

The 2-D Orthogonal Packing Problem with Multiple Levels of Prioritization: A Spatial Optimization Perspective

William K. Kirschenman^{a,*}, H. Sebastian Heese^b, Michael G. Kay^a, Russell E. King^{a,c},
Brandon M. McConnell^{a,c}

^a*Department of Industrial and Systems Engineering, North Carolina State University, Campus Box 7906, Raleigh, 27695-7906, NC, United States*

^b*Poole College of Management, North Carolina State University, 2300 Nelson Hall, Raleigh, 27695-7229, NC, United States*

^c*Center for Additive Manufacturing and Logistics (CAMAL), North Carolina State University, 915 Partners Way, Box 7906, Raleigh, 27695-7906, NC, United States*

Abstract

This paper addresses two-dimensional orthogonal packing within a confined space, integrating bin packing principles with facility layout concepts to address scenarios in which items must not only fit but also be arranged according to spatial priorities. We embed a prioritization matrix into the bin packing framework, enabling items to be clustered with one another or pulled toward certain bin access points based on assigned priority weights. Unlike traditional bin packing, which often minimizes bin count or unused space, our approach balances proximity to bin access points and adjacency among functionally related items already assigned to a given bin, extending the utility of bin packing to applications requiring more nuanced layout preferences.

We introduce a single mixed-integer linear programming (MILP) model and a complementary sliding-window matheuristic that scales effectively to larger problem instances. Numerical experiments illustrate that this matheuristic approach consistently outperforms a direct MILP solve with a commercial solver in both runtime and solution quality, and also performs best among the adapted heuristic and metaheuristic alternatives considered in our study. This computational study underscores the flexibility and effectiveness of embedding multi-level priorities into orthogonal packing.

Keywords: Packing, Facilities planning and design, Combinatorial optimization, Integer programming, Logistics

1. Introduction

The problem of efficiently arranging objects within a confined space is a fundamental challenge in operations research, commonly studied as a bin packing problem. Classical formulations of bin packing primarily focus on minimizing either the number of bins used or the amount of unused space within each bin. However, many real-world applications impose additional requirements beyond spatial efficiency. In certain contexts, the placement of items relative to

*Corresponding author at: Department of Industrial and Systems Engineering, North Carolina State University, Campus Box 7906, Raleigh, NC 27695-7906, United States. Email: wkkirsch@ncsu.edu

Email addresses: wkkirsch@ncsu.edu (William K. Kirschenman), hseese@ncsu.edu (H. Sebastian Heese), kay@ncsu.edu (Michael G. Kay), king@ncsu.edu (Russell E. King), mcconnell@ncsu.edu (Brandon M. McConnell)

other objects or specific reference points within the bin is equally important. In such settings, classical bin packing’s emphasis on space utilization alone becomes insufficient, as operational priorities must also be addressed. These priorities may arise from access sequencing constraints, the need to maintain proximity between functionally related items, or requirements to position high-priority items near designated locations, such as loading doors or off-loading points.

A related stream of research in facility layout planning addresses the optimal arrangement of departments, machines, or storage units within a defined space to minimize transportation costs, congestion, or material handling inefficiencies. These problems often consider adjacency-based cost structures, where minimizing weighted distances between specific objects is a primary objective. By contrast, two-dimensional bin packing formulations primarily focus on ensuring feasibility under geometric constraints, such as non-overlapping placement and orientation restrictions. This research integrates elements of both domains by incorporating a prioritization-weighted cost function into a classical bin packing framework. Specifically, we introduce a prioritization-weight matrix that captures item-to-item and item-to-bin access point placement preferences, thereby extending the bin packing problem to include spatial considerations traditionally found in facility layout optimization.

To illustrate the core challenge and the impact of our approach, consider a small example involving six items belonging to three distinct groups, each with varying priorities for access point proximity (detailed in Table 1). Items are assigned two types of priority values: *Item Priority* reflects an item’s relative importance within its group (lower values indicate higher priority), while *Group Priority* reflects the group’s overall importance relative to other groups (again, lower values indicate higher priority). Together, these determine each item’s global priority ranking for placement near the bin access point, as detailed in Section 3.2. These items are to be placed within a rectangular bin of 700 length and 350 width. Items within a group should ideally be kept close together (group cohesion), while high-priority items should be near a designated bin access point (here, the center-left edge of the bin).

Table 1: Example instance illustrating groups and priorities.

Item	Length	Width	Group	Item Priority	Group Priority
1	258	131	2	1	1
2	312	137	1	1	2
3	210	100	1	2	2
4	202	96	1	3	2
5	341	144	3	1	3
6	136	86	3	2	3

Figure 1 contrasts two layouts for this instance, each solved to optimality under different objectives using the mathematical framework detailed later in this paper. The left panel shows the optimal layout generated when the objective is solely to minimize the total weighted rectilinear distance from each item’s centroid to the access point (white circle). The weights guiding this proximity objective are derived from a combination of the ‘Group Priority’ and ‘Item Priority’ values listed in Table 1, reflecting each item’s overall importance for being near the access point (the specific calculation method is described in Section 3.2). While the resulting arrangement might not perfectly match intuitive visual placement based purely on rank (e.g., item 1 is not

the absolute closest), this layout represents the mathematical optimum found by the solver when only these access point proximity weights are considered. Specifically, the solver determines that positioning items with overall priority ranks 3 and 6 closest to the access point yields a lower total weighted distance than placing item 1 directly adjacent, as the combined weighted benefit of optimally positioning items with these particular priority ranks outweighs having just the single highest-priority item at minimum distance.

The right panel shows the optimal layout produced by our full prioritization approach. This layout minimizes the same weighted distance to the access point as the left panel, but also incorporates an incentive for items belonging to the same group to be placed closer together. This ‘group cohesion’ aspect is driven by relative priorities among items within the same group, based on the ‘Item Priority’ values from Table 1 (as detailed in Section 3.2). The resulting layout demonstrates how our approach encourages items within the same group (indicated by shading: white=highest priority group, gray=medium, black=lowest) to cluster together, while still positioning high-priority items advantageously near the center-left access point. This simple example highlights how incorporating group cohesion via our prioritization matrix leads to operationally superior layouts while still considering access priorities.




Figure 1: Comparison of layouts for the 6-item example. Left: Layout optimizing only for access point proximity (No Groups). Right: Layout optimizing for both access point proximity and group cohesion (With Groups). Items are shaded by group (white, gray, black) and labeled by overall priority rank (1 = highest, 6 = lowest). The access point (center-left edge) is indicated by the white circle.

Prioritization-based packing has wide-ranging applications in commercial settings such as warehouse operations and cargo stowage. In these contexts, a single shipment or customer order may involve multiple items that form a distinct group, which needs to be kept cohesive to simplify downstream handling. Within each group, certain items might also require a specific sequence for order picking or unloading, such as fragile goods that must be accessible first or last. Finally, across multiple groups of items, some groups can be more time-critical than others (e.g., in a vehicle routing solution for delivery truck shipments), creating a tiered prioritization among different groups. Modern e-commerce fulfillment centers exemplify this multi-level prioritization. Optimizing item placement within storage bins or shipping containers is critical to reducing retrieval times and improving operational efficiency. In its warehouse operations, UPS conducts detailed inventory slotting analysis to determine the best placement of items for selection, replenishment, and workload balance, while constantly evaluating facility layout designs (UPS, 2023). Both UPS and DHL increasingly deploy automation and robotics technologies in

warehouses to increase process efficiency (DHL, 2023), yet must also juggle the trade-off between placing items near loading points for faster dispatch and preserving group cohesion. Similar challenges arise in freight transportation, where cargo must be positioned strategically to minimize handling time and the number of trips necessary to move goods. FedEx uses truck-loading robots to help determine ideal package layouts to make the most of available space and subsequently load the packages, saving time and effort (Dexterity AI, 2023). A similar arrangement-stage decision also arises in commercial parcel delivery once packages have already been assigned to a vehicle. UPS has described package cars being loaded in delivery order with assigned shelf positions, and Amazon has described Vision-Assisted Package Retrieval (VAPR) technology designed to reduce stop-by-stop package search and retrieval inside the van (Witt, 2005; Levine, 2024). These examples illustrate why sequence-aware access and retrieval can matter even after assignment and capacity decisions have already been made. The transportation industry increasingly employs automated load-planning systems that optimize for weight distribution and space efficiency, yet many of these systems lack mechanisms to account for adjacency-based priorities. In warehouse operations, there has been little research studying the impact of human factors in order-picking efficiency (Grosse et al., 2017). Analyzing optimal storage positions of items could help derive guidelines that balance ergonomics and performance goals. Taken together, these examples underscore how packing decisions often require careful sequencing across multiple groups, proper ordering of items within each group, and close proximity among group members—three interconnected levels of priority that go beyond simple space minimization.

A primary motivation for this study arises in military logistics, where many of the same multi-level considerations apply. Combat loading in a contested environment demands rapid and prioritized off-loading, but also close grouping of functionally related assets. As outlined in Joint Publication 3-02 (Joint Chiefs of Staff, 2019) and historical embarkation planning doctrine, the arrangement of cargo on transport platforms must align with the operational priority of deployment on the ground. Optimization of military logistics scheduling and resource allocation continues to be an active research area (Brown et al., 2017). For example, artillery assets may need to be off-loaded first while remaining in close proximity to supporting vehicles. Existing load planning tools, such as the Integrated Computerized Deployment System (ICODES) Single Load Planner, ensure feasibility in load planning by performing structural checks, such as weight distribution and stability constraints (Goodman and Pohl, 2003). However, these systems typically require extensive manual intervention to satisfy operational priorities, such as enforcing a specific off-loading sequence or maintaining adjacency between mission-critical assets. Without a formal mechanism to incorporate these spatial priorities, planners must resort to trial-and-error efforts that often fail to capture all desired operational nuances (Kirschenman et al., 2024). In particular, groups of functionally related equipment must remain cohesive to rapidly organize upon off-loading, the equipment may have a desired group-internal sequence (e.g., combat vehicles before support equipment), and certain groups may need to off-load before others (e.g., to establish a defensive perimeter at a contested beach or port). Balancing these layered requirements—with compact placement and quick access to off-loading or access points—demands an approach that integrates all three levels of priority. Historical military operations underscore the criticality of these principles. For instance, an after-action report regarding embarkation sched-

ules for the Inchon Landing campaign highlighted that in challenging over-the-beach operations with limited off-loading facilities and restricting conditions, “debarkation priorities might well be considered more extensively” and stressed that “it is paramount that the combat elements and their combat material go over the beaches together” (X Corps Headquarters, 1951). This directly mirrors the need to manage both item/group off-loading priorities and maintain group cohesion, which are central tenets of our approach.

In this paper, we concentrate on a single-bin environment rather than addressing which bin to choose or whether items fit in multiple bins. We assume that an upstream process has already determined a suitable bin; if items do not fit, the problem is simply infeasible. We further assume that the bin’s access point width is at least as large as the smallest dimension of any item to be packed, ensuring that all items assigned to the bin can physically pass through the access point. Our primary objective, therefore, is to model and optimize how items are arranged inside the chosen bin under multiple spatial priorities. Importantly, this decision stage does not assign items to bins, minimize bin count, or minimize unused space; it takes the upstream bin-selection and feasibility determination as given and focuses on the internal arrangement within the selected bin. In the baseline prioritization-matrix construction used in this paper, access-point proximity is the dominant component of the diagonal weights, which encourages dense packings adjacent to the access point, while the off-diagonal weights preserve group cohesion. This assumption is particularly well-suited to military combat loading, where vessels are typically loaded with deliberate operational excess capacity (United States Navy, 2021), and the critical challenge is not whether equipment fits, but whether it can be accessed in the required priority order during contested off-loading operations. While our approach naturally extends to multi-bin or multi-vessel selection in future work, here we highlight the fundamental challenge of prioritization-aware packing within a single bin.

The core contribution of this work is a flexible framework for encoding spatial preferences in two-dimensional bin packing through a prioritization matrix. Unlike standard bin packing formulations where the objective is purely geometric, our framework enables practitioners to express different types of spatial relationships: how strongly items should be placed near a designated access point, and how strongly certain items should be clustered together. By adjusting how the matrix encodes these preferences, different operational scenarios can be modeled without reformulating the underlying optimization. In this paper, we study a specific instantiation that balances item-to-access-point proximity with group cohesion, but the framework generalizes to other configurations—for example, emphasizing clustering over access-point proximity, treating all access-point priorities equally, or disabling item-to-item preferences entirely. We demonstrate the flexibility of this framework by applying it to scenarios with multiple levels of prioritization.

The underlying non-overlap and feasibility constraints remain those of a standard orthogonal packing model. The methodological novelty lies in embedding the prioritization matrix directly into the objective function so that item-to-access-point and item-to-item preferences are represented within a single hybrid packing–layout problem, which in turn motivates the tailored exact and matheuristic solution approaches and lower bounds developed in this paper.

The remainder of this paper is structured as follows. Section 2 reviews the relevant literature, covering classical bin packing formulations, facility layout models, and prioritization-based

packing approaches. Section 3 presents the mathematical formulation of the proposed model, adapting prior three-dimensional bin packing research to our two-dimensional prioritization-aware setting. Section 4 discusses solution approaches and lower bounding techniques, while Section 5 presents a computational performance evaluation of these approaches, including an equal-budget comparison against adapted heuristic and metaheuristic alternatives. Finally, Section 6 offers managerial insights, outlines potential extensions, and discusses how the proposed framework can be adapted for broader applications in transportation and logistics.

By integrating prioritization considerations into the orthogonal packing problem objective, this study provides a structured approach to optimizing item placement in scenarios where spatial feasibility and operational priority must be jointly considered.

2. Literature Review

This section summarizes the key bodies of literature relevant to our work. We begin with classical bin packing in Section 2.1, then examine the facility layout literature in Section 2.2, and finally review studies that incorporate prioritization or weighted objectives into packing problems in Section 2.3.

2.1. Bin Packing Literature

Bin packing problems have been widely studied in combinatorial optimization, with applications in logistics, manufacturing, and storage optimization. The fundamental problem consists of packing a set of items into one or more bins while ensuring non-overlapping placement and minimizing either the number of bins used or the amount of unused space in an input minimization context (Wäscher et al., 2007). Even in its simplest one-dimensional form, bin packing is NP-hard in the strong sense (Garey and Johnson, 1979). The complexity increases significantly in higher-dimensional cases where width, height, orientation, and other spatial constraints must be considered (Bortfeldt and Wäscher, 2013).

In the two-dimensional variant (2D-BPP), items must be arranged within rectangular bin boundaries while preventing overlaps, optionally allowing a restricted set of rotations (Lodi et al., 2002). Classical approaches to the 2D-BPP emphasize optimizing space utilization, with many studies focusing on heuristics and exact algorithms for minimizing the number of bins required (Christensen et al., 2017; Lodi et al., 2002; Iori et al., 2021). The geometric nature of the problem has led to a variety of solution methods, including mixed-integer linear programming (MILP) formulations (Castro and Oliveira, 2011; Lodi et al., 2004), heuristic placement strategies (Chazelle, 1983; Coffman et al., 1999), and metaheuristic techniques such as tabu search, Greedy Randomized Adaptive Search Procedure (GRASP), or genetic algorithms (Hopper and Turton, 2001; Alvarez-Valdés et al., 2008, 2013; Lodi et al., 1999).

While much of the bin packing literature focuses on minimizing bin count or maximizing space usage, certain variants introduce additional considerations, such as weight balancing, load stability, sequence-based objectives (Li and Zhang, 2018), and spatial arrangement with loading sequence optimization in container terminal operations (Zhu et al., 2019). Recent developments include the colored bin packing problem, where items have color attributes and no two items of the same color may be placed consecutively within any bin (Borges et al., 2024). In particular,

residual bin packing problems, which involve strongly heterogeneous items and bins, have received comparatively less attention (Wäscher et al., 2007). The problem considered in this work can be classified as a single-bin residual bin packing problem, where all items must be packed into a single, predefined bin, while the objective extends beyond space utilization to include prioritization-weighted distances between items and an access point on the bin’s edge.

To model the necessary geometric constraints, we follow prior three-dimensional bin packing formulations by Paquay et al. (2016) and Fontaine and Minner (2023), adapting them to a two-dimensional setting. Fontaine and Minner (2023) brought further computational enhancements, such as distance normalization and specialized symmetry-breaking constraints, but these improvements are tailored to standard bin packing objectives and do not readily apply to our distance-based prioritization. We retain the core non-overlapping and orientation constraints from these studies while introducing a prioritization-weighted distance objective to handle various types of spatial priorities—such as pulling certain items closer to a bin access point, or encouraging adjacency among high-priority pairs of items—even though the underlying bin packing constraints remain similar to classical 2D-BPP approaches.

2.2. Facility Layout Planning Literature

Facility layout planning (FLP) focuses on the spatial arrangement of departments, workstations, or operational units within a bounded facility to optimize metrics such as material handling costs or adjacency relationships. Traditionally, FLP has been formulated as a combinatorial optimization problem where the objective is to minimize total weighted rectilinear distance while ensuring non-overlapping orthogonal arrangement (Meller et al., 1998; Pérez-Gosende et al., 2021). Classical formulations broadly categorize FLP into discrete or continuous models, distinguished by whether department placements are confined to predefined grid positions or allowed to move freely within the facility space (Drira et al., 2007).

One of the earliest mathematical formulations of FLP is the Quadratic Assignment Problem (QAP) (Koopmans and Beckmann, 1957), which assigns departments to fixed grid locations. While subsequent work extended QAP formulations with improved decomposition methods (Bazaraa and Sherali, 1980), QAP-based methods are often computationally prohibitive for large problems, leading to various heuristic and metaheuristic approaches including construction and local search heuristics (Drezner, 1987) and GRASP-based approaches (Cravo and Amaral, 2019). More flexible representations, such as the Unequal-Area Facility Layout Problem (UA-FLP), allow departments to have non-uniform sizes and aspect ratios (Bozer and Meller, 1997).

These UA-FLP formulations share geometric constraints with the 2D-BPP, particularly where departments or items must fit within a finite space. However, UA-FLP typically focuses on minimizing total distance-based cost (e.g., flow costs between departments). Most facility layout models do not explicitly capture multi-tiered prioritization: for instance, certain items might need strict adjacency or sequential access—scenarios found in contested military settings. Studies on multi-floor layout problems address both horizontal and vertical adjacency (Ahmadi et al., 2017), but these typically focus on assigning departments to floors rather than embedding more granular priority relationships among individual departments. Similarly, the literature on warehouse design sometimes accounts for pick frequency or accessibility (Gu et al., 2007), yet rarely incorporates a flexible adjacency matrix that assigns varying weights to pairwise item and

item-to-access-point proximities to assist with actual item placement in the storage location assignment problem. This gap leaves many real-world applications outside the scope of traditional facility layout approaches.

Accordingly, bridging 2D-BPP constraints with distance-centric FLP objectives offers a way to handle multi-level prioritization. In our work, we place a structured prioritization-weight matrix into a bin packing formulation, thus capturing adjacency-based costs akin to FLP while retaining classical constraints on item placement and orientation.

2.3. Spatial Optimization in Bin Packing and Related Problems

In many real-world applications, the placement of items within a bin is influenced by priorities beyond simple space utilization. In particular, item accessibility, load balancing, sequencing, and adjacency preferences introduce additional complexity into bin packing formulations. Unlike studies that consider prioritizing which bin an item is assigned to, the present work focuses on prioritizing the internal arrangement of items within the bin. These priority-driven ideas have been explored in multiple domains, including container loading, multi-drop delivery, and facility layout problems.

Container loading problems (CLP) provide one of the most direct extensions of bin packing in practical logistics, particularly in settings where item placement influences the efficiency of sequential unloading. In multi-drop container loading, for instance, cargo must be arranged to allow staged unloading at different destinations, ensuring that items for the first drop-off points are more accessible than those required later (Bischoff and Ratcliff, 1995). Several studies explicitly incorporate prioritization: Altarazi (2013) proposes a two-step heuristic for truck loading with priority weights based on demand, while do Nascimento et al. (2021) decompose the CLP into subproblems to deal with feasibility, ensuring a strict prioritization hierarchy for packing considerations similar to Vancroonenburg et al. (2014).

A broader review of container loading constraints highlights that prioritization is rarely explicitly modeled as an optimization objective. Bischoff and Ratcliff (1995) mention that relative priorities could be handled by adjustments of the coefficients in the objective function, but this only relates to the inclusion of an item in a solution — not relative distances between items or to a bin’s access point. A survey by Bortfeldt and Wäscher (2013) found that less than 2% of container loading papers at the time explicitly considered loading priorities in the design of algorithms for container loading. Filella et al. (2023) take an innovative approach by modeling unloading constraints as soft constraints, introducing penalties in the objective function and thus more flexibly optimizing both for space utilization and desired loading configurations. However, pairwise item interactions are not considered.

Another related stream of research emphasizes load balancing and stability in bin packing. Trivella and Pisinger (2016) introduce the load-balanced multi-dimensional bin packing problem, formulating the objective function to keep the center of mass of the loaded bins as close as possible to a desired barycenter location while also minimizing the number of total bins used. Erbayrak et al. (2021) extend this idea into a multi-objective approach that also promotes “family unity,” meaning items from the same product family end up in the same bin. However, such family unity does not ensure close adjacency within the bin, nor does it account for sequential ordering or varying strengths of cohesion within a family of items. By contrast, the method we

propose encourages more nuanced relationships: certain groups of items might need to cluster tightly, others might require moderate proximity, and some items might need to lie near a bin edge for rapid access rather than purely for space efficiency.

Overall, as far as we know, existing bin packing or container loading models do not incorporate a flexible distance-based prioritization matrix that simultaneously captures item-to-item and item-to-bin relationships. By integrating these priorities directly into the objective, our work unifies key ideas from both bin packing and facility layout planning. This provides a framework for important use cases such as warehouse design, multi-drop delivery scenarios, and military combat loading, where the relative positioning of items critically impacts retrieval efficiency and overall operational effectiveness.

3. The 2-D Orthogonal Packing Problem with Prioritization

3.1. Mathematical Formulation

We consider a single rectangular bin of fixed length L and width W , into which a set of n rectangular items $i \in I = \{1, 2, \dots, n\}$ must be placed. Each item i has length p_i and width q_i , potentially oriented in either dimension (rotated 90°). Assuming that an upstream process has already assigned the items to the bin and established that a feasible packing exists, we optimize their spatial arrangement within the bin using a prioritization-based objective that draws certain items together or pulls them toward a designated bin access point on the bin’s boundary. We formulate this as a MILP, using efficient formulations introduced by Paquay et al. (2016) and extended by Fontaine and Minner (2023), but simplifying to a single-bin, 2-D setting, with a customized distance-based objective. We highlight the Fontaine and Minner (2023) formulation’s parameter and variable naming, as the naming differs slightly between the two papers. Table 2 summarizes the notation used in the formulation.

Throughout this paper, all item dimensions (p_i, q_i) and bin dimensions (L, W) are positive integers. Consistent with the integer-coordinate convention inherited from Paquay et al. (2016) and Fontaine and Minner (2023) and used in our instance generation (Section 5.1.1), the “+1” terms in constraints (6) and (8) encode strict separation between item pairs.

The 2-D orthogonal packing problem with prioritization is subject to fundamental constraints: (1) all items must be placed entirely within the boundaries of the assigned bins; (2) no two items assigned to the same bin may overlap in any dimension; and (3) items must be placed orthogonally to each other, such that each dimension of an item is parallel to one of the bin dimensions. While additional constraints can be incorporated (e.g., item rotation restriction or load balancing), they are excluded from this study to focus on a generalized formulation of the problem.

We define a prioritization matrix π_{ik} for each item $i, k \in I$ where $i \leq k$. We encode a prioritization weight structure within π_{ik} that allows for different packing configurations. We compute π_{ik} ahead of time and include it as a parameter to our model. For π_{ik} where $i < k$, π_{ik} encodes the priority weight for items i and k to be close to each other. For π_{ii} , we encode the priority weight for placing item i close to the bin access point. Let (x^o, y^o) denote the coordinates of the bin access point, which is treated as a known parameter of the model.

Each item $i \in I$ has a continuous x_i (y_i) variable designating the coordinate position of the left (bottom) side of the item; a continuous x_i^r (y_i^r) variable designates the coordinate position of the right (upper) side of the item. The binary variables a_{ik}^x (a_{ik}^y) are equal to 1 if item k is placed entirely to the left (bottom) of item i in the x - (y -) dimension, and 0 otherwise. The binary orientation variables $t_{i,c,d}$ employ a dimension-mapping approach where $c \in \{1, 2\}$ indexes the bin dimensions (x, y) and $d \in \{1, 2\}$ indexes the item dimensions (length, width). Here, $t_{i,c,d} = 1$ if item dimension d maps to bin dimension c , and $t_{i,c,d} = 0$ otherwise. While four binary variables encode two orientations (0° and 90° rotation), the formulation ensures each item dimension maps to exactly one bin dimension, providing a general framework for orientation control. We introduce auxiliary decision variables ρ_{ik}^x and ρ_{ik}^y to represent rectilinear distances, whether between centroids of each item i and k where $i < k$, or between each item i and the associated bin access point (x^o, y^o) when $i = k$. We use these rectilinear distances in the objective function to spatially optimize the placement of items within the bin.

Table 2: Notation used in the mathematical formulation.

Category	Notation	Definition
Sets/indices	$I = \{1, 2, \dots, n\}$	Set of items assigned to the bin
	$i, k \in I$	Item indices
	$c \in \{1, 2\}$	Bin-dimension index for x and y
	$d \in \{1, 2\}$	Item-dimension index for length and width
Parameters	n	Number of items
	L, W	Bin length and width
	p_i, q_i	Length and width of item i
	(x^o, y^o)	Coordinates of the bin access point
	π_{ik}	Prioritization weight for item pair i, k ; diagonal entries weight access-point proximity and off-diagonal entries weight item-to-item proximity
Decision variables	x_i, y_i	Coordinates of the left and bottom sides of item i
	x_i^r, y_i^r	Coordinates of the right and upper sides of item i
	a_{ik}^x, a_{ik}^y	Binary relative-placement variables for separating item k from item i in the x - or y -dimension
	$t_{i,c,d}$	Binary orientation variable indicating whether item dimension d maps to bin dimension c
Auxiliary variables	ρ_{ik}^x, ρ_{ik}^y	Rectilinear-distance components for item pair i, k or item i and the access point when $i = k$

The mathematical model is then defined as follows:

$$\min \sum_{i \in I} \sum_{k \in I, i < k} \pi_{ik} (\rho_{ik}^x + \rho_{ik}^y) \quad (1)$$

$$\text{s.t. } x_i^r - x_i = t_{i,1,1}p_i + t_{i,1,2}q_i \quad \forall i \in I \quad (2)$$

$$y_i^r - y_i = t_{i,2,1}p_i + t_{i,2,2}q_i \quad \forall i \in I \quad (3)$$

$$a_{ik}^x + a_{ki}^x + a_{ik}^y + a_{ki}^y \geq 1 \quad \forall i, k \in I, i < k \quad (4)$$

$$x_k^r \leq x_i + (1 - a_{ik}^x)L \quad \forall i, k \in I \quad (5)$$

$$x_i + 1 \leq x_k^r + a_{ik}^x L \quad \forall i, k \in I \quad (6)$$

$$y_k^r \leq y_i + (1 - a_{ik}^y)W \quad \forall i, k \in I \quad (7)$$

$$y_i + 1 \leq y_k^r + a_{ik}^y W \quad \forall i, k \in I \quad (8)$$

$$\sum_{c \in \{1,2\}} t_{i,c,d} = 1 \quad \forall i \in I, d \in \{1,2\} \quad (9)$$

$$\sum_{d \in \{1,2\}} t_{i,c,d} = 1 \quad \forall i \in I, c \in \{1,2\} \quad (10)$$

$$\frac{x_i + x_i^r}{2} - \frac{x_k + x_k^r}{2} \leq \rho_{ik}^x \quad \forall i, k \in I, i < k \quad (11)$$

$$\frac{x_k + x_k^r}{2} - \frac{x_i + x_i^r}{2} \leq \rho_{ik}^x \quad \forall i, k \in I, i < k \quad (12)$$

$$\frac{y_i + y_i^r}{2} - \frac{y_k + y_k^r}{2} \leq \rho_{ik}^y \quad \forall i, k \in I, i < k \quad (13)$$

$$\frac{y_k + y_k^r}{2} - \frac{y_i + y_i^r}{2} \leq \rho_{ik}^y \quad \forall i, k \in I, i < k \quad (14)$$

$$\frac{x_i + x_i^r}{2} - x^o \leq \rho_{ii}^x \quad \forall i \in I \quad (15)$$

$$x^o - \frac{x_i + x_i^r}{2} \leq \rho_{ii}^x \quad \forall i \in I \quad (16)$$

$$\frac{y_i + y_i^r}{2} - y^o \leq \rho_{ii}^y \quad \forall i \in I \quad (17)$$

$$y^o - \frac{y_i + y_i^r}{2} \leq \rho_{ii}^y \quad \forall i \in I \quad (18)$$

$$t_{i,c,d}, a_{ik}^x, a_{ik}^y \in \{0,1\} \quad \forall i, k \in I, c \in \{1,2\}, d \in \{1,2\} \quad (19)$$

$$0 \leq x_i \leq x_i^r \leq L \quad \forall i \in I \quad (20)$$

$$0 \leq y_i \leq y_i^r \leq W \quad \forall i \in I \quad (21)$$

$$\rho_{ik}^x, \rho_{ik}^y \geq 0 \quad \forall i, k \in I, i \leq k \quad (22)$$

We refer to this formulation as the 2-D Orthogonal Packing Problem with Prioritization (2DOPP-P).

The objective (1) sums the weighted rectilinear distances among items or between each item and the bin's access point, where weights come from π_{ik} . This structure pulls specified items closer together when their pairwise importance warrants adjacency, while simultaneously drawing each item toward the bin access point according to its individual priority level. For each item, Constraints (2) and (3) collectively ensure each item's coordinates reflects the actual item dimensions given the chosen item orientation. Constraint (4) ensures no overlapping of items i and k . Constraints (5)–(8) work together with Constraint (4) to enforce non-overlap. Constraint (4) requires that at least one of the four binary variables $a_{ik}^x, a_{ki}^x, a_{ik}^y, a_{ki}^y$ equals 1, so a separating relation must hold in at least one direction. When $a_{ik}^x = 1$, Constraint (5) forces $x_k^r \leq x_i$, meaning item k lies entirely to the left of item i . When $a_{ik}^x = 0$, Constraint (6) instead rules out item k lying entirely to the left of item i by requiring $x_i + 1 \leq x_k^r$ under the same unit-grid convention; by itself, it does not force item i to lie entirely to the left of item k . Constraints (7)–(8) operate analogously in the y -dimension, so overlap may occur in one dimension but not in both simultaneously. Constraints (9) and (10) ensure that each item's dimension is aligned to one of the bin's dimensions, effectively allowing for 90° rotation. Constraints (11)–(14) define ρ_{ik}^x and ρ_{ik}^y as absolute rectilinear distances between the centroids

of items i and k , and (15)–(18) likewise measure the distances ρ_{ii}^x and ρ_{ii}^y between the centroid of each item i and the access point (x^o, y^o) . Constraints (19)–(22) establish decision variable domains.

Although many bin packing and facility layout formulations incorporate explicit symmetry-breaking constraints (e.g., forcing an item to one half of the bin), we forgo such constraints in this work. In most cases, the bin’s access point already provides sufficient asymmetry unless it lies exactly on a midline, and even then, small- ε offsets for a single item can locally break symmetry without unduly restricting the feasible region. Hence, we accept the small risk of mirrored layouts in exchange for a simpler model that preserves advantageous placements along the midline when they are beneficial.

Even though our formulation reduces the three-dimensional model from Fontaine and Minner (2023) into two dimensions, the introduction of the prioritization structure in our objective function significantly complicates the optimization. Specifically, the interplay between group-level priorities, item-level sequencing preferences, and spatial proximity requirements leads to complex trade-offs. These trade-offs prevent the use of traditional relaxations and drastically increase computational difficulty, making even moderate-sized problem instances challenging or impossible to solve to optimality within reasonable time limits.

3.2. Constructing the Prioritization Matrix

In many operational scenarios, there is a need to place higher-priority items close to a bin access point for rapid retrieval or dissemination, while also grouping related items together. To reflect these goals, we use a prioritization matrix π that encodes both group cohesion and access-point proximity. For the model to remain valid more generally, any prioritization matrix must satisfy basic requirements including non-negativity of weights ($\pi_{ik} \geq 0$) and appropriate scaling to ensure meaningful objective function behavior. Minor adjustments to π can accommodate certain loading or stowage contexts with different adjacency or distance requirements.

Each diagonal entry π_{ii} is an inverted global priority level, capturing how strongly item i should be placed near the bin’s access point, while each off-diagonal entry π_{ik} (for $i < k$) encodes how strongly items in the same group should be drawn together based on their relative priorities. This single matrix π thus unifies proximity to the access point and group cohesion under one set of objective coefficients and can be recalculated whenever group structures or item-level priorities change. The construction method we present below represents one of potentially many approaches for building a prioritization matrix that could be tailored to achieve various operational goals and objectives. To illustrate this construction, using the data for the example instance shown in Table 1, we present the process applied in our computational study. In the baseline construction used here, assigning unique global ranks and then inverting them yields distinct diagonal priorities π_{ii} . The deterministic lower bound in Section 4.2.1, however, only requires nonnegative prioritization weights, so ties in the diagonal priorities do not affect its validity. Furthermore, this construction emphasizes access-point proximity within the objective function itself, primarily through the relatively larger magnitude of the diagonal π_{ii} entries compared to the off-diagonal π_{ik} entries.

Overview of Key Elements. π_{ii} (diagonal entries) represent how strongly item i should be placed near the bin’s access point. A higher diagonal value indicates a stronger pull toward that location. π_{ik} for $i < k$ (off-diagonal entries) represent how strongly items i and k should be close to each other. In the baseline structure presented here, $\pi_{ik} > 0$ only if i and k belong to the same group.

Step-by-Step Construction. We first sort all groups in ascending order of group-level priority, then sort the items in each group by ascending item-level priority. In this scheme, an item labeled with a smaller numeric value is considered higher priority. This sorting yields a provisional rank for each item, which we then reverse—that is, we flip the ordered list of raw ranks $(1, 2, \dots, |I|)$ to $(|I|, |I| - 1, \dots, 1)$ while keeping each item’s index attached to its position—so that the item with priority level 1 (i.e., the highest-priority item overall) ends up with the largest rank value. Concretely, if each item i is assigned a raw rank $\alpha[i]$ such that $\alpha[i] = 1$ corresponds to the highest-priority item and $\alpha[i] = 2$ the next, and so on, we set $\alpha^{\text{rev}}[i] = |I| - \alpha[i] + 1$, making $\alpha^{\text{rev}}[i]$ larger for items of higher priority. The reversed rank $\alpha^{\text{rev}}[i]$ becomes the diagonal entry π_{ii} , thereby assigning stronger pull toward the bin’s access point for items with higher overall priority.

Next, we define off-diagonal entries π_{ik} for pairs (i, k) belonging to the same group. Let each item i have an integer item-level priority within its group (again, a smaller number indicates higher priority), and define Δ_{ik} to be the absolute difference between those two item-level priorities. We identify G_{max} as the number of items in the largest group (i.e., the group containing the most items) across the entire problem instance. For any two items i and k in the same group, we then set $\pi_{ik} = \frac{G_{\text{max}} - \Delta_{ik}}{G_{\text{max}}}$ which is always strictly between 0 and 1 for nonzero Δ_{ik} . Two items whose item-level priorities differ only slightly thus receive a larger off-diagonal weight, while items farther apart in priority receive a smaller weight. By using the same denominator G_{max} for all groups, we standardize the off-diagonal scale throughout the instance. Items in different groups simply receive $\pi_{ik} = 0$, indicating no built-in incentive for them to be close. Note that assigning integer ranks (≥ 1) to π_{ii} and fractional values (< 1) to π_{ik} inherently ensures $\pi_{ii} > \pi_{ik}$ for all i, k where $i < k$. This deliberate scaling maintains a strong emphasis on access-point proximity, promoting dense configurations near the access point, while the smaller π_{ik} weights encourage group cohesion without dominating the primary placement objective.

Illustrative Example Data. The data used to generate the prioritization matrix for the illustrative example is shown in Table 1. This involves six items across three groups with assigned item and group priorities. Using this information, we obtain the prioritization matrix π_{ik} shown in Table 3, where the diagonal entries π_{ii} (in bold) reflect global item-level ranks derived from both item and group priorities, while off-diagonal entries encode intra-group closeness based on the formula above (using $G_{\text{max}} = 3$ for this instance).

Higher-priority items (diagonal entries) tend to appear closer to the access point, while items in the same group cluster based on off-diagonal values. In larger instances, balancing these competing priorities often necessitates trade-offs between dense packing and group cohesion; however, including off-diagonal values helps promote group cohesion with little to no detriment to access-point proximity.

Table 3: Prioritization matrix π for the example instance (derived from data in Table 1).

Item	1	2	3	4	5	6
1	6	0	0	0	0	0
2	0	5	2/3	1/3	0	0
3	0	0	4	2/3	0	0
4	0	0	0	3	0	0
5	0	0	0	0	2	2/3
6	0	0	0	0	0	1

4. Solution Techniques and Lower Bounds

In this section, we present the core methods we use to solve and evaluate the two-dimensional orthogonal packing problem with prioritization. Section 4.1 outlines two primary solution approaches: a direct single-MILP model and a sliding-window matheuristic strategy that exploits group-based and item-based rankings. Then, in Section 4.2, we discuss two lower-bound techniques designed to assess solution quality in a more rigorous manner. We highlight both the algorithmic and theoretical aspects of our framework before presenting numerical experiments in Section 5.

4.1. Solution Techniques

Having defined how we incorporate different priority levels, we now present and compare two primary methods for solving the prioritized 2-D orthogonal packing problem introduced in Section 3. The first is an exact method that guarantees global optimality (the Single MILP approach), while the second is a matheuristic that provides high-quality approximate solutions efficiently (the sliding-window approach). We focus on a single-bin environment where items can belong to different groups, and each group or item has an associated rank indicating its relative importance. These priorities guide how we structure subproblems in the following approaches.

2DOPP-P is structurally difficult to solve exactly. The formulation combines the geometric feasibility burden of classical bin packing with a distance-based objective more typical of facility layout planning, creating a hybrid model that is harder than either setting in isolation. The big-M constraints (5)–(8) used to enforce non-overlap produce weak LP relaxations, and the $O(n^2)$ binary separation variables create a combinatorial search space that grows rapidly, as shown by the runtime analysis in Section 5.2. In the facility layout planning literature, exact methods for similarly weighted distance objectives are typically limited to approximately 12 departments (Meller et al., 1998), whereas our instances contain up to 52 items. Decomposition techniques such as column generation are effective for multi-bin packing problems in which a master problem assigns items to bins and subproblems generate feasible single-bin packings. That structure is absent here because all items are already assigned to one bin, leaving no natural master-subproblem decomposition.

In preliminary testing, we also explored a “groupwise matheuristic” that solved entire groups in descending order of priority. While intuitive and conceptually simpler, that strategy did not allow items from different groups to be rearranged together, leading to consistently worse solutions and longer run times compared to the sliding-window approach. Thus, we omit further details here.

4.1.1. Single MILP Approach

A straightforward approach to finding a solution is to solve the 2DOPP-P formulation as a single, unified MILP. Conceptually, one feeds the entire set of items into the model outlined in Section 3 and utilizes a general-purpose solver (e.g., Gurobi). This method is referred to here as the Single MILP Approach. We construct the prioritization-based objective for all items $i, k \in I$, activate all item-placement and non-overlap constraints, and then run the solver until it either proves optimality or reaches a prescribed time limit.

A main advantage of the single MILP approach is that, without time or memory limits, it can prove a globally optimal solution. Even with time limits, it yields a bounded optimality gap if the solver can establish a dual bound. However, computation times grow rapidly as the number of items increases. In practice, for instances beyond 7 items, the solver may require hours to confirm optimality, even if it finds a feasible solution quickly. While this direct method is conceptually straightforward and advancements in mixed-integer programming techniques continue to improve solver performance (Pryor and Chinneck, 2011; Henry et al., 2023), the combinatorial complexity of the prioritization-aware packing problem often proves intractable for larger problems in real-world planning.

4.1.2. Sliding-Window Matheuristic Approach

Given the computational challenges of the single MILP approach, matheuristic methods offer a way to break the problem into smaller, more manageable subproblems. One natural attempt is to consider groups individually based on priority. However, as noted earlier, fixing entire groups sequentially can lead to suboptimal global arrangements. Therefore, we focus on a sliding-window matheuristic approach.

This sliding-window matheuristic is a sequential optimization strategy that iteratively solves a series of smaller overlapping subproblems, each containing a fixed number of consecutive items from a globally prioritized list. By concentrating on only a handful of items in each subproblem, it avoids ever having to solve a large single MILP.

We begin by ordering all items in descending priority, using a two-level sorting procedure: first by group priority (ascending numeric order), then by item-level priority within each group (also ascending numeric order), where smaller numeric values indicate higher priority. This produces a single global list where the highest-priority item from the highest-priority group appears first, followed by remaining items in their combined priority order. We choose a window size w and solve for the top w items under the usual constraints and objective. Once this subproblem terminates, we fix the single highest-priority item among those w items and keep the other $(w - 1)$ items “free.” We then slide the window forward by introducing the next item in the global priority list (i.e., the $(w + 1)$ -th item), maintaining the subproblem size at w . Because we fix one old item and add one new item each iteration, only w items remain truly variable, making it less computationally taxing than a full-instance solve.

A key parameter is the sliding-window size, w . Our experiments vary w from 1 to 9. While instances beyond 7 items become computationally intractable to prove global optimality, near-optimal solutions can often be found quickly within our imposed time limits. Preliminary testing revealed that window sizes beyond 9 showed diminishing returns or decreased performance under our time constraints (detailed analysis provided in Section 5.2). Naturally, $w = 1$ is an extreme

that places items one by one (akin to a simple greedy strategy), whereas a larger w captures more local packing trade-offs at the cost of potentially longer solve times per iteration. Because these smaller subproblems can each be run with a time limit, the method can progress quickly through the globally ordered priority list, freeing and fixing items iteratively.

In this matheuristic, each subproblem is kept small, typically leading to shorter runtimes. By sliding through a single global priority list, one can occasionally re-optimize items from different groups in the same window, offering cross-group flexibility. However, each iteration fixes only the top-priority item and locks its coordinates, potentially “freezing” suboptimal placements early if the window is too small. Performance thus depends heavily on choosing a window size that balances computational effort and solution quality. We note that theoretically, with very small window sizes and tightly constrained bin utilization, the matheuristic could potentially freeze placements that leave insufficient space for remaining items. However, we did not observe this behavior in our experiments, as the extended bin length (see Section 5.1.1) provides sufficient flexibility for prioritized placement.

4.1.3. Adapted Heuristic and Metaheuristic Methods

To place the proposed sliding-window matheuristic in a broader algorithmic context, we compare it against representative heuristic and metaheuristic methods from related packing and facility layout literatures after adapting them to the prioritization objective studied here. The goal is to assess whether the matheuristic’s advantage persists against established alternative search paradigms under a common equal-budget design. The adapted set includes nine hybrid methods formed by pairing three base layout-generation methods with three placement-space improvement routines. The base methods are a classical bottom-left packing heuristic adapted with priority-ordered item sequencing (Chazelle, 1983), Maximum Rectangles (MaxRects) (Jylänki, 2010), which builds layouts through a dynamic set of free rectangles, and Sequence Pair with Simulated Annealing (SP+SA) (Murata et al., 1996), a floorplanning-style layout-search method based on a sequence-pair representation. The improvement routines are Local Search, Variable Neighborhood Search (Mladenović and Hansen, 1997), and Tabu Search (Glover, 1989). We also test two Greedy Randomized Adaptive Search Procedure (GRASP) variants (Alvarez-Valdés et al., 2008), both implemented as randomized MaxRects-style construction followed by local search, together with three additional metaheuristics: a genetic algorithm with a MaxRects decoder (Bortfeldt and Gehring, 2001), Iterated Greedy (Ruiz and Stützle, 2007) initialized from a MaxRects layout and then refined through destroy-repair search, and Filtered Beam Search (Chen et al., 2025), which keeps several promising partial layouts at each stage of construction. Each adapted method is evaluated on the same 90 instances as the sliding-window method under an equal per-instance time budget. The comparative evaluation of these adapted methods against the sliding-window matheuristic is presented in Section 5.3. Detailed method descriptions, adaptation notes, and pseudocode are provided in the Supplemental Materials, where Table S1 serves as a legend mapping the benchmark labels used in the manuscript to the corresponding method descriptions.

4.2. Lower Bound Methods

We now discuss two techniques to establish lower bounds for the prioritized 2-D orthogonal packing problem. In the first, a rectangular-liquid relaxation yields an aggregate deterministic lower bound aligned with the prioritization objective. In the second, a statistical approach applies extreme value theory to infer a lower bound from the distribution of observed feasible minima. Both methods provide benchmarks for assessing how close our solution approaches come to optimality, especially when exact proofs of optimality are computationally infeasible.

4.2.1. Rectangular-Liquid Lower Bound

We next develop a rectangular-liquid lower bound (RL-LB) for the proposed problem. The key idea is to first drop the nonnegative vertical-offset term in the access-point objective and then relax each item into a full-height strip of equal area $A_i = p_i q_i$. This reduces the access-point component to a one-dimensional weighted-centroid problem whose strip order is determined by decreasing density π_{ii}/A_i , consistent with the adjacent-swap criterion underlying Smith's rule (Smith, 1956), while retaining a conservative non-overlap floor for the pairwise cohesion terms. This yields an aggregate deterministic lower bound on the full objective.

Theorem 1 (Rectangular-Liquid Lower Bound). *Consider a feasible instance of 2DOPP-P with bin dimensions $L \times W$ and access point $o = (0, W/2)$ on the center of the left edge. For each item i , let $A_i = p_i q_i$. Relabel the items, if necessary, so that*

$$\frac{\pi_{11}}{A_1} \geq \frac{\pi_{22}}{A_2} \geq \dots \geq \frac{\pi_{nn}}{A_n}.$$

Define

$$LB_{\text{access}} = \sum_{j=1}^n \pi_{jj} \left(\sum_{\ell=1}^{j-1} \frac{A_\ell}{W} + \frac{A_j}{2W} \right), \quad LB_{\text{pair}} = \sum_{i < k} \pi_{ik} \frac{\min\{p_i, q_i\} + \min\{p_k, q_k\}}{2},$$

and

$$LB_{\text{RL}} = LB_{\text{access}} + LB_{\text{pair}}.$$

Then $LB_{\text{RL}} \leq OPT$, where OPT is the minimum objective value of 2DOPP-P.

Figure 2 illustrates this access-term relaxation on the running six-item example. The labels inside the rectangles indicate the relative priority ranks from Figure 1, but the strips in panel (b) are ordered by decreasing density π_{ii}/A_i rather than by raw rank.

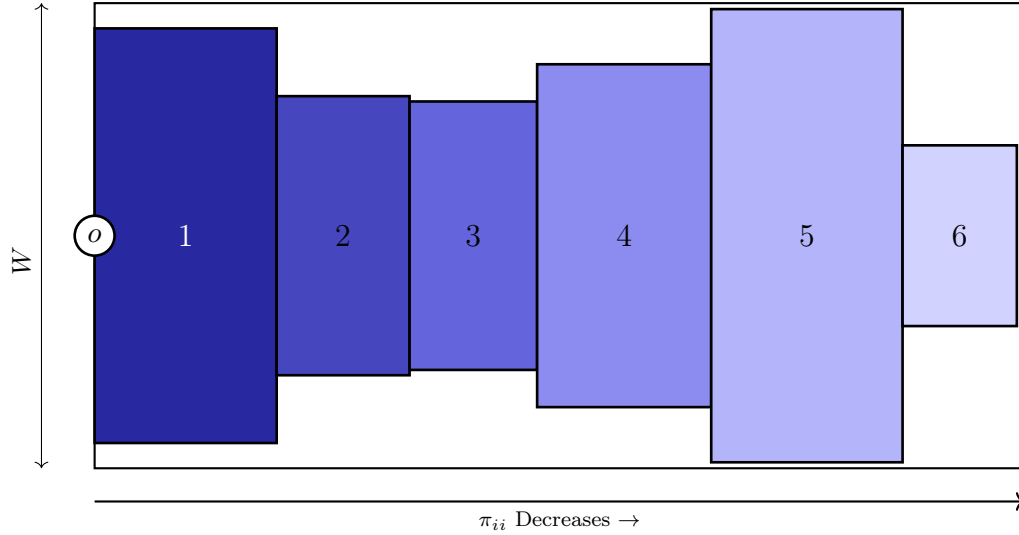
This schematic explains only the access-point portion of the RL-LB; the pairwise cohesion floor is handled separately through LB_{pair} .

Proof. Fix any feasible packing. For each item i , the access-point contribution is


$$\pi_{ii} \left(\frac{x_i + x_i^r}{2} + \left| \frac{y_i + y_i^r}{2} - \frac{W}{2} \right| \right),$$

so

$$\sum_i \pi_{ii} \left(\frac{x_i + x_i^r}{2} + \left| \frac{y_i + y_i^r}{2} - \frac{W}{2} \right| \right) \geq \sum_i \pi_{ii} \frac{x_i + x_i^r}{2}.$$



(a) Actual 2-D Placement



(b) Rectangular-Liquid Relaxation

Figure 2: Rectangular-liquid lower bound illustration for the access-point term. Top: a feasible 2-D placement of the running example. Bottom: each item is reshaped into a full-height strip of length A_i/W and ordered by decreasing priority density π_{ii}/A_i , yielding the relaxed strip-centroid distances used in LB_{access} .

The right-hand side is the weighted horizontal centroid-distance component of the true access-point objective, obtained by dropping the nonnegative vertical-offset term.

Now further relax the geometry by ignoring the original rectangle shapes and retaining only each item's area. In this relaxed access-term problem, an item of area A_j can be spread over the full bin height W , so its leftmost possible placement is a full-height strip of length A_j/W . Since the access contribution depends only on horizontal centroid location, the minimum relaxed access term is attained by assigning the leftmost strip positions to the largest densities π_{jj}/A_j ,

then proceeding in decreasing order. In that relabeled order, the centroid of strip j lies at

$$\sum_{\ell=1}^{j-1} \frac{A_\ell}{W} + \frac{A_j}{2W},$$

which gives the access contribution stated in LB_{access} for that ordering. To determine the best strip order, compare two adjacent strips i then k versus k then i , holding fixed the preceding strips $1, \dots, j-1$. The first order gives

$$\pi_{ii} \left(\sum_{\ell=1}^{j-1} \frac{A_\ell}{W} + \frac{A_i}{2W} \right) + \pi_{kk} \left(\sum_{\ell=1}^{j-1} \frac{A_\ell}{W} + \frac{A_i}{W} + \frac{A_k}{2W} \right),$$

whereas the second gives

$$\pi_{kk} \left(\sum_{\ell=1}^{j-1} \frac{A_\ell}{W} + \frac{A_k}{2W} \right) + \pi_{ii} \left(\sum_{\ell=1}^{j-1} \frac{A_\ell}{W} + \frac{A_k}{W} + \frac{A_i}{2W} \right).$$

Subtracting the second expression from the first shows that placing i before k is no worse exactly when

$$\frac{\pi_{ii}}{A_i} \geq \frac{\pi_{kk}}{A_k}.$$

This adjacent-swap criterion confirms that the minimum relaxed access term is attained by ordering the strips so that

$$\frac{\pi_{11}}{A_1} \geq \frac{\pi_{22}}{A_2} \geq \dots \geq \frac{\pi_{nn}}{A_n},$$

which yields LB_{access} . Since the strip construction solves the relaxed access problem, we have

$$LB_{\text{access}} \leq \sum_i \pi_{ii} \frac{x_i + x_i^r}{2} \leq \sum_i \pi_{ii} \left(\frac{x_i + x_i^r}{2} + \left| \frac{y_i + y_i^r}{2} - \frac{W}{2} \right| \right).$$

For pairwise terms, any two non-overlapping orthogonally packed items must be separated along at least one axis. Along that separating axis, the centroid distance is at least half the sum of the two item dimensions on that axis. The most conservative case occurs when the items are adjacent along their longest sides, so the separation is measured across their smaller dimensions. Since either possible dimension of item i is at least $\min\{p_i, q_i\}$, the rectilinear centroid distance between items i and k is at least

$$\frac{\min\{p_i, q_i\} + \min\{p_k, q_k\}}{2}.$$

Therefore

$$\sum_{i < k} \pi_{ik} \frac{\min\{p_i, q_i\} + \min\{p_k, q_k\}}{2} \leq \sum_{i < k} \pi_{ik} (\rho_{ik}^x + \rho_{ik}^y).$$

Summing the access and pairwise inequalities yields $LB_{\text{RL}} \leq OPT$. \square

Note that for an arbitrary point on the same edge, $o = (0, y^o)$, the perpendicular strip term

remains unchanged and the tangential component for item i is lower-bounded by

$$\max \left\{ 0, \frac{\min\{p_i, q_i\}}{2} - y^o, y^o - \left(W - \frac{\min\{p_i, q_i\}}{2} \right) \right\}.$$

The resulting extension is

$$LB_{\text{RL}}^{\text{edge}} = LB_{\text{access}} + \sum_i \pi_{ii} \max \left\{ 0, \frac{\min\{p_i, q_i\}}{2} - y^o, y^o - \left(W - \frac{\min\{p_i, q_i\}}{2} \right) \right\} + LB_{\text{pair}}.$$

The right-edge case is symmetric, and top/bottom-edge variants follow by swapping x and y .

The RL-LB is computationally inexpensive: the access term requires an $O(n \log n)$ sort by π_{ii}/A_i , and the pairwise term requires an $O(n^2)$ pass, giving total complexity $O(n \log n + n^2)$.

4.2.2. Statistical Bound using Extreme Value Theory

Having discussed a deterministic bound, we also use a complementary statistical lower bound (Wilson et al., 2004) to further evaluate solution quality. In essence, Wilson’s method leverages extreme value theory, drawing upon the fundamental result by Fisher and Tippett (1928) regarding the asymptotic distribution of extreme values, to fit a three-parameter Weibull distribution to observed minimum objective values, thereby constructing confidence intervals for the global optimum z . The three-parameter Weibull cumulative distribution function (CDF) is given by $F(x; a, b, c) = 1 - \exp\left(-\left(\frac{x-a}{b}\right)^c\right)$ for $x \geq a$, where a is the location (threshold) parameter representing the minimum possible value, $b > 0$ is the scale parameter, and $c > 0$ is the shape parameter.

Our implementation proceeds as follows. First, for each large problem instance we create n “slightly perturbed” variants by altering only the per-group priority hierarchies, keeping all other aspects of the problem constant. Each variant is solved using whichever solution method performed best on that instance, and we record the minimum objective value obtained, denoted x_i for $i = 1, \dots, n$. Crucially, we then recalculate that solution’s objective value using the original problem’s prioritization hierarchy, so that all minima are assessed on the same baseline. Let $x_{(k)}$ denote the k -th order statistic such that $x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(n)}$, hence $x_{(1)} = \min\{x_i\}$. Next, we fit the Weibull curve via least-squares estimation to the n observed minima. We use the analytical parameter estimates from Zanakis (1979) as starting values for the location (\hat{a}), scale (\hat{b}), and shape (\hat{c}) parameters:

$$\hat{a} = \frac{x_{[1]}x_{[2]} - x_{[2]}^2}{x_{[1]} + x_{[n]} - 2x_{[2]}}, \quad (23)$$

$$\hat{b} = x_{[\lceil 0.63n \rceil]} - \hat{a}, \quad (24)$$

$$\hat{c} = \frac{\ln\{\ln(1 - q_1)/\ln(1 - q_2)\}}{\ln\{(x_{[\lceil nq_1 \rceil]} - \hat{a})/(x_{[\lceil nq_2 \rceil]} - \hat{a})\}}, \quad (25)$$

where $q_1 = 0.97366$ and $q_2 = 0.16731$. These initial estimates are then refined using the Nelder and Mead (1965) algorithm to obtain the final fitted parameters $\hat{a}, \hat{b}, \hat{c}$. We then evaluate goodness of fit with both the Kolmogorov–Smirnov (KS) and Anderson–Darling (AD) tests, the latter being critical for tail accuracy. Following Golden and Alt (1979), using the fitted

parameters, we construct an approximate $100(1 - e^{-n})\%$ confidence interval for the threshold parameter a . This interval takes the form $[\hat{x}^*, x_{(1)}]$, where the lower limit \hat{x}^* is given by $\hat{x}^* = x_{(1)} - \hat{b}$. Since the threshold parameter a represents the theoretical minimum objective value, \hat{x}^* provides a statistical lower bound for the true global optimum z . This process largely mirrors the heuristic-based method in [Wilson et al. \(2004\)](#) but adapts naturally to our mixed-integer optimization setting. This general methodology, leveraging extreme value theory to estimate confidence intervals for the optimal objective value, represents a well-established approach in combinatorial optimization for evaluating solution quality where tight deterministic lower bounds may be unavailable or computationally prohibitive ([Derigs, 1985](#); [Altinel et al., 2000](#); [Akyüz et al., 2010](#)).

In [Section 5.2.2](#), we illustrate how this method, in tandem with the RL-LB, offers a more comprehensive view of solution quality.

5. Performance Evaluation of Solution Approaches

This section presents a computational study to demonstrate how the proposed 2-D orthogonal packing problem with prioritization performs under realistic logistics scenarios. We describe the generated data and instances, then evaluate the main computational study, and finally benchmark a selected sliding-window configuration against adapted heuristic and metaheuristic alternatives under equal time budgets.

5.1. Data and Instances

We first describe our instance-generation procedures, including group and item sizes, item aspect ratios, and bin dimensions. We then outline our computational settings and experimentation process.

5.1.1. Instance Generation

To motivate the selection of parameters for our computational study, we base our instance generation on characteristics observed in realistic logistics scenarios, particularly military combat loading onto large vessels, while ensuring applicability to other settings where possible. We generate a diverse set of 90 problem instances to evaluate the proposed 2-D orthogonal packing problem with prioritization; these instances are available from [Kirschenman et al. \(2025\)](#). Our experimental design provides sufficient statistical power to evaluate solution technique parameters across varied problem characteristics, with observed performance trends remaining consistent even with smaller instance subsets. Each instance includes between 1 and 5 groups of items, capturing a variety of possible scenarios. Within each group, items share a common mission or operational function and thus belong together for off-loading or unloading purposes, but they can differ in size or priority. Our instance generation process generally proceeds as follows:

Group and Item Counts. First, we randomly select the number of groups to be between 1 and 5. Each group then has a size drawn from one of three categories: (i) small (1–4 items), (ii) medium (5–8 items), or (iii) large (9–12 items). We cap groups at a maximum of 12 items for two reasons: (1) from a computational standpoint, 12 is large enough to be challenging yet still meaningful

within our solution framework, and (2) this accommodates a typical platoon-sized echelon of vehicles in military contexts. This range also aligns with observations from online retail data, where 84% of "orders" (analogous to groups) include four or fewer items (Hübner et al., 2015). Fontaine and Minner (2023) limit order sizes to 10 items based on a Brazilian e-commerce data set (Olist, 2019) in which the maximum observed order size is 21 items, but over 97% contain just one or two. Consequently, each generated problem instance could contain as few as 1 item or as many as 60 items. These extremes reflect a broad range of scenarios—covering the upper limits typically seen in e-commerce order sizes and matching the capacity of the largest roll-on/roll-off (RO/RO) vessel used by the U.S. Navy (United States Navy, 2021) when we consider up to 5 groups of 12 items each.

Homogeneity Levels for Item Sizes. To capture the varied nature of equipment footprints in real logistics settings, we categorize each group as (i) homogeneous, (ii) weakly heterogeneous, or (iii) strongly heterogeneous in terms of item dimensions. A homogeneous group assigns identical length/width to all items, mirroring situations where equipment is standardized. Weak heterogeneity means that two distinct item sizes appear within the group. Strong heterogeneity means that no more than two items share the same size. These settings let us move from standardized equipment to mixtures with substantially more dimensional variation while keeping all items rectangular and allowing 90° rotation.

Item Dimensions and Aspect Ratios. All items are rectangular, with length and width drawn so that the aspect ratio lies between 1:1 and 1:3. This range mirrors the footprints of typical military wheeled/tracked vehicles—for example, an M1 Abrams tank has length-to-width close to 2.7:1 (General Dynamics Land Systems, 2025). Some items are small (e.g., a generator set or small package), others large (a tank or large crate).

Bin Dimensions and Extended Length. To ensure adequate space for priority-based placement, we start with a sufficient bin area for feasibly packing assigned items and then inflate the bin length by an additional factor. This approach is conceptually similar to the “open dimension problem” in Wäscher et al. (2007)’s typology, allowing extra space to accommodate loading priorities rather than purely minimizing empty area. In practice, the extended bin length grants more freedom for positioning items around the bin access point (for instance, along one edge or corner), ensuring the solver can focus on optimizing the prioritization objective—pulling items toward the access point—rather than being constrained by feasibility considerations that would force suboptimal packing decisions, making it easier to isolate the effect of prioritization in the objective.

Bin Access Points. We considered two primary bin access points: a center-left (CL) edge and a bottom-left (BL) corner on the bin’s plane. These locations represent common configurations found in logistics, such as centered ramps for beach landings or corner ramps for pier docking in RO/RO operations (United States Navy, 2023, 2021). In our experiments, the two access points produced nearly identical results regarding the effectiveness of the prioritization. Therefore, for brevity, we report only CL-based results in subsequent sections. Furthermore, the CL location

can represent centered access points common in other settings, like multi-drop delivery trucks, enhancing the broad applicability of the presented results.

All models were implemented in Julia 1.11.0 and solved by Gurobi 12.0. All experiments were conducted on an Intel Xeon (Skylake) processor, 2.30 GHz, with up to 1 terabyte RAM operating at 2400 MHz, and access to up to 32 cores through the attached virtual machine. For each problem instance, we performed a single MILP solve under a time limit defined as the number of items multiplied by 60 seconds. For the matheuristic approach, we applied time limits of 5 seconds, 15 seconds, and 60 seconds per windowed solve. Consequently, for each problem instance, a total of 27 solves is carried out, one for each window size ($w \in \{1, 2, \dots, 9\}$) and time limit combination. This time allocation strategy ensures fair comparison between the two approaches: both the single MILP and matheuristic receive equivalent maximum computational budgets. The single MILP approach receives $(60 \times |I|)$ seconds total time, while the matheuristic approach could theoretically use the same total time in its most computationally intensive configuration (window size $w = 1$ with 60 seconds per subproblem, yielding $60 \times |I|$ seconds). However, in practice, smaller subproblems in the matheuristic approach often solve near-instantaneously, allowing the method to achieve superior solution quality while using less computational time than its allocated budget.

5.2. Performance of Solution Approaches

This section evaluates our prioritized orthogonal packing framework and solution methods, providing insights into quantitative metrics of runtime, solution quality, and optimality gaps. We analyze runtime behavior for the matheuristic approach, evaluate the effectiveness of the lower bounds, compare the objective values achieved by different solution methods, and explore the resulting trade-off between computation time and solution quality.

5.2.1. Runtime Analysis

As discussed in Section 4.1, the sliding-window matheuristic approach applies time limits to subproblems because the time required to prove optimality grows rapidly with the number of items. To illustrate this computational challenge, we extracted 22 seven-item groups, 30 eight-item groups, and 23 nine-item groups from our 90 problem instances and solved them as standalone problems to proven optimality. Table 4 shows the resulting solution times, demonstrating exponential growth as instance size increases. In contrast, our matheuristic applies per-window time limits (5, 15, or 60 seconds), often achieving near-optimal solutions within these constraints for windows up to 7 items.

Table 4: Summary statistics for time (in seconds) to prove optimality across all 7–9 item groups treated as standalone instances.

Group Size	Mean	Std	Min	Median	Max
7 Items	9.72	3.88	4.26	8.38	17.32
8 Items	94.56	39.91	30.55	88.51	184.75
9 Items	1279.50	613.31	352.67	1281.83	2853.02

Figure 3 further shows how the sliding-window approach scales with increasing window size (w) and different per-subproblem time limits (5, 15, or 60 seconds). Subproblems with up to 5

items often solve almost instantly, whereas larger windows (8–9 items) may benefit from longer time limits but then risk hitting diminishing returns. Ultimately, windows of 6–7 items under short time limits emerge as a solid compromise between speed and solution quality. These observations guide our choice of matheuristic parameters in the subsequent evaluations.

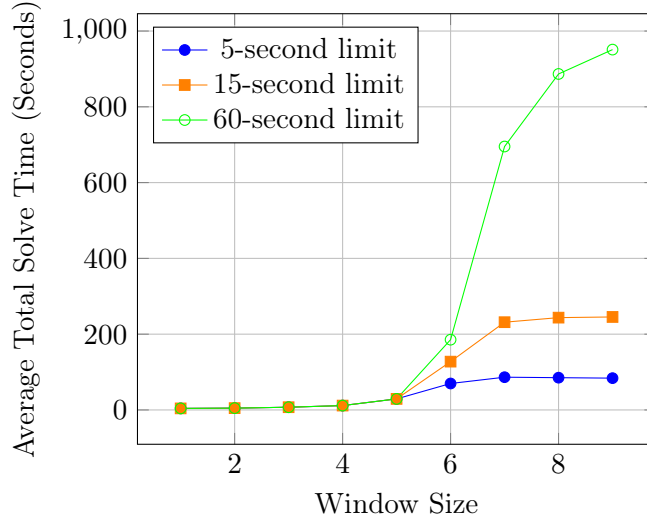


Figure 3: Average Total Time vs. Window Size for Windowed Solve by Time Limit.

5.2.2. Assessing Solution Quality with Lower Bounds

We next investigate how closely our best solutions approach optimality by comparing them against three sets of lower bounds: (1) Gurobi’s best lower bound from the single MILP solution approach, (2) the deterministic rectangular-liquid lower bound described in Section 4.2.1, and (3) the statistical bound described in Section 4.2.2. We focus here on the largest problem instances since they present the most challenging scenarios and highlight where any bound weaknesses are most pronounced. We note, however, that similar trends hold for smaller problems as well.

Rectangular-Liquid Deterministic Bound. Recall from Section 4.2.1 that our rectangular-liquid lower bound (RL-LB) explicitly incorporates prioritization through access-point density ordering and conservative pairwise separation terms. This approach directly addresses weaknesses inherent to traditional bin packing or facility layout bounds—particularly their inability to account for distance-based prioritization or group-placement constraints. As Table 5 shows, the RL-LB provides a deterministic benchmark that complements Gurobi’s built-in bound (LB), showing that structure-aware lower bounds can have increased utility as problem size grows.

Statistical Bound. While the Gurobi and RL-LB bounds provide provable lower bounds, they often exhibit substantial optimality gaps of 15–80% as shown in Table 5. The statistical lower bound, though not deterministically provable, leverages the distribution of high-quality feasible solutions to estimate where the true optimal likely lies. This complementary perspective is particularly valuable for large instances where traditional bounds are loose but solution quality assessment remains critical for operational decision-making.

Table 6 presents results for our three largest instances using the statistical lower bound approach described in Section 4.2.2. As detailed there, we fit a three-parameter Weibull distri-

Table 5: Comparison of Best Known Objective and Lower Bounds for 8 Largest Problem Instances. Values for Obj, LB, and RL-LB are scaled by 10^5 (e.g., 11.3 represents 1.13e6). % Gap columns show $((\text{Obj} - \text{Bound}) / \text{Obj}) \times 100\%$. The “Group Size” codes (s = small, m = medium, l = large) denote the size category of each group, and the “Heterogeneity” codes (h = homogeneous, w = weakly heterogeneous, s = strongly heterogeneous) reflect item dimension variation within groups.

Items	Groups	Group Size	Heterogeneity	Obj	LB	RL-LB	LB % Gap	RL-LB % Gap
52	5	mlllll	swhhw	11.3	2.16	7.42	80.80%	34.34%
45	5	mllml	wsss	8.58	1.55	6.01	81.95%	29.93%
44	5	mmlll	wshh	7.97	1.53	5.29	80.82%	33.58%
42	5	slmll	shsws	7.11	1.38	5.20	80.62%	26.89%
42	4	lmll	hhhw	6.31	1.40	3.41	77.87%	45.97%
41	4	lmll	hwws	7.23	1.49	4.91	79.36%	32.05%
40	5	mllmm	sshws	6.31	1.27	4.54	79.94%	28.06%
38	5	lmlsl	hshhw	5.24	1.11	3.59	78.90%	31.40%

bution to a set of perturbed minima (Eqs. (23)–(25)) and construct a Golden–Alt confidence interval for the theoretical optimum. Table 6 presents the results for our three largest instances. The “Min” and “Max” columns shows the smallest and largest of the 100 observed minima, while the “KS (Stat, p -val)” and “AD (Stat, p -val)” columns confirm a satisfactory Weibull fit (none of the p -values are near conventional significance thresholds). The resulting Golden–Alt CI provides a statistical lower bound for each instance, with the upper bound set by the minimum observed value. Finally, the “Opt. Gap (%)” column compares our best-known feasible solution against that statistical lower bound. Overall, these results suggest that the best solutions are far closer to the global optimum than implied by either Gurobi’s or our rectangular-liquid lower bound. While the procedure remains sensitive to sample size, distribution fit, and feasible minima generation, in practice, it provides a useful complementary perspective on solution quality—especially when deterministic bounds are weak. Next we shift focus from “how close to optimal?” to “which solution method performs best?” across all problem instances.

Table 6: Wilson’s Statistical Lower Bound Results for the Three Largest Instances. Min, Max, and Golden-Alt CI values are scaled by 10^5 .

Instance	Min	Max	KS (Stat, p -val)	AD (Stat, p -val)	Golden-Alt CI	Opt. Gap
42 items	6.31	6.74	(0.056/0.432)	(0.401/0.847)	(6.15, 6.31)	2.61%
45 items	8.58	9.31	(0.034/0.189)	(0.117/0.999)	(8.24, 8.58)	4.58%
52 items	11.3	12.0	(0.076/0.644)	(0.534/0.712)	(10.5, 11.3)	7.62%

5.2.3. Comparative Objective Performance

Figure 4 illustrates the percentage deviation of each solution technique’s final objective value relative to the best objective value obtained for that instance, across all 90 problem instances. In the figure, solution approaches are ordered left to right by increasing median deviation. Solution approaches that start with “win” designate a sliding-window matheuristic with a set window size and per-subproblem time limit. For example, “win_8it_60” refers to an 8-item window, 60-second time limit, sliding-window matheuristic. As a benchmark, we use the single MILP approach (labeled “single”): any method with median performance below that of single

MILP is retained, while those that failed to improve on single MILP were excluded for clarity (e.g., sliding-window sizes of 3 or fewer and certain sliding-window combinations that struggled under 5-second time limits).




Figure 4: Boxplot of Percentage Deviation from Best Known Objective, Ordered by Increasing Median.

Matheuristic methods consistently outperform the single MILP with visibly tighter deviation spreads, as seen in Figure 4. Allowing more time per window (e.g., 15 vs. 5 seconds) naturally reduces median deviation, but increasing the window size alone does not guarantee improvement. For example, a 9-item window with a 60-second limit ranks among the top performers with a median deviation less than 1%, yet its 9-item, 5-second variant performs worse than the single MILP and is not shown. The 8-item, 60-second sliding-window matheuristic achieves the strongest median performance overall, although it requires an average total runtime of roughly 900 seconds (see Figure 3). By contrast, the most effective method at a 5-second time limit is the 7-item window, which finishes processing most instances in just over a minute while still outperforming single MILP.

These observations highlight the relationship between solution quality and computational effort. Figure 5 highlights this balance more explicitly. Here, each data point indicates a method’s average percent deviation from best solution (y-axis) versus its average solve time (x-axis), across all problem instances. For single MILP, we record intermediate solver states at 100-second increments (though solver callbacks may align slightly off these intervals), then average the deviations. As the figure shows, matheuristic methods define the Pareto front, with single MILP solutions lagging well behind. In particular, 5–7 item windows under a 5-second limit appear to strike an excellent compromise between speed and solution quality, rarely

exceeding a few percent deviation.




Figure 5: Pareto curve showing the relationship between mean solution time and mean percentage deviation from the best obtained objective values across all instances. Black squares show the Pareto-efficient frontier, which is entirely composed of matheuristic solutions; black diamonds represent single MILP solutions recorded at various time intervals. Single MILP points represent mean percentage deviations averaged across all 90 instances, with completed instances carrying forward their final objective values to subsequent time points, ensuring consistent population comparison.

Overall, the matheuristic approach significantly outperforms a single full-model solve in both time and quality. Larger window sizes can improve solution quality, but diminishing returns arise as subproblem size grows. Meanwhile, short time limits suffice for small windows, achieving excellent solutions in mere seconds or minutes. By choosing window sizes strategically, planners can solve large instances quickly while preserving high alignment with prioritization goals. [Appendix A](#) provides side-by-side visualizations for three representative instances in [Figure A.6](#), and [Table A.8](#) summarizes the corresponding objective values for those examples.

5.3. Comparison with Adapted Heuristic and Metaheuristic Methods

Building on the time-quality tradeoff analysis in [Section 5.2.3](#), we use the sliding-window configuration with $w = 7$ and a 5-second time limit per window as the reference matheuristic for a broader equal-budget comparison. That configuration lies near the elbow of the observed Pareto curve: among the short-budget settings, it offers a strong balance between objective-value quality and total solution time. We therefore compare this reference configuration against the 14 adapted methods summarized in [Section 4.1.3](#). Each adapted method receives a per-instance time budget of $T = 5 \times \max(1, n_{\text{items}} - 6)$ seconds, which matches the total budget of the reported

$w = 7$, 5-second sliding-window configuration. The added heuristic study therefore contributes 1,260 method-instance runs (14×90). When the 90 reported windowed-matheuristic runs are included, the comparison table summarizes 1,350 method-instance outcomes.

For each instance in this equal-budget heuristic comparison only, we define the comparison-study best-known solution, hereafter BKS, as the lowest objective value observed among the 15 methods reported in Table 7. Table 7 reports each method’s average and median percentage deviation from that BKS and the percentage of instances on which the method attains the BKS.

Table 7: Performance comparison of 15 solution methods on 90 benchmark instances under equal time budgets, ordered by increasing mean deviation. “Mean Dev” and “Med Dev” report percentage deviation from the BKS, defined for each instance within this equal-budget heuristic comparison as the best objective observed among the 15 reported methods. “%BKS” is the percentage of instances on which the method attains that best value. The three SP+SA-based methods returned feasible solutions on 83 of the 90 instances, so their deviation summaries are computed over those 83 completed runs.

Method	Mean Dev	Med Dev	%BKS
Windowed Matheuristic	0.2%	0.0%	84.4%
GA + MaxRects	3.8%	3.1%	5.6%
Iterated Greedy	4.2%	3.3%	7.8%
MaxRects + VNS	6.5%	5.7%	1.1%
SP+SA + VNS	7.2%	6.1%	2.2%
BL-PS + VNS	8.8%	7.9%	2.2%
MaxRects + Tabu	9.3%	9.4%	1.1%
MaxRects + LS	9.3%	9.4%	1.1%
Filtered Beam Search	9.8%	9.5%	0.0%
GRASP (standard)	9.9%	10.3%	2.2%
GRASP (reactive)	10.1%	10.3%	2.2%
SP+SA + Tabu	10.2%	9.6%	1.1%
SP+SA + LS	10.3%	9.7%	1.1%
BL-PS + LS	16.5%	15.6%	1.1%
BL-PS + Tabu	21.4%	19.9%	1.1%

The windowed matheuristic attains the BKS on 76 of the 90 instances (84.4%), far more often than any adapted alternative, while also showing the strongest overall consistency. Its mean deviation from the BKS is 0.2%, and its median deviation is 0.0%. Among the adapted methods, the strongest average results come from GA+MaxRects and Iterated Greedy, with mean deviations of 3.8% and 4.2%, respectively. The best-performing construction-plus-improvement combinations are MaxRects+VNS and SP+SA+VNS, but both remain materially behind the windowed matheuristic in mean deviation, median deviation, and BKS frequency.

This table also highlights the matheuristic’s consistency across the instance set. Its combination of a 0.2% mean deviation, a 0.0% median deviation, and BKS attainment on 84.4% of instances indicates that its advantage is not limited to average performance alone. Its worst observed gap from the BKS is only 6.3%, whereas the next-best worst observed gap among the adapted methods is 29.7%. Looking beyond the summary statistics in Table 7, a per-instance review of the 90 benchmark outcomes shows that the windowed method remains within the top two methods on 85 of the 90 instances. On the five instances where it falls outside the top two, its average margin from the BKS is only 2.8%, with a worst such margin of 6.3%. These exceptions occur on medium-to-large instances and are driven by GA+MaxRects or Iterated Greedy rather than by the simpler adapted methods.

Several additional patterns are worth noting. First, the strongest adapted alternatives are the more elaborate metaheuristics, especially GA+MaxRects and Iterated Greedy, whereas the BL-PS-based hybrids remain substantially weaker in the aggregate. Second, among the constructor-plus-improvement families, the MaxRects- and sequence-pair-based variants outperform the simpler bottom-left starts, suggesting that both the underlying layout representation and the quality of the improvement phase matter for this problem class. Third, Variable Neighborhood Search is the strongest improvement routine within the nine hybrid methods: the three VNS variants occupy the fourth through sixth positions overall, ahead of the corresponding Local Search and Tabu Search variants for every constructor. The saved best-so-far traces collected for this benchmark study also indicate that most improvements occur well before the allocated time budget is exhausted, rather than appearing predominantly in the final portion of the run.

Overall, this equal-budget comparison supports the same practical conclusion suggested by the preceding Pareto analysis. The selected $w = 7$, 5-second sliding-window configuration provides a strong balance of solution quality and computational effort, and within this benchmark set it also delivers the strongest overall empirical performance against the adapted heuristic and metaheuristic alternatives.

6. Conclusion

In this paper, we introduce a two-dimensional orthogonal packing model that embeds multiple levels of prioritization into the objective for arranging an assigned item set within a selected bin, thereby unifying classical packing constraints with facility-layout-inspired distance preferences. Our approach centers on a prioritization matrix that captures both item-to-item adjacency needs and item-to-access-point proximity requirements. Through computational experiments, we demonstrate how this formulation can yield spatially dense configurations close to a bin access point while maintaining strong group unity and partial sequencing alignment when desired.

The proposed framework applies to combat loading in military logistics, warehouse and cargo stowage, multi-drop deliveries, or any application where fine-grained placement priorities influence overall layout. We show that single-level and multi-level prioritization can be accommodated by simply adjusting weights within a unified matrix, making the model highly flexible in handling diverse real-world scenarios, from roll-on/roll-off ship embarkation planning to automated pick-and-pack warehouses.

Managerial Implications. Our computational study suggests that the sliding-window metaheuristic with window sizes 6–7 and 5-second per-window time limits offers the strongest practical speed-quality tradeoff, typically completing most instances in about one minute while remaining close to the best-known solutions. The results also show that optimization becomes far more valuable as problem size grows. On very small cases, a domain expert may produce a reasonable layout by inspection, whereas on larger instances the simultaneous demands of group cohesion, access-point proximity, and non-overlap make automated optimization much more important.

The study also indicates that group cohesion and access-point proximity can be pursued together rather than treated as competing goals, since adding intra-group cohesion weights generally improves operational grouping with little loss in access-point performance. More broadly,

practitioners can explore different operational priorities by adjusting the prioritization matrix without reformulating the model, and the benchmarking results suggest that the matheuristic’s reliability across instance types is important for planning use, even when it is compared against a set of adapted heuristic and metaheuristic alternatives.

Practical Considerations. In implementing the two optimization-based methods discussed (single MILP and sliding-window matheuristic), we note several common trade-offs. First, proving optimality for even moderately sized instances can be prohibitively time-consuming, since bin packing and layout problems are known to exhibit exponential complexity (Meller et al., 1998). Consequently, imposing relatively short time limits on each subproblem in the sliding-window matheuristic often yields high-quality solutions without requiring the solver to close large optimality gaps. Second, seamless termination in the sliding-window method means each smaller subproblem finishes quickly, ensuring the overall procedure completes in a predictable manner. Although the matheuristic strategy does not provide a global MILP gap, we observe that practical performance can be strong when matheuristic parameters (e.g., subproblem sizes or time limits) are chosen judiciously.

To complement these practical observations, we briefly note what can and cannot be said analytically about the sliding-window method.

To further characterize the matheuristic, we provide three implementation-level guarantees. Let n be the number of items and let w denote the window size. The number of windowed MILP solves is exactly $N_{\text{win}} = \max(1, n - w + 1)$, matching the loop structure in our implementation. Thus, when each window solve terminates within the prescribed per-window time limit, total solver time is bounded by that limit multiplied by N_{win} (plus model construction overhead). In practice, all tested window subproblems in our computational study found feasible solutions well within the prescribed per-window time limit.

We also note a boundary-case equivalence: when $w \geq n$, the sliding-window method performs a single solve with no fixed items, so the executed MILP is the same full-instance model solved by the single-MILP baseline under the same solver settings. Hence, any performance differences between the two approaches necessarily arise in the regime $w < n$, where sequential fixing trades global re-optimization freedom for computational tractability.

At the same time, we do not claim a general worst-case approximation ratio for the full sliding-window matheuristic. The combination of non-overlap constraints, weighted pairwise and item-to-access-point distances, and irrevocable fixing decisions creates path dependence that can produce substantial quality variation across instances when w is small. Our targeted calibration on 27 small instances ($n \leq 9$), each solved to proven optimality, confirms this behavior: the greedy case $w = 1$ had a mean gap of 7.9% and a worst observed gap of 64.4%, whereas $w \geq 4$ reduced the mean gap below 0.5% (worst case 2.2%), and $w = 6$ matched the exact optimal on all applicable instances. Solution quality is empirically non-decreasing in w when the per-window time limit is sufficient to solve each subproblem near-optimally; for very large windows on harder instances, the per-window time limit itself becomes the binding constraint rather than the window size.

Taken together, these observations clarify the scope of the method’s theoretical characterization: the procedure admits simple implementation-level bounds and informative boundary-case

behavior, but not a general worst-case approximation guarantee.

Notably, both solution approaches are applicable whenever there is a well-defined bin access point and prioritized items or groups. If a future application requires purely intra-group clustering without a bin access point consideration, the group ordering logic would need to be revised to account for the lack of a distance-based pull toward a bin edge. Nonetheless, the broader concept of decomposing large layout problems into manageable subproblems still applies, illustrating the extensibility of our proposed framework to alternate objectives and constraints.

Future Directions. Our computational results demonstrate the effectiveness of the matheuristic approach, offering planners a practical method to balance solution quality and runtime by tuning parameters (e.g., window size, time limits). Building on this, future algorithmic work could explore specialized heuristics or metaheuristics which may further reduce solve times or handle problem variants with additional constraints such as load balancing, non-adjacency requirements for hazardous items, or orientation restrictions (e.g., if certain cargo cannot be rotated).

Finally, the single-bin assumption could be lifted to investigate multi-bin or multi-vessel scenarios, incorporating a bin-selection phase that weighs how to distribute items among multiple bins while still respecting spatial priorities. Likewise, a scenario with multiple access points or varying item-level sequencing could be modeled by suitably extending the prioritization matrix. Overall, this study highlights the power of integrating spatial and priority-based considerations in orthogonal packing. By illustrating a flexible distance-weighted formulation and showing how it can be efficiently solved via matheuristic techniques, we provide a basis for future research in military, warehouse, and other logistics domains where both feasibility and strategic placement goals must be satisfied.

Funding

This work was partially supported by the Omar N. Bradley Officer Research Fellowships in Mathematics.

Conflict of Interest Statement

The authors declare no conflict of interest.

Data Availability Statement

The data that support the findings of this study are openly available in Dryad ([Kirschenman et al., 2025](#)) at <https://doi.org/10.5061/dryad.8pk0p2p1z>.

Appendix A. Solution Visualization Examples

This appendix presents side-by-side solution visualizations for three representative instances: a small instance (2 groups, 14 items), a medium instance (3 groups, 29 items), and a large instance (5 groups, 52 items). Figure A.6 shows these visualizations. In each row, the left panel shows the single MILP solution and the right panel shows the sliding-window matheuristic

solution. Table A.8 reports the corresponding objective values for these same examples. Using the computational time-limit rule described in Section 5.1, the corresponding single-MILP time limits were 840, 1,740, and 3,120 seconds for the small, medium, and large instances, respectively.

Table A.8: Objective comparison for the three representative visualization instances. Here, SW denotes the sliding-window matheuristic with $w = 7$ and a 5-second time limit per window. The final column reports the percentage by which the single-MILP objective exceeds the corresponding SW objective.

Instance	Single MILP Obj.	SW Obj.	MILP Gap vs. SW
Small (2 groups, 14 items)	49,812.4	49,376.9	0.9%
Medium (3 groups, 29 items)	282,691.3	269,157.8	5.0%
Large (5 groups, 52 items)	1,357,198.4	1,142,394.5	18.8%

In all three examples, the sliding-window matheuristic attains the lower objective value. Accordingly, the percentages in Table A.8 can be read as the improvement the single-MILP objective would need in order to match the lower sliding-window objective.

The small instance provides a straightforward baseline. The two methods produce visually similar layouts, and the objective values are correspondingly close: the single-MILP objective is only 0.9% above the sliding-window matheuristic objective.

The medium instance begins to show more visible divergence between the two methods, but the main group structures remain relatively easy to follow. Here, the single-MILP objective is 5.0% above the sliding-window matheuristic objective, making it a useful intermediate comparison between the small and large cases.

The large instance provides the clearest view of the trade-off between group cohesion and access-point proximity. In the sliding-window solution, the lowest-priority group remains almost entirely adjacent, while one smaller item is pulled slightly closer to the access point. The second-lowest-priority group is also somewhat more dispersed, suggesting that thinner items can exploit narrow gaps that larger items cannot. This row also shows the clearest overall difference between the two methods: the single-MILP solution exhibits substantially weaker group cohesion, and its objective is 18.8% above the lower sliding-window matheuristic objective.

Taken together, these three examples complement the 6-item illustration in Figure 1. They show that the visual and objective-value differences between the two methods are modest on the small instance, begin to widen on the medium instance, and become most pronounced on the large instance.

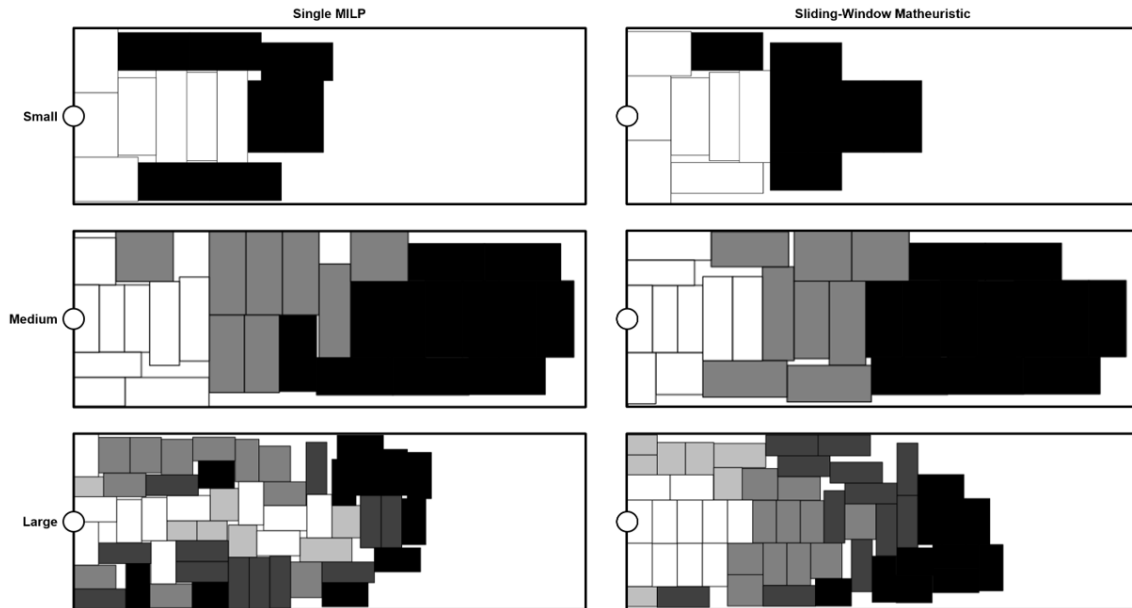


Figure A.6: Side-by-side solution visualizations for three representative instances. Rows correspond to a small instance (2 groups, 14 items), a medium instance (3 groups, 29 items), and a large instance (5 groups, 52 items). The left column shows the single MILP solution and the right column shows the sliding-window matheuristic solution. Items are shaded by group priority from white (highest-priority group) to black (lowest-priority group), and the bin's access point is designated by a white circle on the center-left edge. For presentation consistency, the displayed bin widths are normalized across rows and excess bin length is cropped; these visual adjustments do not change the underlying layouts or objective values.

References

- Ahmadi, A., Pishvae, M.S., Jokar, M.R.A., 2017. A survey on multi-floor facility layout problems. *Computers & Industrial Engineering* 107, 158–170. doi:[10.1016/j.cie.2017.03.015](https://doi.org/10.1016/j.cie.2017.03.015).
- Akyüz, M.H., Öncan, T., Altinel, I.K., 2010. The multi-commodity capacitated multi-facility weber problem: Heuristics and confidence intervals. *IIE Transactions* 42, 825–841. doi:[10.1080/0740817X.2010.491504](https://doi.org/10.1080/0740817X.2010.491504).
- Altarazi, S., 2013. A new prioritizing-stacking heuristic algorithm for the inner-city truck loading problem. *International Journal of Business and Management* 8, 137. doi:[10.5539/ijbm.v8n13p137](https://doi.org/10.5539/ijbm.v8n13p137).
- Altinel, I.K., Aras, N., Oommen, B.J., 2000. Fast, efficient and accurate solutions to the hamiltonian path problem using neural approaches. *Computers & Operations Research* 27, 461–494. doi:[10.1016/S0305-0548\(99\)00065-9](https://doi.org/10.1016/S0305-0548(99)00065-9).
- Alvarez-Valdés, R., Parreño, F., Tamarit, J.M., 2008. Reactive GRASP for the strip-packing problem. *Computers & Operations Research* 35, 1065–1083. doi:[10.1016/j.cor.2006.07.004](https://doi.org/10.1016/j.cor.2006.07.004).
- Alvarez-Valdés, R., Parreño, F., Tamarit, J.M., 2013. A grasp/path relinking algorithm for two-and three-dimensional multiple bin-size bin packing problems. *Computers & Operations Research* 40, 3081–3090. doi:[10.1016/j.cor.2012.03.016](https://doi.org/10.1016/j.cor.2012.03.016).

- Bazaraa, M.S., Sherali, H.D., 1980. Benders' partitioning scheme applied to a new formulation of the quadratic assignment problem. *Naval Research Logistics Quarterly* 27, 29–41. doi:[10.1002/nav.3800270104](https://doi.org/10.1002/nav.3800270104).
- Bischoff, E.E., Ratcliff, M., 1995. Issues in the development of approaches to container loading. *Omega* 23, 377–390. doi:[10.1016/0305-0483\(95\)00015-G](https://doi.org/10.1016/0305-0483(95)00015-G).
- Borges, Y.G., Schouery, R.C., Miyazawa, F.K., 2024. Mathematical models and exact algorithms for the colored bin packing problem. *Computers & Operations Research* 164, 106527. doi:[10.1016/j.cor.2023.106527](https://doi.org/10.1016/j.cor.2023.106527).
- Bortfeldt, A., Gehring, H., 2001. A hybrid genetic algorithm for the container loading problem. *European Journal of Operational Research* 131, 143–161. doi:[10.1016/S0377-2217\(00\)00055-2](https://doi.org/10.1016/S0377-2217(00)00055-2).
- Bortfeldt, A., Wäscher, G., 2013. Constraints in container loading – a state-of-the-art review. *European Journal of Operational Research* 229, 1–20. doi:[10.1016/j.ejor.2012.12.006](https://doi.org/10.1016/j.ejor.2012.12.006).
- Bozer, Y.A., Meller, R.D., 1997. A reexamination of the distance-based facility layout problem. *IIE Transactions* 29, 549–560. doi:[10.1080/07408179708966365](https://doi.org/10.1080/07408179708966365).
- Brown, G.G., DeGrange, W.C., Price, W.L., Rowe, A.A., 2017. Scheduling combat logistics force replenishments at sea for the US Navy. *Naval Research Logistics* 64, 677–693. doi:[10.1002/nav.21780](https://doi.org/10.1002/nav.21780).
- Castro, P.M., Oliveira, J.F., 2011. Scheduling inspired models for two-dimensional packing problems. *European Journal of Operational Research* 215, 45–56. doi:[10.1016/j.ejor.2011.06.001](https://doi.org/10.1016/j.ejor.2011.06.001).
- Chazelle, B., 1983. The bottom-left bin-packing heuristic: An efficient implementation. *IEEE Transactions on Computers* 32, 697–707. doi:[10.1109/TC.1983.1676307](https://doi.org/10.1109/TC.1983.1676307).
- Chen, M., Peng, X., Tang, X., 2025. Filtered beam search algorithm for the two-dimensional rectangular packing problem. *International Transactions in Operational Research* 32, 3729–3755. doi:[10.1111/itor.70010](https://doi.org/10.1111/itor.70010).
- Christensen, H.I., Khan, A., Pokutta, S., Tetali, P., 2017. Approximation and online algorithms for multidimensional bin packing: A survey. *Computer Science Review* 24, 63–79. doi:[10.1016/j.cosrev.2016.12.001](https://doi.org/10.1016/j.cosrev.2016.12.001).
- Coffman, E.G., Galambos, G., Martello, S., Vigo, D., 1999. Bin packing approximation algorithms: Combinatorial analysis, in: *Handbook of Combinatorial Optimization: Supplement Volume A*. Springer, pp. 151–207. doi:[10.1007/978-1-4757-3023-4_3](https://doi.org/10.1007/978-1-4757-3023-4_3).
- Cravo, G.L., Amaral, A.R.S., 2019. A GRASP algorithm for solving large-scale single row facility layout problems. *Computers & Operations Research* 106, 49–61. doi:[10.1016/j.cor.2019.02.009](https://doi.org/10.1016/j.cor.2019.02.009).

- Derigs, U., 1985. Using confidence limits for the global optimum in combinatorial optimization. *Operations Research* 33, 1024–1049. doi:[10.1287/opre.33.5.1024](https://doi.org/10.1287/opre.33.5.1024).
- Dexterity AI, 2023. Dexterity AI and FedEx unveil first-of-its-kind robotics trailer loading technology. URL: <https://www.dexterity.ai/blog/dexterity-ai-and-fedex-unveil-first-of-its-kind-robotics-trailer-loading-technology/>. accessed: 2025-01-30.
- DHL, 2023. To optimize supply chains, optimize the warehouse. *The Wall Street Journal* URL: <https://partners.wsj.com/blue-yonder/the-new-supply-chain/to-optimize-supply-chains-optimize-the-warehouse/>. accessed: 2025-01-30.
- Drezner, Z., 1987. A new heuristic for the quadratic assignment problem. *Naval Research Logistics* 34, 321–333. doi:[10.1002/1520-6750\(198706\)34:3<321::AID-NAV3220340304>3.0.CO;2-H](https://doi.org/10.1002/1520-6750(198706)34:3<321::AID-NAV3220340304>3.0.CO;2-H).
- Drira, A., Pierreval, H., Hajri-Gabouj, S., 2007. Facility layout problems: A survey. *Annual Reviews in Control* 31, 255–267. doi:[10.1016/j.arcontrol.2007.04.001](https://doi.org/10.1016/j.arcontrol.2007.04.001).
- Erbayrak, S., Özkır, V., Yıldırım, U.M., 2021. Multi-objective 3d bin packing problem with load balance and product family concerns. *Computers & Industrial Engineering* 159, 107518. doi:[10.1016/j.cie.2021.107518](https://doi.org/10.1016/j.cie.2021.107518).
- Filella, G.B., Trivella, A., Corman, F., 2023. Modeling soft unloading constraints in the multi-drop container loading problem. *European Journal of Operational Research* 308, 336–352. doi:[10.1016/j.ejor.2022.10.033](https://doi.org/10.1016/j.ejor.2022.10.033).
- Fisher, R.A., Tippett, L.H.C., 1928. Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society* 24, 180–190. doi:[10.1017/S0305004100015681](https://doi.org/10.1017/S0305004100015681).
- Fontaine, P., Minner, S., 2023. A branch-and-repair method for three-dimensional bin selection and packing in e-commerce. *Operations Research* 71, 273–288. doi:[10.1287/opre.2022.2369](https://doi.org/10.1287/opre.2022.2369).
- Garey, M.R., Johnson, D.S., 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, San Francisco.
- General Dynamics Land Systems, 2025. M1A2 Abrams Main Battle Tank Specifications. <https://international.gdls.com/english/products/ABRAMS/M1A2.pdf>. Accessed: 2025-03-01.
- Glover, F., 1989. Tabu search—part I. *ORSA Journal on Computing* 1, 190–206. doi:[10.1287/ijoc.1.3.190](https://doi.org/10.1287/ijoc.1.3.190).
- Golden, B.L., Alt, F.B., 1979. Interval estimation of a global optimum for large combinatorial problems. *Naval Research Logistics Quarterly* 26, 69–77. doi:[10.1002/nav.3800260108](https://doi.org/10.1002/nav.3800260108).
- Goodman, S., Pohl, J.G., 2003. ICODES: A Ship Load-Planning System, in: *Proceedings of the 13th International Ship Control Systems Symposium (SCSS)*, Orlando, FL. pp. 1–10.

- Grosse, E.H., Glock, C.H., Neumann, W.P., 2017. Human factors in order picking: a content analysis of the literature. *International Journal of Production Research* 55, 1260–1276. doi:[10.1080/00207543.2016.1186296](https://doi.org/10.1080/00207543.2016.1186296).
- Gu, J., Goetschalckx, M., McGinnis, L.F., 2007. Research on warehouse operation: A comprehensive review. *European Journal of Operational Research* 177, 1–21. doi:[10.1016/j.ejor.2006.02.025](https://doi.org/10.1016/j.ejor.2006.02.025).
- Henry, S.M., Hoffman, M.J., Waddell, L.A., Muldoon, F.M., 2023. Holistic fleet optimization incorporating system design considerations. *Naval Research Logistics* 70, 675–690. doi:[10.1002/nav.22115](https://doi.org/10.1002/nav.22115).
- Hopper, E., Turton, B.C., 2001. An empirical investigation of meta-heuristic and heuristic algorithms for a 2D packing problem. *European Journal of Operational Research* 128, 34–57. doi:[10.1016/S0377-2217\(99\)00357-4](https://doi.org/10.1016/S0377-2217(99)00357-4).
- Hübner, A., Holzapfel, A., Kuhn, H., 2015. Operations management in multi-channel retailing: an exploratory study. *Operations Management Research* 8, 84–100. doi:[10.1007/s12063-015-0101-9](https://doi.org/10.1007/s12063-015-0101-9).
- Iori, M., De Lima, V.L., Martello, S., Miyazawa, F.K., Monaci, M., 2021. Exact solution techniques for two-dimensional cutting and packing. *European Journal of Operational Research* 289, 399–415. doi:[10.1016/j.ejor.2020.06.050](https://doi.org/10.1016/j.ejor.2020.06.050).
- Joint Chiefs of Staff, 2019. Joint Publication 3-02: Amphibious Operations. United States Department of Defense. URL: <https://www.jcs.mil/Doctrine/Joint-Doctrine-Pubs/3-0-Operations-Series/>.
- Jylänki, J., 2010. A thousand ways to pack the bin: A practical approach to two-dimensional rectangle bin packing. URL: <https://github.com/juj/RectangleBinPack/blob/master/RectangleBinPack.pdf>. technical report.
- Kirschenman, W.K., Heese, H.S., Kay, M.G., King, R.E., McConnell, B.M., 2025. Problem instances for the two-dimensional bin packing problem with multiple levels of prioritization. Dryad Repository. URL: <https://doi.org/10.5061/dryad.8pk0p2p1z>, doi:[10.5061/dryad.8pk0p2p1z](https://doi.org/10.5061/dryad.8pk0p2p1z). dataset.
- Kirschenman, W.K., McConnell, B.M., King, R.E., 2024. Automated vessel selection and combat load planning. *Army Sustainment* 56, 58–61. <https://www.lib.ncsu.edu/resolver/1840.20/44301>.
- Koopmans, T.C., Beckmann, M., 1957. Assignment problems and the location of economic activities. *Econometrica* 25, 53–76. doi:[10.2307/1907742](https://doi.org/10.2307/1907742).
- Levine, I., 2024. New AI tech in amazon vans spotlights packages to save drivers effort and time. About Amazon. URL: <https://www.aboutamazon.com/news/transportation/amazon-van-delivery-van-packages>. accessed: 2026-03-22.

- Li, X., Zhang, K., 2018. Single batch processing machine scheduling with two-dimensional bin packing constraints. *International Journal of Production Economics* 196, 113–121. doi:[10.1016/j.ijpe.2017.11.015](https://doi.org/10.1016/j.ijpe.2017.11.015).
- Lodi, A., Martello, S., Monaci, M., 2002. Two-dimensional packing problems: A survey. *European Journal of Operational Research* 141, 241–252. doi:[10.1016/S0377-2217\(02\)00123-6](https://doi.org/10.1016/S0377-2217(02)00123-6).
- Lodi, A., Martello, S., Vigo, D., 1999. Heuristic and metaheuristic approaches for a class of two-dimensional bin packing problems. *INFORMS Journal on Computing* 11, 345–357. doi:[10.1287/ijoc.11.4.345](https://doi.org/10.1287/ijoc.11.4.345).
- Lodi, A., Martello, S., Vigo, D., 2004. Models and bounds for two-dimensional level packing problems. *Journal of Combinatorial Optimization* 8, 363–379. doi:[10.1023/B:JOCO.0000038915.62826.79](https://doi.org/10.1023/B:JOCO.0000038915.62826.79).
- Meller, R.D., Narayanan, V., Vance, P.H., 1998. Optimal facility layout design. *Operations Research Letters* 23, 117–127. doi:[10.1016/S0167-6377\(98\)00024-8](https://doi.org/10.1016/S0167-6377(98)00024-8).
- Mladenović, N., Hansen, P., 1997. Variable neighborhood search. *Computers & Operations Research* 24, 1097–1100. doi:[10.1016/S0305-0548\(97\)00031-2](https://doi.org/10.1016/S0305-0548(97)00031-2).
- Murata, H., Fujiyoshi, K., Nakatake, S., Kajitani, Y., 1996. Vlsi module placement based on rectangle-packing by the sequence-pair. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 15, 1518–1524. doi:[10.1109/43.552084](https://doi.org/10.1109/43.552084).
- do Nascimento, O.X., de Queiroz, T.A., Junqueira, L., 2021. Practical constraints in the container loading problem: Comprehensive formulations and exact algorithm. *Computers & Operations Research* 128, 105186. doi:[10.1016/j.cor.2020.105186](https://doi.org/10.1016/j.cor.2020.105186).
- Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. *The Computer Journal* 7, 308–313. doi:[10.1093/comjnl/7.4.308](https://doi.org/10.1093/comjnl/7.4.308).
- Olist, 2019. Brazilian e-commerce public data set by Olist. <https://www.kaggle.com/olistbr/brazilianecommerce>. Accessed: 2025-03-27.
- Paquay, C., Schyns, M., Limbourg, S., 2016. A mixed integer programming formulation for the three-dimensional bin packing problem deriving from an air cargo application. *International Transactions in Operational Research* 23, 187–213. doi:[10.1111/itor.12111](https://doi.org/10.1111/itor.12111).
- Pérez-Gosende, P., Mula, J., Díaz-Madroño, M., 2021. Facility layout planning. an extended literature review. *International Journal of Production Research* 59, 3777–3816. doi:[10.1080/00207543.2021.1897176](https://doi.org/10.1080/00207543.2021.1897176).
- Pryor, J., Chinneck, J.W., 2011. Faster integer-feasibility in mixed-integer linear programs by branching to force change. *Computers & Operations Research* 38, 1143–1152. doi:[10.1016/j.cor.2010.10.025](https://doi.org/10.1016/j.cor.2010.10.025).

- Ruiz, R., Stützle, T., 2007. A simple and effective iterated greedy algorithm for the permutation flowshop scheduling problem. *European Journal of Operational Research* 177, 2033–2049. doi:[10.1016/j.ejor.2005.12.009](https://doi.org/10.1016/j.ejor.2005.12.009).
- Smith, W.E., 1956. Various optimizers for single-stage production. *Naval Research Logistics Quarterly* 3, 59–66. doi:[10.1002/nav.3800030106](https://doi.org/10.1002/nav.3800030106).
- Trivella, A., Pisinger, D., 2016. The load-balanced multi-dimensional bin-packing problem. *Computers & Operations Research* 74, 152–164. doi:[10.1016/j.cor.2016.04.020](https://doi.org/10.1016/j.cor.2016.04.020).
- United States Navy, 2021. Large, Medium-Speed, Roll-on/Roll-off Ships T-AKR Fact File. <https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2226172/1arge-medium-speed-roll-onroll-off-ships-t-akr/>. Accessed: 2025-03-01.
- United States Navy, 2023. Landing Craft, Air Cushion (LCAC) Fact File. <https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2170004/landing-craft-air-cushion-lcac/>. Accessed: 2025-03-01.
- UPS, 2023. Warehouse optimization: A guide to improving efficiency and reducing costs. <https://www.ups.com/us/en/supplychain/logistics-solutions/distribution.page>. Accessed: 2025-01-30.
- Vancroonenburg, W., Verstichel, J., Tavernier, K., Berghe, G.V., 2014. Automatic air cargo selection and weight balancing: a mixed integer programming approach. *Transportation Research Part E: Logistics and Transportation Review* 65, 70–83. doi:[10.1016/j.tre.2013.12.013](https://doi.org/10.1016/j.tre.2013.12.013).
- Wäscher, G., Haußner, H., Schumann, H., 2007. An improved typology of cutting and packing problems. *European Journal of Operational Research* 183, 1109–1130. doi:[10.1016/j.ejor.2005.12.047](https://doi.org/10.1016/j.ejor.2005.12.047).
- Wilson, A.D., King, R.E., Wilson, J.R., 2004. Case study on statistically estimating minimum makespan for flow line scheduling problems. *European Journal of Operational Research* 155, 439–454. doi:[10.1016/S0377-2217\(02\)00910-4](https://doi.org/10.1016/S0377-2217(02)00910-4).
- Witt, C., 2005. Package flow technology at UPS. MHL News. URL: <https://www.mhlnews.com/transportation-distribution/article/22046526/package-flow-technology-at-ups>. accessed: 2026-03-22.
- X Corps Headquarters, 1951. G-4 Summary, 15 Aug to 30 Sep 1950. Unpublished official record in X Corps War Diary Summary for Operation Chromite, 15 August to 30 September 1950, Part II, Section IV. Pagination approximately 40, relative to start of Part II. Year inferred from related map dated 8 June 1951. Approved by LTG E. M. Almond and COL J. S. Guthrie. Archival document accessed at the MacArthur Memorial Museum.
- Zanakis, S.H., 1979. A simulation study of some simple estimators for the three-parameter weibull distribution. *Journal of Statistical Computation and Simulation* 9, 101–116. doi:[10.1080/00949657908810302](https://doi.org/10.1080/00949657908810302).

Zhu, H., Ji, M., Guo, W., Wang, Q., Yang, Y., 2019. Mathematical formulation and heuristic algorithm for the block relocation and loading problem. *Naval Research Logistics* 66, 333–351. doi:[10.1002/nav.21843](https://doi.org/10.1002/nav.21843).