

The Best of Days, the Worst of Days: Twitter Sentiment Regarding Automated Vehicles

Michael McGurrin
McGurrin Consulting
mike@mcgurrin.com

Abstract

This report analyzes a corpus of tweets posted on Twitter over a period of over one year regarding automated vehicles (AVs). The sentiment for the tweets was determined and the days before and after the most extreme (both positive and negative) sentiment scores were examined, looking at the extent of the movement in the sentiment scores, if and how rapidly the scores regressed to the mean, and the news stories that immediately preceded or coincided with these extreme scores.

While previous studies have explored Twitter-based sentiment analysis in the context of AVs, including the effects of crashes on sentiment, this study is unique in that it examines the impact of a larger number of news events, both positive and negative, and rather than taking only snapshots of sentiment, it analyzes a continuous time series data of daily Twitter sentiment over a period of 421 days. The data was collected between February 28, 2022, through April 24, 2023. In total, 149,252 tweets were analyzed.

In every case, extreme sentiments expressed on Twitter were traced back to one or more events in the news. There were significantly more negative outliers (10) than positive outliers (2), and the sentiment had negative skewness, but the overall average sentiment was positive. While new events could cause major swings in sentiment, the effects were short-lived in every case, with the maximum observed half-life to reversion to the mean of 6 days, and no long-term shift in average sentiment over the 421 days that were analyzed.

Introduction

This report examines a corpus of tweets posted on Twitter over a period of over one year regarding automated vehicles and examines the positive and negative news that immediately preceded or coincided with the most extreme sentiment scores. The sentiment score for the day is the average score over all relevant tweets, with each negative tweet counting as -1, each neutral tweet as 0, and each positive tweet as +1.

The study examined the days with some of the most positive and most negative sentiment scores regarding AVs, reviewed samples of the actual tweets, the most frequently found words appearing in tweets on those days, and the news stories that were published just before or at the same time.

This paper is one of a series of three related papers. The first paper examines how two AV-related crashes impacted Twitter sentiment toward automated vehicles (McGurrin M. F., 2023). The second

paper examines the seasonality patterns in the sentiment scores regarding both EVs and AVs, as well as the long-term trends, variance, and skewness of the data (McGurrian M. , 2022).

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, especially since there is often disagreement among human judges regarding sentiment categorization, so multiple human reviewers are needed for each sample. 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). 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, also faces many challenges. These include negated words (e.g., “not good” vs. “good”), slang, emojis, and sarcasm. This study used the Valence Aware Dictionary and sEntiment Reasoner (VADER) to determine the sentiment of each tweet (Hutto & Gilbert, VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text, 2014). VADER was specifically designed to analyze short social media posts and can handle many (though not all) of these challenges. 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.

Many researchers have used Twitter as a data source, 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 AVs, the volume still amounts to hundreds of new relevant tweets per day. 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 used custom filtering code to attempt to remove these duplicates.

Related Work

A good number of studies have examined sentiment regarding AVs, including several that examined the impact of news on sentiment scores. Several different types of social media have been used to provide the corpus for study. For example, one study used comments posted regarding YouTube videos about AVs rather than tweets (Li, Lin, Choi, & et.al., 2018). They developed a corpus of over 30,000 comments posted on videos about self-driving vehicles.

¹ This study was conducted using the free tier of Twitter’s API that was in existence for years up to the spring of 2023. 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.

Another study took the unusual approach of analyzing five, rather than two or three, sentiment categories (highly negative, negative, neutral, positive, and highly positive) (TrivyZa, 2021).

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

Several previous studies have examined the impact of crashes on public sentiment regarding AVs. One such 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 the crash, but a significantly larger decrease in positive tweets. The choice of a three-day time frame 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 (Penmetza & et.al., 2021). The latter study examined the twitter sentiment separately for tweets with the key words “Uber,” “Volvo” (the Uber vehicle in the crash in question was a Volvo), “Pedestrian Crash,” “Tesla,” and “Self-driving” or “Autonomous.” The sentiment shifted negative for each word, with the largest shift occurring for the “Self-driving” or “Autonomous” keywords and almost no shift for the “Volvo” keyword. This study also examined the most frequently used words found in tweets posted three days before each crash and compared them with the most frequently found words in tweets posted three days after the crash.

Another study used a combination of Facebook, Twitter, and Reddit data to look at the impact of the same two crashes on public opinion towards AVs, but also looked at the impact of the COVID-19 pandemic (Chen, Zeng, Xu, & Di, 2021). They found that “After the Uber and Tesla crashes, the proportion of people with a negative attitude toward AVs increases by 10% due to their concerns about safety issues regarding AVs.”

A paper by Othman showed that surveys conducted shortly after fatal AV accidents show a negative shift in public attitudes and that there is, at least for the four years between 2016 and 2019, a correlation between the number of fatal accidents involving AVs and the percentage of drivers afraid of AVs in that year’s survey (Othman, 2023). That article also references several other papers that show the negative impact that crash reports have on sentiment regarding AVs.

Jing, Cai, et.al., used over 20,000 comments from Chinese social media platforms (Sina Weibo and Tik Tok) to conduct topic modeling, text network analysis, and sentiment analysis regarding the public perception of AVs before and after crashes (Jing, et al., 2023). They found a spike in the volume of comments the day after an AV crash, but the volume dropped back in just two days. However, they found that the overall negative influence of the crash on overall sentiment lasted for 7 days. The number of negative comments, which had surged, dropped quickly (in about 2 days), but the number of positive comments took longer to return to their average values.

Ding, Korolov, et.al. also used twitter data to examine sentiment regarding AVs (Ding, Korolov, Wallace, & Wang, 2021). Consistent with other studies (including this one), they found that the overall sentiment was positive. They also looked at how sentiment varied between tweets using the terms “Self-driving,” “Autonomous, and “Driverless” and found that while all three had overall positive average sentiment,

“Self-driving” was the least positive and “Driverless” the most positive. They also examined the news stories that appear to be the cause of significant positive and negative spikes in sentiment.

This author has also written a related paper that analyzes different aspects of sentiment analysis of both the twitter corpus that was used in this study and a similar corpus regarding electric vehicles. The report analyzes the seasonality of the data, along with long-term trends (McGurrian M. F., 2023).

Approach

A combination of the [Twitter API v2](#) (Twitter, n.d.) and [snsrape](#) (JustAnotherArchivist, n.d.), a social media scraping tool, was utilized to extract tweets discussing AVs based on finding one or more specific keywords in the body of the tweet. 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

The dataset was 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. In total, after filtering, 149,252 tweets remained. These remaining tweets were indexed by their posting dates, and the sentiment expressed in each tweet was analyzed to determine its polarity (positive, neutral, or negative).

Various computer-based approaches exist for sentiment assessment, including machine learning algorithms and neural networks. In this project, VADER, an open-source software package, was used to calculate the AV sentiment indices (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 fast processing.

Two measures are 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 only negative tweets were posted that day) to +1 (indicating only positive tweets were posted). The positive/negative ratio considers only positive and negative tweets, excluding neutral ones, and represents the ratio of positive tweets to negative tweets. Additional information can be found on the [AV Sentiment Index website](#) (McGurrian M. , AV Sentiment Index, n.d.).

Analysis

The mean sentiment score for AVs during the study period was 0.20 (where a score -1 is entirely negative tweets, a score of 0 is equally positive and negative, and a score of +1 is entirely positive tweets). The standard deviation was 0.09. **Error! Reference source not found.** shows a time series plot of the sentiment scores. The red lines depict two standard deviations above and below the means, and the red circles indicate the outlier days when the sentiment score differed from the mean by two or more standard deviations.

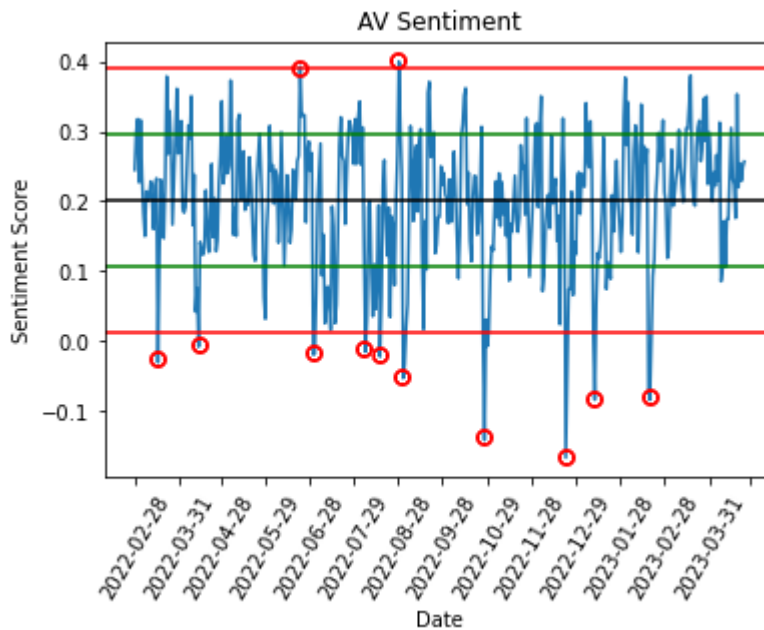


Figure 1 A time series plot of the sentiment scores. The black line shows the mean value, the green lines show one standard deviation from the mean, and the red lines show two standard deviations from the mean.

This analysis examines positive outlier days when the sentiment scores two standard deviations above the mean and the negative outlier days when the scores were two standard deviations below the mean. For each of these, samples of tweets posted on that day are provided, along with relevant news that came out at the time of or immediately preceding these extreme sentiment days.

The study also examines “hot words” and “not words.” “Hot words” are defined as the subset of the one hundred words appearing most frequently in the set of tweets for that day that did NOT appear in the top one hundred words for tweets posted seven days previously. Hot words include both single-word and two-word combinations. “Not words” are the opposite. They are words that were among the one hundred most frequently found words seven days previously, but do not appear in the one hundred most frequently found words on the date in question. Figure 2 is a word cloud that shows the hot words

for June 8, 2023. These hot words and not words can be useful in identifying the possible reasons for sentiment shifts.

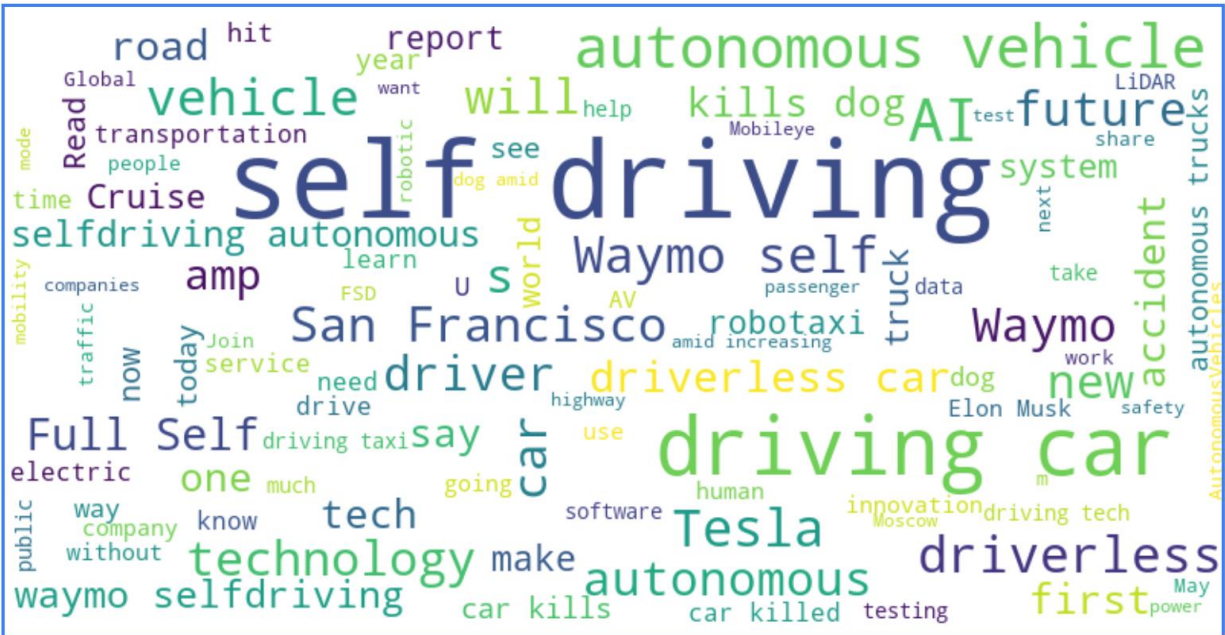


Figure 2 The "hot words" for June 8, 2023

Finally, we examine how quickly the sentiment score reverted to the mean by examining the half-life to return. The half-life is the number of days until the sentiment score returned to halfway between the peak or trough and the mean.

Sentiment Peaks

There are two peaks that are two standard deviations above the mean, occurring on June 22, 2022 and August 29, 2022. Each is discussed below.

June 22, 2022 Peak

On 22nd, the sentiment score peaked at 0.39, two standard deviations above the mean. A review of the tweets and news items does not identify any single stand-out event or news item that drove the spike in positive sentiment. Among the top one hundred words appearing on June 22nd but not the week before were "companies," "industry," and "Researchers." The tweets that included "companies" often referred to the Pennsylvania House passing legislation enabling automated vehicle operations:

- *Self-driving cars use a different number of sensors, radars, cameras, and AI to travel to destinations without needing a human driver. Many auto companies have started manufacturing self-driving cars: including @Tesla, @Audi, @BMW, @Ford, etc. #selfdrivingcars #AI*
- *The legislation passed on Monday, June 20, and would allow #AutonomousVehicle companies to deploy and test #driverless cars and trucks on city streets. The bill will now move to the Senate for consideration.*

- *The Pennsylvania House on Monday passed a bill that would allow autonomous vehicle companies to test and deploy driverless vehicles without a safety driver.*
- *PA needs a clear regulatory environment that allows companies to test & deploy driverless technology on public roads. This policy change is needed for the existing SWPA autonomous vehicle industry to continue to flourish and grow. Read more [↓](#)*

Tweets that included the word “industry” included:

- *Self-driving cars use a different number of sensors, radars, cameras, and AI to travel to destinations without needing a human driver. Many auto companies have started manufacturing self-driving cars: including @Tesla, @Audi, @BMW, @Ford, etc. #selfdrivingcars #AI*
- *A: “It remains to be seen if self-driving trucks will revolutionize the shipping industry in the near future, but many are confident the change to autonomous shipping is inevitable as a slow process.”*
- *How free-wheeling Texas became the self-driving trucking industry's promised land*

There were also several tweets that included the word “industry” that also referred to the action taken by the Pennsylvania House.

The hot word “Researchers” was also associated with positive tweets, such as:

- *“Researchers have developed a way to help autonomous vehicles create 'memories' of previous experiences and use them in future navigation, especially during adverse weather conditions when the car cannot safely rely on its sensors.” @ScienceDaily:*
- *Self-driving vehicles with memory? Researchers have found a way - (Interesting Engineering)*
- *Technology helps self-driving cars learn from own 'memories'. Date: June (22, 2022; Source: Cornell University; Summary: Researchers have developed a)*
- *Andy Vermaut shares: Technology helps self-driving cars learn from own 'memories': Researchers have developed a way to help autonomous vehicles create 'memories' of previous experiences and...*

Practically all the tweets containing the word popular word “Researchers” referred positively to the same news item, about particular research coming out of Cornell University.

It is also interesting in this case to look at the top one hundred words from 7 days before the peak that did *not* appear among the top one hundred words for the peak. These included “crash,” “San Francisco,” “NHTSA,” “Car Crashes,” “Elon Musk,” and “Tesla.” On that date, there were numerous tweets about a New York Times Article, *Tesla Autopilot and Other Driver-Assist Systems Linked to Hundreds of Crashes*. This news had mostly fallen out of the twitter feed by June 22nd.

While the average sentiment on June 15th was a bit below average, it does appear that the sentiment peak on June 22nd was caused *both* by a higher number of positive tweets and by a lower number of negative tweets. The average number of positive tweets received on any given day during the study period was 143, while 183 positive tweets were received on June 22nd. The average number of negative

tweets received during the study period was seventy-six, however only forty-five negative tweets were received on June 22nd.

Figure 3 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the spike in positive sentiment lasted for just a few days, with a half-life of four days (in fact, it dropped to below average on the fourth day).




Figure 3 The sentiment scores for the 20 days centered around the June 22nd Peak. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

August 29, 2022 Peak

On August 29th, the sentiment score peaked at 0.399, over two standard deviations above the mean and the highest sentiment recorded during the study period. As it was with the June 22nd peak, this peak appears to have been driven both by an unusually large number of positive tweets and an unusually smaller number of negative tweets. There were 254 positive tweets, compared with an average of 143, while there were only fifty negative tweets, compared with an average of seventy-six. Words like “raising price,” “price increase,” “price hike,” and “accident” were among the one hundred most frequent one or two-word phrases that had appeared the week prior but did not appear among the top words on August 29th.

However, in this case, there were two positive news reports that clearly drove the increase in positive tweets. Whereas the *absence* of references to Tesla and Elon Musk helped to drive the sentiment higher on June 22nd, both featured heavily in positive tweets on August 29th, with hot words including “Musk tells,” “Tesla aims,” “Musk aims,” “Musk said,” and “Musk wants.” This was driven by announcements by Elon Musk on August 29th that Tesla had a goal of widely releasing self-driving by the year’s end. The second item was the announcement by Waymo, also made on August 29th, that its robotaxi service would expand operations to downtown Phoenix.

Tweets relating to these two subjects included:

- *Elon Musk Says Focused On Wide Release Of Self-Driving Teslas By Year-End*
- *Elon Musk Wants to Launch Tesla Full Self-Driving by 2022 End: Report*

- *Tesla CEO Elon Musk aims to get the automaker's self-driving technology ready by year-end and hopes it could be in wide release in the U.S. and possibly in Europe, depending on regulatory approval.*
- *"The two technologies I am focused on, trying to ideally get done before the end of the year, are getting our Starship into orbit ... and then having #Tesla cars to be able to do self-driving." #ElonMusk.*
- *Elon Musk says Tesla aims to roll out self-driving cars by year end: Musk has previously stated that autonomous vehicles will be safer than cars with human drivers.*
- *Waymo opens up driverless robotaxi service in downtown Phoenix to vetted passengers | TechCrunch*
- *.@Waymo's #driverless vehicles are picking up passengers in downtown #Phoenix*
- *TSLA @elonmusk Waymo opens up driverless robotaxi service in downtown Phoenix to vetted passengers - Yahoo Finance*
- *"Waymo announced earlier this week its fleet of self-driving semitractor-trailers would increase to 60.... In addition, the company will begin operating on public roads in Texas and Arizona."*
- *Check this out! Waymo's driverless vehicles are picking up passengers in downtown Phoenix*

Figure 4 shows the sentiment scores for the 10 days prior to and after the peak. As can be seen in the figure, the significant increase in positive sentiment lasted for just a few days, with a half-life of two days. One of the largest troughs of negative sentiment occurred on Sept. 1, just three days later. This trough is discussed in the next section.




Figure 4 The sentiment scores for the 20 days centered around the August 29th Peak. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

Sentiment Troughs

While only two days had positive sentiment peaks that were at least two standard deviations above the mean during the study period, there were ten negative troughs at least two standard deviations below

the mean, with two falling three standard deviations below the mean. In addition, whereas neither peak remained two standard deviations above the mean for more than a day, several of the troughs last for two to four days. Overall, the sentiment data is negatively skewed. The Fisher-Pearson coefficient of skewness for AV sentiment is -0.75 (McGurrin M. F., 2023).

The December 22 – December 23 trough will first be examined in some detail, then each remaining trough will be discussed in chronological order.

December 22 – December 23, 2022 Trough

On December 22, the sentiment score dropped to its lowest point during the entire study period. Back on Thanksgiving in November of 2022, there was an eight-car pileup on the San Francisco Bay Bridge that resulted in several injuries. The cause of that pileup was not initially known, and it did not cause a jump in sentiment regarding AVs. However, beginning on December 21st, multiple media sources reported that the driver of the vehicle that started the pileup reported that the cause was his Tesla operating in “full self-driving” mode (Shakir, 2022). This news was reflected in tweets beginning the next day, with the number of negative tweets on automated vehicles went up significantly, as shown in **Error! Reference source not found.** The number of negative tweets remained high for several days afterwards.




Figure 5 Number of tweets of each sentiment for the 21 days centered on December 22nd, 2022. The dramatic rise in the number of negative tweets is clearly visible.

Hot word for tweets about AVs for that day included “told police,” “crash told,” “triggered,” “behind eight,” and “Bay Bridge,” indicating that the news of the pileup being caused by self-driving software was the probable cause for the increasingly negative sentiment. Tweets posted that day included:

- *More from the Boy Genius: Tesla 'full self-driving' triggered an eight-car crash, a driver tells police | CNN Business*
- *Tesla Full Self-Driving Linked to Eight-Car Crash*
- *#Tesla 'full self-driving' triggered an eight-car crash, a driver tells police | CNN Business*
- *A self-driving Tesla caused an 8 car accident in the San Francisco area. I think there's still a little work to be done.*

Figure 6 shows the sentiment scores for the 21 days centered around the trough. As can be seen in the figure, the half-life for the extreme drop was 4 days. Two of those days were two standard deviations below the mean.




Figure 6 The sentiment scores for the 20 days centered around the December 22 – December 23 trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

Interestingly, another crash involving a Tesla occurred the first week of January 2023. In that crash, a Tesla ran off a cliff (Griggs & Karimi, 2023). However, it was quickly reported that the crash did not involve the self-driving features, and this crash does not appear to have affected the sentiment on Twitter regarding automated vehicles. Neither “crash,” “cliff,” nor “Tesla” show up among the hot words for January 4th or 5th, indicating that this event did not appear in any significant number of tweets about AVs. The sentiment scores for that time does not show a drop off, as shown in Figure 7, and the number of negative tweets did not rise, as shown in Figure 8. Note that the sentiment trough occurring on January 10th is discussed below in its own section.




Figure 7 The sentiment scores for December 25, 2022 through January 14, 2023. The black line marks the mean score for that particular time period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.




Figure 8 Number of tweets of each sentiment for the 21 days centered on January 4th.

March 16, 2022 Trough

On March 16h, the sentiment had a very sharp, but short, drop. A review of the tweets regarding AVs posted on that day reveal that a very large number of negative tweets all referred to the same news story that came out that day, that Tesla had fired an employee who had posted videos of full self-driving failures on YouTube, e.g., *Tesla fires employee who posted YouTube videos of Full Self-Driving accident* (Brodkin, 2022). Tweets included:

- *Tesla employee was fired after sharing video of a crash on Full Self-Driving Beta – Electrek*
- *Tesla employee fired after posting YouTube video of self-driving Model 3 running off a road*
- *Tesla's "Full Self Driving" cars can't drive themselves. Tesla knows this. Employees who mention it will be fired.*

A look at the “hot words” appearing the most in tweets that day but not appearing frequently 7 days earlier, one sees “posted YouTube,” “Tesla employee,” “Tesla fired,” “fires employee,” and “YouTube video,” confirming this source for the negative sentiments. Figure 9 shows the number of tweets of each sentiment for the 21 days centered on the trough. It shows that the number of positive and neutral tweets remained about the same on March 16th, but the number of negative tweets spike.

Figure 10 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the extreme drop in sentiment lasted for just one day.




Figure 9 Number of tweets of each sentiment for the 21 days centered on the March 16th trough in sentiment.




Figure 10 The sentiment scores for the 20 days centered around the March 16th trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

April 13, 2022 Trough

The April 13th trough was also driven by a single, negative event in the news. Various news sources reported, beginning on April 11th, that the San Francisco police had pulled over a driverless car operated

by Cruise for driving without headlights at night, and that after initially stopping, the vehicle then proceeded ahead (Mendoza, 2022). The most frequently found hot words found that day included “police,” “car stopped,” “puzzles cops,” “San Francisco,” “Cruise,” and “traffic stop.” And the twitter feed was full of references to this event, including:

- *Awkward Video Shows San Francisco Police Pulling Over Driverless Car*
- *‘Ain’t Nobody In It’: Police Pull Over Driverless Car During San Francisco Traffic Stop*
- *San Francisco Cops Try To Pull Over Driverless Cruise Taxi, Hilarity Ensues! #cops #cruise #driverless #ensues #entertainment #francisco #hilarity #over #pull #san #taxi #try* (note that while tweeting on the same event, this tweet expresses positive sentiment, with words such as “hilarity” and “entertainment)
- *Driverless car appears to flee the scene after being pulled over by cops*

However, it is interesting to note that the number of negative tweets actually shot up on April 11th, peaked on April 12th, and then was dropping on April 13th. The sentiment score was lower on the 13th than the previous two days because the number of positive tweets shot up on the 11th and peaked on the 12th, as shown in Figure 11. The number of neutral tweets also peaked during that same period. There does not appear to be any single event that explains this jump, however there was a convergence of tweets reporting positive news, e.g.,

- *IEEE Standard Will Make Autonomous Vehicles Safer - IEEE Spectrum and NASA’s Self-Driving Perseverance Mars Rover Is Breaking Records 👍*
- *Now he’s got self-driving taxi’s. I so want to try that next time I’m in SF. 😊*

These were combined with a large number of positive and neutral tweets regarding the event in San Francisco, such as:

- *Twitter hilariously reacts to police pulling over driverless car.*

The incident was clearly not a positive news story for AVs, however there were no fatalities, injuries, or property damage, and many tweeted positively about the humor to be found in the event, while others posted neutral tweets with just a simple statement of facts.




Figure 11 Number of tweets of each sentiment for the 21 days centered on the April 13th trough in sentiment.

Figure 12 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the drop in sentiment was just one day.




Figure 12 The sentiment scores for the 20 days centered around the April 13th trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

July 1 – July 2, 2022 Trough

The next major drop in sentiment was also due to Cruise vehicles. A few days before, up to 20 Cruise vehicles converged on the same intersection, made the same turn, and then inexplicably stopped. Cruise staff arrived on the scene after about 20 minutes, but it took about 2 hours to clear out the incident (Schneider, 2022). The incident made the newspapers and online media a few days later, on July 1st, and this was reflected in negative tweets, including:

- *A Fleet of Self-Driving Cars Failed Yet Again*
- *Cruise's Confused Self-Driving Cars Clog San Francisco Street*
- *"Fleet of Cruise driverless cars blocked traffic for hours Tuesday night - San Francisco Examiner"*

- *Driverless Robotaxi Fleet Paralyzed for Hours in San Francisco. Unionizing?* 🤖

Figure 13 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 3 days. Two of those days were two standard deviations below the mean.




Figure 13 The sentiment scores for the 20 days centered around the July 1 – July 2nd trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

The number of negative tweets actually dropped off sharply after just a single day, however, for an unknown reason, the number of neutral and positive tweets also dropped sharply on July 2nd, resulting in a second day where the sentiment score was over two standard deviations below the mean. This is shown in Figure 14.

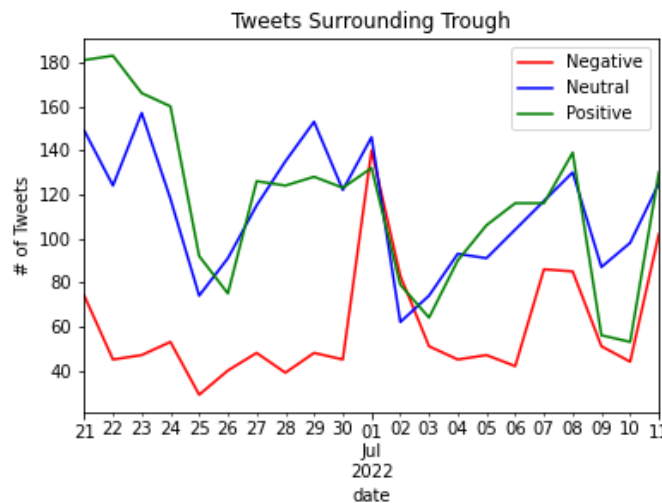


Figure 14 Number of tweets of each sentiment for the 21 days centered on the July1 – July 2 trough in sentiment.

August 6, 2022 Trough

The August 6 trough was not driven by a large spike in negative tweets, but rather by a convergence of a rise in negative tweets and a decline in neutral and positive tweets. In fact, as shown in Figure 15, the number of negative tweets was actually higher on August 10 and 11, but the number of neutral and positive tweets were *far* higher on those days, so the overall sentiment score, while low on those days, did not go below two standard deviations from the mean.




Figure 15 Number of tweets of each sentiment for the 21 days centered on the August 6th trough in sentiment. Note the higher number of negative tweets on August 10th and 11th, as well as the large spike in positive tweets on August 8th and 9th.

The increase in negative tweets on the 6th was driven by news that the California Department of Motor Vehicles (DMV) had filed complaints accusing Tesla of falsely advertising its Autopilot and Full Self-Driving features (Shepardson & Stempel, California regulator claims Tesla falsely advertised Autopilot, Full Self-Driving, 2022). Once again, the news was reflected in negative tweets, such as:

- *\$TSLA NEW ARTICLE : Tesla accused of false advertising of Autopilot, Full Self-Driving: 'Not based on facts'*
- *Tesla Inc (TSLA.O), the electric car company run by billionaire Elon Musk, has been accused by a #California state transportation regulator of falsely advertising its #Autopilot and Full #SelfDriving features as providing #autonomo...*
- *California DMV accuses Tesla of deceptive practices in marketing Autopilot and Full Self-Driving options News Link:...*
- *Regulators have raised questions about Tesla's autopilot and self-driving features, which are a big part of the Texas-based company's lofty valuation.*

While negative tweets continued to be posted on August 8th, e.g., *Teslas full self driving claims are one of the greatest scams in history*, there was a very large spike in positive and neutral tweets reflecting announcements relating to Baidu in China, such as:

- *China made a significant legislative move in the area of driverless taxis, allowing Baidu to run ride-hailing services in two cities without a driver or someone monitoring safety inside the car.*

- *#China's Baidu gets greenlight for country's 1st fully driverless commercial robotaxi service in 2 megacities*
- *Baidu launches China's first fully driverless taxis in Chongqing, Wuhan*
- *China took a notable step in driverless taxis, with two cities approving Baidu to operate ride-hailing services without a driver or a person overseeing safety in the vehicle*

This was also reflected in the “hot words” for August 8th, which included “robotaxi license,” “Baidu wins,” “Baidu gets,” “two Chinese,” and “Baidu,”

Figure 16 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 2 days.




Figure 16 The sentiment scores for the 20 days centered around the August 6th trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

August 16, 2022 Trough

The number of both positive and negative tweets increased around August 16th, but the number of negative tweets increased more, as shown in Figure 17.




Figure 17 Number of tweets of each sentiment for the 21 days centered on the August 16th trough in sentiment. Note that tweets of all three sentiments increased, but negative tweets increased more than positive tweets.

The positive tweets included a significant number about the news that Lyft was about to deploy a self-driving car service in Las Vegas, as well as the announcements of various partnerships in the AV field. On the negative side, there were still a number of tweets regarding the California DMV accusations regarding false advertising that had driven the August 6th trough, but the primary driver for the increase in negative tweets were news stories from Germany concerning a fatal crash involving a BMW (Reuters, 2022). The original police report had stated that the car was self-driving. BMW stated that the car had no self-driving capability, however it did have a Level 2² driving assistance system that involved steering as well as acceleration and braking. Although BMW stated that the car was not “self-driving,” many tweets included this term. Tweets regarding the fatal crash included:

- *#Electric #SelfDriving' #BMW test car veers into oncoming traffic leaving one dead and nine injured in mass pile-up in #Germany*
- *One dead after 'self-driving' BMW test car veers into traffic*
- *TSLA @elonmusk BMW refutes German police report that SUV in fatal crash was self-driving - WJAC Johnstown*
- *#BMW just found out self driving cars are not easy to get right. Their stocks just dropped*

Figure 18 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 2 days.

² SAE International defines six levels of driving automation systems, including level 0, no driving automation (SAE International, 2021). Level 2, “Partial Driving Automation” systems provide both steering and brake / acceleration support to the driver, however these are driving assistance systems, and the driver must constantly supervise such systems.




Figure 18 The sentiment scores for the 20 days centered around the August 16th trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

September 1 – September 2, 2022 Trough

The confluence of two separate negative news stories were picked up on Twitter, generating this trough in sentiment. The first news item was the California legislature passing a bill that made it illegal to advertise a vehicle as “self-driving” that did not meet the state’s definition of “self-driving” (Mitchell, 2022). This was already a violation of a DMV regulation, but the state legislature had run out of patience with the DMV’s lack of enforcement. Examples of tweets about this item included:

- *California lawmakers just passed a bill that could force Tesla to stop calling its beta software 'Full Self Driving' in the state — now, it's up to Gov. Newsom to sign it into law*
- *TSLA @elonmusk California Targets Tesla 'Full Self-Driving' With Newly Passed Bill - The Drive*
- *TSLA @elonmusk California lawmakers just passed a bill that could force Tesla to stop calling its beta software 'Full Self Driving' in the state — now, it's up to Gov. Newsom to sign it into law - Business Insider Africa*
- *The California Department of Motor Vehicles has rules banning the advertisement of cars as self-driving when they are not, but it has not enforced them. Now, the Legislature is making it a state law.*

The second news item concerned the announcement by General Motors (GM)’s Cruise startup that it had recalled and updated the software in 80 self-driving vehicles following their investigation into a crash in June 2022 that had injured two people (Shepardson, GM startup Cruise recalls and revises self-driving software after crash, 2022). The recalled software, in certain circumstances, incorrectly predicted the path of an oncoming vehicle, causing the Cruise vehicle to hard brake to avoid what it perceived as a potential collision. Examples of tweets about this item included:

- *GM's Cruise recalls and updates self-driving software following crash*
- *An unprotected left turn intersection pushes a self-driving car to stop, but not for humans.*
- *GM driverless car unit recalls vehicles after accident –*
- *Cruise updates software for autonomous vehicles after crash*

Figure 19 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 4 days. Two of those days were two standard deviations below the mean.

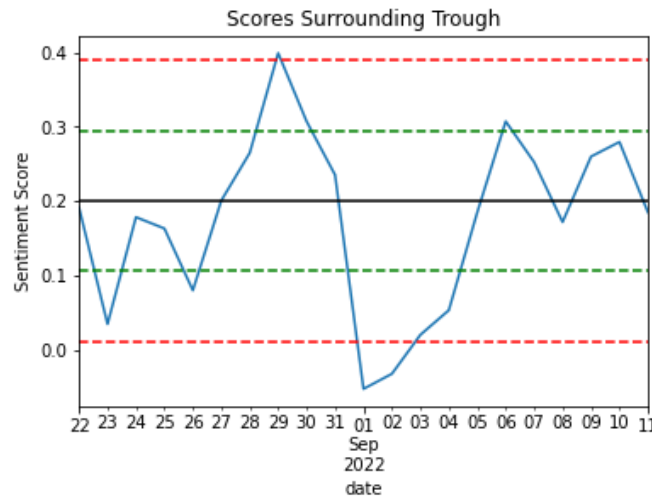


Figure 19 The sentiment scores for the 20 days centered around the September 1 – 2nd trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

October 26 – October 29, 2022 Trough

There were multiple news stories driving Twitter sentiment, both positive and negative, but the convergence of two negative stories dominated, driving the sentiment low and keeping it there for several days. The most frequently found word in tweets on October 26th that was not frequently found on October 19th was “IPO.” Mobileye had a successful IPO, and there were many positive tweets relaying the news (Capoot, 2022), including:

- Intel \$INTC has official spun out its self-driving technology unit Mobileye \$MBLY. It's one of the biggest U.S. IPOs of the year, and investors ate it up, with shares popping over 30% on debut.
- Mobileye Global, the self-driving technology firm spun off by Intel, gains in its trading debut after raising \$861 million in one of the few US IPOs to exceed its goals this year
- Shares of Intel's self-driving car business, Mobileye, soars following an IPO that valued the unit at roughly \$17 billion

However, at the same time, Argo AI, a joint self-driving venture involving Ford and Volkswagen, announced that it was shutting down (Korosec, 2022). In addition, news broke that Tesla was facing a criminal investigation over their claims regarding their “Autopilot” and “Full Self-Driving” systems (Etherington, 2022). Tweets about the Argo AI shutdown included:

- Ford Posts Loss as It Takes \$2.7 Billion Charge on Argo Driverless-Venture
- Ford reins in hopes for self-driving cars as Argo AI shuts down
- Argo AI, Ford's Self-Driving Venture With Volkswagen, Is Shutting Down

Tweets about the criminal investigation included:

- Tesla’s self-driving claims are reportedly under criminal investigation
- Tesla is reportedly target of DoJ criminal probe of self-driving claims
- Tesla Could Face Criminal Charges Under Reported Self-Driving Investigation

As show in Figure 20, the total number of tweets increased during this time, and the number of tweets of each sentiment increased. However, the increase in the number of negative tweets significantly exceeded the increase in the number of positive or neutral tweets.

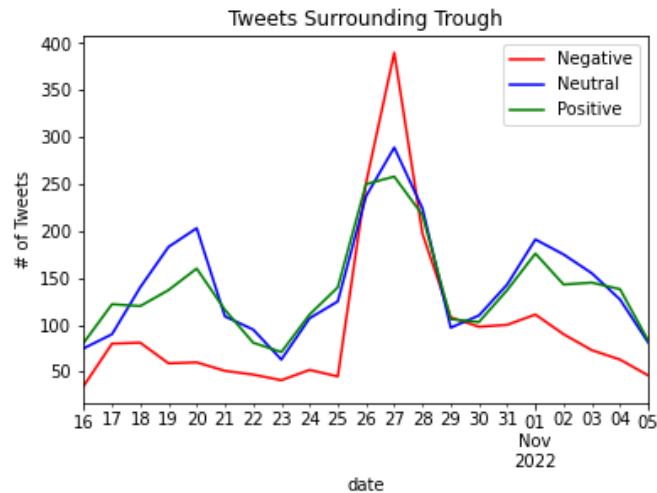


Figure 20 Number of tweets of each sentiment for the 21 days centered on October 26th. Note that tweets of all three sentiments increased, but negative tweets increased more than positive tweets. The number of tweets of each sentiment increased, however the increase in negative tweets significantly exceeded the increase in positive or neutral tweets.

Figure 21 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 6 days, the longest observed. Three of those days were two standard deviations below the mean.




Figure 21 The sentiment scores for the 20 days centered around the October 26 – October 29th trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

January 10 – January 11, 2023 Trough

The second week of January saw a mix of news, both positive and negative, and this was reflected in tweets. Positive tweets included many reflecting Mercedes’ introduction of Level 3 self-driving into the U.S. and, more unusually, the announcement of a “self-driving” baby stroller at the Consumer Electronics Show (CES). However, negative news dominated, including a potential federal investigation after Elon Musk indicated (in a tweet) that Tesla might allow some owners to disable alerts and video from an Austin cyclist showing a self-driving car veering into the bike lane. On January 11th, the twitter feeds also picked up on the release of the surveillance footage from the Thanksgiving Day crash (Klippenstein, 2023), further driving down the sentiment scores. Figure 22 shows the number of negative tweets, having increased significantly on January 10th, increased further on January 11th.




Figure 22 Number of tweets of each sentiment for the 21 days centered on January 10th.

Figure 23 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 3 days. Two of those days were two standard deviations below the mean, with the second day showing a lower sentiment score than the first.




Figure 23 The sentiment scores for the 20 days centered around the January 10 – January 11 trough. The black line marks the mean score for the overall study period, the green lines mark +/- one standard deviation, and the red lines mark +/- two standard deviations.

February 17 – February 19, 2023 Trough

One positive piece of news had a large presence in tweets posted on February 17th, concerning a successful trial of self-driving cars, with an engineer reporting that traffic lights could be obsolete within 20 years (Sky News, 2023). However, this was outweighed by an even larger rise in neutral and negative tweets, driven by news that Tesla was recalling over 360,000 vehicles to fix flaws in the behavior of its “Full Self-Driving” feature (Associated Press, 2023). These increased in the volume of tweets of all sentiments can be seen in Figure 24.




Figure 24 Number of tweets of each sentiment for the 21 days centered on February 17th. Note that tweets of all three sentiments increased sharply, but negative tweets increased more than positive tweets.

Tweets that reflected this news included:

- *Tesla's Recall Targets a 'Fundamental' Flaw - Self driving.* 🚗
- *If 'self-driving' Teslas are defective, why are regulators letting them stay on the road? #LAAngels #TheHaloWay #LTBU #Angels #LAA*
- *"Tesla "Recalls" Over 300,000 Vehicles Over Self Driving Danger" from @Timcast*
- *An @NHTSAGov report found @Tesla "Full Self-Driving" systems can malfunction in and around intersections and can fail to follow the speed limit, a sharp contrast to @ElonMusk's suggestions that the system can operate safer than humans.*

Figure 25 shows the sentiment scores for the ten days prior to and after the peak. As can be seen in the figure, the half-life for the extreme drop was 3 days. Three of those days were two standard deviations below the mean, with the second day showing the lowest sentiment score.




Figure 25 The sentiment scores for the 20 days centered around the February 17 – February 19 trough. The black line marks the mean score for the overall study period, the green lines mark \pm one standard deviation, and the red lines mark \pm two standard deviations.

Summary

Each of the outlier days in terms of sentiment can be related to one or more specific news events that occurred at the same time, as summarized in Table 1. The two positive peaks appear to have been driven by multiple positive news stories coming out at the same time and were also accompanied by a decrease in the number of negative tweets (the reason for the latter is unknown). In contrast, single negative news stories appear adequate to cause negative outliers.

There were ten negative outliers during the period studied and only two positive outliers. And while the overall average sentiment was positive, with a mean of 0.20, the sentiment scores had a negative skew, with a Fisher-Person coefficient of skewness of -0.75 (McGurrin M. F., 2023).

The effect on the twitter stream was short-lived in every case. The longest half-life was six days, however the rest of the outliers had half-lives between one and four days. There has been discussion regarding the long-term effects of the reports of crashes on the sentiment regarding AVs. Two studies had each taken snapshots at two points in time and found that over time, the sentiment regarding AVs

had grown more negative, however Abdelgawad and Othman showed that in both cases, the later snapshots were taken shortly after news had appeared of the first fatal accident involved an AV (Othman, 2023). Therefore, the result could have been caused by a short-lived shift in sentiment, rather than a long-term trend. This current study did not rely on two snapshots, but rather continually monitored the sentiment daily for 421 days. This study found no long-term change in sentiment, with the best fit slope for the sentiment score being essentially zero (2.2×10^{-5}) (McGurrin M. F., 2023).

Table 1 A summary of each outlier, listing the date(s), the causes, and the half-life for the duration of the extreme sentiment scores.

Outlier Date(s)	Cause(s)	Half-Life (days)
Positive Outliers		
June 22, 2022	Multiple, including Pennsylvania legislative action and successful research from Cornell	4
August 29, 2022	Announcement by Elon Musk that Tesla had a goal of widely releasing self-driving by the year’s end and the announcement by Waymo that its robotaxi service would expand operations to downtown Phoenix.	2
Negative Outliers		
March 16, 2022	News that Tesla had fired an employee who had posted videos of full self-driving failures on YouTube	1
April 13, 2022	News stories reporting that the San Francisco police had pulled over a driverless car operated by Cruise for driving without headlights at night, and that after initially stopping, the vehicle then proceeded ahead.	1
July 1, 2022 – July 2, 2022	News reports of an incident, which had occurred a few days before, that up to 20 Cruise vehicles converged on the same intersection, made the same turn, and then inexplicably stopped. Cruise staff arrived on the scene after about 20 minutes, but it took about 2 hours to clear out the incident.	3
August 6, 2022	News that the California Department of Motor Vehicles (DMV) had filed complaints accusing Tesla of falsely advertising its Autopilot and Full Self-Driving features. This negative news outweighed several positive stories that were also tweeted about.	2
August 16, 2022	News stories from Germany concerning a fatal crash involving a BMW (Reuters, 2022). The original police report had stated that the car was self-driving. BMW stated that the car had no self-driving capability.	2
September 1, 2022 – September 2, 2022	Two negative news stories: the California legislature passed a bill that made it illegal to advertise a vehicle as “self-driving” that did not meet the state’s definition of “self-driving” and the announcement by GM’s Cruise startup that it had recalled and updated the software in 80 self-driving vehicles following their investigation into a crash in June 2022 that had injured two people.	4

October 26, 2022 – October 29, 2022	Two negative news stories: Argo AI, a joint self-driving venture involving Ford and Volkswagen, announced that it was shutting down and news that Tesla was facing a criminal investigation over their claims regarding their “Autopilot” and “Full Self-Driving” systems.	6
December 22, 2022 – December 23, 2022	Reports that the Tesla that had caused an 8-car pileup with injuries the previous month had been in self-driving mode at the time.	4
January 10, 2023 – January 11, 2023	Three items: dominated, announcement of a potential federal investigation after Elon Musk indicated that Tesla might allow some owners to disable alerts; video from an Austin cyclist showing a self-driving car veering into the bike lane; and on January 11 th , the release of the surveillance footage from the Thanksgiving Day crash.	3
February 17, 2023 – February 19, 2023	Recall announcement related to Tesla’s Full Self-Driving feature	3

A Note on the Data

The terms of use for the Twitter API prohibit the author from publishing the corpus of twitter feeds. However, the corpus, along with the calculated daily sentiment-scored tweet count and sentiment scores and the lists of hot words and not word can be made available privately to other researchers. A similar set of data, covering the same time period, is available for tweets regarding electric vehicles.

References

- Alhari, M. I., Lubis, M., & Pratiwi, O. N. (n.d.). Sentiment Analysis of The Public Perspective Electric Cars in Indonesia Using. *2022 International Conference of Science and Information Technology in Smart Administration (ICSINTESA)*. doi:10.1109/ICSINTESA56431.2022.10041604
- Associated Press. (2023, February 16). *Tesla recalls nearly 363,000 cars with 'Full Self-Driving' to fix flaws in behavior*. Retrieved April 21, 2024, from NPR: <https://www.npr.org/2023/02/16/1157521492/tesla-full-self-driving-recall-fsd>
- Brodkin, J. (2022, March 16). *Tesla fires employee who posted YouTube videos of Full Self-Driving accident*. Retrieved April 16, 2024, from Ars Technica: <https://arstechnica.com/tech-policy/2022/03/tesla-fires-employee-who-posted-youtube-videos-of-full-self-driving-accident/>
- Capoot, A. (2022, October 26). *Mobileye pops more than 37% in market debut after spinning out of Intel*. Retrieved April 19, 2024, from CNBC: <https://www.cnbc.com/2022/10/26/mobileye-pops-more-than-30percent-in-ipo-after-spinning-out-of-intel.html>
- Chen, X., Zeng, H., Xu, H., & Di, X. (2021). Sentiment Analysis of Autonomous Vehicles After. *2021 IEEE Intelligent Transportation Systems Conference (ITSC)* (pp. 1211-1216). Indianapolis: IEEE. doi:10.1109
- Costello, F. J., & Lee, K. C. (2020, February). Exploring the Sentiment Analysis of Electric Vehicles Social Media. *Journal of Digital Convergence*, 18(2), 249-259. doi:10.14400/JDC.2020.18.2.249

- Ding, Y., Korolov, R., Wallace, W., & Wang, X. (2021, October). How are sentiments on autonomous vehicles influenced? An analysis using Twitter Feeds. *Transportation Research Part C*, 131(103356). doi:10.1016/j.trc.2021.103356
- Etherington, D. (2022, October 26). *Tesla said to face criminal investigation by the Department of Justice over self-driving claims*. Retrieved April 19, 2024, from Tech Crunch: <https://techcrunch.com/2022/10/26/tesla-criminal-investigation-doj-autopilot-self-driving/>
- Griggs, B., & Karimi, F. (2023, January 4). *Tesla plunges 250 feet off a California cliff, all 4 occupants survive*. Retrieved April 20, 2024, from CNN: <https://www.cnn.com/2023/01/03/us/tesla-crash-cliff-california-cec/index.html>
- 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
- Jena, R. (2020, October). An empirical case study on Indian consumers' sentiment towards electricvehicles: A big data analytics approach. *Industrial Marketing Management*, 90, 605-616. Retrieved April 28, 2023, from https://www.academia.edu/82678374/An_empirical_case_study_on_Indian_consumers_sentiment_towards_electric_vehicles_A_big_data_analytics_approach
- Jing, P., Cai, Y., Wang, B., Wang, B., Huang, J., Jiang, C., & Yang, C. (2023). Listen to social media users: Mining Chinese public perception of autonomous vehicles after crashes. *Transportation Research Part F*, 93, 248-265. doi:10.1016/j.trf.2023.01.018
- JustAnotherArchivist. (n.d.). *snsrape*. Retrieved April 27, 2023, from GitHub: <https://github.com/JustAnotherArchivist/snsrape>
- Klippenstein, K. (2023, January 10). *EXCLUSIVE: SURVEILLANCE FOOTAGE OF TESLA CRASH ON SF'S BAY BRIDGE HOURS AFTER ELON MUSK ANNOUNCES "SELF-DRIVING" FEATURE*. Retrieved April 21, 2024, from The Intercept: <https://theintercept.com/2023/01/10/tesla-crash-footage-autopilot/>

- Korosec, K. (2022, October 26). *Ford, VW-backed Argo AI is shutting down*. Retrieved April 19, 2024, from Tech Crunch: https://techcrunch.com/2022/10/26/ford-vw-backed-argo-ai-is-shutting-down/?guccounter=1&guce_referrer=aHR0cHM6Ly9hcnN0ZWNobmljYS5jb20v&guce_referrer_sig=AQAAA3KrBuE9-HVidXSyBOXSJVEzaPrgGuh87mld0zx_1IDF8q3oxO4WBO6yAmfFeEF5QPctV34xehXaj5Ai0fodeBeRrNIGH5Un7
- 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. (2022). *Observing the Effect of a Crash on Twitter Sentiment: Early Results from Time Series Data*. doi:10.13140/RG.2.2.36194.61123
- McGurrin, M. (n.d.). *AV Sentiment Index*. Retrieved April 27, 2023, from AV & EV Sentiment Indices: <https://www.mcgurrin.info/sentiment/AVsentiment.shtml>
- McGurrin, M. F. (2023). *Seasonality and Variance of Twitter Sentiment Regarding Electric and Automated Vehicles*. doi:0.13140/RG.2.2.21739.14883/1
- Mendoza, D. (2022, April 13). Driverless car got pulled over by the police, then bolted. The company said it was by design. *13 News Now*. Retrieved April 17, 2024, from <https://www.13newsnow.com/article/news/nation-world/driverless-car-pulled-over/507-3bde6611-8f9f-4285-bc8f-382f66d1a6cc>
- Mitchell, R. (2022, August 31). *Bill targeting Tesla's 'self-driving' claims passes California Legislature*. Retrieved April 18, 2024, from LA Times: <https://www.latimes.com/business/story/2022-08-31/newly-passed-bill-could-force-tesla-to-scrap-the-name-full-self-driving-in-california>
- Othman, K. (2023). Public attitude towards autonomous vehicles before and after crashes: A detailed analysis based on the demographic characteristics. *Cogent Engineering*. doi:10.1080/23311916.2022.2156063
- Penmetsa, P., & et.al. (2021). Effects of the autonomous vehicle crashes on public perception. *IATSS Research*, 45(4), 485-492. doi:S0386111221000224
- Reuters. (2022, August 16). *BMW confirms fatal crash, says car wasn't self-driving*. Retrieved April 18, 2024, from Reuters: <https://www.reuters.com/business/autos-transportation/bmw-confirms-fatal-crash-says-car-wasnt-self-driving-2022-08-16/>
- Ruan, T., & Lv, Q. (2022). Public perception of electric vehicles on reddit over the past decade. *Communications in Transportation Research*, 2(0). Retrieved April 28, 2023, from <https://www.sciencedirect.com/science/article/pii/S2772424722000208>
- SAE International. (2021, April 3). *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_202104*. Retrieved April 18, 2024, from SAE International: https://www.sae.org/standards/content/j3016_202104/
- Schneider, B. (2022, July 1). Cruise driverless cars blocked traffic for hours Tuesday night. *San Francisco Examiner*. Retrieved April 18, 2024, from <https://www.sfexaminer.com/news/cruise-driverless->

[cars-blocked-traffic-for-hours-tuesday-night/article_57e8d2c8-f8c4-11ec-b99d-33fd6c31cb3e.html](https://www.theverge.com/2022/12/22/23523201/tesla-fsd-braking-crash-bay-bridge-california-chp)

Shakir, U. (2022, December 2022). *Tesla's Full Self-Driving is blamed for eight-car pileup in California*. Retrieved April 20, 2024, from The Verge: <https://www.theverge.com/2022/12/22/23523201/tesla-fsd-braking-crash-bay-bridge-california-chp>

Shepardson, D. (2022, September 1). *GM startup Cruise recalls and revises self-driving software after crash*. Retrieved April 18, 2024, from Reuters: <https://www.reuters.com/business/autos-transportation/gms-cruise-recalls-revises-self-driving-software-after-crash-2022-09-01/>

Shepardson, D., & Stempel, J. (2022, August 5). California regulator claims Tesla falsely advertised Autopilot, Full Self-Driving. *Reuters*. Retrieved 04 18, 2024, from <https://www.reuters.com/business/autos-transportation/california-regulator-claims-tesla-falsely-advertised-autopilot-full-self-driving-2022-08-05/>

Sky News. (2023, February 17). *Traffic lights 'could be obsolete in 20 years' after UK driverless cars trial*. Retrieved April 21, 2024, from Sky News: <https://news.sky.com/story/traffic-lights-could-be-obsolete-in-20-years-after-uk-driverless-cars-trial-12812875>

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/>

TrivyZa, M.-F. (2021). *Autonomous Vehicles*. Masters Thesis, National Technical University of Athens, SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING.

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