Sales Prediction with Parametrized Time Series Analysis

Michael Schaidnagel*, Christian Abele*, Fritz Laux†, Ilia Petrov†
Fakultät Informatik
Reutlingen University
D-72762 Reutlingen, Germany
{*michael.schaidnagel|christian.abele*}*@student.reutlingen-university.de**
{*fritz.laux|ilia.petrov*}*@reutlingen-university.de*†

***Abstract*—When forecasting sales figures, not only the sales history but also the future price of a product will influence the sales quantity. At first sight, multivariate time series seem to be the appropriate model for this task. Nontheless, in real life history is not always repeatable, i.e. in the case of sales history there is only one price for a product at a given time. This complicates the design of a multivariate time series. However, for some seasonal or perishable products the price is rather a function of the expiration date than of the sales history. This additional information can help to design a more accurate and causal time series model. The proposed solution uses an univariate time series model but takes the price of a product as a parameter that influences systematically the prediction. The price influence is computed based on historical sales data using correlation analysis and adjustable price ranges to identify products with comparable history. Compared to other techniques this novel approach is easy to compute and allows to preset the price parameter for predictions and simulations. Tests with data from the Data Mining Cup 2012 demonstrate better results than established sophisticated time series methods.**

***Keywords*-sales prediction, multivariate time series, price-sales correlation, parametrized predictor.**

I. Introduction

Sales prediction is an important goal for any time series based analysis [1], [2]. The task consists of forecasting sales quantities given the sales history. This can be achieved by extending the time series into the future.

The prolongation of the time series into the future is determined by the underpinning time series model [3]. If this model is not well supported by the empirical data it is likely that the accuracy of the forecast is low. So the challenge is to find data build "similar" products or situations (e.g. time or price) and form clusters of products for building a prediction model. If a major sales factor like the product price changes, the model based solely on previous sales will lead to wrong forecasts. Therefore, it is important to include the price as parameter into the model in addition to the sales history.

Standard solutions to this problem have been to supply a long history of sales with sufficient data to validate the model and to correlate the sales data with the variable product price. The mathematical tools of choice for analyzing multiple time series simultaneously are multivariate statistical techniques like Vector AutoRegressive (VAR) models [4], [5] or such as the Vector ARIMA (AutoRegressive Integrated Moving Average) [6]. The model parameters are estimated with least square or Yule-Walker functions [4, p. 181]. The accuracy of the estimator depends on the number of observations and its "strength" of correlation.

To illustrate the process consider an excerpt from the Data Mining Cup 2012 [7] dataset (Table I:

Table I
Sample Data, Data Mining Cup 2012

day	Prod#	price	quantity
1	1	4.73	6
1	2	7.23	0
1	3	10.23	1
1	4	17.90	0
...	...	...	...
1	570	7.91	0
2	1	4.73	12
2	2	7.23	1
...	...	...	...
42	569	9.83	2
42	570	7.84	0
43	1	5.35	?
43	2	7.47	?
...	...	...	...
43	570	7.84	?
...	...	...	...
56	570	8.12	?

The information provided comprises a collection of 570 products whose history of sales and prices are given over a period of 42 days. The task was to predict the sales quantities for the next 14 days where the daily sales price was preset. The majority of products produced only low quantity sales. Comparisons with other sales data showed a similar distribution [8] [9] which indicates that the sample is typical for larger collections.

When we tried to predict the future sales with commercial ARIMA products we experienced a low prediction quality with a relative accuracy of only 47%.

The disappointing results from professional tools implementing ARIMA encouraged us to look for a simpler and better prediction model. First of all we assumed that the future price is causally influenced and should not be treated as stochastic variable. Second, it might be helpful to filter out cyclic behavior from the "white noise" in case of low volume sales. Third, the proposed approach is suitable for online scenarios: (i) it has low computational overhead, the underlying model is simple and maintainable; (ii) the algorithm can operate on partial datasets and can be used for incremental forecasting.

A. *Structure of the Paper*

In the next subsection follows a discussion of related work and we contrast it with our contribution.

The rest of the paper is structured as follows: The data profile under investigation and the research problem will be described formally in Section II. In Section III we present our parametrized time series algorithm that predicts sales volumes with variable product prices and low data support. The following Section IV gives a description of the technical framework for the implementation of the prototype. The results are discussed in Section V and compared with standard methods found in commercial products like ARIMA. From these experiences we draw our conclusion in the last section.

B. *Related Work*

Adaptive correlation methods for prognostic purposes have been proposed early in the 1970^{th} by Griese [10] and more specifically as AutoRegressive Moving Average (ARMA) method by Box and Jenkins [11]. As ARMA is constrained to a stationary stochastic process the ARIMA is of more practical use as it can handle time series with a linear trend and is therefore widely implemented.

The idea behind ARMA and ARIMA is that the model adapts automatically to a given history of data. A natural extension is to include other influential factors beside the prognostic value itself. This leads to multivariate models, namely Vector Autoregressive (VAR) models [6]. The development of the model was influenced and motivated by critiques of Sims [12] and Lucas [13]. In essence, their statement is: every available data is potentially correlated.

If the model is extended to cover the influence from correlated data this leads to a vectorial stochastic model $(\mathbf{X}_t(\pi, \pi_r))$ that allows not only the serial time dependence t of each component but also the interdependence of products π and product prices π_r. We prefer to use parentheses () for a stochastic process instead of braces $\{\}$ because it is rather a sequence of stochastic variables than a set.

To estimate the parameters of such a multivariate ARMA process the following equation ([14, p. 417], [4, p. 167]) has to be solved:

$$\Phi(L)X_t(\pi, \pi_r) \quad = \quad \Theta(L)Z_t \tag{1}$$

where L denotes the backshift (lag) operator and

$$\Phi(x) \quad := \quad I - \Phi_1 x - \Phi_2 x^2 - \ldots - \Phi_p x^p \tag{2}$$

$$\Theta(x) \quad := \quad I + \Theta_1 x + \Theta_2 x^2 + \ldots + \Theta_q x^q \tag{3}$$

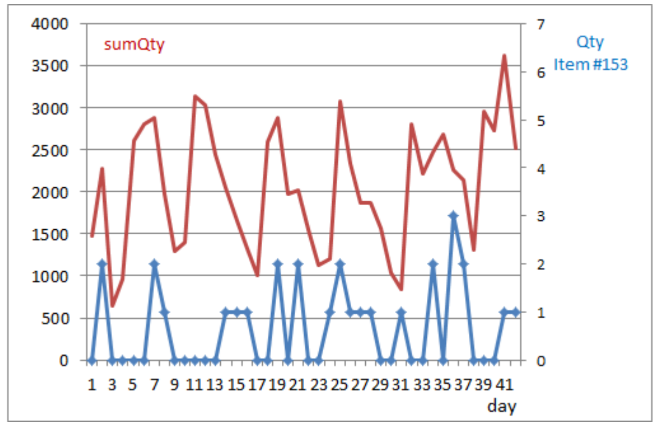

Figure 1. Seven day periodicity for the overall sales data (DMC2012) and a typical low selling product (item # 153)

are matrix-valued polynomials with dimensions of p (order of regression) and q (order of moving average). Z_t denotes a multivariate "white noise" process.

There is one major drawback to this approach in our problem setting. The model treats all historical input values as stochastic variables. However, the product price does not vary *stochastically*, its value is preset by the vendor. Economic models assume a causal dependency between the price of a product and its sales quantities (see Arnold [15, chap. 17]). Variations in consumer demand are caused by various factors like price, promotions, etc [16]. This causal dependency is not modeled by VAR methods. This is an issue for the multivariate model.

Another complication arises from the noisy periodicity resulting from low volume sales. The low sales quantity introduces a kind of random pattern that makes it hard to find even a known periodicity. In the sample data the overall sales history shows a clear 7-day periodicity (see Figure 1) but not for individual products.

Cyclic sales quantities are a typical behavior for short shelf-life products and are important for building a causal sales model. Doganis et al. [17] investigated the sales quantity of fresh milk (a short shelf-life product) in Greece. They used a genetic algorithm applied to the sales quantities of the same weekday of last year. Our approach is only similar in that we take corresponding weekdays but it differs in how we analyse the weekly periodicity and correlate it with the sales prices.

To recapitulate, there are two general arguments against the multivariate VAR approach sketched above: Granger and Newbold [18] showed that simpler models often outperformed forecasts based on complex multivariate models. And Lucas [13] criticized that the economic models are too static and that "any change in policy will systematically alter the structure of the econometric model". Applied to the sales forecast situation the variation of the price does not play a

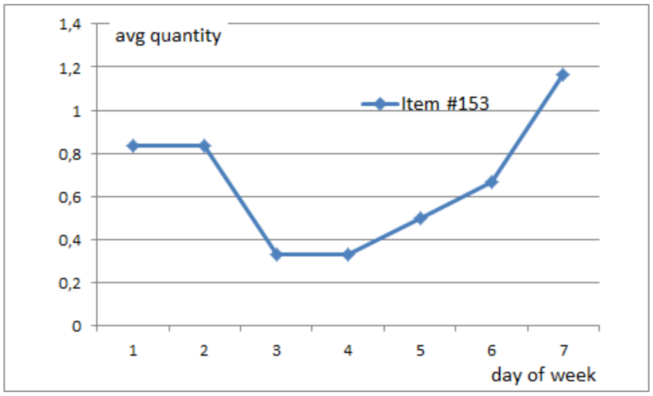

Figure 2. Seven day periodicity for the sales of item # 153

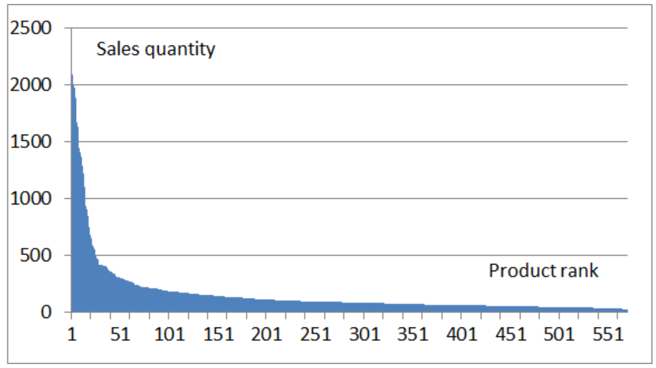

Figure 3. Sales quantity ranking of sample data (DMC2012)

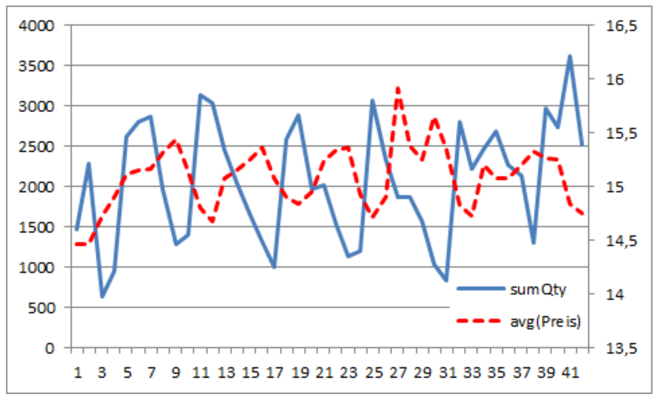

Figure 4. Sales quantity and average price time series of sample data (DMC2012)

stochastic, but a *systematic*, i.e. a functional, role.

Our idea is to filter the seasonality by a period-based "folding" of the sales quantity, i.e. the aggregation of sales quantities for the same weekdays. This cancels the stochastic variation and accumulates the seasonal effect. Applying such a model improves the prediction coverage and accuracy for low volume data with a cyclic behavior.

II. Problem Description and Contribution

In Section I we pointed out that the nature of the data and its sales profile play an important role for the time series analysis. In particular, the influence of price and periodicity are dominant factors as we will see in the following. For this reason we present first the characteristics of the data.

A. Data Profile

The 570 products of the sample data realized a total of 86641 units sold. The average price ranged between 14.46 and 15.92 over the period of 42 days. The maximum price variability of a single product is $\pm 48\%$, but on average the price varies only by $\pm 9\%$. However, for high selling products (> 500 units) the variability stands at $\pm 15\%$.

The total sales quantity per product ranged from 17 to 2083 over the 6 weeks. Broken down to the day level the product price ranged from 0.24 to 152.92 and the sales quantities between 0 and 193. The sample had average sales per product of 152 units with a standard deviation of 257 which indicates a high sales variability of the items.

This conjecture is confirmed by the product sales ranking that roughly follows a shifted hyperbolic distribution (see Figure 3) which supports that low volume sales contribute significantly to the overall sales and may not be neglected. 506 products out of 570 sell less than 250 units in total but contribute with approximately the same quantity sold (43991 units) as the 64 high selling products.

The low volume sales (sum of sales < 250) showed a strong positive trend ($\approx 40\%$ increase over 42 days) whereas the high volume sales (sum of sales ≥ 250) had a more stationary behavior. In the sample data are more than 100 products that sell less than six units a day. Nearly all of them sell none at half of the time.

The above properties require that an adequate forecasting algorithm, which has to allow for low volume sales with high variability and be able to accustom for some price variability.

A 7 day sales periodicity was highly visible over the whole sample (see Figure 1) but not on the product level. Tests with sample low selling products with an assumed periodicity of 7 days produced better results than without this assumption. Again, the cumulative values of sales quantities and prices plotted over time give clear indication for a negative price-quantity correlation and a 7 day periodicity except for the first week (Figure 4). These results indicate that our algorithm has to deal with hidden periodicity.

In order to avoid the risk that our method is tailored for a particular data set we investigated a second data set with a completely different profile. The data records contain Piglet Market Data (PMD), including price and quantity sold for weekly auctions over a period of six years. The PMD does not show any clear periodicity but the market tended lower towards the end of the year. It is characterised by the following quantities: (i) Min. Qauntity - 701; (ii) Max. Quantitiy - 2063; (iii) Min. Price - 20 €; (iv) Max. Price -

71.5 €; (v) the trace sequence length spans 311 data points; (vi) Qantity standard deviation - 207.97; (vii) price standard deviation - 10.38 € which depicts a similar variance as the previous time series but without periodicity. We will show in Section V that our method produces a more accurate prediction.

B. Formal Problem Description

The problem consists of developing a parametrized time series model that is able to forecast future sales quantities depending on the given sales history and a price parameter. The solution of the stochastic Equation (1) is a multidimensional mapping

$$\begin{aligned} F: \quad (\mathbf{\Pi}, \mathbf{T}) &\longrightarrow (\mathbb{R}^+, \mathbb{N}_0) \\ (\pi, t) &\longmapsto (\hat{\pi}_r, \hat{x}_t) \end{aligned} \tag{4}$$

where $\mathbf{\Pi}$ is the set of products and $\mathbf{T}$ are consecutive time intervals. A product $\pi \in \mathbf{\Pi}$ is described by its identification number π_i and its price π_r. The mapping F computes sales quantity $\hat{x}_t$ and price $\hat{\pi}_r$ for every product π and time interval t .

The vector time series $(\hat{\pi}_r, \hat{x}_t)$ is a concrete realisation of the stochastic process (X_t) of Equation (1). The mapping F has to be adjusted so that the process (X_t) explains best a given realisation. This can be done by various estimator functions: least square error, Yule-Walker, maximum-likelihood, or Durbin-Levison algorithms. This is where our approach differs from the traditional because in real business the price is not a stochastic variable but is preset by the vendor. Instead of predicting the future price $\hat{\pi}_r$ we use the price as input parameter.

Having fixed the model in this way it is possible to transform the mapping F to the following form:

$$\begin{aligned} F_r: \quad (\mathbb{N}, \mathbb{R}^+, \mathbf{T}) &\longrightarrow \mathbb{N}_0 \\ (\pi_i, \pi_r, t) &\longmapsto \hat{x}_t \end{aligned} \tag{5}$$

With this predictor $F_r(\pi_i, \pi_r, t)$ it is possible to forecast the sales quantities for future time periods $t > T$ (T is the present time) of a product $\pi \in \Pi$ using the future price π_r as input.

C. Contribution

By restricting our approach to model a linear trend, seasonality, and using historic and future prices as causal parameter leads to a predictor function that is easy to compute and explain. It yields higher accuracy for data with hidden periodicity and variable prices than the ARIMA model. The novelty of our contribution comprises:

- a model that has a causal explanation
- where the future price is a major input factor and
- the overall periodicity is respected by individual items.

The prediction function can also be used for simulation to see how the price will influence the sales quantity.

III. The Parametrized Time Series Model

For a causal predictor function F_r we need to identify and quantify all influencing factors. Therefore we analyzed different correlations of the attributes quantity, price and time. We used the standard Pearson Correlation [19] as a measure to determine the linear dependence between two time series. It is widely used and can range between -1 and $+1$. This section will present the relations which have been analyzed.

A. Price-Sales Correlation

The main conjecture was that the price has a causal influence on the quantity. This is justified by the price elasticity of demand theory by Alfred Marshall [20]. As the correlation coefficients of all 570 products ranged from -0.6515 to $+0.3471$, we expected that the products with strong correlation exhibit a better prediction accuracy. Surprisingly this seemed not to be the case.

A systematic analysis with three synthetic time series lead to an explanation. The first series had a growing price trend, the second and third had a cyclic price development where one product responded immediately and the other responded with a delay. ARIMA did recognize the price trend but forecasted a constant quantity instead of a decreasing one. This was the result of the low integer sales numbers that produced a monotone decreasing step function. Our approach managed to forecast the right quantities as long as a matching price was present in the history.

Surprisingly ARIMA could not deal well with the systematic cyclic price development and a detailed analysis showed that the step function of the price (price was kept constant for two days) was the reason. Figure 5 shows the result of the ARIMA compared to our F_r algorithm (see equation 5). The lower Qty (Figure 5) values at the extrema produced by our algorithm results from the delay in the response to the price change. Without lag, no damping of extrema occurs in F_r.

B. Price Similarity

We analyzed correlations between the price development of different products. The assumption was to find product bundles which are linked together via their price development. For the analysis the prices were first normalized such that we could easily compare the different price levels. Several bunches of products were linked together via their prices. But the corresponding sales figures of these products were not related. This is why we ignored the possible cross price influence from other products for the forecast.

C. Sales Periodicity

One of the most interesting properties of the given data was the periodicity of the total sales curve. It showed a clear 7 day period (Figure 1). This period was not directly observable in most sales time series of individual products.

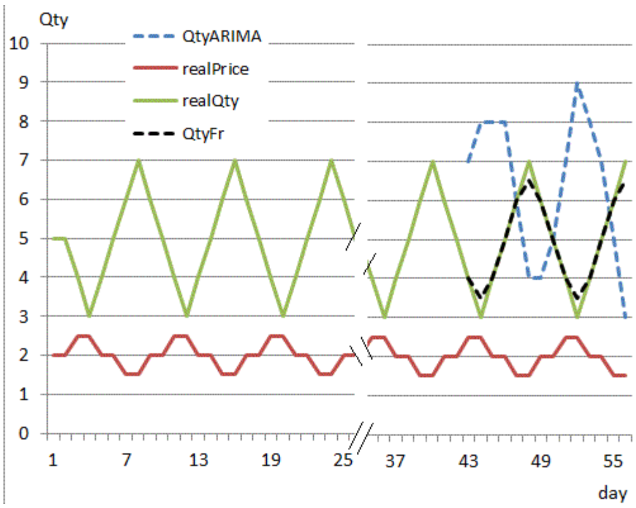

Figure 5. Forecast of synthetic time series with delayed price-sales dependency

Input: $t > T$ // prediction day
  π_r (input) // price at day t
  δ (input) // price range (e.g. $\pm$ 10%)
  $period$ (input)
Def: u, l // upper & lower price limit
  QtyList // list of sales quantities
  w // = t - n * period
  $\hat{x_t}$ // predicted quantity at day t
for each $\pi \in \Pi$ {
  $u := \pi_r(1 + \delta/100)$; $l := \pi_r(1 - \delta/100)$
  $w := t - period$
  while $(w \geq 1)$ {
   if $(\pi_r(w) < u)$ & $(\pi_r(w) > l)$
    QtyList.add($p(w)$) // $p(w)$ is Qty on day w
   $w := w - period$
  }
  if (QtyList $\neq \emptyset$)
   $\hat{x_t} := trend$(QtyList)
   return $\hat{x_t}$
  else
   return nil
}

Figure 7. Parametrized Sales Prediction Algorithm F_r

Also the Pearson Correlation between the sum curve and the single products was too low to draw any conclusions. Nevertheless, since the total sales curve consists of all products, there must be a hidden periodicity within the individual products.

When we used this 7-day periodicity and summed up the sales of corresponding days most products confirmed the 7-day periodicity. Cumulative sales quantities over matching weekdays reveal hidden cyclic pattern even for low volume sales products. Figure 2 shows the periodicity pattern of product #153 that cannot be seen from its original time series (Figure 1).

Using this 7-day periodicity we were able to improve the forecast quality of our method. The algorithm takes only prices from the same days of each period (see Figure 7 for an example) and computes its trend ($trend(Qtylist)$). The trend computation can be performed in different way; for our purses we use linar trend. A more detailed explanation will follow in the next subsection.

A systematic spectral analysis discovered not only the dominant weekly patterns but weaker 4, 5 and 14 days patterns.

D. The Parametrized Predictor Function

Putting all correlation observations together the result is a function F_r whose pseudocode is shown in Figure 7. As an input, it takes the price π_r of a product π at the prediction day t, the periodicity and a price range δ. The upper and lower price limits are set to $\pm\delta$ percent. Using the periodicity from the previous analysis the algorithm looks for prices $\pi_r(w)$ that occur on days $w = t$ modulo $period$. For example, if the prediction day is 43 and the periodicity is 7 days, only the information from the days 36, 29, 22, 15, 8 and 1 will be considered (see Figure 6). If the price on such a day is outside of the upper or the lower limit (day 1 in our example), the sales quantity is ignored. If the price is within the bounds, the corresponding quantity on that day ($p(w)$) is inserted into the quantity list ($QtyList$). After all matching quantities have been selected, the forecast quantity is computed as linear trend ($trend(Qtylist)$) of these quantities.

If all prices are outside of the upper and lower limit, no forecast is produced. The procedure may be repeated with enlarged upper and lower limits if needed. This algorithm defines a simple forecasting model that takes into account the sales trend, the periodicity, and the price influence to predict sales quantities.

IV. Technical Framework and Infrastructure

This section covers some technical details about execution and implementation of the two models discussed in this paper.

A. ARIMA Model Execution

The Microsoft Visual Studio 2008 and Microsoft SQL Server 2008 were used to apply the ARIMA model on the given data set. In order to run the ARIMA mining models a OLAP cube was build. It consists of the dimensions *price*, *product* and *time*. In the corresponding time series mining model we used *itemId* and *day* as key attributes and the *price* attribute as input. The *quantity* was set as predictable attribute. Most model parameters were left as default, except the minimum series value and the periodicity hint. The

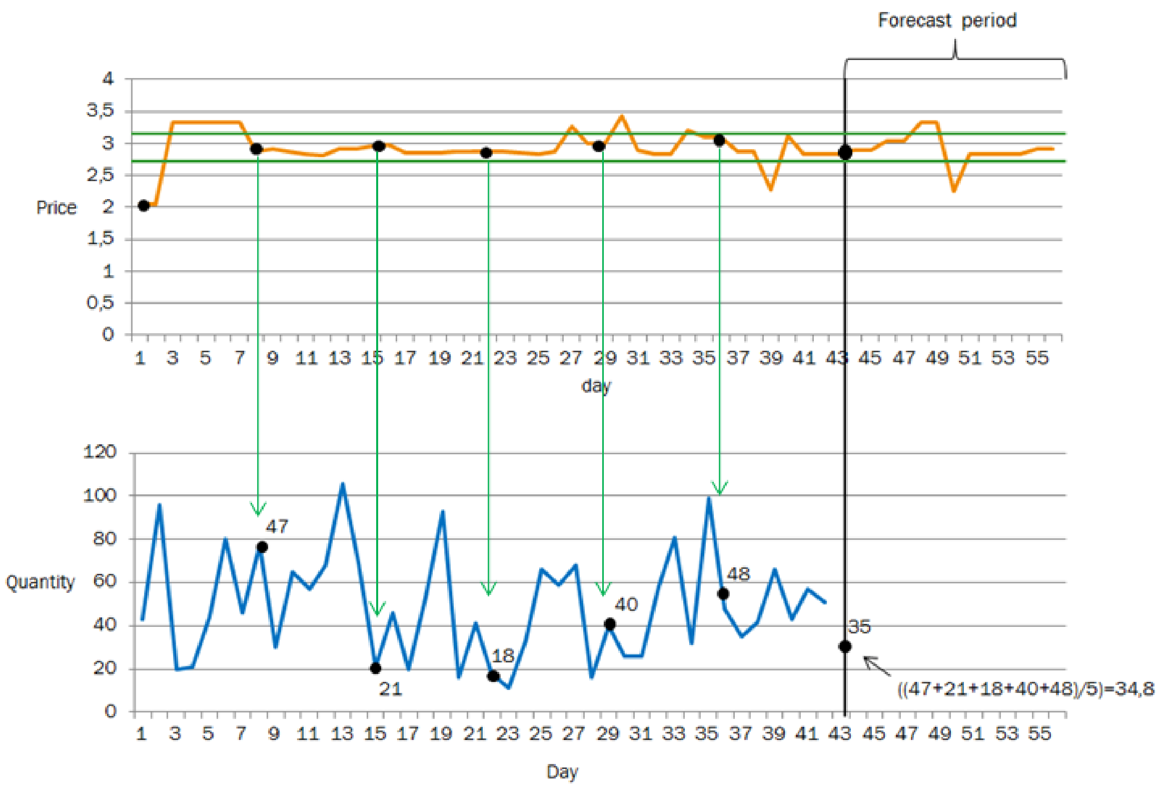

Figure 6. Illustration of prediction concept using trend, periodicity, and a parametrized price

following listing shows all model parameters used:

AUTO_DETECT_PERIODICITY	0.6
FORECAST_METHOD	ARIMA
HISTORIC_MODEL_COUNT	1
HISTORIC_MODEL_GAP	10
INSTABILITY_SENSITIVITY	1.0
MAXIMUM_SERIES_VALUE	$+1E308$
MINIMUM_SERIES_VALUE	0
MISSING_VALUE_SUBSTITUTION	None
PERIODICITY_HINT	7
PREDICTION_SMOOTHING	0.5

B. Implementation of the Suggested Model in Java

Our own approach is implemented in Java. We used Eclipse (Version: Indigo Service Release 1) with Java Platform Standard Edition 6.0 (JRE6). The data was stored in a MySQL database on an Apache web server (2.2.21). During execution time the data is queried from the database, the model parameters computed and the forecast results are instantly stored in the corresponding result table in the database. The model was developed using the standard java.sql.* package which was used to interface with the database and for SQLException handling.

V. Results

The absolute prediction error was measured as $|realQty - predictQty|$. The performance of F_r was measured against ARIMA and the results are given as sum of absolute error numbers over 14 prediction days. The F_r algorithms benefited from two input parameters: the hidden periodicity that was calculated in a previous step and the predefined future price. The hidden sales periodicity contributed for an improvement of about 20%. The overall forecast was improved by 26.7%. The price influence was less dominant than expected, but was determinant for a cluster of 26 products. Cluster characteristics: (i) correlation < -0.25; (ii) relative standard error < 0.25; (iii) sales quantity > 160; and (iv) price variation $((max(\pi_r) - min(\pi_r)) > 4)$. In total, F_r could forecast this cluster 36.4% better than ARIMA.

A. Comparison of Results with ARIMA

In this section we compare the results of ARIMA and our F_r method. The following table shows the prediction error points of both ARIMA and F_r. The prediction error is calculated as sum of the absolute difference between real and predicted values ($\sum \mid realQty - predictedQty \mid$, $predictedQty \in \{predictedARIMA, predictedF_r\}$). The total error of all 570 products was 30152 for the ARIMA and 22093 for F_r. This is an improvement of 26.7% compared to ARIMA (Table II). For further analysis we splitted the products into disjoint sets according to different criteria. This allowed us to find the strengths and weaknesses of our algorithm in terms of total sales quantity, sales sparsity, and price.

In the case of PMD we predicted the sales quantities for a full year. The result analysis for this forcast yields an improvement of 20% for the F_r alogrithm over ARIMA. This is less than for the DMC set but with more then 20% still substantial.

Table II
COMPARISON OF ARIMA PREDICTION ERROR WITH F_r ALGORITHM

Class	ARIMA	F_r	improvement
All products	30512	22093	26.7%
quantity < 500	24338	17248	29.1%
quantity ≥ 500	6174	4845	21.5%
(quantity = 0) in $< 1/3$ time	19178	15942	16.9%
(quantity = 0) in $\geq 1/3$ time	11334	6151	45.7%
avg(π_r) < 20	26756	18711	30.1%
avg(π_r) ≥ 20	3756	3382	10.0%
top 100 items	15805	6131	61.2%
least 470 items	16167	15962	1.2%
PMD	10841	8612	$20,6\%$

VI. CONCLUSION AND FUTURE WORK

The broad range of products with its hidden periodicity made the analysis difficult. The low volume sales further complicated the analysis of the influence of the price on the sales quantities. The conclusions drawn from the above results can be reduced to the following three statements:

1) Data profiling is crucial for choosing the best time series model
2) Low sales volume can hide a cyclic sales behavior and the price should be treated as input parameter
3) Simple models for sales forecasting based on causal parameters can outperform some sophisticated stochastic models.

Cross influence from other products should be further investigated. This could lead to clusters of products that behave coherently. Sometimes products are complementary. This has not been investigated. The overall sales had a significant periodicity length of 7, the available data covered only 56 days. Hence, a longer time history would be needed to verify this and to look for overlaying periodicities. A spectral analysis on a individual product level could further improve the prediction accuracy. For products with a strong monotone price development our approach to look for similar prices is not well suited. The price trend should be computed instead.

The F_r algorithm can be used with incomplete time series. This is particularly useful for real-time analysis used in recommender systems and should be further investigated.

REFERENCES

[1] August-Wilhelm Scheer, *Absatzprognosen* [engl. Sales Forecasting], Springer Verlag, Berlin, 1983

[2] Manfred Hüttner, *Markt- und Absatzprognosen* [engl. Market and Sales Forecasting], Kohlhammer, Stuttgart, 1982

[3] Yang Lan and Daniel Deagu, "A New Approach and Its Applications for Time Series Analysis and Prediction Based on Moving Average of n^{th}-Order Difference", in: Dawn E. Holmes and Lakhmi C. Jain (Eds.), *Data Mining: Foundations and Intelligent Paradigms*, Vol 2: Statistical, Bayesian, Time Series and other Theoretical Aspects, pp. 157 - 182, Springer Berlin Heidelberg, 2012

[4] Klaus Neusser, *Zeitreihenanalyse in den Wirtschaftswissenschaften* [engl. Time Series Analysis in Economic Sciences], B. G. Teubner Verlag, Wiesbaden, 2006

[5] Alan Julian Izenman, *Modern Multivariate Statistical Techniques - Regression, Classification, and Manifold Learning*, Springer Science + Business Media, New York, 2008

[6] Helmut Lütkepohl, *New Introduction to Multiple Time Series Analysis*, corr. repr., Springer Verlag, Berlin Heidelberg, 2007

[7] N. N., DATA MINING CUP 2012, prudsys AG, Zwickauer Str. 16, D-09113 Chemnitz, [Online] http://www.data-mining-cup.de/en/review/dmc-2012, last access: 13.01.2013

[8] Sam Oches, "Top 50 Sorted by Total Units", Special Report of QSR Magazine, Journalistic Inc., August 2011, [Online] http://www.qsrmagazine.com/reports/top-50-sorted-total-units

[9] Economist Intelligence Unit, "Denmark: Market Indicators and Forecasts", [Online] http://datamarket.com/data/set/1wmo/ (indicators: Private consumption, Consumer goods), last access: 30.08.2012

[10] Joachim Griese, *Adaptive Verfahren im betrieblichen Entscheidungsprozess* [engl. Adaptive Methods in Operational Decision-Making], Physica Verlag, Würzburg - Wien, 1972

[11] Georg E. P. Box and Gwilym M. Jenkins, *Time Series Analysis: Forecasting and Control*, 1^{st} rev. ed., Holden Day,Oakland, San Francisco, 1976

[12] Christopher A. Sims, "Macroeconomics and Reality", in: Econometrica, Vol. 48, No. 1, pp. 1 - 48, 1980

[13] Robert E. Lucas, "Econometric policy evaluation: A critique", In: K. Brunner and A. H. Meltzer (Eds), *The Phillips Curve and Labor Markets*, Vol. 1, Carnegie-Rochester Conference Series on Public Policy, pp. 19 - 46, Amsterdam, North-Holland, 1976

[14] Peter J. Brockwell and Richard A. Davis, *Time Series: Theory and Methods*, 2^{nd} Edition, Springer Science + Business Media, New York, 2006

[15] Roger A. Arnold, *Economics*, 9^{th} ed., South-Western College Publ., 2008

[16] Jack G. A. J. van der Vorst, Andrie J. M. Beulens, W. de Wit, Paul van Beek, "Supply chain management in food chains: Improving performance by reducing uncertainty", in: International Transactions in Operational Research, Vol. 5(6), pp 487 - 499, 1998

[17] Philip Doganis, Alex Alexandridis, Ranagiotis Patrinos, and Haralambos Sarimveis, ”Time series sales forecasting for short shelf-life food products based on artificial neural networks and evolutionary computing”, Journal of Food Engineering 75, pp. 196 - 204, Elsevier Ltd., 2006, [Online] http://www.elsevier.com/locate/jfoodeng, last access: 30.08.2012

[18] Clive W. J. Granger and Paul Newbold, ”Economic Forecasting: The Atheist's Viewpoint”, in: G.A. Renton (ed.), *Modelling the Economy*, pp. 131 - 148, Heinemann, London, 1975

[19] Horst Rinne and Katja Specht, *Zeitreihen - Statistische Modellierung, Schätzung und Prognose* [engl. Time Series - Statistical Modelling, Estimation and Prediction], Verlag Vahlen, München, 2002

[20] Alfred Marshall, *Principles of Economics*, 8^{th} ed., Cosimo Classics, 2009, first publ. 1890