

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Voxel-based solution approaches to the 3D irregular packing problem

Carlos Lamas-Fernandez

Southampton Business School, University of Southampton, Southampton, UK C.Lamas-Fernandez@soton.ac.uk

Julia A. Bennell

Leeds University Business School, University of Leeds, Leeds, UK J.Bennell@leeds.ac.uk

Antonio Martinez-Sykora

Southampton Business School, University of Southampton, Southampton, UK A.Martinez-Sykora@soton.ac.uk

Research on the three-dimensional (3D) packing problem has largely focused on packing boxes for the transportation of goods. As a result, there has been little focus on packing irregular shapes in the operational research literature. New technologies have raised the practical importance of 3D irregular packing problems and the need for efficient solutions. In this work we address the variant of the problem where the aim is to place a set of 3D irregular items in a container, while minimizing the container height, analogous to the strip packing problem. In order to solve this problem we need to address two critical components; efficient computation of the geometry and finding high quality solutions. In this work, we explore the potential of voxels, the 3D equivalent of pixels, as the geometric representation of the irregular items. In this discretised space we develop a geometric tool that extends the concept of the nofit polygon to the 3D case. This enables us to provide an integer linear programming formulation for this problem that can solve some small instances. For practical size problems, we design metaheuristic optimisation approaches. Since the literature is limited, we introduce new benchmark instances. Some are randomly generated and some represent realistic models from the additive manufacturing area. Our results on the literature benchmark data and on our new instances show that our metaheuristic techniques achieve the best known solutions for a wide variety of problems in practical computation times.

Key words: 3D irregular packing, open dimension problem, voxel, metaheuristics

1. Introduction

The three-dimensional (3D) irregular packing problem consists of the efficient placement of arbitrary 3D items within a designated volume without overlapping. While 3D packing of regular shapes (boxes) such as the container loading problem and 2D irregular packing problems have received a lot of attention in the literature, the 3D irregular variants have not. This problem is of great interest to researchers due to the maturity of the field, making this problem the next significant frontier, and also due to the rise in new applications, most notably those motivated by additive layer manufacturing and 3D printing.

Problems in which the space available is scarce so not all the items can be packed are equivalent to the well-known knapsack problem. For example, prioritizing items to be carried for emergency services such as an ambulance or for humanitarian aid by helicopter.

Alternatively, the space may not be constrained but using more space increases the cost, so the problem is to pack all the items using the least space. One example arises when packaging customer orders for delivery, when the sender needs to decide what is the most suitable box or set of boxes to pack a given set of items. Fuelled by e-commerce, the business to customer deliveries are increasing. In this context, efficient packaging naturally plays a financial role, but also a sustainable one (Mangiaracina et al. 2015). . In 2019 the worldwide parcel traffic increased to over 21 billion parcels from 12 billion parcels delivered in 2015 (Mazareanu 2020), and being able to reduce waste generated would have an economic benefit and a significant positive environmental impact.

A major industry that benefits from efficient packing algorithms in three dimensions is the additive manufacturing (AM) industry. AM is the process of generating physical objects from computer designs, usually by adding layers of a certain material. Among other techniques it includes 3D printing. If more than one item is printed in a single print run, a packing software is needed to generate a layout that ensures items do not overlap and that the layout is contained within the print area. AM is a flourishing industry, that had an estimated worth of \$4 billion in 2014 and is expected to grow to over \$21 billion by 2020 (Thompson et al. 2016). The resulting configuration of

the printing layouts plays an important role in the estimation of the material cost and production time for additive manufacturing (Baumers et al. 2013). One of the issues being addressed in recent studies is focused on managing the 3D printers efficiently, which involves making decisions about minimizing the extra support required, the height of the build and other aspects that depend on the materials being used (see Griffiths et al. (2019) or Attene (2015) for examples).

The main reason for not having efficient algorithms available for solving problems that consider 3D irregular items is the lack of geometric tools available in the literature to find good placements and ensure the feasibility of the layouts, i.e., the non-overlapping condition of the items.

In this article we examine the problem of packing all items, referred to as small items, into a single box-shaped container, referred to as the large object, with a fixed width and length but undetermined height. The small items are of an irregular shape and have a fixed orientation. The objective is to place all the items within the container in non-overlapping positions, while minimizing height of the layout. According to the typology in Wäscher et al. (2007), this is an irregular three-dimensional open dimension problem (I3DODP). This problem is also commonly called the strip packing problem.

The earliest papers tackling 3D packing are motivated by additive manufacturing applications and use a polygonal mesh for the item representation. An early example is by Ikonen et al. (1997) who proposes a genetic algorithm to pack arbitrary shaped items. Their algorithm places items according to a certain order, selecting one of a given finite set of orientations and ‘attachment’ points. These points are the points where one item might make contact with another item, and they are provided as part of the input data. This problem is also tackled by Dickinson and Knopf (1998), who propose a constructive algorithm where each item is packed in a position that is selected with the aim of maximizing the density of the packing.

Egeblad et al. (2007, 2009) also use a polygonal mesh representation in what they call a general purpose algorithm. The heuristic works with an initial layout and explores the solution space by translating items either vertically or horizontally. Items are permitted to lie in overlapping

positions during the search. A more specific application is presented in Egeblad et al. (2010), where a combination of pre-processing techniques and bespoke heuristics is proposed to solve a practical knapsack problem appearing in the furniture industry. Also based on polygonal mesh, Liu et al. (2015) develops an efficient constructive algorithm that allows the rotation of items based on a minimal potential energy placement. More recently, Ma et al. (2018) consider a slightly different problem where the container has a fixed size and the objective is to maximize its utilization by selecting and packing a subset of items from a collection. Their approach consists of shrinking items to a small percentage of their original size and placing them in a starting non-overlapping position, then they define a non-overlapping space for each item, and allow them to grow, move and rotate within that space, until all objects achieve their original size.

A different line of research is the one explored in Stoyan et al. (2005), where they use phi-objects to represent the geometry (Stoyan and Yaskov 1983, Scheithauer et al. 2005, Bennell et al. 2010). Analogous to the two-dimensional case, the phi-objects are mathematical descriptions based on parametric functions. In principle, this approach can accurately represent the irregular item including curved surfaces. The theory is based on a few simple basic shapes called primary phi-objects. These have formal mathematical descriptions, which are combined into parametric functions (called phi-functions) that describe the interaction of the items given a position and orientation. As a result they can test for overlap efficiently in any orientation. More complex items are constructed through combining primary phi-objects and the associated phi-functions, which form non linear constraints in their model. Deriving these models is complex and can be simplified by using quasi-phi-functions, described in Romanova et al. (2018). With this representation, it is possible to formulate most packing problems as non-linear programs. The strength of this approach relies on their precise analytical description of the objects. However, items may need many primary phi-objects to be represented and the complexity of the resulting models makes them hard to solve. In general, for the instances available in the literature, the model can only find a local optimum and the resulting packing is not very successful compared to other simpler metaheuristic approaches, as we will see in Section 5.1.1.

Phi-objects and polygonal mesh representations both suffer from increasing complexity with the irregularity of the item. For phi-objects, complicated shapes require the combining of many primary objects to represent the final shape, while the polygonal mesh approach needs more faces and vertices with the more features one object has (irregularities, holes, etc.) in order to provide a good representation. Both, the number of primary objects in a phi-object and the number of vertices or faces in a mesh have a direct impact in the computational cost of the packing algorithms using them. To overcome this, some researchers opt for approximating the three-dimensional models by discrete sets of smaller regular shapes. One frequently adopted way of discretizing the shapes is to approximate them by cubes, usually referred to as voxels.

Approximating the small items using voxels is used in Hur et al. (2001), where the authors develop a bottom-left approach for cylinder containers, and use a genetic algorithm to search over the space of packing sequences. Also using voxels, Jia and Williams (2001) propose a simulation-based packing algorithm motivated by particle packing. In their algorithm, particles can randomly move and rotate as long as they do not overlap with each other. In a later work, this algorithm was made more computationally efficient by Byholm et al. (2009). They take advantage of the discretized space to employ some computational tricks with the shape representation. This includes, for example, removing some voxels that are not going to play a role in the final packing result; resulting in reduced computational time. de Korte and Brouwers (2013) also deal with particle packing in their investigation of the packing of spheres, which they approximate by voxels.

One important parameter affecting the decomposition or discretization is the resolution, or the size of the basic units used to approximate the item. If it is too large, the fidelity of the approximation will be poor and some geometrical aspects can be lost. For example, holes or concavities smaller than a voxel cannot be represented. On the other hand, if the resolution is small, giving high fidelity, the model will require a large amount of computer memory and the packing algorithms will be very slow due to more operations being performed. Some approaches (Cagan et al. 1998, Edelkamp and Wichern 2015) use tree structures as a way to overcome this problem. The

tree structure has very coarse representations at the top, but can be made finer later in the tree if the features require it. If the basic volume units used are cubes, these trees are called octrees and each cube is divided in eight identical smaller cubes at each level.

In Edelkamp and Wichern (2015) they approximate shapes by spheres that are organised in a tree structure. This representation is then used in a simulated annealing algorithm that finds 3D printing layouts of irregular objects. The algorithm takes advantage of the use of spheres to allow free rotation of items. Cagan et al. (1998) use rectangular solids (not necessarily cubes) for the decomposition and develops a simulated annealing solution method based on simple movements and rotations in order to find efficient component layout for various industrial designs (e.g. a car engine or a heat pump).

Voxelization and octrees are used in a variety of applications due to their power for dealing with complicated shapes, and are an active topic of research in their own right (Baert et al. 2013, Schwarz and Seidel 2010), including closely related variations such as chain codes (Lemus et al. 2015, Sánchez-Cruz et al. 2014).

Our chosen approach is discussed in Section 2, where we introduce the problem formally, describe our geometric approach and introduce an integer programming model. Since this model is too complex for practical use, we also investigate metaheuristic approaches. In Section 3 we provide the building blocks for developing metaheuristic algorithms, including a constructive algorithm and different local search neighborhoods. These neighborhoods explore two different problem representations, one based on an item sequence that is decoded by different placement rules, and one based on the item position in the layout. Based on these building blocks, we develop three different packing algorithms that are presented in Section 4.

Along with the lack of previously published work on the problem, there is also a lack of benchmark data. We have gathered together the main data sets found in the literature, and added instances found in online 3D printing data sets. Furthermore, we developed an instance generator to provide a wide range of data set with specific characteristics. The generator and the benchmark instances are described in Section 5 along with a full discussion of the results of our computational experiments.

The contributions of this work are the following. We explore the potential of voxels to represent irregular items for 3D packing problems. We develop effective solution methods for the I3DODP that can solve large problems. We address the lack of data by providing a rich set of benchmark data sets and an instance generator. In adopting a voxel representation of the items to be packed, we develop a new geometric tool analogous to the nofit polygon for two-dimensional irregular packing, called the nofit voxel. Finally we provide a new mathematical programming formulation for the I3DODP and test its capabilities.

2. Voxelised three-dimensional packing

In this section we introduce formally the optimization problem we are solving, based on a discrete voxel space. We also discuss geometric considerations and introduce tools for handling the non-overlap constraints.

2.1. Problem description

We have a set of n small items $\mathcal{I} = \{p_1, p_2, \dots, p_n\}$ that need to be placed inside a large object \mathcal{C} (the container). The container has a rectangular base and a variable height (z -coordinate) and our objective is to find a set of n placements $L = \{l_1, l_2, \dots, l_n\} \subset \mathbb{Z}^3$ for all the items in \mathcal{I} ensuring that no two items overlap and that all items lay within the container, while minimizing the container height. The orientation of the items is fixed.

2.2. Voxel representation

In order to represent the irregular items from \mathcal{I} , we divide their bounding box into identical small cubes which are called voxels and use a binary code to distinguish which ones are part of the item and which ones are not. Having a rectangular shape (bounding box) as a support, enables us to encode this information in a structure analogous to a binary matrix. For each item, the voxel in the bottom-left-back corner of its bounding box is called the reference voxel. To locate the items in the container, it is sufficient to determine the position of the reference voxels. In Figure 1 we show an example of an item and its reference voxel. We also consider the bottom-left-back corner

of the container as the origin of the coordinate system, so all the points inside the container have positive coordinates.




Figure 1 Example of an irregular item represented by voxels with its reference voxel highlighted (solid black)

2.3. Constraint handling

There are two main constraints for the strip packing problem; the containment of small items within the boundary of the container (containment constraints) and no overlap between small items (no-overlap constraints). In this case, we also add the constraint of fixed orientation of the items. Since the containers we use are rectangular, the containment constraints are trivial. The reference voxel of each item has a maximum permitted x and y positions, which corresponds to the size of the container in the x or y dimension minus the magnitude of the item along that dimension. For convenience, we set a maximum value for z that is large enough to easily contain all the small items. We denote the positions that lay inside the container for the item $p_i \in \mathcal{I}$ as the set $\mathcal{IFV}_{p_i, \mathcal{C}}$, the inner-fit voxel of the item p_i and the container \mathcal{C} . Note that this is a similar concept to the inner-fit rectangle in two dimensions, see Gomes and Oliveira (2006) for details.

Non-overlapping constraints between items are, however, more difficult to model, as both of their shapes are arbitrary. If we look at the two-dimensional literature, the nofit polygon (NFP) is a concept that has been consistently used in 2D irregular packing problems since it was first developed in the sixties by Art (1966) who described it as the envelope. Informally, the NFP is derived from two polygons and represents the relative positions of these polygons in terms of

whether they overlap, represented by the interior of the NFP, touch, represented by the boundary of the NFP, or are separated, represented by the outside of the NFP. A more in-depth definition and a review of various algorithms for calculating the NFP can be found in Bennell and Oliveira (2008). It is worth noting that NFPs have a close relationship with Minkowski sums. For two arbitrary polygons A and B, it can be shown that $A \oplus -B = \{a - b : a \in A, b \in B\} = NFP_{AB}$ (See for example Milenkovic et al. (1992) or Bennell et al. (2000)).

To test the overlap between items, we define the nofit voxel (NFV), which is analogous to the union of the NFP and its interior. In general, dividing items into discrete cubes allows us to perform very quick intersection tests to identify overlap between items. Before formally defining the NFV, we first describe the simple overlap test for items represented by voxels.

Given two items $p, q \in \mathcal{I}$, with their reference voxels located at $l_p = (p_x, p_y, p_z)$ and $l_q = (q_x, q_y, q_z)$ respectively, we say that they are in a non-overlapping position if either:

1. Their bounding boxes do not intersect.
2. In the intersection of their bounding boxes, no two voxels from the items coincide in the same position. For a binary matrix where an entry is positive if it is a voxel of an item, this is equivalent of saying no two positive values of their binary matrices are in the same position.

The second condition can be simply tested by checking the corresponding elements of the matrices that represent the voxels of the two items and stopping at the first coincidence. Testing overlap in this order leads to an efficient implementation as, unless two items are very close in the layout, the first condition, which is very simple to test, will be sufficient to conclude that there is no overlap between them.

Building on this idea, it is trivial to see that if l_p and l_q lead to feasible positions for p and q , then for any $\alpha = (\alpha_x, \alpha_y, \alpha_z) \in \mathbb{Z}^3$, the points $l_p + \alpha = (p_x + \alpha_x, p_y + \alpha_y, p_z + \alpha_z)$ and $l_q + \alpha = (q_x + \alpha_x, q_y + \alpha_y, q_z + \alpha_z)$ also lead to feasible positions for p and q .

In other words, a non-overlapping position is maintained as long as the relative position of the items does not change. Following this idea, we define the NFV for two items p and q , which are

represented by voxels, as a set of the overlapping relative positions of the items p and q . More formally,

DEFINITION 1. The **Nofit voxel** of p and q is a set $\mathcal{NFV}_{p,q} \in \mathbb{Z}^3$ with the property that, if $l_p = (0, 0, 0)$ and $l_q \in \mathcal{NFV}_{p,q}$, p and q intersect.

In Figure 2 we illustrate the nofit voxel of two irregular items.




Figure 2 Two arbitrary items, p and q and their nofit voxel $\mathcal{NFV}_{p,q}$

If the NFV of two items is known, it can be used to determine if two items overlap by testing whether the reference point of q is inside the NFV. However, in Definition 1, we require that the location of p is fixed at the origin. Let us now consider the case where p is located at an arbitrary point l'_p and q is located at l'_q . We can still use the NFV to test for overlap. In this case they intersect only if $l'_p - l'_q \in \mathcal{NFV}_{p,q}$. This test is more efficient than the simple overlap test described above. Since this is a recurrent test during the remainder of the paper, we introduce the simplified notation $\mathcal{NFV}_{p,q}(p_x, p_y, p_z)$ for it, defined in equation (1).

$$(q_x, q_y, q_z) \in \mathcal{NFV}_{p,q}(p_x, p_y, p_z) \iff (q_x - p_x, q_y - p_y, q_z - p_z) \in \mathcal{NFV}_{p,q} \quad (1)$$

To compute the NFV of two items p and q , $\mathcal{NFV}_{p,q}$, we can simply keep p fixed at the origin and perform an overlap test between p and q in all of their possible overlapping positions (i.e. the positions where the bounding boxes of p and q overlap). When a test is positive, the point is added to the $\mathcal{NFV}_{p,q}$ and the process continues until all points are tested.

Note that, if $\mathcal{NFV}_{p,q}$ is known, it is trivial to calculate $\mathcal{NFV}_{q,p}$ using the following property:

$$\mathcal{NFV}_{q,p}(i, j, k) = \{-\mathbf{a} : \mathbf{a} \in \mathcal{NFV}_{p,q}(i, j, k)\} \quad (2)$$

Thanks to the property in equation (2), the calculation of NFVs only needs to be done once for each pair of item types. Once the NFVs are available, the overlap checks are merely a matter of checking if the reference voxel is in a set; provided the orientation of the items remains unchanged.

2.4. ILP formulation

In this section we describe an integer linear programming model for the problem stated in Section 2.1. The model uses binary variables to determine the reference voxel of the items and defines constraints based on the NFV information to ensure non-overlapping.

For each item p , we define a region, $\mathcal{B}(p) \subseteq \mathcal{IFV}_{p,C}$ where it can be placed. This region is a discrete set of voxels that are identified in the model by the following binary variables:

$$x_{pijk} = \begin{cases} 1, & \text{if } p \text{ is placed in } (i, j, k) \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

We mean ‘placed’ in the sense that, if x_{pijk} is set to 1, the reference voxel of the item p is located at (i, j, k) . As mentioned earlier, the reference voxel is chosen to be the corner voxel of the item’s bounding box with smallest coordinates.

The full model is as follows:

$$\text{minimize } H, \quad (4)$$

subject to

$$\sum_{(i,j,k) \in \mathcal{B}(p)} x_{pijk}(h_p + k) \leq H \quad \forall p \in \mathcal{I} \quad (5)$$

$$\sum_{(i,j,k) \in \mathcal{B}(p)} x_{pijk} = 1 \quad \forall p \in \mathcal{I} \quad (6)$$

$$x_{pijk} + \sum_{\substack{(l,m,n) \in \\ \mathcal{NFV}_{p,q}(i,j,k) \cap \mathcal{B}(q)}} x_{qlmn} \leq 1 \quad \forall p, q \in \mathcal{I}, p \neq q, \forall (i,j,k) \in \mathcal{B}(p) \quad (7)$$

$$H \in \mathbb{N} \quad (8)$$

$$x_{pijk} \in \{0, 1\} \quad \forall p \in \mathcal{I}, \forall (i,j,k) \in \mathcal{B}(p) \quad (9)$$

Where h_p is the height of item p . The first constraint (5) sets the value of the variable H , which is the total height of the container in voxels. It is defined as being higher than the highest point of any of the placed pieces (which is their reference voxel position plus their height). In equation (6) we make sure that all items are placed. The sum of the placement variables for each item p is restricted to the set $\mathcal{B}(p)$, thus this constraint also guarantees containment. Finally, equation (7) uses the NFV of two items to ensure that there is no overlap between each pair of items. There is a constraint for each position (i,j,k) where p can be placed. Then for all other items, q , we constrain the sum of the binary variables of the position of p with the binary variables for the position of q that belong to the $\mathcal{NFV}_{p,q}(i,j,k)$. In principle, these would be the points of $\mathcal{NFV}_{p,q}(i,j,k)$; but we concern ourselves only with the ones that are also part of $\mathcal{B}(q)$, which also ensures containment. Constraint (8) restricts the domain of H to natural numbers and constraints (9) defines the placement variables as binary.

This model is an extension of the Dotted Board Model from Toledo et al. (2013) for the two-dimensional problem and, to the best of our knowledge, it is the first integer linear programming model available for the discretized I3DODP.

If we define the regions $\mathcal{B}(p)$ to be $\mathcal{B}(p) = \mathcal{IFV}_{p,c}$ for all items, the optimal solution of the model would be the optimal solution of the discretized packing problem. Unfortunately, solving such a model requires a very high computational effort, which makes solving reasonable sized instances impractical, as we will show in Section 5.2, where we provide some computational experiments evaluating the capabilities of the model.

However, we can reduce the number of variables by constraining the placement options for each item. Given a starting layout, bounds on the number of voxels an item can move away from its current position provide an easier to solve model. This is a common strategy known as compaction and separation, which has been used successfully in the past in 2D irregular packing literature in conjunction with metaheuristics (see for example Bennell and Dowsland (2001) or Gomes and Oliveira (2006)). Models that use phi-functions follow a similar approach, see for example Stoyan et al. (2005).

In order to use this approach to compact the layout, the model needs to be given an initial solution and a region $\mathcal{B}(p)$ for each item p that bounds the movement, which is substantially smaller than $\mathcal{IFV}_{p,C}$. Given $\mathcal{B}(p)$, the next step is to solve the model (4)-(9) to optimality. If the height is reduced, the procedure has been successful. At this point, new $\mathcal{B}(p)$ regions can be defined again, with the aim to further reduce the overall height (H).

As a separation procedure, the initial solution must be infeasible (contain overlap). In this case the height of the container is fixed at H thus removing the constraints in (5). The regions $\mathcal{B}(p)$ are defined to respect the fixed height and the objective function is to remove overlap and find a feasible solution. This separation procedure is used as one of our neighborhoods in the Variable Neighborhood Search presented in section 4.3.

3. Building blocks of the 3D packing heuristics

In this section we examine a number of components that are the building blocks for the solution algorithms we design to solve the I3DODP. We first describe a constructive procedure to generate initial solutions. Since the procedure represents the solution as a sequence of items, and determines an item's position using placement rules, we are able to use sequence-based neighborhoods within a search heuristic. Given an initial solution, an alternative approach is to define neighborhoods that work with a solution representation that uses the position of the items in a layout. This approach usually requires a relaxation of the overlap constraint or containment constraint. For this case, we describe neighborhoods that work on the layout of complete solutions, where a strategic oscillation

technique and a new objective function related to overlap guide the search.

The two approaches that we investigate align with the strategies presented in Bennell and Oliveira (2009) for 2D irregular packing, working with sequence representations and with layout representations.

3.1. Constructive algorithm

We propose a constructive algorithm (CA) based on a bottom-left-back strategy. The idea is to place the items one at a time, ensuring their placement minimizes first the z coordinate, followed by the x and then the y coordinates. For a packing problem with a collection of items $\mathcal{I} = \{p_1, p_2, \dots, p_n\}$ and a container \mathcal{C} , let $S = \{s_1, s_2, \dots, s_n\}$ be an ordered sequence of the items to be placed. This sequence could be a random permutation of \mathcal{I} , or a specific sorting of the items, for example by decreasing volume.

The algorithm starts by placing the first item in the sequence, s_1 , at location $l_1 = (0, 0, 0)$. This placement is always feasible, because of our definition of the reference points and the coordinate system of the container. Then, the next item in the sequence, s_i , is placed in the bottom left back position $l_i = (x, y, z)$ such that the item lays within the container, i.e. $l_i \in \mathcal{IFV}_{s_i, \mathcal{C}}$ and it does not overlap with any other placed item, i.e. $l_i \notin \mathcal{NFV}_{s_j, s_i}(l_j)$, $\forall j < i$. To find position l_i , the algorithm tests points from $\mathcal{IFV}_{s_i, \mathcal{C}}$ until it reaches a valid placement. If an item of the same type has been placed earlier, the search starts from the last valid placement to avoid retesting the same points of $\mathcal{IFV}_{s_i, \mathcal{C}}$ previously found to be infeasible.

The points are evaluated starting from the lowest z values, followed by the lowest x and the lowest y , so the first valid position found is the one that lays at the bottom-left-back-most possible position. This step can be modified easily if the x or y position are sought in reverse order (i.e. from highest value to lowest); in that case the algorithm rule would be called ‘right’ instead of ‘left’ or ‘front’ instead of ‘back’. In fact, each item in the sequence can be packed using a different rule. To acknowledge this, we introduce the notation $R = \{r_i\}$ for a set of rules, where r_i represents one of the following: ‘left-back’, ‘left-front’, ‘right-back’ or ‘right-front’. A solution of the algorithm is then

given by $L = CA(S, R)$, where items have been placed in the order determined for the sequence S , in the bottom-most possible position and according to the corresponding rule from R .

While the algorithm is deterministic, different sequences S and different placement rules R can result in different packing layouts. This opens the possibility to search over the sequence of items to be placed. We explore this possibility in the next section with the sequence-based neighborhoods.

3.2. Sequence-based neighborhoods

Searching over the sequence of items has been a common topic of research in 2D irregular packing problems, see for example Gomes and Oliveira (2002), who implement a 2-exchange (swap) heuristic to search over a item sequence. Dowsland et al. (1998) and Abeysooriya et al. (2018) use a technique called Jostle, that constructs solutions iteratively and extracts the placement sequence from the layout generated on the previous step.

In this work, we identify two types of sequence based neighborhoods, the sequence swap neighborhood and the rule change neighborhood. In section 4.1 we describe a procedure to exploit them in an Iterated Local Search algorithm.

The sequence swap neighborhood consists of all the solutions obtained by swapping two items of different type in the sequence S . If $L = CA(S, R)$ and $S = \{s_1, s_i, \dots, s_j, \dots, s_n\}$, $i < j$

$$N_S(L) = \{\{CA(S', R)\} : S' = \{\dots, s_j, \dots, s_i, \dots\} \forall s_i, s_j \in S, s_i \neq s_j, i < j\} \quad (10)$$

The rule change neighborhood consists in changing of one rule for one item in the packing sequence. More formally, if $L = CA(S, R)$, then the neighborhood is defined by

$$N_R(L) = \{\{CA(S, R')\} : R' = \{\dots, r_{i-1}, r'_i, r_{i+1}, \dots\} : r'_i \in R \setminus r_i\} \quad (11)$$

Each rule is effectively creating a different construction heuristic and this space of solutions can be searched algorithmically. This kind of search is sometimes referred to as a hyper-heuristic, see for example Burke et al. (2010), Terashima-Marín et al. (2010) for examples in 2D irregular packing using a polygonal representation or Mundim et al. (2018) for a framework using a discretized approach.

3.3. Layout-based neighborhoods

A layout-based neighborhood of a solution L is a set of solutions $N(L)$ that are reachable from L by applying a perturbation or move. This movement is no longer in the sequence or in the set of rules, but a change of position of one or more items in the packing layout. As opposed to the sequence-based neighborhoods, the solutions in these neighborhoods might not always be feasible, as some moves might introduce overlap between items.

We distinguish two types of layout-based neighborhoods. The first type involves moving an item to a nearby position and is called *single item neighborhoods*. The second type involves more dramatic changes, such as swapping the positions of two items in the layout and is called *full solution neighborhoods*. The reason for this distinction is that the single items neighborhoods are usually quicker to evaluate and provide solutions of similar quality to the current one, since only one item is moved. However, the full solution neighborhoods are necessary to create kicks in the solution that help to escape local optima.

Axis aligned direction neighborhood

The axis aligned neighborhood of an item p , $N_A(L, p, \delta)$ includes all the solutions of the form $L = \{l_1, l_2, \dots, l_n\}$ where the reference point of the item p , l_p , is substituted by another point from $\mathcal{IFV}_{p,c}$ that is reachable by translating l_p by a fixed amount of voxels between 0 and δ , in an axis aligned direction. More formally,

$$N_A(L, p, \delta) = \left\{ \{l_1, \dots, l_{p-1}, l'_p, l_{p+1}, \dots\} : l'_p \in \bigcup_{d=-\delta}^{\delta} \bigcup_{i=1}^3 \{l_p + de_i\} \cap \mathcal{IFV}_{p,c} \right\} \quad (12)$$

where e_i denotes the vectors of the standard basis in \mathbb{R}^3 . The parameter δ determines the number of voxels the reference point can be displaced. The special case where we consider all the points within the container in any axis aligned direction is denoted by $\delta = \delta_{max}$. In this case, the neighborhood is similar to the translation neighborhood proposed by Egeblad et al. (2007, 2009) for polyhedral items. The Axis aligned direction neighborhood is illustrated in Figure 3.

Note that the neighbors of a feasible solution might be infeasible with respect to the overlap constraint, but they will not violate the containment constraint, as we impose the condition that items can only move within their inner-fit voxel.




Figure 3 neighborhood of axis aligned directions, $\delta = 1$ (left) and $\delta > 1$ (right). The red point is the original reference point l_p and the grey points are the possible reference points in the neighborhood.

Enclosing cube neighborhood

The enclosing cube neighborhood contains all the valid points in a cube centred at the reference point and with sides of $2\delta + 1$ voxels, as shown in Figure 4. More formally,

$$N_C(L, p, \delta) = \left\{ \{l_1, \dots, l_{p-1}, l'_p, l_{p+1}, \dots\} : l'_p \in \bigcup_{d \in \{-\delta, \dots, \delta\}^3} \{l_p + d\} \cap \mathcal{IFV}_{p,c} \right\} \quad (13)$$




Figure 4 neighborhood of enclosing cube, $\delta = 1$. The grey points are the reference points in the neighbors, while the original reference point is located in the centre of the cube.

For a large δ , this is equivalent to the complete inner-fit voxel of the item. In this case, the neighborhood can be seen as an insert move. We denote this situation by $\delta = \delta_{max}$.

item swap neighborhood

The swap neighborhood of an item is a full solution neighborhood. It represents all the possible swaps the item can make with other items. Some of the swaps might violate the containment constraint, since the swapped items could end up with parts outside the container. If this happens, the item is moved inwards along the axis (one voxel at a time) until the point is in its inner-fit voxel. This position, denoted by $l_p^{(q)}$ (the resulting position of p after swapping with q), is defined as follows,

$$l_p^{(q)} = \arg \min_{l' \in \mathcal{IFV}_{p,c}} \|l_q - l'\|_1 \quad \forall p, q \in \mathcal{I}; p \neq q \quad (14)$$

Then, the neighborhood contains all the possible swaps of a item p with others,

$$N_{\text{swap}}(L, p) = \left\{ \{ \dots, l_{p-1}, l_q^{(p)}, l_{p+1}, \dots, l_{q-1}, l_p^{(q)}, l_{q+1}, \dots \} : q \in \mathcal{I}, q \neq p \right\} \quad (15)$$

3.4. Strategic Oscillation

A search algorithm using a layout-based neighborhoods is likely to produce solutions that contain overlap. While allowing overlap helps navigate the solution space, it makes the problem effectively a bi-objective problem, since the overlap constraint is relaxed and added to the objective function. Solutions containing overlap are not feasible with respect to the original problem so we need to produce solutions with zero overlap without compromising on solution quality. Strategic oscillation (Glover and Marti 2011), has been used extensively in the two and three-dimensional irregular packing literature, see for example Bennell and Dowsland (1999), Umetani et al. (2009), Imamichi et al. (2009) or Egeblad et al. (2009), to navigate this problem. The idea is that the algorithm oscillates between focusing primarily on completely removing the overlap and then primarily on reducing the height of the layout. Certain criteria allow the switch in focus during the course of the algorithm.

In each iteration of the strategic oscillation the algorithm works on an overlapping layout with fixed height, and tries to remove overlap. Once this has been achieved or the maximum number of local search iterations is reached, the height is modified as follows. If the solution still has overlap, the height will increase by one voxel and the algorithm will try to resolve the overlap with more space. If it succeeds, the height is reduced by a certain percentage (controlled by the parameter `dec_p` in our computational experiments) and the algorithm is run again, with the aim of finding non-overlapping positions at a lower height.

Since the reduction of height leaves some items protruding above the top of the container, their position needs to be adjusted. We have implemented two strategies to achieve this: *move down* and *random reinsert*. With the *move down* strategy, protruding items are simply moved downwards the minimum amount until they fully fit inside the new container. Instead, with the

random reinsert protruding items are removed from the layout and inserted in a random position of their inner-fit voxel. In the computational experiments, the strategy is selected by means of the algorithm parameter `S0_mode`.

3.5. Objective function

While our objective is to minimize the height of the final layout, if we search over the layout and use strategic oscillation, intermediate solutions will have overlap. For this case, we propose an objective function whose main goal is to understand when a solution is better than another, in terms of leading to a non-overlapping solution. In the nesting literature, typical objective functions include measuring the penetration depth of a pair of items (Bennell and Dowsland 1999, Umetani et al. 2009, Imamichi et al. 2009) or the exact amount of overlap (area in 2D or volume 3D) for each pair of items (Egeblad et al. 2009), among others. Our objective function is calculated based on the volume of the overlap of the bounding boxes,

$$F(\mathcal{I}, l_1, l_2, \dots, l_n) = 1 - \frac{1}{\binom{n}{2}} \sum_{i,j \in \mathcal{I}, i < j} \frac{Vol(B(i, l_i) \cap B(j, l_j))}{V_{\max}}, \quad (16)$$

where $B(i, l_i)$ is the bounding box of item i located in position l_i and $Vol()$ is a function that returns the voxel volume (number of voxels) of its argument. The constant $V_{\max} = \max_{i \in \mathcal{I}} Vol(B(i))$ is the maximum volume of any bounding box present in the instance and is used to scale the function value between 0 and 1. Dividing the sum by $\binom{n}{2}$, the number of pairs of items in the solution, is used to average the overlap of the items. This function is very quick to evaluate and provides an intuitive idea of ‘how difficult’ it is to separate a pair of overlapping items.

4. Search algorithms

Using the building blocks described in the last section, we propose three different metaheuristic algorithms. The first is an iterated local search, which is based on searching over the item sequence and generating the layout by applying the constructive algorithm. The next two approaches are a

tabu search and a variable neighborhood search, which are based on searching over the layouts of complete solutions.

4.1. Iterated local search

Iterated local search is a metaheuristic algorithm that iteratively searches neighborhoods, accepting improving moves, to find local optima. At each local optima the search moves to a different part of the solution space by performing a kick (Lourenço et al. 2010).

We implement this algorithm to take advantage of the different rules that can be used in our constructive algorithm. The initial solution starts with a sequence of items sorted in order of decreasing height and the solution is constructed using the left-back placement rules to place all of the items. For the local search we use the rule change neighborhood, $N_R(L)$, that involves changing one of the two placement rules (left/right or back/front) of an item and repacks using the constructive algorithm.

The kicks consist of performing a number of position swaps in the item order sequence (and their placement rules) at random. The number of swaps to be performed is controlled by an algorithm parameter, n_{swaps} , which is determined in Section 5.3.

While searching over the sequence can produce quick efficient layouts, the solution space does not contain all possible solutions and some good layouts might be unreachable using these techniques. To overcome this, we propose two other algorithms that work on the layout of full solutions.

4.2. Tabu Search

Tabu search (TS), introduced by Glover (1989), is one of the first metaheuristic and has been successfully applied to various combinatorial optimization problems, including irregular packing (Błazewicz et al. 1993, Bennell and Dowsland 1999). It uses neighborhoods to search the solution space by testing all solutions in the neighborhood of the current solution and moving to the best neighbor, whether it improves the objective function value or not. The solution resulting from the

move is added to a tabu list so it cannot be revisited for a certain number of moves. In practice, an attribute of the move is stored as this is more efficient and achieves the desired outcome.

TS lends itself to the layout representation since the voxel neighborhoods are finite and can be evaluated in full with a low computational effort. Maintaining a tabu list is an effective way to prevent items moving back and forth to the same positions. Our TS works within the strategic oscillation framework searching for non-overlapping solutions between each adjustment of the height constraint. Each height oscillation occurs after a predefined maximum number of accepted neighborhood moves or if a zero overlap solution is found by the TS.

The initial solution, L_0 , is produced using the constructive algorithm (Section 3.1) as $L_0 = CA(S_0, R_0)$, where S_0 is the sequence of items from the instance ordered by decreasing volume and R_0 is the set of rules that assigns ‘left-back’ to all the items. In the first step the algorithm oscillates to create some overlap, as described in Section 3.4. In each iteration, the algorithm explores the axis aligned neighborhood of the solution, with parameter $\delta = \text{d_axis}$, whose value is determined in the computational experiments. When a movement is performed, the opposite direction of the movement for that item is added to the tabu list, to avoid the same item moving backwards. The length of the tabu list, TL is also a parameter of the algorithm.

Note that to evaluate each move it is not necessary to calculate the objective function from equation (16). Instead, it is sufficient to perform an update to evaluate the affect of the change in the solution quality. This test is very quick to perform, since it only involves recalculating the overlap terms for the item that has moved. The efficiency of this step allows for the testing of a large number of movements in a reasonable computational time.

To control the height, the tabu search will run on a given solution for a number movements (determined by parameter `nMoves`) or until it completely removes the overlap, and then perform the oscillation.

4.3. Variable Neighborhood Search

Having designed a number of neighborhoods, it is natural to consider implementing a Variable Neighborhood Search algorithm (VNS). This metaheuristic, introduced by Mladenović and Hansen

(1997), makes use of a list of neighborhoods that are explored in order, and works on the assumption that a local optimum for a neighborhood might not be a local optimum in an alternative neighborhood. VNS is implemented with strategic oscillation in the same way as TS.

The algorithm starts from an initial solution (created in the same way as in the previous section) that contains some overlap, and searches the first neighborhood in the list. In our implementation, we apply a *steepest descent* rule, and therefore, all the solutions in the neighborhoods are evaluated before moving to the best improving solution.

If there are no improving neighbors, the incumbent solution is labelled as a local optima for the current neighborhood. Then, VNS selects the next neighborhood in the list and explores it until it is again stuck in a local optima. If an improvement is found, the incumbent solution moves to it and VNS returns to the first neighborhood and continues the search.

When a solution is a local optima for all the neighborhoods, the algorithm applies a kick (or a shake) and returns to the first neighborhood in the sequence. After a certain number of kicks (parameter `nK`), the algorithm performs strategic oscillation in order to increase or decrease the maximum allowed height and the process is repeated.

We have implemented VNS with three neighborhoods that are searched in this order:

- Enclosing cube neighborhood, N_C with $\delta = \text{d_cube}$
- Axis aligned neighborhood, N_A with $\delta = \text{d_axis}$
- ILP model, with $B_p = N_C$ and a small δ value.

The first two neighborhoods are explored for all the overlapping items. We have chosen to use the cube neighborhood with a small delta first, as it can provide very similar solutions by moving only one item by one voxel. The axis aligned neighborhood has the capacity to reach further away points and can often provide more dissimilar solutions, especially at early stages of the algorithm. Our experiments found this order of the first two neighborhood definitions to be more efficient because it finds the best local optima through comprehensively searching a small local neighborhood of the current solution before searching more widely. When the second neighborhood definition finds

a better solution, the search returns to the focused optimization using the first neighborhood definition. If these neighborhoods cannot improve the solution, we run the ILP model (3)-(9). The δ values used for the cube neighborhoods of the model are a parameter of the algorithm. We use a different parameter for the items that are overlapping ($\delta = \text{Ov_d}$) and for the items that are not overlapping ($\delta = \text{NOv_d}$). Both of these values are determined in Section 5.3. To avoid spending excessive computational time running the model, and taking into account that is only practical when using only small δ values, we have restricted its use to only once per strategic oscillation and only in layouts where there is at most one pair of overlapping items. The running time of the model is restricted to a maximum of 300s.

Once we have hit a local optimum for all the neighborhoods and before oscillating, we disrupt the solution with the aim of exploring a different part of the solution space. This disruption, also called a kick, is performed in two steps: a shaking of the layout and an item swap.

To perform shaking, we move all overlapping items to a random location within a neighborhood N_C . The number of voxels, δ , that defines the size of the neighborhood is set such that it increases linearly with the number of disruptions used so far in the search. This technique allows us to have smaller disruptions in the early stages in order to guide the algorithm quickly to good solutions. However, in later stages of the algorithm, the disruptions are larger, allowing the search to escape from local optima. The parameter will vary between a minimum value of δ_{min} , which is the size of the starting neighborhood, and a maximum value of δ_{max} . The actual parameter used in the iteration is calculated as $\delta = \min(\delta_{min}, \lfloor \rho(\delta_{max} + 1) \rfloor)$, where ρ is the ratio between the kicks performed and the maximum kicks allowed for the current height. Suitable values for δ_{min} and δ_{max} are explored and selected in the computational experiments in Section 5.3.

The second step of the kick consists of swapping two overlapping items. To avoid the costly procedure of testing all possible swaps, we limit the amount of tests. We test the swap of each item in the layout with a number of items from the overlap list, determined by the parameter `test_swaps`. The overlapping items are chosen at random and avoiding items of the same type. In total we test a

number between 1 and $n \times \text{test_swaps}$ swaps, and perform the one that returns the best objective function, even if it does not improve the incumbent solution.

These two steps complement each other. The shaking part creates many overlapping positions and enables the search to explore movements of items that might otherwise stay in non-overlapping positions and prevent the solution improving. The swap part produces a larger change in the layout, that might be difficult to obtain with moving items one at a time. Note that creating many overlaps will create a greater opportunity for swapping items to succeed than if we apply these two steps in a opposite order.

Each iteration of the VNS consists of a number of local searches followed by a kick. Therefore, in each iteration we will find many local optima, possibly all of them containing overlap, for the same container height. Since the kick involves swapping one or more pairs of items, there is a risk that these swaps might not be beneficial in the long term. To avoid this situation, the kick is not always applied to the incumbent solution, but to the best of the local optima found for the current height. By doing this, item swaps that do not lead to an improvement are not carried forward in the search.

5. Computational experiments

In this section we evaluate the performance of the proposed voxel-based algorithms. In Section 5.1 we describe the instances used for testing. In Section 5.2 we explore the limits of the ILP model to solve instances to optimality. In Section 5.3 we describe how we have performed the parameter configuration for the metaheuristic algorithms, along with the best configuration for each algorithm. Finally in Section 5.4 we compare the performance between the algorithms and also against the state of the art in the literature.

5.1. Test Instances

We test our algorithms across a range of different instances to compare their performance in different scenarios. We found that in the 3D literature researchers mainly use simple shapes, such as

convex polyhedra (Stoyan et al. 2005, Pankratov et al. 2015), or non-convex shapes that are defined by very few vertices (Stoyan et al. 2004, Egeblad et al. 2009, Liu et al. 2015). In our experiments, we solve some of these instances and adapt one popular instance from the two-dimensional literature with similar properties, *shapes0* from Oliveira et al. (2000). In order to test our algorithms more thoroughly, we also propose a set of randomly generated instances, called *blobs*, and two examples of complex realistic instances. Unlike standard academic instances, the *blobs* data sets are designed to include more challenging shapes, such as “smooth” surfaces or holes.

In Table 1 we summarise some of the key differences of these instances

Table 1 Instances solved and their features

Instance	Source	Convex	Holes	Items (types)
Stoyan2	Stoyan et al. (2005)	✓	✗	12 (7)
Stoyan3	Stoyan et al. (2005)	✓	✗	25 (7)
Merged i	Egeblad et al. (2009)	✗	✓	15 - 75 (15)
Experiment1	Liu et al. (2015)	✗	✓	36 (5)
Experiment i	Stoyan et al. (2004)	✗	✓	20 - 50 (10)
Blobs i	Generated	✗	✓	20 (10)
Shapes_3D	Adapted	✗	✗	43 (4)
Engine	Realistic	✗	✓	97 (56)
Chess	Realistic	✗	✗	32 (6)

5.1.1. Instances from the literature To test our algorithms, we have collected the available instances from the literature, which tend to be simpler shapes. The first instance is a set of convex polyhedra that has been introduced by Stoyan et al. (2005) and later used in Egeblad et al. (2009). It consists of seven convex shapes of similar volume that are repeated a number of times each. Stoyan et al. (2005) defined three sets of items *Example1*, *Example2* and *Example3*, but we only solve the latter two, since *Example1* contains one item that has a dimension coinciding exactly

with the width of the container and it becomes infeasible with our voxelization parameters. We call these instances *Stoyan2* and *Stoyan3*

In order to find more challenging instances, Egeblad et al. (2009) mixed this set with some items available in the work of Ikonen et al. (1997), creating the *Merged i* instances, that we solve as well. The other instances we consider are a set of non-convex polyhedra from Stoyan et al. (2004) (*Experiment2* to *Experiment5*), that has been later solved by Liu et al. (2015) and a similar set of polyhedra proposed by Liu et al. (2015), *Experiment1*. However in the original work *Experiment1* has been solved only when allowing rotation, and therefore we do not solve exactly the same problem, since our orientations are fixed.

These instances are originally represented using the polygonal mesh approach and therefore it is necessary to convert them to voxels before applying our algorithms. To perform this process, we have used the *binvox* software package (Nooruddin and Turk 2003, Min 2004 - 2019), where we apply the *exact* setting, to ensure the voxelized representation overestimates the original one and therefore our solutions are always feasible. One important parameter used in the conversion is the resolution, i.e. the size of the voxels used in the converted shapes. In order to solve each instance with the same resolution, we have chosen a voxel size such that the largest side of the container is voxelized with 150 voxels. Where there are several instances containing common items, we chose the largest side of the largest container among the instances so that the items have a common voxelization. This means that the smaller containers in the group of instances might not get an exact voxelization and therefore we round the dimensions down, ending up with a slightly smaller container. 150 voxels is chosen as it seemed a reasonable compromise between accuracy and solving speed. Nevertheless, as it will be shown in Section 5.4 the resolution chosen plays a key role in both the solution time and the quality of the results.

5.1.2. A simpler instance To complete the tests with simpler shapes, we have also adapted a classical instance from the two-dimensional irregular packing literature, the set *shapes0* from Oliveira et al. (2000). This set consists of four simple distinct shapes (one convex and three non-convex) and has been widely used in the research of the strip packing problem. To adapt these

figures to 3D we maintained the original proportions of the original shapes and voxelized them in a grid of one unit. For two of the shapes, we preserved their symmetry along the third axis. For the remaining two, we extruded the shape along the third axis so they had a size of 6 voxels. We call our instance *Shapes_3D* and the final result can be seen in Figure 5. The container for this set has sides 20×20 and we maintained the same item description (43 items in total) and fixed the orientation.




Figure 5 Instance *Shapes_3D*

5.1.3. Realistic instances We also solve two instances based on models downloaded from a popular 3D printing online community. In order to solve a meaningful problem, we chose two designs that have shapes with a certain complexity and that are themselves a collection of different parts. We chose a realistic model of an engine by Harrell (2015) (See Figure 6 for an example of an item already voxelized) and a chess set by Edwards (2014).

Both of these instances were available as triangular mesh files. We converted them using the same process as described in Section 5.1.1, but as there was no container size provided, we chose the resolution differently. Based on the size of the items and since they were available in different scales, we decided to pack the *Engine* in a container with a base of 500×500 voxels, where one voxel represents 6 units and the *Chess* set in a container of size 50×50 , where one voxel is mapped to one unit of the original model.




Figure 6 Example item from the *Engine* instance

5.1.4. Randomly generated instances Our random instance generator extends to three dimensions the shape generation approach presented in Robidoux et al. (2011). We illustrate some of its steps in Figure 7 in an example to generate a 2D shape, since it is easier to appreciate in two dimensions. The following steps below describe the method for a 3D item:

Step 1 Start with a three-dimensional matrix of zeros, that will be the container of the future item.

Step 2 Randomly assign 1 to the value of some elements in the matrix (Figure 7 (a)).

Step 3 Connect the points in a random order with discretized lines of a certain thickness. The last point is connected with the first in order to create a loop (Figure 7) (b)

Step 4 Apply a Gaussian blur to the matrix. This step changes the values of the elements so they are between 0 and 1 (Figure 7) (c)

Step 5 To convert to binary values again, apply a threshold, where values less than a parameter (around 0.005) are set to zero and values over that parameter are set to one (Figure 7) (d).

There are a few parameters that can be controlled in this method. The first of them is the size of the matrix where the points are drawn. A larger size will mean a smoother shape, but more costly to compute. We also need to decide how many and how to select the random elements that are initially set to one. More points give more opportunities for the shape to have irregularities, but too many will fill up the space of the shape too much and it can lose some interesting features, such as holes. For sampling the points, we have chosen a random uniform distribution in a cube,




Figure 7 Example of the steps to generate a ‘blob’ in two dimensions. In (a) 9 points are drawn randomly, in (b) they are connected in a closed loop, in (c) a Gaussian blur is applied (the values between 0 and 1 are represented by the shade of grey) and in (d) values are set to either 0 or 1, depending on a threshold value.

but there are other possible distributions. Robidoux et al. (2011) suggest using distributions in a disc (uniform or normal) if one wishes to get rotation invariant shapes, but this is not relevant for our purposes.

The way the points are connected can also impact the shape. In our case we use straight lines, but other methods such as splines could give more smooth results. Another decision is how to add ‘thickness’ to these lines. Looking again at Figure 7 (b), we see that the lines in that case have more than one pixel in thickness. To explain this, let us think about the matrix as a grid of squares (or cubes, in 3D). When connecting two points, we first paint the squares of the grid intersected by the segment that connects the centre of the two points. By painting here we mean setting to 1 the corresponding matrix elements. To add ‘thickness’, we also paint the squares surrounding this initial line, if they are within a certain distance from it in any axis aligned direction. This distance is another parameter, that we call the thickness parameter. This results in painting with a kind

of cross-hairs, the same concept as depicted in the neighborhood from Figure 3. Finally, both the standard deviation of the Gaussian blur and the threshold have also an impact on the outcome.

Following this technique, we have generated nine instances in total, with three different sizes and three shape styles for each size. We hand-picked the parameters in order to define the three different styles, that we have called ‘round’, where items are closer to spheres or ellipsoids; ‘peaked’, where shapes have more pronounced concavities and more frequently holes; and ‘neutral’, which is somehow a middle ground between the two. In Figure 8 we show one item representing each different style. The full description of the parameters used is reported in Table 2, where intervals indicate that the parameter has been randomly chosen in that interval.




Figure 8 Three items from the medium instances with style ‘round’ (left), ‘neutral’ (centre) and ‘peaked’ (right)

The matrix size in Table 2 is indicative, as applying the Gaussian blur might set to positive some values outside the original cube where the points were sampled (this happens, for example, if one of the points was in the corner of the cube). To acknowledge this fact and ensure the shapes remained smooth, we allowed a maximum size of $100 \times 100 \times 100$ for any shape, which seemed sufficient to avoid such cases.

Based on the size of the largest item from each set, we chose the size of the container base to be (in voxels) 78×78 for the small instances, 84×84 for the medium instances and 90×90 for the large instances.

Table 2 Parameters used to generate the blobs instances

<i>Name</i>	<i>Size</i>	<i>Style</i>	<i>Matrix size</i>	<i>Points</i>	<i>Thickness</i>	<i>Filter σ</i>	<i>Threshold</i>
<i>blobs1</i>	Small	Round	$25 \times 25 \times 25$	[2,20]	3	[3,5]	0.05
<i>blobs2</i>	Small	Neutral	$25 \times 25 \times 25$	[2,15]	[1,3]	[1,2]	0.01
<i>blobs3</i>	Small	Peaked	$25 \times 25 \times 25$	[3,7]	1	[1,2]	0.005
<i>blobs4</i>	Medium	Round	$35 \times 35 \times 35$	[4,30]	[4,5]	[4,5]	0.05
<i>blobs5</i>	Medium	Neutral	$35 \times 35 \times 35$	[4,25]	[3,4]	[2,3]	0.01
<i>blobs6</i>	Medium	Peaked	$35 \times 35 \times 35$	[3,10]	[1,3]	[1,2]	0.005
<i>blobs7</i>	Large	Round	$45 \times 45 \times 45$	[4,40]	[4,5]	[4,5]	0.05
<i>blobs8</i>	Large	Neutral	$45 \times 45 \times 45$	[4,3]	[3,4]	[2,3]	0.01
<i>blobs9</i>	Large	Peaked	$45 \times 45 \times 45$	[3,14]	[1,3]	[1,2]	0.005

5.2. Testing the ILP model

The model presented in Section 2.4 can provide an optimal solution if solved to optimality. However, the large number of variables for the proposed instances makes it often impractical to solve. In order to test its boundaries, we have chosen the four items from the *Shapes_3D* instance and generated all of the 155 instances that feature each item type repeated between 0 and 3 times, and whose total number of items is between four and seven. The container size was set to 20×20 voxels. We have then solved them for a maximum of 10 hours computational time, using IBM[®] ILOG[®] CPLEX[®] 12.6.1 and providing the bottom-left-back upper bound.

In eight of the instances the bottom left back algorithm found a solution such that the height used corresponds to the height of the largest item, and therefore, the solution is optimal. In Table 3 we can observe that the percentage of solved instances clearly decreases when adding items, to the point that only 43% of the instances with seven items can be solved in the given time. In Figure 9 we show one instance with seven items solved to optimality. The gap obtained when the

Table 3 Summary of the results achieved by solving the ILP model on instances with between four and seven items. **Opt** denotes the number of instances solved to optimality in each set.

Items	Instances	Opt	Percentage
4	31	27	87%
5	40	30	75%
6	44	26	59%
7	40	17	43%
Total	155	100	65%

instances were not solved to optimality is usually large, which suggest that the difficulty of solving the derived models increases quickly.




Figure 9 Initial solution from bottom left back (left, 16 voxels height) and optimal solution (right, 12 voxels height) for one instance with seven items. Solving time was 78 s.

5.3. Metaheuristic algorithms configurations

In this section we describe the parameters we have used to configure our algorithms for the computational experiments. The configuration was done by using the software package “irace” López-Ibáñez et al. (2016) with a budget of 1000 one-hour runs for each method. The single parameter tested for ILS is reported in Table 4.

In Tables 5 and 6 we report the tests performed for TS and VNS respectively. Note that the parameters for Strategic Oscillation (`S0_mode` and `dec_p`) have been tested independently for the

Table 4 Parameter tested for the ILS and the best values found by irace.

Parameter	Description	Tested Values	Best
n_swaps	Number random item swaps performed between iterations	$[1, 25] \subset \mathbb{N}$	4

both methods, as they could (although they do not, in this case) differ on their best value for different search strategies.

Table 5 Parameters tested for the TS algorithm and the best values found by irace. In strategic oscillation MD stands for *move down*, and RR stands for *random reinsert*

Parameter	Description	Tested Values	Best
nMoves	Number of tabu search movements performed before strategic oscillation	$10^i, i = 2, 3, 4, 5, 6$	10^4
d_axis	Parameter δ used for the axis aligned neighborhood (N_A)	1,3,5,10, ∞	∞
TL	Lenght of the tabu list	$[1, 30] \subset \mathbb{N}$	13
SO_mode	Strategy used to relocate protruding items during strategic oscillation	MD, RR	MD
dec_p	Percentage of height reduced in strategic oscillation	$[1, 25] \subset \mathbb{R}$	7%

The large value of `test_swaps` suggests that it is beneficial testing as many swaps as possible during the kick in VNS. The values of the parameters found by irace are the ones we use in the remainder of the computational experiments.

Table 6 Parameters tested for the VNS and the best values found by irace. In strategic oscillation *MD* stands for *move down*, and *RR* stands for *random reinsert*

Parameter	Description	Tested Values	Best
nK	Number of kicks before strategic oscillation	$[1, 100] \subset \mathbb{N}$	78
K_dmin	Value of δ_{min} in VNS kicks	$[1, 3] \subset \mathbb{N}$	1
K_dmax	Value of δ_{max} in VNS kicks	$[2, 10] \subset \mathbb{N}$	8
S0_mode	Strategy used to relocate protruding items during strategic oscillation	<i>MD</i> , <i>RR</i>	<i>MD</i>
dec_p	Percentage of height reduced in strategic oscillation	$[1, 25] \subset \mathbb{R}$	7%
test_swaps	Number of swaps tested for each item during the kick	$[1, 50] \subset \mathbb{N}$	48
d_cube	Parameter δ used for the cube neighborhood (N_C)	$[1, 10] \subset \mathbb{N}$	5
d_axis	Parameter δ used for the axis aligned neighborhood (N_A)	1,3,5,10, ∞	∞
0v_d	When running the model, δ value used for the cube neighborhood of overlapping items	$[1, 4] \subset \mathbb{N}$	1
NOv_d	When running the model, δ value used for the cube neighborhood of non overlapping items	$[1, 2] \subset \mathbb{N}$	2

5.4. Algorithm comparison

We have implemented all our algorithms in C++ and run them in a single 2.6 GHz core with 16 GB of memory, part of the IRIDIS HPC facility. All of our experiments allow one hour of computational time and we report both the best and the average results obtained after 10 runs. Since all the methods tested use nofit voxels, they have been pre-calculated and the generation and loading time is not included in the reported time of the algorithms.

It is worth noting that as result of the voxelization of items, we may be packing slightly larger items into a slightly smaller container. This has two consequences. First, some configurations of the items that may be feasible in a mesh representation may not be feasible in the voxel representation with a given fidelity, which may include the best layout in the literature. Second, comparing solutions with the literature on the basis of volume utilisation no longer makes sense as the total volume of the items is not the same. For example, for the instance *Merged1*, the total volume of our voxelized version of the items is 15% larger than the original. To address the second point, all our experiments compare the height of the packing layout.

The results for the *blobs* instances are shown in Table 7. We can observe that VNS obtains on average the best results. The average height is 162.64, which is 1.7% better than the TS and 8% better than the ILS. As for the best results, in seven out of nine instances the VNS obtained the best solution. This pattern is also observed in the *Experiment* set (Table 9), where VNS is again superior to all the other methods. All three voxel-based methods obtain the best result in the literature for these instances. The same trend is repeated on the realistic instances and the adapted *Shapes*, reported in Table 8. All methods here obtain the same value for *Shapes-3D*, although ILS seems to be more consistent. For the more realistic instances *Chess* and *Engine* VNS obtains both the best and most consistent results.

In Figure 10 we show some of the best results found for this group of instances and in Table 9 we report the results for the *Experiment* instances.

All voxel based methods perform better than both HAPE3D (Liu et al. 2015) and the phi-function based methods (Stoyan et al. 2004). Among them, VNS has the best and most consistent results. In Table 10 we report the results for the *Stoyan* and *Merged* instances.

Table 7 Comparison of results (final height, in voxels) for the blobs instances

Instance	ILS		TS		VNS	
	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>
<i>blobs1</i>	79	79.64	75	76.90	74	74.90
<i>blobs2</i>	64	65.91	59	60.90	59	59.60
<i>blobs3</i>	47	47.73	44	45.20	44	44.50
<i>blobs4</i>	211	214.18	202	205.60	187	191.70
<i>blobs5</i>	221	228.91	206	211.30	202	205.00
<i>blobs6</i>	91	94.82	88	89.10	86	88.00
<i>blobs7</i>	337	343.36	309	323.40	312	323.50
<i>blobs8</i>	357	368.82	329	338.10	330	340.50
<i>blobs9</i>	146	149.36	136	138.60	133	136.10
Av.	176.97		165.46		162.64	

Table 8 Comparison of results (final height, in voxels) for the blobs instances

Instance	ILS		TS		VNS	
	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>
<i>Shapes</i>	52	52.73	52	53.00	52	53.20
<i>Engine</i>	78	79.36	77	78.70	75	78.10
<i>Chess</i>	144	145.55	139	141.40	131	135.10

While the voxel based methods obtain better results than the phi-functions, the algorithm 3DNest, based on a polygonal mesh obtains slightly better results for all instances but one. We note that 3DNest algorithm permits a small overlap between items (0.1%). In order to distinguish whether the differences in quality for the *Merged* instances could be linked to the resolution chosen




Figure 10 Best result obtained by ILS for *Shapes_3D* (left, height of 52 voxels); by TS for *Engine* (centre, height of 77 voxels); and by VNS for *Chess* (right, height of 131 voxels).

Table 9 Comparison of results (final height) for the instances found by St04 (Stoyan et al. 2004) and HAPE3D (Liu et al. 2015)

Instance	ILS		TS		VNS		St04	HAPE3D
	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Best</i>
<i>Exp1</i>	36.40	38.58	34.13	35.12	34.27	35.37	-	31.2*
<i>Exp2</i>	29.40	31.10	28.47	30.17	28.47	29.40	32	36.9
<i>Exp3</i>	42.70	46.05	41.77	43.54	41.30	42.65	49	55.6
<i>Exp4</i>	62.30	65.35	56.23	59.62	54.83	59.52	63.21	71.8
<i>Exp5</i>	77.00	81.41	74.20	76.16	72.33	76.81	78.7	92

* The result of 31.2 for the instance *Exp1* is obtained by allowing 8 different rotations, which are not allowed in the other algorithms.

when voxelizing the items, we performed another set of tests where VNS uses finer resolutions and longer search times. We considered two extra resolutions, *fine* and *extra-fine*, where the largest side of the largest container base is voxelized with 200 and 250 voxels respectively, and a maximum time computational time of four hours. The results for these extended tests are shown in Table 11.

These illustrate how increasing the accuracy of the representation closes the gap between 3DNest and VNS, and with finer resolutions VNS is able to improve the average solution quality by 1.4%. As a result, we are finding the best known solutions for the largest instances (60 and 75 items)

Table 10 Comparison of results (final height) for the instances found in St05 (Stoyan et al. 2005) and 3DNEST (Egeblad et al. 2009)

Instance	ILS		TS		VNS		St05	3DNEST
	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Avg</i>	<i>Best</i>	<i>Best</i>
<i>Example2</i>	22.72	23.30	21.33	21.60	21.23	21.70	30.92	19.83
<i>Example3</i>	34.99	35.65	30.93	32.35	29.76	31.53	45.86	29.82
<i>Merged1</i>	20.52	21.19	18.47	19.72	18.30	18.67	-	16.99
<i>Merged2</i>	24.62	25.58	22.40	23.27	22.40	22.64	-	21.97
<i>Merged3</i>	28.73	29.48	27.02	28.16	24.62	26.26	-	24.04
<i>Merged4</i>	30.61	31.86	28.73	30.27	28.56	29.44	-	27.80
<i>Merged5</i>	34.20	35.65	30.61	31.70	30.61	31.12	-	30.13

Table 11 Extended results for the *Merged* instances, using more time and finer resolutions for VNS

Instance	VNS - <i>Standard</i>	VNS - <i>Fine</i>	VNS - <i>Extra-fine</i>	3DNest
<i>Merged1</i>	18.30	17.95	17.13	16.99
<i>Merged2</i>	22.40	22.19	22.06	21.97
<i>Merged3</i>	24.62	24.24	24.32	24.04
<i>Merged4</i>	28.56	28.21	26.47	27.80
<i>Merged5</i>	30.61	30.27	29.24	30.13
Av.	24.90	24.57	23.84	24.17

while being within less than 1% of the height in the three smaller instances. Since VNS does better on larger problems, this suggests our approach is more scalable.

The best layouts obtained for each resolution for the instance *Merged5* can be seen in Figure 11.

6. Conclusion

In this article we explore the potential of voxel based approaches for solving the three-dimensional irregular open dimension problem. We develop the concept of nofit voxel, in which we base our




Figure 11 Best layouts found for the *Merged5* instance with three different resolutions: *standard* (left, height of 30.61), *fine* (centre, height of 30.27) and *extra-fine* (right, height of 29.24, best known value for this instance).

algorithms. The usage of voxels and the nofit voxel allow for a quick and robust way of dealing with complex items in packing problems.

With these tools at hand, we formulate the problem as an ILP model that is solved to optimality for small instances, but can also be used with a compaction/separation strategy. We propose a constructive algorithm and an Iterated Local Search that extends the quality of the constructive algorithm. To further improve the packing quality, we propose two metaheuristic techniques that search over the space of solutions containing overlap. We collected a set of instances from the literature and added our own to create a benchmark set of instances with different characteristics where we could test our algorithms. We show that both metaheuristic techniques are very competitive, even when compared to more complex and accurate geometrical representations, such as the polygonal mesh or the phi-functions. The choice of voxels and the nofit voxel to represent and pack the items is backed by our computational experiments, where we show that our results outperform most of the pre-existing literature. The only exception is the algorithm 3DNest, which finds similar quality solutions than Variable neighborhood Search with a fine resolution. This does also show, however, that the ability of voxels to use different representation accuracy depending on the quality/speed of packing required, makes voxel-based approaches more suitable for real-world problems, where the items to be packed might be of higher complexity than the standard

academic instances. Finally, we are able to solve efficiently real-life instances containing complex geometries as typically used in the 3D printing industry, where we show that both the Tabu Search and the Variable Neighborhood Search algorithms can find competitive results, with significant improvement over our constructive-based Iterated Local Search.

References

- Abeysooriya RP, Bennell JA, Martinez-Sykora A (2018) Jostle heuristics for the 2D-irregular shapes bin packing problems with free rotation. *International Journal of Production Economics* 195:12–26, ISSN 09255273, URL <http://dx.doi.org/10.1016/j.ijpe.2017.09.014>.
- Art R (1966) An approach to the two dimensional irregular cutting stock problem. *Tech. Rep. 36.Y08, IBM Cambridge Scientific Centre* URL <http://hdl.handle.net/1721.1/97782>.
- Attene M (2015) Shapes In a Box: Disassembling 3D Objects for Efficient Packing and Fabrication. *Computer Graphics Forum* (8).
- Baert J, Lagae A, Dutré P (2013) Out-of-core construction of sparse voxel octrees. *Proceedings - High-Performance Graphics 2013, HPG 2013*, 27–32, ISBN 9781450321358 (ISBN), ISSN 14678659, URL <http://dx.doi.org/10.1111/cgf.12345>.
- Baumers M, Tuck C, Wildman R, Ashcroft I, Rosamond E, Hague R (2013) Transparency Built-in: Energy Consumption and Cost Estimation for Additive Manufacturing. *Journal of Industrial Ecology* 17(3):418–431, ISSN 15309290, URL <http://dx.doi.org/10.1111/j.1530-9290.2012.00512.x>.
- Bennell J, Oliveira J (2008) The geometry of nesting problems: A tutorial. *European Journal of Operational Research* 184(2):397–415, URL <http://www.sciencedirect.com/science/article/pii/S0377221706012379>.
- Bennell J, Scheithauer G, Stoyan Y, Romanova T (2010) Tools of mathematical modeling of arbitrary object packing problems. *Annals of Operations Research* 179(1):343–368, ISSN 02545330, URL <http://dx.doi.org/10.1007/s10479-008-0456-5>.
- Bennell JA, Dowsland KA (1999) A tabu thresholding implementation for the irregular stock cutting problem. *International Journal of Production Research* 37(18):4259–4275, ISSN 0020-7543, URL <http://dx.doi.org/10.1080/002075499189763>.
- Bennell JA, Dowsland KA (2001) Hybridising Tabu Search with Optimisation Techniques for Irregular Stock Cutting. *Management Science* 47(8):1160–1172, ISSN 0025-1909, URL <http://dx.doi.org/10.1287/mnsc.47.8.1160.10230>.

- Bennell JA, Dowsland KA, Dowsland WB (2000) The irregular cutting-stock problem - a new procedure for deriving the no-fit polygon. *Computers and Operations Research* 28(3):271–287, ISSN 03050548, URL [http://dx.doi.org/10.1016/S0305-0548\(00\)00021-6](http://dx.doi.org/10.1016/S0305-0548(00)00021-6).
- Bennell Ja, Oliveira JF (2009) A tutorial in irregular shape packing problems. *The Journal of the Operational Research Society* 60:S93–S105, ISSN 0160-5682, URL <http://dx.doi.org/10.1057/jors.2008.169>.
- Błazewicz J, Hawryluk P, Walkowiak R (1993) Using a tabu search approach for solving the two-dimensional irregular cutting problem. *Annals of Operations Research* 41(4):313–325, ISSN 02545330, URL <http://dx.doi.org/10.1007/BF02022998>.
- Burke E, Hyde M, Kendall G, Woodward J (2010) A genetic programming hyper-heuristic approach for evolving 2-D strip packing heuristics. *IEEE Transactions on evolutionary computation* 14(6):942–958, URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5491153.
- Byholm T, Toivakka M, Westerholm J (2009) Effective packing of 3-dimensional voxel-based arbitrarily shaped particles. *Powder Technology* 196(2):139–146, ISSN 00325910, URL <http://dx.doi.org/10.1016/j.powtec.2009.07.013>.
- Cagan J, Degentesh D, Yin S (1998) A simulated annealing-based algorithm using hierarchical models for general three-dimensional component layout. *Computer-Aided Design* 30(10):781–790, ISSN 00104485, URL [http://dx.doi.org/10.1016/S0010-4485\(98\)00036-0](http://dx.doi.org/10.1016/S0010-4485(98)00036-0).
- de Korte A, Brouwers H (2013) Random packing of digitized particles. *Powder Technology* 233:319–324, ISSN 00325910, URL <http://dx.doi.org/10.1016/j.powtec.2012.09.015>.
- Dickinson JJK, Knopf GKG (1998) Serial packing of arbitrary 3D objects for optimizing layered manufacturing. *Proc. SPIE*, volume 3522, 130–138, ISSN 0277786X, URL <http://dx.doi.org/10.1117/12.325756>.
- Dowsland KA, Dowsland WB, Bennell JA (1998) Jostling for position: Local improvement for irregular cutting patterns. *Journal of the Operational Research Society* 49(6):647–658, ISSN 14769360, URL <http://dx.doi.org/10.1057/palgrave.jors.2600563>.
- Edelkamp S, Wichern P (2015) Packing Irregular-Shaped Objects for 3DPrinting. Hölldobler S, Krötzsch M, Peñaloza R, Rudolph S, eds., *KI 2015: Advances in Artificial Intelligence: 38th Annual German Confer-*

- ence on AI, volume 9324 of *Lecture Notes in Computer Science* (Cham: Springer International Publishing), ISBN 9783319244884, ISSN 16113349, URL <http://dx.doi.org/10.1007/978-3-319-24489-1>.
- Edwards T (2014) Classic Chess Set from glChess. URL <https://www.thingiverse.com/thing:322616>.
- Egeblad J, Garavelli C, Lisi S, Pisinger D (2010) Heuristics for container loading of furniture. *European Journal of Operational Research* 200(3):881–892, ISSN 0377-2217, URL <http://dx.doi.org/https://doi.org/10.1016/j.ejor.2009.01.048>.
- Egeblad J, Nielsen BK, Brazil M (2009) Translational packing of arbitrary polytopes. *Computational Geometry* 42(4):269–288, ISSN 09257721, URL <http://dx.doi.org/10.1016/j.comgeo.2008.06.003>.
- Egeblad J, Nielsen BK, Odgaard A (2007) Fast neighborhood search for two- and three-dimensional nesting problems. *European Journal of Operational Research* 183(3):1249–1266, ISSN 03772217, URL <http://dx.doi.org/10.1016/j.ejor.2005.11.063>.
- Glover F (1989) Tabu search-part I. *ORSA Journal on computing* URL <http://pubsonline.informs.org/doi/abs/10.1287/ijoc.1.3.190>.
- Glover F, Marti R (2011) Tabu search. Panos M Pardalos DZD, Graham R, eds., *Handbook of Combinatorial Optimization*, chapter Tabu Searc, URL http://link.springer.com/chapter/10.1007/0-387-33416-5_3.
- Gomes A, Oliveira JF (2002) A 2-exchange heuristic for nesting problems. *European Journal of Operational Research* 141(2):359–370, ISSN 03772217, URL [http://dx.doi.org/10.1016/S0377-2217\(02\)00130-3](http://dx.doi.org/10.1016/S0377-2217(02)00130-3).
- Gomes AM, Oliveira JF (2006) Solving Irregular Strip Packing problems by hybridising simulated annealing and linear programming. *European Journal of Operational Research* 171(3):811–829, ISSN 03772217, URL <http://dx.doi.org/10.1016/j.ejor.2004.09.008>.
- Griffiths V, Scanlan JP, Eres MH, Martinez-Sykora A, Chinchapatnam P (2019) Cost-driven build orientation and bin packing of parts in Selective Laser Melting (SLM). *European Journal of Operational Research* 273(1):334–352, ISSN 0377-2217, URL <http://dx.doi.org/10.1016/J.EJOR.2018.07.053>.
- Harrell E (2015) Toyota 4 Cylinder Engine 22RE, Complete working model. URL <http://www.thingiverse.com/thing:644933>.

- Hur SM, Choi KH, Lee SH, Chang PK (2001) Determination of fabricating orientation and packing in SLS process. *Journal of Materials Processing Technology* 112(2-3):236–243, ISSN 09240136, URL [http://dx.doi.org/10.1016/S0924-0136\(01\)00581-7](http://dx.doi.org/10.1016/S0924-0136(01)00581-7).
- Ikonen I, Biles W, Kumar A, Wissel J, Ragade R (1997) A Genetic Algorithm for Packing Three-Dimensional Non-Convex Objects Having Cavities and Holes. *Proceedings of the 7th International Conference on Genetic Algorithms, East Lansing, Michigan, Morgan Kaufmann Publishers* 591–598, URL http://www.academia.edu/download/39497284/A_Genetic_Algorithm_for_Packing_Three-Di20151028-15046-1htxzp4.pdf.
- Imamichi T, Yagiura M, Nagamochi H (2009) An iterated local search algorithm based on nonlinear programming for the irregular strip packing problem. *Discrete Optimization* 6(4):345–361, ISSN 15725286, URL <http://dx.doi.org/10.1016/j.disopt.2009.04.002>.
- Jia X, Williams R (2001) A packing algorithm for particles of arbitrary shapes. *Powder Technology* 120(3):175–186, ISSN 00325910, URL [http://dx.doi.org/10.1016/S0032-5910\(01\)00268-6](http://dx.doi.org/10.1016/S0032-5910(01)00268-6).
- Lemus E, Bribiesca E, Garduño E (2015) Surface trees – Representation of boundary surfaces using a tree descriptor. *Journal of Visual Communication and Image Representation* 31:101–111, ISSN 10473203, URL <http://dx.doi.org/10.1016/j.jvcir.2015.06.004>.
- Liu X, Liu JM, Cao AX, Yao ZL (2015) HAPE3D—a new constructive algorithm for the 3D irregular packing problem. *Frontiers of Information Technology & Electronic Engineering* 16(5):380–390, ISSN 2095-9184, URL <http://dx.doi.org/10.1631/FITEE.1400421>.
- López-Ibáñez M, Dubois-Lacoste J, Pérez Cáceres L, Birattari M, Stützle T (2016) The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* 3:43–58, ISSN 22147160, URL <http://dx.doi.org/10.1016/j.orp.2016.09.002>.
- Lourengo HR, Martin OC, Stützle T (2010) Iterated Local Search: Framework and Applications. 363–397 (Springer, Boston, MA), URL http://dx.doi.org/10.1007/978-1-4419-1665-5_{_}12.
- Ma Y, Chen Z, Hu W, Wang W (2018) Packing Irregular Objects in 3D Space via Hybrid Optimization. *Computer Graphics Forum* 37(5):49–59, ISSN 0167-7055, URL <http://dx.doi.org/10.1111/cgf.13490>.

- Mangiaracina R, Marchet G, Perotti S, Tumino A (2015) A review of the environmental implications of B2C e-commerce: a logistics perspective. *International Journal of Physical Distribution & Logistics Management* (6).
- Mazareanu E (2020) Parcel traffic worldwide by sector 2015-2019 . *Statista*, URL <https://www.statista.com/statistics/737418/parcel-traffic-worldwide-by-sector/>.
- Milenkovic V, Daniels K, Li Z (1992) Placement and Compaction of Nonconvex Polygons for Clothing Manufacture. *Fourth Canadian Conference on Computational Geometry*, 1–8, number August (St. John's, Newfoundland), URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.49.4310&rep=rep1&type=pdf>.
- Min P (2004 - 2019) binvox. URL <http://www.patrickmin.com/binvox> or <https://www.google.com/search?q=binvox>.
- Mladenović N, Hansen P (1997) Variable neighborhood search. *Computers & Operations Research* 24(11):1097–1100, ISSN 03050548, URL [http://dx.doi.org/10.1016/S0305-0548\(97\)00031-2](http://dx.doi.org/10.1016/S0305-0548(97)00031-2).
- Mundim LR, Andretta M, Carravilla MA, Oliveira JF (2018) A general heuristic for two-dimensional nesting problems with limited-size containers. *International Journal of Production Research* 56(1-2):709–732, URL <http://dx.doi.org/10.1080/00207543.2017.1394598>.
- Nooruddin FS, Turk G (2003) Simplification and repair of polygonal models using volumetric techniques. *IEEE Transactions on Visualization and Computer Graphics* 9(2):191–205.
- Oliveira JF, Gomes AM, Ferreira JS (2000) TOPOS – A new constructive algorithm for nesting problems. *OR Spektrum* 22(2):263, ISSN 01716468, URL <http://dx.doi.org/10.1007/s002910050105>.
- Pankratov AV, Romanova TE, Chugay AM, Problems E, Acad N (2015) Optimal packing of convex polytopes using quasi-phi-functions. *Journal of Mechanical Engineering* 18(2):55–64.
- Robidoux N, Stelldinger P, Cupitt J (2011) Simple random generation of smooth connected irregular shapes for cognitive studies. *Proceedings of The Fourth International C* Conference on Computer Science and Software Engineering - C3S2E '11*, 17–24 (New York, New York, USA: ACM Press), ISBN 9781450306263, URL <http://dx.doi.org/10.1145/1992896.1992899>.

- Romanova T, Bennell J, Stoyan Y, Pankratov A (2018) Packing of concave polyhedra with continuous rotations using nonlinear optimisation. *European Journal of Operational Research* 268(1):37–53, ISSN 0377-2217, URL <http://dx.doi.org/10.1016/J.EJOR.2018.01.025>.
- Sánchez-Cruz H, López-Valdez HH, Cuevas FJ (2014) A new relative chain code in 3D. *Pattern Recognition* 47(2):769–788, ISSN 00313203, URL <http://dx.doi.org/10.1016/j.patcog.2013.08.010>.
- Scheithauer G, Stoyan YG, Romanova TY (2005) Mathematical Modeling of Interactions of Primary Geometric 3D Objects. *Cybernetics and Systems Analysis* 41(3):332–342, ISSN 1060-0396, URL <http://dx.doi.org/10.1007/s10559-005-0067-y>.
- Schwarz M, Seidel HP (2010) Fast parallel surface and solid voxelization on GPUs. URL <http://dx.doi.org/10.1145/1882261.1866201>.
- Stoyan YG, Gil N, Pankratov A, Scheithauer G (2004) Packing non-convex polytopes into a parallelepiped. Technical Report June, URL <http://www.math.tu-dresden.de/~scheith/ABSTRACTS/PREPRINTS/04-non-conv.pdf>.
- Stoyan YG, Gil NI, Scheithauer G, Pankratov A, Magdalena I (2005) Packing of convex polytopes into a parallelepiped. *Optimization: A Journal of Mathematical Programming and Operations Research* 54(2):215–235, ISSN 0233-1934, URL <http://dx.doi.org/10.1080/02331930500050681>.
- Stoyan YG, Yaskov GN (1983) Mathematical methods for geometric design. *Proceedings of International Conference PROLAMAT 82*, 67–86, URL https://scholar.google.co.uk/scholar?cluster=12557682819935042194&hl=en&as_sdt=2005&sciodt=0,5#0.
- Terashima-Marín H, Ross P, Farías-Zárate CJ, López-Camacho E, Valenzuela-Rendón M (2010) Generalized hyper-heuristics for solving 2D Regular and Irregular Packing Problems. *Annals of Operations Research* 179(1):369–392, ISSN 0254-5330, URL <http://dx.doi.org/10.1007/s10479-008-0475-2>.
- Thompson MK, Moroni G, Vaneker T, Fadel G, Campbell RI, Gibson I, Bernard A, Schulz J, Graf P, Ahuja B, Martina F (2016) Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints. *CIRP Annals - Manufacturing Technology* 65(2):737–760, ISSN 17260604, URL <http://dx.doi.org/10.1016/j.cirp.2016.05.004>.

- Toledo FM, Carravilla MA, Ribeiro C, Oliveira JF, Gomes AM (2013) The Dotted-Board Model: A new MIP model for nesting irregular shapes. *International Journal of Production Economics* 145(2):478–487, ISSN 09255273, URL <http://dx.doi.org/10.1016/j.ijpe.2013.04.009>.
- Umetani S, Yagiura M, Imahori S, Inamichi T, Nonobe K, Ibaraki T (2009) Solving the irregular strip packing problem via guided local search for overlap minimization. *International Transactions in Operational Research* 16(6):661–683, ISSN 09696016, URL <http://dx.doi.org/10.1111/j.1475-3995.2009.00707.x>.
- Wäscher G, Haubner H, Schumann H (2007) An improved typology of cutting and packing problems. *European Journal of Operational Research* 183:1109–1130, ISSN 03772217, URL <http://dx.doi.org/10.1016/j.ejor.2005.12.047>.