Mixing machine learning and optimization for the tactical capacity planning in last-mile delivery

Mixing machine learning and optimization for the tactical capacity planning in last-mile delivery

Edoardo Fadda
DAUIN
Politecnico di Torino
Torino, Italy
edoardo.fadda@polito.it

Stanislav Fedorov
DAUIN & CARS@Polito
Politecnico di Torino
Torino, Italy
stanislav.fedorov@polito.it

Guido Perboli
DIGEP & CARS@Polito
Politecnico di Torino
Torino, Italy
guido.perboli@polito.it

Ivan Dario Cardenas Barbosa
University of Antwerp
Antwerpen, Belgium
IvanDario.CardenasBarbosa@uantwerpen.be

***Abstract*—Tactical Capacity Planning (TCP) is becoming a crucial part of logistics in the current environment of demand-driven economics. This paper proposes an innovative approach in the TCP setting, consisting of using the collected historical data of the geographical position and the volume of the orders to plan the capacity requirements for the next day. To this end, the clustering of the city to microzones is introduced using K-means clustering. Then, four different methods (Gaussian Process regression, ARIMA model, Neural Network regression, and Long Short Term Memory network) are used to forecast the next day order volume for each of the clusters. Finally, the Variable Cost and Size Bin Packing problem solved with the predicted demand to outline the usage of a heterogeneous fleet required to serve the next time period. Through experiments on the real data, we conclude, that the proposed algorithm is satisfying the decision safety framework with completely unknown demand and could also be used for other demand forecast applications.**

***Index Terms*—Capacity planning, Variable Sized and Cost Bin packing, last-mile delivery, Machine Learning**

I. Introduction

In the field of last-mile logistics, the main critical resource to manage is the fleet of vehicles. In the last years, this resource has become more difficult to handle since, with the advent of the on-demand economy, very high levels of delivery services are requested: the time windows for the delivery are becoming very short (up to 2 hours in urban delivery) and the number of parcels to deliver is increasing exponentially. Moreover, the orders from e-commerce platforms are hardly predictable, due to their dependency on many external variables such as discounts in the shops, weekends, holidays, and other customers' reasons. Operating in these difficult conditions leads the last mile logistic operators to use third-party services and rented vehicles in order not to buy a fleet that is going to be used only partially. Thus, to decrease the operative costs, it is of paramount importance to estimate the capacity required to satisfy the next day delivery demand [1, 2, 3].

That question justifies the recent investigations on the Tactical Capacity Planning (TCP) problem, which consists of identifying the requirements in terms of quantity and types of vehicles to serve all the deliveries in the next day. To tackle this issue variants of the Bin Packing Problems have been developed in the past decade [4]. In particular, one of the approaches that have proven to be very effective in the management of the heterogeneous fleet is the Variable Cost and Size Bin Packing Problem (VCSBPP), which has been defined both in its deterministic and stochastic counterpart [5]. Nevertheless, introducing stochastic factors in mathematical programming is a nontrivial task since it leads to more complex problem formulations that become computationally intractable in the real setting. However, it is impossible to completely disregard the stochastic factors without introducing heavy errors in the decision process [6]. We then consider the issue of having an accurate forecast of the delivery demand by presenting an innovative technique mixing machine learning and optimization that enables companies to have a precise estimation of the next day vehicle needs.

In more details, since predicting this quantity by forecasting the deliveries for every single customer is difficult due to its intrinsic variability it is reasonable to consider the cumulative number of deliveries in a certain area of the city. Thus it is possible to take advantage of statistical results (such as the Central Limit Theorem) that guarantees the nice properties of cumulative data. This generates the question on the introduction of such urban areas, usually called *microzones*. A possible choice is to divide the city dynamically [7]. Nevertheless, as we can observe from Figure 1 which represents locations of delivery orders done for two different days, the delivery distribution follows the structure of the city (e.g. the central districts have more delivery orders than the residential one) and it is similar from one day to the other. For this reason, we consider a static city division into geographical microzones [8], and then we forecast the delivery demand relying on the historical data of delivery of each of them. Having the predicted delivery demand, it is possible to estimate the capacity and number of vehicles needed by the company in the next day by solving the deterministic VCSBPP instead of its stochastic version.

This work presents the implementation of the outlined idea and its application to the real-case scenario of the TCP problem, based on the collected historical delivery data. Besides, different recent forecasting methods suitable to our problem setting are applied, and their performances are compared. To the best of the author‘s knowledge, this is the very first attempt to access the demand forecast through the introduction of the microzones in the city and based on it to create the framework

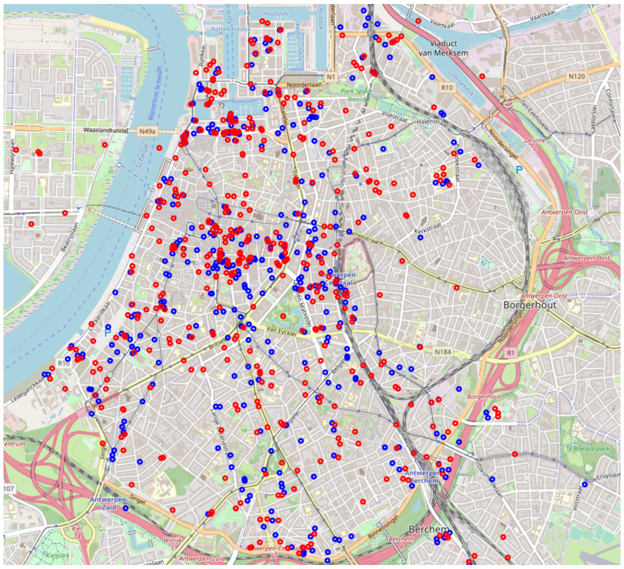

Fig. 1. Order locations in the city of Antwerp for two days highlighted with red and blue colors for the first and the second day correspondingly. The overall pattern remains the same.

able to recommend the quantity of the vehicles required to serve the next day delivery.

The paper is organized as follows. Section II covers the existing approaches similar to the one outlined in the current work; Section III describes the methodology used in the study, results of which are outlined in Section IV. Finally, Section V outlines the conclusions and future work.

II. State of the Art

The TCP problem is relevant in many industrial fields such as engineering [9], medicine [10], semiconductor manufacturing [11], and many others. In logistics applications as last-mile delivery, its application is more recent [12]. The existing literature of the problem usually focuses on the operational side [4], and only a few papers highlight the importance of the tactical decision level in logistics [13, 14]. Due to the lack of literature in this particular field, we focus our attention on a brief review of the approaches concerned with the crucial parts of the proposed method, i.e. the demand forecasting and clustering.

Demand prediction is important for many fields of application such as the food industry, perishable goods supply or crowd shipping [15]. Nowadays, e-commerce is presenting a new challenge for demand forecast since it is characterized by highly fluctuated data. Not many methods can capture such a complex data structure due to the presence of high uncertainty and no evident trend. Thus, more sophisticated and robust approaches have to be applied. For example, in [16] authors successfully used the Long Short Term Memory (LSTM) network to predict the logistics delivery demand in the manually defined sub-region of the city. The result of this work is a two-dimensional LSTM which enables the company to have a reliable support decision system for future decisions. However, this approach is characterized by many empirical parameters and requires manual division of the city, hence further work is required towards this direction.

In our search of the attempts of improving this result, we observe that the main lack of papers dealing with delivery forecasting is that they just consider time dependency and do not consider spatial correlation. In other words, it is crucial to account for the variables connected to the spatial allocation of the order to improve the performance of the time series forecast. However, a straightforward introduction of the multivariate regression complicates the problem [17]. Thereby, the existing approaches adopt clustering to access the spatial information of the demand, as hierarchical clustering approach in [18]. One more example of clustering is introduced in [19], where has been applied the division of the customers into different logical segments in accordance to the monthly volume of product delivered for the further applications of the demand forecasting methods to each of them.

Despite clustering techniques received little attention in the context of the TCP problem, it has been widely used in the operational level logistics problems, as Two-Echelon Vehicle Routing Problem [20, 21]. More recently, it is used in the on-demand economy conditions for the usage of crowd drivers and automated vehicles. Notice that, clustering the city in the geographical area implies the introduction of the microzones, which has been proven to have a direct impact on the operational level [22], due to the straightforward assignment of each of the vehicles to serve a predefined area. In this paper, we unify the advantages of the aforementioned approaches by accessing the demand forecast tools through coupling them with the automatic clustering of the city regarding the previous orders' historical data.

III. Methodology

In this section, we describe step by step the proposed approach to the scenario of the TCP problem, where we assume the completely unknown demand and, based on the historical data of orders done, have to give a decision about the fleet amount required to serve the next day. This problem setting corresponds to the real managerial decision about the quantity of the vehicles to rent in the next time period, that is in our case become the daily decision due to the lack of data. To this end, we aim to construct the decision support framework, in which we run the algorithm "on the fly", collecting the data, then forecasting demand and giving recommendations about the fleet requirements.

As stated in Section I, for supporting accurate demand forecast we consider spatial information of the previous orders by introducing geographical microzones, computed by a clustering algorithm, and then predicting the total delivery demand volume for each of them. Thus, our starting goal is to group the orders in different clusters according to their geographical location. Hence we deal with the case of unsupervised learning, and among all the possible clustering algorithms in this field, we select the K-means because it is easy to tune state-of-the-art algorithm [23, 24]. After the data being clustered, we can forecast the total next day demand for each of the defined microzones. In this work we explore four different forecasting approaches, comparing the performances on the data provided.

Firstly, we fit the Gaussian Processes (GP) regression model [25]. To apply it we introduce the theoretical assumption of the jointly Gaussian distribution of the random variables, which allows us to use the closed-form solution, which is quickly adaptable to changes in the data trends. The fitting process in the context of the GP regression is the optimization of the parameters of the mean and kernel functions, which are chosen to be linear mean and Radial basis function, respectively. The output from the fitted model is the mean $m(t+1)$ and the variance $\sigma^2(t+1)$ in the next-day (time $t+1$), that provides the safety of choosing the predicted value. More precisely, we recall that the 2σ bounds around the predicted mean guarantee the 99% probability of finding the next sample in the considered boundaries, which is proven in [26]. We use this fact to introduce safety in the next-day demand by choosing the prediction for the next time slot to be $v(t+1) = m(t+1) + \frac{1}{2}\sigma(t+1)$.

Another model to compare the performance is the ARIMA model, which consists of the three basic components, which are differencing order, auto-regressive (AR) and moving-average (MA) terms; and characterized by associated parameters, which are d, q, p correspondingly [27, 28, 29]. To set the parameters, we used the python package *pmdarima* model[1], able to fit the ARIMA model automatically defining the coefficients.

The Neural Network-based forecast recently showed high accuracy and robustness to highly chaotic data, which is the case of demand delivery history. Hence, we choose to use two models based on it. As the basic case, we feed the 1 hidden layer NN with the data, that have the input lag of one day. In such a way, we move the window over the delivery history data for each cluster to train NN, and for the last step, we predict the expected volume. As for the LSTM network, it is the type of the basic Recurrent Neural Network, that has a certain cell structure, the details of which can be found in [16, 30, 31]. For our purpose, we train the simple LSTM network without the hidden layers by feeding it with the data with the input lag.

Having the predicted total volume of orders for the each of defined clusters we can move to the next step. The TCP setting requires not just the conclusion on the overall volume of goods to move, but the number and types of vehicles to use. To compute these quantities we use the VCSBPP. In the problem formulation, the bin volumes are assigned in correspondence to the vehicles presented in the given data. In particular, the types of vehicles used are three types of vans with big capacity (150-500 kg) and one type of cargo bicycle with a small capacity of 20 kg. For the test case, the bin costs are defined, with the help of the company, to be $P_b = C_b + 0.2 * \max(C_b)$, where P_b is the cost of the selected type b of the bin and C_b is the capacity of associated bins. This choice reflects the real case scenario of increased cost in terms of rent, fuel, and emissions for bigger vehicles, and the addition of the 20% of the capacity of the biggest bin represents the fixed costs as, for example, the salary of a driver.

Concerning the definition of the items in the VCSBPP formulation, the predicted total parcel volume for the whole cluster is considered as one item, which is the essential feature of our approach. In this way, we are assigning each of the vehicles to a certain geographical microzone, since each of the items corresponds to the cluster and hence, to the microzone. However, the case of a small number of clusters leads to obtaining the items with a volume higher than any capacity available, since they represent the sum of the volumes of the parcels for the whole cluster. In such a case, the VCSBPP is infeasible by definition, and to deal with it, we increase the number of clusters since the K-means clustering algorithm introduces the new centers by separating the biggest existing cluster, which is clearly outlined in Figure 2. Thereby, increasing the number of clusters we lowering the demand for each of them, and thus balance the available vehicles volumes and the demand. This allows us to tune the number of clusters in a very easy manner, by the simple increment of it whenever we receive the infeasibility of the VCSBPP. We can also observe, that increasing the number of clusters leads to the better accuracy of the predictions but increases the computational burden since the forecast has to be made once for each cluster and it increases the size of the VCSBPP to solve. Moreover, in the limit case of the same number of clusters and customers, it is impossible to obtain an accurate prediction. The outlined approach is a practical way to balance the performance of the algorithm in terms of accuracy and computational complexity.

To test the proposed algorithm, we follow the approach in [32] and run the algorithm "on-the-fly", starting the simulation with the assumption of the knowledge of the orders made in the last 5 days. After receiving this information, we perform a clustering step with a low number of starting clusters. Then, we perform the forecast of the summed-up volume of the orders for each of the clusters. Each predicted quantity is considered as an item, that we want to pack into the vehicles available. We do this by optimizing the VCSBPP with the existing solver package[2]. If the problem is infeasible, we increase the number of clusters and repeat the procedure. The pseudo-code of the overall Capacity Forecast algorithm is outlined below as the Algorithm 1, where the n is the starting number of clusters, V_b is the set vehicle capacities given and b is the indicator set of bin type for each vehicle available. The $Data$ consists of the historical data of the orders done in the form of the geographical position of the order and the volume of the corresponding parcels. Further, the clustered $Data(c)$ is the total volume of orders separated by each cluster c.

The output of the algorithm is the number and types of vehicles to book, the predicted volume for each cluster v_n,

[1]http://alkaline-ml.com/pmdarima/

[2]http://www.gurobi.com

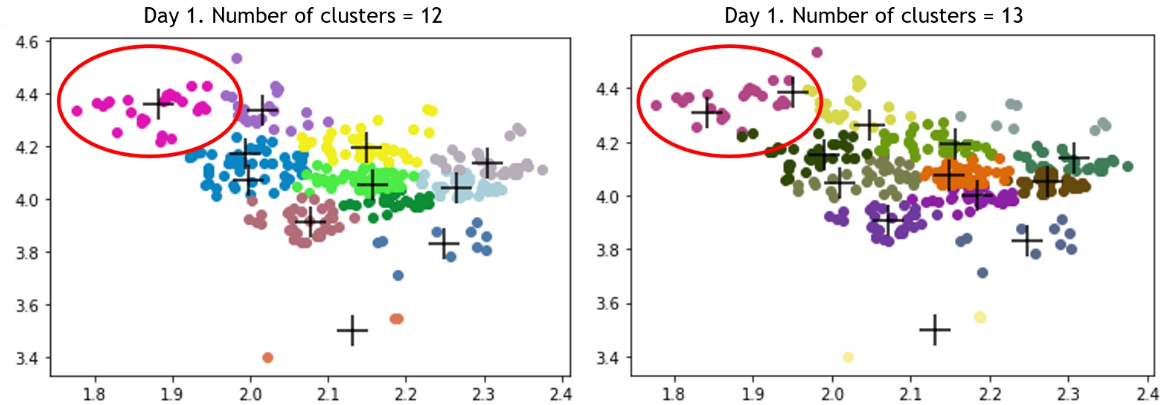

Fig. 2. First day microzones with different numbers of clusters. The K-means clustering introduce the new cluster by "separating" the biggest one of the precious run.

Algorithm 1 Capacity Forecast algorithm

procedure CAPACITY FORECAST(n, $Data$, V_b, b)
  while True **do**
    $Clusters \leftarrow Clustering(n, Data)$
    for c in $Clusters$ **do**
      $v_n \leftarrow Forecast(Data(c))$
    end for
    $Vehicles \leftarrow$ VCSBPP $(items = v_n, bins = V_b(b))$
    if Infeasible **then**
      $n \leftarrow n + 1$
    else
      return $(Vehicles, v_n)$
    end if
  end while
end procedure

and the total cost incurred, computed by the value of the objective function of the VCSBPP. To test the performance of the proposed method, we compare the cost obtained by using the outlined capacity forecast algorithm and the cost received from the additional solution of the VCSBPP using the real order volumes observation of the next day. By doing so we can estimate the relative loss of the resources for the hypothetical company. It is worthwhile noting that the usage of the VSCBPP is crucial to get an estimation of the cost of the assignment. In fact, since the VSCBPP is a discrete problem the relation between the parameter of the instance and the final cost is nonlinear. Thus, evaluating our method just by using the error in the forecast demand may lead to an unfair evaluation. This characteristic follows from the choice to consider discrete bins, hence if the solution of the VSCBPP is tight (the company has almost all the vehicle full) a small prediction error can lead us to rent one more vehicle thus to wrongly estimate the cost, while if the solution is not tight (the company has some free capacity), even if we made big errors in the demand, we are not going to wrongly estimate the cost.

IV. Experimental results

In this section, we present the results obtained by applying our algorithm to the 45 days demand prediction scenario. This data comes from 50 working days of parcel delivery history from February 1, 2019 to 17 of April, 2019. The observations are from the city of Antwerp (Belgium) and the records are provided by one of the biggest delivery companies, which desired to remain confident. The data frame consists of the geographical locations of the orders (the longitude and latitude), the day and time the delivery is done, the weight of the parcels, and the type of vehicle used for the delivery. Thus, using these data we construct a time series of the daily volumes of the parcels, assuming the knowledge of the previous 5 days delivery history.

As it is mentioned above, the principle step of the algorithm validation is the comparison of the objective function cost values of the VCSBPP solution obtained by passing into it the predicted order volumes and the real volumes. For simplicity, the obtained solutions are called the predicted and real or actual demand VCSBPP solutions accordingly. Since the cost and capacity of the bins in our setting are linearly dependent, it is possible to consider the difference in the cost as the difference in the capacity required to rent predicted by our algorithm and the actual required capacity of the next day. We expect to have equal or higher values of the predicted solution due to considering the total volume for each cluster as one item, and the safety treatment in the case of the GP regression. In Figure 3 is plotted the daily comparison of the VCSBPP objective function values for the demand predicted with GP regression and actual data. This forecasting method assumes the maximal safety in the decision about the required capacity, hence the predicted solution is way higher than the actual. From the TCP point of view, there must be no points where it is lower, i.e. there are no scenarios in which the company can not satisfy the demand. Hence, using this forecasting tool it is required to carefully choose the higher bound for the predicted values to be safe in satisfying the demand. In this case, we observe that the predicted solution is distinguished by a van with the highest capacity from the actual one, which is acceptable, but further improvement in the forecasting mechanics would decrease this gap.

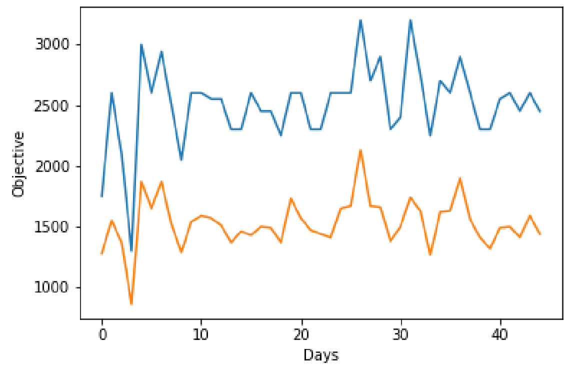

Fig. 3. The daily comparison of the values of the objective function for the solutions of the VCSBPP for the predicted with GP regression (blue) and actual (brown) demand.

Concerning the test of the other forecasting methods, their output accurately follows the pattern of the actual demand solution objective function values, as it is outlined in Figure 4, but there are multiple days, where it is lower than the real solution. Thus, the introduction of some safety to the decision is required. The fixed threshold for the predicted volumes for each cluster could be one of the options. However, the introduction of such a threshold is a highly empirical task and it should be treated regarding the financial perspective of the problem, which is out of scope for this paper.

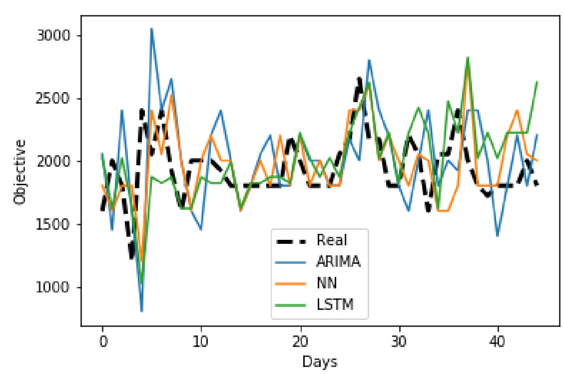

Fig. 4. The daily comparison of the values of the objective function for the actual solutions of the VCSBPP and the three testing models

The averaged among clusters demand prediction RMSE for all the forecasting models is outlined in Figure 5. This comparison shows the pure performance of forecasters, but in the condition of such complex correlations and constantly changing trends of demand forecasting data, it is not possible to expect a convergence of any forecast to a stable error level. We can observe, that GP regression is the most stable method, the LSTM network and ARIMA model shows the trend to improve the performance with the data enrichment; the one-layer NN-based forecast appears to be very accurate in the beginning and hence fast to train, but loses the performance with time, that could be due to the overfitting.

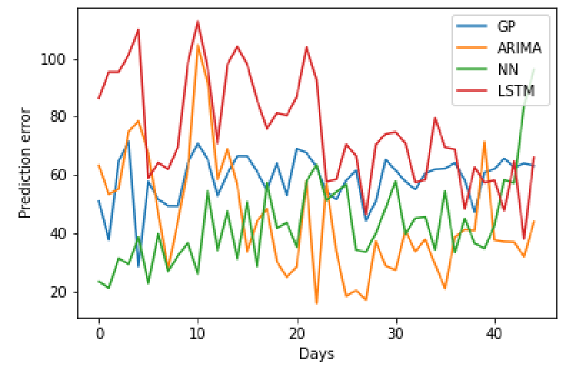

Fig. 5. The averaged among clusters demand prediction Root Mean Square Error for the four different forecasting methods

The VCSBPP solution structure for all four forecasting methods is outlined in Table I. It shows the percentage of the different capacities chosen by different algorithms, corresponding to the different sizes of the vehicles. This percentage is averaged over all the solutions of the 45 days, hence, the sum of it for each of the methods is not supposed to be 100%. From the table we can check which capacity is preferred by the logic of the overall framework, and the influence of the particular forecasting approach in particular. The results can be compared to the actual VCSBPP solution structure with the full knowledge of the real demand. As well, the bin costs and the mean of Root Mean Squared Errors (RMSE) for the total demand volume prediction in the 45 days are outlined.

TABLE I
THE VCSBPP SOLUTION STRUCTURE COMPARISON.

Capacities, kg	150	500	300	20	RMSE
Costs, eu	250	600	400	120	-
Real demand solution	52.86%	23.21%	37.05%	68.75%	-
GP prediction	9.29%	79.46%	35.98%	2.68%	58.5
ARIMA prediction	2.75%	61.76%	15.69%	0.98%	44.8
NN prediction	1.18%	61.27%	16.18%	0.49%	43.4
LSTM prediction	3.53%	52.49%	23.24%	0.49%	75.4

From the solution structure, it is clear that the proposed approach is prone to choose the biggest vehicle regardless of the forecasting tool. This because we consider the collected volumes of the orders inside each cluster as one item. However, this is not a problem for the case when the set of bin types does not differ much in capacities, as in the typical supply chain problem [5] and in last-mile and e-commerce applications, where the parcel sizes are small [33]. The best overall performance in terms of the prediction error shows the NN prediction, but the most error decrease with the data collection is showed by the LSTM network, which is more promising in the long time perspective.

V. Conclusions and future work

In this paper, we propose an innovative decision support system able to capture the daily orders demand behavior in the city by combining the regression, classification, and optimization methods. We developed a capacity forecast algorithm, which relies on the knowledge about the recent delivery made and provides a forecast about the capacity required for the next time period. The simulation of demand and capacity prediction for the 45 days shows the acceptable results for the concept of the TCP. Among four different regression models, the GP regression shows the capability to safely satisfy the next day demand, but the LSTM shows the better capability to capture the daily deviations.

This work shows the capability of the proposed system to work in the conditions of the completely unknown demand, which enables the whole variety of directions for future investigation. Some of them include the introduction of more sophisticated regression models and the application to additional urban scenarios and application domains.

References

[1] T. G. Crainic, N. Ricciardi, G. Storchi., and G. Storchi, "Models for evaluating and planning city logistics systems," *Transportation Science*, vol. 43, no. 4, pp. 432–454, 2009.

[2] G. Perboli, S. Fedorov, and M. Rosano, "The European Concept of Smart City: A Taxonomic Analysis," in *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*, 2020, pp. 1725–1730.

[3] E. Fadda, P. Plebani, and M. Vitali, "Optimizing monitorability of multi-cloud applications," in *Advanced Information Systems Engineering. CAISE 2016. Lecture Notes in Computer Science*, vol. 9694. Springer, 2016, pp. 411–426.

[4] G. Perboli, L. Gobbato, and F. Perfetti, "Packing problems in transportation and supply chain: new problems and trends," *Procedia-Social and Behavioral Sciences*, vol. 111, pp. 672–681, 2014.

[5] T. G. Crainic, L. Gobbato, G. Perboli, and W. Rei, "Logistics capacity planning: A stochastic bin packing formulation and a progressive hedging meta-heuristic," *European Journal of Operational Research*, vol. 253, no. 2, pp. 404–417, 2016.

[6] T. G. Crainic, G. Perboli, W. Rei, and R. Tadei, "Efficient lower bounds and heuristics for the variable cost and size bin packing problem," *Computers & Operations Research*, vol. 38, no. 11, pp. 1474–1482, 2011.

[7] E. Fadda, G. Perboli, and G. Squillero, "Adaptive batteries exploiting on-line steady-state evolution strategy," in *Applications of Evolutionary Computation. EvoApplications 2017. Lecture Notes in Computer Science*, G. Squillero and K. Sim, Eds., vol. 10199. Springer, Cham., 2017, pp. 329–341.

[8] J. Beckers, I. Cárdenas, and A. Verhetsel, "Identifying the geography of online shopping adoption in belgium," *Journal of Retailing and Consumer Services*, vol. 45, pp. 33–41, 2018.

[9] A. N. Carvalho, F. Oliveira, and L. F. Scavarda, "Tactical capacity planning in a real-world eto industry case: A robust optimization approach," *International Journal of Production Economics*, vol. 180, pp. 158–171, 2016.

[10] A. Larsson and A. Fredriksson, "Tactical capacity planning in hospital departments," *International journal of health care quality assurance*, 2019.

[11] S. J. Hood, S. Bermon, and F. Barahona, "Capacity planning under demand uncertainty for semiconductor manufacturing," *IEEE Transactions on Semiconductor Manufacturing*, vol. 16, no. 2, pp. 273–280, 2003.

[12] G. Perboli, R. Tadei, and L. Gobbato, "The multi-handler knapsack problem under uncertainty," *European journal of operational research*, vol. 236, no. 3, pp. 1000–1007, 2014.

[13] T. G. Crainic, G. Perboli, W. Rei, and M. Rosano, *Capacity Planning with Uncertainty on Contract Fulfillment*. CIRRELT, 2020.

[14] M. M. Baldi, D. Manerba, G. Perboli, and R. Tadei, "A generalized bin packing problem for parcel delivery in last-mile logistics," *European Journal of Operational Research*, vol. 274, no. 3, pp. 990–999, 2019.

[15] E. Fadda, D. Mana, G. Perboli, and R. Tadei, "Multi Period Assignment Problem for Social Engagement and Opportunistic IoT," in *2017 IEEE 41st Annual Computer Software and Applications Conference (COMPSAC)*, vol. 2, 2017, pp. 760–765.

[16] Y.-S. Lin, Y. Zhang, I.-C. Lin, and C.-J. Chang, "Predicting logistics delivery demand with deep neural networks," in *2018 7th International Conference on Industrial Technology and Management (ICITM)*. IEEE, 2018, pp. 294–297.

[17] M. Li, S. Ji, and G. Liu, "Forecasting of chinese e-commerce sales: an empirical comparison of arima, nonlinear autoregressive neural network, and a combined arima-narnn model," *Mathematical Problems in Engineering*, vol. 2018, 2018.

[18] J. Huber, A. Gossmann, and H. Stuckenschmidt, "Cluster-based hierarchical demand forecasting for perishable goods," *Expert systems with applications*, vol. 76, pp. 140–151, 2017.

[19] P. W. Murray, B. Agard, and M. A. Barajas, "Forecasting supply chain demand by clustering customers," *IFAC-PapersOnLine*, vol. 48, no. 3, pp. 1834–1839, 2015.

[20] G. Perboli, R. Tadei, and E. Fadda, "New valid inequalities for the two-echelon capacitated vehicle routing problem," *Electronic Notes in Discrete Mathematics*, vol. 64, pp. 75–84, 02 2018.

[21] E. Fadda, L. F. Tiotsop, D. Manerba, and R. Tadei, "The stochastic multipath traveling salesman problem with dependent random travel costs," *Transportation Science*, vol. 54, no. 5, pp. 1372–1387, Sep. 2020. [Online]. Available: https://doi.org/10.1287/trsc.2020.0996

[22] J. Janssens, J. Van den Bergh, K. Sörensen, and D. Cattrysse, "Multi-objective microzone-based vehicle routing for courier companies: From tactical to operational planning," *European Journal of Operational Research*, vol. 242, no. 1, pp. 222–231, 2015.

[23] F. Corno, L. De Russis, and T. Montanaro, "Estimate user meaningful places through low-energy mobile sensing," in *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2016, pp. 003 039–003 044.

[24] F. Corno, L. De Russis, A. Marcelli, and T. Montanaro, "An unsupervised and non-invasive model for predicting network resource demands," *IEEE Internet of Things Journal*, vol. PP, pp. 1–1, 07 2018.

[25] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[26] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.

[27] P. Castrogiovanni, E. Fadda, G. Perboli, and A. Rizzo, "Smartphone data classification technique for detecting the usage of public or private transportation modes," *IEEE Access*, vol. 8, pp. 58 377–58 391, 2020.

[28] A. Cuzzocrea, E. Mumolo, E. Fadda, and M. Tessarotto, "A novel big data analytics approach for supporting cyber attack detection via non-linear analytic prediction of ip addresses," in *Computational Science and Its Applications – ICCSA 2020*. Cham: Springer International Publishing, 2020, pp. 978–991.

[29] A. Cuzzocrea, C. K. Leung, D. Deng, J. J. Mai, F. Jiang, and E. Fadda, "A combined deep-learning and transfer-learning approach for supporting social influence prediction," *Procedia Computer Science*, vol. 177, pp. 170–177, 2020.

[30] Z. Zhao, W. Chen, X. Wu, P. C. Chen, and J. Liu, "Lstm network: a deep learning approach for short-term traffic forecast," *IET Intelligent Transport Systems*, vol. 11, no. 2, pp. 68–75, 2017.

[31] Y. Li, S. Carabelli, E. Fadda, D. Manerba, R. Tadei, and O. Terzo, "Machine learning and optimization for production rescheduling in industry 4.0," *The International Journal of Advanced Manufacturing Technology*, pp. 1–19, 2020.

[32] X. Zhou and H. S. Mahmassani, "A structural state space model for real-time traffic origin–destination demand estimation and prediction in a day-to-day learning framework," *Transportation Research Part B: Methodological*, vol. 41, no. 8, pp. 823–840, 2007.

[33] L. Brotcorne, G. Perboli, M. Rosano, and Q. Wei, "A Managerial Analysis of Urban Parcel Delivery: A Lean Business Approach," *Sustainability*, vol. 11, no. 12, p. 3439, 2019.