

Adaptive Mechanisms in an Airline Ticket Demand Forecasting System

Silvia Riedel^{*,**} and Bogdan Gabrys^{**}

* Lufthansa Systems Berlin GmbH, Fritschstrasse 27-28, 10585 Berlin Germany

Phone: +49-30-34007141, Fax: +49-30-34007100

email: silvia.riedel@lhsystems.com

** Computational Intelligence Research Group, School of Design, Engineering & Computing

Bournemouth University, Poole House, Talbot Campus

Poole, BH12 5BB United Kingdom

Phone: +44-1202-595298, Fax: +44-1202-595314

email: bgabrys@bournemouth.ac.uk

ABSTRACT: Adaptivity is a very important feature for industrial forecast systems. In the airline industry, a reliable forecasting of a demand for tickets at different fare levels forms a crucial step in a global optimization process, the objective of which is to sell a restricted number of available seats in a plane with a maximized revenue. Due to continuously changing demand caused by seasonality, special events like holidays or fairs, changes in the flight schedules or changes of the political or cultural situation of a country, there is a need for robust, adaptive forecasting techniques able to cope with such changes. In this paper an overview of various adaptive mechanisms used in the new forecasting system of the Lufthansa Airline is presented.

KEYWORDS: Adaptive Forecasting, Airline, Revenue Management, Exponential Smoothing

1 INTRODUCTION

The product of the airline industry are tickets offered with different booking conditions and for different levels of comfort. To maximize revenue the priority is given to high revenue booking classes. Capacity must be protected for high revenue passengers usually arriving shortly before a plane departure. Based on the size of the protected capacity, the capacity of low revenue classes needed to fill up the aircraft can be determined. Therefore, the central question of revenue management is: How much of the overall capacity should be made available for low-yield customers? or in other words: How much space should be reserved for the high-yield segment?

To answer this question, the following technical components are used: a) an inventory to control capacity; b) a forecasting for assessing the demand in advance; and c) an optimization to maximize the revenue by capacity control.

While in this paper the focus is placed on forecasting of the ticket demand, more detailed information about all revenue management components can be found in [4].

The forecast component has to produce predictions for each fare level and departure date. The ticket demand at a fare level (e.g. per departure date) can be modelled as a time series. So formally one can say that given the time series (b_t) , $t = 1..n$ denoting the ticket bookings (in the following representing the demand) per departure date with departure date= n representing the most recent departure, the general problem is to forecast the demand in (b_{n+h}) , $h \in \mathcal{N} \geq 1$. An example of the demand values and a one step ($h = 1$) ahead forecast are shown in Figure 1.

Since a large number of techniques for forecasting time series can be found in the literature [2],[6],[8], [9],[10],[12], there is seemingly a big pool of potential forecasting models to choose from.

Unfortunately, in practice only few methods have been found to produce adequate forecasts because of the structure and quality of the existing data. For our application the world is changing so quickly that in general only a small number of historical data are available and frequently a number of relevant values are missing. Multiple studies on this topic have shown that for our data the simple and robust time series forecasting models like simple average, different versions of exponential smoothing or regression models are significantly better than a number of well known more sophisticated methods. The reason for this lays in the simple methods' ability to make adequate forecasts even on a small number of historical data and to be able to adapt more quickly to new situations.

Some of the potential problems with the ticket demand data and forecasting are illustrated in Figure 2 which shows a typical demand changes per departure date. One of the common problems illustrated in Figure 1 and Figure 2 is that of predicting small numbers which results from the very fine level on which the forecasts have to be performed. On this level, the data are

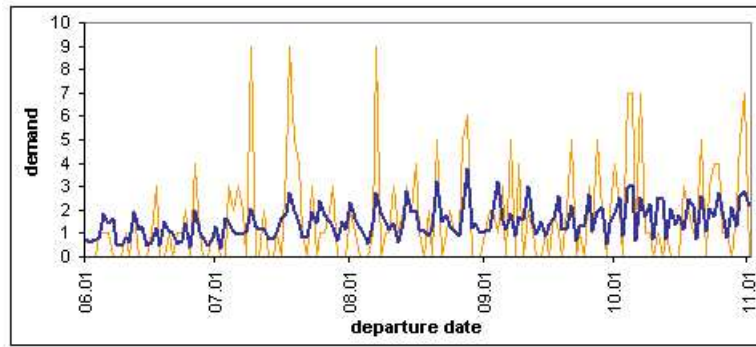


Figure 1: Example of the demand values per departure date (thin line) with one step ($h=1$) ahead forecasts (thick line).

extremely noisy and they frequently exhibit multiple structural breaks. These structural breaks in the time series data reflect the changes in booking behaviour caused by seasonal changes, special events like holidays or fairs, changes in the flight schedules of both the airlines for which the predictions are made and the competitors, or changes of the political or cultural situation of a country. All these changes have to be handled in the forecasting process and are the focus of the following overview of the adaptation mechanisms used within the new forecasting system of the Lufthansa Airline.




Figure 2: Bookings (demand) per departure date (thin line) with seven days average (thick line).

The remaining of the paper is organised as follows. Section 2 first describes the abstract concepts of attractivity and changes to attractivity forming a general framework for adaptation. This is followed by description of different approaches to adaptation to: the flight specific demand, the season and special events, schedule changes and changes to market situation. In section 3 experimental results comparing the performance of various base forecasting techniques and illustrating the impact of adaptation on the forecasting quality are discussed. Finally, conclusions and further research will be presented in the last section.

2 MODELS OF ADAPTATION¹

Generally there are two principal options in order to adapt to new situations. One is to adapt the historical data to the new situation and by doing so producing a direct forecast based on the adapted historical data. The other is to produce forecasts based on the historical data and then to adapt the forecasts to the new situation. We have decided for a combination of both. To handle the needs for adaptation in a general sense, we have built a general model of attractivity and attractivity changes which will be described first. However, it is also necessary to treat certain types of changes in a specialized manner which will be explained in the subsequent sections.

2.1 THE MODEL OF ATTRACTIVITY AND ATTRACTIVITY CHANGES

The general model of adaptation used within the Lufthansa Airline forecasting system is based on the abstract terms of attractivity, attractivity changes and short term influences. The proposed model assumes that the changes requiring adaptation can

¹Please note that due to the commercial sensitivity of the covered material both the data and the description of the methods in the sections 2 and 3 could not be described in any greater detail.

be categorised into two major groups: *permanent changes* like market changes or long term schedule changes and *short term changes* like seasonal behaviour, events or schedule changes influencing only some departures.

The attractivity a_n is the number of bookings at the departure date n if there would be no random noise, no flight specific behaviour and no quickly changing influences like season, events and short time schedule changes in the data. In other words the attractivity represents the stable world behaviour.

Figure 3 shows the bookings together with the attractivity estimate.




Figure 3: Demand per departure date with an estimate of attractivity. Between June and August the level of the raw data changes significantly. This effect is modelled as an attractivity change. The peaks in March and April are too short to represent attractivity changes and are modelled as short term influences.

An attractivity change $c_n = a_n - a_{n-1}$ is a change of the attractivity compared to the previous departure dates. In our example attractivity changes can be seen between the months of June and August.

A short term influence t_n is the influence of the bookings caused by influences like seasonal behaviour, events or short term schedule changes. In figure 4 the short term influences are visualized together with the attractivity.




Figure 4: Demand per departure date with attractivity and short term influence estimation representing an estimation for the general development of the demand.

Attractivity changes are represented as a sum of changes caused by different changes like season, events, flight schedule changes etc. The changes can be represented as absolute changes a_n^j (also used for initialization purposes) or relative changes r_n^k . The attractivity change is then defined by

$$c_n = (a_{n-1} + \sum_j a_n^j) * (1 + \sum_k r_n^k) - a_{n-1} \quad (1)$$

The same approach is used to model short term influences.

As already indicated earlier, there are two principal approaches to adaptation to changes in our application: to adapt the historical data or the forecasts themselves. Our approach is using both adaptations. The adaptation to attractivity changes lying in the future or to short term influences is carried out as a task of adaptation of the basic forecast. It does not influence the historical data and the production of the basic forecasts. The adaptation to historical attractivity changes is performed on the historical data. The historical data are adapted to the current situation to represent a stable basis for the basic forecasts. Short term influences are removed from the historical data. The advantage of this approach is that the adapted historical data represent

always a history matching to the most recent departure and that it can then be used to produce all basic forecasts without any other need for adaptation. Another advantage is that the repeated short term influences are isolated from the historical data and can therefore be also learnt for future adaptations.

Beside the adaptation to attractivity changes and short term influences, there is a possibility to use the the most recent bookings for a future departure for the estimation of the flight specific behaviour, which is interpreted as noise in the time series models. This type of adaptation is described in the following subsection.

2.2 ADAPTATION TO THE FLIGHT SPECIFIC DEMAND

A more detailed view of the bookings is the visualization of the so called booking curves b_t^d (Figure 5) where d denotes the days prior to the departure date showing the development of the number of bookings related to the days before the departure for one single departure date t .

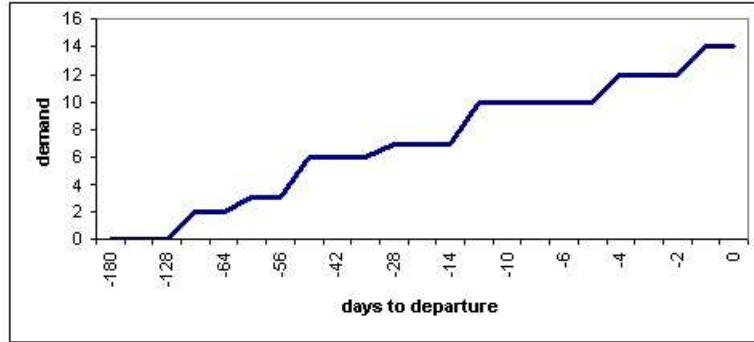


Figure 5: Bookings(demand) per days to departure.

Given a process date n (date to calculate forecasts), a departure date $n + h$ with $h > 0$ (the departure is in the future) and the days to departure h , all data (b_t^d) with $t + d \leq n$ are already known and can theoretically be used to forecast b_{n+h}^0 (the number of bookings at the departure time for the departure date $n + h$).

If a forecast $\widehat{b_{n+h}^0}$ is made for b_{n+h}^0 using time series models on (b_t^0) , $t = 1..n$, the information about (b_{n+h}^d) , $d \geq h$ which represent the booking development of the flight to be forecast until the given process date n is not used to produce the forecast. Figure 6 shows the bookings at the departure together with the bookings 7 days before the departure. A high correlation between the number of already existing bookings with the number of bookings at the departure is easy to see.




Figure 6: Demand per departure date (thin line) with the demand 1 week before the departure (thick line).

The following adaptation of the forecast $\widehat{b_{n+h}^0}$ is done: Even if the booking value b_{n+h}^d is already known, we also produce forecasts $\widehat{b_{n+h}^d}$ for b_{n+h}^d using the same time series model and based on the same departures as for b_{n+h}^0 to extract the flight specific demand from the expected demand.

Then the forecast $\widehat{b_{n+h}^0}$ can be adapted to the already existing flight specific demand by

$$\widetilde{b_{n+h}^0} = \widehat{b_{n+h}^0} + b_{n+h}^d - \widehat{b_{n+h}^d}. \quad (2)$$

Figure 7 shows the adapted forecasts.




Figure 7: Demand per departure date (thin line) with a forecast adapted to the flight specific demand (thick line).

2.3 ADAPTATION TO THE SEASON AND SPECIAL EVENTS

There is not only a strong seasonal behaviour in the booking data, special events like fairs, holidays, cultural or political events or strikes are influencing the data, too. There are a number of known and documented approaches to handle seasonal behaviour in time series (see e.g. [5],[7],[10]). These approaches could be used and produce better forecasts compared to the case when no seasonal model is used. However, in our application not only the bookings at a departure b_{n+h}^0 are influenced, but also the developments of the booking curves b_{n+h}^d should be adapted to the seasonal behaviour.

The seasonal behaviour can be explained with the general rule that if the demand is growing, the bookings are arriving earlier. Figure 8 shows typical booking curves in the high and low season.

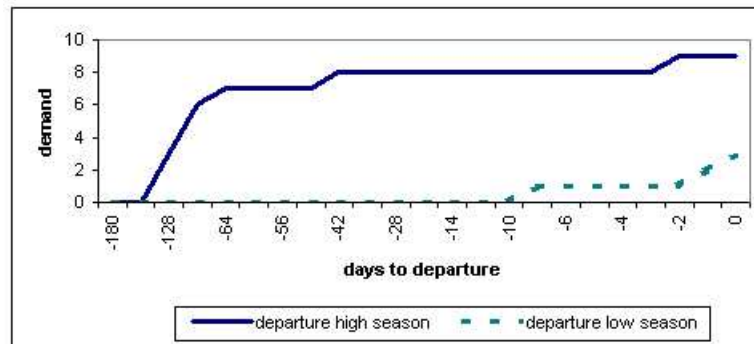


Figure 8: A typical demand for high season and low season. It can be seen that in a high season the bookings are arriving earlier before the departure.

In our model for processing seasonality and event related information first the bookings are cleaned in order to apply the time series models on data without seasonal behaviour interpreted as a short time influence. This is followed by using the extracted seasonality information to produce forecasts as a part of the adaptation of the forecasts to short time influences. And finally the quality of the seasonality information is assessed and retained for future references.

One of the indicators used in the process of extracting current seasonal influences is the seasonal behaviour of the data during the previous years which is stored in separate profiles. However, the past seasonal data on their own are usually not accurate enough for predicting the current seasonal behaviour. It is therefore first supplemented by the newest booking data.

There again is a need for adaptivity of a different kind in order to ensure forecast stability and the process could be summarised in the following four steps:

- seasonal information: extract the current year's seasonal and event information by taking into account the already existing bookings b_{n+h}^d
- historical data: separate the season and event influences from the historical data
- forecast adaptation : adapt the time series forecasts to the season and the existing events
- learning : learn historical seasonal and event behaviour

Let (s_t) be a time series indicating the seasonal information per departure date as a factor related to the yearly average (1 means no seasonal deviation, 0.5 means low season with only half of the normal demand etc.). The values (s_t) are a weighted sum of the seasonal behaviour of the previous years and the seasonal information contained in the current booking values.

The model of adaptation of the forecasts to the season is carried out with forecasts for every $\widehat{b_{n+h}^d}$ based on the cleaned data (i.e. with removed influences of estimated seasonal behaviour and events). The basis for the adaptation of the forecasts are the following two simple models:

- model 1: all seasonal bookings are present (or are missing) on the first day of the booking period
- model 2: all seasonal influences arrive at the same time as the "normal" bookings

The adapted forecast $\widetilde{b_{n+h}^d}$ is a weighted linear combination of the two models

$$\widetilde{b_{n+h}^d} = w_1 * (\widehat{b_{n+h}^d} + \widehat{b_{n+h}^0} * (s_{n+h} - 1)) + w_2 * (\widehat{b_{n+h}^d} * s_{n+h}) \quad (3)$$

The weighting factors depend on the the number of days to the departure and the season and event information. Figure 9 illustrates the adaptation of the basic forecasts to seasonal information.




Figure 9: Adaptation to the season. The figure shows the adaptation based on two adaptation models (additive adaptation and multiplicative adaptation). A weighted sum of these two models gives the adapted forecast. The adapted forecast is shown for the high and low season.

2.4 ADAPTATION TO SCHEDULE CHANGES

Flight schedule changes are a routine part of the business process at Lufhansa. At least two times per year a new flight schedule plan is published. Frequently flights are changed also between these two schedule plan releases.

In general flight schedule changes can be of the following types: a) a flight is cancelled; b) a new flight is introduced; c) a day of week frequency of a flight is changed; and d) a departure/arrival time is changed. As in the previously discussed cases all these changes have to be taken into account and adapted to when producing the demand forecasts.

If the schedule attributes of a flight change, the adopted approach is to look to the "adjacent" flights in order to estimate the new situation. In that way information about the initialization of a new flight scheduled for an additional day of week can be retrieved from the booking behaviour in the already existing flights scheduled for other days of the week. For that purpose an abstract definition of neighborhood between flights has been adopted. This neighborhood can then be used to approximate missing historical values or the changes of the historical values in comparison to the new situation.

Figure 10 shows the above described procedure as an example for the initialization of a new flight.

The decision to interpret a flight schedule change as an attractivity change or a short term influence depends on the number of departures which are influenced by the change. If only some departures are influenced by the schedule change, it is not worth adapting all the historical data. Initializations for new flights can be interpreted as short term influences, too.




Figure 10: Adaptation to a new flight. The figure shows an example of a new flight arriving on a day where two other flights are already existing on the same routing. The new flight is initialized by a weighted sum of the history for the other flights.

2.5 ADAPTATION TO CHANGES OF THE MARKET SITUATION

Market changes can be permanent changes of the political situation of a country, changes related to competitive airlines or general changes in the passenger behaviour. Market changes lead to structural breaks in the time series. As the changes are permanent, they are interpreted as attractivity changes, not as short term influences. In a lot of situations market changes have an adaptation period, i.e. the passenger behaviour is adapting with a various speed to the new situation.

A market change m_t can be defined by a start date t_0 , a date where the adaptation period is finished t_1 and a factor m representing the influence of the market change on the number of bookings at the departure.

The market change m_t is then defined by

$$m_t = \frac{\max(t - t_0, t_1 - t_0)}{t_1 - t_0} * m \quad (4)$$

The adaptation can then be performed as follows:

$$\widetilde{b}_{n+h}^0 = \widehat{b}_{n+h}^0 * m_{n+h} \quad (5)$$

As a special case of attractivity changes, the adaptation of the historical data is performed to the influence m_n of the current departure date. Market changes in the future are realized by an adaptation of the forecasts.

Figure 11 illustrates how the data is adapted to the new situation.

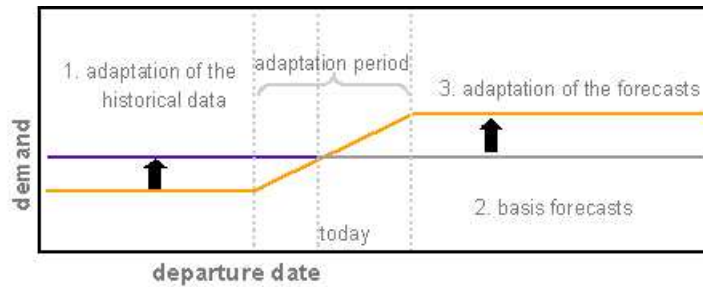


Figure 11: Adaptation to a market change. Market changes are typical examples of attractivity changes. First the history is adapted to the attractivity of the most recent departure, then the forecasts are adapted to future attractivity changes.

3 EXPERIMENTAL RESULTS

The experiments (forecasts) have been performed on ten different origin-destination pairs for 18 fare levels, distinguished also by the point of sale of the ticket, "country of origin", "country of destination" and "other". For learning purposes, 20 weeks of demand data have been available. The forecasts have been produced for another, separate set of 20 weeks. All the reported forecasts are 8 weeks ahead out-of-sample forecasts. The learning process has been adaptive, that is, the models' parameters have been re-estimated as new observations become available.

Forecasting method	Mean absolute error
Simple exponential smoothing	42.04
Brown model (without season)	43.31
Holt winters model (without season)	43.43
Linear regression	52.11
Autoregression	54.54
Moving average	56.27
ARMA	55.49

Table 1: Mean absolute error for different time series forecasting techniques without adaptation.

Adaptation technique	Mean absolute error
Flight specific demand (see 2.2)	24.14
Season + events (previous years) and flight specific demand (see 2.2, 2.3)	28.37
Season + events (previous years + current year) and flight specific demand (see 2.2, 2.3)	21.52
Market situation (see 2.5)	39.35
Market situation, flight specific demand (see 2.5, 2.2)	23.78
All adaptation techniques	20.14

Table 2: Mean absolute error for simple exponential smoothing forecasts supplemented by various adaptation techniques discussed in Section 2.

In order to illustrate the importance of adaptation in the discussed application firstly seven different, well known time series forecasting techniques have been applied to the above described data and their performance compared. Both the methods and their performance on the testing data are shown in Table 1. Mean absolute error per departure shown in the second column of the table have been used as the performance measure. As the mean demand at the departure for the examined data is 489.37, the forecast errors vary between about 8% to 12% of the demand. As can be seen from Table 1 the best results have been obtained using simple exponential smoothing method.

In order to measure the impact of different adaptation mechanisms discussed in Section 2 on an overall forecast quality, in the second batch of experiments the best performing individual forecasting method (i.e. the simple exponential smoothing method) has been combined with different adaptations. As in the case of basic forecasts reported in Table 1, the forecast is performed out-of-sample 8 weeks before the departure, but in some of the adaptations the already existing demand of the departures to forecast is used. The adaptations indicated in Table 2 have been investigated. The forecast errors of the adapted forecasts are shown in the second column of Table 2 . The only adaptation for which the experiments have not been performed for the considered data is the adaptation to schedule changes. The reason for this omission is that most of the schedule changes result in initializations or deletions of time series values which in turn would require significantly more historical data than was available for the individual forecasting methods to perform adequately. Therefore, origin- destination pairs without relevant schedule changes have been chosen for the experiments.

As it can be seen from Table 2 all types of adaptation have resulted in improved forecasts in comparison to the forecasts without adaptation. Not surprisingly, the most significant improvement have been noticed for cases with the adaptation to the flight specific demand and seasonal behaviour. It can also be seen that even the information about variations of short term influences should be taken into account, learnt and adapted to. This is illustrated in the results of the adaptation to the season and events. If the seasonality and event information is not adapted to the current years behaviour (Table 2, row 2) the forecast is worse than when only adaptation to the flight specific demand is used (Table 2, row 1). On the other hand, if we supplement the historical seasonal model by adapting to the current year's seasonal and event related changes (Table 2, row 3) a significant improvement can be observed.

When all models of adaptation have been utilized the basic forecast error has been reduced by over 50% which also represented the best of all forecasting results.

4 CONCLUSIONS AND FURTHER RESEARCH

Accurate forecasting of the demand for airplane tickets forms an important part of a revenue management process in the airline industry. In this paper we have presented an overview of adaptive mechanisms utilized in the new forecasting system of the

Lufthansa Airlines. The experiments show that for our application adaptivity is a very important feature and that the forecast quality can be significantly improved when using adaptive methods. Especially the adaptation to the flight specific demand and the season have been of crucial importance.

The parametrization of the adaptations is currently based on fixed rules. The basic forecasts produced by different time series models are combined by a fixed rule depending on the booking level. The weighting parameters to adapt the forecasts to the season and events are calculated by fixed rules too, depending on the number of days to departure and the seasonality and event information.

And though all these rules have been optimized with respect to the performance, the intention is to go away from the fixed rules and to adopt a dynamic approach to combination. The combination should adjust the used basic methods as well as the adaptation to the events/ season/ market situation.

Our current research is focused on combining individual forecasts and the replacement of the existing rules by dynamically calculated weights depending on different influence factors. The first tests have been executed using linear combination methods [1],[11] and the preliminary results are promising.

References

- [1] Bates, J., Granger, C.W.J., 1969. "The combination of forecasts". Operational Research Quarterly
- [2] Brockwell, P.J., Davis, R.A., 1987. "Time Series Theory and Methods". Springer Verlag
- [3] Brown, R.G., 1963. "Smoothing, Forecasting and Prediction". Prentice- Hall
- [4] Cross, R.G., 1997. "Revenue Management". Broadway Books
- [5] Edel, K., Schaeffer, W.S. 1997. "Analyse saisonaler Zeitreihen".Physica Verlag
- [6] Franses, P. H., 1963. "Time series models for business and economic forecasting". Cambridge University Press 1998
- [7] Godfrey, G.A., Powell, W.B., 2000. "Adaptive estimation of daily demands with complex calendar effects fro freight transportation". Transportation Research Part B 34, 451-469
- [8] Granger, C.W.J., Newboldt, P., 1986. "Forecasting Economic Time Series". Academic Press, Inc.
- [9] Kennedy, P., 1992. "A guide to Econometrics". The MIT Press
- [10] Masters, T., 1995. "Neural, Novel & Hybrid Algorithms for Time Series Prediction". John Wiley & Sons
- [11] de Menezes, L.M., Bunn, D.W., Taylor, J.W., 2000. "Review of guidelines for the use of combined forecasts". European Journal of Operational Research 120, 190-204
- [12] Rudolph, A., 1998. "Prognoseverfahren in der Praxis". Physica- Verlag 1998