UNIVERSITÀ POLITECNICA DELLE MARCHE
FACOLTÀ DI ECONOMIA "GIORGIO FUÀ"

Corso di Laurea Magistrale in Business Organization and Strategy

# Previsione della domanda nel settore della moda

# Forecasting demand in the fashion industry

Relatore: Chiar.mo
Prof. Pigini Claudia

Tesi di Laurea di:
Boccia Chiara

Anno Accademico 2022 – 2023

# 1 INTRODUCTION

This thesis aims to examine the application of big data analysis in demand forecasting for the fashion retail industry. As part of management's decision-making processes, forecasting is a common statistical tool in business. Given all the historical information, forecasting makes the most precise predictions, providing knowledge of possible future events. The competitiveness of fashion companies largely depends on their ability to make accurate forecasts. In particular, retailers' profit is affected by the level of accuracy of sales forecasts, as fashion retailers must always supply the right products at the right time while maintaining a good stock. Accurate forecasting of customer demand allows apparel manufacturers to increase or adjust their production, also affecting the performance of the entire supply chain. Sales forecasting is critical in the fashion industry because of all the uncertainty associated with supply and demand.

This thesis presents, in the second chapter, an analysis of the various forecasting models proposed in the literature for the fashion products. At first, the classical statistical models were studied which are fast in making predictions, these do not work well when it comes to predictions that require the use of highly complex data. Consequently, forecasting systems based on ANN models have been proposed in the literature. Although these have proven to be versatile, they require a longer computational time; these, for example, can take hours to complete a simple forecasting task. Consequently, although the ANN model offers accurate predictions, the time required is an obstacle to its real-world application. Therefore, nowadays, how to develop more accurate and timely sales forecasting methods becomes an important research topic. Recently, a relatively new learning algorithm for neural networks called extreme learning machine (ELM) has been proposed. The ELM model dramatically reduces the learning time of ANN and even makes real-time application possible. Although ELMs and ANNs are very accurate, they remain black boxes that do not highlight time series patterns, which can be useful for management to shed light on the behaviour of quantities relevant to them. In this sense, traditional forecast methods allow to identify, for example, whether there are trends and seasonality in sales.

In the third chapter, however, the study focus on forecasting time series methods, in particular, the most common classical statistical methods are Exponential smoothing and ARIMA. Exponential smoothing is often used for recognizing elements in time series such as trend and seasonality and how these are incorporated into the smoothing process, whereas, ARIMA seek to capture the autocorrelations in the data.

The fourth chapter illustrates the dataset related to the monthly sales of various clothing items of an Italian fashion store called "Nunalie" ", that will be later used for my empirical analysis. The sales are considered for a period of 4 years, specifically from January 2016 to December 2019. In particular, the monthly sales graphs of both individual products and the categories into which they were grouped, were analysed, focusing mainly on seasonality and the study of sub-series graphs.

In the fifth chapter, the forecasting of future sales of the 4 fashion categories into which the individual items were aggregated is carried out, using the Holt-Winter and Arima methods. Finally, the sixth chapter contains the main conclusions and the seventh contains the replication material for the analysis.

# 2 RETAIL SALES FORECAST IN THE FASHION INDUSTRY

## 2.1 General overview of retail sales forecast

Retail sales forecasting plays an increasingly important role nowadays since estimating the future demand for a retail product is an essential task for business planning. Sales forecasting is not only the basis on which business plans are built, but it is also essential for improving a company's competitive ability. The history of sales forecasting can be traced back more than 50 years; a large number of documents relating to sales forecasting have been published since its inception, involving a wide variety of applications in real-world industries, including the fashion industry.

Demand uncertainty exists in all sectors depending on a number of factors and no single forecasting model can account for all of them (Fildes, Ma and Kolassa, 2022). Forecasting has usually been used as a basis for the design, implementation of supply chain functions and for dealing with this uncertainty. Sales forecasting is generally a very complex problem due to the influence of internal and external environments, especially for the fashion industry.

A study of the literature on fashion retail sales forecasting shows that it is a rather difficult task since consumer tastes are constantly changing and, in addition, the life of the fashion product is relatively short. It is possible to add that sales of fashion products are influenced by a number of 'influencing factors' including seasonal factors (both temperature and weather strongly influence demand), fashion trend factors, weather, marketing strategy, political climate, article characteristics (such as colour), calendar (it has been tested that sales increase significantly during holidays) and macroeconomic trends. Other factors that contribute to make sales forecasts more difficult are the large number of stock-keeping units (SKUs) and limited historical data (this drawback occurs when considering many influencing factors at the same time or forecasting the sale of a new product). This leads to the search for increasingly sophisticated forecasting methods capable of evaluating all possible variables.

### 2.2 Evolution of quantitative forecasting methods in the fashion industry

For the study of different demand and sales forecasting models for fashion products, the literature has undergone several revisions over the years. Traditionally, companies used simple algorithms such as linear regressions to make forecasts based on an almost limited amount of available data. The rise of big data posed a great challenge to traditional algorithms and, thus, it led to a revolution in forecasting the demand for fashion products. From the research carried out, it is interesting to note that pure statistical methods have not been studied in the literature in the last 15 years not only because they have already been explored in the past but also because they are not sufficient to provide sophisticated forecasting results on their own. New studies first shifted to AI, but as these models are not sufficient to generate more accurate prediction results, research has shifted to the study of new hybrid models. These can utilise the strengths of different models together to form a new forecasting method.

Current forecasting techniques are generally divided into two groups: classical methods based on mathematical and statistical models (e.g. linear regression, exponential smoothing, ARIMA, SARIMA), and modern heuristic methods using artificial intelligence techniques including artificial neural networks (ANN) and evolutionary computation.

Classical methods use a linear functional form for modelling time series, but as they are unable to capture certain variables influencing trends, they may not be suitable for different real-world time series such as fashion retail forecasting. The forecasting performance of such classical methods depends above all on the existence of reliable historical data which guarantee both the correct identification of model structures and the efficient tuning of parameters. Therefore, these traditional models are inappropriate for a short historical data perturbed by variables. Recently, modern heuristic methods based on artificial intelligence have developed; these include both neural networks and hybrid models. These methods are used to build non-linear prediction models. Over the years, several authors such as Frank et al. or Thomassey et al. have studied sales forecasting in the fashion industry by comparing these new methods with traditional forecasting models. These comparisons were made since historical sales data are often short and particularly disrupted by numerous factors, which are neither controlled nor identified, such as weather, fashion trends and the economic environment.

The studies carried out over the years, on the evolution of forecasting methods in the field of fashion, has shown that each method has its own limitations and disadvantages, which are analysed below.

Analysing the forecasting horizon, it can be stated that most existing models make medium- and long-term forecasts, whereas short-term or real-time forecasts are not yet fully explored. In the world of fashion, however, a short-term forecast is very important because fashion trends are unpredictable and delivery times are getting shorter and shorter. The study on speed terms showed that while statistical methods can deliver forecast results very quickly, artificial intelligence methods take much longer. Over the years, we have seen how delivery times have evolved to the present day, for instance, in the fashion industry, fast fashion companies such as ZARA that are adopting a strategy with very short delivery times. Consequently, forecasting results must be available in a very short time for these companies. From the literature reviewed, we observe that due to the high speed of ELM, it can be a good candidate to work under fast fashion forecasting in conjunction with statistical methods. In addition, the fuzzy method can be used for fast fashion forecasting, this generates forecasts by combining those of different methods through fuzzy logic. This theory was proposed by Zadeh, (1965), but in the context of fashion retail forecasting, it was Sztandera et al. (2004) who created a new multivariable fuzzy logic model based on different variables such as colour, time and size. As the factors that create uncertainty increase, fuzzy systems prove to be more suitable because they do not have the same precision as analytical tools and are not cost-effective when the process follows a precise mathematical model. Good performance comes from the ability to identify non-linear relationships in the inputs. Moreover, it has been verified that the multivariate model performs better than univariate models. The study of the literature shows us that Frank et al. (2003) and Sztandera et al. (2004) used the combination of fuzzy methods with ANN because they are considered complementary in fact, while the former use expert linguistic and verbal input, the latter extract information from the systems.

General study of the literature shows that traditional statistical methods depend on time series data and that most statistical models are fast in making predictions. The ARIMA model, for example, makes time series predictions in seconds based on hundreds of historical data (Box et al., 2008). Statistical models are found to be effective and efficient provided users have the knowledge to choose the right model and select the appropriate parameters; this is highly dependent on the expert's knowledge. This can be considered a limitation especially if the one who has to make the decision is not sufficiently informed.

For this reason, companies have started to use artificial neural network (ANN) because it has attractive features such as the ability to learn and model data with any feature but most importantly, it does not require any specific expert knowledge of the problem to be studied. From the study of the literature, systems based on ANN models have been well proposed and the results are impressive (Chang, Wang and Liu, 2007; Chang, Wang and Tsai, 2005; Hamzaçebia, Akayb and Kutayb, 2009; Ling, Leung, Lam and Tam, 2003). With the emergence of ANN, prediction systems based on it have been widely developed and studied (cf. Hansen and Nelson, 1997). While we can say that artificial intelligence (AI) methods can perform better in terms of accuracy than traditional statistical prediction models, it should be pointed out that they commonly require much longer time and more computing power, this is a major obstacle for real-world application. Therefore, many researchers propose to combine multiple methods together to form a new "hybrid method" to achieve efficient and effective forecasting. Because for many industries, predictions must be made very quickly, a relatively new learning algorithm called extreme learning machine (ELM) has recently been proposed (Huang, Zhu, Siew, 2006a). This not only learns faster and with higher performance than traditional gradient-based learning algorithms. The ELM model dramatically reduces the learning time of ANN and even makes it possible to apply ELM in real time (Huang, Zhu, Siew, 2006b). It should be pointed out, however, that this algorithm is not perfect because it is known that ELM results, when there is a single run prediction, are unstable when compared with traditional statistical and ANN models. As a result, an averaging approach was proposed (Sun, Au and Choi, 2007) with the aim of improving the stability of

ELM predictions. This averaging step has been shown to be effective but still requires the ELM to be run repeatedly for many times, which increases the prediction time. As a result, it is not as fast as the original ELM model.

From the study of the literature, therefore, it was found that traditional statistical methods, AI systems and hybrid models were compared with each other by examining a number of parameters such as data processing speed, functional form, computational power, prediction accuracy interpretability and volume of data. By examining these values, results emerge that have been listed in the Table 1.

| | Models analyzed | | |
|---|---|---|---|
| Parameters | Statistic | AI/ML | Hybrid |
| Speed for data processing | Very fast | Slow due to high data to process | Faster than AI/ML |
| Functional Form | Linear | No linear | No linear |
| Computing power | Low | High | Higher than AI/ML |
| Prediction accuracy | Low | Moderate | High |
| Interpretability | Easily interpreted | Relatively difficult | Complex |
| Volume of data | Low | Great | Very wide |

Table 1: Summary comparison according to different parameters studied in the literature

## 3 FORECASTING TIME SERIES

In this chapter, I briefly describe the two traditional methods for predicting time series are exponential smoothing and ARIMA models, both of which offer alternatives approaches to the issue. I follow (Hyndman and Athanasopoulos, 2018), to which I refer the reader for further details.

While ARIMA models seek to capture the autocorrelations in the data, exponential smoothing methods are based on a description of the trend and seasonality in the data (Hyndman and Athanasopoulos, 2018).

### 3.1 Exponential Smoothing

Exponential smoothing, which was first presented in the late 1950s, served as the inspiration for some of the most effective forecasting techniques. Predictions made with the use of exponential smoothing techniques are weighted averages of earlier observations, with the weights exponentially decrementing with time. The associated weight is higher the more recent the observation.

The method is often chosen based on recognizing the time series' essential patterns such as trend and seasonality and how these are incorporated into the smoothing process, for instance, in an additive or multiplicative manner.

With this statistical approach, it is possible to perform model selection among competing models.

The simplest exponential smoothing technique is known as simple exponential smoothing (SES). This approach is appropriate for predicting data without a distinct trend or seasonal pattern.

Simple smoothing method is between the extreme cases of naïve method (it makes the assumption that the most recent observation is the only relevant one and that all earlier observations don't have any bearing on the future) and the average approach (it bases its forecasting on the premise that all observations are equally important and should be given similar weights).

Exponential smoothing models are calculated by:

$$F_{t+1} = F_t + \alpha\,(A_t - F_t) = \alpha A_t + (1-\alpha)F_t$$

where $F_t$ is the forecast calculated at time t (current) while $F_{t+1}$ is the future forecast, $A_t$ is actual observation and $0 < \alpha < 1$ is the smoothing parameter with higher values leading to greater responsiveness and lower values to greater stability. If forecasts are required for several periods in the future, such as t+1, the forecast for the first period t is used for all subsequent periods.

This can be observed by expanding the calculation above to the equivalent form to:

$$F_{t+1} = \alpha A_t + \alpha (1-\alpha) A_{t-1} + \alpha (1-\alpha)^2 A_{t-2} + \alpha (1-\alpha)^3 A_{t-3} + \cdots$$

This expression shows that higher weights are given to more recent demand observations.

The parameter α regulates how quickly the weights diminish.

Exponential smoothing refers to the fact that the weights assigned to past observations fall off exponentially as we go back in time.

More weight is given to observations from the farther past if α is low, or close to 0. The more recent data are given more weight if α is large, meaning it is close to 1.

Simple smoothing method requires using smoothing parameter and from this value it is possible to compute the values of all forecasts.

Sometimes smoothing parameters can be chosen trough subjective method because it is based on the previous experience.

However, estimating the unknown parameters from the observable data is a more dependable and impartial technique to obtain values for the parameters.

By extending simple exponential smoothing, Holt made it possible to forecast data that had a trend.

The following models are suitable for nonstationary time series, or time series with a trend. The first is double exponential smoothing, just considers base and trend, whereas the other additionally take seasonality into account.

### 3.2 Double Exponential Smoothing

Using this method, it is possible to have an estimate of the base or expected level of demand in period t (denoted by Et) and an estimate of the trend, in particular, the increase or decrease per period (denoted by Tt). Considering these values, the forecast for period t+n (n periods after the current period) is given by:

$$F_{t+n} = E_t + nT_t$$

The values of base and trend are updated by the following rules:

$$E_t = \alpha\, A_t + (1-\alpha)\,(E_{t-1} + T_{t-1})$$

$$T_t = \beta\,(E_t - E_{t-1}) + (1-\beta)T_{t-1}$$

The parameter $\alpha$ ($0 < \alpha < 1$) is used to smooth the base and the parameter $\beta$ ($0 < \beta < 1$) is used to smooth the trend. As in the exponential smoothing model, higher values of the smoothing parameters lead to more responsiveness while smaller values lead to more stability.

### 3.3 Holt-Winters Additive Method

In addition to the base and trend estimates of double exponential smoothing (E and T), the Holt-Winters methods (Holt, 1957; Winters, 1960) include estimates of the seasonal factors for periods (denoted by S). The parameters p, represents the number of seasonal periods in a year. For example, p = 12 would correspond to monthly seasonal adjustments and p = 4 would correspond to quarterly seasonal adjustments. In the additive version, the forecast for period t+n (n periods after the current period) is given by:

$$F_{t+n} = E_t + nT_t + S_{t+n-p}$$

The values of base and trend and the seasonal factors are updated by the following

$$E_t = \alpha\,(A_t - S_{t-p}) + (1-\alpha)\,(E_{t-1} + T_{t-1})$$

$$T_t = \beta\,(E_t - E_{t-1}) + (1-\beta)T_{t-1}$$

$$S_t = \gamma\,(A_t - E_t) + (1-\gamma)S_{t-p}$$

As before, α and β smooth base and trend while the parameter $0 < \gamma < 1$ is used to smooth trend.

The forecasting formula:

$$F_{t+n} = E_t + nT_t + S_{t+n-p}$$

can be used for any number of future periods as long as one interprets the final term to be the most recent estimate of seasonality for the period being forecasted.

### 3.4 Holt-Winters Multiplicative Method

The multiplicative version of the Holt-Winters method uses seasonal factors as multipliers rather than additive constants. The forecast for period t+n is given by:

$$F_{t+n} = (E_t + nT_t)\, S_{t+n-p}$$

The values of base and trend and the seasonal factors are updated by the following rules:

$$E_t = \alpha \frac{A_t}{S_{t-p}} + (1-\alpha)(E_{t-1} + T_{t-1})$$

$$T_t = \beta\,(E_t - E_{t-1}) + (1-\beta)T_{t-1}$$

$$S_t = \gamma \frac{A_t}{E_t} + (1-\gamma)S_{t-p}$$

**3.5 ARIMA**

ARIMA models provide another approach to time series forecasting and they aim to describe the autocorrelation in the data.

A subgroup of linear regression models called ARIMA models makes an effort to predict the future values of the target variable using the previous observations of the variable.

The fact that simple ARIMA models ignore external variables is a crucial component of these models. Instead, only previous values of the target variable (or features constructed from those past values) are used to make the forecast.

ARIMA stands for Autoregressive Integrated Moving Average, particularly, it is composed by several elements:

- ***Auto-regression (AR):*** describes a model where a variable is changing and regressing on its own lagged, or prior, values.
- ***Integrated (I):*** the integrated part of the model can be applied one or more time in order to allow the time series to become stationary.

- ***Moving average (MA):*** incorporates the relationship between a residual error from a moving average model applied to lagged observations.

  In ARIMA functions each component as a parameter with a standard notation. This notation is explained by parameters p, d, and q. These parameters are replaced with integer values to specify the type of ARIMA model being used. They can be defined as:

- ***p:*** The lag order, commonly known as the number of lag observations in the model.
- ***d:*** the degree of differencing, it represents the number of times the raw observations are differenced;
- ***q:*** the size of the moving average window also known as the order of the moving average.

  There are various terms in a linear regression model. A parameter that can have a value of zero (0) would indicate that the model should not include that specific component.

  The data are differenced in an autoregressive integrated moving average model to make it stationary.

  The ARIMA model offers different advantages, particularly it requires the prior data of a time series instead of generalize the forecast. Plus, it is useful for short-term forecasts.

The model offers also disadvantages, for instance the parameter (p,d,q) could be define in a quite subjective way, it no-performs well long term forecasts, its computation is expensive, finally, it cannot be used for performing seasonal time series.

### 3.6 Seasonal ARIMA model

A variety of seasonal data can be modelled using ARIMA models.

The ARIMA models we've seen so far are transformed into seasonal ARIMA models by adding more seasonal factors. The notation for an SARIMA model is specified as:

$$\text{ARIMA} \quad \underbrace{(p,d,q)}_{\substack{\uparrow \\ \text{Non-seasonal part} \\ \text{of the model}}} \quad \underbrace{(P,D,Q)_m}_{\substack{\uparrow \\ \text{Seasonal part of} \\ \text{of the model}}}$$

m is the number of observations made per year. This parameter influences P, D and Q parameters. For the model's seasonal components, it uses uppercase notation, while for its non-seasonal components, it uses lowercase notation.

The model's seasonal portion is made up of terms that are comparable to its non-seasonal elements but involve backshifts of the seasonal period.

There may be many different parameters and term combinations in seasonal ARIMA models. Therefore, when fitting models to data, it is appropriate to test out a variety of models and select the best fitting model using a suitable criterion.

# 4 DATA DESCRIPTION

## 4.1 Dataset description

The dataset is based on data from a real Italian fast fashion company, Nunalie, and consists of 5577 new products and sales for about 45 millions related to fashion seasons from January 2016 to December 2019. For each product, the dataset contains multimodal information: its image, textual metadata, sales after the first release date, and three related Google Trends describing category, colour, and fabric popularity. From all the different data made available, the following clothes were selected:

- long sleeves;
- short sleeves;
- sleeveless;
- tracksuit;
- shorts;
- miniskirt;
- trapeze dress;
- medium cardigan.

Figures 1 - 8 show monthly sales derived from the sum of weekly sales taken from the database available on Kaggle.

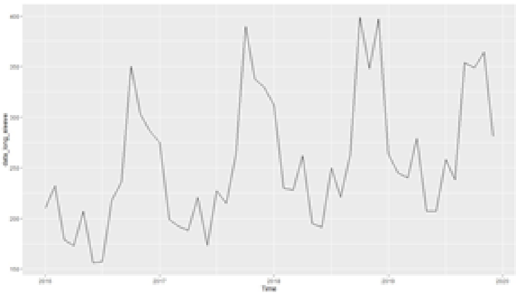

Figure 1: monthly sales long sleeves

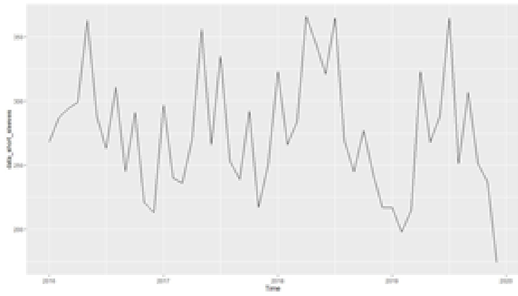

Figure 2: monthly sales short sleeves

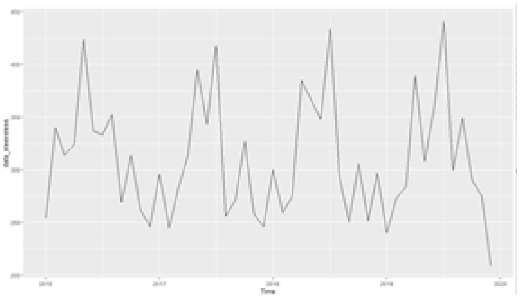

Figure 3: monthly sales sleeveless

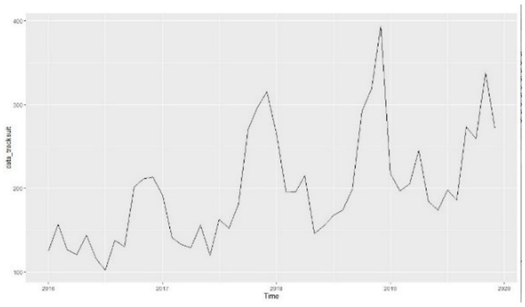

Figure 4: monthly sales tracksuit

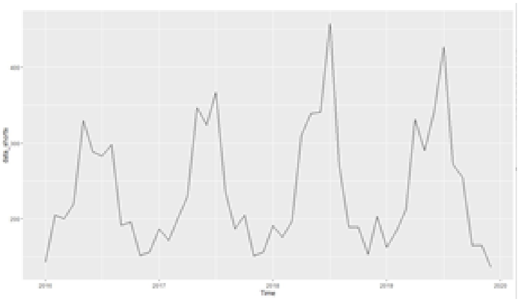

Figure 5: monthly sales shorts

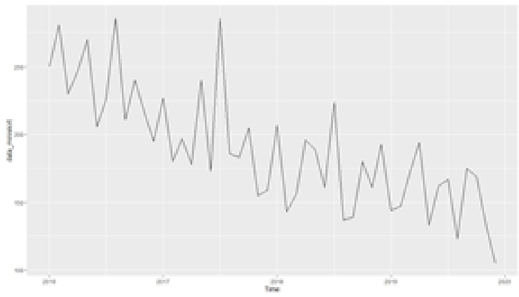

Figure 6: monthly sales miniskirt

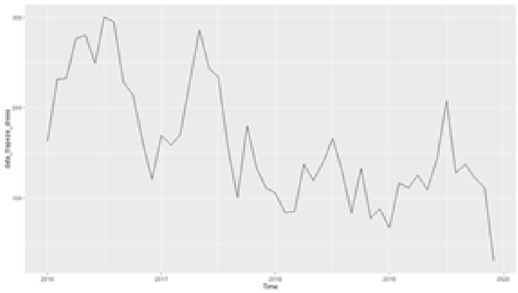

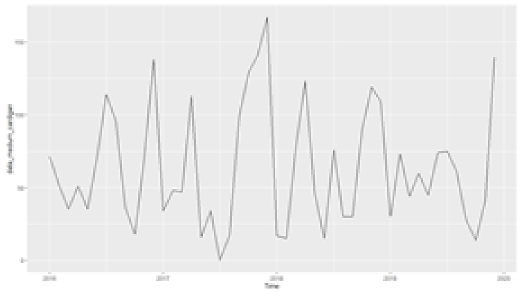

Figure 7: monthly sales trapeze dress      Figure 8: monthly sales medium cardigan

Next, the clothing items that belong to the same fashion style were aggregated: in group one under the name "SUMMER" (Figure 9) there are short sleeves and shorts, in group two called "CHIC" (Figure 10) there are miniskirts and trapeze dresses, in group three called "CASUAL" (Figure 11) are included sleeveless and medium cardigans, and finally in group four called "SPORT" (Figure 12) are tracksuits and long sleeves. Below are graphs of the monthly sales of the following groups:

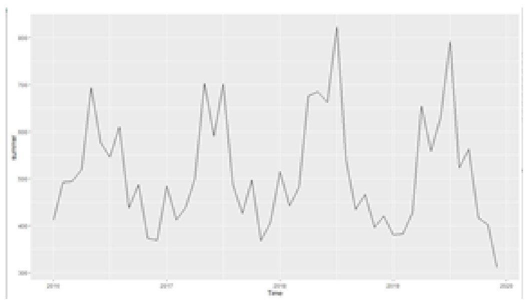

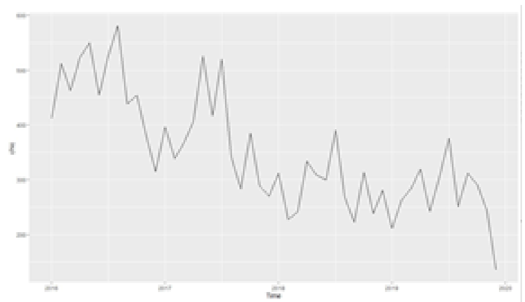

Figure 9: monthly sales summer      Figure 10: monthly sales chic

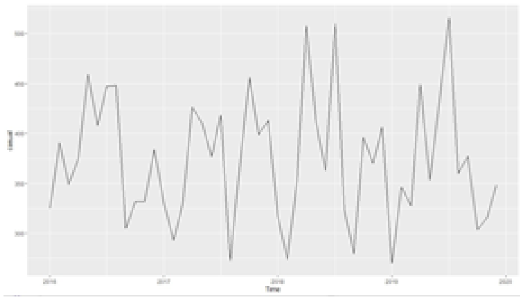

Figure 11: monthly sales casual

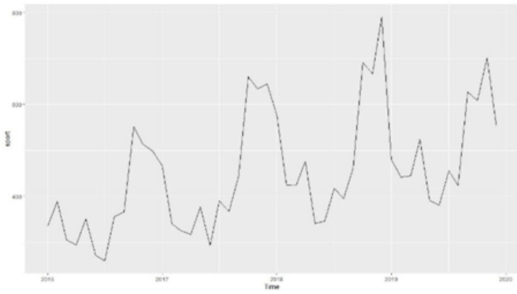

Figure 12: monthly sales sport

**4.2 Seasonality**

The charts autocorrelations function (ACF) have been calculated in order to analyze the characteristics of sales more in depth. These types of graphs are usually used to see graphically how the correlations vary as lag varies. The graph used for this purpose is the “correlogram”.

In Figure 13, which represents the ACF of summer products, it is clear that seasonality prevails. The values vary, but are positive in correspondence of values that allow to characterize an annual periodicity. This means that, the values of a given moment or period of the year are strongly correlated with those of the same instants or periods of previous years. So, it is said that the seasonal component prevails because the phenomenon varies during each year and in a similar way from one year to another.

In Figure 14, which represents the sales characteristics for chic products, the trend prevails because the autocorrelation for small lags tends to be large and positive since observations close in time are also close in size. This graph tends to have positive values that slowly decrease as lag increases. In this case, the present is influenced by the recent past, this from the more remote past and, in general, that the series presents a trend.

In Figure 15, which represents the trend of casual products, the values always fluctuate within a narrow band, this means that the series is not significantly correlated with the delayed series, or that the past does not explain the present and that the variations from one period to another are random, for this reason it is said that the accidental component or stochastic part prevails.

Graph 16, which includes products belonging to the sports category, shows a trend very similar to summer products (Figure 14) where seasonality prevails. As in Figure 14, there are very similar values from year to year. Since these variations are repeated over the years, it is clear that the seasonal component prevails.

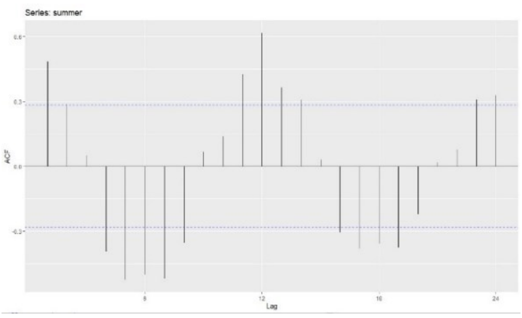


Figure 13: ACF summer

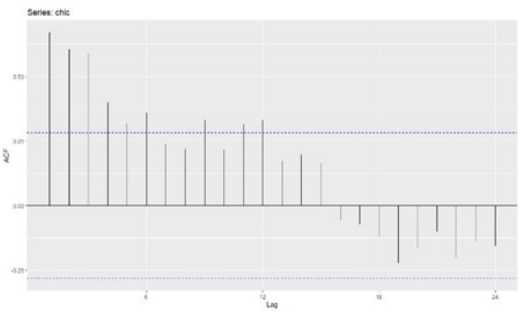


Figure 14: ACF chic

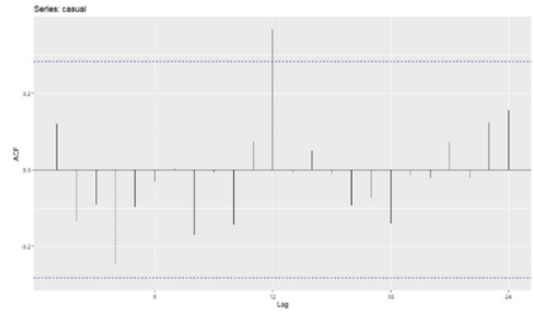


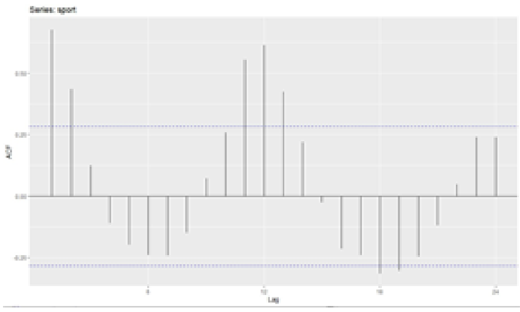


Figure 15: ACF casual

Figure 16: ACF sport

Subsequently, sub-series charts were developed to highlight any hidden trends or seasonality. The sub-series chart brings together the data for each season in small, separate time charts so, it is possible to see the seasonal components. Horizontal blue lines indicate the average monthly value for each month. These types of charts show both the underlying seasonal patterns and the changes in seasonality over time. They are useful for identifying changes that occur within particular seasons.

In Figure 17, it is possible to see the seasonality of summer products, in fact, there is a noticeable increase in values during the summer months.

As for chic products, Figure 18, the sub-series chart shows an almost flat trend during the months of the year. Moreover, it is possible to affirm that for these items there is no seasonality.

The monthly values of the casual products, Figure 19, show an almost constant path with a peak in July.

In Figure 20 it is possible to notice that the average monthly values of sports products change during the winter months therefore these tend to increase from October to December; for this reason, there is seasonality.

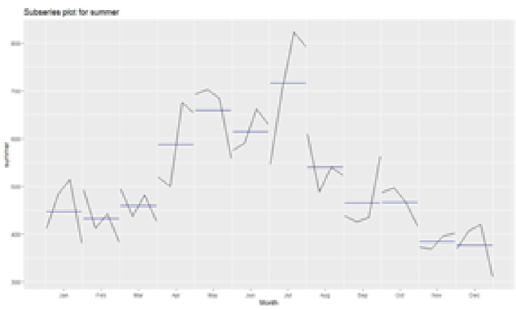


Figure 17: subseries plot for summer

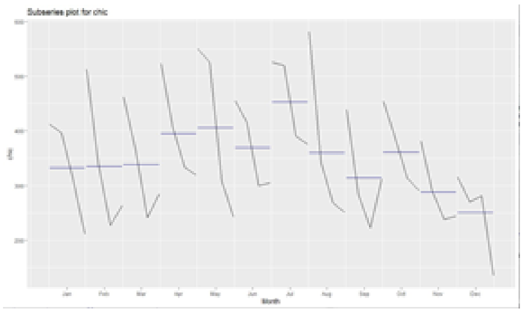


Figure 18: subseries plot for chic

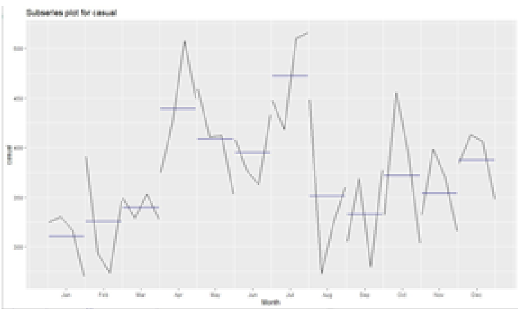


Figure 19: subseries plot for casual

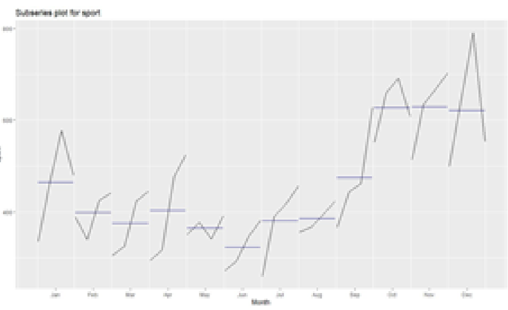


Figure 20: subseries plot for sport

Then the seasonal graphs were made where the data are analysed with respect to the individual seasons. A seasonal chart is able to show more clearly the underlying seasonal component, and it is particularly useful for identifying years when the dynamics of the component change. In these graphs the values of the individual seasons are overlying.

Figure 21, representing sales of summer products, shows a certain seasonality. In particular, it is evident that there is a growth from the month of May to a peak in July. In contrast, the months of November and February show the lowest values.

With regard to Figure 22, it can be seen that over the years there has been a general decline in sales for chic products, in particular the lowest values are in December.

For casual products, Figure 23, it is possible to see that there are positive peaks in the months of April and July, whereas the lowest values emerge in the months between August and September.

In Figure 24 for sports products, it is possible to note a stable trend in the first part of the year with a positive peak from October to December.

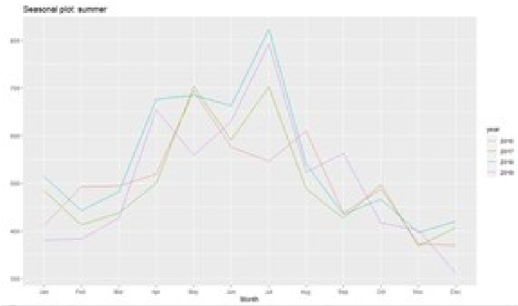


Figure 21: seasonal plot summer

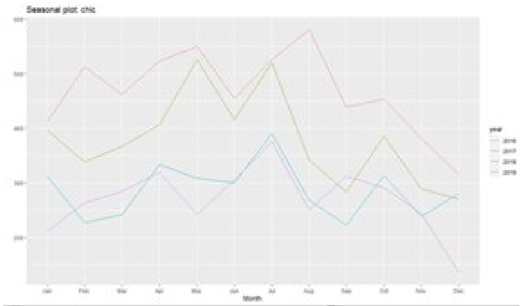


Figure 22: seasonal plot chic

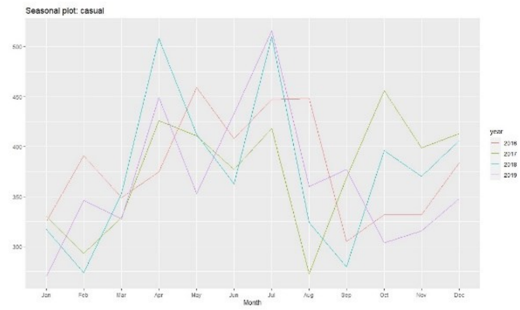


Figure 23: seasonal plot casual

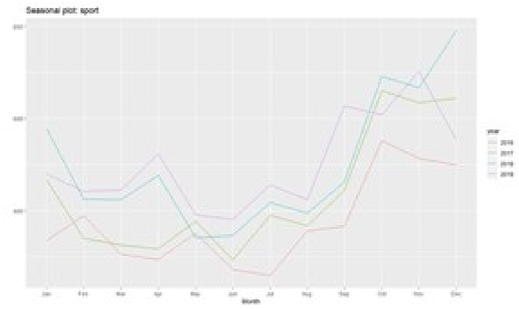


Figure 24: seasonal plot sport

# 5 FORECAST ANALYSIS

## 5.1 Holt-Winters Method

The 'hw' command was used to calculate the forecast for both additive and multiplicative methods for the next 24 months. In the following graphs, the red line identifies the additive method while the blue line is the multiplicative one. In order to assess the accuracy of the forecast, the residual distribution, together with its ACF plot, and the Ljung-Box test (Table 2) are reported. In general, a predictive model has a good performance if the distribution of the residuals resembles a white noise process, following a Gaussian distribution and with values of the ACF within the acceptance region for the null hypothesis of no autocorrelation.

Starting from the prediction of the clothes belonging to the summer category, it can be seen that seasonality is also maintained in the prediction (Figure 25). From the results obtained, it is evident that the distribution of the residuals, both for the additive method (Figure 26) and the multiplicative method (Figure 27), is centred around zero and the values in the ACF plots in both cases never touch the confidence bands. Moreover, the Ljung-Box test (Table 2) shows good p-values close to 1%, with the multiplicative method behaving better than the additive one.

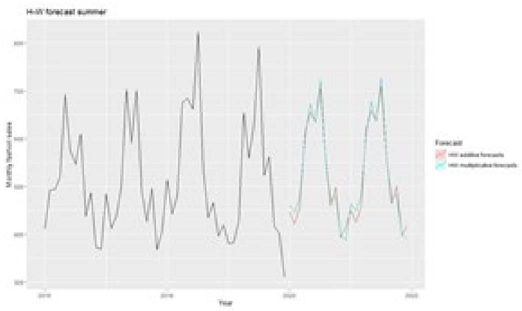


Figure 25: H-W forecast summer

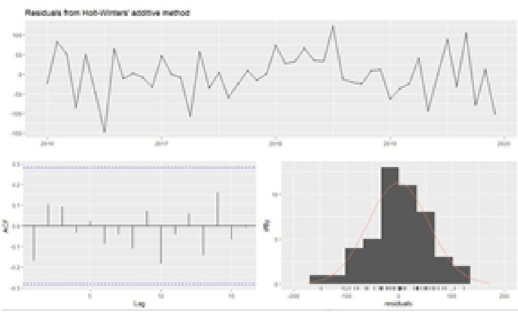


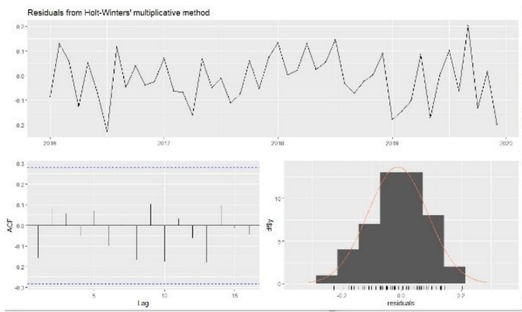


Figure 26: residuals from H-W additive Method

Figure 27: residuals from H-W multiplicative method

The forecast for chic products (Figure 28) shows a very similar situation for both methods with the residuals centred around zero and falling within the confidence interval for the additive (Figure 29) and the multiplicative (Figure 30) methods. The Ljung-Box test (Table 2) shows, in this case, a better result for the multiplicative method.

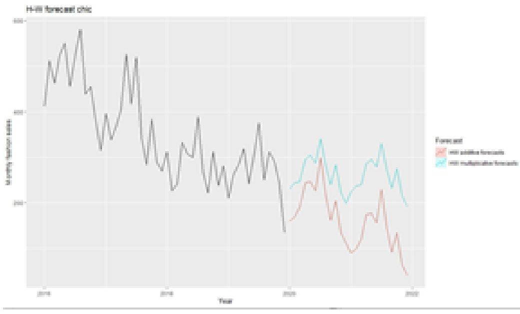


Figure 28: H-W forecast chic

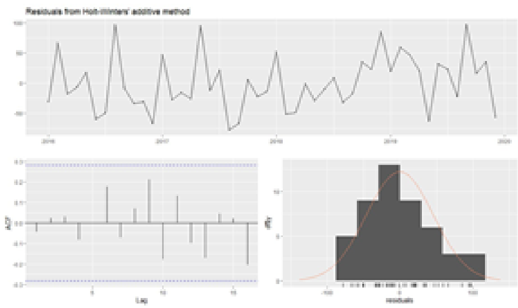

Figure 29: residuals from H-W additive Method

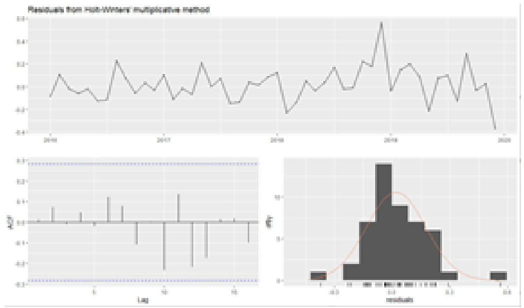

Figure 30: residuals from H-W multiplicative method

The prediction results for casual items (Figure 31), show a very similar situation in both the additive and multiplicative method. Moreover, the distribution of the residuals is more centred around zero in the additive method (Figure 32) than in the multiplicative one (Figure 33) and, when looking at the ACF plots, it is evident that all values fall within the confidence band. The Ljung-Box test (Table 2) reveals low and almost identical p-values in both cases (the multiplicative is slightly lower than the additive).

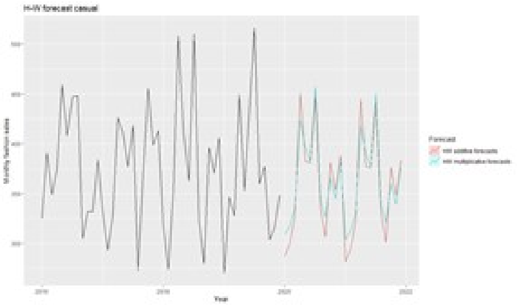

Figure 31: H-W forecast casual

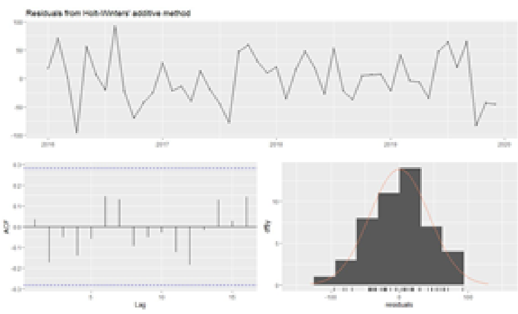

Figure 32: residuals from H-W additive Method

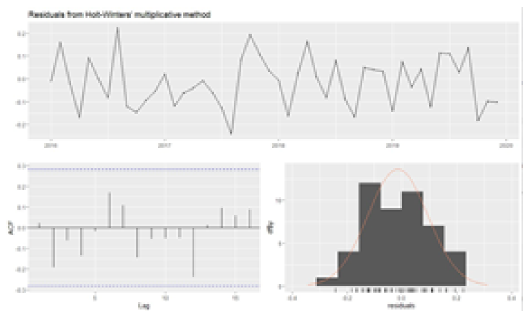

Figure 33: residuals from H-W multiplicative method

The results of the forecast for sports products (Figure 34), like summer products, show that seasonality was maintained in the forecast. Moreover, the forecast shows a better fit with the additive model (Figure 35) than with the multiplicative (Figure 36), which can be seen both from the distributions of the residuals and also from the ACF plots, where in the multiplicative one value for the autocorrelation touches the confidence band but also from the Ljung-Box test

(Table 2), where the p-value for the multiplicative is smaller than for the additive one.

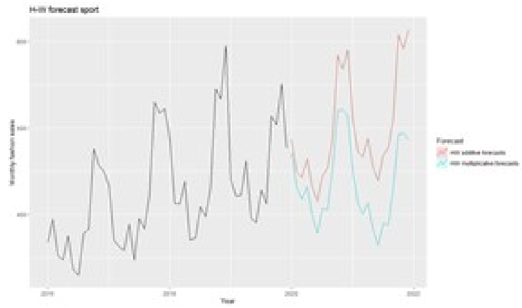


Figure 34: H-W forecast sport

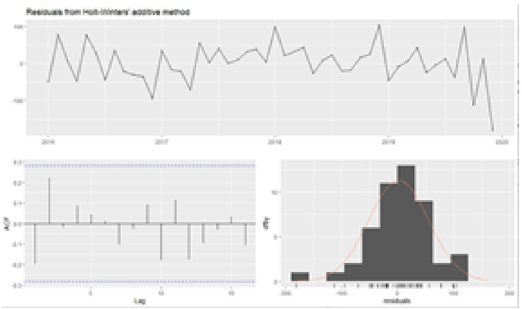


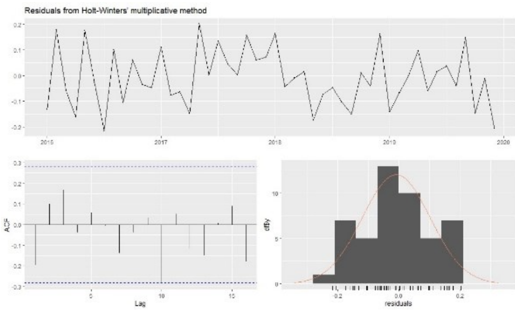


Figure 35: residuals from H-W additive Method

Figure 36: residuals from H-W multiplicative method

Furthermore, looking at the values in table Ljung Box, it can be concluded that any method is not able to eliminate completely the autocorrelation from the residuals, as the Liung Box test always rejects, which is evident from the p-value always being less than 5%. It is therefore worth looking at other methods.

Table 2 Ljung Box test HW

| HW- ADDITIVE | Q* | df | p-value |
|---|---|---|---|
| summer | 10.883 | 3 | 0.01238 |
| chic | 20.136 | 3 | 0.0001591 |
| casual | 17.358 | 3 | 0.0005964 |
| sport | 14.89 | 3 | 0.001913 |

| HW- MULTIPLICATIVE | Q* | df | p-value |
|---|---|---|---|
| summer | 10.614 | 3 | 0.01401 |
| chic | 14.253 | 3 | 0.00258 |
| casual | 17.46 | 3 | 0.0005684 |
| sport | 21.717 | 3 | 7.471e-05 |

Tables 3 and 4 report the results on several metrics commonly used to assess the accuracy of time series forecasting models, such as RMSE (mean square error), MAE (mean absolute error), MPE (mean percentage error), MAPE (mean absolute percentage error), MASE (mean absolute scaled error) and ACF1 (autocorrelation of residuals at lag 1).

In particular, MAPE measures the average percentage deviation of forecasts from actual values, and MPE measures the average deviation of forecasts from actual values as a percentage of actual values. This values are taken into account for identifying the forecast accuracy, in general, lower values indicate better forecasts.

The ETS additive model seems to perform well in terms of MAPE and MPE values in all four products, with MAPE values ranging from 8.619% to 11.913% and MPE values ranging from -1.395% to -0.104%. Summer products have the lowest MAPE value of 8.619%, indicating that the predictions of the ETS additive model for this product are accurate on average. With regard to MPE, summer products also have the lowest MPE value of -1.395%, indicating that the predictions of the ETS additive model for this product have, on average, a greater deviation, of the actual values as a percentage, than the other products.

| ETS-ADDITIVE | ME | RMSE | MAE | MPE | MAPE | MASE | ACF1 |
|---|---|---|---|---|---|---|---|
| summer | -1.178 | 56.923 | 44.233 | -1.395 | 8.619 | 0.801 | -0.170 |
| chic | -0.352 | 45.511 | 37.757 | -1.109 | 11.913 | 0.492 | -0.045 |
| casual | -0.274 | 42.762 | 35.283 | -1.152 | 9.603 | 0.670 | 0.035 |
| sport | 1.796 | 53.150 | 39.614 | -0.104 | 8.804 | 0.560 | -0.197 |

Table 3 – ETS additive forecast accuracy indicators.

The ETS multiplicative, as for additive, seems to perform well in terms of MAPE and MPE values, in fact, MAPE values range from 8.724% to 11.201% and MPE values range from -2.781% to -0.190%. The casual products have the lowest MPE value of -2.781%, indicating that the predictions of the ETS multiplication model for this product had, on average, a much higher deviation from the actual values, as a percentage, than the other products. The products in the summer category had the lowest MAPE value of 8.724%, indicating that the predictions of the ETS multiplicative model for this product were, on average, relatively accurate.

| ETS-MULTIPLICATIVE | ME | RMSE | MAE | MPE | MAPE | MASE | ACF1 |
|---|---|---|---|---|---|---|---|
| summer | -6.422 | 55.108 | 44.043 | -2.374 | 8.724 | 0.798 | -0.164 |
| chic | 3.468 | 44.687 | 35.820 | -0.190 | 11.201 | 0.467 | -0.069 |
| casual | -5.544 | 41.366 | 34.415 | -2.781 | 9.494 | 0.653 | 0.034 |
| sport | -4.801 | 52.208 | 41.321 | -2.144 | 9.332 | 0.584 | -0.156 |

Table 4 – ETS multiplicative forecast accuracy indicators

Comparing the results of the multiplicative method with those of the additive method, the differences in accuracy metrics between the two methods are not substantial. In summary, for both additive and multiplicative methods, the ACF1 values are negative for summer, chic and sport indicating that there is some autocorrelation in the model residuals. Interestingly, the MPE values are all negative so for all categories there is a tendency to underestimate sales, in fact the more the value deviates from 0, the greater the overestimation or underestimation (depending on the positivity or negativity of the value). In particular, in the multiplicative, the model has an MPE value that deviates more from 0 for casual products than for the other categories, suggesting that it tends to underestimate sales more severely for this category while for the additive this is the case for products in the summer category.

## 5.2 ARIMA

The graphs below show the forecasts made using the ARIMA method. The "auto.arima" command was used to automatically find the best order model. The function auto.arima() in R uses a variation of the Hyndman-Khandakar algorithm (Hyndmann & Khandakar, 2008), which combines unit root testing, AICc-minimisation and MLE to obtain an ARIMA model. The arguments of 'auto.arima()' involve many variations of the algorithm. What is described here is the default behaviour. When fitting an ARIMA model to a (non-seasonal) time series data set, the following procedure provides a useful general approach.

With these notions in mind, it is possible to make the analysis of the different product clusters. The analysis starts with the summer products (Figure 37). The forecasts for these products, with the assigned model (0,0) (0,1,0) (Table 5), produce the best results among all items as the residuals (Figure 38) are well centred around zero and the ACF plots, while showing some residual autocorrelation, have all values fall within the acceptance region.

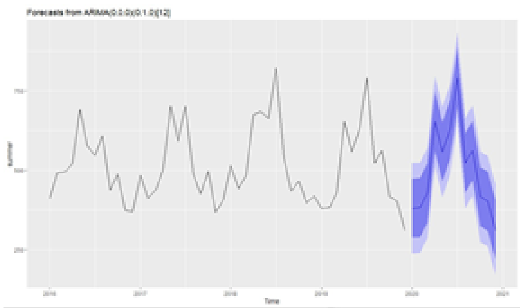

Figure 37: forecasts from ARIMA summer

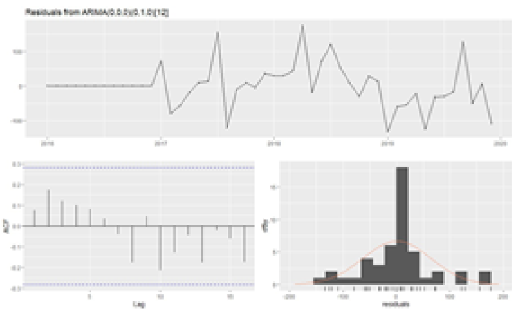

Figure 38: residuals from Arima summer

The chic products (Figure 43), with the model assigned (0,1,1) (0,0,1) (Table 5), show that the residuals are poorly centred, the ACF plots show low autocorrelation over the entire forecast period and all values fall within the confidence band (Figure 44).

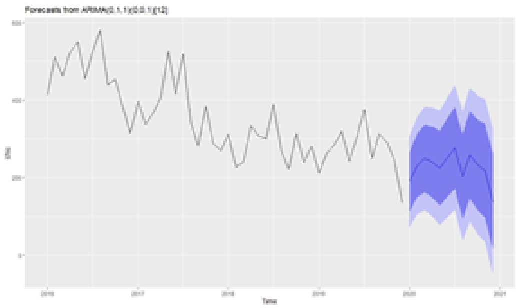

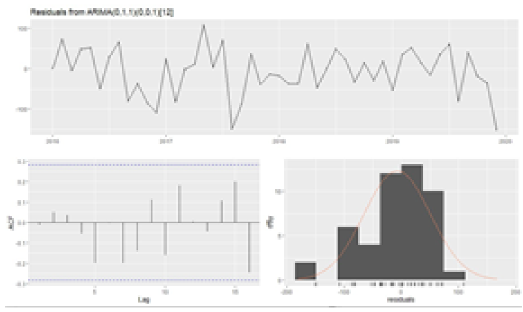

Figure 39: forecasts ARIMA chic

Figure 40: residuals from ARIMA chic

The casual products (Figure 41), with the model assigned (0,0,0) (1,0,0) (Table 5), the residuals (Figure 42) do not look good, in fact, they are not centred and from the ACF graph, we can see that there is a low autocorrelation of the last predicted observations, but all values fall within the confidence band.

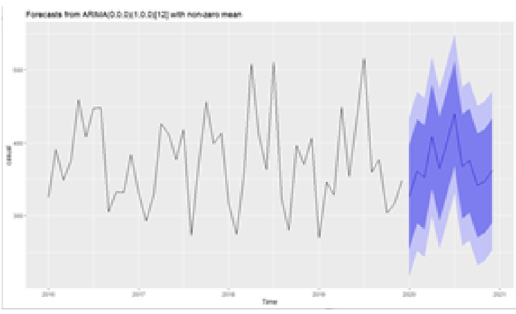

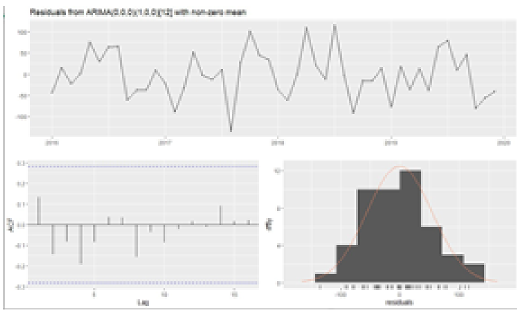

Figure 41: forecast ARIMA casual

Figure 42: Forecast ARIMA casual

The sport products (Figure 43), with the model assigned (1,1,1) (0,1,1) (Table 5), the residuals (Figure 44) do not appear to be centred and from the ACF graph, it is evident that there is low autocorrelation between the values and that all fall within the confidence band.

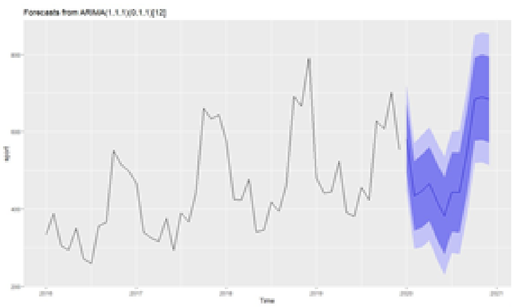


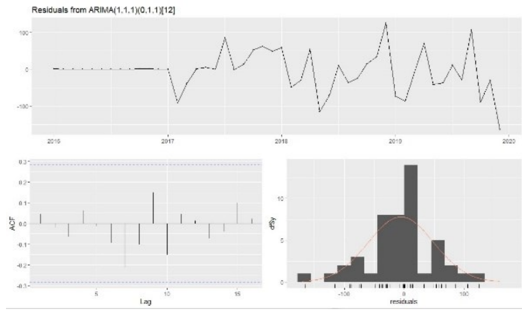


Figure 43: forecast ARIMA sport

Figure 44: forecast ARIMA sport

The Ljung-Box test (Table 5) for all 4 product categories (summer, chic, casual, sport) shows a p-value greater than 0, 05 so the null hypothesis on absence of autocorrelation cannot be rejected.

| | | | | | | | Residuals | | |
|---|---|---|---|---|---|---|---|---|---|
| | model | Sigma^2 | Log likelihood | AIC | AICc | BIC | Q* | df | p-value |
| SUMMER | (0,0,0)(0,1,0) | 5277 | -205.36 | 412.72 | 412.84 | 414.3 | 8.5145 | 10 | 0.5787 |
| CHIC | (0,1,1)(0,0,1) | 3561 | -261.03 | 528.06 | 528.62 | 533.61 | 8.4481 | 8 | 0.391 |
| CASUAL | (0,0,0)(1,0,0) with non-zero mean | 3103 | -261.49 | 528.99 | 529.53 | 534.6 | 6.9792 | 9 | 0.6393 |
| SPORT | (1,1,1)(0,1,1) | 4451 | -200.43 | 408.85 | 410.19 | 415.07 | 7.1872 | 7 | 0.4097 |

Table 5 - ARIMA Ljung box test

Looking at the results on the forecast error indicators (Table 6), the performance of the ARIMA varies between the different product categories. The model seems to perform well for sports, with a low RMSE and MAE, indicating that it can make accurate sales predictions for this category. Looking at the indicators provided, we can see that the MPE values for all clothes are negative, suggesting that the ARIMA model tends to underestimate actual values. The MPE values range from -5.326 for chic products to -1.058 for summer products, with the lowest MPE value for the chic category. As mentioned earlier, the MAPE values give an indication of how accurate the ARIMA model is in predicting the percentage error and a lower MAPE value indicates a more accurate prediction. In this case, it is possible to see that MAPE values range from 7.736 for sport products to 14.946 for chic products, with casual products having a MAPE value of 12.037. Overall, the ARIMA model seems to perform reasonably well for all clothes, with the lowest MPE value for chic clothes and the lowest MAPE value for sports products.

The results suggest that the ARIMA method can be used to make reasonably accurate sales forecasts for certain product categories.

| ARIMA | ME | RMSE | MAE | MPE | MAPE | MASE | ACF1 |
|---|---|---|---|---|---|---|---|
| SUMMER | 0.750 | 62.910 | 41.542 | -1.058 | 8.264 | 0.752 | 0.076 |
| CHIC | -7.453 | 57.779 | 46.057 | -5.326 | 14.946 | 0.601 | -0.012 |
| CASUAL | -0.822 | 54.528 | 43.393 | -2.484 | 12.037 | 0.823 | 0.134 |
| SPORT | -5.902 | 54.473 | 36.858 | -2.325 | 7.736 | 0.521 | 0.045 |

Table 6 - ARIMA forecast accuracy indicators

## 6 CONCLUSIONS

After analysing the dataset using various statistical techniques, including seasonal ARIMA, multiplicative and additive Holt-Winters methods, the results indicate that the chosen methods generated valid and reliable predictions. Each method demonstrated its strengths in predicting specific clusters within the dataset. For example, for the summer cluster, both Holt-Winters methods give more accurate predictions when considering MAPE whereas, for the MPE parameter, the additive fit better for the sport cluster whereas the multiplicative for the chic one. The ARIMA method give better predictions for clothes belonging to the sport and chic groups with reference to the MAPE and MPE parameters.

Furthermore, the MPE and MAPE values suggest that each method tends to underestimate summer sales for the additive method and casual sales for the multiplicative one, whereas ARIMA underestimates sales of chic products. In general, among all three methods, the most one which underestimates the demand is ARIMA method since the difference between forecast and demand is greater.

Comparing the RMSE, MAE and MASE values of the different methods, it can be seen that the ETS methods seem to perform slightly better than the ARIMA method, since the values are slightly lower for most categories. However, the differences in accuracy metrics between the two methods are not substantial so, it is difficult to identify which method is better because the results suggest that each method is better suited for different categories of fashion products.

Furthermore, the conclusions derived from the Ljung-Box test, ACF graphs and the distribution of residuals further support the reliability of the models as these statistical tests adequately capture the underlying relationships in the data and patterns in a way that increases their credibility and usefulness.

The study's insights into the potential applications of these models for forecasting time-series data in various contexts are particularly relevant in the case of the fashion industry. The fashion industry is a very sensitive sector as it is influenced by numerous variables and the continuous change in consumer tastes. For this reason, good demand forecasting is crucial to make companies as profitable as possible, since without it, they carry excessive inventories that are unnecessary and costly, but more importantly, they miss opportunities because they have failed to predict different customer needs, preferences or purchasing intentions.

In conclusion, the results obtained from this study suggest that the chosen forecasting methods, such as Holt-Winters and seasonal ARIMA, are reliable and effective in predicting time series data for fashion analysis. The outcomes from the study provide valuable insights into the potential application of these models in a number of contexts where their use helps to improve supply chain management and brings better results for customers.

## 7 REPLICATION MATERIAL

In this section, I am going to share the link for the dataset and the script, which I created to build in R for the statistical computation and the graphs.

The original dataset, which I changed for my study, is available on Kaggle at the link: https://www.kaggle.com/datasets/konradb/visuelle-complete-dataset, whereas the code I used is written below.

```
# Install required packages (example)
# install.packages("ggplot2", "fpp2", dependencies = TRUE)
library(fpp2, tidyverse)
#leaps model selection
# Load packages
# This package allows us to create publication-quality plots
require(ggplot2, tidyverse)
# Monthly Sales dataset from Kaggle
data = read.csv("monthly fashion sales.csv", head=TRUE)
head(data)
data_long_sleeve = ts(data[,2], frequency=12, start = c(2016,1))
autoplot(data_long_sleeve)
data_short_sleeves = ts(data[,3], frequency=12, start = c(2016,1))
autoplot(data_short_sleeves)
```

```
data_sleeveless = ts(data[,4], frequency=12, start = c(2016,1))

autoplot(data_sleeveless)

data_tracksuit = ts(data[,5], frequency=12, start = c(2016,1))

autoplot(data_tracksuit)

data_shorts = ts(data[,6], frequency=12, start = c(2016,1))

autoplot(data_shorts)

data_miniskirt = ts(data[,7], frequency=12, start = c(2016,1))

autoplot(data_miniskirt)

data_trapeze_dress = ts(data[,8], frequency=12, start = c(2016,1))

autoplot(data_trapeze_dress)

data_medium_cardigan = ts(data[,9], frequency=12, start = c(2016,1))

autoplot(data_medium_cardigan)

# look for seasonality and trends #

ggAcf(data_long_sleeve)

ggAcf(data_short_sleeves)

ggAcf(data_sleeveless)

ggAcf(data_tracksuit)

ggAcf(data_shorts)

ggAcf(data_miniskirt)

ggAcf(data_trapeze_dress)

ggAcf(data_medium_cardigan)
```

```
# group prodoct's characteristics #
data_short_sleeves + data_shorts
summer= data_short_sleeves + data_shorts
autoplot(summer)
ggAcf(summer)
data_miniskirt + data_trapeze_dress
chic = data_miniskirt + data_trapeze_dress
autoplot(chic)
ggAcf(chic)
data_sleeveless + data_medium_cardigan
casual = data_sleeveless + data_medium_cardigan
autoplot(casual)
ggAcf(casual)
data_long_sleeve + data_tracksuit
sport = data_long_sleeve + data_tracksuit
autoplot(sport)
ggAcf(sport)
# sales curve for each year #
ggseasonplot(summer)
ggseasonplot(chic)
ggseasonplot(casual)
```

```
ggseasonplot(sport)
#we look for the averges and the possible trends
ggsubseriesplot(summer)
sub_plot_su <- ggsubseriesplot(summer)
plot(sub_plot_su) + ggtitle("Subseries plot for summer")
ggsubseriesplot(chic)
sub_plot_ch <- ggsubseriesplot(chic)
plot(sub_plot_ch) + ggtitle("Subseries plot for chic")
ggsubseriesplot(casual)
sub_plot_ca <- ggsubseriesplot(casual)
plot(sub_plot_ca) + ggtitle("Subseries plot for casual")
ggsubseriesplot(sport)
sub_plot_sp <- ggsubseriesplot(sport)
plot(sub_plot_sp) + ggtitle("Subseries plot for sport")
# holt-winters chic
chic <- window(chic,start=2016)
fit1 <- hw(chic,seasonal="additive")
fit2 <- hw(chic,seasonal="multiplicative")
autoplot(chic) +
  autolayer(fit1, series="HW additive forecasts", PI=FALSE) +
  autolayer(fit2, series="HW multiplicative forecasts",
```

```
          PI=FALSE) +
  xlab("Year") +
  ylab("Monthly fashion sales") +
  ggtitle("H-W forecast chic") +
  guides(colour=guide_legend(title="Forecast"))
checkresiduals(fit1)
checkresiduals(fit2)
# holt-winters summer
summer <- window(summer,start=2015)
fit3 <- hw(summer,seasonal="additive")
fit4 <- hw(summer,seasonal="multiplicative")
autoplot(summer) +
  autolayer(fit3, series="HW additive forecasts", PI=FALSE) +
  autolayer(fit4, series="HW multiplicative forecasts",
          PI=FALSE) +
  xlab("Year") +
  ylab("Monthly fashion sales") +
  ggtitle("H-W forecast summer") +
  guides(colour=guide_legend(title="Forecast"))
checkresiduals(fit3)
checkresiduals(fit4)
```

```
# holt-winters casual
casual <- window(casual,start=2016)
fit5 <- hw(casual,seasonal="additive")
fit6 <- hw(casual,seasonal="multiplicative")
autoplot(casual) +
  autolayer(fit5, series="HW additive forecasts", PI=FALSE) +
  autolayer(fit6, series="HW multiplicative forecasts",
            PI=FALSE) +
  xlab("Year") +
  ylab("Monthly fashion sales") +
  ggtitle("H-W forecast casual") +
  guides(colour=guide_legend(title="Forecast"))
checkresiduals(fit5)
checkresiduals(fit6)
# holt-winters sport
sport <- window(sport,start=2016)
fit7 <- hw(sport,seasonal="additive")
fit8 <- hw(sport,seasonal="multiplicative")
autoplot(sport) +
  autolayer(fit7, series="HW additive forecasts", PI=FALSE) +
  autolayer(fit8, series="HW multiplicative forecasts",
```

```
          PI=FALSE) +
  xlab("Year") +
  ylab("Monthly fashion sales") +
  ggtitle("H-W forecast sport") +
  guides(colour=guide_legend(title="Forecast"))
checkresiduals(fit7)
checkresiduals(fit8)
#ACCURACY ETS
accuracy(fit1)
accuracy(fit2)
accuracy(fit3)
accuracy(fit4)
accuracy(fit5)
accuracy(fit6)
accuracy(fit7)
accuracy(fit8)
# ARIMA
# SEASONAL -->
auto.arima(summer)
fit.arima.season.SM <- auto.arima(summer , seasonal = TRUE)
fit.arima.season.SM %>% forecast(h=12) %>% autoplot(include=80)
```

```
checkresiduals(fit.arima.season.SM)

auto.arima(chic)

fit.arima.season.CH <- auto.arima(chic , seasonal = TRUE)

fit.arima.season.CH %>% forecast(h=12) %>% autoplot(include=80)

checkresiduals(fit.arima.season.CH)

auto.arima(casual)

fit.arima.season.CA <- auto.arima(casual , seasonal = TRUE)

fit.arima.season.CA %>% forecast(h=12) %>% autoplot(include=80)

checkresiduals(fit.arima.season.CA)

auto.arima(sport)

fit.arima.season.SP <- auto.arima(sport , seasonal = TRUE)

fit.arima.season.SP %>% forecast(h=12) %>% autoplot(include=80)

checkresiduals(fit.arima.season.SP)

#ACCURACY ARIMA

accuracy(fit.arima.season.SM)

accuracy(fit.arima.season.CH)

accuracy(fit.arima.season.CA)

accuracy(fit.arima.season.SP)
```

**REFERENCES**

Forecasting: Principles and Practice Rob J Hyndman and George Athanasopoulos Monash University, Australia (2018).

Fildes, Robert, Shaohui Ma, and Stephan Kolassa. "Retail forecasting: Research and practice." *International Journal of Forecasting* 38.4 (2022): 1283-1318.

Frank, Celia, et al. "Forecasting women's apparel sales using mathematical modeling." *International Journal of Clothing Science and Technology* 15.2 (2003): 107-125.

Sébastien, Thomassey, Happiette Michel, and Castelain Jean-Marie. "Mean-term textile sales forecasting using families and items classification." *Studies in Informatics and Control* 12.1 (2003): 41-52.

Zadeh, Lotfi A. "Fuzzy sets." *Information and control* 8.3 (1965): 338-353.

Sztandera, Les M., Celia Frank, and Balaji Vemulapali. "Predicting women's apparel sales by soft computing." *Artificial Intelligence and Soft Computing-ICAISC 2004: 7th International Conference, Zakopane, Poland, June 7-11, 2004. Proceedings 7*. Springer Berlin Heidelberg, 2004.

Sztandera, Les M., Celia Frank, and Balaji Vemulapali. "Prediction of women's apparel sales using soft computing methods." *Knowledge-Based Intelligent Information and Engineering Systems: 8th International Conference, KES 2004, Wellington, New Zealand, September 20-25, 2004, Proceedings, Part III 8*. Springer Berlin Heidelberg, 2004.

Box, George EP, Gwilym M. Jenkins, and Gregory C. Reinsel. "Forecasting." *Time Series Analysis* (2008): 137-191.

Chang, Pei-Chann, Yen-Wen Wang, and Chen-Hao Liu. "The development of a weighted evolving fuzzy neural network for PCB sales forecasting." *Expert Systems with Applications* 32.1 (2007): 86-96.

Chang, Pei-Chann, Yen-Wen Wang, and Chi-Yang Tsai. "Evolving neural network for printed circuit board sales forecasting." *Expert Systems with Applications* 29.1 (2005): 83-92.

Hamzaçebi, Coşkun, Diyar Akay, and Fevzi Kutay. "Comparison of direct and iterative artificial neural network forecast approaches in multi-periodic time series forecasting." *Expert systems with applications* 36.2 (2009): 3839-3844.

Ling, Sai-Ho, et al. "Short-term electric load forecasting based on a neural fuzzy network." *IEEE Transactions on industrial electronics* 50.6 (2003): 1305-1316.

Hansen, James V., and Ray D. Nelson. "Neural networks and traditional time series methods: a synergistic combination in state economic forecasts." *IEEE transactions on Neural Networks* 8.4 (1997): 863-873.

Huang, Guang-Bin, Qin-Yu Zhu, and Chee-Kheong Siew. "Extreme learning machine: theory and applications." *Neurocomputing* 70.1-3 (2006a): 489-501.

Huang, Guang-Bin, Qin-Yu Zhu, and Chee Kheong Siew. "Real-time learning capability of neural networks." *IEEE Trans. Neural Networks* 17.4 (2006b): 863-878.

Sun, Zhan-Li, Kin-Fan Au, and Tsan-Ming Choi. "A neuro-fuzzy inference system through integration of fuzzy logic and extreme learning machines." *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37.5 (2007): 1321-1331.

Holt, C. E. "Forecasting seasonals and trends by exponentially weighted averages (ONR Memorandum No. 52)." *Carnegie Institute of Technology, Pittsburgh USA* 10 (1957).

Winters, Peter R. "Forecasting sales by exponentially weighted moving averages." *Management science* 6.3 (1960): 324-342.

Hyndman, Rob J., and Yeasmin Khandakar. "Automatic time series forecasting: the forecast package for R." *Journal of statistical software* 27 (2008): 1-22.

Swaminathan, Kritika, and Rakesh Venkitasubramony. "Demand forecasting for fashion products: A systematic review." *International Journal of Forecasting* (2023).

Xia, Min, et al. "Fashion retailing forecasting based on extreme learning machine with adaptive metrics of inputs." *Knowledge-Based Systems* 36 (2012): 253-259.

Xia, Min, and Wai Keung Wong. "A seasonal discrete grey forecasting model for fashion retailing." *Knowledge-Based Systems* 57 (2014): 119-126.

Stahl, Christina, Nikolai Stein, and Christoph M. Flath. "Analytics applications in fashion supply chain management—A review of literature and practice." *IEEE Transactions on Engineering Management* (2021).

Liu, Na, et al. "Sales forecasting for fashion retailing service industry: a review." *Mathematical Problems in Engineering* 2013 (2013).

Choi, Tsan-Ming, et al. "Fast fashion sales forecasting with limited data and time." *Decision Support Systems* 59 (2014): 84-92.

Yu, Yong, Tsan-Ming Choi, and Chi-Leung Hui. "An intelligent fast sales forecasting model for fashion products." *Expert Systems with Applications* 38.6 (2011): 7373-7379.

Guo, Z. X., Wai Keung Wong, and Min Li. "A multivariate intelligent decision-making model for retail sales forecasting." *Decision Support Systems* 55.1 (2013): 247-255.

**ACKNOWLEDGMENTS**

I would like to dedicate this final space to all those who supported me during my university career and during the writing of my thesis.

First of all, I would like to thank my teacher, Claudia Pigini, who was always very helpful, followed me and helped me for the whole time. Her support and helpfulness made all the difference and made the work more enjoyable to do.

Next, I would like to thank my family, who supported me throughout my university career, stood by me during difficult times and rejoiced with me at every success. Their support was crucial, despite the various problems, not only scholastically speaking, they gave me the strength to go on, not to get down on myself and to focus only on my studies.

I would like to continue by thanking my second family, who, although they have only recently entered my life, supported me, believed in me and were always there for me when I needed them. They cheered with me for every achievement and helped, advised and supported me when I was in difficulty.

I would like to thank my boyfriend who was an important part in the writing of the paper, helped me throughout these months and made me feel his presence and help even during the time away. His support was crucial both practically, in researching, studying the material, carrying out the project and writing, but above all, his closeness gave me the strength to think that despite everything I would make it through the last and dreaded exam and complete my thesis. Thanks to his

support and strength, I am only a few steps away from finishing my university degree.

My thanks go to my colleagues, those who were closest to me and with whom I spent most of my time between highs and lows, exam preparation, group work, laughter and late-night video calls to finish assigned projects.

Finally, I would like to thank all the professors on the course who trained me and enabled me to acquire skills and knowledge that I hope to put into practice in the workplace soon.