UNIVERSIDADE FEDERAL DE VIÇOSA

A Machine Learning Approach for Predicting Pallet Demands in Ceramic Tile Production

Matheus Aguilar de Oliveira
Magister Scientiae

VIÇOSA - MINAS GERAIS
2025

MATHEUS AGUILAR DE OLIVEIRA

A Machine Learning Approach for Predicting Pallet Demands in Ceramic Tile Production

Dissertation submitted to the Computer Science Graduate Program of the Universidade Federal de Viçosa in partial fulfillment of the requirements for the degree of *Magister Scientiae*.

Adviser: Andre Gustavo dos Santos

Co-adviser: Manuel Iori

VIÇOSA - MINAS GERAIS
2025

Ficha catalográfica elaborada pela Biblioteca Central da Universidade Federal de Viçosa - Campus Viçosa

T

O48m
2025

Oliveira, Matheus Aguilar de, 2000-
A machine learning approach for predicting pallet demands in ceramic tile production / Matheus Aguilar de Oliveira. – Viçosa, MG, 2025.
1 dissertação eletrônica (57 f.): il. (algumas color.).

Texto em inglês.
Orientador: André Gustavo dos Santos.
Dissertação (mestrado) - Universidade Federal de Viçosa, Departamento de Informática, 2025.
Referências bibliográficas: f. 53-57.
DOI: https://doi.org/10.47328/ufvbbt.2025.817
Modo de acesso: World Wide Web.

1. Aprendizado do computador. 2. Paletes (Transporte, armazenagem, etc.). 3. Manuseio e transporte de materiais. 4. Cerâmica - Indústria. 5. Logística. 6. Heurística. I. Santos, André Gustavo dos, 1974-. II. Universidade Federal de Viçosa. Departamento de Informática. Programa de Pós-Graduação em Ciência da Computação. III. Título.

CDD 22. ed. 006.31

Bibliotecário(a) responsável: Bruna Silva CRB-6/2552

MATHEUS AGUILAR DE OLIVEIRA

A Machine Learning Approach for Predicting Pallet Demands in Ceramic Tile Production

Dissertation submitted to the Computer Science Graduate Program of the Universidade Federal de Viçosa in partial fulfillment of the requirements for the degree of *Magister Scientiae*.

APPROVED: September 15, 2025.

Assent:

Matheus Aguilar de Oliveira
Author

Andre Gustavo dos Santos
Adviser

To my family, friends, colleagues, professors, and everyone who helped me on this journey. I thank you immensely and dedicate this dissertation to you.

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to the individuals and institutions whose contributions were essential to the successful completion of this work.

To my advisor, Dr. André Santos, for his unwavering support and guidance throughout this study. His knowledge, dedication, and insightful advice were fundamental to the development and execution of this work.

To my co-advisor, Manuel Iori, for his guidance, encouragement, and for facilitating the initial connections with the company that made this study possible.

To my colleagues Marco and Thiago, for their invaluable collaboration and partnership during this journey.

To the professors of UFV's Departamento de Informática, for their guidance and inspiration.

To the case study company, for generously providing the data and information necessary for this research.

To my parents, Alencar and Alzira, my brother Lucas, and my entire family, for their love, encouragement, and constant support.

To Jéssica, for the love, companionship, and encouragement in every step of the way.

To my friends, especially Lucas, Gabriel, and Natan, for their companionship and support.

To the Universidade Federal de Viçosa, for granting me the opportunity to pursue this Master's degree.

This work has been sponsored by the following Brazilian research agencies: Coordination for the Improvement of Higher Education Personnel (CAPES; Financing code 001), Minas Gerais State Foundation for Research

Aid (FAPEMIG) and National Council of Scientific and Technological Development (CNPq).

Finally, I extend my heartfelt thanks to all those who, in one way or another, contributed to this work, whether directly or indirectly.

"If I have seen further it is by standing on the shoulders of Giants."

(Isaac Newton)

ABSTRACT

OLIVEIRA, Matheus Aguilar de, M.Sc., Universidade Federal de Viçosa, September, 2025. **A Machine Learning Approach for Predicting Pallet Demands in Ceramic Tile Production**. Adviser: Andre Gustavo dos Santos. Co-adviser: Manuel Iori.

This dissertation develops and evaluates a predictive approach for a variant of the Distributor's Pallet Loading Problem (DPLP) applied in the context of an Italian ceramic tile manufacturer. In this scenario, orders from clients are composed by boxes of different sizes and weights that must be loaded onto pallets while meeting constraints on weight, volume and stability, as well as operational uncertainties in the warehouse. The objective is to estimate quickly and accurately the total number of pallets, and in an extended setting, the number of pallets of each type, required to fulfill an order. The proposed solution is a hybrid method that combines machine learning (ML) with heuristics. Historical data of the company is used to extract relevant features, which are then enriched with heuristic bounds to improve model accuracy. Three ML models, namely XGBoost, LightGBM, and Random Forest were trained and tuned using company's data. Experiments on thousands of past orders and a separate set of recent, unseen orders show that the hybrid approach consistently outperforms PackVol (the company's current software), achieving mean squared and absolute errors around 5.3 and 2.3 times smaller, respectively, while producing predictions a much faster time. The approach is also extended to a multi-output regression setting to predict pallet quantities by type, maintaining high accuracy and efficiency even under this more complex objective, with squared and absolute errors around 3.4 and 1.8 times smaller, respectivelly. The results demonstrate that combining ML with heuristic features yields a practical, scalable, and accurate predictive tool for the company's scenario, with potential applicability to other packing, loading, and logistics problems in this industry and many others.

Keywords: distributor's pallet loading problem; machine learning; heuristics ; tile production; real-world instances

RESUMO

OLIVEIRA, Matheus Aguilar de, M.Sc., Universidade Federal de Viçosa, setembro de 2025. **Uma Abordagem de Aprendizado de Máquina para Prever a Demanda de Paletes na Produção de Cerâmica**. Orientador: Andre Gustavo dos Santos. Coorientador: Manuel Iori.

Esta dissertação desenvolve e avalia uma abordagem preditiva para uma variante do *Distributor's Pallet Loading Problem* (DPLP) aplicada no contexto de uma fabricante italiana de cerâmicas e azulejos. Nesse cenário, pedidos de clientes são compostos por caixas de diferentes tamanhos e pesos que devem ser carregadas em paletes, atendendo a restrições de peso, volume e estabilidade, bem como a incertezas operacionais no armazém. O objetivo é estimar, de forma rápida e precisa, o número total de paletes e, em uma extensão do problema, o número de paletes de cada tipo necessários para atender a um pedido. A solução proposta é um método híbrido que combina *machine learning* (ML) com heurísticas. Dados históricos da empresa são utilizados para extrair características relevantes, que são posteriormente enriquecidas com limites heurísticos para melhorar a precisão dos modelos. Três modelos de ML, nomeadamente XGBoost, LightGBM e Random Forest, foram treinados e ajustados usando dados da empresa. Experimentos mostraram que a abordagem híbrida supera consistentemente o software atualmente utilizado pela empresa, chamado PackVol, alcançando erros quadráticos médios e absolutos cerca de 5,3 e 2,3 vezes menores, respectivamente, além de produzir previsões em um tempo muito mais rápido. A abordagem também foi estendida para um cenário de regressão de múltiplos valores, a fim de prever quantidades de paletes por tipo, mantendo alta precisão e eficiência mesmo nesse objetivo mais complexo, com erros quadrados e absolutos cerca de 3,4 e 1,8 vezes menores, respectivamente. Os resultados demonstram que a combinação de ML com métodos heurísticos usados para alimentar o modelo com estimativas, limites inferiores e limites superiores gera uma ferramenta preditiva prática, escalável e precisa para o cenário da empresa, com potencial de aplicação em outros problemas de empacotamento, carregamento e logística nessa indústria e em diversas outras.

Palavras-chave: distributor's pallet loading problem; machine learning; heurísticas; produção de cerâmicas; instâncias reais

LIST OF ILLUSTRATIONS

LIST OF TABLES

SUMMARY

1 INTRODUCTION

Packing problems are some of the most commonly studied problems in the area of operations research and optimization. These problems usually consist in packing a set of smaller objects (items) into one or more larger objects (containers) satisfying some required constraints (IORI et al., 2021b; SILVEIRA, 2022). Such problems possess many theoretical (CRAINIC et al., 2012) as well as practical applications (ALONSO et al., 2016) and are recurrent in important industries like logistics (ZUO; LIU; CHAN, 2022) and manufacturing (AO et al., 2023). Therefore, solutions for these problems are desirable due to their impact in such core industries. However, packing problems have a strong tendency to be NP-Hard, so an efficient and exact solution for them is not currently known. This characteristic encourages the development of approximate solutions like metaheuristic and machine learning approaches for packing problems.

This project addresses a packing problem called Distributor's Pallet Loading Problem (DPLP) (BISCHOFF; JANETZ; RATCLIFF, 1995), which consists in packing a set of items into layers and then stacking such layers onto a pallet. It is a complex problem with many constraints such as the pallet's weight limit, volume and stability. The objective of this problem is to pack all the items into pallets using the smallest number of pallets as possible. The DPLP is present in many situations that require the packing of items. Some examples of these are food packages, soft drink bottles (as shown in Figure 1), cans and boxes for shipping goods. Therefore, the DPLP has many practical applications and a solution that minimizes the number of pallets can save costs of the companies that build them (SILVEIRA, 2022).

The case study of this work is an Italian company that produces ceramic and porcelain tiles. This company deals with the DPLP in its production process, since it packs boxes of tiles into pallets (as illustrated in Figure 2). The company has an e-commerce system that receives orders from clients, that are composed of different types of ceramic tiles. The process of building the clients' orders begins when the operators responsible for picking and loading request the necessary tiles through the warehouse's automated system, which then sends them to be placed onto shipping pallets. Some tiles may be already stored in the warehouse and can be dispatched quickly, while others are not immediately available because they must first arrive from other facilities, sometimes even from abroad. Additionally, even for tiles that are in stock, the internal

Figure 1 – Real-world application of the DPLP. Source: SILVEIRA, 2022

handling performed by automated machines along with human operators introduces further variability, since the system does not release them in a predictable sequence or time frame. As a result, operators cannot know in advance when each tile will reach the loading area. Since the space in the loading area is very limited, they cannot afford to wait idly for missing tiles and must start loading available tiles as soon as they arrive. These supply and handling uncertainties, combined with the spatial constraints of the loading area, make it very difficult to design a pallet-building procedure that is both systematic and practical for operators to follow in real time. Therefore, a direct solution to the DPLP is not viable to be implemented in the company's scenario.

However, another characteristic of the company's shipping process is that the clients are responsible to pick the pallets directly at the warehouse after they are ready. Since they do it by using vehicles of their own or rented ones, having a precise or close enough estimation of the number of pallets used can help clients organize and reduce their shipping costs. Therefore, a solution that can estimate the number of pallets that an order will require is desirable in the context of this company, and this is much simpler and viable to implement than a model that defines a way to pack the boxes and accounts for all dynamic factors of the warehouse.

A machine learning (ML) approach was selected to build this model to predict the number of pallets that an order will require. This ML approach learns from historical data from the company how to infer the number of pallets and do it for new orders. The reason to choose ML for this task was due to some advantages that it has against some

Figure 2 – Company's pallets of ceramic tiles boxes. Source: Case study company

more traditional methods (e.g. exact models and heuristics) in the packing context:

- ML models are able to do inferences much faster than an exact model or a metaheuristic, something that aligns with the company's need to make provide quick predictions to the clients in their e-commerce system;
- ML models have a more data-driven characteristic than traditional methods (WU et al., 2023), something that would allow them to adapt to gradual changes in the production process of the company;
- ML models are less reliant on specific knowledge of the domain than traditional methods (WU et al., 2023), something that helps in the company's context since their scenario is very complex and dynamic;

Besides that, some heuristics were built to enrich the ML models, in a way to improve the solution's accuracy.

1.1 Motivation

The initial motivation for this dissertation is the theoretical and practical importance of the Distributor's Pallet Loading Problem, since it appears and can impact many relevant industries like beverages (AKKAYA et al., 2018), food (MORABITO; MORALES; WIDMER, 2000) and logistics (ALONSO et al., 2016; IORI et al., 2020; IORI et al., 2021b). Furthermore, reviews such as that by WU et al. (2023) highlight the potential of machine learning (ML) and hybrid approaches that combine ML with heuristics for addressing packing problems like the DPLP. This potential served as a key motivation for the present work, which explores the application of ML techniques in the context of the DPLP.

Previous studies have explored the use of optimization techniques for packing problems in the ceramic industry. For example, a matheuristic approach was applied to the Two-Dimensional Irregular Bin Packing Problem in a Spanish ceramic tile company (MARTINEZ-SYKORA et al., 2017), while heuristic methods were developed to address ceramic bowl placement in a Thai company (MAUNGMEESRI; THWE, 2024). Despite these contributions, optimization-based solutions in the ceramics domain remain limited. In this work, we extend the research by investigating the potential of machine learning for tackling such problems.

Additionally, another significant motivation to build this approach over a ceramic company scenario was the weight of this industry over the global market that has grown so much in the last decade. According to the ACIMAC Research Department (2023), world production has reached almost 17 billion square meters, with an estimated value of 248.89 billion USD in 2023 and is projected to reach 359.35 billion USD by 2030, growing at a compound annual growth rate of 5.6% from 2024 to 2030 (Grand View Research, 2023). Moreover, Italy (the country of the case study company) is one the world's greatest ceramic exporters and this sector has great influence over the country's economy with an estimated value of 7.5 billion euros in 2024 (Ceramic World Web, 2025). Besides that, a model built from the ceramic industry has the possibility to be adapted to other fields that deal with the DPLP.

1.2 Hypotheses

The following hypotheses have been developed based on the problem and objectives of the project:

Hypothesis 1: An ML model trained on historical company orders can predict the number of pallets required for new orders based on their general characteristics (e.g., number of tiles, weight, and volume);

Hypothesis 2: Heuristics can be used to generate additional features to enrich the ML models and improve their performance;

Hypothesis 3: An ML solution is able to provide estimations much faster than the current company approaches and with better or close enough accuracy;

These hypotheses are going to be tested and validated over the following sections of this dissertation.

1.3 Objectives

The main objective of this dissertation is to build a viable machine learning approach to estimate the pallet demands in the Distributor's Pallet Loading Problem applied over a real world scenario. To achieve this primary objective, a set of specific objectives must completed:

Objective 1: Extract an useful set of features from the case study company historical data;

Objective 2: Enumerate, test and choose the best machine learning models to build the solution;

Objective 3: Build heuristics to generate additional features over the dataset and enrich the ML models;

Objective 4: Compare the ML solution with the current company one to assess the potential to improve the prediction and consequently the transportation process;

1.4 Contributions

The main contributions of this dissertation can be summarized as:

- A model capable of estimating pallet demands of the DPLP in a real-world scenario.
- A feature analysis to understand which characteristics of the order are the most relevant for the prediction.
- Heuristics to compute bounds of the pallet demands that are able to improve the ML models.
- An investigation of which ML models have the best accuracy in this task.
- A faster and more accurate solution for the company than their current methods to predict pallet demands.

1.5 Dissertation structure

The structure of this dissertation is as follows. Chapter 2 describes the case study, including the company dataset and the existing solutions. Chapter 3 presents the development of our initial approach to predicting the total number of pallets in an order. This process involved selecting relevant company features, choosing the most suitable ML model, implementing heuristics to enhance the models, and comparing the new solution with the company's current methods. The resulting solution, based on the XGBoost model, was published as a paper in the International Conference on Optimization and Learning (OLA2025). Chapter 4 extends the work from Chapter 3 by evaluating additional ML models beyond XGBoost to assess the generality of the proposed strategy. This study was published as an extended abstract in the LVII Brazilian Symposium on Operations Research (SBPO 2025). Chapter 5 introduces another extension of the problem addressed in Chapter 3, where the model performs multi-value predictions to estimate not only the total number of pallets required but also their respective types. The results of this multi-output prediction were published in the International Conference on Optimization and Decision Science (ODS2025). Finally, Chapter 6 provides concluding remarks and discusses potential directions for future research.

2 CASE STUDY DATASET AND CURRENT SOLUTIONS

This chapter presents an overview of the dataset provided by the company and the current methods used for predicting pallet requirements. It describes the structure and content of the data, including details on the tiles, pallets, and historical orders, highlighting the variety and characteristics relevant to packing and shipping operations. The chapter also discusses solutions currently used at the company. In summary, it presents the foundation to understand the approaches presented in the following chapters.

2.1 Dataset

The company provided a dataset containing various operational attributes, including characteristics of tiles and pallets as well as some orders received in previous years. The main training dataset consists on previous orders received by the company in the years of 2022 and 2023, totaling approximately 207,000 entries covering a wide range of tile types and order sizes.

A detailed description of the provided data is presented in the following subsections.

2.1.1 Tiles

For each ceramic tile produced by the company, the following information was provided: the tile's dimensions (width, height, and thickness), its weight, and a set of company-defined limits. These limits specify how many tiles of a given type can be packed into a single box and how many boxes can be stacked on a single pallet. In addition, the dimensions of the boxes used for packaging were also included in the dataset.

Tiles are organized in 22 types based on their width and height ranging from small ones measuring 20x20 centimeters to big ones measuring 120x120 centimeters. The most common type in orders is the 60x60 size.

2.1.2 Pallets

For each pallet used by the company, the following information was provided: the pallet type (an id represented by a letter, e.g. F, X, etc.), its dimensions (width, height, and depth), the maximum supported weight, and the maximum allowable height for stacked items.

Pallet's maximum weight range from 500 to 1300 kilograms and the maximum allowed height is 1 meter for all types.

2.1.3 Orders

For each order received by the company, the dataset includes the types and quantities of tiles requested (e.g., 10 units of 30×30 tiles and 20 units of 60×60 tiles), as well as the number and types of pallets used for shipment (e.g., two pallets of type F and one pallet of type B). The specific allocation of tiles within each pallet was also provided.

Orders can include multiple tile types, ranging from a single type (homogeneous orders) to up to 14 different types. On average, each order contains 1.35 tile types, or 2.71 when homogeneous orders are excluded.

2.2 Current solutions

Currently, the company relies on estimates made by human operators to predict the number and types of pallets required for each order. In addition to these manual estimations, the company also uses an heuristic software called PackVol for making such predictions.

2.2.1 PackVol

PackVol is a specialized optimization software designed to maximize space utilization in containers, trucks, and other transport vehicles (PACKVOL, 2024). The company configured the software to match its operational constraints and solve the DPLP heuristically. However, there are some issues with this approach. First, PackVol does not account for dynamic warehouse factors, so the proposed solution may not be feasible to implement. Second, PackVol focuses primarily on minimizing the number of pallets rather than providing estimations, which can lead to it underestimating the number of pallets required for an order. This underestimation poses a more significant operational risk to the company than occasional overestimation, as it can disrupt the client's shipping plan.

3 A HYBRID APPROACH FOR PALLET LOADING IN CERAMIC TILE INDUSTRY

This chapter presents a hybrid approach in which a machine learning model is trained over a company dataset comprising two years of customer orders, with the aim of predicting the number of pallets required by an order. The accuracy of the machine learning technique is largely improved by including additional features, such as lower and upper bounds, in the dataset, obtained using quick optimization algorithms. The resulting hybrid algorithm has been compared with the model-based software currently used at the company, consistently providing better-quality results in shorter computing times. [1]

3.1 Introduction

The global ceramic tile market has grown significantly over the past decade, with worldwide production reaching 16.8 billion square meters (ACIMAC Research Department, 2023). Italy is a leading exporter, generating €7.2 billion in revenue, highlighting the sector's importance to the national economy. Efficient logistics optimization in the ceramic tile industry, including loading of tiles on pallets and transportation of pallets, is crucial for cost reduction and company competitiveness. Loading operations involve transferring specified quantities of ceramic tile boxes, either of the same size or different sizes, from production pallets to shipping pallets to fulfill customers' orders. When loading pallets, the company needs to carefully consider weight and volume constraints and keep pallet stability and balance to ensure the safe transport of tiles while using the minimum number of shipping pallets.

In this paper, we handle a variant of the distributor's pallet loading problem (DPLP) (MUNGWATTANA; PIYACHAYAWAT; JANSSENS, 2023) that emerges from an international ceramic tile company headquartered in Italy. The company has an e-commerce platform to receive customers' orders. As the company does not provide the transportation of the shipping pallets, customers need to collect their pallets by using their vehicles or third-party carriers. Consequently, the company should provide the number of shipping pallets associated with each order or at least predict this number

[1] The results of this chapter appear in: Taccini, M.; Oliveira, M.A.; Santos, A.G.; Queiroz, T.A.; Iori, M. (2025). A Hybrid Approach for Pallet Loading in Ceramic Tile Industry. In: Proceedings of the International Conference on Optimization and Learning (OLA2025).

accurately, especially for customers who need to contract a third-party carrier. This prediction is complex, even if using specialized software to address pallet loading problems (SILVA; OLIVEIRA; WÄSCHER, 2016), because the company faces specific constraints in its operational daily routine, necessitating extensive customization of such software.

In this paper, we propose a hybrid approach to predict the number of shipping pallets required for a given customer order. The approach combines machine learning (ML) techniques to gain information from a large dataset of previously shipped pallets with techniques from the optimization field to enrich the dataset by quickly providing additional features such as lower and upper bounds on the number of pallets. The hybrid approach is then compared with solutions currently in use at the company, including manual operators' decisions and outputs from a given specialized software (PACKVOL, 2024). The computational experiments show that the proposed approach is a valuable tool for the company.

The remainder of the paper is organized as follows. Section 5.2 contains a brief review of the literature related to the DPLP. Section 5.3 describes the company problem with its constraints and requirements. Section 5.4 presents our proposed hybrid approach. Section 3.5 contains the numerical experiments on real-world instances. Finally, Section 5.6 concludes the paper and provides several directions for future research.

3.2 Literature Review

The literature on cutting and packing problems is extensive. According to (SWEENEY; PATERNOSTER, 1992), the first scientific papers on these problems date back to the 1940s. The most recent surveys concern problems in one-, two-, and three-dimensions and contain detailed discussions on solution methods based on exact and heuristic approaches (see, e.g., (IORI et al., 2021a) and (YAGIURA; IMAHORI; HU, 2025)). An essential characteristic of these problems is given by the number of practical constraints they may include, such as complete shipment, conflicting items, cargo stability, load-bearing, multi-drop, and load-balancing (NASCIMENTO; QUEIROZ; JUNQUEIRA, 2021). Concerning pallet loading problems, (SILVA; OLIVEIRA; WÄSCHER, 2016) surveyed the manufacturer's variant, in which all items are identical, and the objective is to pack the maximum number of items onto a single pallet. Besides commenting on the different data sets and solution methods, the authors emphasized that the computational complexity of this problem is still an open question. The DPLP, on the other hand, assumes that items are not necessarily identical and is an NP-hard problem. Its objective is to find the minimum number of pallets necessary to load all items. The

DPLP is also referenced as the multi-pallet loading problem (TERNO et al., 2000) and the pallet building problem (CALZAVARA et al., 2021).

One of the first authors to handle the DPLP was (HODGSON, 1982), considering the transport of palletized cargo by the U.S. Air Force. The author proposed a dynamic programming-based algorithm and tested it on instances having 30 boxes. In (TERNO et al., 2000), the DPLP was solved using an algorithm based on branch-and-bound. The authors considered practical constraints such as weight limits, load-bearing, stability, and grouping of items. In (BIRGIN; LOBATO; MORABITO, 2012), the problem had no limit on the number of boxes to load, and boxes could have multiple orientations. The authors developed a recursive partitioning algorithm, which integrated a recursive five-block heuristic and an L-approach. The authors could not find an instance for which this algorithm failed in finding an optimal solution (such an instance was later found in (QUEIROZ; MIYAZAWA; WAKABAYASHI, 2015)). In (CALZAVARA et al., 2021), the DPLP emerged from a robotized task of loading boxes into layers, which were then stacked while meeting load-bearing requirements. Constraints like contiguity, visibility, and multiple orientations were considered. The authors proposed a mathematical model and a reactive greedy randomized adaptive search procedure for solving realistic, large-sized instances. A new DPLP variant in which pallets may have different sizes was handled in (MUNGWATTANA; PIYACHAYAWAT; JANSSENS, 2023). The problem emerged from a lamp and lighting manufacturing company that needed to load many carton boxes onto pallets of multiple sizes. Constraints such as load-bearing and multiple orientations were taken into consideration.

Using artificial intelligence techniques to handle combinatorial optimization problems is a promising area of research. The survey in (BENGIO; LODI; PROUVOST, 2021) discussed recent advances in combining ML techniques with optimization methods to efficiently handle hard combinatorial optimization problems. The authors also commented on the benefits (e.g., speed and generalization) and challenges (e.g., training and accuracy) of using ML purely. Concerning cutting and packing problems, an increasing number of publications combine ML or artificial intelligence techniques with exact or heuristic algorithms to handle, e.g., bin packing (FANG et al., 2023) and container loading problems (QUE; YANG; ZHANG, 2023).

3.3 Problem Description

The problem under investigation is a variant of the DPLP with specific operational and product-related constraints. After being produced, the ceramic tiles are first loaded on production pallets, which are only used for inbound logistics. An automated system then delivers the production pallets to the loading area. The company operators load the products onto shipping pallets using hydraulic machines. The automated system is also responsible for other logistics activities inside the warehouse. As a result, products are not available simultaneously, as their arrival sequence depends on the system's decisions and the warehouse's current state. For instance, machinery downtime may alter the order in which production pallets are sent to the loading areas. Moreover, the loading area cannot have more than four production pallets at a time. Due to this space limitation, operators need to load the products onto the pallets as soon as possible. As other products arrive, operators continue loading until the shipping pallets of that order are completed before moving to the next order. When loading a shipping pallet, operators must also satisfy load requirements such as vertical stability, load-bearing, and restricted box orientations.

Our DPLP can be formally described as follows. A customer order consists of a set B of three-dimensional boxes, each containing a given number of tiles. Each box belongs to a type $t \in T = \{1, 2, ..., m\}$, corresponding to a specific tile size. A box of type t has width $w_t \in \mathbb{Z}_+^*$, depth $l_t \in \mathbb{Z}_+^*$, height $h_t \in \mathbb{Z}_+^*$, and weight $p_t \in \mathbb{Z}_+^*$. Let $d_t \in \mathbb{N}$ denote the number of boxes of type t required by the order. Let P represent the set of identical shipping pallets, each with a two-dimensional loading surface of width $W \in \mathbb{Z}_+^*$ and depth $L \in \mathbb{Z}_+^*$. The DPLP's objective is to load all boxes requested in the order by using the minimum number of pallets. Constraints on the maximum pallet height $H \in \mathbb{Z}_+^*$ and weight $W_{\max} \in \mathbb{Z}_+^*$ need to be satisfied.

Due to these complex company constraints, obtaining a solution for this DPLP variant is not easy, especially for large-sized and heterogeneous orders (i.e., composed of boxes of different sizes). Even if a DPLP solution satisfying all constraints is available, it could suffer from warehouse uncertainties, possibly making it impractical for the company operators. Besides that, the company should give the number of pallets associated with each customer's order at the moment of purchase on its e-commerce platform. Therefore, the company is currently interested in a fast and accurate tool capable of predicting the number of shipping pallets, rather than solving a dedicated but time-consuming optimization problem.

3.4 Hybrid Model

We propose a hybrid approach to the DPLP under investigation that considers optimization algorithms to enhance the predictive performance of an ML model. The company provided data on two years of sales and the corresponding number of pallets generated for each order. This data includes the maximum number of boxes per shipping pallet for each box size t, denoted by $F_t^{\max}$, and the maximum weight ($W_{\max}$) and volume (WLH) a shipping pallet can support.

An initial training dataset is created from the company-provided data. We start calculating the frequency $F_t = d_t$ (i.e., the number of boxes of type t), the weight $W_t = d_t p_t$, and the volume $V_t = w_t l_t h_t d_t$ for each type $t \in T$. The dataset is labeled with the number of shipping pallets Z_j generated by each order $j \in \{1, ..., n\}$. Since training on the initial dataset can take a long time, all intermediate training steps (Sections 3.4.1 - 3.4.3) are conducted using a reduced dataset. The complete dataset is only used on the final computational experiments (Section 3.5). The reduced dataset is created by randomly selecting a subset of orders from the initial dataset, resulting in a training set with about $\frac{1}{8}$ of the total number of orders.

3.4.1 Machine Learning Model Selection

We developed a pipeline to train and test four different machine learning models. These models, which were chosen based on the literature review and our experience, consist of: XGBoost (CHEN; GUESTRIN, 2016), LightGBM (LGBM) (SHI et al., 2024), Random Forest, and Support Vector Machine (SVM) from scikit-learn (PEDREGOSA et al., 2011). Each model passes through a hyperparameter tuning phase on the Mean Squared Error (MSE) and Mean Absolute Error (MAE) metrics, obtained by performing a standard grid search. Table 1 reports the hyperparameters and their possible values, besides the best values found along with MSE and MAE metrics. The possible values were defined using a trial-and-error approach and following the indications in each model documentation and the literature (see, e.g., (BANERJEE, 2020)).

Table 1 – Grid search results with the hyperparameter values and errors.

Model	Hyperparameter	Values	Best Value	MSE	MAE
XGBoost	learning_rate max_depth gamma	0.1, 0.3 4, 5, 6 0, 0.25, 1	0.3 4 0.25	4.78	1.17
LGBM	learning_rate num_leaves	0.005, 0.01 6, 8, 12, 16	0.01 16	16.49	2.58
Random Forest	max_features max_depth max_leaf_nodes	sqrt, log2, None 3, 6, 9 3, 6, 9	None 9 9	13.61	2.26
SVM	C gamma	0.1, 1, 10, 100, 1000 1, 0.1, 0.01, 0.001, 0.0001	100 0.0001	35.89	3.45

The results in Table 1 show that XGBoost outperforms the other models regarding the MSE and MAE values. Consequently, we decided to use only XGBoost in the following steps. We also decided to perform a deeper tuning of this model's hyperparameters. Table 2 contains the outcome of this additional tuning test, which leads to better MSE and MAE values for the XGBoost model.

Table 2 – XGBoost tuning with the best hyperparameter values and errors.

Hyperparameter	Values	Best Value	MSE	MAE
learning_rate	0.01, 0.1, 0.2, 0.3	0.3	1.11	0.43
gamma	0, 0.1, 0.5	0.1		
max_depth	3, 6, 10	6		
min_child_weight	0, 1, 3, 5	0		
subsample	0.5, 1	0.5		
colsample_bytree	0.5, 1	1		

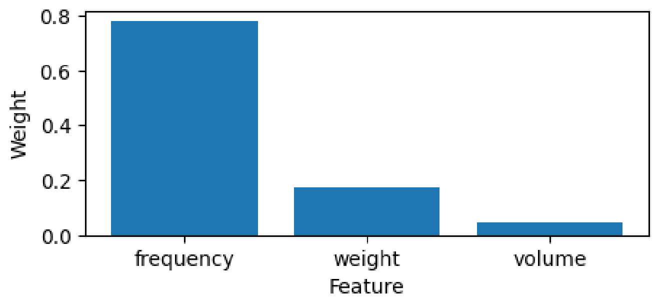

Figure 3 – Importance of the feature weights in the XGBoost model.

3.4.2 Features Optimization

After successfully fine-tuning the XGBoost model, we investigated the importance of each feature in the reduced dataset. This analysis reveals that columns W_t and V_t, for $t \in T$, are marginally contributing to the model's predictions, as shown in Figure 3. Furthermore, tile sizes that appear less frequently in the dataset exhibit higher error rates, as shown in Figure 4.

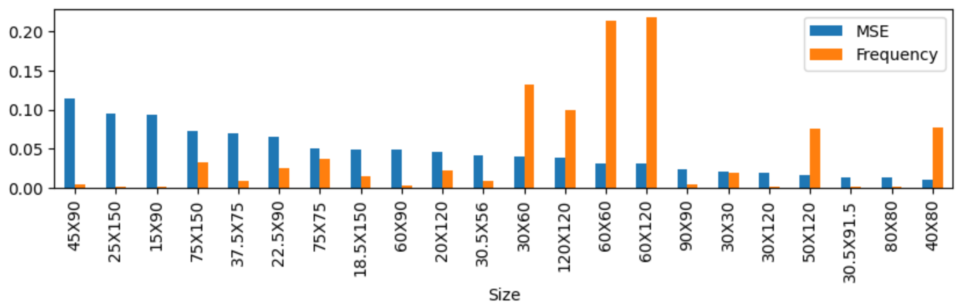

Figure 4 – MSE and frequency values (F_t), normalized and sorted by decreasing MSE.

These results allow the removal of less informative features, such as the size-grouped weight W_t and volume V_t. At the same time, we introduce more general order-aggregated features, as follows. Let O be the set of orders we need to process, and let $o \in O$ be a given order. We determine the total number of boxes $F_o = \sum_{t=1}^{m} F_t$, the total weight $W_o = \sum_{t=1}^{m} W_t$, the total volume $V_o = \sum_{t=1}^{m} V_t$, the average box weight $W_{avg} = \frac{W_o}{F_o}$, and the average box volume $V_{avg} = \frac{V_o}{F_o}$. Figure 5 shows how these new features impacted the predictive capability of the XGBoost model.

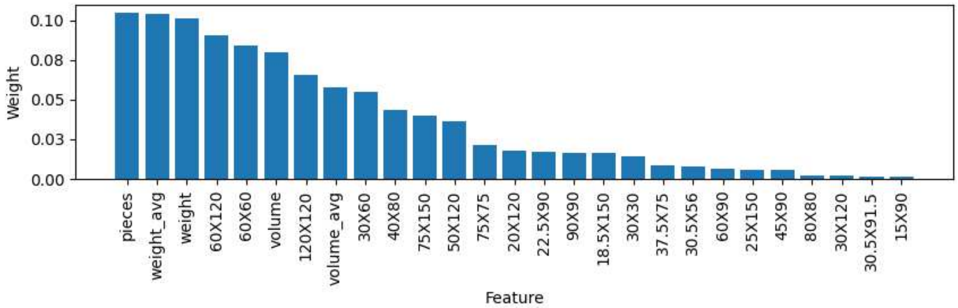

Figure 5 – Importance of the new order-aggregated features in the XGBoost model.

In the set O of orders, we observe that predicting the number of shipping pallets for homogeneous orders, defined O_h, which consist of only one type of box, and O_{full} ones, which require only production pallets regardless of the types, is straightforward for the model. In this way, we divided results into two groups: one containing all available orders O and another excluding these simpler orders $O_{mix} = O - (O_h \cup O_{full})$. With this strategy, we evaluated the model in both scenarios, preventing simple orders to improve the metrics and resulting in premature conclusions. Table 3 summarizes the model's performance before and after including the new features.

Table 3 – MSE and MAE values for each group of orders and the dataset structures S_1 (original) and S_2 (enriched with additional features).

	MSE		MAE	
Dataset structure	O	O_{mix}	O	O_{mix}
S_1	1.01	1.38	0.42	0.53
S_2	0.36	0.47	0.24	0.31

3.4.3 Incorporation of Optimization-based Features

Aiming to reduce the errors further, three additional features (i.e., HS_1, HS_2, and HS_3) computed with quick algorithms are introduced into the training dataset. To ensure computational efficiency, the proposed algorithms have the worst-case time complexity of at most $O(F_o^2)$, where F_o is the total number of boxes in the customer's order. The proposed algorithms are as follows:

- HS_1: calculated as $HS_1 = \max\left(\lceil\frac{\sum_{t=1}^m W_t}{W_{\max}}\rceil, \lceil\frac{\sum_{t=1}^m V_t}{WLH}\rceil\right)$. This can be interpreted as the value of the linear relaxation lower bound (LB);
- HS_2: calculated as $HS_2 = \sum_{t=1}^m \lceil\frac{F_t}{F_t^{\max}}\rceil$, where $F_t^{\max}$ is the maximum number of boxes of type t that can be loaded onto a pallet. This is, instead, an upper bound (UB) for the DPLP instance, computed by considering that all pallets are built with only homogeneous tiles (and this can easily be done by satisfying vertical stability and load bearing constraints);

- HS_3: described in Algorithms 1 and 2. This algorithm loads pallets by stacking layers of same-sized boxes, starting with those having the largest area until the pallet limits are reached. Once full, it begins a new pallet and continues until all boxes are packed. This may be seen as a quick heuristic solution approach, but, in fact, it does not consider vertical stability and load bearing constraints, and hence it only provides an approximation function of the minimum number of pallets.

Algorithm 1 HS3 - constructive heuristic to give the number of shipping pallets

$pallets \leftarrow 1$ ▷ Pallet estimation
for $t \in$ SORTED(T) **do** ▷ Iterate over box types sorted by decreasing base area
    while $d_t > 0$ **do** ▷ Boxes of type t left to pack
        if $w - p_t < 0$ **OR** $h - h_t < 0$ **OR** no space to pack one box **then**
            $x, y, w, h \leftarrow W, L, W_{\max}, H$ ▷ Restart using an empty pallet
            $pallets \leftarrow pallets + 1$
        end if
        $b \leftarrow \lfloor \frac{w}{p_t} \rfloor$ ▷ How many boxes can be packed based on weight
        $q, d_x, d_y \leftarrow$ TRYPACK(b, w_t, l_t, x, y) ▷ Try to pack directly
        $q_2, d_{x2}, d_{y2} \leftarrow$ TRYPACK(b, l_t, w_t, x, y) ▷ Try rotating to pack
        if $q < q_2$ **OR** $(q = q_2$ **AND** $p_x p_y < p_{x2} p_{y2})$ **then**
            $q, d_x, d_y \leftarrow q_2, d_{x2}, d_{y2}$ ▷ Choose to pack more boxes, then more area
        end if
        $d_t, w, h \leftarrow d_t - q, w - p_t q, h - h_t$ ▷ Pack boxes, update weight and height
        $x, y \leftarrow d_x, d_y$ ▷ Update dimensions for next layer
    end while
end for
return $pallets$

Algorithm 2 TRYPACK

Input: b, b_x, b_y, l_x, l_y ▷ Boxes to pack, box and layer dimensions
$X \leftarrow \min(b, \lfloor \frac{l_x}{b_x} \rfloor)$ ▷ Maximum number of boxes in x-dimension
$Y \leftarrow \min(b, \lfloor \frac{l_y}{b_y} \rfloor)$ ▷ Maximum number of boxes in y-dimension
$q, d_x, d_y \leftarrow 0, 0, 0$ ▷ Best quantity and new layer dimensions
for $x \in \{1 \dots X\}$ **do** ▷ Test all number of boxes in x-dimension
    $y \leftarrow \min(Y, \lceil \frac{b}{x} \rceil)$ ▷ How many rows in y-dimension will be needed
    if $q < \min(b, xy)$ **OR** $(q = \min(b, xy)$ **AND** $d_x d_y < b_x b_y xy)$ **then**
        $q, d_x, d_y \leftarrow \min(b, xy), b_x x, b_y y$
    end if
end for
return q, d_x, d_y

Table 7 describes the final structure S_3 of the training dataset, which now includes the three additional optimization-based features. Compared with Table 3, the MSE and MAE values for the dataset structure S_3 are 0.24 and 0.19 for O, and 0.32 and 0.26

for O_{mix}, respectively, representing a considerable improvement in the accuracy of the model.

Table 4 – Final training dataset structure S_3.

j	F_1	$\dots$	F_m	F_o	W_o	V_o	W_{avg}	V_{avg}	HS_1	HS_2	HS_3	Z_j
1	10	$\dots$	3	25	250	500	10	20	5	9	8	7
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$
n	7	$\dots$	0	13	195	65	15	5	2	4	3	3

3.5 Computational Experiments

All the codes were implemented in Python 3.10 and executed in a computer with an Intel Core i5-1135G7 2.40 GHz processor, 8 GB of RAM, and Linux Mint 21 Cinnamon as the operating system. We evaluate the XGBoost model considering the structure S_3 (i.e., the one enriched with the additional optimization-based features) and the complete two-year dataset with all orders. This corresponds to our hybrid approach. In Table 5, we observe the impact of the two groups O_h and O_{full} on the training phase. Consequently, we train the model using the structure S_3, first including all orders O and then keeping only mixed orders O_{mix} (which we recall is equal to $O - (O_h \cup O_{full})$).

Table 5 – MSE and MAE values for each group of orders and the dataset structure S_3.

			MSE		MAE	
Groups of orders	Orders	Training Time (min)	O	O_{mix}	O	O_{mix}
O	166391	54.79	0.22	0.73	0.11	0.33
O_{mix}	40533	26.90	0.25	0.76	0.12	0.33

As reported in Table 5, applying the model to all orders O yields slightly better results compared to using O_{mix}. However, using only O_{mix} can reduce the training time by around 51% while keeping satisfactory predictive results.

We next compare in Table 12 the proposed hybrid approach with PackVol (PACKVOL, 2024), a specialized software the company uses to solve the DPLP heuristically. PackVol was configured by the company employees to satisfy the specific operational constraints of the DPLP as closely as possible. This means that PackVol is not able to capture some constraints. For example, as mentioned in Section 5.3, ceramic tiles of a given order are not all available simultaneously due to the warehouse's current state. As a result, PackVol could underestimate the number of pallets. We also present the number of pallets obtained from the company operators for comparison purposes. This experiment considers 30 real instances, consisting of 30 orders randomly selected among the heterogeneous orders made in 2024 (which are not included in the training set). The real instances I are all heterogeneous, and are sorted by increasing number of box types m. Its maximum value is 6 since 99.32% of the orders in the training dataset contain six or fewer types. Table 12 contains the instance number and the solution value (i.e., number of shipping pallets) obtained with PackVol, the proposed hybrid approach, and the company operators. We assume the operators' solution is the basis for comparison, so the rows "Diff." contain the difference with such a solution for each instance. Positive values indicate the usage of more pallets, which could be acceptable. Negative values indicate fewer pallets in the solution and could represent additional logistic issues for customers (e.g., the need for contracting "on the flight" an extra vehicle).

Table 6 – Results for the 30 real instances and comparison among the different methods.

Inst.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
# types	2	2	2	2	2	2	3	3	3	3	3	3	4	4	4
Company	2	6	4	1	6	2	22	5	5	8	10	7	5	8	9
PackVol	1	6	4	1	6	2	22	5	3	8	10	6	4	6	8
Diff.	-1	0	0	0	0	0	0	0	-2	0	0	-1	-1	-2	-1
Hybrid	2	7	4	1	6	2	22	5	5	9	11	7	6	7	9
Diff.	0	1	0	0	0	0	0	0	0	1	1	0	1	-1	0

Inst.	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
# types	4	4	4	5	5	5	5	5	5	6	6	6	6	6	6
Company	7	15	6	26	9	28	2	13	10	22	10	22	21	10	16
PackVol	7	15	6	25	8	27	4	13	6	18	8	21	19	10	14
Diff.	0	0	0	-1	-1	-1	2	0	-4	-4	-2	-1	-2	0	-2
Hybrid	7	16	7	26	9	27	3	13	10	21	10	22	20	12	16
Diff.	0	1	1	0	0	-1	1	0	0	-1	0	0	-1	2	0

Observing the results in Table 12, PackVol achieves an MSE of 2.13 and an MAE of 0.93, whereas our hybrid approach achieves significantly better results, with 0.50 and 0.43 values of MSE and MAE, respectively. These outcomes highlight that our approach outperforms PackVol in both metrics, demonstrating its capability to predict the number of pallets following the real context of the company with its operational requirements and workflows. Another important aspect is the computation time required by the proposed approach. Although the training phase may be time-consuming, after training, it solved all instances in just 0.016 seconds, while PackVol required around one second per instance. From Table 12, we can also observe that PackVol returns solutions different from those by the company in 16 cases, returning fewer pallets for 15 cases and achieving a largest deviation of -4 pallets. This may bring issues for customers. On the other hand, the hybrid approach returns solutions that differ from the company ones for just 12 instances, returning fewer pallets for only 4 of them and having the largest deviation of just 2 pallets. These values confirm that our approach, trained on real-world data, effectively models the problem constraints and results in accurate predictions.

3.6 Concluding Remarks

We proposed a hybrid approach composed of machine learning and optimization components to predict the number of pallets required to ship an order of ceramic tiles. We evaluated four machine learning models, and after fine-tuning, we selected XGBoost. Its performance was improved by adding new features that capture general information about the company's operational constraints. We also incorporated three measures obtained by quick optimization algorithms, leading to significant reductions in both MSE and MAE values. The resulting hybrid approach obtained better results than PackVol, the specialized model-based software currently used at the company.

The computational experiments highlight that our hybrid approach, once it is trained on real-world data, is fast and well-suited for the company-specific operational constraints. It allows more accurate and practical predictions than PackVol, resulting in more solutions equal to the ones obtained by the company operators and a reduced number of solutions requiring fewer pallets. This is remarkably achieved without feeding the approach with any direct information about the pallet loading constraints (i.e., stability, load-bearing requirements, and orientation), making the approach a very flexible tool.

Future research could propose new features to the hybrid approach (e.g., solutions obtained with heuristics based on local search). Another interesting research avenue is to work on a more general problem, considering pallets of different sizes, which will lead to a more challenging objective function. We finally plan to test the approach to other related pallet loading, cutting, and packing problems, as well as scheduling problems (BORGES et al., 2024; QUEIROZ; MUNDIM, 2020).

4 COMPARISON OF PREDICTIVE MODELS FOR AN APPLIED DISTRIBUTOR'S PALLET LOADING PROBLEM

This chapter presents an extension of the work of chapter 3, that defined a hybrid approach for predicting pallets demands based on XGBoost. This chapter addresses the behavior of this approach when other machine learning models are used (LightGBM and Random Forest, in this case study). The new models were able to surpass significantly model-based software currently used at the company as well as provide a slightly better result than XGBoost. [1]

4.1 Introduction

The global ceramic production market has grown considerably in the last decade, with world production reaching almost 17 billion square meters according to (ACIMAC Research Department, 2023). Among these markets, the Italian market stands out as one of the main global ceramic exporters and, consequently, has a great influence on the country's economy. Among the different problems that ceramic industries face, some are related to the logistics and transport of their products. In Italian industries, it is common for ceramic boxes to be organized on pallets and for customers to be responsible for picking up these pallets for their locations.

The task of loading pallets consists of moving ceramic boxes, which have different dimensions and weights, from production pallets to shipping pallets for customers. Production pallets are usually homogeneous in their composition, while shipping pallets are heterogeneous, formed according to each customer's orders. While performing this task, industry operators, who are specialized personnel, need to be aware of restrictions such as pallet weight limits, vertical loading stability, allowed box orientation, aiming to minimize the number of pallets used.

The problem investigated in this work arises from the real context of a ceramic manufacturing industry located in the Emilia-Romagna region of Italy. Due to the given characteristics, the industry deals daily with a distributor's pallet loading problem (DPLP)

[1] This chapter is a full version of the following extended abstract: Oliveira, M.A.; Santos, A.G.; Queiroz, T.A.; Taccini, M.; Iori, M. (2025). Comparação de Modelos Preditivos para um Problema Aplicado de Carregamento de Paletes do Distribuidor. Submitted to the LVII Brazilian Symposium on Operations Research (SBPO 2025).

(SILVA; OLIVEIRA; WÄSCHER, 2016). Generally, it has an e-commerce system that receives customer orders. Orders can include different types of ceramics, which are produced, if not available in stock, and then loaded onto pallets by operators. Customers are responsible for picking up the pallets directly from the industry and thus transporting them to their locations using their own or rented vehicles. Therefore, the industry must provide customers, when placing an order, with the precise number of pallets or a good estimate of this number so that they, in turn, can organize the transport.

Pallet loading problems are part of cutting and packing problems, which have been extensively studied in the operations research literature. Systematic reviews have recently been published by (IORI et al., 2021a), who focused their review on exact methods for two-dimensional problems, and (SALEM; SILVA; OLIVEIRA, 2023), who considered problems with uncertainties and thus proposed a classification for them. In addition to systematic reviews, entire books dedicated to cutting and packing problems can be found, as in the case of (SCHEITHAUER, 2018) and (YAGIURA; IMAHORI; HU, 2025), both providing an introduction to the main problems and solution methods applied to them.

Due to the large number of applications, cutting and packing problems come with different restrictions to reflect real scenarios. Issues such as weight limits, fragility, vertical and horizontal stability, multiple orientations, loading complexity, grouping, priority, conflict, delivery order, among others, are commonly considered when proposing solution methods. (BORTFELDT; WÄSCHER, 2013) and, more recently, (NASCIMENTO; QUEIROZ; JUNQUEIRA, 2021), wrote about them, considering the container loading problem, although such restrictions also appear in a wide variety of problems. For example, in (NASCIMENTO; QUEIROZ; JUNQUEIRA, 2021), mathematical formulations can be found for a wide variety of practical restrictions, such that these formulations can be easily extended to other cutting and packing problems.

Regarding pallet loading problems, it is possible to mention the systematic review conducted by (SILVA; OLIVEIRA; WÄSCHER, 2016). These authors paid more attention to the problem that considers loading only one type of item onto a single pallet. They discussed aspects of computational complexity and the main methods developed by the specialized literature. Another review of articles on the same version of the problem was conducted by (BOSCH, 2016). The author focused his review on approaches that have been published in the last 30 years and how they could be used in practice by industries, especially considering the use of robots to carry out pallet loading. In (GUNAWARDENA; WIJAYANAYAKE; KAVIRATHNA, 2021), a two-phase heuristic can be found, initially constructing layers of items that respect stability and weight restrictions. Subsequently, the layers are used to obtain the complete pallet loading in a way that respects vertical stability restrictions.

Regarding the DPLP, the interest of the present article, a considerable number of publications can also be found in the literature. In general, the DPLP considers that there are different types of items and the objective is to minimize the number of pallets needed for this task. (BISCHOFF; JANETZ; RATCLIFF, 1995) proposed a heuristic that first creates layers of items and then loads these layers onto pallets, respecting vertical stability restrictions. In (TERNO et al., 2000), the problem was solved using an enumerative branch-and-bound algorithm. The author also considered the inclusion of practical restrictions such as weight limits and vertical stability. (QUEIROZ; MIYAZAWA; WAKABAYASHI, 2015) investigated the properties of methods proposed in the literature for the problem, identifying instances for which such methods failed to obtain an optimal solution.

A real application of the DPLP in a logistics company was studied by (GZARA; ELHEDHLI; YILDIZ, 2020). The company employs automation resources to deal with the problem, using robots to load pallets. Dealing with thousands of items and pallets every day, each pallet needed to be loaded in about two minutes. The proposal of these authors was a column generation-based heuristic approach to construct layers of items that respect weight limit, sequencing, and vertical stability restrictions. (CALZAVARA et al., 2021) also dealt with a real robotic application, for which it was necessary to respect vertical stability, fragility, weight limit, contiguity, visibility, and item rotation restrictions. The authors developed a heuristic for the construction of item layers and, subsequently, pallets, a metaheuristic based on a reactive and adaptive randomized greedy search, and a metaheuristic that uses an integer linear programming model.

In the literature, it is also possible to find proposals for solution methods for the DPLP and other pallet loading problems that use artificial intelligence techniques. In (LAYEB; OMRI, 2024), a reinforcement learning model was applied to deal with the problem in the presence or absence of the vertical stability constraint. The authors were able to obtain solutions with an average pallet volume occupancy of 99%. (THEODOROPOULOS et al., 2025) dealt with a real problem in an aluminum industry that uses robots to load pallets. The items do not necessarily have a regular shape, dealing with the general case of the problem, in addition to including vertical stability and weight limit constraints. In addition to proposing a priority-based constructive heuristic, the authors developed a reinforcement learning model to determine the next item to be loaded onto the pallet. In turn, (MAGNANI et al., 2025) combined a constructive heuristic with random forest and support vector regression machine learning techniques. The heuristic is responsible for creating layers with the items, while the learning techniques select the layers to build pallets that respect constraints such as vertical stability and fragility.

The version of the DPLP under study differs from those present in the reviewed literature due to the characteristics of the ceramic production industry. It has very specific restrictions linked to its daily operations, which includes dependence on an already active automation system in its warehouse. There is also the range of products offered, with very diverse characteristics, which requires many practical restrictions to be respected. This would require a highly specialized computational system to solve the DPLP and provide a solution to customers even when they are placing orders. To circumvent this situation and, at the same time, provide quick and accurate solutions, we propose a hybrid approach. This approach uses different machine learning models, in addition to considering bounds calculated by constructive heuristics, in order to predict the number of pallets needed for a given customer order. We emphasize that some results of this approach, limited to the XGBoost model, were published in (TACCINI et al., 2025).

This article is organized as follows. Section 4.2 contains the definition of the problem under study, including a description of the current scenario faced by the industry that motivates the need for an approach that uses artificial intelligence techniques. Section 4.3 describes the hybrid approach, with details on how historical data provided by the industry and heuristic solutions are used to train machine learning models. Section 4.4 describes computational experiments performed on real-world instances and discusses their results. Section 4.5 contains conclusions and considerations for future work.

4.2 Problem Definition

The problem addressed in this article is a version of the DPLP that arises from a real situation experienced by a ceramic industry. Using its e-commerce platform, the customer places their order containing one or more types (of boxes) of ceramics. At the conclusion of the order, the customer is informed about the (estimated) quantity of pallets they need to pick up at the industry. In turn, the industry considers the ceramics available in its warehouse or the need to produce according to the orders received. Once produced, they are loaded onto production pallets, generally homogeneous (i.e., with the same type of ceramic). These pallets are then positioned within the industry's warehouse by an automated system. It is important to emphasize that the term *ceramic* is being used to refer to a box containing one or more ceramic pieces, depending on their dimensions, volume, and weight.

The fulfillment of a customer's order occurs as follows. Operators request the ceramics that make up the order from the automated warehouse system. The system locates the production pallets capable of fulfilling the order and, via conveyors, sends them to the location (operation bay) where operators load, with forklifts, the shipping pallets. Each bay receives up to four production pallets simultaneously, while also having a limited operating area. As the bays are limited by their area, operators need to immediately load the shipping pallets with the production pallets that are available, also to free up space for other production pallets to arrive at the bay.

Another important point to highlight is that operators have no control over which production pallets arrive at their bay, as this depends on the automated system, the location of production pallets within the warehouse, and the other bays, which are also operating in parallel. Even with this uncertainty, operators aim to minimize the number of pallets used to fulfill the customer's order, respecting vertical stability constraints, i.e., items cannot fall due to gravity after being positioned on the pallet, maximum weight limit supported per pallet, allowed orientations for positioning ceramics, and fragility of ceramics regarding stacking others on top of them. Operators load pallets in horizontal layers, observing the mentioned restrictions, until the maximum allowed pallet height is reached.

Due to these factors and the dynamism of warehouse operations, it becomes too costly for the industry to obtain an exact solution or even a heuristic solution that meets all the constraints and uncertainties present in this variant of the DPLP. Even a solution that respected the problem's constraints could be unfeasible for operators to implement due to the uncertainties propagated by the automated warehouse system (which they cannot directly interfere with). The industry, then, is concerned with having an approach that allows it to quickly and accurately estimate the number of pallets needed to fulfill each customer's order. In this context, a hybrid approach is proposed that combines machine learning techniques with constructive heuristics, in addition to considering historical data.

4.3 Proposed Hybrid Approach

The hybrid approach uses historical industry data, in addition to using constructive heuristics to generate valid bounds. All of this generates information that is organized into columns of a matrix, which serves as input for the machine learning models. Unlike previous work in (TACCINI et al., 2025), which focused solely on the XGBoost model (CHEN; GUESTRIN, 2016), it is extended to include other models, such as Light Gradient-Boosting Machine (LGBM) (SHI et al., 2024) and Random Forest (RF). The goal is to generalize the approach by including other machine learning models.

4.3.1 Column Set

The set of columns presented in Table 7 was extracted from the ceramic industry's historical data and the results of the heuristics used to provide bounds for the problem.

These historical records contain orders placed by different customers since 2022 and include the following information:

- the quantity of each type of ceramic in the order, denoted by F_i;
- the total quantity, weight, and volume of the ceramics, denoted by F_o, W_o, and V_o, respectively;
- average weight and volume of the ceramics, denoted by W_{avg} and V_{avg}, respectively;
- the actual quantity of shipping pallets used, denoted by Z_j and used only in the training set for the machine learning models.

From the information provided by the industry, the inclusion of three other columns containing the results obtained by constructive heuristics for each order was considered:

- Heuristic HS_1, which consists of dividing the total weight and total volume of the order by the maximum weight and volume limit of the shipping pallets. In other words, it is a heuristic that estimates a lower bound for the number of pallets needed to fulfill the order;
- Heuristic HS_2, which separates the types of ceramics into homogeneous shipping pallets, calculates the number of shipping pallets needed for each type, and sums them at the end. In other words, it is a heuristic that estimates an upper bound for the number of pallets.
- Heuristic HS_3, which considers loading the shipping pallet layer by layer, starting with the largest ceramics until the pallet's height limit is reached. It continues loading the pallets in this way until the customer's order is complete. In other words, it is a simple constructive heuristic that provides the number of pallets for each order, but without guaranteeing that all restrictions are respected at the end of loading, such as vertical stability.

Table 7 – Set of columns to be used as input by the hybrid approach.

j	F_1	$\cdots$	F_m	F_o	W_o	V_o	W_{avg}	V_{avg}	HS_1	HS_2	HS_3	Z_j
1	7	$\cdots$	2	21	170	525	12	25	6	10	9	8
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
n	10	$\cdots$	0	12	200	75	15	5	1	3	3	3

4.3.2 Machine Learning Models

The machine learning models considered in the hybrid approach are XGBoost, LGBM, and RF. These models were chosen after preliminary computational experiments on the columns presented in Table 7. Then, their hyperparameters were calibrated through a grid search. The values used for these parameters during calibration were obtained from empirical analyses, standard values from libraries that implement such models, and based on the literature (ZOUININA, 2024; AHMED, 2022; SHARMA, 2024; SAXENA, 2024). Table 8 contains the values considered for each model and the best value obtained after calibration.

Table 8 – Set of values adopted during calibration by grid search.

Model	Parameter	Values	Best Value
XGBoost	learning_rate	0.01, 0.1, 0.2, 0.3	0.1
	gamma	0, 0.1, 0.5	0.1
	max_depth	3, 6, 10	10
	min_child_weight	1, 3, 5	3
	subsample	0.5, 1	0.5
	colsample_bytree	0.5, 1	0.5
LGBM	learning_rate	0.001, 0.01, 0.1	0.1
	max_depth	3, 5, 7, 10, ∞	∞
	num_leaves	7, 31, 127, 1023	1023
	min_data_in_leaf	20, 50, 100	50
	feature_fraction	0.8, 1.0	1.0
RF	max_depth	10, 50, 100, ∞	50
	min_samples_split	2, 5, 10, 20	20
	min_samples_leaf	1, 3, 6, 10	1
	max_features	1, 3, 6	6
	bootstrap	True, False	False

4.4 Computational Experiments

The hybrid approach, considering the three machine learning models and the three heuristics, was implemented in Python 3.10 using the *scikit-learn* library (PEDREGOSA et al., 2011). Experiments were conducted on a computer with an Intel Core i5-1135G7 2.40 GHz processor, 8 GB of memory, and a *Linux Mint 21 Cinnamon* operating system.

The machine learning models were initially trained using a portion of the historical data representing a set of approximately 167,000 orders. The orders were randomly split with 80% allocated for the training phase. Furthermore, only the more elaborate orders, i.e., those containing more than one type of ceramic, were considered for the experiments to ensure the quality of the final results. Thus, homogeneous orders, i.e., those with only one type of ceramic and that produced only complete pallets, were disregarded.

The evaluation of the models considers two metrics: mean squared error (MSE) and mean absolute error (MAE). The results of the initial training phase are presented in Table 9.

Table 9 – Results of the models on the training order set.

Model	MSE	MAE
XGBoost	0.728	0.329
LGBM	0.726	0.326
RF	0.779	0.332

The results in Table 9 show that LGBM and RF achieved very similar results to XGBoost in the training phase. LGBM stands out with MSE and MAE values lower than those of the XGBoost model, while the RF model presents a relatively higher MSE value than the other two models.

From the trained models, the second phase of the experiments considers their validation on the orders not used in the previous phase. The results for 30 orders (instances) received by the industry in 2024 are presented below, representing a period entirely different from that of the orders used in the models' training phase. This allows the models to be evaluated in a context where the orders are completely new, similar to what would be done by a forecasting system developed for the industry. In addition, these instances include heterogeneous orders in which more than one type of ceramic is requested, with this value varying from 2 to 6 different types.

The results of the models are compared with PackVol (PACKVOL, 2024), a software used by the industry to heuristically solve the DPLP. It is important to note that this software does not take into account the uncertainties present in the industry's warehouse, as discussed in Section 4.3, for example, the fact that not all ceramics are available to operators at the time of loading shipping pallets. The results obtained for the 30 instances are summarized in Table 10 in terms of MSE and MAE values.

Table 10 – Results on the 30 real industry instances.

Model	MSE	MAE
PackVol	2.133	0.933
XGBoost	0.500	0.433
LGBM	0.467	0.400
RF	0.400	0.400

The results in Table 10 show that all models achieved significantly better results than PackVol. Regarding execution time, PackVol took 1 second to solve each of the 30 instances, whereas the models solved all 30 in less than 0.5 seconds. Thus, it is possible to observe that the hybrid approach, even considering different machine learning models, is more accurate and faster than the approach (i.e., PackVol) currently used by the industry. The RF model stands out with the lowest error values.

The LGBM and RF models present very similar results to each other and better than the results of the XGBoost model. The LGBM and RF models return the same MAE value, with the RF model being superior in terms of MSE. This advances the conclusions pointed out in (TACCINI et al., 2025), which presented only the XGBoost model as part of the hybrid approach. In addition, PackVol results in MSE values up to five times worse and MAE values up to two times worse when compared to the machine learning models.

4.5 Conclusion and Future Work

This article presented a hybrid approach for predicting the number of pallets in the context of the DPLP that arises in a ceramic manufacturing industry. The hybrid approach consists of three machine learning models and three simple heuristics for the problem. The models were calibrated and trained on an extensive set of industry orders. This process allowed their application to real instances, as well as comparison with software currently used by the industry.

The results of the computational experiments indicate the superiority of the hybrid approach, especially as it is not limited to the XGBoost model but also considers LGBM and RF. All models returned superior results to PackVol based on MSE and MAE metrics. In the comparison between them, the RF model presents the best overall results, surpassing the proposal previously presented in (TACCINI et al., 2025) and indicating that the hybrid approach can be an accurate and fast alternative for the industry.

Future work could evaluate other machine learning models, especially outside the family of tree-based methods. Possible options would be an approach using neural networks or support vector regressors. Another point to explore is extending the prediction by the hybrid approach to orders involving different types of shipping pallets. It may also be interesting to investigate new heuristics for the problem, providing more precise bounds for the machine learning models.

5 COMBINING LEARNING AND HEURISTICS FOR PALLET PREDICTION IN CERAMICS DISTRIBUTION

This chapter also presents an extension of the work of Chapter 3. The model was extended to make predictions considering the types of pallets that will be required instead of only the total amount. It considers two groups of pallets: standard pallets and extra large pallets. Diverse approaches were built to make this multiple value prediction, all built upon the hybrid model for single pallet prediction. Even in this more challenging scenario, the models were able to outperform the company's current model-based solution and keep their predictions accurate and efficient. [1]

5.1 Introduction

The ceramic tile industry has experienced steady global growth in recent years, with worldwide production reaching 16.8 billion square meters (ACIMAC Research Department, 2023). Italy ranks among the top exporters, generating €7.2 billion in revenue, which highlights the strategic role of the sector in the national economy. The management of pallet loading operations and transportation is an important component of competitiveness in this industry. These operations involve transferring boxes of ceramic tiles, which often vary in size and weight, from production pallets (used in internal logistics operations) to shipping pallets (used for external logistics) to fulfill customer orders. To ensure safe transportation and minimize logistics costs, each pallet must comply with specific weight and volume limits and maintain adequate structural integrity.

This study addresses a real-world pallet loading problem faced by a major international ceramic tile company headquartered in Italy. Customers' orders are submitted through the company's e-commerce platform or via standard commercial channels. Each customer is responsible for collecting the ceramics, either using their vehicle or relying on a third-party carrier. Consequently, the company must inform the customer in advance of both the number of pallets required for each order and the types

[1] The results of this chapter appear in: Taccini, M.; Oliveira, M.A.; Santos, A.G.; Queiroz, T.A.; Iori, M. (2025). Combining Learning and Heuristics for Pallet Prediction in Ceramics Distribution. In: Proceedings of the International Conference on Optimization and Decision Science (ODS2025).

of pallets that will be used. Since pallet types vary in size and capacity, they impact transportation planning and cost estimation for the customer.

To provide customers with this information, the company must solve a variant of the distributor's pallet loading problem (DPLP) (SILVA; OLIVEIRA; WÄSCHER, 2016), aiming to minimize the number of pallets required for loading boxes. Providing an optimal solution for such a problem is not a straightforward task, as it involves not only its computational complexity (i.e., an NP-hard problem), but also numerous constraints and uncertainties arising from the company's warehouse and daily practices (e.g., not all production pallets are available at the same time to build the shipping pallets). Such limitations underscore the need for a tailored solution method that can learn from historical data and provide accurate, real-time predictions for the customers' orders.

In this work, we propose an approach that predicts both the quantity and types of shipping pallets required to fulfill a customer's order. It integrates machine learning with optimization heuristics to obtain accurate predictions, extending the preliminary study in (TACCINI et al., 2025). Its results are compared with the current solutions proposed by the company. These include decisions made by the warehouse's expert operators as well as the output produced by a specialized pallet loading software tool (PACKVOL, 2024). The results of the computational experiments demonstrate that the proposed approach provides practical benefits within the company's logistics processes. Additionally, this study contributes to the literature by bridging the gap between machine learning and operations research to solve a complex, highly constrained logistics problem in a dynamic industrial environment.

The structure of the paper is as follows. In Section 5.2, we briefly describe the literature about the DPLP and related problems. In Section 5.3, we formally present the problem under investigation. In Section 5.4, we describe the proposed approach and detail each of its parts. In Section 3.5, we present and discuss the computational experiments on the company's instances. In Section 5.6, we conclude our research by listing the main findings and providing directions for future studies.

5.2 Literature Review

There are different variants of pallet loading problems. These are usually classified as cutting and packing problems, for which we refer to the typology proposed by (WÄSCHER; HAUßNER; SCHUMANN, 2007). General surveys on cutting and packing problems were provided by (SALEM; SILVA; OLIVEIRA, 2023; IORI et al., 2021a), while real-world constraints have been discussed in works such as (BORTFELDT; WÄSCHER, 2013; NASCIMENTO; QUEIROZ; JUNQUEIRA, 2021). Pallet loading problems can be categorized into two main variants: those involving manufacturers and those involving distributors. The survey by (SILVA; OLIVEIRA; WÄSCHER, 2016) addressed the manufacturer's pallet loading problem. This variant assumes that all items are identical and seeks to maximize the number of items packed onto a single pallet.

The problem we address here is a generalization of the distributor's pallet loading problem (i.e., the DPLP), also known in the literature as the pallet building problem and the multi-pallet loading problem. It assumes that one or more pallets are used to pack a set of items of various types, thereby minimizing the number of pallets required.

The number of contributions on the DPLP in the literature is significant. The earliest works date back to the 1980s, such as (HODGSON, 1982), which proposed an algorithm based on dynamic programming. In (BISCHOFF; JANETZ; RATCLIFF, 1995), a heuristic that packs items in layers was introduced, allowing the inclusion of practical constraints such as vertical stability. Computational experiments yielded promising results; however, performance declined in instances with greater variability in box dimensions. In (QUEIROZ; MIYAZAWA; WAKABAYASHI, 2015), the authors reviewed previous algorithms, including that of (BIRGIN; LOBATO; MORABITO, 2012), and presented instances where such methods fail to find optimal solutions. In (GZARA; ELHEDHLI; YILDIZ, 2020), the authors addressed a DPLP for a logistics company requiring the packing of items onto pallets every two minutes in an automated warehouse processing thousands of items and hundreds of pallets daily. They proposed a layer-based column generation algorithm that incorporates constraints such as vertical stability, load-bearing capacity, pallet weight limits, and planogram sequencing. However, the load-bearing constraint was modeled using a simplified load transfer based on surface area percentages.

Another practical DPLP was addressed in (CALZAVARA et al., 2021), involving a robotized application with constraints including load-bearing, contiguity, visibility, and multiple orientations. Most recently, (YAO et al., 2025) introduced a DPLP variant that incorporates a buffer area, which is useful for dealing with uncertainty related to partial information about items that becomes available only before packing. They proposed a greedy heuristic combining block generation, open space representation, and a selection strategy to improve pallet utilization. The resulting effectively balances efficiency and real-time decision-making, accounting for robot interference and practical constraints such as vertical stability.

Regarding the use of artificial intelligence techniques for pallet loading problems, (LAYEB; OMRI, 2024) solved instances of the manufacturer's variant using deep reinforcement learning, achieving an average volume utilization of approximately 99%. In contrast, (MAGNANI et al., 2025) addressed the DPLP with practical constraints such as vertical stability and load-bearing. They proposed a hybrid approach that combines a heuristic, which constructs layers with items, with machine learning techniques (e.g., random forest and support vector regression) to identify promising layers for pallet construction.

5.3 Problem Description

The problem addressed in this study is a DPLP variant obtained from a large-scale ceramic tile manufacturer in Italy. We are given a set of customer orders. Each order requires a set of boxes. The boxes have three dimensions, and a demand is associated with each box. To supply these orders, the company has been using two types of shipping pallets: the first with a larger base area and higher load capacity, while the second is smaller and used for lighter or more compact products. The boxes are loaded in levels, respecting the pallet's rectangular base, height limit, and weight capacity, as well as its vertical stability, load-bearing requirements, and the allowed orientations of the boxes. There are also compatibility constraints, meaning that not both pallets can be used to transport all types of boxes. The problem's objective is to load all boxes of the given order using the minimum quantity of pallets, while satisfying all the previously mentioned constraints.

Additional details about the company's warehouse and loading operations are provided below to better situate the problem. Ceramic tiles are first loaded onto standardized production pallets, which are not used for shipping but act as intermediate supports during internal logistics transport. These pallets are transported to the loading area via an automated material handling system, which has a limited capacity. Consequently, production pallets cannot be simultaneously available for the company's operators during the loading of shipping pallets. Their arrival is influenced by the scheduling choices of the warehouse's automated system and its current operational status. For instance, machinery failures can delay pallet delivery, disrupting the sequence in which production pallets reach the loading area, which itself faces strict spatial limitations. As a result, operators must begin transferring boxes from each production pallet to the appropriate shipping pallet immediately upon its arrival, without delay. Only after finalizing an order can operators proceed to the next one.

When loading the shipping pallets, operators must adhere to the previously mentioned constraints. Additionally, the two types of shipping pallets that the company handles are tailored to the specific dimensions of different products. This further complicates the loading operations since not every product can be loaded onto every type of pallet. The selection of a type affects not only the feasibility and efficiency of the loading process but also the information provided to customers for transportation planning purposes.

Solving this DPLP variant is particularly challenging. Even when a feasible loading plan is found, it may still be impractical for operators to implement the corresponding solution due to the unpredictability of day-to-day operations and the potential presence of packing movements that are difficult to execute in practice. Additionally, the company must inform customers, at the time of purchase, about the number and type of shipping pallets required for their order. Therefore, the company is currently interested in having a fast and accurate estimator for these values, rather than relying solely on an optimization procedure that, while correct, may require significant computation time.

5.4 Proposed Approach

The proposed approach is an extension of our previous study presented in (TACCINI et al., 2025). That work did not account for the different types of shipping pallets the company uses to supply its customers' orders. The company is indeed interested in a more precise prediction, not only obtaining a total expected number of pallets but also accounting for their sizes (i.e., how many pallets are needed for each available size). To this end, we developed a new approach that utilizes independent regressors and regressor chain strategies to achieve multiple outputs.

5.4.1 Single-Output Regressor Model

The approach developed in (TACCINI et al., 2025) involved a selection phase in which four different machine learning (ML) models were trained and compared on both the mean squared error (MSE) and the mean absolute error (MAE). At the end of this selection phase, XGBoost (CHEN; GUESTRIN, 2016) was identified as the most effective model. Then, the XGBoost model was enhanced through a feature optimization process. This was performed by analyzing the structure of the input dataset. At the end of this second phase, the training dataset was improved by adding more general and order-based features, such as the total number of boxes in the order, total weight and volume of the order, and average weight and volume of boxes in the order, which reduced the MSE and MAE values by 66% and 42%, respectively.

To further improve the XGBoost model, three additional optimization-based features were introduced, namely the results obtained with heuristics HS_1, HS_2, and HS_3 (TACCINI et al., 2025). The first feature, HS_1, computed as the maximum between total weight divided by max pallet weight, and total volume divided by max pallet volume, can be interpreted as a lower bound (LB) for the problem. On the other hand, the second feature, HS_2, represents an upper bound (UB) for the DPLP instance, assuming that all pallets are composed exclusively of homogeneous boxes. The third added feature, HS_3, is based on a greedy heuristic that loads pallets by stacking layers of boxes with identical dimensions, starting with those having the largest area and continuing until all constraints are respected.

The inclusion of all these features resulted in a dataset structure, called S_3, which further decreased the MSE and MAE errors. The approach developed in (TACCINI et al., 2025) achieved superior performance, yielding better predictions within shorter computation times.

5.4.2 Multi-Output Regressor Model

Building on the approach in (TACCINI et al., 2025), we extend it to a multi-output regression model. We consider the following strategies to comply with the company's current objectives:

1. In the first scenario, we use two independent single-output regressor models, namely M_1 and M_2. Each one consists of the XGBoost model and was trained separately to predict each type of available shipping pallet, namely, pallet type $j = 1$ and pallet type $j = 2$. The training process assumed the dataset structure S_3. Despite its straightforward implementation, this strategy does not account for potential correlations between outputs for mixed orders, i.e., orders for which both types of shipping pallets, $j = 1$ and $j = 2$, may be present. The number of orders in this situation is around 30% of the total. We refer to this model as *independent regressors* (IR).

2. The second scenario involves training the models sequentially, with each one receiving as input both the original features, i.e., S_3, and the predictions, which we denote as P, from one of the previous models. This strategy has allowed the modeling of interdependencies between the types of pallets. We considered two strategies, since the final results/predictions are influenced by the order in which a model is applied to predict a given type of pallet. The first strategy involves using the first model, M_1, with input S_3, resulting in the first output, P_1, for the shipping pallet of type $j = 1$. Next, the second model, M_2, has as input $S_3 \cup P_1$, for the shipping pallet of type $j = 2$. We refer to this strategy as *regressor chain one* (RC1). The second strategy involves starting with the second model, M_2, which takes S_3 as input to predict the pallet of type $j = 2$. Next, we use the first model, M_1 with the input $S_3 \cup P_2$, for predicting the shipping pallet of type $j = 1$. This is named *regressor chain two* (RC2).

Figure 6 illustrates each of the three strategies. We consider $input = S_3$, while M_1 and M_2 are the XGBoost models for each type of pallet, $j = 1$ and $j = 2$, respectively.

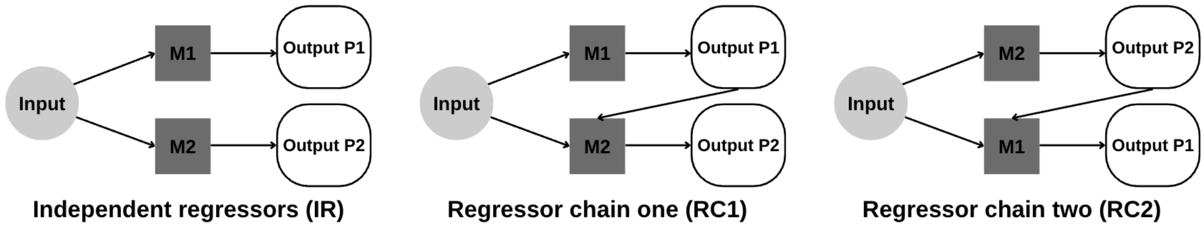

Figure 6 – The three strategies developed to predict the solution of the DPLP variant under investigation.

It is worth noting that we tested various strategies to enhance the approach's accuracy; however, the ones we present here yielded the best overall results.

The tuning of the hyperparameters and the training phase for the three strategies were carried out according to the procedure described in (TACCINI et al., 2025). The orders were randomly split, with 80% allocated to the training phase, using a dataset containing approximately 166,000 orders from the years 2022 and 2023.

Table 11 reports the MSE and MAE values calculated for the test dataset, which consists of 20% of the total initial dataset. The table compares the three different strategies, IR, RC1, and RC2, based on their prediction performance. For the test dataset, while all strategies behave similarly, IR slightly outperforms the others globally and particularly for the pallet of type $j = 1$. However, for $j = 2$, we observe RC1 and RC2 show advantages depending on which measure of error we would like to prioritize, MSE or MAE.

Table 11 – MSE and MAE values for each strategy, calculated on the test dataset. The first two columns report the average values across the two pallet types $j = 1$ and $j = 2$, while the remaining columns show the type-specific metrics.

Strategy	MSE	MAE	MSE $j = 1$	MAE $j = 1$	MSE $j = 2$	MAE $j = 2$
IR	0.659	0.256	1.133	0.407	0.185	0.105
RC1	0.663	0.257	1.129	0.414	0.198	0.100
RC2	0.660	0.267	1.140	0.430	0.181	0.104

5.5 Computational Results and Discussions

The proposed approaches were implemented in Python 3.10. We performed computational experiments on a computer equipped with an Intel Core i5-1135G7 2.40 GHz CPU, 8 GB of RAM, and Linux Mint 21 Cinnamon as the operating system. The experiments utilized 30 real instances from the company, which were randomly selected from all the heterogeneous orders received in 2024 and represent different levels of heterogeneity. These orders were not considered during the training phase to guarantee a fair comparison of the strategies.

The instances are ordered by their heterogeneity, i.e., the number m of different types of boxes the order has, with a maximum of $m = 6$. The three strategies are compared with PackVol (PACKVOL, 2024), the software already in use by the company. It is configured to reflect the specific operational constraints of the DPLP variant under investigation; however, it does not account for the uncertainties reported in Section 5.3. We also compare our solutions with the company's expert operator solutions, i.e., the shipping pallets the operators load to fulfill each order. These values are used as a baseline for comparing all the approaches.

Table 12 has the results obtained for all the 30 instances (i.e., the number of pallets $j = 1$ and $j = 2$ required by each instance). In contrast, Table 13 reports the MAE and MSE values for each approach, with the company's operator solutions serving as the baseline for comparison. Firstly, we observe that the difference in errors between the independent regressor and the chained ones is very small, suggesting that the correlation between the two types of pallets j is weak. This could indicate that a simpler approach, such as IR, without inter-dependencies between the two types, is already able to return satisfactory predictions.

Concerning Table 12, as the heterogeneity of the boxes increases (ranging from 2 to 6 in instances 1–30), values for all approaches tend to increase slightly, suggesting that the problem becomes more complex as the orders' heterogeneity grows. PackVol shows a tendency to overestimate the number of pallets of type $j = 2$, which are typically associated with larger sizes. Additionally, IR, RC1, and RC2 often perform as well as, or better than, PackVol's solutions. RC2 often produces very competitive results, particularly in instances with more heterogeneity (instances 20–30).

Considering the type of pallet $j = 2$, we notice that the predictions better approximate the company's solutions. It is worth noting that PackVol gives poor results for pallets of type $j = 1$ compared to the other approaches.

Regarding the MAE and MSE values reported in Table 13, we conclude that RC2 is the most balanced and robust strategy, achieving the lowest overall errors and better handling the types of pallets $j = 1$. It achieves the lowest overall MSE (0.533) and lowest MAE (0.467), indicating the highest prediction accuracy globally. On the other hand, RC1 has the best results specifically for the type $j = 2$. IR remains a simple and interesting strategy, as its predictions are more accurate than those of PackVol. Notice that PackVol performs much worse than the others, especially for $j = 1$, with an MSE of 3.100 and MAE of 1.300.

Table 12 – Results of the 30 real-world instances for each approach. For each instance, the number of shipping pallets is reported for each type $j = 1$ and $j = 2$. The values are shown for the company's operator solutions, Packvol, and the proposed strategies: IR, RC1, and RC2.

Instance	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
m	2	2	2	2	2	2	3	3	3	3	3	3	4	4	4
Company: $j = 1$	2	5	3	1	2	1	17	5	2	8	5	1	5	5	7
Company: $j = 2$	0	1	1	0	4	1	5	0	3	0	5	6	0	3	2
PackVol: $j = 1$	1	6	3	1	4	1	17	5	0	8	5	0	4	1	6
PackVol: $j = 2$	0	0	1	0	2	1	5	0	3	0	5	6	0	5	2
IR: $j = 1$	1	6	3	1	2	1	18	5	2	8	5	1	5	3	6
IR: $j = 2$	0	1	1	0	4	1	5	0	3	0	4	6	0	2	2
RC1: $j = 1$	1	5	3	1	1	1	18	5	2	9	5	1	5	3	7
RC1: $j = 2$	0	1	1	0	4	2	5	0	3	0	4	6	0	3	2
RC2: $j = 1$	1	6	3	1	2	2	19	5	2	9	6	1	6	3	7
RC2: $j = 2$	0	1	1	0	4	1	5	0	3	0	4	7	0	3	2

Instance	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
m	4	4	4	5	5	5	5	5	5	6	6	6	6	6	6
Company: $j = 1$	3	15	5	19	6	22	1	12	10	20	8	14	14	9	14
Company: $j = 2$	4	0	1	7	3	6	1	1	0	2	2	8	7	1	2
PackVol: $j = 1$	2	14	4	17	4	21	2	13	6	16	5	13	12	9	12
PackVol: $j = 2$	5	1	2	8	4	6	2	0	0	2	3	8	7	1	2
IR: $j = 1$	3	15	6	18	6	18	2	14	9	19	9	14	12	9	13
IR: $j = 2$	4	1	1	8	3	7	1	0	0	1	2	7	6	1	2
RC1: $j = 1$	3	16	5	18	7	18	2	16	10	19	9	12	10	9	14
RC1: $j = 2$	4	0	1	7	3	7	1	0	0	1	2	8	7	1	2
RC2: $j = 1$	3	15	6	18	7	22	2	13	9	19	9	15	13	10	14
RC2: $j = 2$	4	0	1	8	3	7	1	1	0	1	2	7	6	1	2

Table 13 – MSE and MAE values for each approach in the 30 instances.

Model	MSE	MAE	MSE $j = 1$	MAE $j = 1$	MSE $j = 2$	MAE $j = 2$
PackVol	1.833	0.867	3.100	1.300	0.567	0.433
IR	0.800	0.500	1.300	0.700	0.300	0.300
RC1	1.183	0.517	2.200	0.867	0.167	0.167
RC2	0.533	0.467	0.833	0.700	0.233	0.233

5.6 Concluding Remarks

This study addressed a real-world variant of the distributor's pallet loading problem faced by a major ceramic tile manufacturer in Italy. The complexity of the problem lies not only in the geometric and operational constraints inherent to the palletization process, but also in the uncertainties associated with warehouse logistics and day-to-day operations. In this context, providing customers with accurate predictions regarding the number and types of pallets required for their orders is a critical component of efficient logistics management.

To tackle this challenge, we proposed a machine learning-based approach that integrates independent and sequential multi-output regression strategies with features derived from both historical data and optimization-based heuristics. Three prediction strategies were developed: independent regressors (IR), regressor chain one (RC1), and regressor chain two (RC2).

Computational experiments conducted on 30 real instances from the company demonstrated the effectiveness of the proposed approach. In particular, the RC2 strategy consistently achieved the best performance, obtaining the lowest overall mean squared error and mean absolute error among all tested approaches. While RC1 provided slightly better results for pallets of type $j = 2$, RC2 offered the most balanced and robust performance across different order heterogeneity. Furthermore, the IR strategy also proved to be a simple yet effective alternative, outperforming the specialized PackVol software, particularly for pallets of type $j = 1$.

Overall, the results confirm that integrating machine learning and optimization-based features can significantly enhance pallet prediction accuracy, offering practical advantages for real-world industrial logistics. The proposed approach provides an effective and scalable solution for dynamic and highly constrained environments. Future research directions include the exploration of more advanced machine learning models, such as deep learning architectures, incorporating real-time warehouse operational data to further improve prediction accuracy, and extending the approach to support dynamic re-optimization as new information becomes available during the loading process.

6 CONCLUSIONS

In this dissertation, we proposed and implemented a machine learning based solution to predict the total pallet demands for the DPLP in a real world ceramic company context. The model uses the historical data from the company to estimate the pallet demands of new orders. A set features was constructed and refined over this data, and the XGBoost ML model was selected to built our strategy. Heuristics were proposed and used to enrich the ML model, something that improved the its accuracy. The overall solution was compared against PackVol, the current software of the company, and it presented a both a faster performance and better accuracy.

The strategy built over XGBoost was extended to be used in other ML models: LightGBM and Random Forest. Both approaches were also able to beat PackVol's efficiency and accuracy by a far margin and were slightly better than XGBoost. This result indicated that the strategy is general enough to support other ML models and not limited to XGBoost, its original model.

The problem and models were then extended to make the prediction of the quantity of multiple pallet groups instead of only the total quantity of the order. Two groups of pallets were considered: standard pallets and extra large pallets. Different approaches were built to make the prediction of multiple values, using the single prediction model as a base. Even with this additional complexity in the problem, all the proposed approaches were able to beat PackVol again by a far margin in both terms of performance and accuracy.

We can revisit the proposed hypothesis of Section 1.2 and note that they here validated:

Hypothesis 1: An ML model trained on historical company orders can predict the number of pallets required for new orders based on their general characteristics (e.g., number of tiles, weight, and volume);

Validation: Multiple solutions were able to be implemented using different ML models (XGBoost, LightGBM and Random Forests) and all were able to present accurate predictions as observed in Chapter 4. The solutions were accurate in both the single regression as shown in chapters 3 and 4 and multiple regression cases as shown in Chapter 5.

Hypothesis 2: Heuristics can be used to generate additional features to enrich the ML models and improve their performance;

Validation: As observed in Subsection 3.4.3, the addition of 3 heuristic values (HS_1,HS_2 and HS_3) to the model's input columns were able to improve the MSE of the model by around 1.5 times and the MAE by around 1.2 times.

Hypothesis 3: An ML solution is able to provide estimations much faster than the current company approaches and with better or close enough accuracy;

Validation: As observed in sections 3.5, 4.4 and 5.5, the hybrid solution is able to consistently beat PackVol in both accuracy and speed for all the ML models used (XGBoost, LightGBM and Random Forest) as well as for both single and multiple regression cases.

In conclusion, an effective solution was proposed for the single prediction case, shown to work for multiple ML models and extended for the multiple prediction case keeping its effectivity. Therefore, the packing problems study area and the company can be benefit from this project.

As a future research, there are some points that can be extended in the current work:

- Make deeper analysis over the order's features and their impact in the prediction. This can help improving the feature set (and consequently the model accuracy) as well as give the company some insights on how to make their production process more effective to reduce pallet usage.
- The ML models used were all based on decision trees. Models from other families could also be considered and tested, like rule-based regression models.
- The multiple value prediction model could use an strategy that was not based on the single value case. Models that support multiple value regression directly, like neural networks, could be considered. Additionally, heuristics specifically made for the multiple pallet case could be proposed to enrich the models in this.
- The model could be extended and tested in other ceramic company scenarios, or even in companies from other industries that also deal with the DPLP (e.g. foods and beverages). A totally general model that did not consider the industries' characteristics could also be proposed.

REFERENCES

ACIMAC Research Department. *World Production and Consumption of Ceramic Tiles*. 2023. <https://www.mecs.org/it/shop/cat/16-mercato-ceramico/prod/111-produzione-e-consumo-mondiale-di-piastrelle-ceramiche>. Accessed: 2025-05-17. Citado 4 vezes nas páginas 16, 21, 32, and 41.

AHMED, Awadelrahman M. A. *LightGBM Cheat Sheet*. 2022. <https://awadrahman.medium.com/lightgbm-cheat-sheet-ddde43f9eaf1>. Accessed in: 2025-05-04. Citado na página 38.

AKKAYA, Selen Burçak; GÜL, Aykut; COŞKUN, Zeynep; KARAMAN, Coşku; ÖZTOP, Hande; MULLAOĞLU, Gizem. The distributor's pallet loading problem: A case study. In: SPRINGER. *The International Symposium for Production Research*. [S.l.], 2018. p. 937–948. Citado na página 16.

ALONSO, Maria Teresa; ALVAREZ-VALDES, Ramon; PARREÑO, Francisco; TAMARIT, Jose Manuel. Algorithms for pallet building and truck loading in an interdepot transportation problem. *Mathematical Problems in Engineering*, Wiley Online Library, v. 2016, n. 1, p. 3264214, 2016. Citado 2 vezes nas páginas 13 and 16.

AO, Wenxuan; ZHANG, Guozhen; LI, Yong; JIN, Depeng. Learning to solve grouped 2d bin packing problems in the manufacturing industry. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. [S.l.: s.n.], 2023. p. 3713–3723. Citado na página 13.

BANERJEE, Prashant. *A Guide on XGBoost hyperparameters tuning*. 2020. <https://www.kaggle.com/code/prashant111/a-guide-on-xgboost-hyperparameters-tuning>. Accessed: 26 of September of 2024. Citado na página 25.

BENGIO, Yoshua; LODI, Andrea; PROUVOST, Antoine. Machine learning for combinatorial optimization: a methodological tour d'horizon. *European Journal of Operational Research*, Elsevier, v. 290, n. 2, p. 405–421, 2021. Citado na página 23.

BIRGIN, E. G.; LOBATO, R. D.; MORABITO, R. Generating unconstrained two-dimensional non-guillotine cutting patterns by a recursive partitioning algorithm. *Journal of the Operational Research Society*, Taylor & Francis, v. 63, n. 2, p. 183–200, 2012. Citado 2 vezes nas páginas 23 and 43.

BISCHOFF, E.E.; JANETZ, F.; RATCLIFF, M.S.W. Loading pallets with non-identical items. *European Journal of Operational Research*, v. 84, n. 3, p. 681–692, 1995. Citado 3 vezes nas páginas 13, 34, and 43.

BORGES, Yulle G. F.; LIMA, Vinícius L. de; MIYAZAWA, Flávio K.; PEDROSA, Lehilton L. C.; QUEIROZ, Thiago A.; SCHOUERY, Rafael C. S. Algorithms for the bin packing problem with scenarios. *Journal of Combinatorial Optimization*, v. 48, n. 4, p. 34, Oct 2024. Citado na página 31.

BORTFELDT, Andreas; WÄSCHER, Gerhard. Constraints in container loading - a state-of-the-art review. *European Journal of Operational Research*, v. 229, n. 1, p. 1–20, 2013. ISSN 0377-2217. Citado 2 vezes nas páginas 33 and 42.

BOSCH, Thomas Georg. *A review of the manufacturer's pallet loading problem research and experiments with an emphasis on the aspect of implementation*. Sacramento, USA: [s.n.], 2016. Citado na página 33.

CALZAVARA, Gabriele; IORI, Manuel; LOCATELLI, Marco; MOREIRA, Mayron C. O.; SILVEIRA, Tiago. Mathematical models and heuristic algorithms for pallet building problems with practical constraints. *Annals of Operations Research*, 2021. ISSN 1572-9338. Citado 3 vezes nas páginas 23, 34, and 43.

Ceramic World Web. *Italian ceramics: total 2024 turnover reaches €7.5 billion*. 2025. Accessed: 2025-09-28. Disponível em: <https://ceramicworldweb.com/en/economics-and-markets/italian-ceramics-total-2024-turnover-reaches-eu75-billion>. Citado na página 16.

CHEN, Tianqi; GUESTRIN, Carlos. XGBoost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY, USA: ACM, 2016. (KDD '16), p. 785–794. ISBN 978-1-4503-4232-2. Disponível em: <http://doi.acm.org/10.1145/2939672.2939785>. Citado 3 vezes nas páginas 25, 36, and 45.

CRAINIC, Teodor Gabriel; PERBOLI, Guido; TADEI, Roberto et al. Recent advances in multi-dimensional packing problems. *New technologies-trends, innovations and research*, IntechOpen, v. 1, p. 91–110, 2012. Citado na página 13.

FANG, Jie; RAO, Yunqing; ZHAO, Xusheng; DU, Bing. A hybrid reinforcement learning algorithm for 2D irregular packing problems. *Mathematics*, v. 11, n. 2, 2023. ISSN 2227-7390. Citado na página 23.

Grand View Research. *Ceramics Market (2024 - 2030)*. 2023. Accessed: 2025-09-28. Disponível em: <https://www.grandviewresearch.com/industry-analysis/ceramics-market>. Citado na página 16.

GUNAWARDENA, Kavindu; WIJAYANAYAKE, Annista; KAVIRATHNA, Chathumi. Solve manufacturer's pallet loading problem (mplp) with practical warehouse constraints. In: *2021 6th International Conference on Information Technology Research (ICITR)*. [S.l.: s.n.], 2021. p. 1–6. Citado na página 33.

GZARA, Fatma; ELHEDHLI, Samir; YILDIZ, Burak C. The pallet loading problem: Three-dimensional bin packing with practical constraints. *European Journal of Operational Research*, v. 287, n. 3, p. 1062–1074, 2020. Citado 2 vezes nas páginas 34 and 43.

HODGSON, Thom J. A combined approach to the pallet loading problem. *A I I E Transactions*, Taylor & Francis, v. 14, n. 3, p. 175–182, 1982. Citado 2 vezes nas páginas 23 and 43.

IORI, Manuel; de Lima, Vinícius L.; MARTELLO, Silvano; MIYAZAWA, Flávio K.; MONACI, Michele. Exact solution techniques for two-dimensional cutting and packing. *European Journal of Operational Research*, v. 289, n. 2, p. 399–415, 2021. ISSN 0377-2217. Citado 3 vezes nas páginas 22, 33, and 42.

IORI, M; LOCATELLI, M; MOREIRA, MCO; SILVEIRA, T et al. A mixed approach for pallet building problem with practical constraints. *LECTURE NOTES IN BUSINESS INFORMATION PROCESSING*, Springer Science and Business Media Deutschland GmbH, v. 417, p. 122–139, 2021. Citado 2 vezes nas páginas 13 and 16.

IORI, Manuel; LOCATELLI, Marco; MOREIRA, Mayron CO; SILVEIRA, Tiago. Reactive grasp-based algorithm for pallet building problem with visibility and contiguity constraints. In: SPRINGER. *International conference on computational logistics*. [S.l.], 2020. p. 651–665. Citado na página 16.

LAYEB, Safa Bhar; OMRI, Oussema. Solving the pallet loading problem with deep reinforcement learning. In: ____. *Handbook of Formal Optimization*. Singapore: Springer Nature Singapore, 2024. p. 807–825. Citado 2 vezes nas páginas 34 and 43.

MAGNANI, Matteo; FRANCHINI, Giorgia; DELL'AMICO, Mauro; ZANNI, Luca. A machine learning approach to speed up the solution of the distributor's pallet loading problem. In: SERGEYEV, Yaroslav D.; KVASOV, Dmitri E.; ASTORINO, Annabella (Ed.). *Numerical Computations: Theory and Algorithms*. Cham: Springer Nature Switzerland, 2025. p. 108–120. ISBN 978-3-031-81241-5. Citado 2 vezes nas páginas 34 and 43.

MARTINEZ-SYKORA, Antonio; ALVAREZ-VALDÉS, Ramón; BENNELL, Julia A; RUIZ, Rubén; TAMARIT, José Manuel. Matheuristics for the irregular bin packing problem with free rotations. *European Journal of Operational Research*, Elsevier, v. 258, n. 2, p. 440–455, 2017. Citado na página 16.

MAUNGMEESRI, Benchalak; THWE, Yamin. Heuristic solutions for organized ceramic bowl placement in manufacturing environments. *International Journal of Engineering Trends and Technology*, v. 72, n. 6, p. 19–28, 2024. Disponível em: <https://ijettjournal.org/Volume-72/Issue-6/IJETT-V72I6P103.pdf>. Citado na página 16.

MORABITO, Reinaldo; MORALES, Silvia Regina; WIDMER, Joao Alexandre. Loading optimization of palletized products on trucks. *Transportation Research Part E: Logistics and Transportation Review*, Elsevier, v. 36, n. 4, p. 285–296, 2000. Citado na página 16.

MUNGWATTANA, Anan; PIYACHAYAWAT, Thanakrit; JANSSENS, Gerrit K. A two-step evolutionary algorithm for the distributor's pallet loading problem with multi-size pallets. *Flexible Services and Manufacturing Journal*, v. 35, n. 4, p. 1256–1275, 2023. Citado 2 vezes nas páginas 21 and 23.

NASCIMENTO, Oliviana Xavier; QUEIROZ, Thiago Alves; JUNQUEIRA, Leonardo. Practical constraints in the container loading problem: Comprehensive formulations and exact algorithm. *Computers & Operations Research*, v. 128, p. 105186, 2021. ISSN 0305-0548. Citado 3 vezes nas páginas 22, 33, and 42.

PACKVOL. *Optimization Software for Load Planning*. 2024. <https://www.packvol.com/index.html>. Accessed: 2024-12-10. Citado 6 vezes nas páginas 20, 22, 29, 39, 42, and 48.

PEDREGOSA, Fabian; VAROQUAUX, Gaël; GRAMFORT, Alexandre; MICHEL, Vincent; THIRION, Bertrand; GRISEL, Olivier; BLONDEL, Mathieu; PRETTENHOFER, Peter; WEISS, Ron; DUBOURG, Vincent et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, JMLR. org, v. 12, p. 2825–2830, 2011. Citado 2 vezes nas páginas 25 and 38.

QUE, Quanqing; YANG, Fang; ZHANG, Defu. Solving 3D packing problem using transformer network and reinforcement learning. *Expert Systems with Applications*, v. 214, p. 119153, 2023. ISSN 0957-4174. Citado na página 23.

QUEIROZ, T. A.; MIYAZAWA, F. K.; WAKABAYASHI, Y. On the L-approach for generating unconstrained two-dimensional non-guillotine cutting patterns. *4OR*, v. 13, n. 2, p. 199–219, Jun 2015. Citado 3 vezes nas páginas 23, 34, and 43.

QUEIROZ, Thiago Alves; MUNDIM, Leandro Resende. Multiobjective pseudo-variable neighborhood descent for a bicriteria parallel machine scheduling problem with setup time. *International Transactions in Operational Research*, v. 27, n. 3, p. 1478–1500, 2020. Citado na página 31.

SALEM, Khadija Hadj; SILVA, Elsa; OLIVEIRA, José Fernando. Cutting and packing problems under uncertainty: literature review and classification framework. *International Transactions in Operational Research*, v. 30, n. 6, p. 3329–3360, 2023. Citado 2 vezes nas páginas 33 and 42.

SAXENA, Sharoon. *A Beginner's Guide to Random Forest Hyperparameter Tuning*. 2024. <https://www.analyticsvidhya.com/blog/2020/03/beginners-guide-random-forest-hyperparameter-tuning>. Accessed in: 2025-05-04. Citado na página 38.

SCHEITHAUER, Guntram. *Introduction to Cutting and Packing Optimization*: Problems, modeling approaches, solution methods. [S.l.]: Springer Cham, 2018. Citado na página 33.

SHARMA, Kalpit. *Mastering Random Forest Hyperparameter Tuning for Enhanced Machine Learning Models*. 2024. <https://medium.com/@kalpit.sharma/mastering-random-forest-hyperparameter-tuning-for-enhanced-machine-learning-models-2d1a8c6
Accessed in: 2025-05-04. Citado na página 38.

SHI, Yu; KE, Guolin; SOUKHAVONG, Damien; LAMB, James; MENG, Qi; FINLEY, Thomas; WANG, Taifeng; CHEN, Wei; MA, Weidong; YE, Qiwei; LIU, Tie-Yan; TITOV, Nikita; CORTES, David. *lightgbm: Light Gradient Boosting Machine*. [S.l.], 2024. R package version 4.5.0.99. Disponível em: <https://github.com/Microsoft/LightGBM>. Citado 2 vezes nas páginas 25 and 36.

SILVA, Elsa; OLIVEIRA, José F; WÄSCHER, Gerhard. The pallet loading problem: a review of solution methods and computational experiments. *International Transactions in Operational Research*, Wiley Online Library, v. 23, n. 1-2, p. 147–172, 2016. Citado 4 vezes nas páginas 22, 33, 42, and 43.

SILVEIRA, Tiago. Algorithms for cutting and packing problems. Università degli studi di Parma. Dipartimento di Ingegneria e architettura, 2022. Citado 3 vezes nas páginas 9, 13, and 14.

SWEENEY, Paul E.; PATERNOSTER, Elizabeth Ridenour. Cutting and packing problems: A categorized, application-orientated research bibliography. *Journal of the Operational Research Society*, Taylor & Francis, v. 43, n. 7, p. 691–706, 1992. Citado na página 22.

TACCINI, Marco; OLIVEIRA, Matheus Aguilar de; SANTOS, André Gustavo dos; QUEIROZ, Thiago Alves de; IORI, Manuel. A hybrid approach for pallet loading in ceramic tile industry. In: *Proceedings of the International Conference on Optimization and Learning (OLA2025)*. [S.l.: s.n.], 2025. p. 1–12. Citado 7 vezes nas páginas 35, 36, 40, 42, 45, 46, and 47.

TERNO, Johannes; SCHEITHAUER, Guntram; SOMMERWEIß, Uta; RIEHME, Jan. An efficient approach for the multi-pallet loading problem. *European Journal of Operational Research*, v. 123, n. 2, p. 372–381, 2000. ISSN 0377-2217. Citado 2 vezes nas páginas 23 and 34.

THEODOROPOULOS, Nikolaos; ANDRONAS, Dionisis; KAMPOURAKIS, Emmanouil; MAKRIS, Sotiris. Automated pallet loading of irregularly shaped objects: A deep reinforcement learning and multi-criteria optimization method. *Journal of Manufacturing Systems*, v. 80, p. 916–932, 2025. Citado na página 34.

WÄSCHER, Gerhard; HAUßNER, Heike; SCHUMANN, Holger. An improved typology of cutting and packing problems. *European Journal of Operational Research*, v. 183, n. 3, p. 1109–1130, 2007. ISSN 0377-2217. Citado na página 42.

WU, Wenjie; FAN, Changjun; HUANG, Jincai; LIU, Zhong; YAN, Junchi. Machine learning for the multi-dimensional bin packing problem: Literature review and empirical evaluation. *arXiv preprint arXiv:2312.08103*, 2023. Citado 2 vezes nas páginas 15 and 16.

YAGIURA, Mutsunori; IMAHORI, Shinji; HU, Yannan. *Cutting and Packing Problems*: From the perspective of combinatorial optimization. [S.l.]: Springer Tokyo, 2025. Citado 2 vezes nas páginas 22 and 33.

YAO, Shaowen; ZHANG, Tai; ZHANG, Hao; QIU, Jian; LENG, Jiewu; LIU, Qiang; WEI, Lijun. The semi-online robotic pallet loading problem. *Computers & Operations Research*, v. 174, p. 106889, 2025. Citado na página 43.

ZOUININA, Sarah. *A Deep Dive into LightGBM: How to Choose and Tune Parameters*. 2024. <https://medium.com/@sarahzouinina/a-deep-dive-into-lightgbm-how-to-choose-and-tune-parameters-7c584945842e>. Accessed in: 2025-05-17. Citado na página 38.

ZUO, Qiruyi; LIU, Xinglu; CHAN, Wai Kin Victor. A constructive heuristic algorithm for 3d bin packing of irregular shaped items. In: SPRINGER. *INFORMS International Conference on Service Science*. [S.l.], 2022. p. 393–406. Citado na página 13.