Reliability Engineering, Data Science, Machine Learning and Bias

Howard W Penrose, Ph.D., CMRP
President, MotorDoc® LLC
www.motordoc.com

Abstract: When dealing with statistical analysis and probabilities its often the missing information that can be an issue. This can result in bias when working in areas such as reliability, data science and machine learning that result in incorrect solutions. We've seen that in such things as the percentage of type of motor failures, faulty IoT devices, false positives and negatives in healthcare and prognostic software, and other studies. The assumptions make sense from what appears to be common sense, but often that remains part of the bias that is used to reinforce our perception of reality and repetition of the same errors.

INTRODUCTION

While perusing LinkedIn, I ran into another article on bias using the same image that I've seen (Figure 1) representing a military study on the survival of aircraft during wartime. The article entitled "Punching Holes in Data Analysis,"[1] by Joel Manzer was an excellent study on *Survivorship Bias*. This type of bias relates not to the presented data and data interpretation, but what the missing data represents.

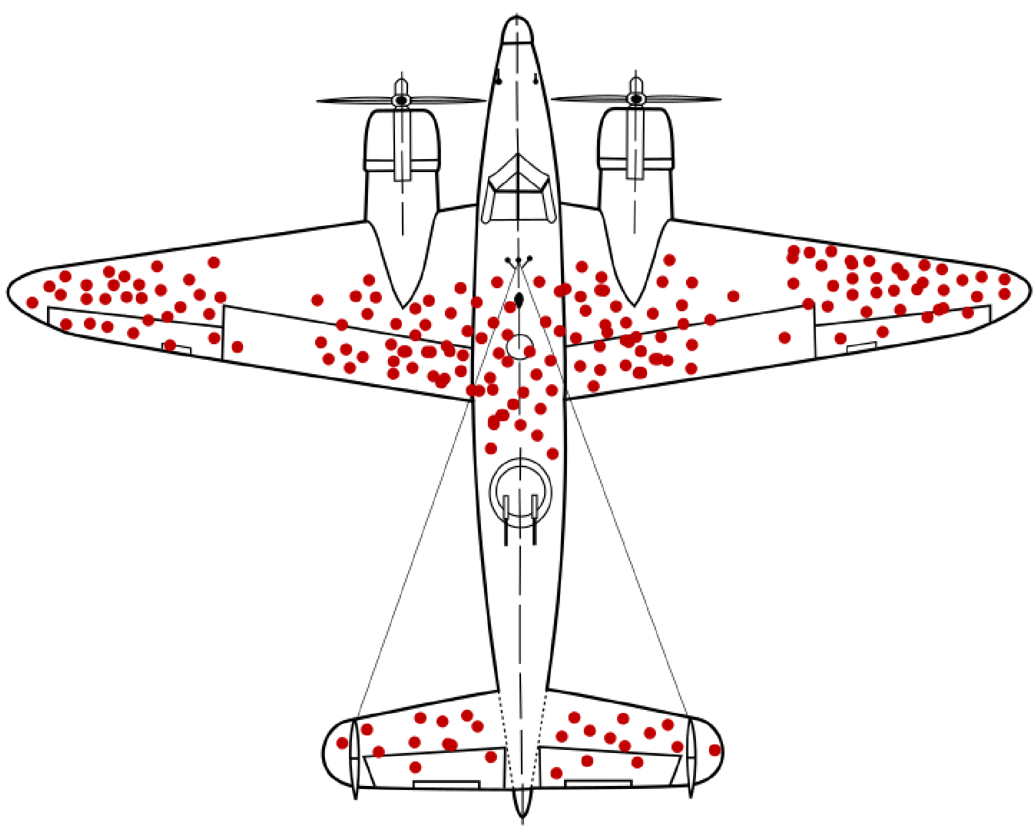

Figure 1: relates to the Abraham Wald study on bomber aircraft survivability. Credits: By Martin Grandjean (vector), McGeddon (picture), Cameron Moll (concept) - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=102017718

[1] (1) Punching Holes in Data Analysis | LinkedIn

The study on aircraft entitled, "A Method of Estimating Plane Vulnerability based on Damage of Survivors,"[2] was developed in 1943 by the Statistical Research Group at Columbia University for the National Defense Research Committee in 1943. It was created to generate "… methods of estimating the vulnerability of various parts of an aircraft based on damage to surviving planes."[3] What made this study particularly important was that they tied missing data and existing data in order to determine the probability of survival about the number of hits, location of hit, and hit by type of weaponry by making educated assumptions related to damage to returning aircraft. This approach is often lacking in modern reliability engineering and data science where we are taught to primarily review the existing data only and that other data that is not present, or worse, does not support our conclusions, is the data that is removed.

Having audited a number of courses related to machine learning and data science the primary teaching is to first perform data manipulation and select features that support the type of conclusions that you are looking for. The intention is to select useful data to provide accurate results within some level of error. Unfortunately, bias is included and taught, such that the data scientist or reliability engineer determines a level of acceptable error to generate false negatives or positives.

For instance, I have visited facilities where unplanned failures are obviously frequent and having an impact on production, yet I was told that the availability is 99%. When probing deeper you discover that the definition of unplanned versus planned failure is set up in such a way that incredibly important data is discarded because corrective action was performed on the same shift, within 24 hours, or some other criteria that does not capture true unplanned outage time. The bias towards excluding such data in order for the facility, or an individual, to meet some goal results because expectations are presented without identifying exactly what data is to be included, or excluded.

A FEW CASES AND EXPERIENCES

None of us are invulnerable to bias, especially when that bias is supported through marketing efforts, industry materials, and articles. Yes, even this one. I am biased towards realizing there is bias in everything.

I discussed this particular issue on electric motor reliability in an article Large Electric Motor Reliability What the Studies Really Say – MotorDoc LLC in which we reviewed the oft-quoted 1983 EPRI study on electric motor failure rates. Prior to the writing of the article, the data was quoted based upon prior quotes and presentations that no one was able to trace back to the real data. IEEExplore[4] finally published the original papers associated with the studies from 1962 to 1995 in 2010 which provided new insights into the original data. The challenge that we faced after the real data became available is that the results contradict what companies have been presenting to sell their reliability and maintenance offerings for decades. Therefore, the industry bias to maintain that bearings represent the largest cause of failure and rotor defects are 1/10 of motor faults continues to be promoted. A number of items had been omitted from the presentations including that 1/3rd of the provided data was not included and that motor repair shops automatically replace bearings through the repair process even when no defects are found. The study only covered failed motors and the repair reports provided and parts replaced without context or other confirmation.

---

[2] ADA091073.pdf (dtic.mil) (original report – public)
[3] Report documentation page for ADA091073.pdf.
[4] https://ieeexplore.ieee.org

In a recent data science and machine learning project for a facility that we are using EMPATH™ to monitor utility-supplied data we noted some unusually high voltages and voltage unbalances from 13% and 28% respectively of the time across a four month period (Figures 2 and 3). These are well outside the normal range of what should be expected. However, local rules, regulations and standards have specific definitions on what is reportable. The result is that the end-user was never made aware of these conditions including one in which a high voltage distribution line was failing and the resulting phase to phase shorts each cycle were impacting production ([Part IV, More Thoughts on Electrical Reliability (Specifically, Is Your Electricity Clean?) | THE RAM REVIEW]).

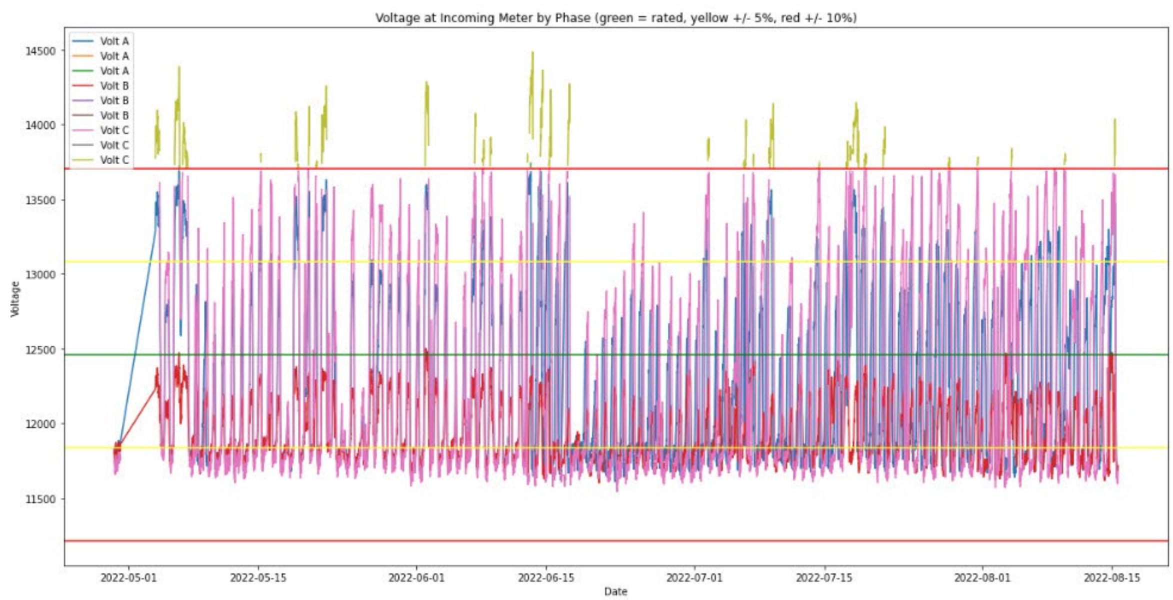

Figure 2: Overvoltage up to 114% of rated supply voltage per phase. Max is +/-10%.

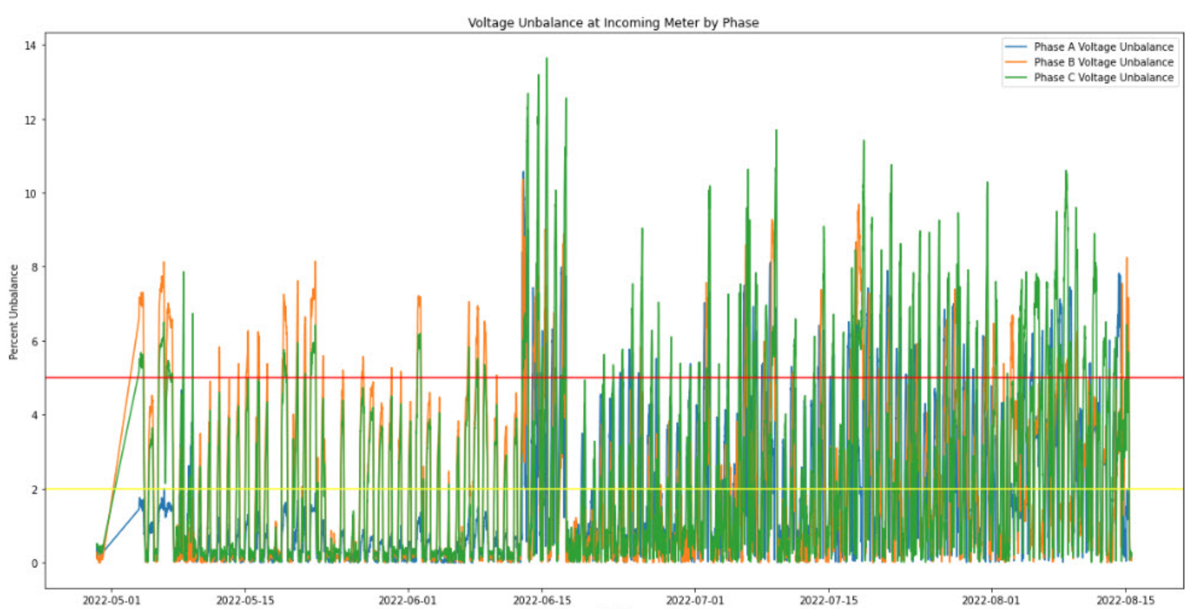

Figure 3: Voltage unbalances of up to 14% phase to phase by phase. Max is 5%.

In the attempt to rapidly produce desktop-developed with limited real-world experience IoT devices (read my bias there?) we are observing an increase in the number of false positive and negative results. This has led to the associated McKinsey report that cited 18% of IoT device outputs for predictive maintenance were accurate. As noted in [How Accurate Is IoT-Based CBM? | THE RAM REVIEW](), the bias tends towards not missing anything, so technology that is not capable of detecting certain defects is producing false positives that have significant impact on maintenance and questioning of vendors.

In a rather personal situation, in 2021, I had a reaction to a medication. This resulted in a serious condition that required emergency action. During the procedure surgeons reviewed my medical records and found data that conflicted with what they were seeing in tests and visually. The result of the bad data was very nearly fatal except that the surgeon relied upon his experience and not the other data on the computer. We later also found that another member of the family had diagnosis in their medical records that were profoundly incorrect. It appears, reportedly, that the medical organization was purchased by an investment company that introduced a new software system that utilized machine learning to review all cases. The ML system was allowed to update medical records with very little supervision and a bias towards false positives. There's more to the story, but ultimately it led us to more research into the general accuracy of machine learning systems and IoT.

TWIST OF EVENTS

There are numerous other examples of bias where either data is left out or a system is biased towards presenting in a specific way. In the case of the original premise behind this story, Abraham Wald, who led the project, specifically identified in their statistical methodology a method for estimating and working with missing data. The objective of the study was to eliminate bias based upon damage to returning aircraft and the number of 'hits' an aircraft could take. Several of the urban legends detail the story in a way that suggests that the military was going to armor the areas that were hit and that Wald was a military analyst that noted that these were aircraft that survived and that the exposed areas (the original purpose of Figure 1) could take hits and didn't require armor. In this case the bias was based upon an image and not the actual report.

The actual report was based upon the number of hits that an aircraft could take based upon experiments and damaged returning aircraft. They then divided the aircraft up into sections in order to determine the probability of which sections and by which type of ammunition, as well as angle of penetration, an aircraft could absorb before failing. In effect, a detailed study on n-1 hits by different ammunition and the probability of loss of an aircraft in order to justify type and location of armor. As those in the military know, nothing happens without a study.

In effect, a lot of the articles and information presented on the Abraham Wald story about bias was… well… biased.

CONCLUSION

Bias is insidious. It leaks into all areas of reliability, data science and machine learning. The key is to be aware of that bias and constantly challenge the data. Otherwise the consequences can be dire. One of the most challenging areas of bias is the missing data or context without which you may be dealing with 'survival bias,' or a bias towards data and information that is included while other data is excluded because it is not present. We also recommend the points presented by Joel Manzer's LinkedIn article.