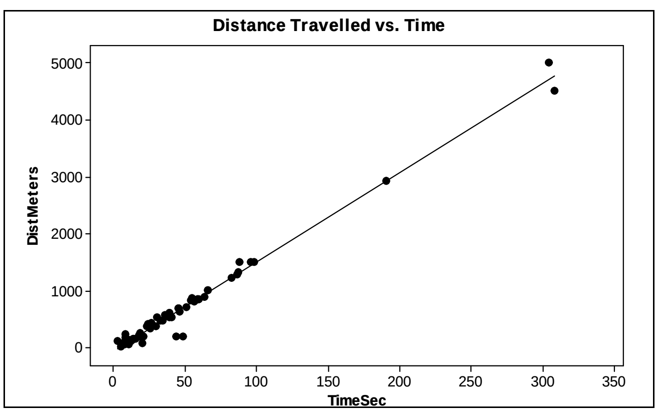

Figure 3. Distance vs. Time

On the other hand, the two long-distance outliers may merit further adjustments because we can find reason to suspect both values. The November 9 flight was the one we estimated at 5,000 meters. According to the regression equation (i.e. the fitted line), given a duration of 304 seconds, the distance was more probably closer to about 4,710 meters. The December 1 flight log indicates uncertainty about the flight duration. Again, using the regression equation we would estimate that the flight duration was 291.58 seconds or about 4 minutes 52 seconds.

The underlying relationship between distance and duration provides another approach to checking the internal consistency of the data. We can compute an average speed for each flight (distance divided by duration; see Figure 4) and inspect for outliers. We might wonder about the consistency of airspeed during the Huffman flights, and the average speeds can give us a glimpse into the development of that consistency. Doing so reveals two unusual observations: # 13 and #15, both second flights piloted by Orville on August 5 and 6. Notably, observation #13 had an estimated speed of 39.46 meters per second, or 88 miles per hour. Inasmuch as the mean speed was 13.2 meters per second, it is hard to believe that Orville achieved such high speeds in flights that lasted less than 10 seconds each according to the log. The 88 mph flight was recorded as having been 108.5 meters in length and lasting 2.75 seconds. Compared to the other flights of similar lengths or durations, this is a very implausible event. One might doubt the accuracy of the raw data, or conclude that Orville spent three thrilling seconds on August 5, 1904.

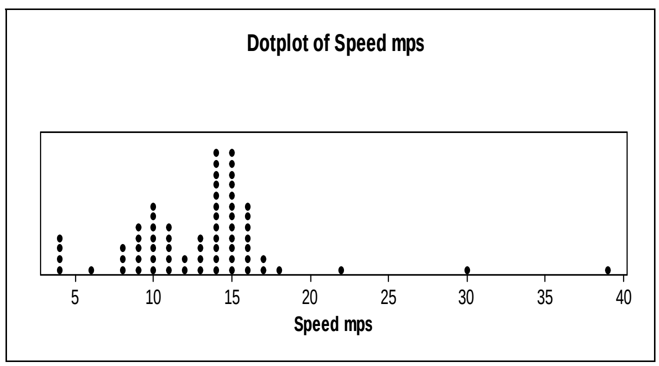

Figure 4. Distribution of Average Flight Speeds

If we speculate that it was easier to record distance more accurately than time, we might suspect that promptness with starting or stopping the timer might have been at play. This provides a rationale for either adjusting the two flight duration values or dropping the two observations as probably inaccurate.

Before deciding how to treat these outliers we also note that the distribution of average speeds (Figure 4) is bimodal. Simple graphical analysis can rule out pilot differences and takeoff method as underlying causes[2], but further investigation indicates that short-distance flights tended to have lower average speeds than the longer-distance flights(see Figure 5). The boxplots also clearly display outliers, and would be particularly useful with a large dataset.

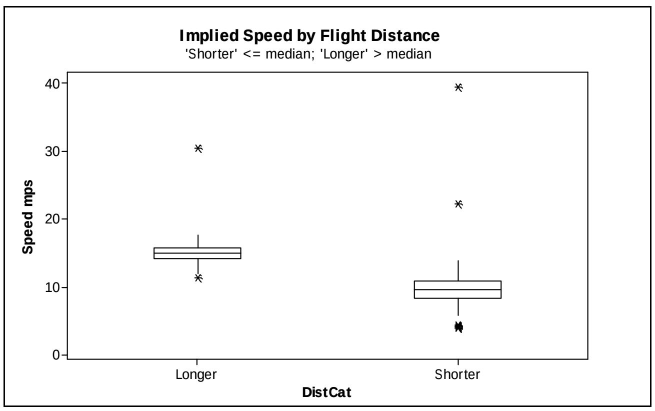

Figure 5. Boxplots of Implied Speeds

Figure 6 allows us to investigate the relationship between speed and total flight distance more precisely. We see that the very short flights also varied considerably in average speed. Flights longer than about 500 meters consistently averaged about 13 meters per second. Shorter flights varied far more, and there seems to be a strong curvilinear pattern to the data.

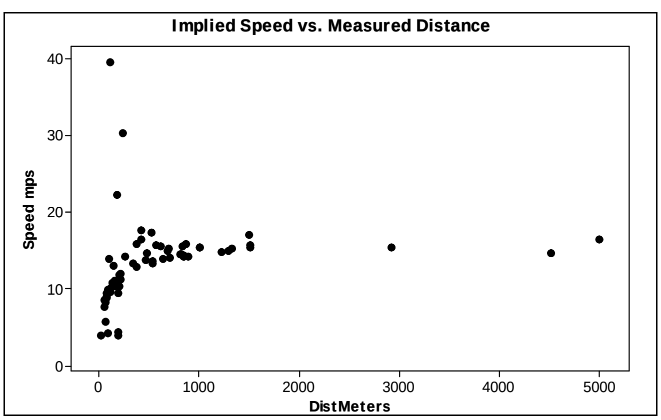

Figure 6. Computed Speed vs. Flight Length

We should also note that the three outliers in the upper left of this graph occurred in three consecutive flights (observations 13 through 15) on August 5 and 6. If we temporarily drop those observations from the dataset and produce the scatter plot once again, we see the result shown in Figure 7; the graph also includes a log-linear fitted line to describe the overall pattern (adj. R^2=.65, a moderately good fit).

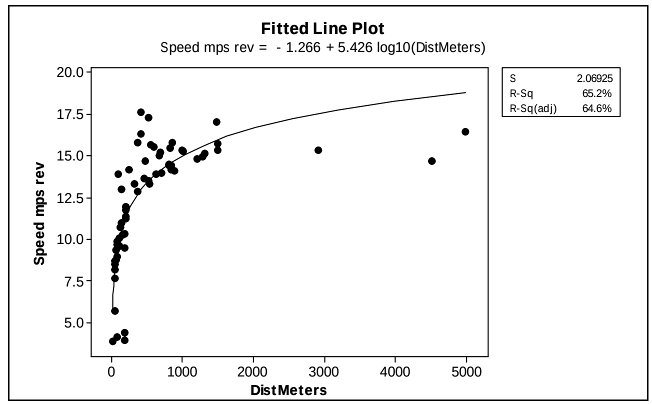

Figure 7. Average Speed vs. Distance, deleting outliers

4.5. How to Treat Extraordinary and Missing Values

If we have sound theoretical and empirical reasons to question the accuracy of the short-flight measurements and we can find a discernable pattern in the data, we might be justified in (a) deleting the outlying values, (b) using a regression model to adjust the suspect duration values, and/or (c) using a regression model to impute the missing duration values. The choice among these alternatives generally depends upon the purpose of the analysis and the needs of the decision-maker or researcher. For the purposes of this case, let's assume that we prefer to adjust the three outlying values rather than delete them, and we wish to impute approximate durations to replace the missing data in the original flight log.

First we use the estimated model shown in Figure 7 to estimate (or fit) values for average speed for all flights with recorded distances. The actual distances are then divided by the fitted average speeds to compute estimated durations, which can replace the suspected outliers as well as the missing values.

After adjusting six questionable observations of duration and estimating plausible times for the other 12 missing values, we can summarize the impact on our dataset. In the case of distances, we adjusted a single value which was the maximum and the change has a small impact on the distribution of the data.

The time variable has changed more substantially as a result of this work. We've added 12 observations and ad

[2] Readers are encouraged to verify this for themselves.

Table 6. Summary Statistics for Cleaned Data

Variable	N	N*	Mean	StDev	Minimum	Q1	Median	Q3	Maximum
DistMeters	80	8	537.4	838.8	7.6	86.0	205.0	701.3	5000.0
AdjDist	80	8	533.8	819.7	7.6	86.0	205.0	701.3	4710.0
TimeSec	68	20	43.43	55.92	2.75	11.10	25.60	55.75	308.00
RevisedTime	80	8	37.45	51.99	2.17	9.35	18.30	49.20	304.00

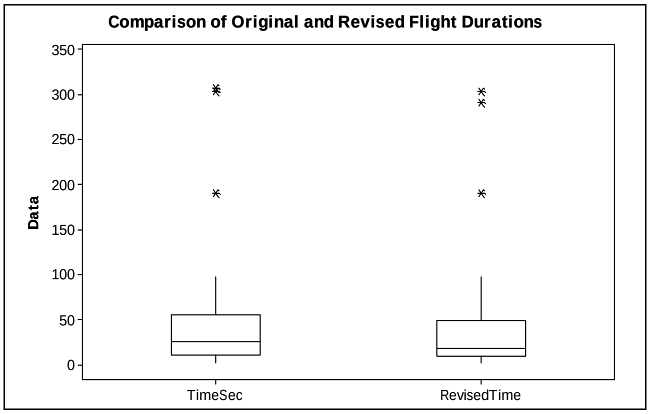

Figure 8. Original and Revised Duration Data

justed six others in all reducing the mean by six seconds and generally reducing variation in the dataset. Figure 8 displays the comparison. But with those adjustments we've addressed the problem of missing data in a way that is consistent with the original data and provides for an 18% increase in useful sample size for this particular variable.

5. Conclusions and Directions for Further Research

This case illustrates five common steps in data preparation and cleaning: standardizing the format of data records, deciding how to treat ambiguously recorded data, conversion of measurements to a single standard unit, detecting and resolving issues with outliers, and imputation of missing data. In so doing it provides an accessible introduction to an important area in the practice of data analysis against the backdrop of a familiar historical project.

The Wright Brother's nine-year odyssey also provides opportunities to comprehend a number of statistical concepts, techniques, and standards of practice. The Wrights engaged in description, estimation, and prediction. They developed hypotheses from prior research by others, tested those hypotheses experimentally, and in turn developed new hypotheses (Box et al., 1978; Box & Liu, 1999). Their approach and methods reflected good hard thinking and logic, as well as theory, and their empirical testing used both physical scale models and full-sized prototypes. They continuously devised appropriate measures and measuring devices, they documented their processes, they recorded the measurements and investigated the accuracy of those measurements, and they constantly worked to control variation. As such the narrative of the case can provide useful examples of statistical thinking and practice.

The dataset presented here is a fraction of the data that the Wrights captured and that has been published. Their other research notes and flight logs are surely fertile ground for additional case development, and classroom research on the effectiveness of this case would be most welcome.

REFERENCES

Adams, N. (2003). *The flyers: In search of the Wright brothers*. New York: Crown.

Ballman, K. (1997). Greater emphasis on variation in an introductory statistics course. *Journal of Statistics Education, 5*(2). http://www.amstat.org/publications/jse/v5n2/ballman.html

Box, G. E., Hunter, W. G., & Hunter, J. S. (1978). *Statistics for experimenters*. New York: Wiley.

Box, G. E., & Liu, P. Y. T. (1999). Statistics as a catalyst to learning by scientific method, part i--an example. *Journal of Quality Technology, 3*(1), 1-15.

Burger, L. W., Jr., & *Leopold*, B. D. (2001). Integrating mathematics and statistics into undergraduate wildlife programs. *Wildlife Society Bulletin, 29*(4), 1024-1030.

Crouch, T. D. (1989). *The bishop's boys: A life of Wilbur and Orville Wright*. New York: *Norton*.

Flight log: Huffman prairie 1904. (2003). from http://www.centennialofflight.gov/chrono/log/1904HuffmanPrairie.htm

Garfield, J., Aliaga, M., Cobb, G., Cuff, C., Gould, R., Lock, R., et al. (2004). *The GAISE project: Statistics education guidelines for college courses*: American Statistical Association.

Hallion, R. P. (2003a). *Taking flight: Inventing the aerial age from antiquity through the first world war*. New York: Oxford University Press.

Hallion, R. P. (2003b). The Wright brothers: How they flew. *Invention and Technology Magazine, 19,* 18-40.

Hobbs, L. S. (*1971*). The Wright brothers' engines and their designs: Smithsonian Institution Press.

Hogg, R. V. (1991). Statistical education: Improvements are badly needed. *The American Statistician, 45*, 342-343.

Hoyle, M. H. (1971). Spoilt data--an introduction and bibliography. *Journal of the Royal Statistical Society. Series A (General), 134*(3), 429-439.

Huffman *prairie*—1904. Born of dreams—inspired by freedom. (2003). from http://www.centennialofflight.gov/essay/Wright_Bros/1904/WR7.htm

Jakab, P. L., & Young, R. (Eds.). (2000). *The published writings of Wilbur and Orville Wright*. Washington DC: Smithsonian Institution.

Rubin, D. B. (1976). Inference and missing data. *Biometrika, 63*(3), 581-592.

Tobin, J. (2003). *To conquer the air: The Wright brothers and the great race for flight*. New York: Free Press.

Wagner, H. M. (2002). And then there were none. *Operations Research, 50*(1), 217-226.

Wright history. (2001). from http://www.wrightredux.org/index.cfm?page=5

Correspondence: rcarver@stonehill.edu