Evaluating order picking efficiency under demand fluctuations

Jens Bürger
Centro de Investigaciones en Nuevas Tecnologías Informáticas, Universidad Privada Boliviana
Institute for Computational Intelligence, Universidad Privada Boliviana

Abstract
Storage location assignment is a dynamic problem due to product lifecycles and time-varying demand patterns. We demonstrate the impact of demand fluctuations on order picking times for frequency-based and genetic algorithm-based storage assignment policies. Our results provide the base for developing re-warehousing strategies to maintain order picking efficiency over time.

Keywords: Order-picking, storage location assignment problem, genetic algorithm

1. Introduction

Order picking, the process of collecting and sorting a set of products according to customer orders, is the main cost driver in warehouse operations (Tompkins et al. 2010). Some estimates go as far as assigning 65% of operational costs of warehouses to order picking (Coyle et al. 1996). Order picking consists of travel between product locations, retrieval of the specific stock keeping unit (SKU), and sorting of SKUs according to customer orders. Travel is estimated to contribute approximately 50% of the total order picking costs (Tompkins et al. 2010). In sum, efforts to reduce warehouse operation costs depend strongly on efficient order-picking. Reducing order picking time and costs can be achieved through a set of mutually dependent strategies. Warehouse layout, with the spatial definition of racks and aisles, predetermines the possibility to optimize travel routes within a warehouse (Bassan et al. 1980). Closely related are routing strategies that aim to traverse the warehouse as to minimize total travel distance required to fulfill orders (Roodbergen and Koster 2001). Efficiency of warehouse traversal can be improved further by order batching (Gademann and Velde 2005). Efficient storage location assignment policies that take product demand statistics into account can further improve order picking (Brynzér and Johansson 1996).

While all these order-picking optimizations present valid approaches for any given moment, they typically do not present insights into the dynamics of warehouse operations. Having a warehouse inventory consisting of a larger number of products, potentially distributed across various different product categories, picking efficiency is a time-dependent function. Due to daily, weekly, monthly, seasonal and yearly demand fluctuations, order picking efficiency can vary greatly under such time-varying dynamics (Kofler et al. 2015). The recognition of order picking as a dynamic problem, for the very same reasons as outlined so far, has been recently addresses by Kofler et al. 2011. For demand patterns obtained over a period of four months the relation between slotting strategies (turnover, affinity) and optimization strategies (re-warehousing, healing) was investigated with respect to finding a compromise between picking efficiency and costs incurred by re-arranging product locations. As the authors mentioned, the objective was to evaluate the effectiveness of the different strategies, not to evaluate robustness over a longer period of time.

In this paper we investigate in more detail the time-dependent dynamics of order picking efficiency in relation to slotting (or the storage location assignment problem (SLAP)). Based on a 5 year demand dataset we investigate how picking efficiency is degrading with time for different slotting policies. Based on a near-optimal storage assignment for an initial period, we show how time-varying demand fluctuations effect order picking efficiency. In connection with different picker capacities we show the trade-off between initially optimized solutions and solutions that are more robust to demand fluctuations.

2. Methodology

2.1. Storage Location Assignment Problem

Storage location assignment searches an order-dependent configuration from a list of products $P = \{p_1, p_2, \ldots, p_{n-1}, p_n\}$ of length n and is akin to the traveling salesman problem (TSP). Due to the time complexity of $O(n!)$ of this class of problems, it is not feasible to exhaustively search the solution space for an optimal configuration. The TSP has been extensively studied and while effective search algorithms exist, none of them can guarantee optimality of the found solutions (Laporte 1992). Similarly, defining a storage location assignment for any realistic warehouse with n products requires the use of constrained or heuristic search methods that produce near-optimal solutions in a reasonable amount of time (Quintanilla et al. 2015). To study order-picking efficiency of a given storage location assignment in response to time-varying product demand we implement three assignment policies: random slotting, turnover-based slotting and a genetic algorithm based slotting policy. The random assignment is completely unaware of product demand or product correlations (affinity). The turnover slotting strategy evaluates product demand over a given period and assigns products to storage locations as a function of their total demand and their distance to the depot. In theory, this optimizes picking efficiency as products in highest demand are located closest to the depot. As a third method we implement a genetic algorithm (GA) for the SLAP. Based on a set of reference orders over an initial period, the GA tries to optimize the total picking distance, intrinsically taking into account product turnover and order correlations. As a heuristic it does not guarantee an optimal solution, but can calculate near-optimal solutions in a reasonable amount of time even for large numbers of products. The GA implements a set of possible storage locations as chromosomes with specific products encoded as genes. In such a genetic representation a chromosome represents a specific permutation of the available genes (products assigned to specific locations). We initialize the GA with a population size of 30 randomly generated individuals defined by their chromosomes. The GA then evolves the population for 200 generations through fitness evaluation, selection, crossover and mutation. According to Mitchell 1998 we use common GA parameters of 0.9 for the crossover rate and 0.01 for the mutation rate. All parts of the GA were implemented with python's *deap* package (Fortin et al. 2012).

2.2. Warehouse modelling

As the demand dataset in our experiments did not contain information about specific product details, other than demand, we model the warehouse in a generic manner assuming homogeneous storage cells and product properties. We define the warehouse as a two-dimensional grid with storage cells on only one level (no vertical storage). We design the warehouse shape parameter r with an x/y ratio of 1 (Hall 1993). The basic distance metric in our model, for x and y dimensions, is a unitless grid cell. Accordingly, each storage cell occupies exactly one grid cell as well as the picker can move one grid cell at a time. In total the warehouse contains

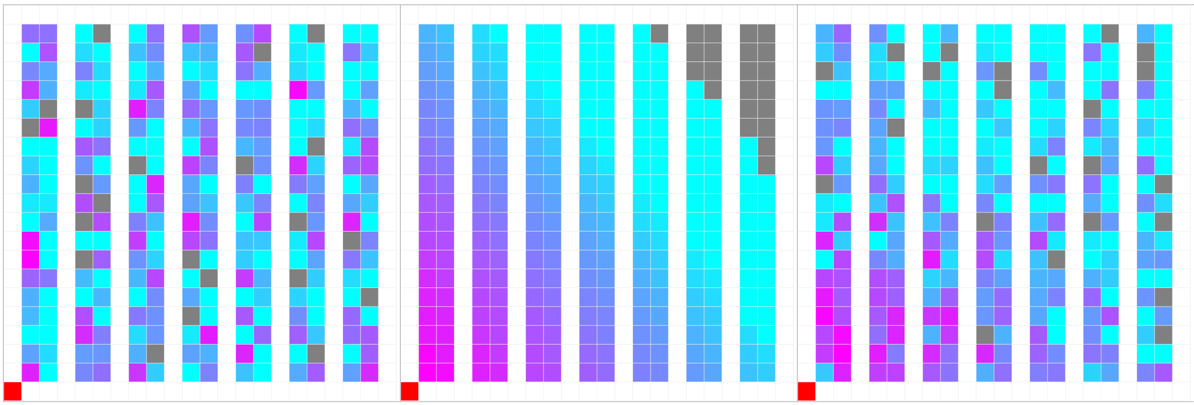

(a) Random Policy (b) Turnover Policy (c) GA Policy

Figure 1: (a) Random storage assignment, (b) turnover storage assignment and (c) an instance of the genetic algorithm storage assignment. Warmer colors indicate products in high demand, cooler colors in low demand. As the product demand differs by orders of magnitude across products, the color represents the log demand.

m racks with s storage cells on each side of the rack as to satisfy $2ms \geq n$, which assures that there are at least as many storage cells as there are products. An illustration of the warehouse and corresponding storage assignments is shown in Figure 1 for the random, turnover and GA slotting policies.

2.3. Order picking

For the process of order picking we define a single picker. The picker receives a daily demand from a set of P products with a total unit demand D. According to the picker capacity i this order is then split into B batches so that $B * i \geq D$. As the used dataset did not contain information on specific customer orders, but only on daily demand we randomized the generation of the batches B and evaluate order picking over the total daily demand D. We model each product as having equal size and weight and the picker has a capacity to collect i items per batch (an item is a single unit of a given product demand). In this project we will evaluate order picking efficiency as a function of picker capacity for i = 100, 250 and 500. When collecting items, the picker follows a mid-point routing strategy (Hall 1993) and overall order picking efficiency is measured as average distance per product for any given order list with n products.

2.4. Product demand

The general hypothesis of our presented work is that order picking efficiency will degrade with time as product demand patterns change due to daily, weekly, monthly, seasonal and/or yearly fluctuations. To test this hypothesis we use a demand dataset obtained from kaggle.com (Zhao 2017). This dataset contains product demand for over 5 years for more than 2.100 products across 4 different warehouses. For our analysis of the operations of a single warehouse we select ”Whse_C” with a total of 244 products. We therefore model a warehouse typical for manual picking and of the size of a small to medium sized enterprise. The patterns present in the dataset provide both, temporal and quantitative differences relevant for our analysis as shown in Table 1. Order frequency ranges from bi-daily to a single time per year with a median of 99 sales days over a 5 year period. Number of products sold per day ranges from 1 to 81 with a median of 26. Order quantity (unit demand) ranges by several orders of magnitude from a single unit up to 33.000 units of a specific product per day. These statistics provide the base for the time-varying demand patterns affecting efficiency of order picking.

Table 1: Demand statistics

	min	median	max
Product order frequency	4	99	674
Products per day	1	26	81
Unit demand per day	1	1727	33195

3. Results

3.1. Single optimization

As a first reference analysis we have computed order picking efficiency for all three slotting policies over a five year range with only a single optimization based on an initial 30 days period. The evaluation was done for the three different picker capacities i of 100, 250, and 500 items and the results of order picking efficiency are shown in Table 2. We can note that the turnover policy generally performs best for all picker capacities i. The data also shows that with increasing picker capacity, the advantage of dedicated slotting policies over a random one becomes smaller. This confirms the results from De Koster et al. 2007 and Lu et al. 2016. What the table does not clearly illustrate are the changes in order picking efficiency over time. To better understand the dynamics of picking efficiency we have to look at the time series plots. In Figure 2 we show order picking efficiency for all three policies with a picker capacity $i = 250$. We can point out two conclusions. First, while the picking distance for the random policy (after an initial period of lower overall demand) does not indicate an apparent long-term upward or downward trend despite seasonal fluctuations, the turnover and GA policies show a continuous degradation of picking efficiency. And second, while in average the turnover policy is preferable, the GA policy showed higher efficiency for the initial period of 30 days. The difference between initial order picking efficiency between the turnover and GA policies becomes more pronounced as we increase the picker capacity. For i=100, 250 and 500, the GA performed 2%, 5.5% and 24% better, respectively. This confirms the assumption that the GA intrinsically exploits product correlations which has increasing benefits the more products are picked along a single route.

3.2. Continuous optimization

With single optimization as reference, we now present results for continuous storage location optimization. Every 30 days we updated product storage locations based on the turnover and GA policy (as random assignment is not to be optimized we leave this policy out for this

Table 2: Average per product picking distance for different slotting policies and picker capacities.

	Capacity = 100				Capacity = 250				Capacity = 500			
	min	mean	max	std	min	mean	max	std	min	mean	max	std
Random	8.00	29.27	72.0	7.43	10.00	24.47	74.0	6.35	6.95	21.79	74.0	8.02
Turnover	**2.00**	**18.45**	**52.0**	5.80	**2.00**	**16.91**	**52.0**	**5.05**	**2.00**	**15.29**	**52.0**	**4.95**
GA	6.75	20.41	74.0	**5.38**	7.33	18.57	66.0	5.64	4.00	17.03	70.0	5.88

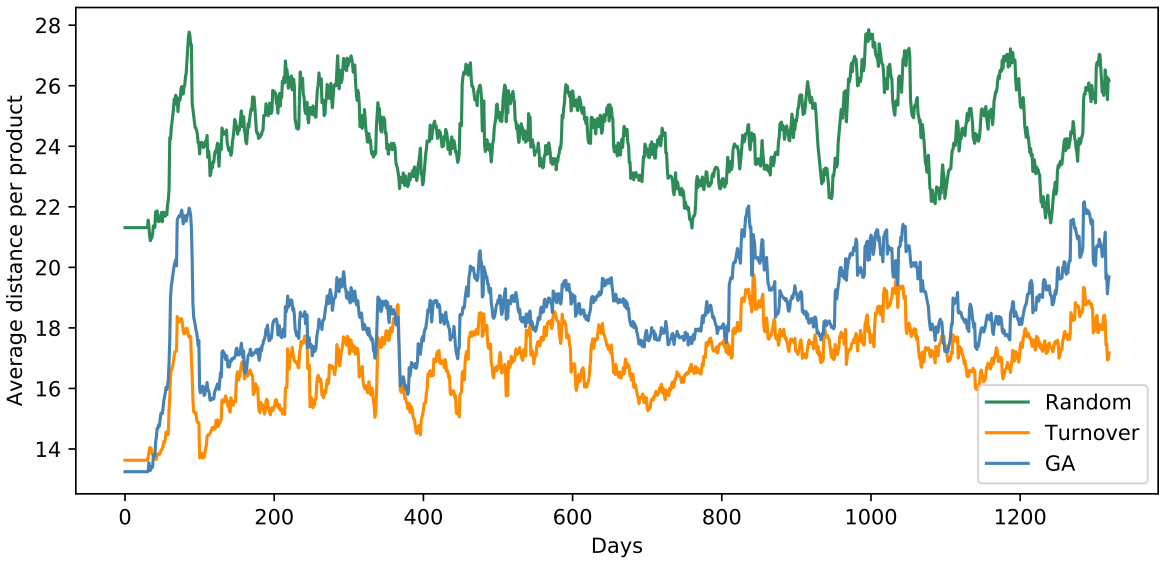

Figure 2: Average per product picking distance for different slotting policies and a picker capacity of 250 with single initial optimization.

Table 3: Average per product picking distance for different slotting policies and picker capacities under continuous optimization

	Capacity i = 100				Capacity i = 250				Capacity i = 500			
	min	mean	max	std	min	mean	max	std	min	mean	max	std
Turnover	4.00	15.20	48.00	4.07	4.00	14.40	48.00	3.96	4.00	13.07	48.00	4.24
GA	**1.78**	**14.54**	**23.35**	**3.08**	**1.71**	**13.01**	**22.64**	**2.92**	**1.77**	**11.81**	**19.99**	**3.07**

analysis). The results for picking efficiency under continuous optimization is shown in Table 3. Confirming the results of the initial optimization from the previous section, the GA significantly outperforms the turnover slotting policy. For all picker capacities the GA exhibits lower mean and standard deviation in the picking distance per product. We can attribute these results to the fitness evaluation of the GA. While the turnover policy only takes the summed demand of each product into account, the GA, based on evaluating fitness through the actual travel distances for a set of reference orders, implicitly evaluates product demand, demand correlations across products and the routing strategy as a multi-objective optimization problem. As one can see in Figure 1 the GA assigns some high demand products further away from the depot. While this implies frequent longer distances for high demand products, it is offset by placing correlated products nearby. The time-series plot for continuously optimized storage assignment (with i = 250) is shown in Figure 3. We can note that with continuous optimization the initial picking efficiency can be roughly maintained (though subject to seasonality in product demand we cannot determine a clear trend).

3.3. Sensitivity to demand changes

Based on the two presented results for single and continuous optimization we can point to an important design criteria impacting re-warehousing strategies (while Kofler et al. 2011 made a distinction between re-warehousing and healing, we use the term re-warehousing in a general sense where healing is considered an incremental re-warehousing). Under single optimization

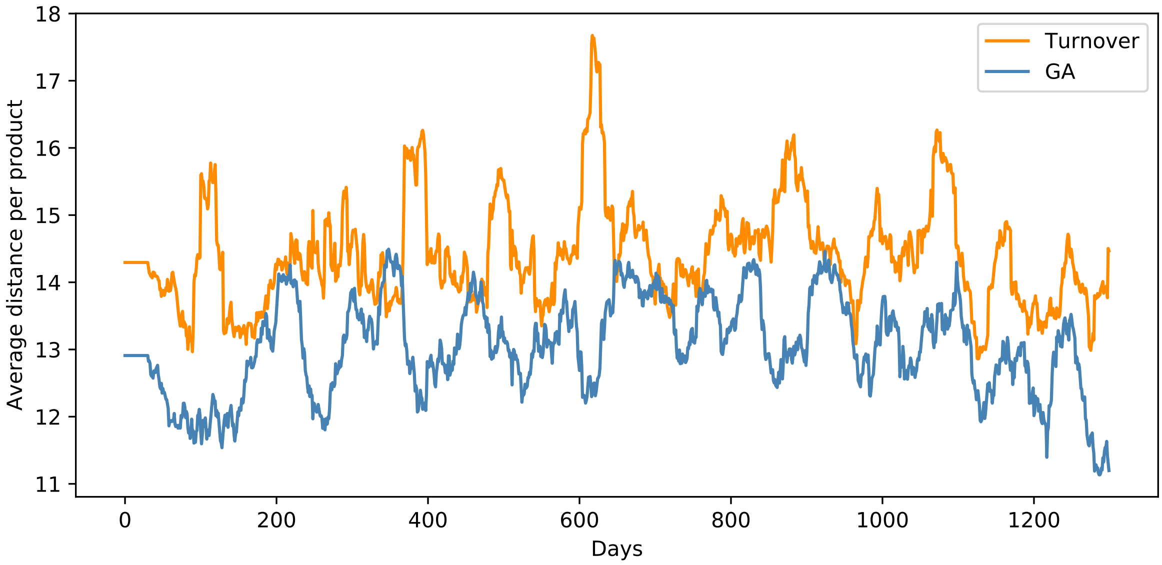

Figure 3: Average per product picking distance for different slotting policies and a picker capacity of 500 with continuous 30 days optimization.

we have seen that the turnover method performs best, if evaluated over a longer period of time. With continuous optimization the GA (as one example of a heuristic search method) shows clear advantages independent of how demand patterns change over time. Therefore warehouse operations are faced with the compromise between efficiency and robustness. The efficiency obtained with the GA was based on exploiting demand correlations across products (in accordance with results from Kofler et al. 2011). As order patterns and therefore correlations change, the GA assignment loses its efficiency. On the other hand, while the turnover policy did not capitalize on some deeper relations between products, it made the approach more robust to demand changes. As discussed earlier, the random policy is the only one fully insensitive to changes in demand patterns, at the expense of being far from efficient, especially for smaller picker capacities. In Figure 4 we plot the cumulative additional picking distance per product between single and continuous optimization as an indicator for the sensitivity to changing demand patterns. For both, the turnover and the GA policy, we see a constant linear increase in distance per product with the difference that the GA increases at a faster rate due to the larger difference in optimized and non-optimized performance. It is interesting to note that, while demand patterns show clear and pronounced seasonal variations in order quantity, the decrease in picking efficiency follows a very linear function. Under considerations of warehouse specific re-warehousing costs, re-warehousing strategies and policies can then be derived based on the picking efficiency baseline (picking distance with continuous optimization), the loss in picking efficiency (cumulative additional distance) and re-warehousing costs.

4. Conclusion

In this work we analyzed order picking efficiency as a time-dependent function that can inform re-warehousing policies. The presented time-series plots have shown in much more detail how changing demand patterns affect order picking efficiency. We have pointed to the differences in efficiency of different slotting policies as well as their robustness to changes in demand patterns. The presented insights into optimality and robustness therefore complement

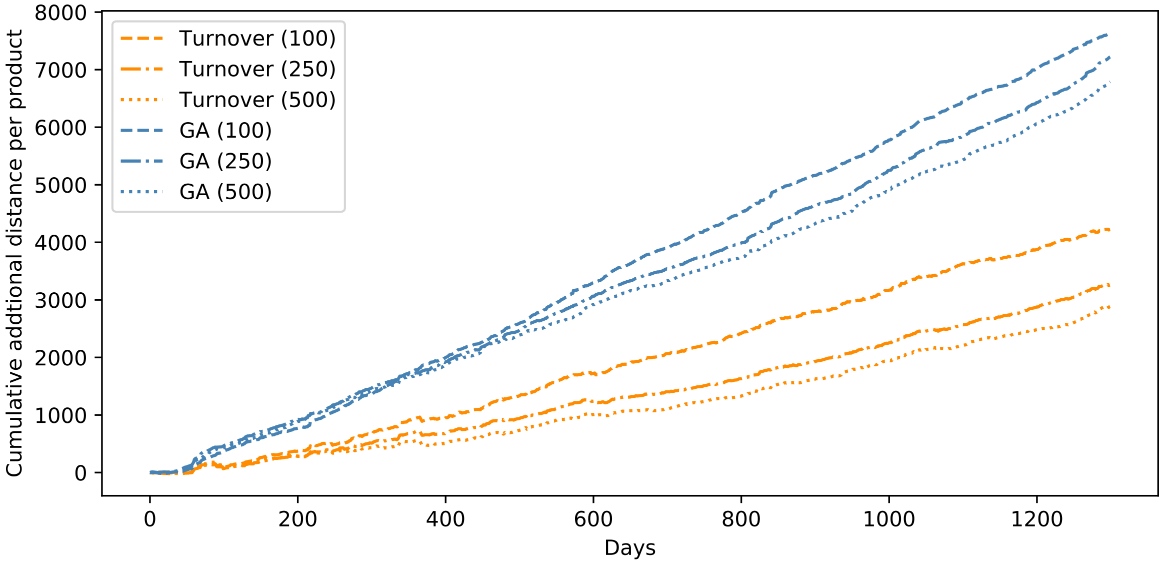

Figure 4: Cumulative additional distance per product of single initial warehouse optimization as compared to continuous optimization. Values in parenthesis are the different picker capacities.

previously published work on the dynamic storage location assignment problem (Kofler et al. 2011; Kofler et al. 2015). A future aspect is the further complexification of the used framework. In this work we have reduced complexity of the warehouse model and the picking process as to study the impact of demand fluctuations and slotting policies on picking efficiency. We expect that the general conclusions of our study remain valid, but inclusion of further aspects such as vertical storage, heterogeneous storage cells, different routing strategies, datasets with different demand patterns, etc. can alter the relative results to some extend. In summary, re-warehousing strategies will have to be designed around the compromise between optimality and robustness. In general, we can say that if re-warehousing can be done at relatively low costs (such as the healing strategy Kofler et al. 2011) than efficiency is preferable over robustness. If re-warehousing is a costly and disruptive process, one might choose more robust slotting policies that can maintain a somewhat stable order picking efficiency.

References

Bassan, Y., Y. Roll, and M. J. Rosenblatt. 1980. Internal layout design of a warehouse. *AIIE Transactions* 12 (4): 317–322.

Brynzér, H., and M. I. Johansson. 1996. Storage location assignment: Using the product structure to reduce order picking times. *International Journal of Production Economics* 46:595–603.

Coyle, J. J., E. J. Bardi, and C. J. Langley. 1996. The management of business logistics. Vol. 6. West publishing company St Paul, MN.

De Koster, R., T. Le-Duc, and K. J. Roodbergen. 2007. Design and control of warehouse order picking: A literature review. *European journal of operational research* 182 (2): 481–501.

Fortin, F.-A., F.-M. D. Rainville, M.-A. Gardner, M. Parizeau, and C. Gagné. 2012. DEAP: Evolutionary algorithms made easy. *Journal of Machine Learning Research* 13 (Jul): 2171–2175.

Gademann, N., and S. Velde. 2005. Order batching to minimize total travel time in a parallel-aisle warehouse. *IIE transactions* 37 (1): 63–75.

Hall, R. W. 1993. Distance approximations for routing manual pickers in a warehouse. *IIE transactions* 25 (4): 76–87.

Kofler, M., A. Beham, S. Wagner, and M. Affenzeller. 2015. "Robust Storage Assignment in Warehouses with Correlated Demand". In *Computational Intelligence and Efficiency in Engineering Systems*, 415–428. Springer.

Kofler, M., A. Beham, S. Wagner, M. Affenzeller, and W. Achleitner. 2011. "Re-warehousing vs. healing: Strategies for warehouse storage location assignment". In *Logistics and Industrial Informatics (LINDI), 2011 3rd IEEE International Symposium on*, 77–82. IEEE.

Laporte, G. 1992. The traveling salesman problem: An overview of exact and approximate algorithms. *European Journal of Operational Research* 59 (2): 231–247.

Lu, W., D. McFarlane, V. Giannikas, and Q. Zhang. 2016. An algorithm for dynamic order-picking in warehouse operations. *European Journal of Operational Research* 248 (1): 107–122.

Mitchell, M. 1998. An introduction to genetic algorithms. MIT press.

Quintanilla, S., Á. Pérez, F. Ballestín, and P. Lino. 2015. Heuristic algorithms for a storage location assignment problem in a chaotic warehouse. *Engineering Optimization* 47 (10): 1405–1422.

Roodbergen, K. J., and R. d. Koster. 2001. Routing methods for warehouses with multiple cross aisles. *International Journal of Production Research* 39 (9): 1865–1883.

Tompkins, J. A., J. A. White, Y. A. Bozer, and J. M. A. Tanchoco. 2010. Facilities planning. John Wiley & Sons.

Zhao, F. 2017. Forecasts for Product Demand. `https://www.kaggle.com/felixzhao/productdemandforecasting/home`. Version 1. Accessed: 2018-09-20.