NAME: Sarah Grobe
DATE: 5/10/2022
Final Project – Scientific article
STAT/CS 387

Title

Evaluating Changes in Amazon Review Volume and Sentiment Over Time

Abstract

Amazon product review data serve as a relatively simple way to gather large amounts of data regarding consumer opinions on products over time. By coupling these review data with a sentiment analysis of the review text, a time series analysis can then be conducted on the sentiment of these reviews over time to learn more about how consumer opinions change as time goes on. In decomposing these data to evaluate their trends and seasonality, it was determined that some product categories do indeed see overall trends in their review sentiment or the volume of reviews made for those products over time. Similarly, some product categories also show a seasonal relationship some of these variables as well. These results could then be used to learn more about when users are more or less likely to write reviews, as well as how their opinions of products may change over time, which may in turn be useful for businesses hoping to increase outreach to customers and gain more feedback from reviewers.

Introduction

The use of the internet in today's world allows for massive volumes of data to be available to anyone. Being able to understand, interpret, and analyze this data can then be an incredibly powerful tool to gain knowledge about the world around us, and understand patterns we hadn't previously been aware of or even anticipated. One project of this type, conducted by Grobe, Kretzler, and Temkin (2021), made use of publicly available Amazon review data to evaluate the overall sentiment of the reviews to better understand which aspects of products, as well as which product types, consumers were most compelled to write reviews about online. Combining the compound sentiment with these classifications, they were able to gain a better sense of which product aspects and product categories had the strongest sentiment and largest volume of reviews associated with them.

Sentiment analysis itself is an incredibly interesting and complex issue itself, as computers are able to extract meaning and value from written words. Since these words can even be difficult for humans to interpret at times, added complications such as misspellings or idioms can make it difficult for a computer to extract the true meaning as well. In this way, sentiment analysis does have quite a few shortcomings, however, it still has value and should not be entirely disregarded (Balaji et al. 2017). Instead, it should be used as a sort of estimate, giving a rough guess of the true sentiment for a given block of text. One such use for this type of sentiment analysis is combining it with time series data in order to forecast sentiment over time (Liapis et al. 2021). Performing this type of analysis on Amazon review data in particular is also not unheard of, and can help provide added information about the content of the review (Mukherjee et al. 2019).

Forecasting is a method of predicting future values in a time series, and has been applied to Amazon review data previously, for instance to predict future ratings of products (Woo & Mishra 2020). There are also a wide range of methods available to achieve this forecasts, which depend in large part upon how far into the future one is planning to forecast (Hagan & Behr 1987, Hippert et al. 2001). These methods range from simple naïve methods such as carrying the last value forward, to neural networks, to even deep learning algorithms (Lara-Benítez et al. 2021, Hippert et al. 2001). In this way, we are able to not only extract meaning through text like the review data, but also attempt to see into the future in regards to future values of these sentiments.

In conducting this forecasting, another possible step in the time series analysis is decomposing the data into three pieces: the trend, seasonality, and remainder. This method allows us to evaluate each of these pieces separately from each other, in order to learn more about the respective relationships (Makridakis 1978). These relationships can also be either additive or multiplicative, depending on whether they increase over time or if the magnitude is relatively constant (Mbuli et al. 2020). In addition to forecasting, this decomposition of the data helps to gain a better understanding of the various relationships, or lack thereof, it may have with time.

Picking up more or less where Grobe, Kretzler, and Temkin (2021) left off, this project aims to combine the aforementioned methods of forecasting and decomposition with the sentiment analysis on Amazon review data in order to gain an understanding any relationships that may exist between review sentiment over time, as well as the number of reviews over time. These variables will be analyzed both for the data set as a whole, and within each of the 24 product review categories to determine if any interesting features such as overall trends or seasonal relationships are apparent, and if so, what those various relationships mean and how they could be used.

Methods

The raw data obtained from He & McAuley (2016) featured nearly 83 million Amazon reviews, resulting in over 18 GB of data. Due to the time and space complexity of this raw data set, it was then reduced to a subset of roughly one hundred thousand reviews, equally distributed across the 24 product categories, resulting in over four thousand reviews per category. This allowed for reasonably large sample sizes in total as well as across categories, while also reducing the time and space complexity of the data to be much more manageable for use on a personal computer.

Once the subset was obtained, a sentiment analysis was conducted using the VADER sentiment analyzer. This allowed for a compound sentiment to be recorded for the content of each review in the subset. The compound sentiment ranges from -1 to 1, where a value of -1 equates to entirely negative sentiment, while 1 equates to entirely positive sentiment. The compound sentiment was determined to be the most useful metric for this investigation due to its ability to provide insight into both positive and negative sentiment in the form of a single value. Further information regarding the data subsetting process and sentiment analysis up to this point can be found in Grobe et al. (2021).

The data were then analyzed across all categories and within each of the 24 categories through exploratory data analysis, or EDA. In particular, the EDA focused on the number of reviews over time, as well as the compound sentiment over time, to help determine whether any relationships may exist between these values and time.

One key feature to the data which revealed itself during the EDA was the drastic increases in the number of reviews over time for all categories, as shown below in Figure 1. Since the data run from roughly 1998 until 2014, it makes logical sense that this type of trend of the number of reviews would be apparent, largely due to the rise in popularity and accessibility of the internet over the course of this timeframe.

As a result of this increase in the total number of reviews, it was determined that the low sample sizes earlier in the dataset compared with the larger sample sizes later in the data set may result in some confounding due to the expected high variability, particularly since these trends became even more visible when viewed for each individual data category, as opposed to the data set as a whole. To help mitigate these effects, the data were then further subsetted to only evaluate the 6 year period ranging from January 1, 2008 through January 1, 2014, when all categories have generally consistently larger sample sizes.

The next step in the EDA was to then check the data for any potential trends and seasonality over time. To evaluate this, scatterplots were created for compound sentiment over time both for the entire subset, as well as for each product category within the aforementioned 6 year timeframe. To aid in the evaluation of seasonality, vertical lines were also included in these plots at each new year, to help make yearly seasonal trends more apparent. To evaluate any trend or seasonality in the number of reviews over time, histograms were created for the number of reviews for the data overall as well as by category over the 6 year timeframe, binned by month. The results of this part of the EDA can be shown in Figures 1 and 2 below.

Using the results of the EDA to help inform the next steps of the analysis, the time series data were then decomposed into the three parts: trend, seasonality, and remainder (also referred to as residual). In particular, all decompositions were assumed to be additive (rather than multiplicative) due to little evidence from the EDA results of seasonal effects changing over time relative to any trend in the data. It was also determined at this point in the investigation that little to no significant effect was apparent on the trend of the number of reviews, either in total or within the product categories, over time beyond the previously mentioned increase. Therefore, the volume of reviews was only examined further for seasonality, with trend and seasonality both examined for the compound sentiment within review categories.

As a sort of secondary exploratory analysis, the sentiment data for each of the 24 categories were all decomposed into these three portions, and scatterplots of the sentiment trend as well as autocorrelation plots were created to help evaluate trend and seasonality, respectively. In product categories which appeared to have a strong trend, as evidenced by the autocorrelation plot, the data were detrended and the analysis was performed again, to aid in isolating any seasonal effect that may be present separate from the trend. These results are shown below in Figure 4.

Finally, the help extract further meaning and interpretability from the results, the F statistics for trend and for seasonality were calculated for each of the 24 product categories. These values were then used to rank the product categories and help quantify which showed the strongest and weakest trends and seasonal effects. The results of this are shown in Tables 1 and 2 below.

Results

As mentioned above, the first step in the EDA was to evaluate the total number of reviews and the compound sentiment of those reviews across all 24 product categories. These overall results for compound sentiment are shown below in Figure 1.

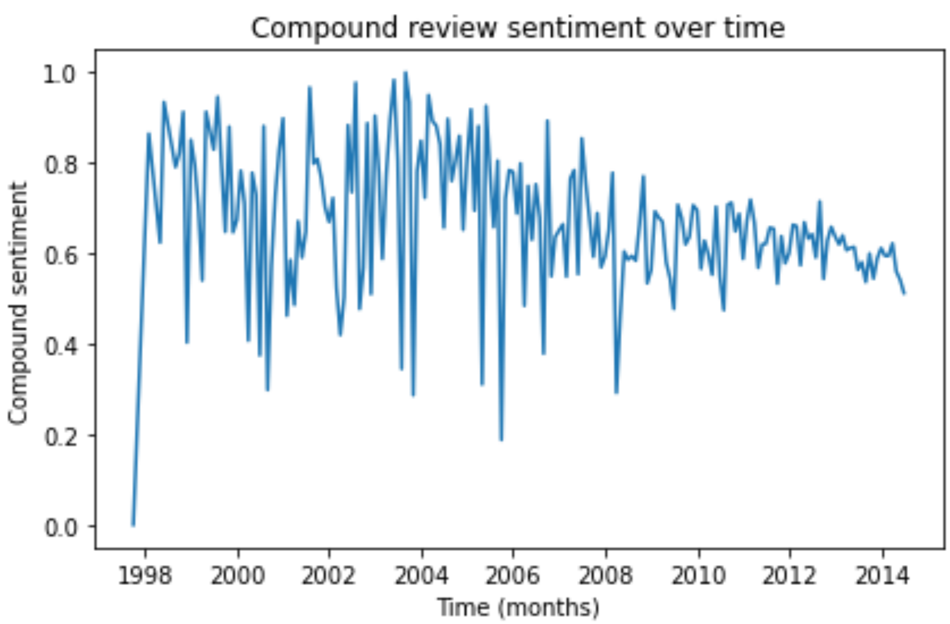

Figure 1. The compound sentiment review for the subset of one hundred thousand reviews across all 24 categories for the full timeframe of the data, from 1998 through 2014.

In the plot above, we see considerable variation in the sentiment values earlier on in the data set, and less variation later on, particularly after 2008. This variation is related to the total number of reviews per month, as shown below in Figure 2.

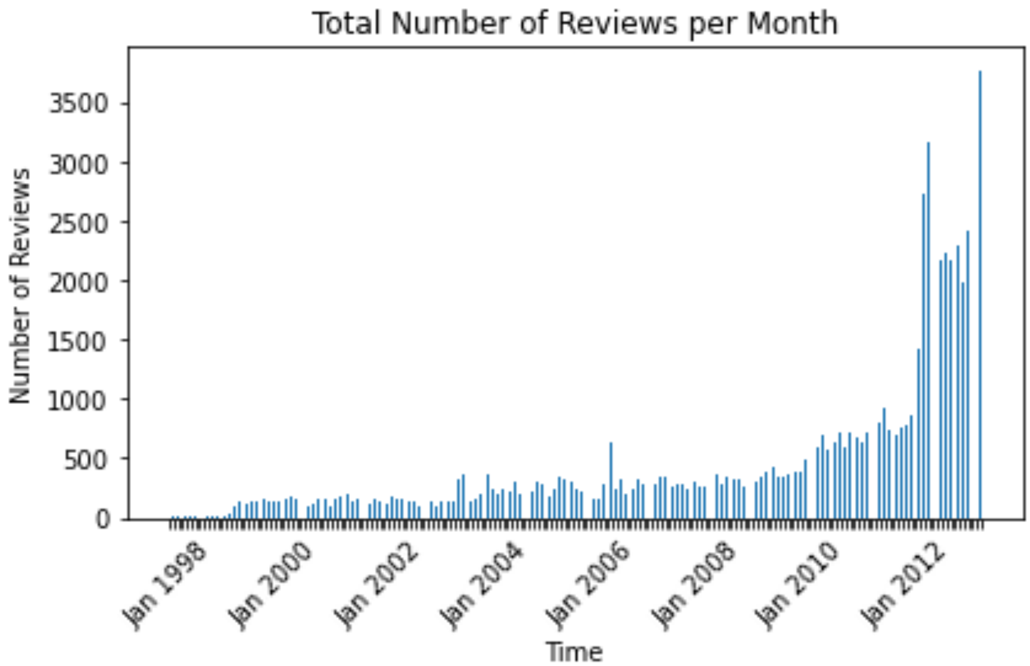

a.

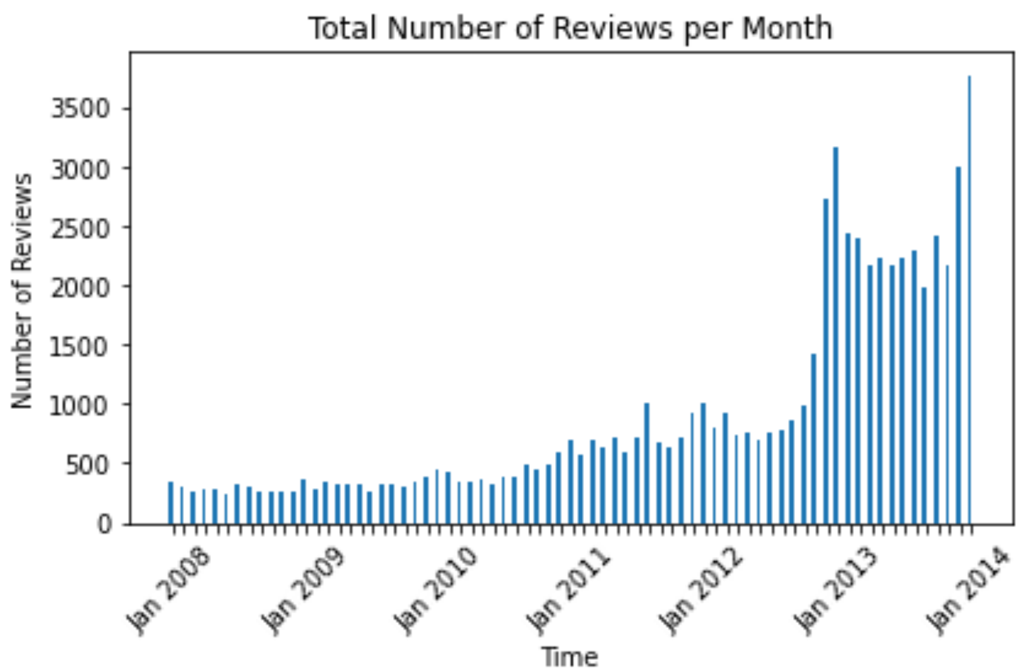

b.

Figure 2. Subplot a depicts the total number of reviews per month across all 24 product categories for the entire range of the data, from 1998 through 2014. Subplot b shows a zoomed in portion of that plot when the review numbers are generally higher, from January 1, 2008 through January 1, 2014.

As is shown in subplot a of Figure 2 above, the total number of reviews are very low earlier on the data set, particularly over the course of the first five years or so. The number of reviews sees a steady increase over the course of the 2000s, however, with generally very high volumes of reviews over the last couple years of the data set. To help account for this, as well as the associated variations in variability depicted in Figure 1, further analysis will be conducted using only the subset shown in subplot b, from 2008 to 2014. Additionally, further analysis of the volume of reviews within categories across this time period yielded results which looked fairly similar to the overall number of reviews as depicted in subplot b. For this reason, it was determined that the presence of any type of trend or seasonality related to product category is very unlikely, and for that reason the remainder of the analysis will instead focus on compound sentiment.

For the next portion of the analysis, the data were evaluated for the presence of trends or seasonality in the compound sentiment within each product category. As a way to quantify the strength of the trends and seasonality, the F statistic for the two respective features were calculated for each of the product categories. Shown below in Table 1 are the 10 strongest and 10 weakest product categories for trends, and in Table 2 the same for seasonality.

Strongest trends:		Weakest trends:	
0.938	Amazon_Instant_Video	0.011	Apps_for_Android
0.913	Beauty	0.469	Books
0.910	Sports_and_Outdoors	0.678	Movies_and_TV
0.901	Cell_Phones_and_Accessories	0.731	Musical_Instruments
0.894	Tools_and_Home_Improvement	0.756	Pet_Supplies
0.886	Grocery_and_Gourmet_Food	0.783	Digital_Music
0.875	Health_and_Personal_Care	0.784	Toys_and_Games
0.866	Patio_Lawn_and_Garden	0.799	Video_Games
0.864	Automotive	0.801	Kindle_Store
0.861	CDs_and_Vinyl	0.807	Baby

Table 1. On the left side of the table are given the product categories with the strongest trends for compound sentiment according to their corresponding F statistic, with the F statistic on the right and the product category name on the left. In the same format are those with the weakest trends, shown on the right.

Strongest seasonalities:		Weakest seasonalities:	
0.668	Apps_for_Android	0.173	Amazon_Instant_Video
0.580	Patio_Lawn_and_Garden	0.209	Baby
0.554	Electronics	0.212	Automotive
0.521	Clothing_Shoes_and_Jewelry	0.252	Kindle_Store
0.453	Movies_and_TV	0.263	Health_and_Personal_Care
0.411	Toys_and_Games	0.263	Digital_Music
0.364	Tools_and_Home_Improvement	0.266	Home_and_Kitchen
0.357	Pet_Supplies	0.268	Sports_and_Outdoors
0.342	Musical_Instruments	0.270	Books
0.341	CDs_and_Vinyl	0.273	Beauty

Table 2. On the left side of the table are given the product categories with the strongest seasonality for compound sentiment according to their corresponding F statistic, with the F statistic on the right and the product category name on the left. In the same format are those with the weakest seasonality, shown on the right.

The results shown here in Tables 1 and 2 are able to reaffirm and quantify much of what was suggested by the EDA, particularly in terms of which product categories showed results of a seasonal effect and which did not.

Finally, to help visualize some of these results shown in the above tables, plots were created which isolated the trend and seasonality for the product categories. Some of the most significant of these plots, as evidenced by the above F statistics, are shown below starting with Figure 3.

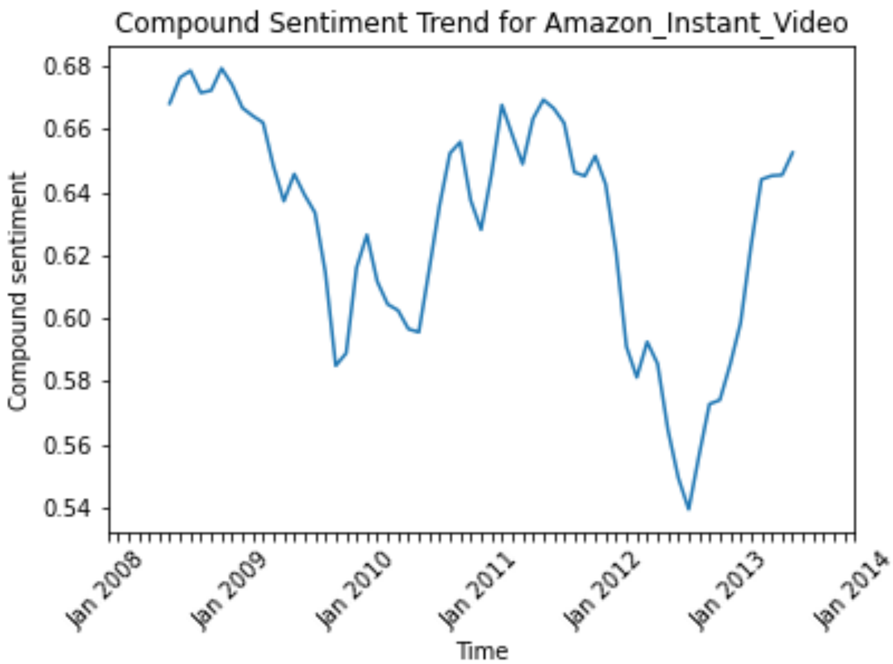

a.

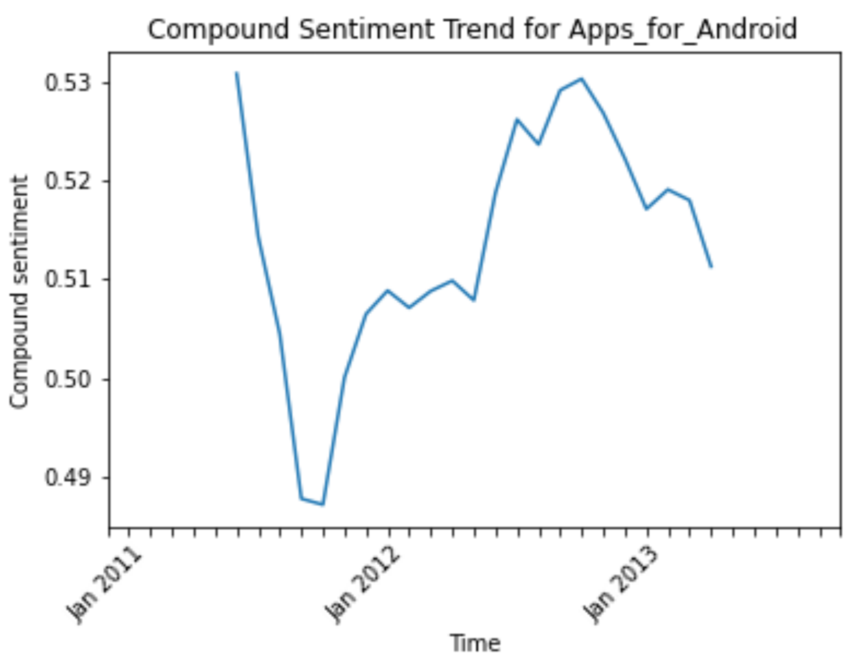

b.

Figure 3. Sample plots of trends for individual product categories. Subplot a shows the trend for Amazon Instant Video, which had the largest trend F statistic as shown in Table 1, while subplot b shows the trend for Apps for Android which had the weakest trend, also shown in Table 1.

The results of the graphs shown above in Figure 3 ultimately match the results present for the F statistics displayed in Table 1. Specifically, it is clear that Electronics in subplot a has a very strong positive trend in the sentiment over time, ranging about 0.2 in the compound sentiment overall and about 0.15 from the beginning of the interval to the end. By contrast, subplot b shows some fluctuation in the compound sentiment, although generally very little change. In particular, the maximum change for this product category is only about 0.1, and has a relatively small net change from the beginning of the interval to the end. These results are consistent with the F statistics given in Table 1.

To evaluate the seasonality of these categories, autocorrelation plots were generated. Some categories, however, featured not only a strong seasonal effect, but a strong trend as well. This was most apparent in evaluating the number of reviews, due to the strong positive trend in the overall number of reviews as time went on. In the event of a strong trend, the data for that product category were "detrended" in order to isolate any seasonal effect that may exist independent of the trend. Figure 4 shows sample results of this process.

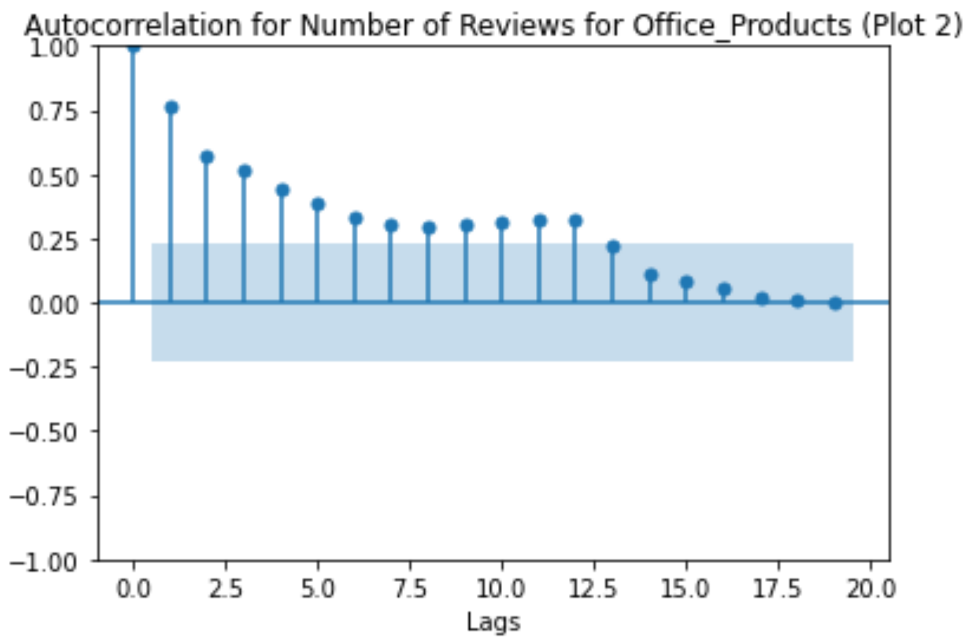

a.

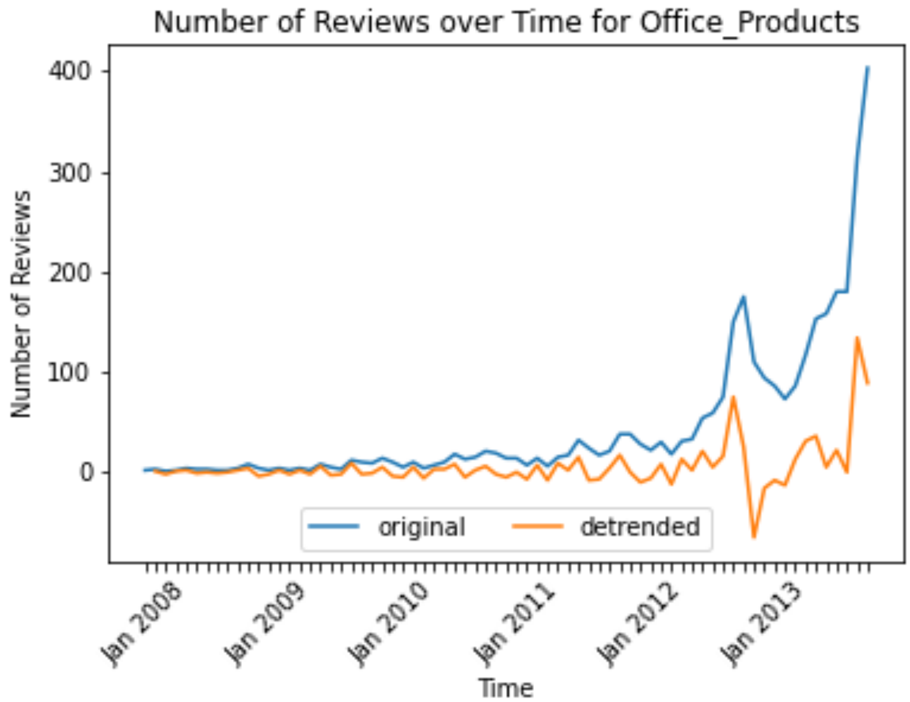

b.

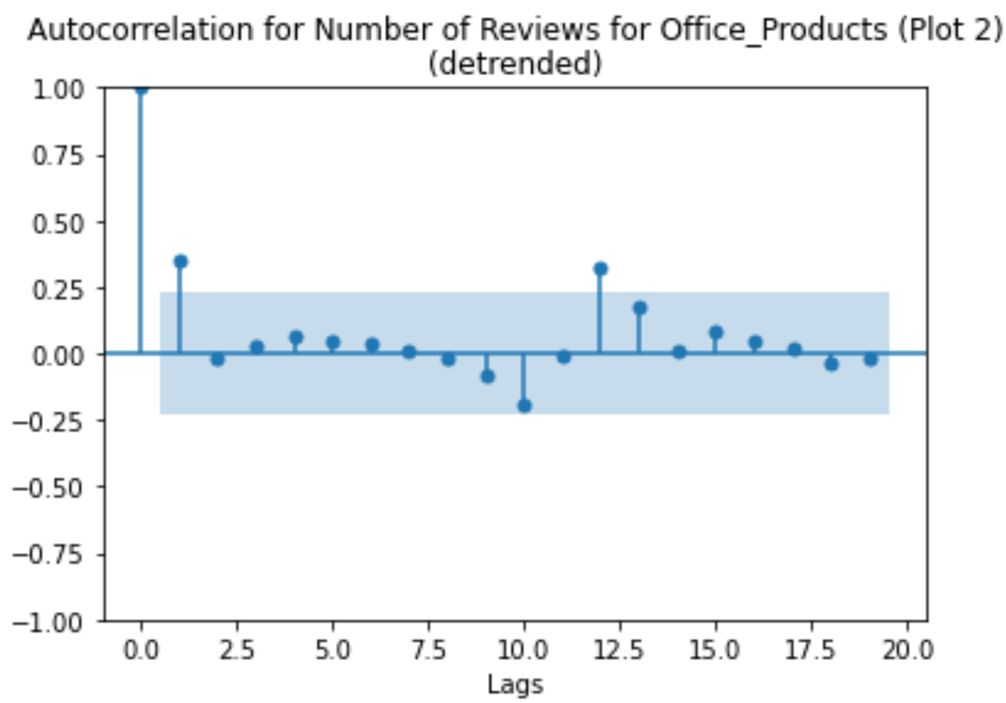

c.

Figure 4. Depicted in these three subplots is the process of detrending the data in order to evaluate seasonality. Subplot a depicts the autocorrelation plot for Office products, which shows a very clear downward trend, which in turn dampens the potential seasonality in the data. Subplot b shows this original data in blue with the detrended data in orange. The detrended data is centered around zero, eliminating any changes in the sentiment over time.

Subplot c finally shows the autocorrelation plot as is shown in subplot a, but run on the detrended data.

The transformation in Figure 4 shown between subplots a and c allows for any seasonality to be isolated from the trend, and account for the increase in number of reviews over this time period. This same process was repeated for all product categories which showed this similar trend. Not all product categories had a significant trend of this type, however, and therefore did not need to undergo this detrending transformation.

Further evaluation of the seasonality of number of reviews showed some significant seasonal effects within some of the categories. One such example is shown below, in Figure 5.

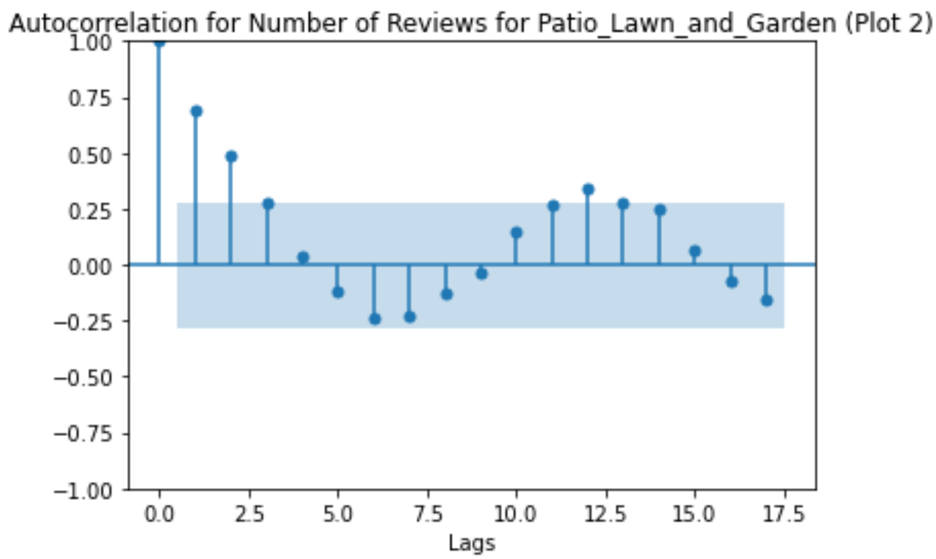

a.

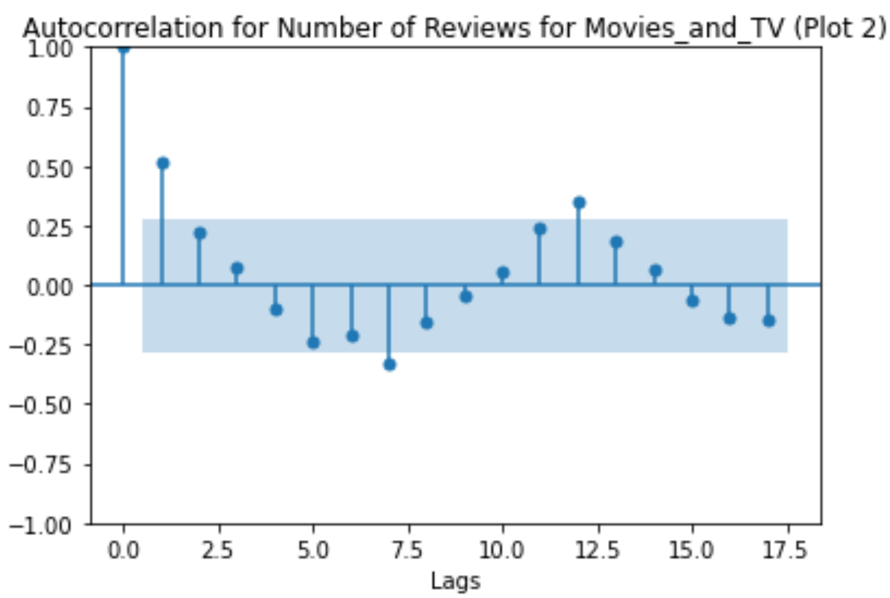

b.

Figure 5. Autocorrelation plots depicting clear seasonal effects for product categories which did not require detrending. Subplot a depicts this for Patio products, while subplot b depicts it for Movies and TV.

As is depicted in both subplots of Figure 5 above, some product categories did exhibit clear seasonal fluctuations in the volume of reviews across the time period examined, which could be indicative of people being more or less likely to write reviews for certain kinds of products during certain points of the year.

Similar plots were created to display those categories with a strong seasonal effect for compound review sentiment as well. Figure 6 below shows plots of some of the stronger seasonal effects for sentiment.

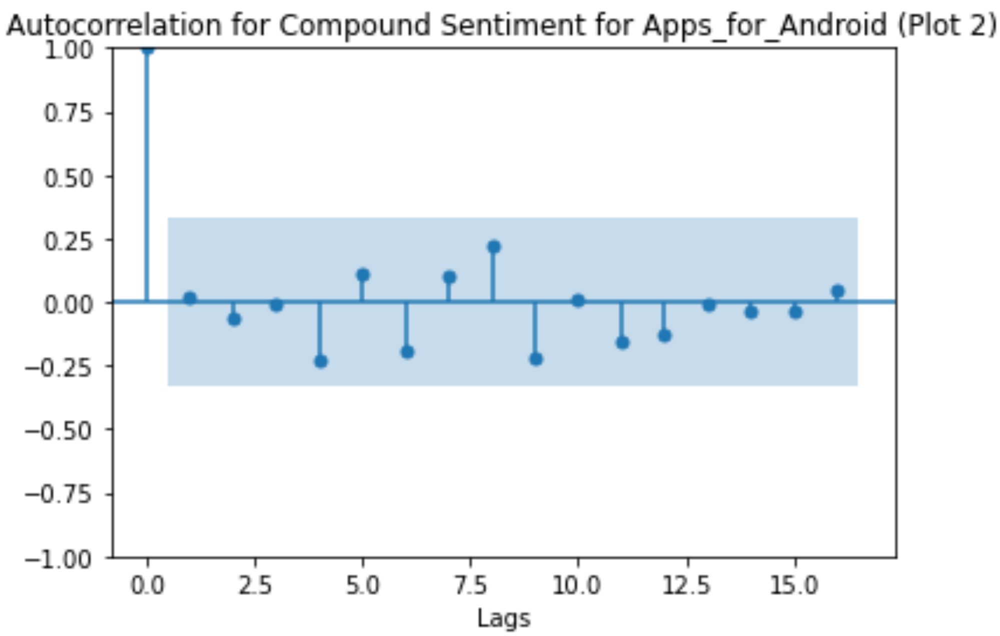

a.

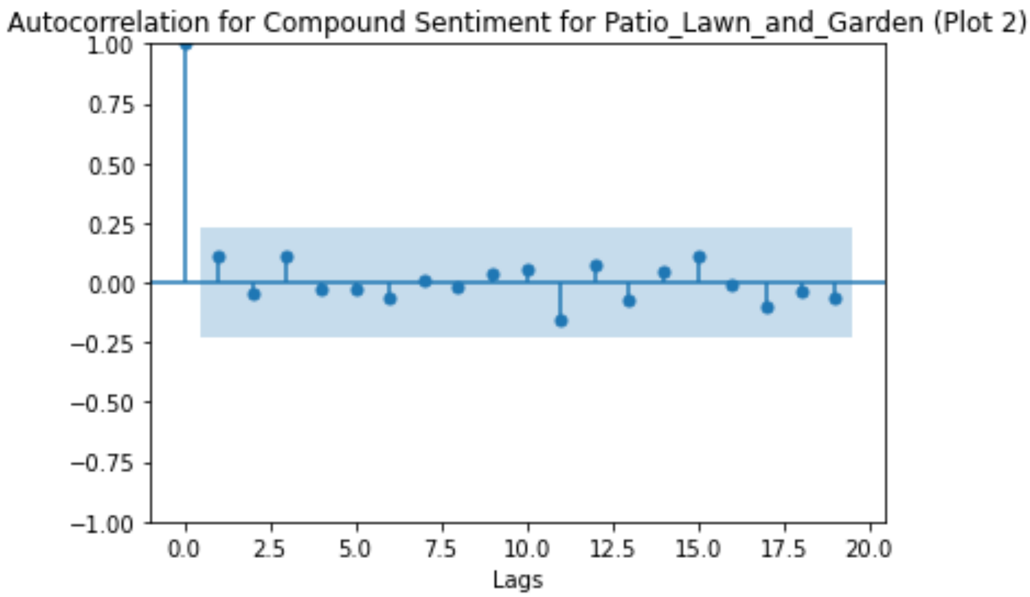

b.

Figure 6. Autocorrelation plots for some of the stronger seasonal effects for compound sentiment review within the six year period. Subplot a depicts the seasonality for Apps for Android, which had the strongest seasonal effect as shown in Table 2. Also shown in subplot b is an autocorrelation plot for Patio products, which had the next strongest seasonality.

While the seasonal results shown above in Figure 6 do not appear to be quite as strong some of the effects for volume of reviews as shown in Figure 5, they do still indicate that there could be possible fluctuations in overall sentiment over the course of a year.

Discussion

As has been shown above, certain product categories do depict clear trends and/or seasonal effects in the volume of Amazon reviews written about them, or in the overall sentiment, positive or negative, for those reviews. Products in the Movies and TV category, for instance showed a very clear seasonal effect, indicating that there are certain points in the year at which people are more or less likely to write a review for products within that category. Similarly, the Electronics category showed a very clear positive trend over the six year period that was analyzed, indicating that in general, the overall sentiment for those types of products has been increasing, possibly indicating an increase in quality and overall customer satisfaction with those types of products. This same product category also showed a potential seasonal relationship between review sentiment and time, indicating that certain times of the year could actually result in generally more positive or more negative sentiments being written about products within that category.

This type of analysis and the results from it could be incredibly useful to businesses, particularly in terms of their marketing teams. If businesses are able to isolate when customers are more likely to write a view using any seasonality that may be present for that type of product, they may be able increase the feedback that they receive on their product, which may also be used not only as a way to help improve the product, but also as a means to increase exposure and others learn about the product, and potentially buy it as well. Similarly, being able to predict when reviewers might leave generally more positive comments may be a way for businesses to help increase something like a product rating, while avoiding reaching out to potential reviewers during a period in which their review is more likely to be negative. This type of review and analysis could be particularly useful to small businesses, which likely do not have the capabilities for any kind of formal marketing team or large-scale outreach. By using this type of analysis, they could help determine, particularly through the use of seasonal effects, when might be the most efficient time period to focus their efforts on outreach in order to receive review feedback from consumers.

As was mentioned previously, it should be noted that sentiment analysis is not perfect, and in fact has a lot of shortcomings associated with it. While these types of shortcomings are essentially unavoidable when working with any kind of sentiment analysis, this should still be kept in mind when interpreting these results. Additionally, these analyses were conducted on a subset of one hundred thousand reviews, which was then further reduced to a subset spanning six years. As mentioned, these subsets were created to aid in time and space complexity, particularly since these analyses were being run on a personal laptop. Ideally, these analyses should be executed on the entire available data set (just under 83 million total reviews), in order to use as much information as possible to learn about the trends and seasonal relationships in the data. While that type of analysis was not practical for this particular project, using the entire data set could have useful and interesting effects, including altering some of the relationships which were discussed here. In this way, it is important to bear in mind that despite efforts to extract a representative sample from the total data set of 83 million reviews, it is possible, and even likely, that important information was left out of the subset, which could ultimately influence the results.

Finally, there are several directions a project of this type could take in the future. In particular, a deeper dive into forecasting to better analyze the trends and seasonal relationships and extrapolate those relationships to future values could be of use for being able to apply these data to something like a marketing strategy. Additionally, it could be worthwhile to examine this data set in conjunction with other data to help evaluate for possible confounders which could contribute to the trends or seasonal relationships which showed up as a result of these analyses. One example of this is using GDP data to evaluate the relationship between review sentiment and GDP, as was examined by Grobe et al. (2021). This type of analysis could also aid in better forecasting results, if that avenue is pursued further as well.

References

1. Balaji, P., Nagaraju, O., & Haritha, D. (2017). Levels of sentiment analysis and its challenges: A literature review. *2017 International Conference on Big Data Analytics and Computational Intelligence (ICBDAC)*, 436–439. https://doi.org/10.1109/icbdaci.2017.8070879

2. Deb, C., Zhang, F., Yang, J., Lee, S. E., & Shah, K. W. (2017). A review on time series forecasting techniques for building energy consumption. *Renewable and Sustainable Energy Reviews*, *74*, 902–924. https://doi.org/10.1016/j.rser.2017.02.085

3. Fry, C., & Manna, S. (2016). Can we group similar Amazon Reviews: A case study with different clustering algorithms. *2016 IEEE Tenth International Conference on Semantic Computing (ICSC)*, 374–377. https://doi.org/10.1109/icsc.2016.71

4. Grobe, S., Kretzler, B., & Temkin, B. (2021, December 15). *Amazon Review Sentiment is Correlated to GDP and General Satisfaction Across Time*.

5. Hajiali, M. (2020). Big Data and sentiment analysis: A Comprehensive and Systematic Literature Review. *Concurrency and Computation: Practice and Experience*, *32*(14). https://doi.org/10.1002/cpe.5671

6. Lara-Benítez, P., Carranza-García, M., & Riquelme, J. C. (2021). An experimental review on Deep Learning Architectures for time series forecasting. *International Journal of Neural Systems*, *31*(03), 2130001. https://doi.org/10.1142/s0129065721300011

7. Liapis, C. M., Karanikola, A., & Kotsiantis, S. (2021). A multi-method survey on the use of sentiment analysis in Multivariate Financial Time Series forecasting. *Entropy*, *23*(12), 1603. https://doi.org/10.3390/e23121603

8. Makridakis, S. (1978). Time-series analysis and forecasting: An update and evaluation. *International Statistical Review / Revue Internationale De Statistique*, *46*(3), 255–278. https://doi.org/10.2307/1402374

9. Mbuli, N., Mathonsi, M., Seitshiro, M., & Pretorius, J.-H. C. (2020). Decomposition forecasting methods: A review of applications in Power Systems. *Energy Reports*, *6*, 298–306. https://doi.org/10.1016/j.egyr.2020.11.238

10. Mukherjee, A., Mukhopadhyay, S., Panigrahi, P. K., & Goswami, S. (2019). Utilization of oversampling for multiclass sentiment analysis on Amazon Review Dataset. *2019 IEEE 10th International Conference on Awareness Science and Technology (ICAST)*. https://doi.org/10.1109/icawst.2019.8923260

11. Mäntylä, M. V., Graziotin, D., & Kuutila, M. (2018). The evolution of sentiment analysis—a review of research topics, venues, and top cited papers. *Computer Science Review*, *27*, 16–32. https://doi.org/10.1016/j.cosrev.2017.10.002

12. Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with Deep Learning : A Systematic Literature Review: 2005–2019. *Applied Soft Computing*, *90*. https://doi.org/10.1016/j.asoc.2020.106181

13. Sharma, S., & Jain, A. (2020). Role of sentiment analysis in social media security and analytics. *WIREs Data Mining and Knowledge Discovery*, *10*(5). https://doi.org/10.1002/widm.1366

14. Syamala, M., & Nalini, N. (2019). A filter based improved decision tree sentiment classification model for realtime Amazon Product Review Data. *International Journal of Intelligent Engineering and Systems*, *13*(1), 191–202. https://doi.org/10.22266/ijies2020.0229.18

15. Woo, J., & Mishra, M. (2020). Predicting the ratings of Amazon products using Big Data. *WIREs Data Mining and Knowledge Discovery*, *11*(3). https://doi.org/10.1002/widm.1400

16. Yadav, A., & Vishwakarma, D. K. (2019). Sentiment analysis using Deep Learning Architectures: A Review. *Artificial Intelligence Review*, *53*(6), 4335–4385. https://doi.org/10.1007/s10462-019-09794-5

17. Hippert, H. S., Pedreira, C. E., & Souza, R. C. (2001). Neural networks for short-term Load forecasting: A review and Evaluation. *IEEE Transactions on Power Systems*, *16*(1), 44–55. https://doi.org/10.1109/59.910780

18. Armstrong, J. S., & Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, *8*(1), 69–80. https://doi.org/10.1016/0169-2070(92)90008-w

19. Hagan, M. T., & Behr, S. M. (1987). The time series approach to Short Term Load Forecasting. *IEEE Transactions on Power Systems*, *2*(3), 785–791. https://doi.org/10.1109/tpwrs.1987.4335210

20. Liu, K., Subbarayan, S., Shoults, R. R., Manry, M. T., Kwan, C., Lewis, F. I., & Naccarino, J. (1996). Comparison of very short-term load forecasting techniques. *IEEE Transactions on Power Systems*, *11*(2), 877–882. https://doi.org/10.1109/59.496169

21. Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*(6088), 533–536. https://doi.org/10.1038/323533a0

22. Mamun, A. A., Sohel, M., Mohammad, N., Haque Sunny, M. S., Dipta, D. R., & Hossain, E. (2020). A comprehensive review of the load forecasting techniques using single and hybrid predictive models. *IEEE Access*, *8*, 134911–134939. https://doi.org/10.1109/access.2020.3010702

23. Singla, P., Duhan, M., & Saroha, S. (2021). A comprehensive review and analysis of Solar Forecasting Techniques. *Frontiers in Energy*. https://doi.org/10.1007/s11708-021-0722-7

24. Nwokolo, S. C., & Ogbulezie, J. C. (2018). A quantitative review and classification of empirical models for predicting global solar radiation in West Africa. *Beni-Suef University Journal of Basic and Applied Sciences*, *7*(4), 367–396. https://doi.org/10.1016/j.bjbas.2017.05.001

25. Tealab, A. (2018). Time series forecasting using Artificial Neural Networks Methodologies: A systematic review. *Future Computing and Informatics Journal*, *3*(2), 334–340. https://doi.org/10.1016/j.fcij.2018.10.003

26. Hirschberg, J., & Manning, C. D. (2015). Advances in natural language processing. *Science*, *349*(6245), 261–266. https://doi.org/10.1126/science.aaa8685

27. Vinodhini, G., & Chandrasekaran, R. M. (2012). Sentiment Analysis and Opinion Mining: A Survey. *International Journal of Advanced Research in Computer Science and Software Engineering*, *2*(6).

28. Das, S. R., & Chen, M. Y. (2007). Yahoo! for Amazon: Sentiment extraction from small talk on the web. *Management Science*, *53*(9), 1375–1388. https://doi.org/10.1287/mnsc.1070.0704

29. Choi, J., Laibson, D., & Metrick, A. (2000). Does the internet increase trading? evidence from investor behavior in 401(k) plans. https://doi.org/10.3386/w7878

30. Khairullah Khan, Baharum B. Baharudin, Aurangzeb Khan, & Fazal_e_Malik. (2010). Automatic Extraction of Features and Opinion-Oriented Sentences from Customer Reviews. https://doi.org/10.5281/zenodo.1074753

31. Zhang, C., Zuo, W., Peng, T., & He, F. (2008). Sentiment classification for Chinese reviews using machine learning methods based on String Kernel. *2008 Third International*

Conference on Convergence and Hybrid Information Technology.
https://doi.org/10.1109/iccit.2008.51

32. He, R., & McAuley, J. (2016). *Amazon product data*. Amazon review data. Retrieved November 29, 2021, from https://jmcauley.ucsd.edu/data/amazon/.