Journal of Environmental Statistics

September 2019, Volume 9, Issue 1. http://www.jenvstat.org

Visualizing Multivariate Time Series of Aerial Fire Fighting Data

Devan G. Becker
Department of Statistical and Actuarial Sciences, University of Western Ontario

W. John Braun
Department of Computer Science, Mathematics, Physics and Statistics, University of British Columbia - Okanagan Campus

C.B. Dean
Department of Statistics and Actuarial Science, University of Waterloo

Douglas G. Woolford
Department of Statistical and Actuarial Sciences, University of Western Ontario

Abstract

Aerial wildland fire fighters have a unique challenge. They are able to fill their tank with water via a nearby body of water, drop this water on a fire, then return to repeat this process. For a given fire, the replications of the fills and drops are multivariate time series measured in space, and this data structure allows us to compare replications over space and time. We use control chart methodologies to determine which time series were unlike the others, then examine the data from both a univariate and a multivariate viewpoint to determine potential causes.

Keywords: visualization, replicated time series, statistical quality control, wildland fire.

1. Introduction

Piloting a fire bombing aircraft over an active wildland fire is an inherently difficult task that is made even more difficult by the long working hours in demanding conditions. Wildfires are fought in order to protect the values at risk (life, property, timber values, to name a few). The safety of the individuals performing these duties is a primary concern.

Many studies have been done on pilot fatigue and performance for commercial airliners. Fatigue has been shown to come about through scheduling (Jackson and Earl 2006; Goode 2003) or length of flight (Strauss 2010). The performance of pilots in stressful (military) environments was investigated by Miller, Matsangas, and Shattuck (2008), which suggested that sleep disruption, sleep restriction, long shift work, or shifts that disrupt the circadian rhythm are especially dangerous in stressful military environments.

The research above has been performed on commercial and military flights, which are inherently different from those faced by aerial wildfire pilots. For example, the most difficult part of a commercial flight is often the take off and landing and these only happen once per shift. As discussed below, aerial fire fighting is considerably more stressful.

There are two types of fire bombing aircraft: a skimmer that scoops water which is dropped on the fire, and an airtanker which drops fire retardant on the fire. The skimmer is able to refill its water tank repeatedly within a shift. The refilling process requires a descent to the surface of a body of water, skimming the water for a short period of time to fill the tank prior to ascending again. Skimmer aircraft, on 4 hours of fuel, can cycle through up to 40 takeoffs and landings, returning for fuel for another 4 hours of landings, takeoffs, and bombing. The airtanker is loaded once each shift with fire retardant at a designated airport (Airtanker Base). Although the airtankers that haul fire retardant make fewer drops than the skimmers the flying conditions are the same. Fire bombing requires both types of aircraft to work in mountainous terrain, windy conditions and often with poor visibility conditions together with numerous other aircraft in a small airspace.

Our objectives are to present an exploratory data analysis (sensu Tukey 1977; Brillinger and Finney 2014) with a focus on evaluating portions of the flights where there may be anomalous flying behaviour. We employ techniques from Statistical Quality Control (see Montgomery 2009) to identify possible outliers in the data, and then investigate these outliers using multivariate time series visualization techniques to assess a univariate time series with vaguely periodic structure. Thus, we effectively extend multivariate visualization techniques to accommodate replicated time series, and introduce techniques for detecting possible changes in ordered replicates of a time series. We have found it useful, as well, to incorporate spatial information into the analysis. All of these techniques are complementary and should be evaluated together.

Our data include 130 aerial fire fighting flights from the 2016 fire season in British Columbia and Alberta, Canada. Each flight contains variables on the plane's location, orientation, and speed recorded each second. The numbers of fill and drop events vary between flights, and there are epochs within many of the flights where the pilot flew in a holding pattern while awaiting specific instructions.

As a case study, we focus on a single flight during the major wildfire in Fort McMurray, Alberta (Nash, St Arnaud, Tithecott, Simpson, Stocks *et al.* 2017). This fire caused an estimated $3.7 billion in insured losses and burned a total of 589,552 hectares. There is an airport in Fort McMurray that many of the aerial fire fighters flew out of, but this also meant that the pilots' refueling downtime was spent in the community that was threatened by the fire. Despite the abundance of risk, there were no reported incidents with any of the pilots fighting this fire. Nonetheless, the Nash *et al.* (2017) report recommends improvements to the planning and management of airspace in order to further increase pilot safety.

The paper is structured as follows. In Section 2, we look at basic plots of the data to determine the obvious outliers and demonstrate the need for more sophisticated visualization. Section 3 briefly outlines the control charts used for this data and provides several visualization techniques for anomalous patterns in univariate and multivariate data as well as an overall learning or fatigue effect. Section 4 summarises the visualizations presented in this paper with a tool that flight controllers could use in real time. Section 5 concludes the paper with a summary of the visualizations presented and a discussion of their applicability to other data.

2. Description of the Data

The raw data consist of unstructured time series with observations recorded each second on a number of variables including 3-dimensional spatial coordinates. An example is displayed in Figure 1, where we consider Angle of Attack. In Panel (a), we see the entire trace, with the times of drops and fills identified. In Panel (b), we display the traces for each of the replicated segments of that same flight. These flight segments are of different duration, since the times between skimming events vary.

The variables recorded include: orientation (pitch, roll, and, yaw), location (latitude, longitude, and altitude), accelerometer readings, and the angle of attack. Pitch, roll, and yaw are all measured relative to level flying. Dipping the nose of the plane up and down changes pitch, dipping one wing down and the other wing upwards changes the roll, and rotating the plane while keeping it level changes the yaw. Angle of attack is defined as the angle between the air flow and the fuselage. The critical angle of attack is where the plane has maximum lift; an angle of attack of 0 means the plane is flying directly into the window (usually indicating level flight), and a higher angle of attack puts the plane at risk of stalling. The critical angle of attack is usually between 15 and 20 degrees. Because this can be measured directly and can be interpreted in terms of lift, regardless of speed, the primary focus of our analysis will be the angle of attack data.

The location of each fire has a different topography which forces the pilots to perform different manoeuvres for fills on different days. Events, such as dropping water on the fire are repeated several times over the course of the flight, but the pilots are often faced with varying wind conditions throughout the flight and may receive new instructions as the fire changes. There are often other planes working on the same fire so the pilots necessarily adjust their behaviour to maintain a safe distance from other planes. There is also a chance that the fire grows unexpectedly in a new direction, requiring the pilots to adopt entirely new strategies mid-shift.

The overall inhomogeneity of the time series poses inferential challenges, but by exploiting the quasi-periodic structure of the fills and drops, we can make some progress. For example, by assuming that flying behaviour at the time of skimming should always follow the same pattern, we can assess the these skimming epochs for anomalies. Through visualization, we can incorporate information from univariate control chart methodology, 3-dimensional spatial locations of the time series, and other variables measured throughout the flight into an insightful exploratory analysis.

Any anomalies that are detected should not be immediately interpreted as pilot error. The flight conditions are challenging and the recorded data contain measurement error. There are no reports of any incidents in any of our data. The ultimate purpose of these analyses is to highlight potential areas of excess risk and present them to the pilots and the flight managers who will then make the appropriate decisions.

The amount of water in a skimmer's tank is one of the recorded variables. If we let $T_{1/2}$ denote a time epoch when the volume of water in the tank is increasing through the halfway point, the interval $T_{1/2} - 15, T_{1/2} + 15$ is a useful observation time window. We can construct

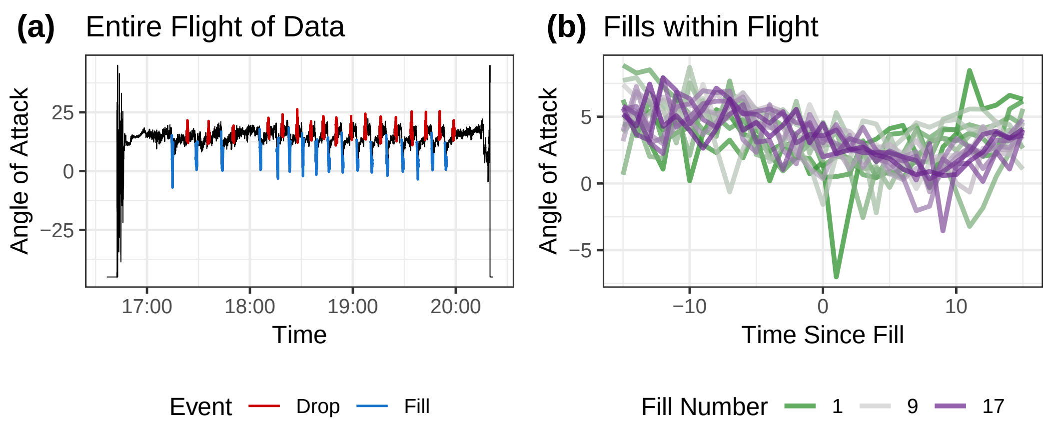

Figure 1: An example flight from the Fort McMurray fire. The raw data **(a)** is a time series with replicated events **(b)**.

an ordered sequence of replicated time series by extracting observations in such windows. Figure 1(b) shows such replicated time series for the Angle of Attack observations. We make the assumption that the pilot was attempting the same approach for every fill, regardless of what happened in other fills. In other words, we assume that these replicated time series are independent and identically distributed. The visualizations that we employ are intended to detect deviations from this assumption.

In addition to the structural inhomogeneity described above, most of the variables in the multivariate time series are correlated. Changing any of the orientation variables will cause a change in the accelerometer readings in accordance with the magnitude of the change. The pitch and roll change smoothly with each other during turns but not during straight sections of flight. The angle of attack is heavily correlated with the pitch and wind direction, but only the pitch is recorded.

3. Identifying and Visualizing Outlying Patterns in Time Series

3.1. Identifying and Visualizing Outliers: Univariate Case

Identifying Outliers

Our approach to the visualization of these time series data exploits control chart methodology where one wishes to determine whether a particular sample is different from the others (Shewhart 1925).

We employ two techniques to attempt to make the data independent and identically distributed so that standard control charts apply. The first technique involves finding the median value of the angle of attack across fills at each time point. This creates a single median fill (the trend) which is then subtracted from all of the other fills. The second technique finds the second difference of the value of the angle of attack for each fill. In both techniques, the mean and standard deviation of the resulting fills are monitored with Shewhart control

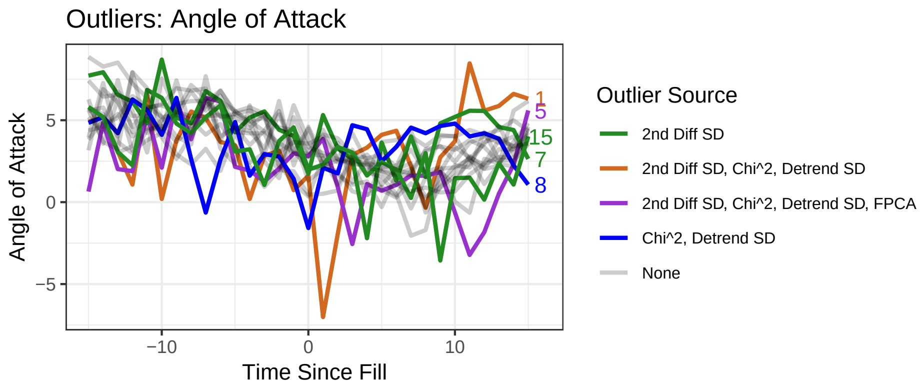

Figure 2: Outliers from the various control chart methods. The non-outliers are transparent so that the overall trend can still be seen. The numbers are the number of the fill within the flight.

charts. We refer to these two methods as detrended and 2nd difference control charts, respectively. We also consider a control chart from Hyndman and Shang (2010) based on spline estimates of each fill. The coefficients of the spline bases are found and then decomposed into principal components. Outliers are labelled based on whether the first two principal components lay outside of a bivariate boxplot (a.k.a. bagplot). We refer to this method as FPCA (Functional Principal Components Analysis). Finally, we use a control chart that monitors a non-parametric estimate of the variance for unusual values based on the χ^2 distribution (Chi^2 control charts, Zhang and Albin 2009).

The different charts have different purposes. The detrended control charts will theoretically detect fills that were far from the median or have high standard deviation relative to the median. The second difference control chart should detect outliers with a strong trend (or a weak trend, since the control limits are symmetric). The FPCA chart should detect fills with a different trend compared to the others, and the Chi^2 chart should detect fills that vary from the median fill in a different way compared to the others.

The control charts for these data are shown in Appendix A. In the sequel, we will be primarily concerned with the binary decision of whether or not a particular fill pattern is an outlier. We will use the identification of an outlier as an indication that the particular fill requires further investigation.

Visualizing Identified Outliers

Figure 2 shows the fills again, now colour-coded according to which control chart methodology identified a given fill as an outlier. With the exception of the two green lines, the particular combination of control chart labels are sufficient to uniquely identify a fill (i.e. there is only one orange, one purple, and one blue fill).

The orange fill was labelled as an outlier by both s charts as well as the Chi^2 chart, which is expected as it has "spikes" that induce variance effects that are different from the other fills. The FPCA control charts did not signal this fill as an outlier since it follows essentially the same trend as the other fills and the large spikes were removed in the spline smoothing step of the algorithm. The large negative spike near $t = 2$ indicates that the aircraft had less lift while taking off after skimming. As this was the first fill in the flight, we might be seeing evidence of a learning effect as the pilot is becoming accustomed to the flying conditions.

The purple fill is the only fill identified as an outlier by the FPCA chart, likely because all of the deviations from the trend are in the negative direction. A possible interpretation is that the nose of the plane was tilted downward excessively, or the airspeed was not high enough to generate lift. Both interpretations are suggestive of difficulty.

At the beginning of the blue fill, the angle of attack was unusually low, while after this fill, the angle of attack was higher than expected. This is likely why the Chi^2 chart signalled it as an outlier, since this form of deviation is different from the deviation of other fills. The detrended s charts also detected this outlier since having observations that are consistently far from the median fill will result in high variance. The pilot may have recognized that he descended to the water too slowly and attempted to correct this on the way back up.

3.2. Identifying and Visualizing Outliers: Multivariate Case

Figure 3 shows examples of multivariate plots for time series with replicates. The column of plots in Figure 3 (a) is an extension of Figure 2 (b) where all of the variables of interest are plotted with a consistent time axis for ease of comparison. The fills are colour-coded as in Figure 2 in order to determine whether the outlying angle of attack observations coincide with anomalies in other variables.

The orange fill was discussed in the previous section, but we now see that all of the other variables appear to have anomalous behaviour during this fill. The large negative spike at $t = 2$ in the angle of attack is also visible as a large positive spike in pitch. For some reason, this fill exhibits a higher degree of variability than the other fills.

The purple fill in Figure 2 was noteworthy for having spikes that were uniformly smaller than in the other fills. Anomalous accelerometer readings indicate that the pilot was taking corrective action.

The plots on the right in Figure 3 provide two perspectives on the same bivariate replicated time series plot. The colour codes indicate whether the fill was an outlier in neither variable, only in angle of attack, only in pitch, or in both. Note that the control charts are not multivariate. This sort of colour scheme could also be used in Figure 3 (a), but we wanted to keep the colour scheme consistent with previous plots and to highlight the orange fill's persistence as an anomaly.

The orange fill appears again in Figure 3 (b) and (c) (now coloured black) as a spike towards the edge of the cube. This indicates that it was a spike in both angle of attack and the pitch of the plane. Contrast this to the fill coloured in red, which had a spike in angle of attack but not in pitch. Since the angle of attack measures the difference between the angle of the wind and the angle of the plane, a spike in angle of attack that is not present in pitch may simply indicate a gust of wind.

Spatial Variation as an Attributable Cause

The spatial location is a special case of multivariate time series. We expect the angle of attack to be different under different wind conditions, which may occur when a plane skims two different bodies of water. We similarly expect slightly different behaviour if one skimming location has a longer approach than another and both are visited in a single flight. This spatial

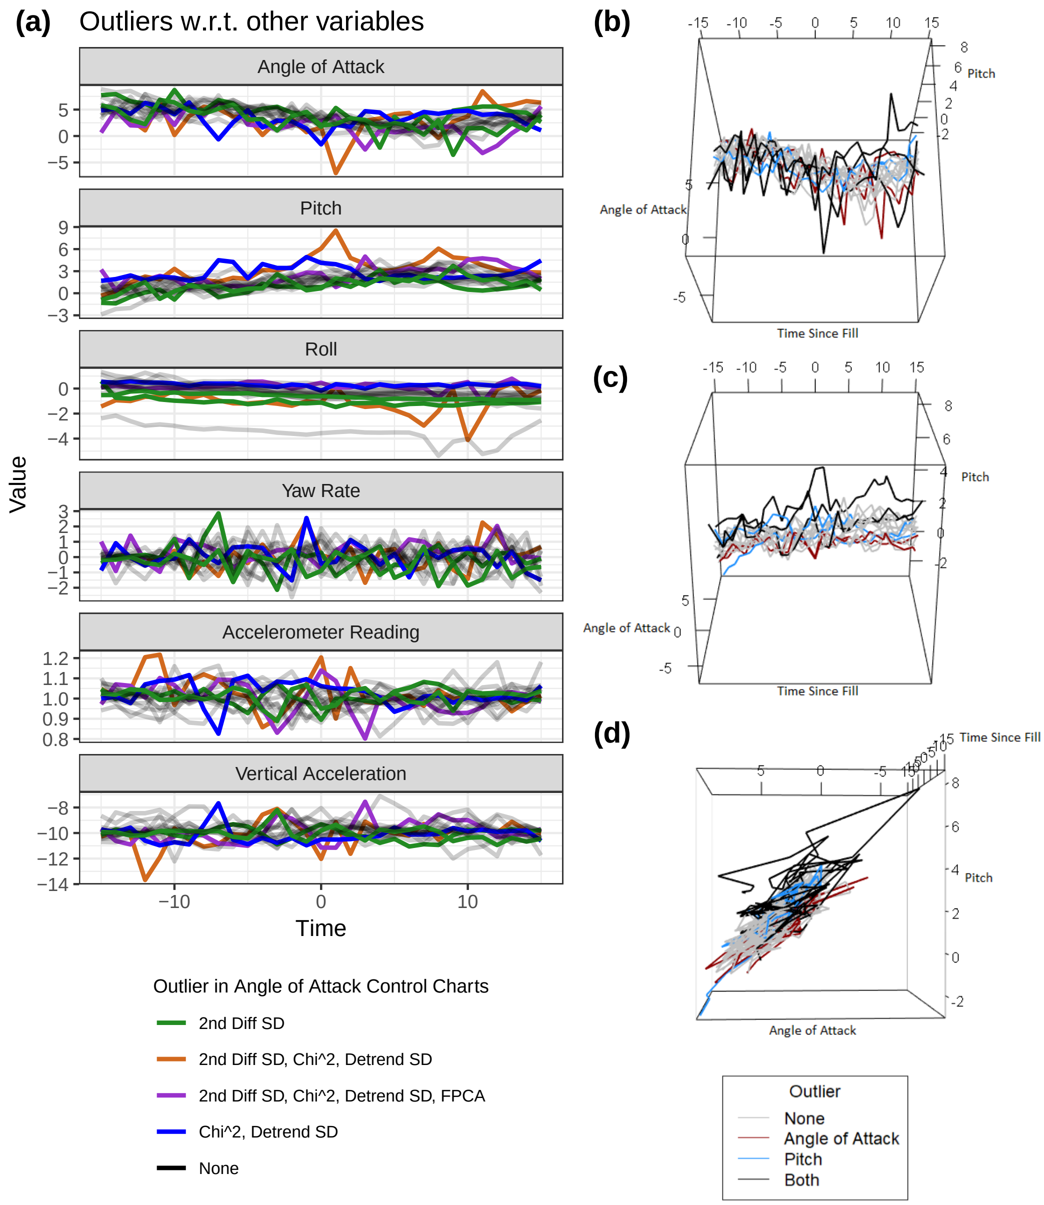

Figure 3: Multivariate plots for evaluating outliers. **(a)** shows the other variables, coloured based on whether they were an outlier in the control charts for angle of attack. The orange fill and the blue fill are noteworthy in all of the other variables, indicating that something went wrong in this fill. Outlier fills that don't show up in other variables may be simply due to a gust of wind. The plots in **(b)**, **(c)**, and **(d)** are different views of a bivariate time series plot of angle of attack and fill, coloured by whether the fill was an outlier in neither, both, or one of the two. These two variables are only sometimes correlated (angle of attack is a function of pitch, roll, yaw, and wind direction).

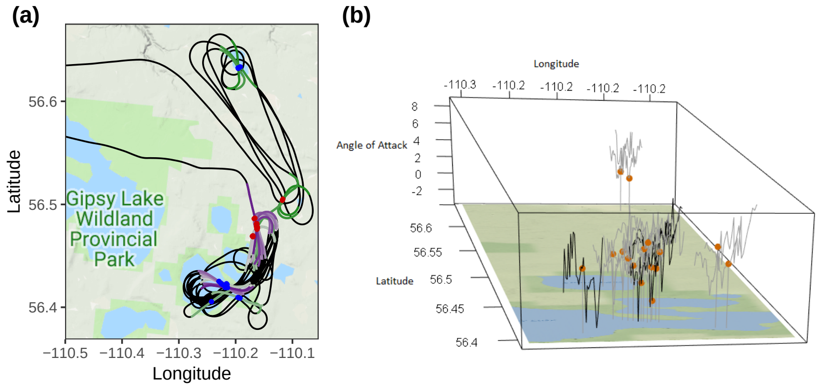

Figure 4: **(a)** Map of the fills (blue dots) and drops (red dots). The colours are the same as those in Figure 1 (b), where green represents earlier fills. The first 3 fills happened at a different lake. Map data ©2018 Google. **(b)** A three dimensional plot where the height of the line is the angle of attack rather than the altitude (positive angle of attack means that the plane is generating lift). Orange dots are placed at $t = 0$ for each fill with grey lines drawn to indicate location on the map. Black lines represent fills that were outliers in at least one control chart. The fill with a large spike near the back of the map (orange in previous sections) was one of the first fills and happened in a different lake from the others. Map data ©OpenStreetMap contributors.

variation is especially pronounced in drop events, where the updrafts from the fire and the height from which the plane drops the water will likely affect the other variables. However, it is not reasonable to create a control chart that labels fills as outliers because they were performed at a different latitude. None of our control charts take spatial variation into account directly. Instead, we can visualize spatial representations of our data to determine whether the variation is due to location.

In Figure 4 we can see that the plane visited two different lakes during this flight. The colours in 4 (a) indicate that it was the earlier fills that occurred in a separate lake. The fill with the large spike (coloured orange in the previous sections) is shown in Figure 4 (b). This large spike occurred in the very first fill in a different lake than most of the other fills, which might explain why it appeared as an outlier.

Exploring Learning and Fatigue as a Possible Attributable Cause

A shortcoming of Figure 2 is that it does not take the order of the fills into account. In long flights with many fill events, it is reasonable to think that the pilot may become fatigued. Conversely, the pilot may take a while to acclimate to the conditions at the start of each flight, i.e. a learning effect. Three visualizations for a possible learning/fatigue effect are given in Figure 5.

Figure 5 (a) groups all of the fills into three temporal categories (first third of fills, middle third, and last third). Within each group, the 1st and 3rd quartiles of angle of attack are computed at each time point. These quartiles are visualized with a ribbon that shows the change in the middle 50% over time. This is intended to visualize broad changes in patterns between the fills rather than individual outliers.

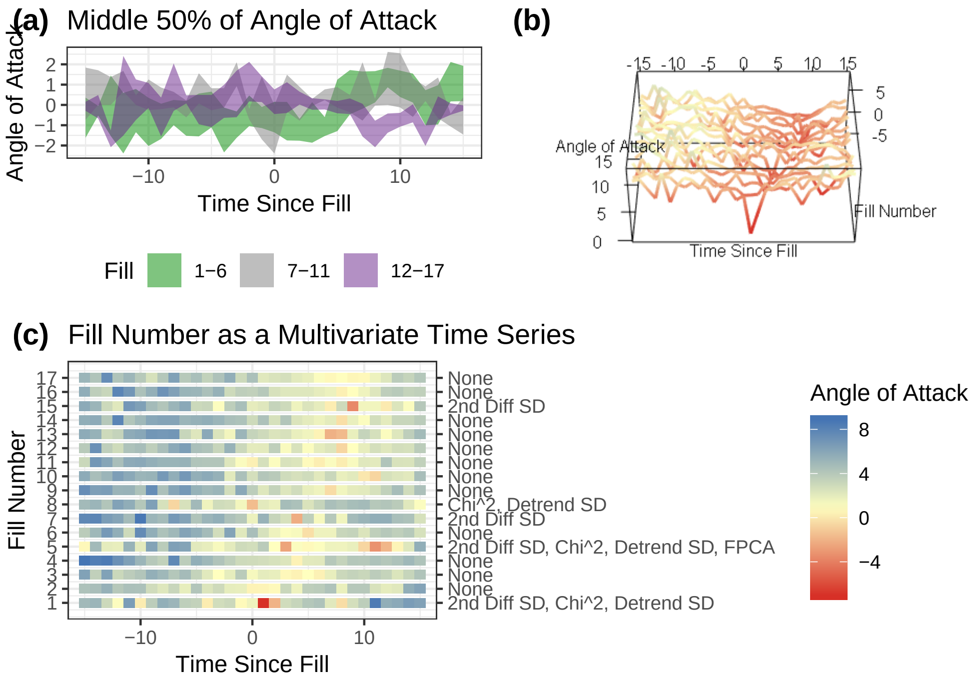

Figure 5: **(a)** Interquartile ribbon plot of detrended fills, showing that later fills (purple) had a different trend than earlier fills. **(b)** A stalagmite plot to determine if there's a smooth learning/fatigue effect for later fills. The z axis is angle of attack and the colour is red for low angle of attack and yellow for high. The first fill is closest to the viewport, the last fill is furthest away. **(c)** A time series strip chart, which is essentially the stalagmite plot shown from above. All three plots are attempting to determine if there's a learning/fatigue effect, but **(b)** also labels outlier fills by the control chart that detected them.

Before the time $t = 0$, the ribbons mostly overlap. Between $t = 5$ and 10, the later fills appear to have a much lower angle of attack. We cannot conclude that this is due to error or if it's intentional, but there appears to be either a learning or a fatigue effect. These plots have been made for several different flights and this pattern remains consistent.

In addition to the trend in Figure 5 (a), the vertical height of each ribbon indicates the interquartile range. The purple ribbon (fills 12 to 17) appears much thinner (less variable) from $t = 0$ to 10, which may indicate that this is indeed a learning effect rather than a fatigue effect.

The second plot is known as a stalagmite plot (see Albert-Green, Braun, Martell, and Woolford 2014). Both the height and colour of the lines are based on the value of the angle of attack. This is essentially Figure 1 (a) with the fills extended onto a third axis. The colour allows us to see a slight diagonal pattern in the red and yellow, which may corroborate the learning/fatigue effect observed in the interquartile ribbons plot.

The figure on the bottom row is a time series strip plot that is essentially the stalagmite plot shown from above. This is adapted from Peng (2008), which used this type of plot to visualize correlation between variables in a single observation of a multivariate time series. The learning/fatigue effect can be seen in the diagonal strip of yellow, but we can also see that the effect is stronger in the first few fills. Another advantage of this plot is the ability to label the outliers, which allows for a visual check of why the fill was flagged as an outlier. It also shows the distribution of outliers over the course of the flight - it appears that there were more outliers earlier in the flight, again indicating a learning effect rather than a fatigue effect.

4. Operationalized Visualizations for Flight Managers

It is perhaps overly ambitious to suppose that an inference procedure will adequately explain the inhomogeneity in one flight, and even more ambitious to suppose that a single inference procedure would be appropriate to all flights in one season. The visualizations that we have presented can be used to address many issues that may occur while flying, but they only apply to one flight at a time.

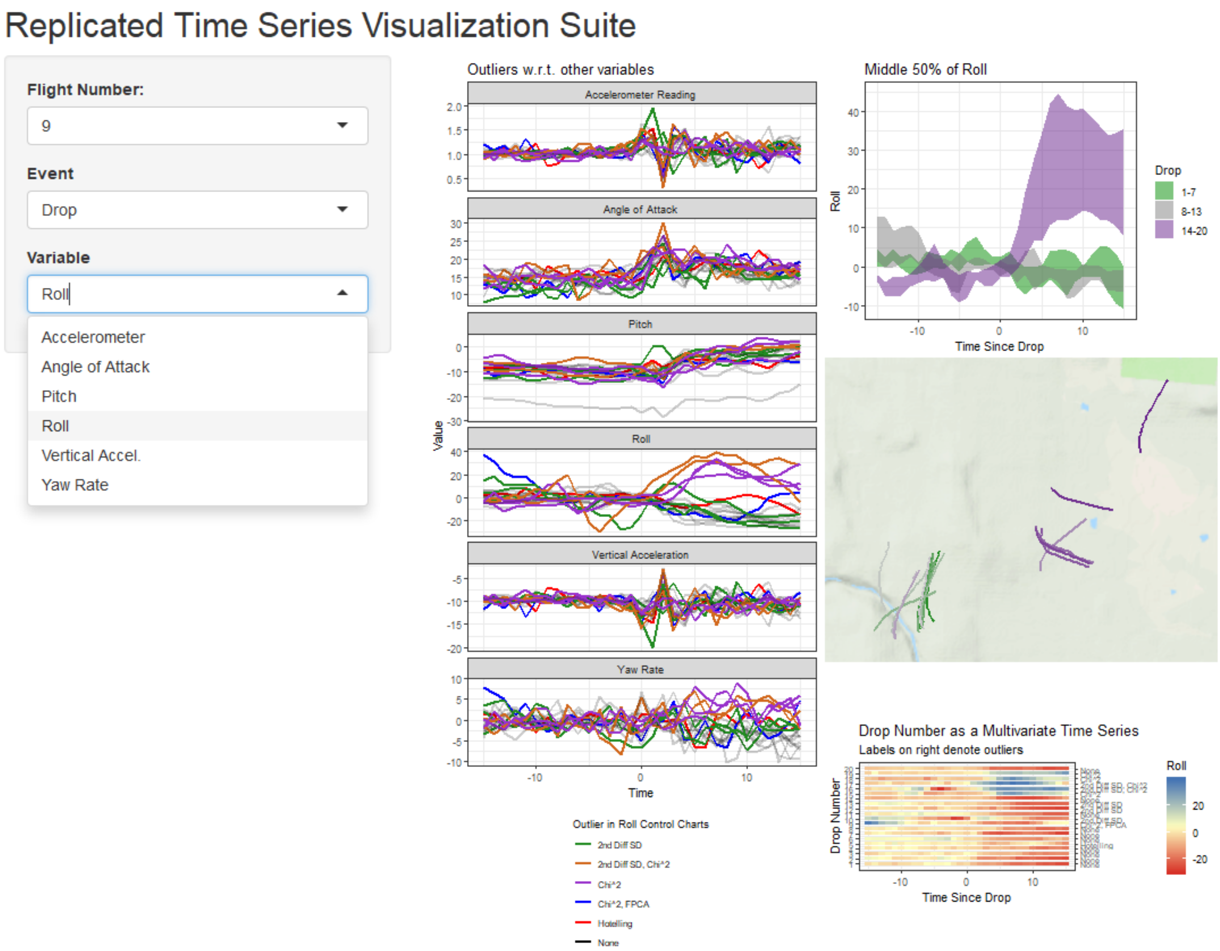

Figure 6: Screenshot of a "dashboard" made using `shiny`. Maps are ©2018 Google.

To investigate a season of flights, the code used to produce the above visualizations can be applied to other flights and variables. Using R's `shiny` and `cowplot` packages (R Core Team 2018; Chang, Cheng, Allaire, Xie, and McPherson 2018; Wilke 2018), the plots can be arranged as an application for researchers.

An example is shown in Figure 6, which shows the same flight as in the previous sections, but looking at roll rather than angle of attack and drops rather than fills. The interquartile ribbon plot in Figure 6 shows a different pattern for later fills. The map indicates that these later drops were performed at a different compass bearing. Thus, the difference is likely due to the wind direction. This insight is confirmed by the multivariate plot, which shows that all of the drops are essentially unremarkable in the other variables.

5. Discussion

In this paper, we have introduced and/or utilized several visualizations for outliers in replicated multivariate time series data that reveal important features. These visualization techniques are applied to aerial wildland fire fighting data, but are applicable to other sequences of time series.

Figure 1 (a) and (b) demonstrate the complex nature of our particular dataset and motivate our reasons for using more sophisticated visualization tools. Figure 2 (b) is the same as Figure 1 but with outliers labelled based on control chart methods. These plots indicate why certain fills were labelled as outliers, but they also highlight the types of outliers that each control chart is able to detect. This plot justifies our use of several different control charts.

Figure 5 (a) - the interquartile ribbon plot - is a novel way to discern the presence of a learning/fatigue effect. We acknowledge that, in our context, the interquartile ribbons are based on a small amount of data (5 or 6 fills). Having more observations would improve the precision when estimating the ribbon. However, they act as an early warning signal that there may be such an effect present. Figure 5 (b) - the stalagmite plot - is a way to interactively visualize this learning/fatigue effect in the raw data. One can rotate the box and zoom-in-and-out when viewing such a plot in R. Figure 5 (c) - the time series strip plot - shows this effect while also indicating which fills were outliers. By presenting the outliers this way, it is easier to discern between a learning effect and a fatigue effect. If there are more outliers in earlier (later) fills, then it is possibly a learning (fatigue) effect.

The multivariate plots in Figure 3 allow us to assign causes to the outliers. By our knowledge of the likely correlations in our data (e.g. angle of attack's relationship with pitch), we are able to differentiate between an outlier due to the pilots actions versus one due to external sources, such as a possible gust of wind.

The same analyses apply to all of the variables for both drops and fills during all flights in our data set. Since this quickly becomes an overwhelming number of plots, work has been done to consolidate this work into a manageable form. Using R's `shiny` (Chang *et al.* 2018) and `cowplot` (Wilke 2018) packages, a matrix of the plots in the paper can easily be constructed such that all of the relevant variables can be investigated through an interactive interface. This interface is shown in Figure 6. With appropriate data collection techniques and sufficiently robust data cleaning code, this application can be used in real time by flight managers during a wildfire. For other applications, this application represents a method for extending the analyses to multiple observations of replicated multivariate spatial time series.

Many different control chart methods were explored in our analysis of the data. We presented results that revealed salient features. However, there are other methods that may be more applicable if one was analyzing multivariate time series replicates in a different context. For instance, analysis of the residuals of time series processes (Psarakis and Papaleonida 2007) and analysing ARMA parameters with a Hotelling's T^2 control chart (Hotelling 1947; Apley and Tsung 2002). The methods presented in this paper were chosen to be contrasted with each other; each method represents a different assumption about the structure of the data. Our emphasis was an exploratory data analysis of these data. Although these control chart methods may not be perfectly suited to such data, they were shown to be useful for identifying outliers which can then be inspected for attributable causes.

We have previously attempted to quantify the learning/fatigue effect by studying the number of outliers early in the flight versus later in the flight. We also considered the time since the pilot's most recent flight as well as the total flight time in the past seven days versus the number of detected outliers. However, the learning/fatigue effect may be obscured by the anomalous fills that can be attributed to causes outside of the pilots' control. For these reasons we were unable to make any firm conclusions about the presence of learning or fatigue.

Buja, Cook, Hofmann, Lawrence, Lee, Swayne, and Wickham (2009) provide a method of statistical inference based on data visualization that could be used to evaluate the learning/fatigue effects. By randomly permuting the order of the fills in Figure 5 (c) we remove any possible learning or fatigue effect. The permuted data can be presented along with the original data to many individuals and we can get them to identify the plot that has a pattern. By showing this to many individuals we can get an estimate of the probability that a flight is identified as having a pattern. After doing this for many different flights, we can compare the probability of a flight having an identified pattern against the time since the most recent flight. If there is a relationship between these two variables then we would conclude that there is indeed a learning or a fatigue effect. Individuals can be contracted for such tasks through Amazon's Mechanical Turk platform (Amazon 2008).

6. Acknowledgements

The authors would like to thank Tim Curtis and David Martin from Latitude Technologies for providing data and guidance, Leigh Barratt from Conair Group for his valuable correspondence, and Jason Loeppky of the University of British Columbia - Okanagan Campus for his guidance and input. Funding was provided through a Mitacs Accelerate internship in collaboration with Latitude Technologies and the University of British Columbia. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding reference numbers RGPIN-2015-04221 and RGPIN-2014-06187]. Additional support was provided by a CANSSI Collaborative Research Team grant and the Institute for Catastrophic Loss Reduction.

This visualization work is done entirely in R (R Core Team 2018). Except for the maps and functional boxplots, all 2D plots were created in `ggplot2` (Wickham 2016) and all 3D plots were created in `RGL` (Adler and Murdoch 2018). Maps in 2D were made using `ggmap` (Kahle and Wickham 2013) and maps in 3D were made with `rgl` and `OpenStreetMap` (Fellows 2016, available under Open Database License, `openstreetmap.org/copyright`). Functional boxplots were made using the package `rainbow` (Shang and Hyndman 2016). The packages `dplyr` (Wickham, François, Henry, and Müller 2018) and `tidyr` (Wickham and Henry 2018) were used extensively for data manipulation.

Appendix A. Control Charts for Angle of Attack (AoA)

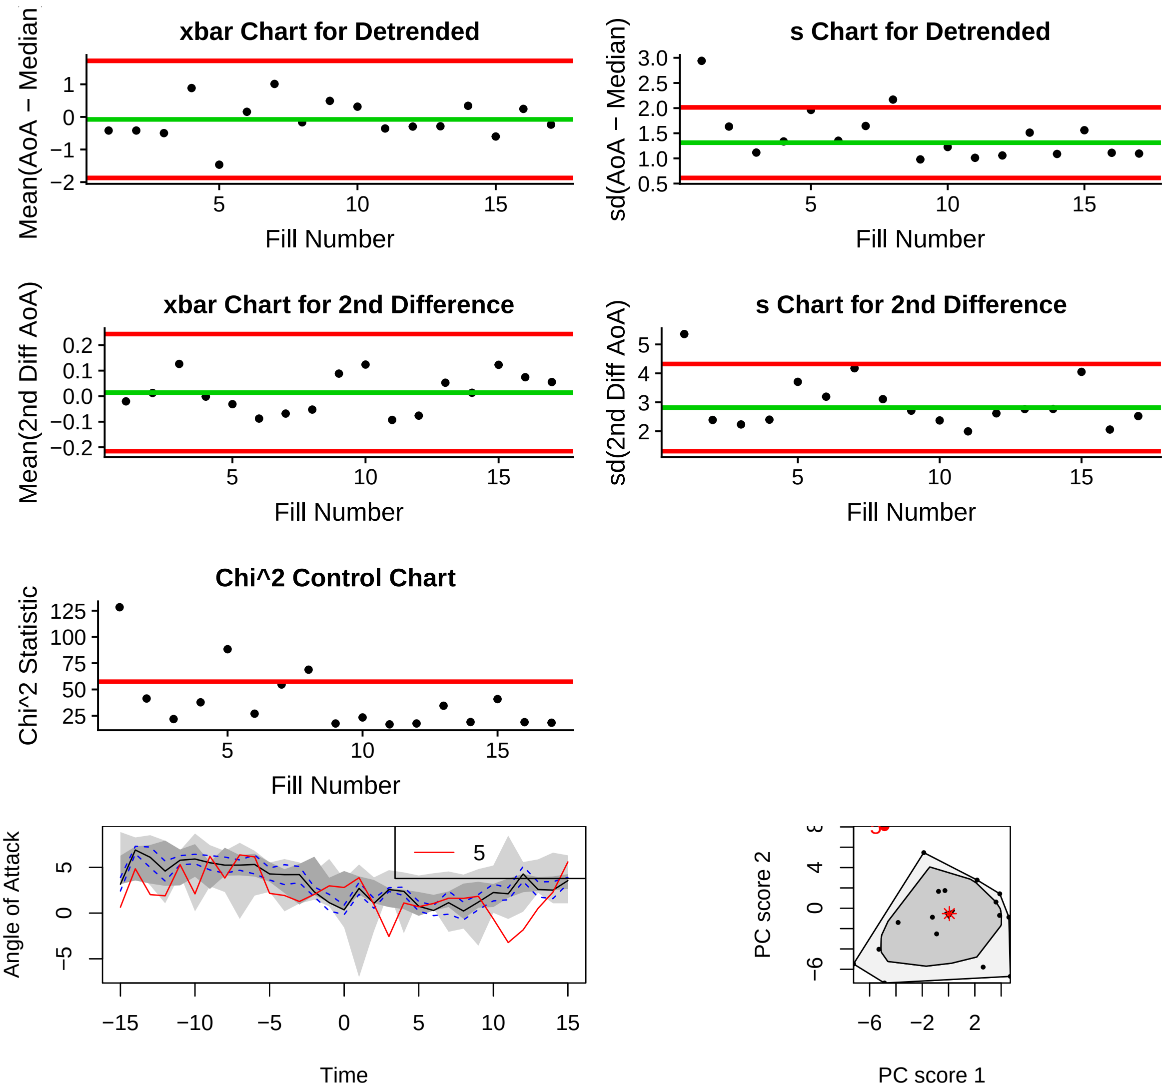

References

Adler D, Murdoch D (2018). "rgl: 3D Visualization Using OpenGL." URL https://cran.r-project.org/package=rgl.

Albert-Green A, Braun WJ, Martell DL, Woolford DG (2014). "Visualization tools for assessing the Markov property: Sojourn times in the forest fire weather index in Ontario." *Environmetrics*, **25**(6), 417–430. ISSN 1099095X. doi:10.1002/env.2237.

Amazon (2008). URL https://www.mturk.com.

Apley DW, Tsung F (2002). "The Autoregressive T^2 Chart for Monitoring Univariate Autocorrelated Processes." *Journal of Quality Technology*, **34**(1), 80–96. ISSN 0022-

4065. doi:10.1080/00224065.2002.11980131. URL https://www.tandfonline.com/doi/full/10.1080/00224065.2002.11980131.

Brillinger DR, Finney MA (2014). "An exploratory data analysis of the temperature fluctuations in a spreading fire." *Environmetrics*, **25**(6), 443–453. ISSN 1099095X. doi:10.1002/env.2279.

Buja A, Cook D, Hofmann H, Lawrence M, Lee EK, Swayne DF, Wickham H (2009). "Statistical inference for exploratory data analysis and model diagnostics." *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, **367**(1906), 4361–4383.

Chang W, Cheng J, Allaire J, Xie Y, McPherson J (2018). "shiny: Web Application Framework for R." URL https://cran.r-project.org/package=shiny.

Fellows I (2016). "OpenStreetMap: Access to Open Street Map Raster Images." URL https://cran.r-project.org/package=OpenStreetMap.

Goode JH (2003). "Are pilots at risk of accidents due to fatigue?" *Journal of Safety Research*, **34**(3), 309–313. ISSN 00224375. doi:10.1016/S0022-4375(03)00033-1.

Hotelling H (1947). "Multivariate quality control." *Techniques of statistical analysis*.

Hyndman RJ, Shang HL (2010). "Rainbow plots, bagplots, and boxplots for functional data." *Journal of Computational and Graphical Statistics*, **19**(1), 29–45. ISSN 10618600. doi:10.1198/jcgs.2009.08158.

Jackson CA, Earl L (2006). "Prevalence of fatigue among commercial pilots." *Occupational Medicine*, **56**(4), 263–268. ISSN 09627480. doi:10.1093/occmed/kql021.

Kahle D, Wickham H (2013). "ggmap: Spatial Visualization with ggplot2." *The R Journal*, **5**(1), 144–161. URL http://journal.r-project.org/archive/2013-1/kahle-wickham.pdf.

Miller NL, Matsangas P, Shattuck L (2008). *Fatigue and its effect on performance in military environments*. Ashgate Publishing Burlington, VT.

Nash T, St Arnaud L, Tithecott A, Simpson B, Stocks B, *et al.* (2017). "A review of the 2016 Horse River Wildfire." *Technical report*, MNP Consulting. URL https://www.alberta.ca/assets/documents/Wildfire-MNP-Report.pdf.

Peng RD (2008). "A method for visualizing multivariate time series data." *Journal of Statistical Software*, **25**(March), 1–17. ISSN 15487660. doi:http://dx.doi.org/10.18637/jss.v025.c01. URL http://www.educonv.com/tw{_}files2/urls{_}246/179/d-178761/7z-docs/1.pdf.

Psarakis S, Papaleonida GEA (2007). "SPC Procedures for Monitoring Autocorrelated Processes." *Quality Technology & Quantitative Management*, **4**(4), 501–540. ISSN 1684-3703. doi:10.1080/16843703.2007.11673168. URL http://www.tandfonline.com/doi/full/10.1080/16843703.2007.11673168.

R Core Team (2018). "R: A Language and Environment for Statistical Computing." URL https://www.r-project.org/.

Shang HL, Hyndman RJ (2016). "rainbow: Rainbow Plots, Bagplots and Boxplots for Functional Data." URL https://cran.r-project.org/package=rainbow.

Shewhart WA (1925). "The application of statistics as an aid in maintaining quality of a manufactured product." *Journal of the American Statistical Association*, **20**(152), 546–548.

Strauss S (2010). "Pilot Fatigue." *Aerospace Medicine NASA/Johnson Space Center, Houston, Texas. http://aeromedical. org/Articles/Pilot_Fatigue. html*, pp. 28–32. URL http://europeanpilotacademy.com/wp-content/uploads/2013/11/Pilot-Fatigue.pdf.

Tukey JW (1977). *Exploratory Data Analysis*. Addison-Wesley Publishing Company.

Wickham H (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-3-319-24277-4. URL http://ggplot2.org.

Wickham H, François R, Henry L, Müller K (2018). "dplyr: A Grammar of Data Manipulation." URL https://cran.r-project.org/package=dplyr.

Wickham H, Henry L (2018). "tidyr: Easily Tidy Data with 'spread()' and 'gather()' Functions." URL https://cran.r-project.org/package=tidyr.

Wilke CO (2018). "cowplot: Streamlined Plot Theme and Plot Annotations for 'ggplot2'." URL https://cran.r-project.org/package=cowplot.

Zhang H, Albin S (2009). "Detecting outliers in complex profiles using a χ^2 control chart method." *IIE Transactions*, **41**(4), 335–345.

Affiliation:

Devan Becker
Department of Statistical and Actuarial Sciences, University of Western Ontario
London, Ontario
E-mail: dbecker7@uwo.ca

Journal of Environmental Statistics http://www.jenvstat.org
Volume 9, Issue 1 Submitted: 11/14/18
September 2019 Accepted: 9/15/19