8 July 2023
Revised 4/21/024

Seasonality and Variance of Twitter Sentiment Regarding Electric and Automated Vehicles

Michael McGurrin
McGurrin Consulting
mike@mcgurrin.com

Abstract

There have been previous studies exploring sentiment analysis of Twitter data concerning Automated Vehicles (AVs), however this study took a unique approach by analyzing Twitter sentiment concerning both AVs and Electric Vehicles (EVs) on a daily basis for a period of over one year. Rather than examining sentiment changes through isolated snapshots in time, this work delved into the seasonality and variance of sentiment expressed in tweets across a span of 421 days.

To obtain the necessary data, the study employed a customized set of keywords and filters to create two distinct corpuses of tweets, one regarding EVs and one regarding AVs. The sentiment analysis tool used in this study was VADER (Valence Aware Dictionary and sEntiment Reasoner), which assigns positive, negative, or neutral polarity values to the tweets.

The sentiment scores for both sets of data revealed clear patterns of weekly seasonality. Specifically, there was a decrease in the number of positive and neutral tweets posted during weekends. Although the long-term trend for sentiment scores for both EVs and AVs were flat over the period studied, the sentiment regarding AVs exhibited significantly higher variance, and both showed a notable negative skewness.

The paper concludes with suggestions for future research, including examining the relationship between sentiment scores and real-world events.

Introduction

While previous studies have used Twitter data for sentiment analysis in the context of Automated Vehicles (AVs), and other research has utilized snapshot-based approaches to examine sentiment changes, this study takes a different approach by analyzing continuous time series data of daily Twitter sentiment over a span of 421 days, from 28 February 2022 through 24 April 2023.

This study created two distinct corpuses of tweets. The first set pertains to electric vehicles (EVs), while the second set pertains to automated vehicles (AVs). This study investigated the presence of seasonality patterns within these tweet collections, as well as the long-term trends and variance in sentiment scores.

This paper is one of two related papers. The second paper, *The Best of Days, the Worst of Days: Twitter Sentiment Regarding Automated Vehicles,* to be published, examines the days with some of the most positive and most negative sentiment scores regarding AVs and the news stories that were published just before or at the same time.

Background and Related Work

Automated Sentiment Analysis

Sentiment Analysis refers to the process of analyzing text samples to determine the sentiment expressed by the author, typically categorized as positive, negative, or neutral. Manual examination of large volumes of text is labor-intensive, and disagreement among human judges regarding sentiment further complicates the task. Therefore, manual assessment is usually performed on small samples to train or evaluate automated methods. Various natural language processing and machine learning techniques have been employed to automate sentiment analysis.

Machine-based sentiment analysis, however, can also be challenging. Reasons for this include the presence of negated words, slang, emojis, and sarcasm. While VADER, the sentiment analysis tool used in this study, is designed to handle many of these cases, it is not perfect. Research has shown that even human raters agree with each other less than 80% of the time, even when dealing with just two sentiment categories (Skystance, 2022). With three sentiment values and an equal number of tweets expressing each sentiment, random chance would yield approximately 33% accuracy. In this study, the algorithm was tested using a similar set of tweets related to automated vehicles (originally generated with different keyword selection criteria), and it agreed with a human annotator 53% of the time. However, the purpose of the sentiment indices is not to obtain precise sentiment polarity values, but rather to track changes over time and correlate them with external events.

Twitter has been commonly used as a data source for sentiment analysis due to its vast number of available tweets, many of which express strong opinions. Additionally, Twitter provided an Application Programming Interface (API) for easy access and filtering of tweets.[1] Despite the fact that only a small percentage of users tweet about automated or electric vehicles, the volume still amounts to hundreds of new tweets per day on these topics. However, there are disadvantages associated with tweets as well. Many bots simply repost news from other sources, resulting in identical or nearly identical tweets. This project uses custom code to attempt to filter these out such duplicates.

Related Work

Numerous studies have examined sentiment regarding automated and electric vehicles, including the impact of reporting of crashes on sentiment. For instance, Li et al. developed a corpus of 50,000 comments posted regarding YouTube videos discussing automated vehicles. Their analysis showed that sentiment scores remained lower for a couple of months following news of a crash (Li, Lin, Choi, & et.al., 2018).

A recent sentiment analysis study of over 30,000 tweets about automated vehicles, Gupta and Sharma (2022) found that a plurality of tweets expressed positive sentiments, with only 16.6% expressing negative sentiments (Gupta & Sharma, 2022). This study also explored the most frequently used words in tweets, classified by sentiment.

Another study analyzed Twitter sentiment three days before, on the day of, and three days after the crash of a Tesla vehicle on February 10, 2019, (Jefferson & Mcdonald, 2019). This research examined both sentiment and word frequencies, revealing a slight decrease in negative tweets during and after

[1] In 2023, Twitter (now X) introduced new policies, removing the free API search capability that was used for this and many other studies. The minimum fee for limited search capability was set at $100 per month.

the crash, but a significantly larger decrease in positive tweets. The choice of a three-day timeframe was based on a previous study involving over 1.7 million tweets analyzed around the time of two vehicle crashes in 2018, one involving an Uber vehicle and a pedestrian, and the other involving a Tesla Model X (Penmetsa & et.al., 2021).

Approach

A combination of the Twitter API v2 (Twitter, n.d.)] and snscrape (JustAnotherArchivist, n.d.), a social media scraping tool, was utilized to extract tweets discussing AVs or EVs based on one or more specific keywords being present in the body of the tweet. For AVs, the keywords used were:

- #selfdriving
- self driving
- driverless
- automated cars
- automated vehicles
- automated trucks
- automated busses
- automated shuttles
- autonomous cars
- autonomous vehicles
- autonomous trucks
- autonomous busses
- autonomous shuttles

For EVs, the keywords used were:

- #EVs
- #ElectricVehicles
- electric vehicles
- ElectricVehicle
- electric cars
- electric trucks
- electric busses

Each of the two datasets were filtered to comprise original tweets, excluding retweets, replies, and quote tweets. Additionally, efforts were made to filter out identical or nearly identical tweets, which are often generated by automated bots on Twitter. The tweets were indexed by their posting dates, and the sentiment expressed in each tweet was analyzed to determine its polarity (positive, neutral, or negative). Separate analyses were conducted for tweets related to AVs and EVs.

Various computer-based approaches exist for sentiment assessment, including machine learning algorithms and neural networks. In this project, VADER (Valence Aware Dictionary and Sentiment Reasoner), an open-source software package, was used to calculate the AV and EV sentiment indices package (Hutto & Gilbert, VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text, 2014) (Hutto & Gilbert, vaderSentiment, n.d.). VADER is a lexicon and rule-based analysis tool specifically designed for analyzing social media posts such as tweets. For this study, some minor additions were made to the lexicon to tailor it to the specific subject matter. Specifically, “advances” and “woot” were added as words expressing positive sentiment, while “dystopia,” “dystopian,” “against,” and “disaster” were added as words expressing negative sentiment. Although this rule-based approach is relatively simple compared to other methods, comparison studies have demonstrated its effectiveness in analyzing social media, often performing on par or better than alternative approaches, with the added advantage of quick processing speed.

Two measures were computed from each daily sample. The score, or average sentiment, represents the simple average of all sampled tweets, with negative tweets assigned a value of -1, neutral tweets assigned 0, and positive tweets assigned +1. Therefore, the score ranges from -1 (indicating entirely negative tweets) to +1 (indicating entirely positive tweets). The positive/negative ratio considers only positive and negative tweets, excluding neutral ones, and represents the ratio of positive tweets to negative tweets. Examples of these metrics, along with additional information, can be accessed on the AV & EV Sentiment Indices website (Michael, 2023).

Analysis

Weekly Seasonality and Trend Lines

Time series data can be analyzed by breaking it down into five components (Date, n.d.):

- Level: The average value of the data
- Trend: The long-term, persistent increase or decrease in the value over time
- Cyclical: Non-periodic changes in the value of the data that repeat but do not follow a fixed pattern
- Seasonal: Patterns in the value of the data that repeat over a fixed period, such as years, seasons, days, or time of day
- Noise: The residual variation that remains in the data after accounting for the other components.

Decomposition of the Number of Tweets

Upon visual inspection, there is a clear weekly pattern observed in the number of tweets posted about EVs, with a decrease in tweets on weekends (Figure 1). The trend is visible but less distinct in the tweets about AVs (Figure 2). In a study by Li et al., a similar drop in the number of comments on YouTube was observed during weekends, but the variation in sentiment was not reported (Li, Lin, Choi, & et.al., 2018).

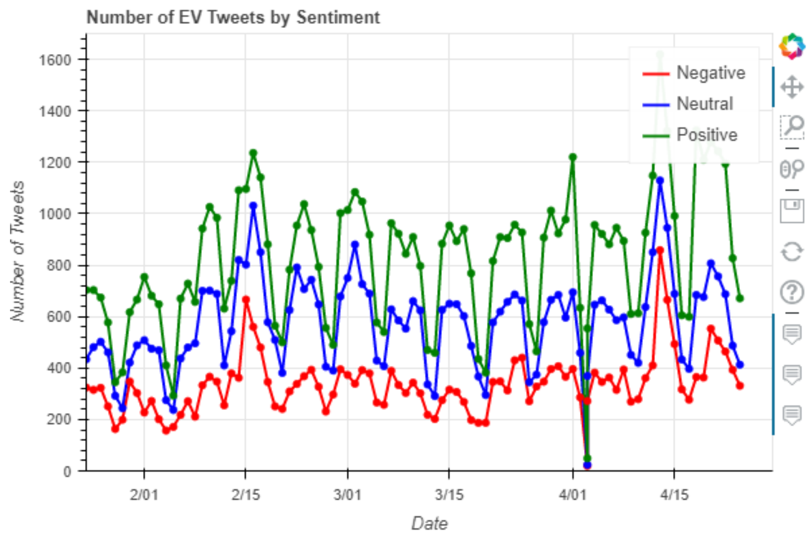

Figure 1. Time Series Plot of Positive, Neutral, and Negative Tweets Regarding Electric Vehicles

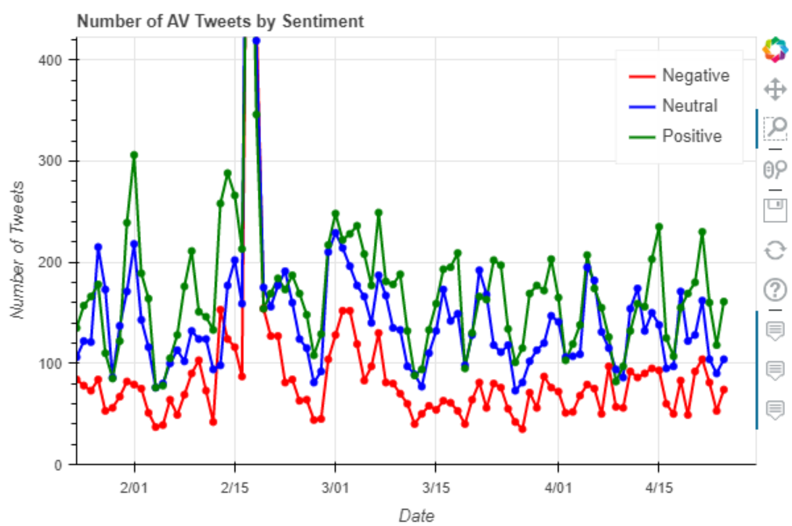

Figure 2. Time Series Plot of Positive, Neutral, and Negative Tweets Regarding Automated Vehicles

Using the seasonal_decompose module (statsmodels.tsa.seasonal.seasonal_decompose, n.d.) from the Python statsmodels library (statsmodels, n.d.) with an additive model with a seven-day period, it is confirmed that a weekly pattern exists, as demonstrated in Figure 3 and Figure 4. Each set of tweets related to EVs or AVs, along with each sentiment value, is presented in four subplots against time. The top subplot displays the original raw count data, followed by the trend subplot (a combination of the level, trend, and cyclical components), which is obtained by applying a moving weighted average to the data. The third subplot presents the calculated values for each component of the period (in this case, each day of the week), obtained by subtracting the trend from the raw data. The fourth subplot illustrates the remaining noise after subtracting the trend.

The weekly component is more prominent in the EV data (with a Pearson autocorrelation coefficient of 0.73 for the total number of tweets) compared to the AV data (coefficient of 0.20 for the total number of tweets). Furthermore, there is a greater difference in the number of positive and neutral tweets posted on weekdays versus weekends than in the number of negative tweets. Table 1 displays the Pearson correlation coefficients for a lag of seven days, illustrating these differences for both EV and AV tweets. The reason behind these discrepancies is unknown.

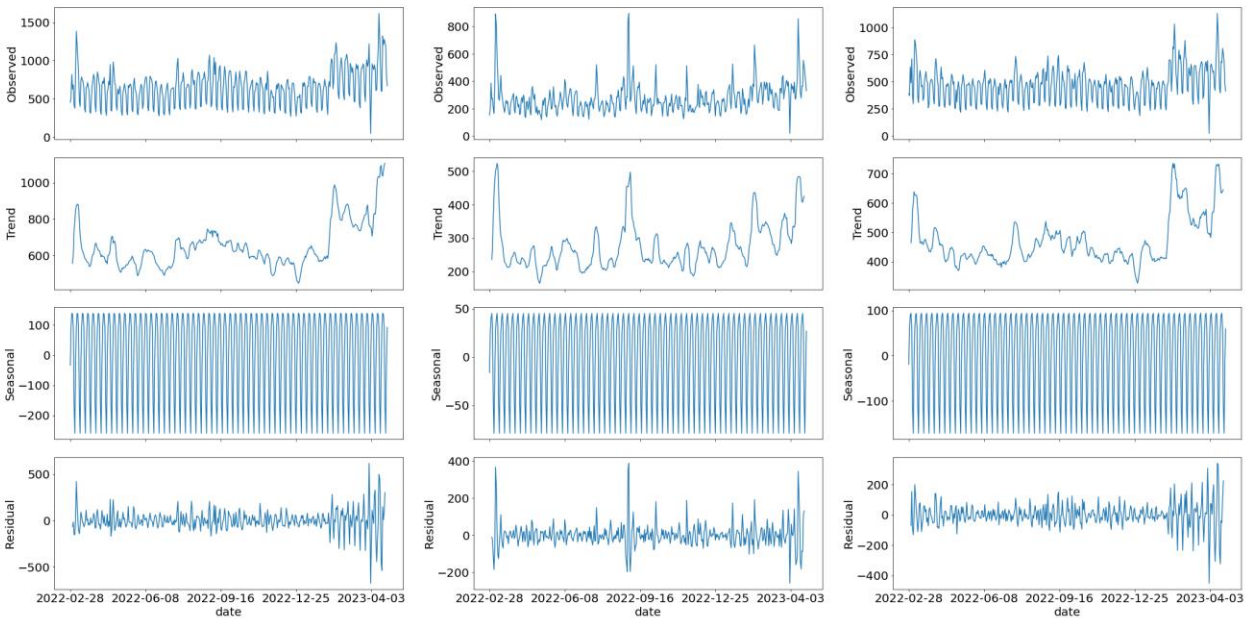

Figure 3. Decomposition of EV Sentiment Data, for Positive (left), Negative (Center) and Neutral (right) Tweets

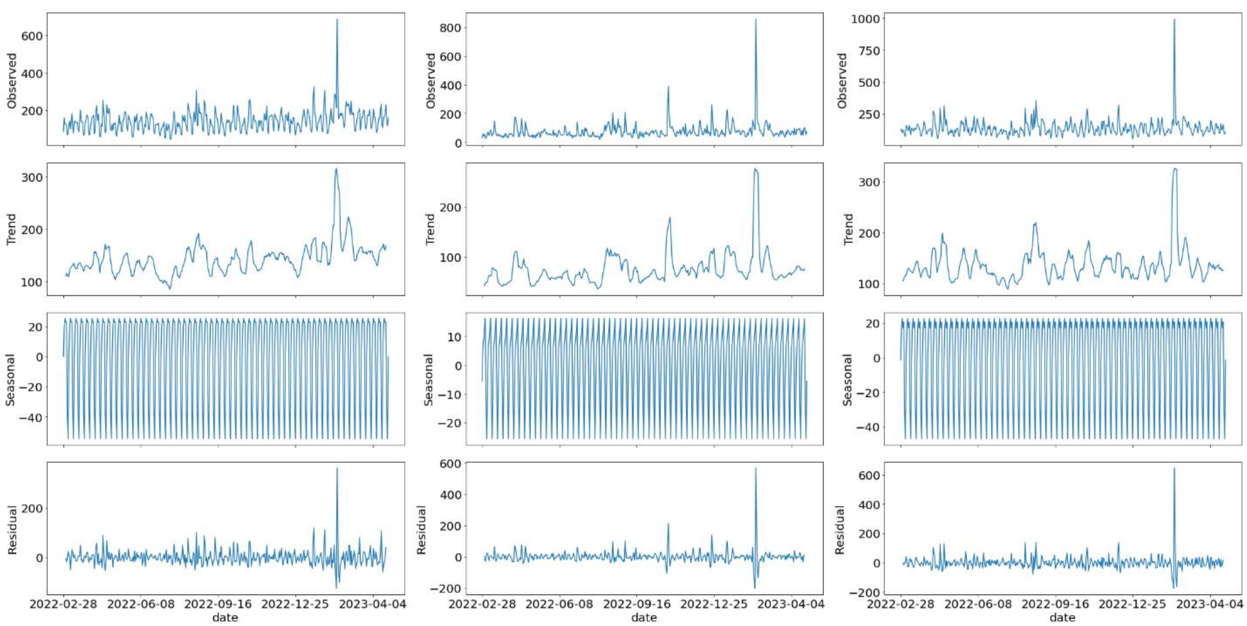

Figure 4. Decomposition of AV Sentiment Data, for Positive (left), Negative (Center) and Neutral (right) Tweets

Table 1. 7-Day Pearson Autocorrelation Coefficients

EV Tweets		AV Tweets	
Sentiment	**7-Day Autocorrelation**	**Sentiment**	**7-Day Autocorrelation**

Positive	0.78	Positive	0.38
Negative	0.41	Negative	0.08
Neutral	0.78	Neutral	0.38
All Tweets	0.73	All Tweets	0.20

Decomposition of the Sentiment Score

The sentiment score exhibits a drop on weekends due to a significant decrease in the number of positive and neutral tweets, while the decline in negative tweets is less pronounced. The 7-day seasonal components for AV and EV sentiment scores are depicted in Figure 5 and Figure 6, revealing that the seasonal component is more substantial for EV sentiment. The weekly autocorrelation factor for EV sentiment is 0.37, while it is only 0.09 for AV sentiment.

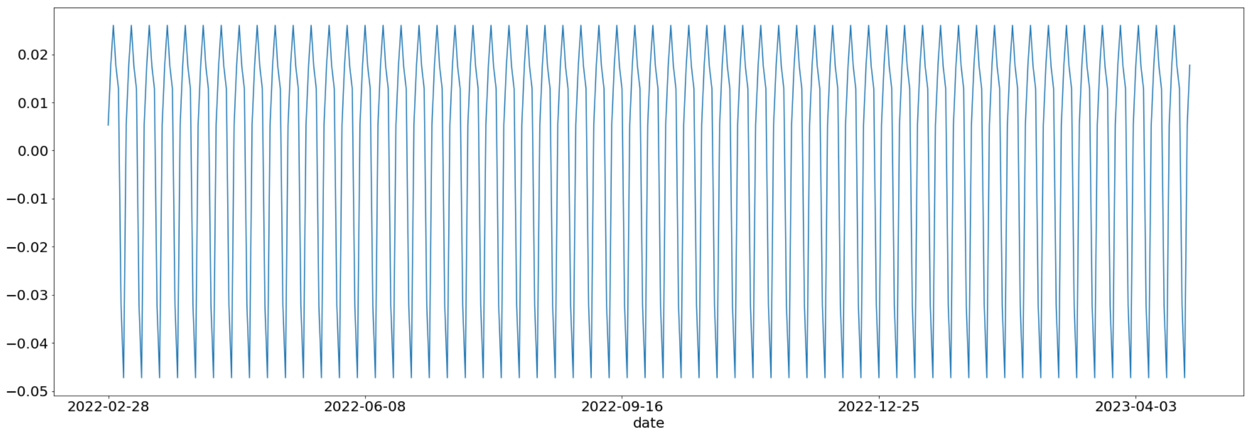

Figure 5. Seasonal Component of EV Sentiment Score, Showing Weekly Seasonality

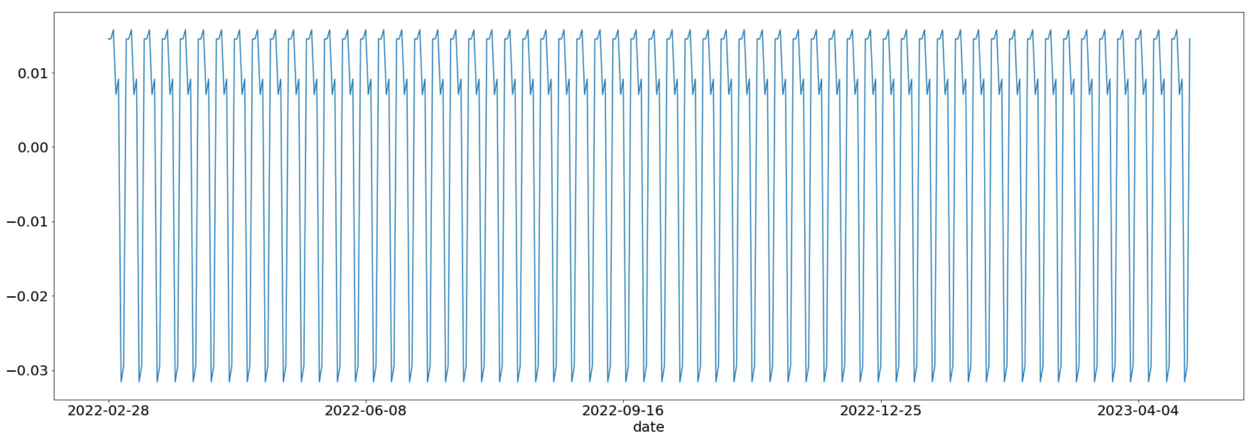

Figure 6. Seasonal Component of AV Sentiment Score, Showing Weekly Seasonality

Long-Term Sentiment Trends, Variance, and Skewness

Plotting a least squares linear fit (represented by the lines in Figure 7 and Figure 8) demonstrates that both EV and AV sentiment trends are flat, indicating no long-term shift towards more positive or negative sentiments on either topic occurred during the period studied. The best fit slope for EV sentiment is 2.2×10^{-5}, while for AV sentiment, it is 9.4×10^{-6}. Concerns have been raised in the literature concerning how negative publicity surrounding crashes involving automated vehicles would impact public acceptance. Over the time covered by this data, there were multiple crashes that received significant publicity. Although these crashes *did* significantly lower the sentiment scores for several days, in each case the sentiment scores rebounded, and at least as far as Twitter sentiment is concerned, did not appear to have any lasting impacts.[2]

[2] For example, on December 21st and 22nd, 2022, multiple media sources reported that the cause of an eight-car pileup on the San Francisco Bay Bridge was a Tesla in "full self-driving" mode. The ratio of positive to negative tweets went from 2.1 to 1 on December 21st down to 0.6 to 1 on December 22nd. But by December 26th, the ratio was back to 2 to 1. More information on sentiment score peaks and valleys, and the half-life for reversion the mean is covered in a forthcoming paper, *The Best of Days, the Worst of Days: Twitter Sentiment Regarding Automated Vehicles.*

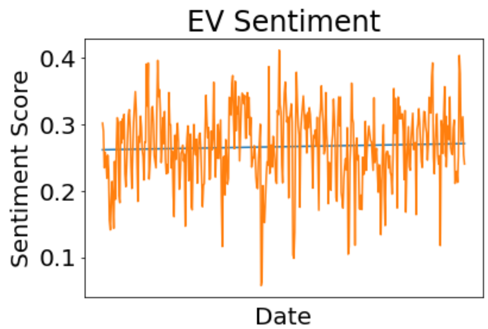

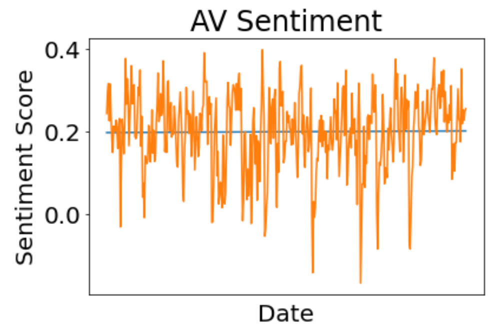

Figures 7 and 7. EV and AV Sentiment Scores over Time, Along with Best Fit Lines

While the sentiment trends regarding both EVs and Avs are flat, the variability of the two data sets are quite different. The variance of the EV sentiment score was 0.0033. while the variance of the AV sentiment score was 0.0089, over two and a half times higher. Both data sets are negatively skewed. The Fisher-Person coefficient of skewness for EV sentiment is -0.47, while for AV sentiment it is -0.75. It appears that negative news has a greater effect on the sentiment for these two topics than positive news. The AV sentiment scores greater variance and negative skewness may indicate both that the topic is more controversial, and that negative news tends to be more negative (e.g., vehicle crashes, sometimes involving injuries, vs. battery and power grid issues and vehicle recalls). However, these are just hypotheses, and more study is needed to link cause and effect.

Summary

This study examines the sentiment expressed on Twitter regarding EVs and AVs, offering a unique contribution by analyzing a continuous time series of daily sentiment data over a span of more than a year. In contrast to previous research, this study provides a comprehensive examination of sentiment patterns over time.

Several statistical measures were analyzed. Both data sets displayed a weekly seasonal pattern, with a decrease in tweet volume during weekends. However, there was a more pronounced decline in positive and neutral tweets compared to negative tweets. The reason behind this phenomenon is unknown.

The long-term trend in the sentiment score was flat, both for EVs and AVs. Despite speculations in the literature that severe crashes could negatively affect public perception of AVs, the findings from tweets posted over the past 421 days suggest otherwise, at least for Twitter users. The negative effects of crashes or other negative news are significant but short-lived.

Notably, the AV sentiment scores exhibited higher variance and greater negative skewness than the EV sentiment scores. This observation suggests that AVs are a more controversial topic, eliciting stronger opinions, and that negative news surrounding AVs tends to be more extreme, occasionally involving accidents resulting in injuries or fatalities. Both sets of scores showed significant negative skewness, indicating that negative news tends to have a more pronounced impact on Twitter sentiment compared to positive news.

A forthcoming technical report, *The Best of Days, the Worst of Days: Twitter Sentiment Regarding Automated Vehicles,* will examine the impact of major news on AV sentiment. Potential future research avenues include comparing time series Twitter sentiment with sentiment measured through alternative methods, such as surveys. A more detailed analysis could delve into the news stories that contribute the most to sentiment fluctuations and explore the reasons behind the negative skewness in scores. Additionally, investigating the phenomenon of reduced negative tweets during weekends and exploring whether this pattern extends to other topics would be of interest.

References

Date, S. (n.d.). *How To Isolate Trend, Seasonality And Noise From A Time Series*. Retrieved April 27, 2023, from Time Series Analysis, Regression, and Forecasting With Tutorials in Python: https://timeseriesreasoning.com/contents/time-series-decomposition/

Gupta, A., & Sharma, S. (2022). Analysis of Public Perception of Autonomous Vehicles Based on Unlabelled Data from Twitter. In M. Tuba, S. Akashe, & A. Joshi (Eds.), *ICT Infrastructure and Computing. Lecture Notes in Networks and Systems,* (Vol. 520). Springer. doi:10.1007/978-981-19-5331-6_7

Hutto, C., & Gilbert, E. (2014). VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. *Eighth International Conference on Weblogs and Social Media (ICWSM-14).* Ann Arbor, MI. Retrieved April 27, 2023, from https://ojs.aaai.org/index.php/ICWSM/article/view/14550/14399

Hutto, C., & Gilbert, E. (n.d.). *vaderSentiment*. Retrieved April 27, 2023, from GitHub: https://github.com/cjhutto/vaderSentiment

Jefferson, J., & Mcdonald, A. (2019). The Autonomous Vehicle Social Network: Analyzing Tweets After a Recent Tesla Autopilot Crash. *The Human Factors and Ergonomics Society's 2019 International Annual Meeting.* Seattle, WA. Retrieved April 26, 2023, from https://www.researchgate.net/publication/333077240_The_Autonomous_Vehicle_Social_Network_Analyzing_Tweets_After_a_Recent_Tesla_Autopilot_Crash

JustAnotherArchivist. (n.d.). *snscrape*. Retrieved April 27, 2023, from GitHub: https://github.com/JustAnotherArchivist/snscrape

Li, T., Lin, L., Choi, M., & et.al. (2018). YouTube AV 50K: An Annotated Corpus for Comments in Autonomous Vehicles. *2018 International Joint Symposium on Artificial Intelligence and Natural Language Processing (iSAI-NLP).* doi:10.1109/iSAI-NLP.2018.8692799

McGurrin, M. (2023). *Observing the Effect of a Crash on Twitter Sentiment:Early Results from Time Series Data.* McGurrin Consulting, LLC. doi:10.13140/RG.2.2.36194.61123

Michael, M. (2023). *AV & EV Sentiment Indices*. Retrieved April 27, 2023, from AV & EV Sentiment Indices: https://www.mcgurrin.info/sentiment/

Penmetsa, P., & et.al. (2021). Effects of the autonomous vehicle crashes on public perception. *IATSS Research, 45*(4), 485-492. doi:S0386111221000224

Skystance. (2022, August 29). *Sentiment Analysis: Everything You Need to Know Text Analysis and Sentiment Analysis Solutions*. Retrieved April 27, 2023, from Skystance: https://skystance.com/sentiment-analysis-everything-you-need-to-know/

statsmodels, 0.13.5. (n.d.). Retrieved April 27, 2023, from statsmodels: https://www.statsmodels.org/stable/index.html

statsmodels.tsa.seasonal.seasonal_decompose, 0.13.5. (n.d.). Retrieved April 27, 2023, from statsmodels: https://www.statsmodels.org/dev/generated/statsmodels.tsa.seasonal.seasonal_decompose.html

Twitter. (n.d.). *Twitter API*. Retrieved April 27, 2023, from Developer Platform: https://developer.twitter.com/en/docs/twitter-api