A DATA MINING APPROACH TO IMPROVE MILITARY DEMAND FORECASTING[1]

Rajesh Thiagarajan, Mustafizur Rahman, Don Gossink and Greg Calbert

Defence Science and Technology Organization, Edinburgh SA 5111, Australia.
firstname.lastname@dsto.defence.gov.au

Abstract

Accurately forecasting the demand of critical stocks is a vital step in the planning of a military operation. Demand prediction techniques, particularly autocorrelated models, have been adopted in the military planning process because a large number of stocks in the military inventory do not have consumption and usage rates per platform (e.g., ship). However, if an impending military operation is (significantly) different from prior campaigns then these prediction models may under or over estimate the demand of critical stocks leading to undesired operational impacts. To address this, we propose an approach to improve the accuracy of demand predictions by combining autocorrelated predictions with cross-correlated demands of items having known per-platform usage rates. We adopt a data mining approach using sequence rule mining to automatically determine cross-correlated demands by assessing frequently co-occurring usage patterns. Our experiments using a military operational planning system indicate a considerable reduction in the prediction errors across several categories of military supplies.

1 Introduction

Identifying and accurately forecasting demand for supply items are vital steps in the planning of a military operation as they facilitate effective decision making. Such demand forecasts, referred to as Advanced Demand Information (ADI) models in the supply chain literature, have been shown to be beneficial in several aspects of supply chain management [2, 3, 4]. In a military context, ADI of supply items should be accurately modelled on the basis of their usage and Rate of Effort (ROE) by military platforms such as ships and aircraft. Consider the logistics planning of a week long military operation O, which involves a ship with a ROE set to 5 hours of sailing per day. Assuming the ship consumes 300 litres of diesel an hour, the diesel ADI model ADI_{diesel} for this operation would consist of a daily demand for 1500 litres of diesel over 7 days resulting in a total demand of 10500 litres of diesel.

Not all stock items in the military inventory have platform-based ROE usage models (referred as ROEM henceforth) that can be used to derive accurate ADI models. A large number of stocks do not have ROEM mainly due to the fact that most stocks in the military inventory are managed through demands aggregated across a number of platforms. For example, demand for a lubricant at a military base over a time period is generally managed by aggregating its demand from a number of military platforms that require the lubricant during that period. While the aggregate demand for supply items at different military bases are captured as a part of the historic demand data, the per-platform usage details are not recorded. Therefore, automatic generation of ROEM is difficult.

---

[1]A preliminary version of this work appeared in ACIIDS 2014 [1]

The lack of ROEM has led to the adoption of autocorrelated demand prediction techniques, such as Simple Exponential Smoothing (SES) [5] and ARIMA [6], in the military operational planning process [7]. In these techniques, the demand of an item for an upcoming operation (e.g., a training exercise) is autocorrelated with the item's usage in the past occurrences of the operation [7]. The demand of an item based on autocorrelation may be assumed to be its ADI model over the operational period. Consider the planning of operation O, which as discussed above, consists of an ADI model of diesel for a ship, but lacks the ADI models of other critical stocks such as lubricants that are essential to the ship's operation. In this case, the demand for a lubricant ADI_l is predicted based on its historic usage by adopting an ARIMA model, denoted as $ADI_l = ARIMA_l$. However, if an impending military operation is (significantly) different from prior campaigns then these prediction models may under or over estimate the demand of critical stocks leading to undesired operational or cost impacts. For example, if the operation O is a high intensity deployment in comparison with previous military campaigns then ADI_l would be underestimated.

To address this, we propose an approach to improve the accuracy of demand predictions by combining autocorrelated predictions with cross-correlated demands of items having known per-platform usage rates. In our approach, the prediction of an ADI model for an item with no ROEM is based not only on the item's historic demand but also takes into account the item's correlation with other items having known ROEM. Our approach builds on the basic premise that the demand for certain supply items are correlated in the context of a military operation. For example, ADI_l and ADI_{diesel} are correlated in the context of operation O because typically there would a surge in lubricant usage as a part of the ship's maintenance routine after a surge in diesel consumption during the operation. In our approach, the ADI model of the lubricant ADI_l required for the ship is predicted by combining $ARIMA_l$ with the lubricant's correlation with ADI_{diesel}.

A key concern of our demand modelling approach is the identification of correlation between demand for supply items. With over a million unique stock items in the military inventory, man-ually identifying correlation in demand between items is infeasible. Therefore, we adopt a data mining approach using sequence rule mining to determine frequently co-occurring usage between supply items from historic demand data. Our approach combines the results from sequence rule mining with time series regression analysis to derive correlated ADI models. We illustrate the effectiveness of our approach by predicting unknown ADI models in a military operational planning system. Our experimental evaluation indicates that incorporating demand correlations in the ADI forecasting process considerably reduces the prediction errors across several categories of military supplies, improving the accuracy of the demand planning process.

The rest of the paper is organised as follows. Section 2 provides an overview of the related work followed by a brief discussion on the existing demand modelling process in Section 3. Our approach to extend the demand modelling process with sequence rule mining is presented in Section 4. The evaluation of our extended demand modelling process within a military operational planning system is discussed in Section 5. We share our results from our experimental evaluation in Section 6. Section 7 summarises our contribution and future work.

2 Related work

A large body of work that deals with the advantages of ADI models in several aspects of supply chain management exists, such as [2, 3, 4]. The approaches in [2, 3] emphasise the utility of ADI models, which are established through market research, inputs from sales managers and advance (but imperfect) orders, to effectively realign manufacturing processes according to the ADI models. Existing works such as [4] have shown that military planning systems are a promising source of ADI models particularly when ROEM exist. However, the problem of generating ADI models for items with no ROEM has not received as much attention. To the best of our knowledge, there are no studies that address the problem of missing ADI model generation in the military context. We address this gap by extending a military operational planning system that generates a few ADI models on the basis of their ROEM. The extension's goal is to leverage

the generated models to improve the ADI prediction accuracy of a large number of items with no ROEM.

In the absence of ADI models, demand prediction techniques that use auto-correlation may be used to estimate the demand for critical military supplies for impending military operation based on their past usage history [7]. However, even items with continuous non-intermittent demands may be under or over estimated if the impending military operation is different from its past occurrences. In our approach, we complement an existing auto-correlated demand prediction technique with ADI correlation and show the resultant reduction in the prediction errors across several categories of military supplies.

3 The military demand modelling process

Accurate ADI models are essential to undertake effective decision making at the planning stage of a military operation. The work in [4] shows that a military operational planning system can be used to estimate the ADI models for an operation. Figure 1 shows the military demand modelling process. The planning system allows a logistics planner to specify the force elements including platforms and personnel to be used in an operation, locations and routes in the operation, a schedule of activities including resource allocations, and the supply chains used to sustain the operation [4].

Apart from troop and equipment movements, the logistics plan of an operation consists of ADI models to effectively procure and distribute supplies to the operation. As shown in Figure 1, the planning system utilises pre-specified ROEM to generate the ADI models for *Item*1 and *Item*2. The ADI models generated by the planning system facilitate a variety of decision support tasks including assessing plan feasibility, logistics sustainability, and risks or weak points in the operational plan. However, due to the lack of ROEM, ADI models for many critical items (e.g., *Item*3 and *Item*4 in Figure 1) are predicted on the basis of their historic demand instead of being generated by the planning system. This inhibits effective and comprehensive military operational planning.

4 Demand modelling process extended with correlated ADI models

The ideal solution to the lack of ADI models would be to establish ROEM for all stock items. However, such an effort would be an enormous, data-intensive and time-consuming undertaking with high costs due to the size of the military inventory. As mentioned earlier, automatic generation of ROEM is difficult as the historic demand data does not record per-platform usage details. Therefore, an approach to improve the accuracy of ADI models available for analyses in the absence of ROEM is required.

In response, we propose an extension to the existing demand modelling process where the accuracy of an item's ADI autocorrelation-based prediction is improved by combining with the item's cross-correlations with known ADI models. Our approach builds on the basic premise that the demand for certain supply items are cross-correlated in the context of a military operation. The correlation in demands may arise from a range of factors including complementary relationships (e.g., a gear lubricant oil and an engine oil), part-of relationships (e.g., an engine and a fuel pump), dependence (e.g., a gun and ammunition), and operational circumstances (e.g., an operation at a tropical region requires both anti-malarial drugs and repellents). As a result, in some cases, when the demands of 2 supply items are correlated and if only one item's ADI model is generated by the military operational planning system then their correlation may potentially be leveraged to improve the accuracy of the second item's predicted ADI model.

Our demand modelling approach presented in Figure 2 extends the basic process discussed in Section 3 with following three additional steps.

1 Identification of cross-correlation in demands using the historic demand data

2 Modelling the cross-correlation in demands

3 Combining autocorrelation with cross-correlation to improve ADI prediction accuracy

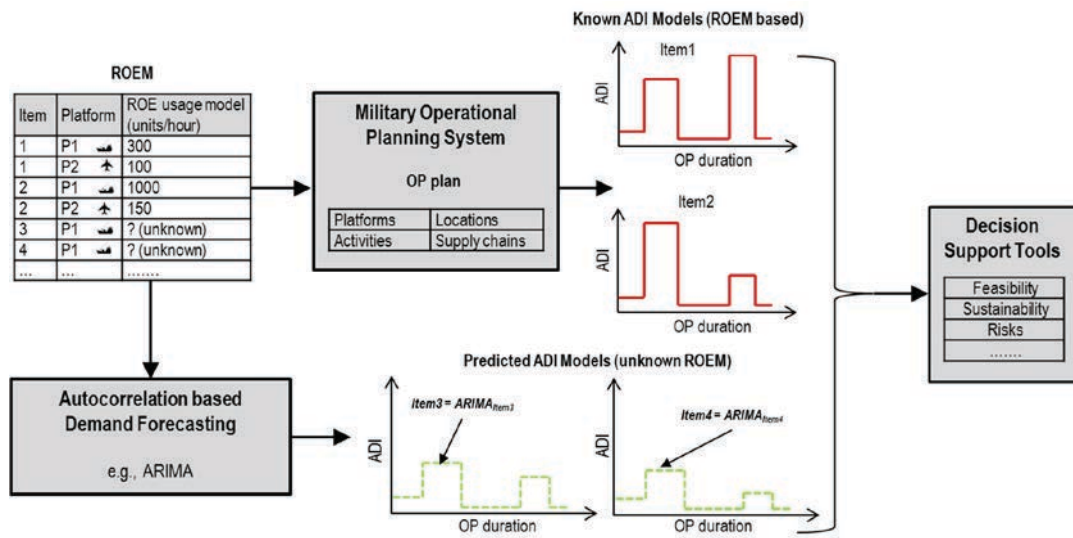The 3 steps in the extended demand modelling approach are detailed below.

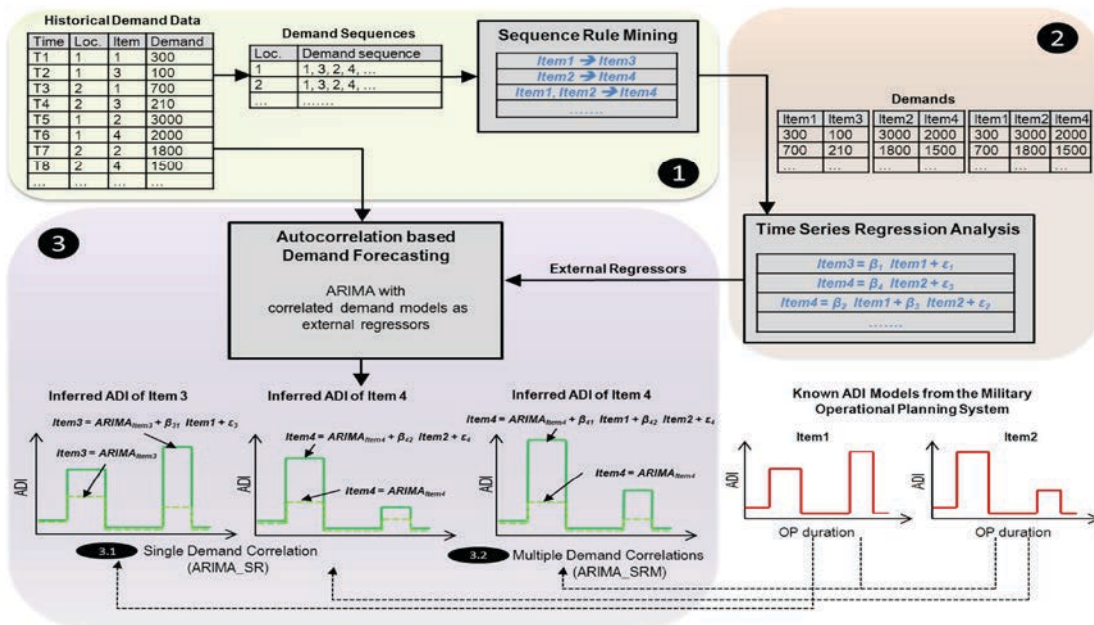**Figure 1**. Military demand modelling process

Figure 2. Demand modelling process extended with correlated ADI models

4.1 Sequence rules to identify correlated demands (Step 1)

The key challenge of our demand modelling approach is the identification of cross-correlation in demands. Specifically, the challenge is to determine frequently co-occurring usage between supply items from the historic demand data. We adopt a data mining approach to address this issue because the size of the military inventory makes manual discovery of correlations infeasible. Widely used data mining methods such as association rule mining [8] and all-pairs correlation techniques [9] are unsuitable to this problem mainly because the historic demand data is not a market basket database. Although conversion to a market basket database using time period based sampling is possible (e.g., all demands on a day can be treated as one market basket transaction), previous works like [9] have noted that such methods are vulnerable to false-positive and false-negative correlations because they do not take into account possible lags between the demands of correlated items. For example, unless the whole duration of the operation O discussed in Section 1 is considered as a single market basket transaction, the lagged correlation between demand for diesel and the lubricant for post-operational maintenance would be ignored.

To address this, we transform the identification of correlation in demands as a sequence rule mining problem, where the goal is to discover sequential rules from a database of sequences [10]. A sequence database is of the form $S = \{s_1, \ldots, s_n\}$, where each sequence s_i consists of chronologically ordered itemsets $\{\{i_1^1, \ldots, i_x^1\}, \ldots, \{i_1^m, \ldots, i_y^m\}\}$. The sequence rule mining process over S returns a set of sequence rules of the form $X \rightarrow Y$, where $X \subset s_i$ and $Y \subset s_i$ are disjoint itemsets such that Y occurs after X in S with a certain support and confidence [10]. While users are allowed to set a minimum support *minsup* value to filter infrequent rules, in most application domains, however, it is difficult to ascertain an optimal *minsup* before the mining process. Therefore, we adopt the rule mining algorithm in [10] that allows the users to efficiently search for *top-k* sequence rules with a certain minimum confidence. The algorithm works by a process called RuleGrowth, where small sequence rules are recursively expanded by adding frequent itemsets to the left and right side of the rules. The process continuously updates a *top-k* rules set when new rules with higher support are found.

Step 1 in Figure 2 shows the adoption of a sequence rule mining process in our extended demand modelling process to identify correlated demands. Firstly, a sequence database is generated from the historic demand data. Each sequence is formed by chronologically ordering the demand transactions at a military base. The rationale behind location-based sequence formation is that if the demand for two items co-occur frequently (with or without lag) across a number of military bases then it is likely that their demands are correlated. The demand sequences table in Figure 2 shows the location-based sequence database generated from the historic demand data. For example, the demand sequence at location 1 consists of items $\{1, 3, 2, 4, \ldots\}$ that occur at times $\{T1, T2, T5, T6, \ldots\}$. The demand sequences are provided as input to the sequence rule mining process [10] to generate sequence rules that are above a user-specified confidence threshold across all locations. Figure 2 shows the sequence rules generated $\{\{Item1\} \rightarrow \{Item3\}, \{Item2\} \rightarrow \{Item4\}, \{Item1, Item2\} \rightarrow \{Item4\}, \ldots\}$. The rule $Item1 \rightarrow Item3$ implies that if a military base demands $Item1$ then it is likely that in the future it will be followed by a demand for $Item3$, indicating a potential correlation in the demands of items 1 and 3. Note that, unlike all-pairs correlation [9], the sequence rules also capture cross-correlations between several items (e.g., $\{Item1, Item2\} \rightarrow \{Item4\}$).

In some data mining domains (e.g., retail) it is sufficient to just identify potential correlations between items as this may already be enough to effectively adapt marketing tactics. In the context of ADI model generation, however, modelling the characteristics of a correlation is equally critical because it is important to quantify how the demand for one or more supply items impact the demand of a correlated item. To this end, we adopt time series regression analysis to model the correlated demands.

4.2 Time series regression to model correlated demands (Step 2)

While a sequence rule indicates a potential correlation, in order to quantify the relationship between items it is important to consider the de-

mand quantities of the correlated items over time. Time period based sampling (e.g., daily demand) can be used to represent the chronologically ordered demands of an item as a time series. Time series regression analysis, a well known statistical method to model the relationship between time series variables, is one way to model the correlated demands. A linear regression model of the form $Y_t = \beta_1 X_t^1 + \cdots + \beta_n X_t^n + \varepsilon$ models the relationship between a dependent variable Y at time t and independent variables $\{X_t^1, \ldots, X_t^n\}$ at t using the regression coefficients $\beta_1, \ldots, \beta_n$, and an error term ε.

The problem of selecting suitable independent variables X from a large pool of potential variables is a common issue in regression analysis. Methods such as stepwise regression [11], which sequentially add/remove variables to a regression model based on a scoring criteria, are suitable if the pool of independent variables is small but do not scale well to larger pools. In the context of ADI model generation, the sequence rules generated in step 1 can be used to substantially reduce the pool of independent variables available to conduct stepwise regression. All items in the antecedent part of a sequence rule are considered as a part of the pool of independent variables to predict the unknown ADI models of items in the subsequent part of the rule. Step 2 of Figure 2 shows the regression analysis process conducted, leveraging the sequence rules to generate linear (or non-linear) models of the correlated demands. For every item in a demand correlation identified by a sequence rule, time period based sampling is performed to generate the item's demand time series. For example, the demands table in step 2 of Figure 2 shows the time series creation of items $\{1,3\}$, $\{2,4\}$, and $\{1,2,4\}$. The regression analysis results include a linear model $Item3 = \beta_1 Item1 + \varepsilon_1$ quantifying the demand for $Item3$ from the demand for $Item1$. A user-defined threshold for the coefficient of determination R^2 [12] is used to filter poorly correlated models.

Recalling the issue of false-negatives and false-positives with time period based sampling raised in Section 4.1, the problem was due to the inability to deal with lagged correlations. We address this issue in step 2 by allowing lagged variants of known ADI models to be part of the pool of independent variables available to the regression analysis. For example, if the monthly demand of a lubricant l_t is dependent on a truck's ROE (i.e., diesel demand) over the last 6 months $\{d_{t-1}, \ldots, d_{t-6}\}$ then it can be modelled as $l_t = \beta_1 d_{t-1} + \cdots + \beta_6 d_{t-6} + \varepsilon$. The weighted sum of the time lagged values has an effect similar to a Finite Impulse Response Filter[13] in removing some frequency components. Due to the additional possible delays and weights there is potential for overfitting in the ADI prediction model. This can be ameliorated by applying regularisation techniques commonly applied in machine learning (regression).

4.3 Combining Auto-correlation and ADI correlation (Step 3)

Recall the ARIMA-based autocorrelated ADI prediction of $Item3$ and $Item4$ from Figure 1. In the final step (Step 3) in our extended demand modelling process in Figure 2, we incorporate the cross-correlated demand models within the ARIMA prediction technique to produce a combined prediction. Specifically, when predicting the ADI model of an item with no ROEM (e.g., $Item3$), we use ADI models discovered in Step 1 and Step 2 in Figure 2 as external regressors in the ARIMA model. It is important to note that both step 1 and 2 may be performed off-line prior to the planning stage.

In Step 3 of the extended demand modelling process, we consider 2 cases of incorporating cross-correlated ADI models in the ARIMA-based prediction process: single demand correlation and multiple demand correlation.

Single Demand Correlation (ARIMA_SR)

Single demand correlation combines an ARIMA-based autocorrelated prediction technique with a single cross-correlated ADI model, referred to as ARIMA_SR [1]. When an ARIMA_SR model is used to predict an item's ADI model then the prediction is not only based on the item's past usage but also on its correlation with one other known ADI model.

The ARIMA_SR process is illustrated in Step 3.1 of Figure 2. The demand for $Item3$, for example, is predicted by combining its ARIMA prediction $ARIMA_{Item3}$ with a single cross-correlated ADI model $Item1$. Candidates for cross-correlation are selected in Step 2 of the extended demand mod-

elling process as discussed in Section 4.2. The ARIMA_SR prediction of *Item*3 is modelled as $Item3 = ARIMA_{Item3} + \beta_{31}Item1 + \varepsilon_3$.

Multiple Demand Correlation (ARIMA_SRM)

The demand for some supply items may be cross-correlated with more than one ADI model. We refer to this case as multiple demand correlation (ARIMA_SRM), shown as Step 3.2 in Figure 2. Let us consider the prediction of a gear lubricant's ADI model ADI_l based on combining $ARIMA_l$ and a cross-correlated diesel ADI model ADI_{diesel}. The dependence that underlies the cross-correlation between diesel and lubricant usage is discussed in Section 1. However, the lubricant's demand may also be cross-correlated with other known ADI models, for example brake fluid ADI_{bf}, because typically a vehicle's maintenance process includes a change of lubricants and brake fluid. Therefore, ADI_l predictions may potentially be improved by also considering the known ADI_{bf}.

The ARIMA_SRM process is illustrated in Step 3.2 of Figure 2. The demand for *Item*4, for example, is predicted by combining its ARIMA prediction $ARIMA_{Item4}$ with multiple cross-correlated ADI models *Item*1 and *Item*2. The ARIMA_SRM prediction of *Item*4 is modelled as $Item4 = ARIMA_{Item4} + \beta_{41}Item1 + \beta_{42}Item2 + \varepsilon_4$.

5 Evaluation of the extended demand modelling process

To evaluate, we used the SES, ARIMA, ARIMA_SR, ARIMA_SRM techniques to predict the demand of items required for military training exercises.

5.1 Implementation

To illustrate our extended demand modelling approach, we have implemented a software system that uses SES, ARIMA, ARIMA_SR, and ARIMA_SRM to infer ADI models for items with no ROEM. The system is implemented as an extension to a military operational planning system.

Consider the planning of a military contingency operation. The mission's logistics plan using the military operational planning system would con-

sist of ADI models of items with ROEM. Based on this initial plan, our implementation offers the planner a set of improved ADI predictions that are derived from a library of autocorrelated and cross-correlated ADI models. The library of correlated demand models used in our system is developed prior to the planning process. As discussed in Section 4.1, the identification of demand correlation is transformed into a sequence rule mining problem that is solved using the *top-k* rule algorithm [10] in the Sequential Rule Mining Framework (SPMF) [14]. The linear models quantifying the cross-correlated demands are generated using the *dynlm* package in the statistical system *R* [15]. The ARIMA models are generated using the *forecast* [16] package in *R*.
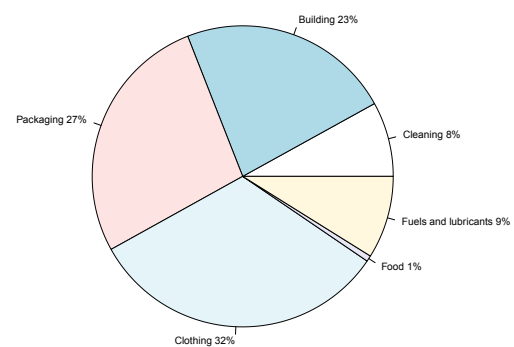

Figure 3. Inventory Profile

5.2 Inventory profile

Our initial analysis indicates that the military operation planning system can only generate ADI models for some frequently used military consumables in the following inventory profiles:

– fuels and lubricants (e.g., engine oil)

– food (e.g., combat ration packs)

– clothing (e.g., combat boots)

– cleaning products (e.g., engine cleaning brush)

– building products (e.g., aircraft sealant)

– packaging (e.g., fuel drums)

The relative proportion of the military inventory profile under consideration is presented in Figure 3.
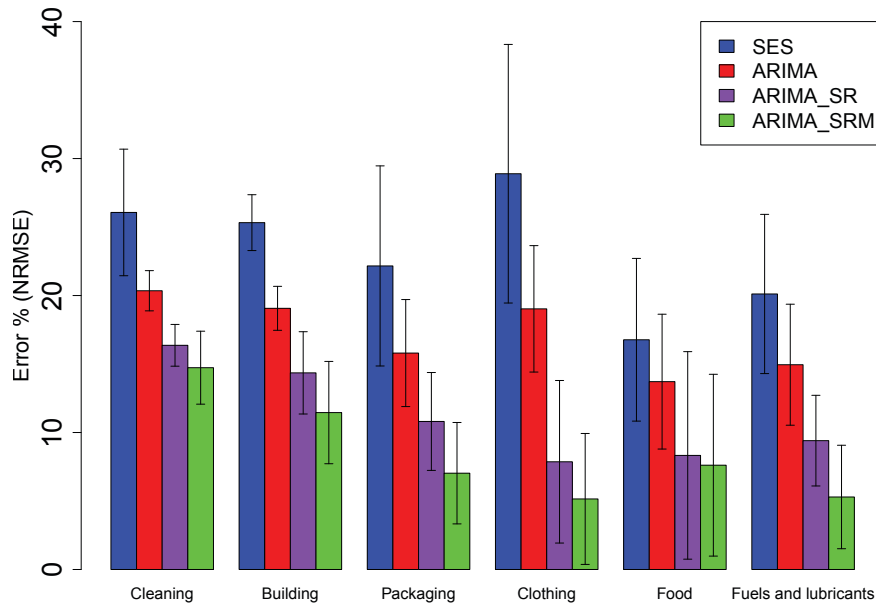
Figure 4. Average Demand Prediction Errors for SES, ARIMA, ARIMA_SR and ARIMA_SRM

5.3 Experimental set-up

We used the SES, ARIMA, ARIMA_SR, ARIMA_SRM techniques to predict the demand of items across the categories of military consumables mentioned in Section 4.3. We used the demand data from military training exercises conducted over the last 20 years in these experiments. The first 10 years of data from the training exercises were used to train the demand predictors, while the remaining data was used for testing the prediction accuracy. To quantify the accuracy of a predicted ADI model, we use the Normalised Root Mean Square Error (NRMSE) metric defined as,

$$NRMSE = \frac{RMSE}{max(ADI_a) - min(ADI_a)}$$

where *RMSE* is the standard Root Mean Square Error [17] and ADI_a is the actual ADI model.

Our evaluation included two sets of experiments. The first set of experiments were conducted to assess whether ARIMA_SR predictions were more accurate than those that are predicted by SES and ARIMA. The second set of experiments evaluated whether adopting ARIMA_SRM over ARIMA_SR improved the accuracy of ADI predictions.

6 Results and discussion

We present the results from our experiments in this section.

6.1 ARIMA_SR evaluation

The results from our experiments to evaluate ARIMA_SR are presented in Figure 4. Each group of bars denotes the average NRMSE resulting from the adoption of SES, ARIMA, ARIMA_SR and ARIMA_SRM for ADI prediction in an inventory segment. Figure 4 shows that on average the ARIMA_SR model's NRMSE is lower than SES and ARIMA across all inventory segments. When the single standard deviations shown as error bars in Figure 4 are taken into consideration, the ARIMA_SR technique performs better than the SES and ARIMA techniques in terms of NRMSE in 3 out of the 6 categories of military consumables.

6.2 ARIMA_SRM evaluation

Figure 4 also presents the results from our experiments to evaluate whether adopting ARIMA_SRM instead of ARIMA_SR improves the accuracy of the ADI predictions. From the results it is evident that on average the ARIMA_SRM model's NRMSE is lower than ARIMA_SR across all inventory segments. The single standard devi-
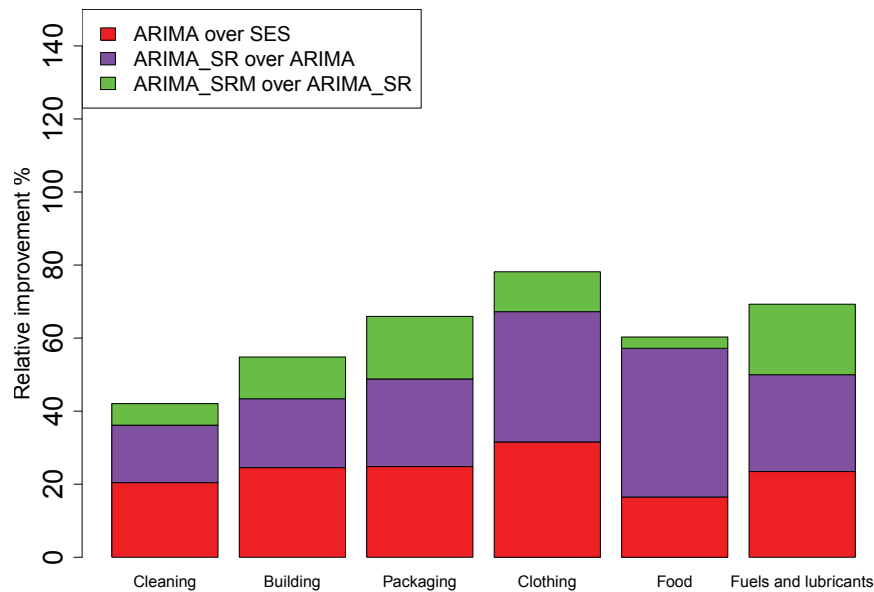
Figure 5. Relative improvements in prediction accuracies

ations indicate that the ARIMA_SRM is more accurate than the SES and ARIMA in all categories of military consumables except 'Food'. We conducted further analysis to explain the results in the 'Food' inventory segment. From our analysis we were able to conclude that as the 'Food' segment consists only 1% items in the experiment's inventory profile the results were more sensitive to variance in the NRMSE values.

Furthermore, we also measured the relative improvement achieved by ARIMA over SES, ARIMA_SR over ARIMA, and ARIMA_SRM over ARIMA_SR. The relative improvement, that is the reduction in NRMSE, of technique T over T' is measured as

$$\frac{\overline{NRMSE}_{T'} - \overline{NRMSE}_{T}}{\overline{NRMSE}_{T'}} \times 100.$$

The relative improvement results shown as a stacked bar chart in Figure 5 indicate that by adopting ARIMA_SRM the ADI prediction accuracy can be improved over ARIMA by at least 20% across all stock categories. Our experiments indicate that in the biggest inventory segment 'Clothing' (32% of the inventory profile as shown in Figure 3) the relative improvement of ARIMA_SRM over ARIMA is about 47%.

From the results presented in Figures 4-5, it can be concluded that combining correlated ADI models discovered using sequence rule mining with autocorrelated demand prediction improves the accuracy of military demand predictions.

7 Summary and future work

We presented an extended demand modelling approach for military operational planning, where the accuracy of ADI prediction for critical stocks is improved by combining autocorrelated predictions with cross-correlated ADI models of items with known ROEM. We highlighted the importance of accurate ADI models to conduct comprehensive and effective operational analysis and decision making. We presented an approach based on sequence rule mining to identify (possibly delayed) correlation in demands. We also showed how regression models are used to quantify the demand correlations. An experimental evaluation of our approach was conducted to show the improvement in the military demand modelling process.

Conducting a full-scale user study to further validate ADI model generation is part of our current and future work. We also plan to explore the prospect of probabilistic cross-correlations with ADI models generated by the military operational planning system.

References

[1] R. Thiagarajan, M. Rahman, G. Calbert, and D. Gossink, "Improving military demand forecasting using sequence rules," in *6th Asian Conference on Intelligent Information and Database Systems (ACIIDS)*, pp. 475–484, 2014.

[2] S. Benjaafar, W. L. Cooper, and S. Mardan, "Production-inventory systems with imperfect advance demand information and updating," *Naval Research Logistics (NRL)*, vol. 58, no. 2, pp. 88–106, 2011.

[3] F. Karaesmen, "Value of advance demand information in production and inventory systems with shared resources," in *Handbook of Stochastic Models and Analysis of Manufacturing System Operations*, vol. 192, pp. 139–165, 2013.

[4] R. Thiagarajan, M. A. Mekhtiev, G. Calbert, N. Jeremic, and D. Gossink, "Using military operational planning system data to drive reserve stocking decisions," in *29th IEEE International Conference on Data Engineering (ICDE) Workshops*, pp. 153–162, 2013.

[5] E. Gardner, "Exponential smoothing: The state of the art Part II," *International Journal of Forecasting*, vol. 22, no. 4, pp. 637–666, 2006.

[6] G. Box, G. Jenkins, and G. Reinsel, *Time Series Analysis: Forecasting and Control*. 2008.

[7] M. Downing, M. Chipulu, U. Ojiako, and D. Kaparis, "Forecasting in airforce supply chains," *International Journal of Logistics Management*, vol. 22, no. 1, pp. 127–144, 2011.

[8] R. Agrawal, T. Imieliński, and A. Swami, "Mining association rules between sets of items in large databases," *SIGMOD Rec.*, vol. 22, no. 2, pp. 207–216, 1993.

[9] H. Xiong, W. Zhou, M. Brodie, and S. Ma, "Top-k correlation computation," *INFORMS Journal on Computing*, vol. 20, no. 4, pp. 539–552, 2008.

[10] P. Fournier-Viger and V. S. Tseng, "TNS: mining top-k non-redundant sequential rules," in *ACM Symposium on Applied Computing (SAC)*, pp. 164–166, 2013.

[11] L. Wilkinson, "Tests of significance in stepwise regression," *Psychological Bulletin*, vol. 86, no. 1, pp. 168–174, 1979.

[12] R. H. Myers, *Classical and modern regression with applications*, vol. 2. 1990.

[13] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*. Prentice–Hall, 1989.

[14] P. Fournier-Viger, A. Gomariz, A. Soltani, and T. Gueniche, "SPMF: Open-Source Data Mining Platform - http://www.philippe-fournier-viger.com/spmf/," 2013.

[15] A. Zeileis, *dynlm: Dynamic Linear Regression*, 2013. R package version 0.3-2.

[16] R. J. Hyndman, G. Athanasopoulos, S. Razbash, D. Schmidt, Z. Zhou, Y. Khan, and C. Bergmeir, *forecast: Forecasting functions for time series and linear models*, 2013. R package version 4.06.

[17] C. J. Willmott, "On the validation of models," *Physical Geography*, vol. 2, no. 2, pp. 184–194, 1981.