

A novel decomposition approach for on-line lot-sizing

Citation for published version (APA):

Aarts, E. H. L., Reijnhoudt, M. F., Stehouwer, H. P., & Wessels, J. (1998). *A novel decomposition approach for on-line lot-sizing*. (Memorandum COSOR; Vol. 9821). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/1998

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

A novel decomposition approach for on-line lot-sizing

E.H.L. Aarts, M.F. Reijnhoudt, H.P. Stehouwer, J. Wessels

September 10, 1998

Abstract

A novel decomposition approach is presented for a class of on-line problems of Wagner and Whitin's economic lot-sizing type. The decomposition is based on the fact that optimal plans contain regeneration points and that the plan between two regeneration points is independent of the rest of the plan. This property is exploited in the following way: first, estimate the next regeneration point and second, determine an optimal plan up to that point. The estimation of a next regeneration point can be done with an artificial neural network or with a statistical classification method.

We present extensive experimental comparisons of the novel approach with more classical ones. The main conclusion from the results is that our approach dominates all other approaches with respect to robustness, performance, and data efficiency. Only in those cases where the demand is known for a large number of periods in advance some classical approaches perform better.

1 Introduction

Most on-line planning problems are difficult to handle satisfactory, and on-line lot-sizing problems of the Wagner-Whitin type are no exception to this rule (Silver, Pyke & Peterson [1998]).

In this paper we present a novel approach for on-line lot-sizing problems which is more immune from demand characteristics, cost calculation, and amount of known future demands. The approach follows a divide-and-conquer method which is based on essential properties of optimal plans. In fact, the plan which we would like to

propose is the one which appears to be optimal when all the demands are known: the optimal off-line plan. For lot-sizing problems this optimal off-line plan is characterized by regeneration points which divide the optimal plan into subplans. The production quantities of each subplan may be determined independently of all other subplans.

We use this property for developing a decomposition approach for the on-line version in the following way. Estimate the next regeneration point, and determine an optimal plan. Repeat the procedure starting just after the first subplan.

The determination of the first optimal subplan turns out to be straightforward, if the next regeneration point is known. However, the estimation of the next regeneration point is a new type of problem. Here, one may exploit the fact that afterwards it is always easy to determine what decision would have been optimal. If a planning horizon (Lundin & Morton [1975]) exists, only a limited number of future demands suffices to locate the optimal next regeneration point. This feature provides the basis for a learning or a statistical approach. In this way, learning examples may be drawn from reality or from simulation as basis for these approaches.

In the present paper we consider three cost structures for Wagner and Whitin type lot-sizing problems in their on-line versions. The first cost structure is the standard Wagner-Whitin cost structure. In the other two cost structures the production capacity for each time period is limited. In the second cost structure there is the possibility of using overtime production for exceeding the capacity limit. The third cost structure has the possibility of purchasing externally.

We use multi-layered perceptrons as learning technique for estimating the next regeneration point. This learning approach is compared with the K -nearest-neighbors method, which is a well-known method for statistical classification. We use several types of myopic approaches and two approaches based on forecasting demands to compare the performance of our approach.

The remainder of the paper is organized as follows. In the next section the lot-sizing model is introduced together with the cost structures. In Section 3 our decomposition approach is presented in more detail. Section 4 gives the experimental set-up and Section 5 contains an overview of the experimental results. Section 6 discusses the results and it gives some concluding comments.

2 Lot-sizing

First we introduce the on-line lot-sizing problem with the different cost structures. For solving this problem, we need solutions to the n -period problem and the off-

line problem, which are described in Subsection 2.2 and Subsection 2.3, respectively.

2.1 The on-line lot-sizing problem

In on-line models it is often assumed that decisions are made without any knowledge of the future (Karp [1992]). Lot-sizing is basically production planning, so it is more realistic to adopt the view that there is some, but limited knowledge of future demands. Therefore, in the on-line model, we suppose that demands are only known for the next m periods, where the constant m is called the length of the *data horizon*. The problem is to find a production quantity (lot size) for the next period, based on the initial inventory level and the demands for period 1 to m only. Then this lot size is produced, incurring certain production cost, and the demand is satisfied. When products have to be kept in inventory to the next period, some holding cost must be paid. After that, the first demand is eliminated and replaced by a new demand, appended to the end of the horizon, and the procedure restarts. Note that at the beginning of each restart, again, demand is known for exactly m periods.

In this on-line model, it is assumed that production for unknown future demands is forbidden. We are looking for approaches that carry out this task with the lowest cost. The cost for one period consists of the two, above mentioned parts, i.e.,

- Production cost, $P : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, for producing a lot size X . P is strictly increasing, so there is no benefit in producing more than necessary.
- Holding cost, $H : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, related to carrying inventory of size I from one period to the next.

The pair (H, P) is called the *cost structure*. In our experiments we use three different cost structures, which are described in the following subsections.

2.1.1 Wagner-Whitin cost structure

The Wagner-Whitin cost structure originates from the single-source model described by Wagner & Whitin [1958]. The pair (H, P) is given by

$$H(I) = hI$$

and

$$P(X) = \begin{cases} 0 & \text{if } X = 0 \\ S + pX & \text{if } X > 0 \end{cases},$$

where h denotes the holding cost per unit product per period, S denotes the setup cost, and p denotes the production cost per unit product.

2.1.2 Cost structure with overtime

The cost structure with overtime originates from the two-source model mentioned by Jagannathan & Rao [1973]. Demand is satisfied either by regular production or by production during overtime. In this cost structure the production cost is given by

$$P(X) = \begin{cases} 0 & \text{if } X = 0 \\ S + pX & \text{if } 0 < X \leq C, \\ S + pC + q(X - C) & \text{if } X > C \end{cases}$$

where C denotes the regular time production capacity, S denotes the setup cost, p denotes the regular time production cost per unit product and q denotes the overtime production cost per unit product. It is assumed that overtime production is more expensive than regular time production ($q > p$). As in the Wagner-Whitin cost structure we take $H(I) = hI$.

2.1.3 Cost structure with purchasing

The third cost structure is based on the model that demand is satisfied either by in-house production with setup cost S and a fixed production cost per unit product p or by purchasing from an outside supplier without setup cost, but at a higher unit price q . Then, it is clear that small production quantities are entirely satisfied by purchasing. The break-even point between purchasing and producing, B , is given by $B = S/(q - p)$. So the production cost is given by

$$P(X) = \begin{cases} qX & \text{if } 0 \leq X \leq B \\ S + pX & \text{if } B < X \leq C, \\ S + pC + q(X - C) & \text{if } X > C \end{cases}$$

where C denotes the in-house production capacity. Analogue to the cost structure with overtime, we assume that purchasing externally is more expensive than in-house production ($q > p$). As in the two other cost structures we take $H(I) = hI$.

2.2 The n -period lot-sizing problem

The basis for solving the on-line lot-sizing problem is the n -period problem. For an overview of results on the n -period lot-sizing problem, see Aggarwal & Park [1993] or Bahl, Ritzman & Gupta [1987]. This problem can be formulated as follows. Given demands $d_t \geq 0$ for the discrete time periods $1, \dots, n$, find a *production plan* with lot sizes X_1, \dots, X_n , such that the cost of that production plan,

$$\sum_{t=1}^n [P(X_t) + H(I_t)],$$

is minimized, where I_t is the *inventory level* at the end of period t given by

$$I_t = I_0 + \sum_{s=1}^t (X_s - d_s) \geq 0,$$

and I_0 is the initial inventory level.

A vector (X_{u+1}, \dots, X_v) of lot sizes with $0 \leq u < v$ is called a *subplan* (of a larger plan) if $I_u = I_v = 0$, and $I_t > 0$ for $t = u + 1, \dots, v$. We say there is a *regeneration point* at the end of period v . So there is a regeneration point whenever the inventory level is zero.

The famous inventory decomposition property of Wagner & Whitin [1958] states that if there exists an optimal n -period plan with a regeneration point at the end of period t with $0 < t < n$, then the n -period plan that is obtained by independently finding optimal production plans for the first t periods and the last $n - t$ periods with $I_t = 0$ is optimal.

The n -period problem can be solved by standard dynamic programming, but using this property a smart dynamic programming formulation can be obtained that is much more efficient to execute. Let $c(s, t)$ be the cost of an optimal subplan (X_{s+1}, \dots, X_t) , then the cost of an optimal n -period plan $f(n)$ can be computed by the dynamic program given by:

$$f(t) = \begin{cases} 0 & \text{if } t = 0 \\ \min\{f(s) + c(s, t) \mid 0 \leq s < t\} & \text{if } 1 \leq t \leq n \end{cases}.$$

This property holds for a large class of cost structures, including the ones of Subsection 2.1. For many interesting cost structures the structure of a subplan in an optimal plan has some nice properties, that allows to calculate the cost of an optimal subplan in polynomial time.

Let (X_{u+1}, \dots, X_v) be a subplan in an optimal plan. For the Wagner-Whitin cost structure the structure of the subplan is given by:

$$X_t = \begin{cases} \sum_{t=u+1}^v d_t & \text{if } t = u + 1 \\ 0 & \text{if } u + 1 < t \leq v \end{cases}.$$

For the cost structure with overtime and for the one with purchasing there exist optimal plans that consist only of *well-formed* subplans. In a well-formed subplan all non-zero lot sizes have a size equal to the production capacity C , except for at most one lot size. That production quantity of special size is the first lot size if it is smaller than C or the last non-zero lot size if it is larger than C .

The number of non-zero lot sizes can be easily determined, but the distribution of these lot sizes over the subplan is found by backward recursion. This recursion ensures that lot sizes are produced as late as possible, so minimizing the holding cost. Details regarding the structure of optimal subplans and proofs of optimality can be found in Stehouwer [1997].

Using the special structure of optimal (well-formed) subplans, it is relatively easy to calculate the cost of an optimal subplan.

2.3 The off-line lot-sizing problem

In off-line lot-sizing problems all information about future demands is given in advance, so it can be compared with n -period lot-sizing when $n = \infty$ or with on-line lot-sizing when $m = \infty$.

In many cases it appears that only a limited number of future demands suffices to determine an optimal first regeneration point. If there exists an integer n such that for all $N \geq n$ there exists an optimal N -period plan with a regeneration point at the end of period t , irrespective of demands in periods $n + 1, n + 2, \dots$, then we call t a *simple planning horizon*.

Experiments done by Lundin [1973] showed that in 97.5% of all cases a simple planning horizon was found within 70 periods using the Wagner-Whitin cost structure. He concluded that the higher the demand variability or the average demand, the smaller the expected planning horizon.

If there exists an upper bound on the length of a subplan in an optimal plan, a *forward algorithm* (Federgruen & Tzur [1994]) that solves subsequent n -period problems will certainly find the smallest simple planning horizon if it exists. Fortunately, for many interesting cost structures such an upper bound can be derived.

In our approach we are looking for a regeneration point within the data horizon, because only production on demand is permissible. So we are more interested in plans in which all subplans have a length not exceeding m . Note that in this case the above mentioned derived upper bounds can be replaced by m .

If we consider optimal plans that consist of subplans with length not exceeding m , we can define the m -simple planning horizon similar to the definition of a simple planning horizon. These plans contain information such that afterwards it is possible to indicate what decision would have been the optimal one in an on-line problem with the same demand sequence.

3 A decomposition approach

Our approach to the on-line lot-sizing problem leads to policies which belong to the class of variable-horizon policies. A *variable-horizon policy* consists of the following steps.

Given are an initial inventory equal to zero ($I_0 = 0$) and finite demands d_1, \dots, d_m for $m > 0$ periods.

1. Choose an optimization horizon $t \in \{1, \dots, m\}$.
2. Determine an optimal plan for the periods $1, \dots, t$.
3. Implement the first subplan of that optimal plan, that is until the inventory is again zero.
4. Update the data horizon, renumber the periods and goto 1.

Different variable-horizon policies differ only in the implementation of step 1, the *horizon-selection rule*.

Recall from Subsection 2.3 that if there exists a simple planning horizon, there will be a moment in time that there is a sufficient amount of demand information available to conclude what would have been the optimal optimization horizon.

We propose a decomposition approach that exploits the fact that solutions to the off-line problem can provide learning examples from which it can be learned how to act in the future in the corresponding on-line problem. A *learning example* consists of the demands in the data horizon together with information containing the (afterwards calculated) optimal first regeneration point.

In this way the problem of finding an optimization horizon can be viewed as a classification problem. A demand vector (d_1, \dots, d_m) belongs to class l if and

only if the posterior probability that l is the optimal first regeneration point, given the demands for period 1 to m , is maximal. These posterior probabilities are usually unknown, so the ideal mapping from demands to a regeneration point, called the *target mapping*, is approximated by our horizon-selection rules.

This classification problem provides the basis for the learning and the statistical approach that we discuss below. For both approaches we derive two variable-horizon policies, so we end up with four policies: two MLP-based policies and two KNN-based policies. For each approach, the first policy maximizes the classification rate, i.e., it chooses the optimization horizon which has the highest probability of being the minimal optimal first regeneration point. This policy works with so-called *zero-one target vectors*, which contains a one in the position equal to the minimal optimal first regeneration point and zeros in the other $m - 1$ positions.

The second policy minimizes the excess cost. The *excess cost* is the extra cost incurred when choosing a (non-optimal) first regeneration point instead of the optimal first regeneration point. This policy works with so-called *cost target vectors* in which each of the m positions contains the excess cost for choosing the first regeneration point equal to that position. Details about the construction of these vectors can be found in Stehouwer [1997].

3.1 A neural network approach

In the learning approach we use *multi-layered perceptrons* (MLP). See Hassoun [1995] for an introduction to multi-layered perceptrons. The demands for the periods 1 to m are the inputs of the units in the first layer. The inputs of the units in a higher layer are the outputs of the units in the preceding layer. The outputs of the units in the highest layer determine a vector. In case we used zero-one target vectors (cost target vectors), we choose the first regeneration point equal to the position with the highest (lowest) value.

Units that are not output units are called *hidden units*. All the multi-layered perceptrons that we use have at most one layer with hidden units. Furthermore, the number of inputs and the number of output units is fixed, so the *topology* of our multi-layered perceptrons is completely determined by the number of hidden units.

Each time a unit receives all its inputs it computes a weighted linear combination of all its inputs, adds a bias weight, applies a response function, and outputs the resulting value. So, if the input is presented to the network all the input units do their computation, then the hidden units, and finally the output units.

The network is trained with learning examples such that after each example,

the weights are adjusted by the back-propagation algorithm to reduce the sum-of-squares error function.

Although the capabilities of multi-layered perceptrons to reproduce outputs for given input vectors may be useful in itself, its primary purpose is to generalize, i.e., to produce outputs for input vectors that are not in the learning examples. Multi-layered perceptrons are known for their ability to generalize.

So, on input of the demand vector (d_1, \dots, d_m) the outputs of a multi-layered perceptron obtained by minimizing the sum-of-squares error function on a finite set of learning examples can be interpreted as estimates of the posterior probabilities that l is the optimal first regeneration point given the demands in the demand vector.

3.2 The statistical approach

A well known supervised learning technique from statistical classification is the *K-nearest-neighbors* technique (KNN).

Let K be some integer smaller than the number of learning examples. Then the KNN-based horizon-selection rule is to take the first regeneration point equal to that number l for which the number of class l learning examples contained in the hyper-sphere around the demand vector exactly enclosing K learning examples is maximal.

This gives an estimate of the posterior probabilities that l is the optimal first regeneration point given the demands in the demand vector.

4 Experimental setup

The literature shows several studies that address the performance of variable-horizon policies. These studies address, without exception, the Wagner-Whitin cost structure. Our design of experiments is inspired by the studies of Baker [1977], Blackburn & Millen [1980], Carlson, Beckman & Kropp [1982], and Zwietering, Van Kraaij, Aarts & Wessels [1991]. In our experiments we vary the data horizon length between 2 and 10.

4.1 Demand processes

We investigate the effect of the characteristics of the *demand process* by considering three types of demand processes.

- We use *uniformly distributed demand* with mean $\mu = 100$ and ranges $R = 75, 150, \text{ and } 200$, i.e., the demand d_t is uniformly distributed on $[\mu - \frac{1}{2}R, \mu + \frac{1}{2}R]$. We label these three distributions U75, U150, and U200, respectively.
- We generate *Erlang distributed demand* with (approximately) equal mean and variance as the above mentioned uniform distributions. We label these three distributions E75, E150, and E200, respectively.
- We generate *seasonal demand* using the formula

$$d_t = (\mu - \frac{1}{2}R) \sin\left(\frac{2\pi t}{T}\right) + u_t,$$

where T denotes the cycle length and u_t is a uniformly distributed random variable with mean $\mu = 100$ and range $R = 75$. We use distributions with cycle length $T = 3$ and 6 , which are labeled S3 and S6, respectively.

4.2 Cost structure parameters

We use the three cost structures introduced in Subsection 2.1. For all the cost structures that we use, we take $h = 1$.

- We select two Wagner-Whitin cost structures; one has $S = 200$ and the other $S = 800$. The difference between these two variants can be measured by computing the order cycle. The *order cycle* depends on the cost structure and on the average demand level of the demand distribution (which is 100 in all our cases). The order cycle, n^* , is chosen such that when there is a constant demand rate, equal to the average demand level, and all the subplans have a length equal to the order cycle, the total cost is minimized (Hax & Candea [1984], Silver, Pyke & Peterson [1998]). The order cycle for the two Wagner-Whitin cost structures is 2 and 4 periods, respectively
- Constructing reasonable cost structures with overtime (or purchasing) is less straightforward than for the Wagner-Whitin case. After some experimentation we choose parameters such that a reasonable number of lot sizes requires overtime: $q = p + 1$, $S = 425$, and $C = 200$. Then the order cycle, n^* , is equal to 2.

- For the cost structure with purchasing we choose (after some experimentation) parameters such that a reasonable number of lot sizes requires purchasing: $q = p + 10$, $S = 100$, and $C = 125$. With these parameters the order cycle equals 1.

4.3 Approaches

For all 288 combinations of eight demand processes, four cost structures, and nine data horizon lengths, we generate two sets of 2,500 combined learning examples, i.e., one learning set and one validation set for monitoring the generalization capability during execution of the learning algorithm.

4.3.1 The neural network approach

Before the training of a multi-layered perceptron starts, a preprocessing step is included in which all the elements of the input and target vectors are scaled to a value between 0 and 1. This enables the use of sigmoidal response functions, identical weight initialization procedures, and identical learning parameter values in all instances.

Next, we remove the difference in interpretation of network outputs between zero-one target vectors and cost target vectors by applying the transformation $x \rightarrow 1 - x$ to the elements of the scaled cost target vectors. This is done because we want to maximize the classification rate and to minimize the excess cost.

The initial weights of the multi-layered perceptrons are drawn from a uniform distribution on $[-1, 1]$. The weights are updated by a back-propagation algorithm applying a sequential gradient-descent based weight update scheme with learning rate $\eta = 0.1$ and momentum term $\mu = 0.9$. The convergence is monitored by a sum-of-squares error measure. The learning algorithm applies 2,500,000 iterations for each multi-layered perceptron. During execution the sum-of-squares error is monitored on the validation set and the network with the lowest error is kept.

Because of the size of our experimental setup, manual tuning of the network topology is infeasible. Therefore, we use a predetermined set of network topologies. For the two Wagner-Whitin cost structures we investigate network topologies with $0, m, m + 2, m + 4$, and $m + 6$ hidden units. For the cost structures with overtime and with purchasing we use $0, m, m + 3, m + 6$, and $m + 9$ hidden units. All networks have one layer with hidden units, except for the networks without hidden units. For each of the five different network topologies we develop ten networks. Learning examples are randomly presented to each network, so after training we

have ten different approximations of the target mapping. These ten networks are combined to form a committee by averaging over network outputs. For each committee we compute the sum-of-squares error on the validation set. The horizon-selection rules are obtained by selecting the committee with the lowest sum-of-squares error.

4.3.2 The statistical approach

A suitable K -value for the K -nearest-neighbors policy is determined by computing the classification rate (or average excess cost) on the validation set for $K = 1, \dots, 20$, and choosing the one with the highest classification rate (or lowest average excess cost).

4.3.3 Classical approaches

We compare our approach with some other approaches that have been proposed in the literature. It appears that most of these approaches can be transformed into variable-horizon policies without loss of performance. Recall from Section 3 that a variable-horizon policy is completely determined by its horizon-selection rule, i.e., by its procedure to select an optimization horizon t .

We consider two types of approaches, i.e., myopic approaches and approaches that use simple demand forecasting. This has led to the following four myopic policies.

- The *fixed-horizon policy* (FIX) chooses t equal to m , that is, it optimizes over the complete data horizon.
- The *least cost per unit time policy* (PUT) is based on Gorham [1968]. In this approach t is chosen in such a way that the cost per unit time, $f(t)/t$, is minimized.
- The *least cost per unit product policy* (PUP) is based on Silver & Meal [1973]. Here t is chosen such that the cost per unit product, $f(t)/\sum_{k=1}^t d_k$, is minimized.
- In the *economic order quantity policy* (EOQ) t is always equal to the order cycle n^* (Harris [1913]).

The idea behind policies with forecasting is to forecast demand for as many periods in the future as necessary for a forward algorithm to be able to detect a simple

planning horizon. We use the following two approaches based on Carlson, Beckman & Kropp [1982].

- Take the overall *average demand* as a forecast for the unknown future demands (AVG).
- Use a 5-period *moving average* to forecast the unknown future demands (MVG).

In both policies the variability in demand decreases in time, so if after 100 iterations still no planning horizon is found, the smallest regeneration point that has been found by the forward algorithm in the last 100 iterations is selected.

5 Empirical results

Most tables in this section present performance numbers in percentages deviation from the optimum. So if a value g is measured, then the value $(g - 1) \times 100$ is used.

A background gray color is added to each table entry. The level of darkness scales with the value of entry as defined by the mapping

$$s(x) = \begin{cases} \frac{x}{5} & \text{if } 0 \leq x \leq 5 \\ 1 & \text{if } x > 5 \end{cases},$$

where 0 and 1 denote white and dark grey, respectively.

The performance of variable-horizon policies is usually evaluated empirically by applying them to different instances of on-line problems with a finite number of periods. A typical set-up is to consider 50 replications of 48-period problems (Baker [1977] and Carlson, Beckman & Kropp [1982]). Such experiments are strongly effected by so-called end-of-horizon effects, caused by the truncation of the infinite horizon to a finite horizon (Blackburn & Millen [1980]). To minimize these effects in our experiments we use one instance of 4000 periods.

Below we present numerical results on the potential, and the off-line and on-line performance of the policies under consideration.

5.1 Potential

An n -period plan is called m -optimal if it consists of subplans with lengths not exceeding m and it has the lowest cost possible. The *potential*, $\alpha_m(n)$, of variable-

Cost structure	m	Demand process							
		Uniform			Erlang			Seasonal	
		U75	U150	U200	E75	E150	E200	S3	S6
<div style="border: 1px solid black; padding: 2px; width: fit-content;"> Wagner-Whitin $h = 1$ $S = 200$ $n^* = 2$ </div>	2	0.66	3.64	7.07	0.61	2.84	5.26	3.89	5.56
	3	0.00	0.22	0.91	0.00	0.08	0.34	0.00	0.71
	4	0.00	0.01	0.09	0.00	0.00	0.02	0.00	0.00
	5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	7	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	<div style="border: 1px solid black; padding: 2px; width: fit-content;"> Wagner-Whitin $h = 1$ $S = 800$ $n^* = 4$ </div>	2	31.58	37.90	42.73	31.64	36.67	41.82	36.24
3		6.75	10.90	14.00	6.83	10.32	13.02	5.40	14.33
4		0.81	2.69	4.22	0.65	2.15	3.76	3.17	5.47
5		0.01	0.52	1.03	0.00	0.35	0.81	0.06	1.40
6		0.00	0.07	0.25	0.00	0.03	0.15	0.00	0.03
7		0.00	0.01	0.08	0.00	0.01	0.02	0.00	0.00
8		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
9		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
10		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
<div style="border: 1px solid black; padding: 2px; width: fit-content;"> Overtime $h = 1$ $q = p + 1$ $S = 425$ $C = 200$ $n^* = 2$ </div>		2	2.65	7.62	10.71	2.63	6.64	9.99	7.98
	3	0.00	0.70	1.65	0.01	0.41	0.99	0.07	2.00
	4	0.00	0.03	0.19	0.00	0.02	0.12	0.00	0.00
	5	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	7	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	<div style="border: 1px solid black; padding: 2px; width: fit-content;"> Purchase $h = 1$ $q = p + 10$ $S = 100$ $C = 125$ $n^* = 1$ </div>	2	1.17	10.99	18.72	1.86	13.81	16.54	17.39
3		0.29	5.03	9.21	0.51	6.52	8.17	0.95	12.44
4		0.09	2.11	4.92	0.11	3.18	4.25	0.30	2.32
5		0.03	0.93	2.42	0.04	1.51	2.28	0.21	0.31
6		0.01	0.46	1.04	0.00	0.89	1.16	0.03	0.06
7		0.00	0.21	0.61	0.00	0.43	0.62	0.01	0.01
8		0.00	0.08	0.26	0.00	0.16	0.24	0.01	0.00
9		0.00	0.01	0.05	0.00	0.04	0.06	0.00	0.00
10		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 1: Estimates of the potential of variable-horizon policies in percentages deviation for different combinations of cost structures, data horizon lengths m , and demand processes.

horizon policies is defined by

$$\alpha_m(n) = \frac{f_m(n)}{f(n)},$$

where $f(n)$ denotes the cost of an optimal plan and $f_m(n)$ denotes the cost of an m -optimal plan. Based on the stationarity of the demand processes under consideration we expect that $\alpha_m(n)$ converges to a limit value as n goes to infinity. This limit can be seen as a measure for the potential of variable-horizon policies with data horizon length m , so no variable horizon policy can do better. The potential for the various cost structures and demand processes is given in Table 1.

5.2 Off-line performance

The *off-line performance* measure of variable-horizon policy π is the deviation from off-line optimality defined by

$$\gamma_m^\pi(n) = \frac{C_m^\pi(n)}{f(n)},$$

where $C_m^\pi(n)$ denotes the cost of the plan obtained by policy π with data horizon length m and n denotes the number of periods for the on-line problem. Based on the stationarity of the demand processes under consideration we expect that $\gamma_m^\pi(n)$ converges to a limit value as n goes to infinity. The off-line performance is presented in Table 2 and Table 3.

5.3 On-line performance

The *on-line performance* measure of variable-horizon policy π is the deviation from the best performance of any variable-horizon policy defined by

$$\phi_m^\pi(n) = \frac{C_m^\pi(n)}{\min\{C_m^{\pi(g)}(n) | g \in H_m\}},$$

where H_m denotes the set of all possible horizon-selection rules that can be written as a function $g : \mathbb{R}^m \rightarrow \{1, \dots, m\}$ and $\pi(g)$ is the variable-horizon policy belonging to horizon-selection rule g .

Because the best performance of any variable-horizon policy is in general unknown, we can only derive lower bounds and upper bound according to

$$\frac{C_m^\pi(n)}{\min\{C_m^{\pi_j}(n) | j = 1, \dots, N\}} = \underline{\phi}_m^\pi(n) \leq \phi_m^\pi(n) \leq \bar{\phi}_m^\pi(n) = \frac{C_m^\pi(n)}{f_m(n)},$$

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	6.81	6.81	7.13	6.90	6.81	6.73	5.88	5.88	5.92	5.72	Average
3	4.08	2.22	6.53	5.75	2.86	2.23	1.55	1.53	1.54	1.49	
4	1.49	1.15	5.20	6.64	1.49	1.02	0.55	0.55	0.59	0.52	
5	0.89	0.89	4.29	6.46	0.86	0.55	0.28	0.41	0.32	0.31	
6	0.37	0.78	4.05	6.75	0.37	0.28	0.19	0.37	0.18	0.18	
7	0.23	0.76	4.18	7.34	0.19	0.15	0.14	0.48	0.10	0.10	
8	0.09	0.72	4.03	7.44	0.09	0.08	0.13	0.54	0.05	0.06	
9	0.03	0.75	4.42	8.27	0.04	0.02	0.13	0.68	0.03	0.04	
10	0.01	0.77	4.51	8.58	0.01	0.01	0.14	0.79	0.02	0.03	
2	12.97	12.97	12.97	13.73	12.97	12.91	11.26	11.77	11.58	11.48	
3	5.75	3.69	11.12	10.57	4.21	3.18	2.73	2.97	2.65	3.01	
4	2.01	2.35	8.33	12.96	2.03	1.55	0.71	0.77	0.73	0.78	
5	1.76	2.00	7.11	11.94	1.43	1.02	0.41	0.52	0.41	0.45	
6	0.86	1.74	6.47	12.51	0.84	0.65	0.27	0.60	0.27	0.29	
7	0.83	1.73	6.02	12.78	0.50	0.40	0.23	0.88	0.18	0.21	
8	0.27	1.69	5.82	13.06	0.27	0.26	0.25	1.00	0.13	0.15	
9	0.07	1.73	5.67	12.77	0.11	0.07	0.26	1.16	0.09	0.09	
10	0.04	1.73	5.62	13.17	0.04	0.03	0.27	1.46	0.06	0.07	

(a) Wagner-Whitin cost structure with order cycle 2

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	37.77	37.77	37.77	37.77	37.78	37.78	37.75	37.71	37.77	37.69	Average
3	11.46	11.31	11.48	11.46	16.85	17.50	11.32	11.40	11.26	11.21	
4	6.80	4.90	7.29	6.39	5.68	5.19	4.25	4.22	4.22	4.17	
5	7.31	2.67	7.81	6.29	2.38	2.38	1.69	1.74	1.70	1.70	
6	4.89	1.48	7.15	6.68	1.33	1.33	0.88	0.93	0.94	0.88	
7	2.40	1.13	6.08	6.98	0.86	0.91	0.54	0.75	0.65	0.67	
8	1.28	0.92	5.30	6.94	0.60	0.57	0.43	0.68	0.48	0.55	
9	1.21	0.79	5.03	6.79	0.49	0.39	0.35	0.66	0.38	0.41	
10	0.96	0.72	4.81	6.70	0.34	0.31	0.32	0.63	0.30	0.38	
2	43.08	43.08	43.08	43.08	43.14	43.14	43.08	43.08	43.08	43.08	
3	16.19	16.19	16.19	16.19	22.34	24.35	16.80	16.79	16.53	16.50	
4	11.04	8.64	12.12	10.27	8.57	8.07	6.51	6.74	6.70	6.64	
5	9.63	6.93	13.36	9.53	5.17	4.82	2.73	3.03	2.73	3.00	
6	5.66	2.72	13.31	10.31	1.99	2.01	1.35	1.50	1.38	1.38	
7	3.41	1.94	12.58	11.29	1.11	1.43	0.73	1.13	0.88	1.03	
8	1.85	1.88	11.14	10.84	0.70	0.62	0.55	1.05	0.61	0.65	
9	1.65	1.83	10.72	11.26	0.58	0.44	0.45	1.09	0.53	0.51	
10	1.31	1.73	10.58	10.89	0.45	0.42	0.53	1.06	0.36	0.46	

(b) Wagner-Whitin cost structure with order cycle 4

Table 2: *Percentage deviation from off-line optimality for ten different variable-horizon policies, using the Wagner-Whitin cost structure with two different order cycles and different data horizons length m . The entries represent average and worst-case percentages of eight different demand processes.*

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	8.38	8.42	8.41	8.37	8.37	8.41	8.34	8.33	8.38	8.29	Average
3	5.39	2.84	6.07	4.82	2.26	2.62	1.94	2.04	1.99	2.00	
4	3.05	1.24	5.42	6.23	1.11	1.20	0.79	0.85	0.83	0.84	
5	1.67	0.88	4.14	6.86	0.73	0.78	0.51	0.63	0.51	0.53	
6	0.93	0.69	3.78	7.30	0.40	0.55	0.35	0.50	0.34	0.34	
7	0.62	0.58	3.57	7.49	0.31	0.31	0.28	0.47	0.22	0.24	
8	0.39	0.53	3.41	7.64	0.19	0.25	0.27	0.52	0.16	0.19	
9	0.21	0.48	3.69	8.04	0.14	0.16	0.25	0.54	0.13	0.16	
10	0.14	0.49	4.22	9.59	0.07	0.08	0.26	0.65	0.11	0.11	
2	12.52	12.52	12.52	12.52	12.45	12.45	12.75	12.76	12.85	12.72	
3	7.99	6.42	11.30	7.41	3.81	4.50	3.49	3.48	3.56	3.62	
4	4.31	1.89	8.81	9.83	1.70	1.67	1.39	1.47	1.46	1.43	
5	2.22	1.59	6.23	9.79	0.90	0.96	0.78	0.95	0.70	0.80	
6	1.08	1.23	6.10	10.89	0.56	0.74	0.47	0.84	0.43	0.48	
7	1.00	0.97	5.87	10.99	0.47	0.40	0.39	0.79	0.28	0.30	
8	0.50	0.96	5.52	10.93	0.30	0.34	0.47	0.99	0.24	0.26	
9	0.32	0.80	5.40	11.07	0.19	0.20	0.47	1.11	0.16	0.17	
10	0.18	0.82	5.31	11.17	0.09	0.09	0.43	1.25	0.13	0.14	

(a) Cost structure with overtime

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	16.34	29.33	20.23	29.04	16.26	16.22	14.54	14.62	14.10	14.01	Average
3	8.26	28.52	12.22	28.60	8.22	8.07	6.38	6.63	6.58	6.17	
4	4.90	28.06	10.77	29.01	4.90	4.80	3.11	3.50	3.13	2.88	
5	2.49	27.81	10.11	29.31	2.49	2.42	1.82	2.48	2.17	1.46	
6	1.17	27.40	9.57	29.45	1.16	1.16	1.26	2.53	1.36	0.85	
7	0.69	27.17	9.29	29.44	0.71	0.68	1.17	2.94	1.10	0.57	
8	0.43	26.98	9.09	29.57	0.44	0.43	1.15	3.39	0.82	0.43	
9	0.18	26.81	8.82	29.68	0.19	0.19	1.16	3.94	0.71	0.31	
10	0.08	26.72	8.63	29.78	0.09	0.09	1.24	4.40	0.59	0.29	
2	26.11	44.82	33.60	44.19	26.15	26.19	24.48	24.45	24.64	24.27	
3	17.71	43.47	24.42	41.83	17.63	17.59	13.64	13.94	13.39	13.31	
4	8.93	42.53	20.29	43.38	8.94	8.93	6.11	6.79	6.47	5.99	
5	5.13	42.29	18.03	43.37	5.16	5.07	3.80	4.96	4.25	3.39	
6	3.12	41.66	17.08	43.18	3.23	3.21	2.49	4.66	2.62	1.99	
7	1.58	40.99	16.79	43.43	1.62	1.61	2.09	5.39	1.79	1.47	
8	0.98	40.72	16.64	44.27	1.03	1.05	2.40	6.20	1.49	1.01	
9	0.52	40.32	15.94	44.51	0.53	0.51	2.66	7.50	1.54	0.70	
10	0.33	40.19	14.98	44.28	0.34	0.33	3.03	8.45	1.32	0.77	

(b) Cost structure with purchasing

Table 3: *Percentage deviation from off-line optimality for ten different variable-horizon policies, using the cost structure with overtime and the one with purchasing and different data horizons length m . The entries represent average and worst-case percentages of eight different demand processes.*

where we have N variable-horizon policies $\pi_i, i = 1, \dots, N$. We expect that the lower bound, $\underline{\phi}_m^\pi(n)$, converges when taking the limit of n to infinity.

Preluding on the experiments, we only observed significant gaps between the upper and lower bound for small values of the data horizon length. Fortunately, it is especially in those cases that the our multi-layered perceptrons generalize excellent and near on-line performance of our new policies can be expected. For these reasons, the lower bounds are used as estimates for the on-line performance. For the sake of completeness, we also give the upper bounds. Numerical values for the bounds are presented in Table 5, Table 6, Table 7, and Table 8.

5.4 Switch-over points

From the literature we know that the off-line performance of variable-horizon policies like FIX, AVG, and MVG can be arbitrarily close to off-line optimality by taking the data horizon length sufficiently large (e.g. Carlson, Beckman & Kropp [1982]). This does not hold for the variable-horizon policies based on supervised learning. In general the data horizon length cannot be enlarged unlimited, because when the dimensionality of the input space is increased by adding new features, it rapidly leads to the point where the learning examples provide a very poor representation of the target mapping.

For these reasons there exists a data horizon length for which FIX, AVG, and MVG begin to exhibit a consistently superior average and worst-case off-line performance compared to that of the MLP-based and the KNN-based variable-horizon policies. This data horizon length is called the *switch-over point* and is calculated for the four cost structures we considered. The results are presented as a multiple of the order cycle in Table 4.

Cost structure	n^*	switch-over point
Wagner-Whitin	2	$4.5n^*$
Wagner-Whitin	4	$3n^*$
Overtime	2	$5n^*$
Purchase	1	$8n^*$

Table 4: *Switch-over points for four different cost structures.*

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	1.06	1.06	1.36	1.14	1.06	0.99	0.19	0.19	0.23	0.04	Average
3	2.63	0.80	5.03	4.27	1.44	0.81	0.13	0.11	0.13	0.08	
4	1.01	0.67	4.70	6.14	1.01	0.54	0.07	0.07	0.11	0.04	
5	0.62	0.63	4.02	6.18	0.59	0.29	0.02	0.14	0.05	0.04	
6	0.21	0.62	3.89	6.58	0.21	0.12	0.03	0.21	0.02	0.02	
7	0.15	0.68	4.09	7.26	0.12	0.07	0.06	0.40	0.02	0.03	
8	0.05	0.68	3.98	7.39	0.05	0.03	0.08	0.49	0.01	0.02	
9	0.01	0.73	4.40	8.25	0.02	0.00	0.11	0.66	0.01	0.02	
10	0.00	0.76	4.50	8.57	0.00	0.00	0.13	0.78	0.01	0.02	
2	1.85	1.85	3.83	2.22	1.85	1.56	0.62	0.46	0.69	0.20	
3	4.68	2.35	8.25	7.72	2.15	1.90	0.42	0.33	0.35	0.35	
4	1.41	2.17	7.60	12.20	1.42	0.95	0.20	0.17	0.21	0.10	
5	1.46	1.87	6.82	11.64	1.10	0.63	0.11	0.28	0.11	0.09	
6	0.62	1.67	6.34	12.38	0.60	0.41	0.11	0.48	0.09	0.05	
7	0.70	1.70	5.97	12.72	0.37	0.27	0.18	0.83	0.07	0.08	
8	0.15	1.68	5.80	13.04	0.15	0.14	0.23	0.98	0.02	0.04	
9	0.03	1.73	5.66	12.75	0.05	0.01	0.25	1.15	0.02	0.04	
10	0.01	1.73	5.62	13.17	0.01	0.00	0.27	1.46	0.04	0.04	

(a) Lower bounds

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	2.98	2.98	3.29	3.06	2.98	2.91	2.10	2.09	2.13	1.94	Average
3	3.79	1.93	6.22	5.45	2.58	1.94	1.26	1.24	1.26	1.21	
4	1.47	1.13	5.19	6.63	1.47	1.01	0.53	0.54	0.58	0.51	
5	0.89	0.89	4.29	6.46	0.86	0.55	0.28	0.41	0.32	0.31	
6	0.37	0.78	4.05	6.75	0.37	0.28	0.19	0.37	0.18	0.18	
7	0.23	0.76	4.18	7.34	0.19	0.15	0.14	0.48	0.10	0.10	
8	0.09	0.72	4.03	7.44	0.09	0.08	0.13	0.54	0.05	0.06	
9	0.03	0.75	4.42	8.27	0.04	0.02	0.13	0.68	0.03	0.04	
10	0.01	0.77	4.51	8.58	0.01	0.01	0.14	0.79	0.02	0.03	
2	5.51	5.51	5.84	6.21	5.50	5.45	3.90	4.39	4.21	4.11	
3	5.37	2.96	10.11	9.57	3.38	2.46	1.80	2.04	1.72	2.07	
4	2.01	2.35	8.24	12.86	2.03	1.55	0.68	0.74	0.71	0.69	
5	1.76	2.00	7.11	11.94	1.43	1.02	0.41	0.52	0.41	0.45	
6	0.86	1.74	6.47	12.51	0.84	0.65	0.27	0.60	0.27	0.29	
7	0.83	1.73	6.02	12.78	0.50	0.40	0.23	0.88	0.18	0.21	
8	0.27	1.69	5.82	13.06	0.27	0.26	0.25	1.00	0.13	0.15	
9	0.07	1.73	5.67	12.77	0.11	0.07	0.26	1.16	0.09	0.09	
10	0.04	1.73	5.62	13.17	0.04	0.03	0.27	1.46	0.06	0.07	

(b) Upper bounds

Table 5: Bounds on the percentage deviation from on-line optimality for ten different variable-horizon policies, using the Wagner-Whitin cost structure with order cycle 2 and different data horizons length m . The entries represent average and worst-case percentages of eight different demand processes.

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	0.06	0.06	0.06	0.06	0.07	0.07	0.04	0.01	0.06	0.00	Average
3	0.27	0.13	0.28	0.27	5.13	5.70	0.14	0.22	0.09	0.06	
4	2.58	0.77	3.04	2.18	1.51	1.04	0.15	0.11	0.12	0.07	
5	5.60	1.03	6.08	4.59	0.75	0.75	0.07	0.12	0.08	0.08	
6	4.02	0.64	6.26	5.80	0.50	0.49	0.04	0.10	0.11	0.04	
7	1.86	0.60	5.52	6.41	0.32	0.38	0.01	0.21	0.12	0.14	
8	0.87	0.50	4.87	6.49	0.19	0.16	0.01	0.26	0.06	0.13	
9	0.88	0.46	4.68	6.44	0.16	0.05	0.02	0.33	0.05	0.08	
10	0.69	0.46	4.53	6.42	0.07	0.04	0.05	0.37	0.03	0.11	
2	0.27	0.27	0.27	0.27	0.26	0.26	0.27	0.06	0.27	0.01	
3	1.28	0.58	1.38	1.28	7.68	8.19	0.53	0.71	0.29	0.28	
4	4.36	2.00	5.37	3.63	2.14	1.46	0.61	0.31	0.35	0.19	
5	6.84	4.32	10.60	6.86	2.61	2.27	0.26	0.29	0.28	0.27	
6	4.81	2.13	12.66	9.01	1.40	1.43	0.15	0.45	0.34	0.16	
7	3.09	1.63	12.24	10.48	0.81	1.12	0.08	0.39	0.33	0.30	
8	1.57	1.60	10.83	10.31	0.27	0.25	0.07	0.56	0.22	0.26	
9	1.38	1.57	10.43	10.92	0.26	0.11	0.14	0.78	0.20	0.18	
10	1.05	1.48	10.31	10.60	0.13	0.11	0.26	0.80	0.12	0.18	

(a) Lower bounds

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	0.23	0.23	0.23	0.23	0.23	0.23	0.21	0.18	0.23	0.17	Average
3	1.12	0.98	1.13	1.12	6.02	6.60	0.99	1.07	0.94	0.91	
4	3.81	1.97	4.28	3.41	2.73	2.25	1.35	1.31	1.32	1.27	
5	6.75	2.13	7.24	5.73	1.84	1.85	1.16	1.20	1.17	1.17	
6	4.81	1.41	7.07	6.61	1.26	1.26	0.81	0.86	0.87	0.81	
7	2.38	1.12	6.07	6.96	0.84	0.90	0.53	0.73	0.64	0.65	
8	1.28	0.92	5.30	6.94	0.60	0.57	0.43	0.68	0.48	0.55	
9	1.21	0.79	5.03	6.79	0.49	0.39	0.35	0.66	0.38	0.41	
10	0.96	0.72	4.81	6.70	0.34	0.31	0.32	0.63	0.30	0.38	
2	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.33	0.55	0.33	
3	2.44	1.93	2.54	2.44	8.14	8.75	2.47	2.47	2.23	2.22	
4	6.55	3.00	7.58	5.81	4.14	3.33	2.10	2.42	2.10	2.24	
5	8.52	5.45	11.79	8.20	3.71	3.37	1.68	1.98	1.68	1.95	
6	5.50	2.69	13.27	10.00	1.95	1.98	1.08	1.27	1.23	1.11	
7	3.41	1.94	12.58	11.20	1.11	1.43	0.65	1.04	0.85	0.94	
8	1.85	1.88	11.14	10.84	0.70	0.62	0.55	1.05	0.61	0.65	
9	1.65	1.83	10.72	11.26	0.58	0.44	0.45	1.09	0.53	0.51	
10	1.31	1.73	10.58	10.89	0.45	0.42	0.53	1.06	0.36	0.46	

(b) Upper bounds

Table 6: Bounds on the percentage deviation from on-line optimality for ten different variable-horizon policies, using the Wagner-Whitin cost structure with order cycle 4 and different data horizons length m . The entries represent average and worst-case percentages of eight different demand processes.

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	0.12	0.15	0.15	0.11	0.12	0.15	0.09	0.07	0.12	0.03	Average
3	3.41	0.90	4.07	2.84	0.34	0.69	0.01	0.12	0.06	0.08	
4	2.25	0.46	4.60	5.40	0.33	0.42	0.01	0.07	0.04	0.06	
5	1.19	0.41	3.65	6.36	0.26	0.31	0.03	0.16	0.04	0.06	
6	0.62	0.38	3.46	6.97	0.09	0.25	0.05	0.19	0.03	0.03	
7	0.41	0.37	3.35	7.27	0.10	0.11	0.08	0.26	0.01	0.04	
8	0.25	0.38	3.26	7.48	0.05	0.10	0.12	0.38	0.02	0.05	
9	0.10	0.37	3.57	7.91	0.02	0.05	0.14	0.43	0.02	0.05	
10	0.07	0.42	4.14	9.51	0.00	0.01	0.19	0.58	0.04	0.04	
2	0.28	0.43	0.31	0.28	0.31	0.38	0.24	0.21	0.33	0.18	
3	4.88	3.36	10.26	6.41	0.81	1.01	0.05	0.21	0.23	0.14	
4	3.51	0.92	7.97	8.98	0.54	0.62	0.03	0.20	0.18	0.13	
5	1.90	1.01	5.66	9.21	0.42	0.36	0.11	0.31	0.13	0.13	
6	0.88	0.88	5.73	10.58	0.23	0.42	0.18	0.43	0.08	0.11	
7	0.77	0.75	5.64	10.74	0.25	0.20	0.17	0.58	0.05	0.07	
8	0.39	0.74	5.29	10.78	0.13	0.19	0.33	0.83	0.03	0.10	
9	0.22	0.66	5.26	10.97	0.06	0.07	0.38	1.02	0.05	0.07	
10	0.12	0.75	5.23	11.12	0.00	0.02	0.38	1.20	0.06	0.09	

(a) Lower bounds

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	0.95	0.99	0.98	0.94	0.95	0.98	0.92	0.90	0.95	0.87	Average
3	4.63	2.09	5.31	4.06	1.53	1.88	1.20	1.30	1.25	1.26	
4	3.01	1.20	5.37	6.18	1.07	1.16	0.74	0.81	0.78	0.79	
5	1.67	0.88	4.14	6.86	0.73	0.78	0.50	0.63	0.51	0.53	
6	0.93	0.69	3.78	7.30	0.40	0.55	0.35	0.50	0.34	0.34	
7	0.62	0.58	3.57	7.49	0.31	0.31	0.28	0.47	0.22	0.24	
8	0.39	0.53	3.41	7.64	0.19	0.25	0.27	0.52	0.16	0.19	
9	0.21	0.48	3.69	8.04	0.14	0.16	0.25	0.54	0.13	0.16	
10	0.14	0.49	4.22	9.59	0.07	0.08	0.26	0.65	0.11	0.11	
2	1.82	2.13	2.01	1.77	1.83	2.07	1.85	1.82	1.94	1.79	
3	5.86	4.33	11.22	7.33	2.11	2.96	1.93	2.09	2.16	2.00	
4	4.31	1.77	8.81	9.83	1.58	1.55	1.27	1.35	1.34	1.31	
5	2.22	1.59	6.23	9.78	0.90	0.95	0.78	0.94	0.70	0.79	
6	1.08	1.23	6.10	10.89	0.56	0.74	0.47	0.84	0.43	0.48	
7	1.00	0.97	5.87	10.99	0.47	0.40	0.39	0.79	0.28	0.30	
8	0.50	0.96	5.52	10.93	0.30	0.34	0.47	0.99	0.24	0.26	
9	0.32	0.80	5.40	11.07	0.19	0.20	0.47	1.11	0.16	0.17	
10	0.18	0.82	5.31	11.17	0.09	0.09	0.43	1.25	0.13	0.14	

(b) Upper bounds

Table 7: Bounds on the percentage deviation from on-line optimality for ten different variable-horizon policies, using the cost structure with overtime and different data horizons length m . The entries represent average and worst-case percentages of eight different demand processes.

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	1.97	13.15	5.28	12.90	1.99	1.96	0.45	0.52	0.08	0.01	Average
3	1.92	20.90	5.56	20.99	1.89	1.75	0.21	0.44	0.39	0.01	
4	1.96	24.32	7.59	25.24	1.97	1.86	0.23	0.60	0.24	0.01	
5	1.02	25.89	8.49	27.36	1.02	0.95	0.36	1.01	0.71	0.01	
6	0.32	26.29	8.63	28.33	0.32	0.31	0.42	1.67	0.52	0.01	
7	0.18	26.49	8.72	28.75	0.19	0.17	0.66	2.41	0.59	0.06	
8	0.10	26.54	8.72	29.12	0.11	0.11	0.82	3.05	0.50	0.10	
9	0.04	26.62	8.66	29.48	0.05	0.05	1.02	3.79	0.57	0.17	
10	0.01	26.62	8.54	29.67	0.01	0.01	1.17	4.32	0.51	0.22	
2	7.32	20.70	13.70	19.90	7.32	7.35	2.37	1.63	0.30	0.03	
3	3.89	37.76	9.83	39.01	3.84	3.81	0.60	0.98	1.29	0.08	
4	5.75	37.37	16.79	39.06	5.83	5.82	0.76	0.99	0.45	0.04	
5	3.16	38.30	17.16	40.54	3.20	3.22	0.53	1.75	2.25	0.05	
6	1.11	38.89	16.78	41.46	1.20	1.19	0.57	2.61	1.05	0.07	
7	0.68	39.27	16.55	41.68	0.70	0.69	1.21	4.10	1.41	0.32	
8	0.51	39.38	16.46	42.89	0.52	0.48	1.43	5.18	1.33	0.34	
9	0.34	40.02	15.73	44.20	0.34	0.29	2.44	7.27	1.36	0.48	
10	0.07	40.09	14.90	44.18	0.08	0.08	2.95	8.37	1.20	0.70	

(a) Lower bounds

m	Myopic				Forecasting		0-1 targets		Cost targets		
	FIX	PUT	PUP	EOQ	AVG	MVG	MLP	KNN	MLP	KNN	
2	2.82	14.11	6.16	13.86	2.83	2.80	1.29	1.35	0.91	0.84	Average
3	2.65	21.79	6.33	21.88	2.62	2.48	0.92	1.16	1.11	0.72	
4	2.67	25.23	8.35	26.16	2.68	2.58	0.93	1.30	0.94	0.70	
5	1.51	26.52	9.02	28.01	1.51	1.44	0.84	1.50	1.19	0.49	
6	0.70	26.79	9.05	28.84	0.70	0.69	0.80	2.06	0.90	0.39	
7	0.45	26.85	9.02	29.12	0.46	0.44	0.93	2.69	0.86	0.33	
8	0.33	26.85	8.99	29.44	0.34	0.34	1.05	3.30	0.73	0.33	
9	0.16	26.78	8.79	29.65	0.17	0.16	1.14	3.91	0.69	0.29	
10	0.08	26.72	8.63	29.78	0.09	0.09	1.24	4.40	0.59	0.29	
2	7.42	21.98	13.80	21.45	7.42	7.45	4.40	3.33	1.98	2.00	
3	4.93	38.41	11.07	39.65	4.95	4.61	1.76	2.23	2.32	1.83	
4	6.45	38.82	17.56	40.53	6.53	6.52	1.83	2.14	1.55	1.40	
5	3.64	38.95	17.66	41.19	3.68	3.70	1.35	2.62	2.68	0.95	
6	2.07	40.20	17.01	41.71	2.16	2.15	1.44	3.58	1.57	0.95	
7	0.96	40.14	16.77	42.56	0.98	0.98	1.48	4.75	1.60	0.84	
8	0.72	40.36	16.64	43.90	0.76	0.78	2.14	5.92	1.49	0.77	
9	0.52	40.26	15.94	44.44	0.53	0.47	2.61	7.45	1.54	0.65	
10	0.33	40.19	14.98	44.28	0.34	0.33	3.03	8.45	1.32	0.77	

(b) Upper bounds

Table 8: Bounds on the percentage deviation from on-line optimality for ten different variable-horizon policies, using the cost structure with purchase and different data horizons length m . The entries represent average and worst-case percentages of eight different demand processes.

5.5 Data efficiency

To compare the different cost structures, we determine the minimal data horizon length such that its worst-case off-line performance is consistently smaller than 1% for all demand processes. Such a length is called the *data efficiency* of a cost structure and it is measured as a multiple of the order cycle. The values are given in Table 9. The entries of the variable-horizon policies that do not yield such a performance are omitted. Note that for PUP and EOQ no such performance is observed for any cost structure. Furthermore we include the minimal data horizon length for which the potential deviates less than 1% from the worst case values. These lengths are denoted m -OPT and can be viewed as an indication of the value of demand information for the corresponding on-line lot-sizing problem.

Cost structure	n^*	m -OPT	Myopic		Forecasting		0-1 targets		Cost targets		
			FIX	PUT	AVG	MVG	MLP	KNN	MLP	KNN	
Wagner-Whitin	2	$1.5n^*$	$3n^*$		$3n^*$	$3n^*$		$2n^*$	$2n^*$	$2n^*$	$2n^*$
Wagner-Whitin	4	$1.5n^*$	$3n^*$		$2n^*$	$2n^*$		$1.75n^*$		$1.75n^*$	$2n^*$
Overtime	2	$2n^*$	$4n^*$	$4n^*$	$2.5n^*$	$2.5n^*$		$2.5n^*$	$2.5n^*$	$2.5n^*$	$2.5n^*$
Purchase	1	$7n^*$	$8n^*$		$8n^*$	$8n^*$					$8n^*$

Table 9: *Data efficiency for four different cost structures and eight different variable-horizon policies.*

6 Discussion

The aim of our investigations was to construct policies for on-line lot-sizing which are more robust than the classical approaches. A variable-horizon policy is said to be *robust* if, irrespective of the cost structure, the demand process, and the data horizon length, it yields a good on-line performance. MLP (0-1 targets), MLP (cost targets), and KNN (cost targets) turn out to be robust policies for the two Wagner-Whitin cost structures and for the one with overtime. Blackburn & Millen [1980] recommended the use of PUT for situations with a small data horizon length. Although the average deviation of PUT is small, it is not a robust policy, because its worst-case deviation is rather large.

For the cost structure with purchasing it can be seen that MLP (cost targets) is outperformed by KNN (cost targets). Apparently, the network topologies we used had too few hidden units. This is supported by the observation that the performance of KNN (0-1 targets) gets worse for large data horizons, which indicates that the

target mapping is complex. The off-line performance characteristics of KNN (cost targets) are an indication of the performance that can be obtained by adding more hidden units. It appears that for the cost structure with purchasing KNN (cost targets) is a robust policy.

For the two Wagner-Whitin cost structures and for the cost structure with overtime it can be observed that KNN (0-1 targets) deteriorates in off- and on-line performance with increasing data horizon length. The same holds to a lesser extent for MLP (0-1 targets). See Subsection 5.4 for an explanation.

Considering data efficiency, it is clear that MLP (0-1 targets), MLP (cost targets), and KNN (cost targets) are rather data efficient. That is, compared to the optimal case (m -OPT) they need only a small extra amount of demand information.

6.1 Conclusions

We conclude that when the data horizons length is smaller than the switch-over point, the MLP-based variable-horizon policies have superior performance characteristics. For data horizon lengths beyond this point, these policies are outperformed by FIX, AVG, and MVG. It appeared that the switch-over point is mainly determined by the cost structure. Furthermore we can conclude that it is larger than the data horizon length needed by FIX to obtain a worst-case off-line performance smaller than 1%.

Furthermore we conclude that, irrespective of cost structure and demand process, the MLP-based variable-horizon policies yield excellent performance which is obtained using only little demand information.

In case the data horizon length is small, demand uncertainty is large and it is important to determine a good optimization horizon. In that case MLP-based variable-horizon policies dominate all other variable-horizon policies through good performance characteristics, great data efficiency, and robustness.

In case the data horizon length is large and the demand uncertainty is small, there is hardly any benefit from determining a suitable optimization horizon by using either forecasting, multi-layered perceptrons or the K -nearest-neighbors technique. The best policy that can be used in this case is the fixed-horizon policy FIX. This is not surprising, because when many demands are known in advance, it is much more likely that a simple m -planning horizon can be found within the data horizon.

6.2 Prior knowledge

In this paper we investigated the performance of two variants of our decomposition approach: a statistical and a learning approach. Our empirical results indicate that MLP-based variable-horizon policies dominate all other variable-horizon policies (including KNN). A problem with these policies is the determination of a good network topology, although more sophisticated learning algorithms exist which add weights and/or units or remove weights and/or units from the network (see Reed [1993]). An other disadvantage is the rather slow learning rate of the multi-layered perceptrons.

In the two-layered perceptrons that we used, the number of hidden units determines the number of decision boundaries the neural net is able to learn. Training of a neural net results in fact in positioning the decision boundaries in such a way that the sum-of-squares error function is minimized.

Further investigation of optimal (well-formed) subplans for the cost structure with purchasing revealed much of the structure of the target mapping. This problem specific prior knowledge can be used in two ways. First, because it says something about the number of decision boundaries of the mapping, it is an indication of how many hidden units are required for a multi-layered perceptron, i.e., it guides the choice of a good network topology. Second, this information can be used to initialize certain weights in such a way that they already represent an existing decision boundary at startup and change the weight update scheme so that it does not alter these values. This will speed up the learning of the multi-layered perceptron.

6.3 Extension to other problems

We developed an approach which exploits the strong points of multi-layered perceptrons for the uncertainty component and which builds upon the numerous results and techniques from the corresponding off-line problems for the combinatorial component. In this way we combine the best of both fields. We conclude that such a two-stage approach may contribute significantly to the performance of on-line problems if the following conditions are satisfied.

- For the learning problem of the first stage there is no efficient solution approach available. Furthermore there is sufficient relevant data available to construct a sufficiently large set of learning examples.
- The combinatorial problem of the second stage can be computed efficiently.

This approach can also be used for other on-line problems. A problem that can be solved with our approach must (at least) satisfy the following requirements:

1. There is a linear ordering in the processing of arriving objects.
2. For the n -object problem there exists a property equivalent to the decomposition property of Wagner and Whitin for lot-sizing. So there exist regeneration points and subplans.
3. Subplans can have different lengths.

An example of a class of problems that satisfies the above requirements are on-line problems where objects that arrive over time are processed in batches of varying length.

References

- AGGARWAL, A., AND J.K. PARK [1993], Improved algorithms for economic lot size problems, *Operations Research* **41**, 549–571.
- BAHL, H.C., L.P. RITZMAN, AND J.N.D. GUPTA [1987], Determining lot sizes and resource requirements: a review, *Operations Research* **35**, 329–345.
- BAKER, K.R. [1977], An experimental study of the effectiveness of rolling schedules, *Decision Sciences* **8**, 19–27.
- BLACKBURN, J.D., AND R.A. MILLEN [1980], Heuristic lot-sizing performance in a rolling-schedule environment, *Decision Science* **11**, 691–701.
- CARLSON, R.C., S.L. BECKMAN, AND D.H. KROPP [1982], The effectiveness of extending the horizon in rolling production scheduling, *Decision Sciences* **13**, 129–146.
- FEDERGRUEN, A., AND M. TZUR [1994], Minimal forecast horizons and a new planning procedure for the general lot sizing model: nervousness revisited, *Operations Research* **42**, 456–468.
- GORHAM, T. [1968], Dynamic order quantities, *Production and Inventory Management* **9**, 75–81.
- HARRIS, F.W. [1913], How many parts to make at once, *Factory: The Magazine of Management* **10**, 135–136.
- HASSOUN, M.H. [1995], *Fundamentals of artificial neural networks*, MIT Press, Massachusetts.
- HAX, A.C., AND D. CANDEA [1984], *Production and Inventory Management*, Prentice-Hall, Englewood Cliffs, New Jersey.

- JAGANNATHAN, R., AND M.R. RAO [1973], A class of deterministic production planning problems, *Management Science* **19**, 1295–1300.
- KARP, R.M. [1992], On-line algorithms versus off-line algorithms: how much is it worth to know the future?, in: J. van Leeuwen (ed.), *Algorithms, Software, Architecture*, Elsevier Science Publishers, 416–429.
- LUNDIN, R.A. [1973], *Planning Horizon Procedures for Production-Inventory Systems with Concave Costs*, Ph.D. thesis, University of Chicago.
- LUNDIN, R.A., AND T.E. MORTON [1975], Planning horizons for the dynamic lot size model: Zabel vs. protective procedures and computational results, *Operations Research* **23**, 711–734.
- REED, R [1993], Pruning algorithms - A survey, *IEEE Transactions on Neural Networks* **4**, 740–747.
- SILVER, E.A., AND J.C. MEAL [1973], A heuristic for selecting lot size quantities for the case of deterministic time-varying demand rate and discrete opportunities for replenishment, *Production and Inventory Management* **14**, 64–74.
- SILVER, E.A., D.F. PYKE, AND R. PETERSON [1998], *Inventory management and production planning and scheduling*, John Wiley & Sons, Chichester.
- STEHOUWER, H.P. [1997], *On-line Lot-sizing with Perceptrons*, Ph.D. thesis, Eindhoven University of Technology.
- WAGNER, H.M., AND T. WHITIN [1958], Dynamic version of the economic lot size model, *Management Science* **5**, 89–96.
- ZWIETERING, P.J., M.J.A.L. VAN KRAAIJ, E.H.L. AARTS, AND J. WESSELS [1991], Neural networks and production planning, *Proceedings of the Fourth International Conference on Neural Networks and their Applications*, Nîmes, France, 529–542.