Early Detection of Critical Faults Using Time-Series Analysis on Heterogeneous Information Systems in the Automotive Industry

Thomas Leitner*, Christina Feilmayr†, Wolfram Wöß‡
Institute for Application Oriented Knowledge Processing, Johannes Kepler University Linz
Linz, Austria
Email: *thomas.leitner@jku.at, †christina.feilmayr@jku.at, ‡wolfram.woess@jku.at

***Abstract*—Beside the manufacturing industry's vision of *industry 4.0*, which is about improving the degree of automation and customisability depending on a huge amount of data, the automotive industry increasingly advances the after-sales market collecting more and more information about the car using sensors and diagnostics mechanisms. This information can be used to earlier reveal malfunctions and faults with rising quantity that customers experience in order to reduce the solving time of the problem. Different heterogeneous data sets exist storing data at various approval stages with different data quality. In order to perform the most accurate detection of critical developing faults it is fundamental to use as much data as possible while weighting their impact by assessing their data quality. For detecting critical performing faults as early as possible time series analysis and forecasting methods are used to analyze their course and predict future values. In this research work, a new approach is proposed, which is subdivided in the following four main tasks: (i) evaluation of data quality metrics of different warranty information systems, (ii) analysis and generation of forecasts on univariate time series based on *Auto-Regressive Integrated Moving Average* (ARIMA), (iii) weighted combination of different predictions, and (iv) improvement of the accuracy by integrating prediction errors. This solution can be used in arbitrary fields of application, in which different information sources should be analyzed using data quality metrics and prediction errors to determine critical courses as early as possible.**

Keywords*–data mining, time series analysis, data quality metrics, automotive industry

I. INTRODUCTION

Quality - abnormality and cause analysis (Q-AURA – *Qualität - Auffälligkeiten und Ursachenanalyse*) is an application developed in cooperation with *BMW Motoren GmbH*, located in Steyr, Austria, with the purpose to decrease the problem solving time for finding causes of engine faults in automobiles in the after-sales market. The system has different goals for supporting the quality management expert in his daily work; (i) automatically finding significant faults that are developing badly, (ii) providing new useful information about the affected engines, and, (iii) analyzing bills of materials and engine components to find technical modifications that potentially provoke a particular fault. Q-AURA has already been evaluated and is used by the quality management experts every day. The first task of detecting significant faults uses fault information from warranty claims of previous weeks, but takes only those information within a specific approval stage into account that originate from a single data source. The goals to additionally use information of other warranty information systems at other approval stages and to earlier determine significant faults leads to the need of an optimisation of the existing system. Methods should be investigated that help to achieve robust results.

In this paper, an approach is presented that uses time series analysis, forecasting methods, and data quality assessment of different information systems. The paper is organized as follows: Section II discusses the central problem and associated challenges. Section III addresses methods that are necessary for the provided approach, while Section IV gives a detailed description of the proposed technique also explaining the integration into Q-AURA. Finally, Section V covers the conclusion.

II. PROBLEM STATEMENT

The contribution of this research work is a new approach, which consists of four parts, whereas each of them discusses a particular problem. The first one is the development of a *method for validating different information systems, which store partially contradictory, complementary, and redundant information*. The business process that is supported by Q-AURA ranges from the engine development department where new engine generations are developed or existing ones are improved to the after-sales market. In case of a technical fault, the car must be checked at a dealer's workshop, where the problem and information about the fault is sent to the automobile manufacturer. Since BMW sells cars in most countries, and since faults are classified differently in various markets, it is necessary to overcome the discrepancies yielding in a consistent view of all faults that occur. Different information systems exist that contain faults at different acceptance levels. These data sets are evaluated using data quality metrics.

The second task is the development of a *method for detecting critical developing faults as early as possible*. In order to identify critical faults, the trend of the most recent weeks is determined. Since enough values are necessary to provide robust results, there is a time delay between the beginning of the fault and its detection. By using prediction methods this delay can be reduced since future values can be predicted, which can then be used to determine earlier whether a particular course is critical or not. Different prediction methods have been evaluated and the best one was selected.

The third task is the development of a *concept to improve the prediction accuracy using forecasts from different information systems*. As explained above different views on fault and warranty information at different approval stages exist. E.g., while one source contains information that is already accepted by the company, but does not include values that are provoked by the customer, another one contains more faults, but those have not been verified by the company. Therefore, these different sources have to be analyzed separately, which results in separate predictions. By consolidating the forecasts

 ISBN: 978-1-61208-358-2

performed on each individual data set a better result can be achieved. Since the information quality of each individual data source has to be taken into account, the quality scoring is used to influence the weighting to get more accurate results.
The fourth task is the *integration into Q-AURA and verification of the proposed concepts* to demonstrate the improvement in comparison to the current applied approach. It is important to describe the established consolidated Q-AURA system to clearly show the benefit of the new approach.
The resulting approach is a set of methods that enable early detection using a weighted combination of forecasts based on data of multiple, heterogeneous information systems that adjusts its parameters using accuracy metrics of previous iterations and quality metrics of each data source.

III. Related Work

This section covers information about the used methods and gives a detailed description of the concepts the proposed approach is based on. Primarily, two basic concepts are discussed, which are data quality metrics including their assessment and analysis of univariate time series containing forecasting methods.

A. Data Quality

The literature provides a wide range of techniques for data quality assessments as well as definitions and descriptions about data quality dimensions and metrics [1], [2]. A detailed comparison is given by Batini et al. [1]. For a compact summary, the data quality metrics that are used in this approach are described in detail.
Completeness describes if all information in the real world within a particular scope is captured by the information system. In other words, a system is complete if it includes the whole truth. For a database scheme D, we assume a hypothetical database instance d_0 that perfectly represents all the information of the real world that is modelled by D. Further, we assume one or more instances $d_i(i \geq 1)$, where each of them is an approximation of d_0. Now we consider some views, where v_0 is an ideal extension of d_0 and $v_i(i \geq 1)$ are extensions of the instances d_i. Further we define completeness as described by (1). In the considered equation, the absolut values represent the number of tuples [3], [4].

$$\frac{|v_i \cap v_0|}{|v_0|} \tag{1}$$

Soundness (similar to completeness) is also determined by comparing the real world and the modelled instances in the information system. It describes if the information system stores the truth, and nothing but the truth, which means that all modelled information also exists in the real world. Equation (2) shows the definition [3], [4].

$$\frac{|v_i \cap v_0|}{|v_i|} \tag{2}$$

Consistency is a metric that focusses on the structural correctness of the represented data. This means that stored information has to meet some conditions, e.g., existing entries have to be unique (no duplicates) or meet assertions. In the literature, different definitions exist, some of them are similar to soundness in others [1].
Correctness is a metric that measures the semantic validity of data. Stored data is correct if it meets semantic rules. By applying those rules it can be determined if the particular entry is in the correct range or has a valid format, for example. Since functional requirements can change over time, it is important to modify these rules if necessary [5].
Integrity is also defined ambiguous in literature. Some definitions declare integrity as the combination of validity and completeness [6]. In this contribution, integrity is treated as the correctness of connections between data structures like tables or views. This means that the connections between data structures are monitored and if too many, wrong or too few results are calculated than expected this metric is decreased. Sometimes *inter-relation integrity* has a similar definition in literature [1].
The data quality metrics used in this approach were chosen carefully. While completeness and soundness measure if the quantities of the basic structural components are correct, integrity proves if the connections between them are valid. Consistency and correctness determine if the entries are semantically correct and in the specified value range.

B. Analysis of Univariate Time Series

This section focusses on time series analysis especially univariate ones. Time series analysis is also a very popular research field since a wide variety of applications exists, ranging from stock analysis and calculations concerning demography to sunspot observations. Forecasts are connected closely to time series analysis, since it is the prediction of future values of a known time series. A popular example are weather forecasts, where former observations are known and based on them (and physical law of course) future values are predicted [7].
The proposed approach focusses on univariate time series, which are time series that solely depend on one variable [7], [8]. Different methods exist for modelling and forecasting univariate time series, under them *Box-Cox Transformation, ARMA errors, Trend, and Seasonality* (BATS) and *TBATS*, *Simple Exponential Smoothing* (SES), *feed forward neural networks with a single hidden layer* (NNETAR), *Croston's method* (CROSTON), and *Auto-Regressive Integrated Moving Average* (ARIMA) [9], [10], [11].These different methods were compared with each other, and the best one was chosen for the proposed approach, which is ARIMA. The comparison was done using example data. For the evaluation which method facilitates the best results, *Goodness-of-Fit* measures (e.g., *Mean Percentage Error*, *Mean Absolute Percentage Error*) were used for comparing the fitting of the curve of each method, while the *Diebold-Mariano Test* was used for comparing the accuracy of the predictions [12].
Nevertheless, time series analysis results in meaningful outcomes if enough values are present, so that prediction methods can be applied.

IV. Improving Early Detection of Significant Faults in Quality Management

This chapter describes the Q-AURA analysis process, the invented concept and its integration into Q-AURA. Q-AURA is a system that monitors faults gathered from warranty information systems, where the data originates at the different car

 ISBN: 978-1-61208-358-2

dealers. The addressed business process is depicted in Figure 1 and encompasses phases from the early development of an engine to the after-sales market.

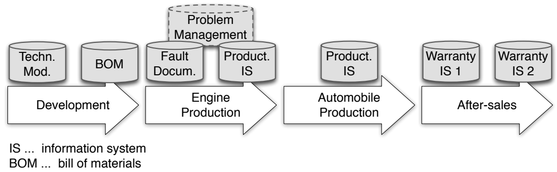

Figure 1: Flow chart of the business process relevant to Q-AURA

It is shown that warranty and fault data of the after-sales market is spread over more than one information system. Since these information systems store partially different data, their integration and combination can provide additional information for determining which faults are developing badly and, thus, have to be investigated further. The figure also presents the involvement of different data sources throughout the whole process in order to get the necessary information of engine components and technical modifications.

A. Q-AURA Approach

First the existing Q-AURA approach is described to explain the underlying analysis method and how the information is processed. Q-AURA provides different steps, each of them is necessary to modify the information in such a way that, finally, potential technical modifications that may cause a particular fault can be determined. Figure 2 illustrates all six steps.

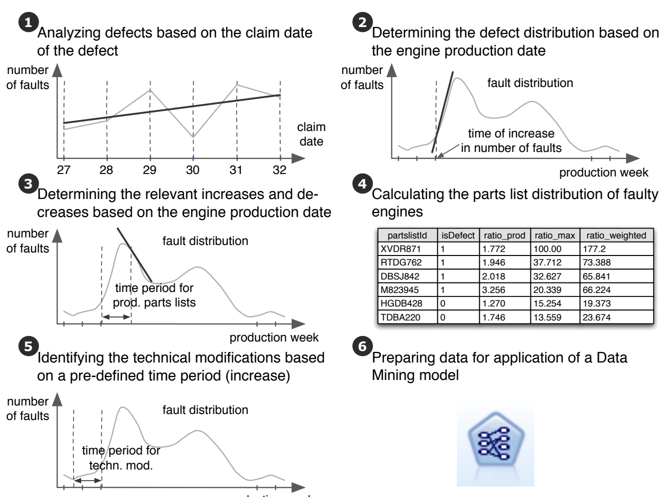

Figure 2: Q-AURA process in detail

The information base for the first step is a set of warranty claims of the last two years from cars produced in the last three years (cf. Figure 2-1). The boundaries were set carefully to determine those cars with corresponding engines that influence the ongoing development process. In the currently used system, faults of warranty claims of the latest six weeks are used to identify current problems with high significance. In order to determine whether the fault is significant or not, a regression analysis is applied [13]. Three different approaches to regression analysis were tested and evaluated containing convex functions, smoothing functions, and a straight line. The evaluation reveals that the straight line approach for regression enables the best results. The *gradient*, *mean value*, and *coefficient of determination* of the regression line are calculated to measure the characteristics of the applied regression. Thresholds, which have been defined and evaluated with the quality management experts, are used to determine if a fault is significant. Significant classified faults are analyzed further. In the second step (cf. Figure 2-2), the production week histogram of engines having a particular fault within the last two years is calculated containing cars produced over the last three years. Next, a normalisation of the fault amount by the total number of produced engines belonging to the same class (cars with the same car brand, fuel type, and engine type) is performed. Afterwards, a 5-point smoothing function is applied to remove the outliers. Significant increases of the resulting smoothed curve are identified to determine those engine production dates, where the ratio between faulty engines and the amount of produced engines of the same type is rising meaning that something changed (e.g., due to a technical modification). In the next step (cf. Figure 2-3), significant decreases are determined. Afterwards, the bill of materials (BOM) distribution of the faulty engines of each period bounded by a significant increase and its subsequent decrease is calculated. The BOM distribution is normalized by the production volumes in order to determine those BOMs that have a bad ratio and, therefore, most likely contain a causing technical modification (cf. Figure 2-4). In the next step (cf. Figure 2-5), the technical modifications of critical BOMs are limited by those that were set operational in a time period before and after the significant increase. Finally, the critical technical modifications are determined using two different methods (cf. Figure 2-6). The first one selects modifications, which most of the critical BOMs contain. The second method uses association rules, in this case the *Apriori* algorithm, for the same task [14], [15].
Q-AURA was already set operational and is used by the quality management experts for their daily work. An evaluation has been done stating that a significant benefit was achieved.

B. Concept of Improving Early Detection

It is obvious that faults that occur anytime during the development and production process should be detected as early as possible. It is not only important because of financial matters, there is also a disadvantage for the reputation of the brand and, consequently, for the company as well. Because of the fact that the analysis of causes can be a time consuming task, even an acceleration by a single day is a big advantage. The proposed approach focusses on detecting faults that develop critically earlier, which means, that Q-AURA can detect potential technical modifications earlier. This results in a reduction of the problem solving time. The proposed improvement uses four main concepts, which are, (i) assessing data quality of information systems storing warranty information, (ii) analyzing and forecasting univariate fault time series, (iii) combining and weighting different predictions and, finally, (iv) evaluating the prediction accuracy followed by an adjustment of the weighting of information sources.
The overall concept is depicted in Figure 3 and consists of the components validator, predictor, combiner, and controller. Each of them has a specific task in the process. First (cf. Figure

 ISBN: 978-1-61208-358-2

3-1), the validator's task is to estimate the data quality of the various information sources. This is done by using different quality metrics. The quality metrics that are used by the presented approach are completeness, soundness, consistency, correctness, and integrity. An overall metric is calculated by a combination of the different factors. Afterwards, the predictor computes forecasts based on the fault and warranty entries of the different information systems (cf. Figure 3-2). As a result an enhanced time series exists containing the predicted value. Linear regression as it is currently used by the Q-AURA approach is applied on the most recent six weeks of the fault time series (containing the newly predicted value). Consequently, the characteristic metrics of the regression analysis performed on each information base is calculated. Subsequently, the combiner's task is to derive the overall prediction whether the fault is significant or not using the parameters of the predictor (cf. Figure 3-3). In order to verify the significance of the prediction, the quality metric of each information system is used. Finally, the controller is necessary to determine the accuracy of the different forecasts (cf. Figure 3-4). This is done by comparing new information system entries of the next week with the predicted values of the predictor. The difference (prediction error) is another weighting factor, which is integrated into the combiner's method. In the next iteration, the combiner uses the newly computed weighting of the controller to adjust the impacts of the information systems. Next, the different concepts of the proposed components are explained in more detail.

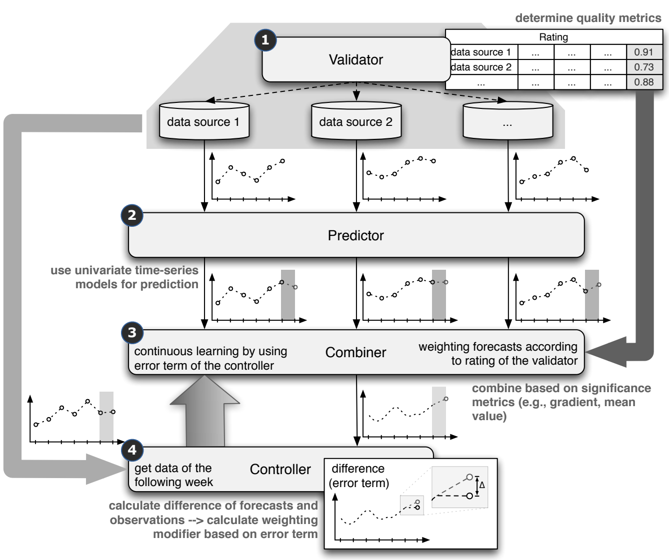

Figure 3: Concept for improving early detection

1) *Validator:* The task of the validator component is determining the overall data quality of the different information systems. Different understanding exists about data quality metrics, therefore, those used are explained. The presented approach uses the quality metrics completeness, soundness, consistency, correctness, and integrity.

Completeness determines if all the information of the real world is captured in the particular data source. The presented approach focusses on closed world assumption which means that any information that is not modelled, is treated as not existing. This means that *Null* values are treated as missing values. Since it is nearly impossible to determine all instances of the real world, an assumption in the proposed approach is made. The real world is approximated by consolidating all instances of the different information systems. The method, which is used for determining the completeness quality metric is relation completeness [2].

Soundness addresses the difference between the real world and the entries of the data sources. This metric indicates if the particular information system stores false values. For an approximation of the real world a combination of the data stored in the different information systems is used. As described in Section III-A, this value is defined as the ratio of the intersection of the information source and the real world and the number of entries in the information source.

Consistency is also a crucial data quality metric and measures the goodness of the entries in the information systems, which is done by proving if duplicates exist or if entries are defined ambiguously. If constraints, referential integrity, and primary keys are applied correctly, this metric can be increased in many cases.

Since consistency does not prove if inserted values are valid, the quality metric correctness is used. This one is very difficult to check, since a software can not automatically prove on its own if a value is correct. Such functional requirements have to be integrated explicitly by implementing concrete rules.

The last quality metric that is applied by the proposed approach is integrity. This metric is assessed by inspecting the join operations between data structures of an information source. This means that if a master data table contains too few entries for a corresponding transaction data table, then this value is decreased.

The overall quality metric is calculated as product of the different single quality metrics, because each feature also influences the other quality metrics (see Figure 4).

data source	completeness	soundness	consistency	correctness	integrity	quality
data source 1	0.85	0.99	0.91	1.00	0.98	0.75
data source 2	0.94	0.94	0.97	0.98	0.96	0.80
data source 3	0.98	0.96	0.89	0.92	0.85	0.65
...	...	...	...	...	...	...

Figure 4: Exemplary results of the validator component

2) *Predictor:* Various time series methods exist to model the behaviour of univariate time series as explained in Section III-B. These methods were applied on warranty and fault information, which is used in Q-AURA. *Goodness-of-Fit* measures and the *Diebold-Mariano Test* were taken for the comparison. The ARIMA method was chosen, since it delivered the best results.

Linear regression is applied on the most recent six weeks of the time series including the predicted value of the next week to determine its characteristics. A straight line is used for fitting, and the features *gradient*, *mean value*, and *coefficient of determination* are determined. The first two parameters can be calculated using following equation (see Equation (3)):

$$y = k * x + d \tag{3}$$

The parameter x determines a point in time on the x-axis of a histogram, while y is the corresponding measured value.

Copyright (c) IARIA, 2014. ISBN: 978-1-61208-358-2

Together these two values determine the coordinates of a data point on the line of regression. The value k is called the gradient and describes the increase between two data points. d is called the offset and is equal to the y-value at the point $x = 0$. The mean value $\bar{y}$ is the average value of the measured points over the six weeks period. In order to describe the steadiness of the curve the *coefficient of determination* is determined [13]. If the regression line depends only on one variable as it is the case if a straight line is used, the coefficient of determination is equal to the square of *Pearson's Correlation Coefficient* r_{xy}^2 (see Equation (4)) [16]:

$$R^2 = r_{xy}^2 = \frac{s_{xy}^2}{s_x^2 s_y^2} \tag{4}$$

After the predictions were performed on each data source, it results in an output as depicted in Figure 5.

data source	gradient	mean value	coeff. o. det.
data source 1	1.92	12.34	0.32
data source 2	2.45	32.91	0.19
data source 3	0.64	24.54	0.98
...	...	...	...

Figure 5: Exemplary results of the predictor component

3) Combiner: In the next step, the combiner uses the previously calculated metrics and determines an overall significance value, which specifies whether the significant fault is critical or not. As explained previously, the characteristic metrics of a regression line can be assessed by using threshold values that have been investigated with experts of the quality management department. There are two different methods available for implementing the combiner component, both based on weighted voting. Weighted voting uses weighting factors to determine the impact of each single data source.

- *Combination of the characteristic metrics*: This option uses the gradient, mean value, and coefficient of determination for the combination task. The relative differences between the gradients of the various data sources and the defined threshold are calculated. Afterwards, the results are combined using weighted voting. The same procedure is applied to the coefficients of determination of the different data sources. Because of the fact that the mean value ensures that sufficient values exist for a robust Q-AURA determination of technical modifications, it is assessed if the mean value of each source exceeds the given threshold. Afterwards, weighted voting is applied on these assessment results.
- *Combination of the significance results*: Before the combination is performed, the characteristic metrics of each data source's regression line is assessed whether the particular fault is significant or not. This results in a significance value for each data source, which is integrated using weighted voting.

Since the first variant (combination of the characteristic metrics) determines the outcome on a more fine-grained basis, it is used in the presented approach.
In the proposed approach, the weighting used by the combiner consists of two components. The first one was already explained earlier and is calculated by the validator component, representing the overall data quality of the various information systems. The second component is explained next and is the accuracy determined by the controller component using the prediction error.

4) Controller: The task of the controller is assessing the prediction quality of the data sources. Since the prediction in the proposed approach is always a one step forecast, which means that only one future value (one week ahead) is calculated, the validation can be done in the following week. Thus, the grading and impact of the actual week's computation is based on the controller's results of previous weeks. In order to reduce the impact of outliers, the assessment of the prediction accuracy is not only based on the last single week. The accuracies of the previous weeks will also influence the calculation, which is achieved by integrating the accuracy metric of the previous value. Since the accuracy value of the previous week was calculated using the observation of the last week and the accuracy of the week before, all the previous accuracies influence the actual value, but with decreasing impact (Equation 5).

$$p_t = \frac{1}{2} * \left(\frac{|f_{t-1} - x_t|}{f_{t-1} + x_t} + p_{t-1}\right) \tag{5}$$

The p values in the equation represent the calculated prediction accuracies, while f_{t-1} represents the forecasted value ($t-1$ shows that it is the forecasted value of the previous week), which is assessed using the actual week's observation x_t. The first term considering the accuracy of the actual week (forecast of the previous week) is a modification to the *Symmetric Mean Absolute Percentage Error (sMAPE)*. Since the highest possible percentage of the standard sMAPE is 200%, a factor is applied to reduce this to 100% for better applicability. After the prediction accuracy values were calculated for each data source, the results can be used for the combiner's computation [17], [18], [19].

C. Q-AURA Integration

This section describes the integration of the improved early detection concept into the currently used Q-AURA approach, which was explained in Section IV-A. Figure 6 shows the overall concept with improved early detection of significant faults. Instead of applying the six week regression analysis the whole process described in Section IV-B is performed to determine whether a fault is significant or not (cf. Figure 6-1). Afterwards, the final outcome is each fault's significance factor. The rest of the Q-AURA approach remains the same (cf. steps 2-6 Figure 6).

Currently, the approach is applied on the data sources of BMW and is still in testing phase. The actual results are very promising, but more extensive investigation have to be done before detailed results can be published.

V. Conclusion

The proposed approach in this research work builds upon an existing system Q-AURA, that is already successfully used by the quality management department in their day-to-day work. Q-AURA is a system that monitors engine faults and

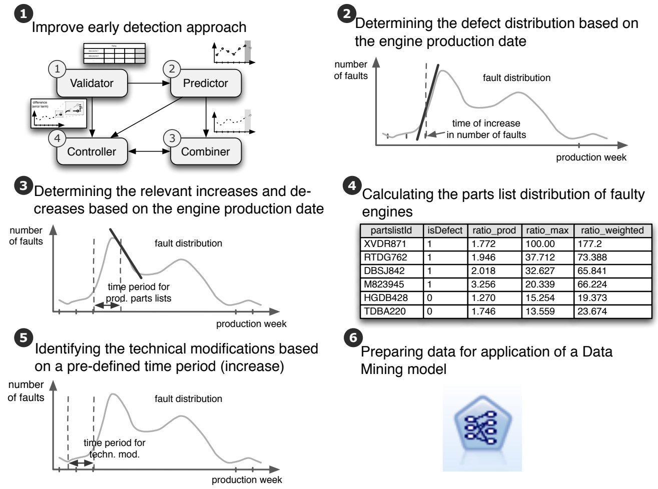

Figure 6: Q-AURA integration of the proposed approach

determines actual problems that develop badly. Afterwards, an automated computation of these faults is performed to find interesting patterns about the cars, resulting in potential technical modifications that may be the cause of faults. Currently, Q-AURA uses linear regression on time series of faults that occurred in the previous six weeks of automobiles that were produced in the last three years.

The contribution addressed in this research work is an approach to detect significant (badly developing) faults earlier by combining predictions of univariate fault time series based on after-sales information, which is stored in different databases. In order to get more accurate results, these different forecasts are weighted according to previous prediction accuracies and data quality metrics of the data sets. The developed technique consists of different components, each of which meets a particular challenge. The validator computes the quality metric for each data source by calculating and combining the metrics completeness, soundness, consistency, correctness, and integrity. Afterwards, the predictor analyzes fault time series based on warranty and claim information of each data source resulting in a forecast of the next week. The controller compares the predictions of the previous week with the observations of the actual week and calculates a weighting factor including the accuracy of previous forecasts. Finally, the combiner integrates the different predictions and determines by weighting and consolidating the values whether the particular fault is significant or not. The approach is already applied and the results are very promising.

REFERENCES

[1] C. Batini, C. Cappiello, C. Francalanci, and A. Maurino, "Methodologies for Data Quality Assessment and Improvement," ACM Computing Surveys, vol. 41, no. 3, July 2009, pp. 1–52.

[2] C. Batini and M. Scannapieco, Data Quality: Concepts, Methodologies and Techniques (Data-Centric Systems and Applications). Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

[3] A. Motro and I. Rakov, "Estimating the Quality of Data in Relational Databases," in In Proceedings of the 1996 Conference on Information Quality. MIT, 1996, pp. 94–106.

[4] Motro, Amihai and Rakov, Igor, "Estimating the Quality of Databases," in FQAS, ser. Lecture Notes in Computer Science, T. Andreasen, H. Christiansen, and H. L. Larsen, Eds., vol. 1495. Springer, 1998, pp. 298–307.

[5] Y. Wand and R. Y. Wang, "Anchoring Data Quality Dimensions in Ontological Foundations," Communications of the ACM, vol. 39, no. 11, November 1996, pp. 86–95.

[6] A. Motro, "Integrity = Validity + Completeness," ACM Transactions on Database Systems, vol. 14, no. 4, December 1989, pp. 480–502.

[7] R. H. Shumway and D. S. Stoffer, Time Series Analysis and Its Applications: With R Examples, 3rd ed. Springer Texts in Statistics, 2011.

[8] P. S. P. Cowpertwait and A. V. Metcalfe, Introductory Time Series with R, 1st ed. Springer Publishing Company, Incorporated, 2009.

[9] A. M. De Livera, R. J. Hyndman, and R. D. Snyder, "Forecasting Time Series With Complex Seasonal Patterns Using Exponential Smoothing," JASA. Journal of the American Statistical Association, vol. 106, no. 496, 2011, pp. 1513–1527.

[10] L. Shenstone and R. J. Hyndman, "Stochastic models underlying Croston's method for intermittent demand forecasting," Journal of Forecasting, 2005.

[11] J. D. Croston, "Forecasting and stock control for intermittent demands," Operational Research Quarterly, vol. 23, no. 3, 1972, pp. 289–303.

[12] F. X. Diebold and R. S. Mariano, "Comparing Predictive Accuracy," Journal of Business & Economic Statistics, vol. 13, no. 3, July 1995, pp. 253–263.

[13] G. U. Yule, "On the Theory of Correlation," Journal of the Royal Statistical Society, vol. 60, no. 4, December 1897, pp. 812–854.

[14] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo, "Fast discovery of association rules," in Advances in Knowledge Discovery and Data Mining, U. Fayyad and et al., Eds. MIT Press, 1996, pp. 307–328.

[15] R. Agrawal, T. Imieliński, and A. Swami, "Mining association rules between sets of items in large databases," in Proceedings of the 1993 ACM SIGMOD International Conference on Management of data, ser. SIGMOD '93. New York, NY, USA: ACM, 1993, pp. 207–216.

[16] M. Mittlböck and M. Schemper, "Explained Variation for Logistic Regression," Statistics in medicine, vol. 15, no. 19, October 1996, pp. 1987–1997.

[17] S. Makridakis and M. Hibon, "The M3-Competition: results, conclusions and implications," International Journal of Forecasting, vol. 16, no. 4, 2000, pp. 451–476.

[18] E. Mangalova and E. Agafonov, "Time Series Forecasting using Ensemble of AR models with Time-Varying Structure," in IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS), May 2012, pp. 198–203.

[19] S. Makridakis, A. Andersen, R. Carbone, R. Fildes, M. Hibon, R. Lewandowski, J. Newton, E. Parzen, and R. Winkler, "The accuracy of extrapolation (time series) methods: Results of a forecasting competition," International Journal of Forecasting, vol. 1, no. 2, 1982, pp. 111–153.

