

Learning based 2D Irregular Shape Packing

ZESHI YANG, LightSpeed Studios, USA
ZHERONG PAN, LightSpeed Studios, USA
MANYI LI, Shandong University, China
KUI WU, LightSpeed Studios, USA
XIFENG GAO, LightSpeed Studios, USA




Fig. 1. Three UV-packing results generated using XAtlas and our learning-assisted method.

2D irregular shape packing is a necessary step to arrange UV patches of a 3D model within a texture atlas for memory-efficient appearance rendering in computer graphics. Being a joint, combinatorial decision-making problem involving all patch positions and orientations, this problem has well-known NP-hard complexity. Prior solutions either assume a heuristic packing order or modify the upstream mesh cut and UV mapping to simplify the problem, which either limits the packing ratio or incurs robustness or generality issues. Instead, we introduce a learning-assisted 2D irregular shape packing method that achieves a high packing quality with minimal requirements from the input. Our method iteratively selects and groups subsets of UV patches into near-rectangular super patches, essentially reducing the problem to bin-packing, based on which a joint optimization is employed to further improve the packing ratio. In order to efficiently deal with large problem instances with hundreds of patches, we train deep neural policies to predict nearly rectangular patch subsets and determine their relative poses, leading to linear time scaling with the number of patches. We demonstrate the effectiveness of our method on three datasets for UV packing, where our method achieves a higher packing ratio over several widely used baselines with competitive computational speed.

CCS Concepts: • **Computing methodologies** → **Computer graphics**; **Machine learning**.

Additional Key Words and Phrases: geometry processing, deep reinforcement learning, combinatorial optimization

ACM Reference Format:

Zeshi Yang, Zherong Pan, Manyi Li, Kui Wu, and Xifeng Gao. 2023. Learning based 2D Irregular Shape Packing. *ACM Trans. Graph.* 42, 6 (December 2023), 16 pages. <https://doi.org/10.1145/3618348>

Authors' addresses: Zeshi Yang, zs243@mail.ustc.edu.cn, LightSpeed Studios, Seattle, WA, USA; Zherong Pan, zrpan@global.tencent.com, LightSpeed Studios, Seattle, WA, USA; Manyi Li, manyili@sdu.edu.cn, Shandong University, Jinan, Shandong, China; Kui Wu, kwu@global.tencent.com, LightSpeed Studios, Los Angeles, CA, USA; Xifeng Gao, xifgao@global.tencent.com, LightSpeed Studios, Seattle, WA, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
0730-0301/2023/12-ART \$15.00
<https://doi.org/10.1145/3618348>

1 INTRODUCTION

2D irregular shape packing lies the theoretical foundation for a wide spectrum of applications across various industrial areas, e.g. material parts assembly [Ke et al. 2020], VLSI module placements [Cheng et al. 2005], and 2D shelf arrangements in automatic warehouses [Önüt et al. 2008]. The most prominent packing application in computer graphics is UV chart packing, where a set of 2D shapes corresponding to patches of 3D models are packed together into a single texture atlas for mapping various appearance details onto the 3D surface for downstream rendering. In the realm of digital games where many 3D models are rendered in the virtual world, improving the UV packing ratio could significantly save the graphical memory, improve the loading speed, and reduce rendering overhead. This problem was recognized decades ago, e.g., by Soucy et al. [1996] and Sander et al. [2001], which has recently revived as an active area of graphic research [Limper et al. 2018; Liu et al. 2019; Zhang et al. 2020].

A packing algorithm needs to search for both the position and orientation of each patch. Due to its combinatorial nature, finding optimal 2D shape packing has been well-understood as being NP-complete and APX-hard. Prior research efforts are focused on heuristic or genetic algorithms to determine the order of objects to be packed, and then search for the pose of each shape, e.g., using no-fit polygon (NFP) and the “Tetris” algorithm. However, the sub-optimality gap of these heuristics can be substantial. On the other hand, when the problem size is small, globally optimal stochastic optimization algorithms have been proposed to search for the near-optimal packing orders. Although these techniques can scale to tens of 2D shapes, UV unwrapping approaches and commercial 3D modeling software can oftentimes generate hundreds of patches for one 3D model, which is far beyond the capabilities of the global search algorithms.

We propose a learning-assisted 2D packing pipeline for general irregular-shaped patches. The design of our approach is based on two observations: 1) Prior works generate sub-optimal solutions by assuming objects are packed sequentially. 2) Prior works try all possible poses and orientations for each to-be-packed object to find the best solution, leading to slow computation. To tackle

these issues, we propose to hierarchically group patches into nearly rectangular super-patches, allowing a larger search space for patch combination and a smaller optimality gap. Instead of exhaustively trying all possible combinations, we train a high-level group selector network (HSN) to efficiently predict how likely a candidate patch subset can be grouped into a rectangular super-patch. Given an identified patch subset, we use the sequential ordering technique, similar to prior works. Specifically, we use a low-level sorter network (LSN) to determine the suitable order of packing within the subset. Next, a low-level pose network (LPN) infers the rough initial pose of each patch. The ultimate pose of these patches is determined by local numerical optimization. These two networks are trained using reinforcement learning to efficiently infer near optimal sequential packing policies.

We have evaluated our method on 3 datasets, containing highly irregular UV patches generated from both organic and man-made 3D models. We show the effectiveness of our approach by comparing against widely used algorithm baselines, including piecewise linear NFP, XAtlas [Young 2023], and [Sander et al. 2003], where we achieve a 5% – 10% packing ratio improvement across all the tested datasets. We further highlight the generality of our learned policies, by training them on one dataset and evaluating them on all three datasets. We find our method still outperforming baselines on unseen datasets, proving that our method does not need to be retrained for each problem domain. We include the full datasets and results with statistics in the supplementary material.

2 RELATED WORK

We discuss representative related works on UV-atlas generation, irregular shape packing, and learning-based packing algorithms.

2.1 UV-atlas Generation

The standard industrial pipeline for generating UV-atlas involves three stages: cutting, parameterization, and packing. Mesh parameterization minimizes various distortion energies as summarized in [Rabinovich et al. 2017]. To realize strict distortion bounds, Sorkine et al. [2002] proposed greedy, simultaneous cutting and parameterization that compromises between the distortion and cut length. Similarly, Lévy et al. [2002] proposed greedy mesh cutting to avoid overlapping. More recent approaches [Poranne et al. 2017] formulate the cut and parameterization under a joint optimization framework, leading to more optimal solutions. Under a strict distortion bound, however, these joint optimization techniques can lead to a large number of patches being handled by the packer.

To handle the resultant large packing problems, Sander et al. [2001] used the bounding box approximation for each patch, simplifying the problem to a bin packing. Lévy et al. [2002] proposed to solve the irregular shape packing by always putting new patches on top of existing ones, where the exact horizontal position is chosen to minimize the wasted space between vertical boundaries. Nöll and Strieker [2011] revised the irregular packing technique [Lévy et al. 2002] and considered both horizontal and vertical boundaries, including holes. Although these techniques are less accurate than exact NFP-based methods [Bennell and Song 2008], their computational costs are much lower for large problem instances. On the

downside, all these packing algorithms are myopic and assume a fixed packing order, potentially limiting their solution quality.

Instead of packing UV patches at a separate stage, several works propose to optimize mesh cuts and patch shapes for higher packing ratio. Limper et al. [2018] proposed to progressively cut along void box boundaries to improve packing ratio. Schertler et al. [2018] generated motorcycle graph for quad-dominant meshes consisting of rectangular patches, leading to perfect packing ratios. Liu et al. [2019] generalized the idea of motorcycle graph to arbitrary meshes by deforming arbitrarily shaped patches to be nearly axis-aligned. Although these techniques are appealing for a high packing ratio, they change either the topology or geometry of the input UV patches. We choose to solve packing by maintaining the original shapes of UV patches for better conformity with existing pipelines.

2.2 General Irregular Shape Packing

Due to the theoretical hardness [Bansal et al. 2006; Hartmanis 1982] and practical importance, the irregular packing problem has garnered research attention over the past decades. In the community of manufacturing design [Attene 2015; Cheng et al. 2005; Önüt et al. 2008; Wang et al. 2022], researchers have developed algorithms to compute highly efficient packing solutions via a bilevel pipeline, where the high-level algorithm computes an order of the shapes for packing and the low-level algorithm determines the affine transformation for each shape regarding the already packed shapes. Early research works focus on low-level geometric algorithms that determine all possible collision-free poses for placing a new shape, resulting in the useful tool of NFP [Bennell and Song 2008] and fast collision checking [Riff et al. 2009]. The most efficient method for computing the NFP is through the Minkowski sum [Bennell and Song 2008], which incurs a complexity of at least $O(N^2)$ with N being the number of edges in each polygon. The ultimate pose of each shape is then chosen from the potential pose set using simple heuristic rules as summarized in [Guo et al. 2022], where prominent rules include bottom-left-first, maximal-packing-ratio, and minimal-boundary-length. Unfortunately, even computing the exact NFP for hundreds and thousands of UV patches is intractable, and practical algorithms [Lévy et al. 2002; Nöll and Strieker 2011] only consider horizontal or vertical boundaries.

Built off of the low-level algorithm, a high-level shape selector further optimizes the packing ratio by manipulating the order of packing. This is a critical step to high-quality packing and researchers propose a series of global optimization algorithms, including brute force search [Crainic et al. 2009], genetic algorithm [Ke et al. 2020], and simulated annealing [Gomes and Oliveira 2006]. However, each search step of these algorithms involves calling the low-level algorithms, which are intractable for UV packing, so existing UV packing algorithms such as [Nöll and Strieker 2011] simply sort the patches in area-descending order.

2.3 Learned Packing Policy

Machine learning bears significant potential in solving combinatorial optimization [Bengio et al. 2021]. Several recent works have applied learning approaches to various packing problems, most of which [Hu et al. 2020; Jiang et al. 2021; Yang et al. 2023; Zhao et al.

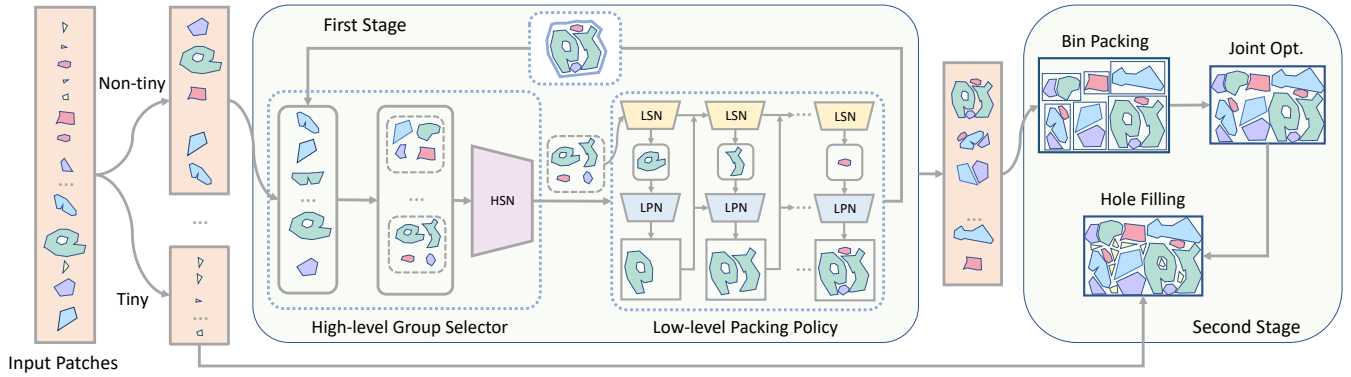


Fig. 2. The overall pipeline of our learning-assisted packing approach.

2021, 2022] deal with 3D bin packing problems. In particular, Zhao et al. [2021] and Yang et al. [2023] only deal with online packing problems where the high-level packing order is fixed. In contrast, Hu et al. [2020] focuses on learning the physically realizable packing order. They trained an attention network to locate the next object to be packed, given their precedence constraints. Similar to our technique, Jiang et al. [2021] and Zhao et al. [2022] presented a complete pipeline involving learned high-level object selectors and low-level pose inference.

Learning irregular packing skills is a much more challenging problem, where the learning model has to recognize the complex object shapes and handle the continuous decision space of object poses. Two recent works [Fang et al. 2023; Goyal and Deng 2020] considered packing of irregular shapes. Goyal and Deng [2020] proposed a problem set for 3D packing problems and showed that a neural shape selector trained via reinforcement learning outperforms heuristic baselines. Fang et al. [2023] also trained a neural shape selector via reinforcement learning. Both algorithms use heuristic rules for the low-level pose computation and Fang et al. [2023] even used NFP-based pose computation. However, only learning the high-level selector is not enough for UV packing, since low-level pose selection can still be a major computational bottleneck.

3 OVERVIEW

Our method takes as input a set of UV patches, each represented as an irregular, planar, manifold triangular mesh in their local frame of reference. We allow an arbitrary number of holes within each UV patch. The output of our algorithm is a set of rigid transformations, one for each patch, such that, after being transformed, the UV patches are tightly packed into a rectangular texture domain in the global frame of reference and in a collision-free manner. The pipeline of our method is illustrated in Fig. 2.

Given the set of UV patches, we filter out the tiny patches and firstly process the non-tiny ones in two stages. Our first stage aims to turn the non-tiny patches into nearly rectangular super-patches (Sec. 4). We maintain the set \mathcal{S} of super-patches that is initialized to be the non-tiny patches. We first employ a sampling-based group selector network (HSN Sec. 4.3) to identify a potential subset \mathcal{S}' of at most H patches that can form a nearly rectangular super-patch.

To compute the super-patch, we utilize a low-level sorter network (LSN Sec. 4.2) and a low-level pose network (LPN Sec. 4.1). LSN re-orders the patches within the subset and sequentially outputs the next patch to be packed. LPN selects the rough initial pose for each incoming patch, whose ultimate pose is adjusted using a local numerical optimization. This procedure is denoted as the low-level function $LL(\mathcal{S}')$. Once the super-patch is generated, it will be inserted back into \mathcal{S} to replace the subset of patches, essentially updating \mathcal{S} to $\mathcal{S} - \mathcal{S}' \cup \{LL(\mathcal{S}')\}$. The above process is repeated until each (super-)patch in \mathcal{S} is nearly rectangular. Our second stage (Sec. 4.4) assembles all the patches using a heuristic bin-packing algorithm. A joint local optimization is then performed to squeeze all the patches as tightly as possible. Finally, the tiny patches filtered out in the beginning will be put into gaps between non-tiny patches. In the following, we describe the two stages in detail. We summarize our network architectures and pseudo-code of our algorithm in our supplementary material.

4 NEAR RECTANGULAR PATCH GENERATION

Although we use high-level selector networks before low-level policies, we need to train low-level policies first due to data dependency and we introduce our method in the order of training.

4.1 Low-Level Pose Network (LPN)

Given a subset of H patches that are selected by HSN, let us assume the first $i - 1$ patches have been packed into a super-patch P_{i-1} in the global frame, and p_i is the geometric domain of the i -th patch in the local frame. Given p_i and P_{i-1} , the low-level packing algorithm needs to select the translation t_i and rotation θ_i for p_i such that the packed shape $P_i \triangleq [R(\theta_i)p_i + t_i] \cup P_{i-1}$ is collision-free with a high packing ratio. Conventional packing algorithms [Guo et al. 2022] would consider each patch independently and evenly sample K rotations and consider all the possible translations under each rotation using NFP algorithm, leading to at least $\mathcal{O}(KN^2)$ complexity with N being the total number of edges in p_i and P_{i-1} , which is a major bottleneck of packing algorithms. And due to its myopic nature, the packing ratio is sub-optimal.

To address the above two shortcomings, we propose to model the packing procedure as a Markov Decision Process (MDP) and train

LPN to maximize the packing ratio via reinforcement learning. Being aware of not only the current patch but also the future incoming patches, our LPN policy exhibits a small optimality gap. Briefly, the MDP is identified with a tuple $\langle S, A, \tau, r \rangle$, which models the procedure of a decision-maker iteratively observing the current system state in the state space S and taking an action in the action space A to change the environment. The environment would then update its state via a state transition function τ and the decision-maker receives a reward r . We refer readers to [Sutton and Barto 2018] for more details on this model. For the packing problem, however, the action space could involve all possible patch translations and rotations, which is notoriously difficult to handle for reinforcement learning. Instead, we propose to restrict the action space to a small discrete subset, and then use local optimization to fine-tune the final pose of each patch.

4.1.1 State Space S . During the i -th iteration, the LPN observes the current system state s_i in the state space S . In our problem, we assume LPN can observe the current packing patch P_{i-1} and a set of at most H future patches to be packed, i.e., $s_i \triangleq (P_{i-1}, p_i, \dots, p_{i+H-1})$. This is because future patches can affect the pose of the current patch in order to achieve joint optimality. Unlike the myopic algorithms that consider a single future patch p_i , we feed the entire ordered sequence of H future patches to the network. Empirically, we find this strategy can effectively guide the network to avoid myopic local minima.

Practically, each patch can be of arbitrary geometric shapes, so we rasterize each patch in s_i (including P_{i-1}) to a 50×50 2D image. For each patch, we move their center-of-mass (COM) to the image center and encode each patch using a shared Fully Convolutional Network (FCN) into a 432-dimensional latent code $\bar{p}_i \triangleq \text{FCN}(p_i)$, which is concatenated as $\bar{s}_i = \text{FCN}(s_i)$ for short. Since patches are of drastically different sizes, we scale all the patches before using the FCN, such that the area of all the H patches is equal to 60% of the area of the 2D image. Note that such global scaling will not change the packing ratio, so it should not change the optimal packing policy. When there are not enough patches to fill up the H channels of patches, we feed the FCN with blank images.

4.1.2 Action Space A . Having observed s_i , our LPN can be represented as a policy function $a_i = \pi^{\text{LPN}}(s_i)$ that maps the state to an action a_i in the action space A . Unlike prior works for learning-based regular shape packing [Zhao et al. 2021] or shape ordering [Fang et al. 2023], the design of action space A for irregular packing is much more challenging. On the one hand, a valid action space should only consist of collision-free patch poses, while identifying these actions can involve extensive collision checks. On the other hand, training




Fig. 3. An illustration of our action space. Given the already packed patch P_{i-1} and the current patch p_i , our action needs to determine the rough rotation θ_i of p_i from local (dashed) to global frame, as well as the relative rotation ϕ_i with respect to P_{i-1} .

a decision-maker in high-dimensional action spaces can be rather data-demanding, and a promising subset of actions should be pre-selected. We tackle these two problems by having the policy π^{LPN} to select a rough initial guess and then use local optimization to revise the pose. Specifically, we re-parameterize the action space A under polar coordinates (Fig. 3). We first compute the COMs for P_{i-1} and p_i , denoted as $\text{COM}(P_{i-1})$ and $\text{COM}(p_i)$. The relative position $\text{COM}(p_i)$ with respect to $\text{COM}(P_{i-1})$ is expressed under polar coordinates with relative angle ϕ_i and we ignore their relative distance. Similarly, the local-to-global rotation of p_i is encoded as another angle θ_i . In summary, we define our action space as: $a_i \triangleq (\theta_i, \phi_i)$. At an early stage of this research, we parameterize our policy to use this continuous action space. However, in experiments we find the optimal state-action distribution can be multi-modal, and Gaussian distributions widely adopted to model the continuous action space cannot capture the multi-modal nature, leading to inferior performances. Therefore, we sample the range of θ_i and ϕ_i at 16 angles and consider the $16 \times 16 = 256$ -dimensional discrete action space.

4.1.3 State Transition. Our state transition function $s_{i+1} = \tau(s_i, a_i)$ computes the next state s_{i+1} from s_i and a_i by converting the action θ_i and ϕ_i into a collision-free, tightly packed pose. Since we use a coarse discretization of the action space, we also use the state transition function to locally revise the action and improve the packing ratio. To this end, we devise the following collision-constrained local optimization:

$$\begin{aligned} \theta_i^*, t_i^* &\triangleq \underset{\theta, t}{\operatorname{argmin}} \|R(\theta)\text{COM}(p_i) + t - \text{COM}(P_{i-1})\|^2 \\ \text{s.t.} \quad &[R(\theta)p_i + t] \cap P_{i-1} = \emptyset, \end{aligned} \quad (1)$$

where we initialize $\theta = \theta_i$ and t_i is initialized from ϕ_i by elongating the relative direction (red line in Fig. 3) between P_i and p_i until they are collision free. Finally, we update the next state as $s_{i+1} \triangleq ([R(\theta_i^*)p_i + t_i^*] \cup P_{i-1}, p_{i+1}, \dots, p_{i+H})$. Note that we use the distance between center-of-mass as a surrogate measure for the packing ratio. We choose not to use the packing ratio as our objective function, because the new patch p_i can oftentimes be entirely contained in the bounding box of P_{i-1} and all the poses of p_i inside the bounding box has the same packing ratio. The collision constraint can be realized in several ways, including the scaffold method [Jiang et al. 2017] and the boundary barrier energy [Smith and Schaefer 2015], and we adopt the latter approach to avoid the costly 2D re-meshing. Although the scaffold method [Jiang et al. 2017] has better solutions in large-scale problems, the barrier energy technique performs better under our small problem sizes with only 3 decision variables. We solve the optimization using Newton's method with line-search to guarantee constraint satisfaction. During the line-search, we implement a bounding volume hierarchy to accelerate the collision check and assembly of barrier energy terms.

4.1.4 LPN Training. We parameterize π^{LPN} as an MLP mapping \bar{s}_i to the Q-value of all 256 actions. After each state transition, the policy receives a sparse reward signal defined as:

$$r(s_i, a_i, s_{i+1}) = \text{pr}(P_i)\mathbb{I}[i = H],$$

where $\text{pr}(P)$ is the packing ratio of super-patch P and $\mathbb{I}[i = H]$ is an indicator function of the last iteration. Note that using sparse

reward signals can significantly slow down policy learning, but such reward does not pose a major problem in our application as we use a short horizon H , i.e. $|H| < 5$. We train our LPN policy via Q-learning algorithm by maximizing the expected cumulative reward:

$$\operatorname{argmax}_{\pi^{\text{LPN}}} \mathbb{E}_{a_i \sim \pi^{\text{LPN}}} \left[\sum_{i=1}^H r(s_i, a_i, s_{i+1}) \right].$$

To solve the stochastic optimization, we adopt the robust double deep Q-learning (DDQN) algorithm [Van Hasselt et al. 2016] and train π^{LPN} to pack randomly sampled batches of at most H patches of arbitrary order from our patch dataset, and we ensure each sampled batch comes from the same 3D model.

4.2 Low-Level Sorter Network (LSN)

Our LSN provides the optimal patch ordering for the LPN to achieve the best packing ratio, as our LPN can only pack the patches in a given order. We model the patch sorting procedure as another MDP denoted as $\langle S, A', \tau', r \rangle$ with the same state space and reward signal as that of LPN. We represent LSN as another policy function $a'_i = \pi^{\text{LSN}}(s_i)$ that selects the next patch to be packed, i.e., a'_i consists of the Q-value of the k future patches. Given the selected next patch, the state transition function $\tau'(s_i, a'_i)$ would invoke LSN and the state transition function τ from Sec. 4.1.3 to yield s_{i+1} .

Neural networks need to understand the relative relationship between the future patches in order to accomplish the sorting task. Therefore, we apply a Graph Attention Network (GAT) module [Velickovic et al. 2017], which is known to be effective in solving sorting tasks [Hu et al. 2020; Zhao et al. 2022]. We organize all the patches into a fully connected graph, where the nodal input of GAT is the patch's feature \bar{p}_i , along with the feature of the already packed super-patch \bar{P}_{i-1} . GAT outputs the high-level graph feature for each of the existing patches. Then the patch features are converted to their corresponding Q-values using an MLP network. This architecture parameterizes our sorting policy, denoted as π^{LSN} . Similar to π^{LPN} , π^{LSN} is trained using DDQN via randomly sampled batches of at most H patches coming from the same 3D model. The LSN and LPN combined define our low-level function $\text{LL}(S')$.

Algorithm 1 Iterative Selection of Super-Patches

- 1: $S \leftarrow$ non-tiny set
 - 2: Sample 400 subsets $S'_{1, \dots, 400}$
 - 3: Sort $S'_{1, \dots, 400}$ in $\text{HSN}(S'_i)$ -descending order
 - 4: Sort $S'_{1, \dots, 10}$ in $\text{pr}(S - S'_i \cup \text{LL}(S'_i))$ -descending order
 - 5: **if** $\text{pr}(S - S'_1 \cup \text{LL}(S'_1)) > \text{pr}(S)$ **then**
 - 6: $S \leftarrow S - S'_1 \cup \text{LL}(S'_1)$, goto Line 2
 - 7: **else** Return S
-

4.3 High-Level Group Selector Network (HSN)

Our low-level policies can only sort and pack a small subset S' of H patches. In order to solve practical UV packing problems with hundreds of patches, we need to iteratively select S' from S . To this end, we design a weighted average packing ratio $\text{pr}(\bullet)$ to evaluate the quality of the updated configuration $\text{pr}(S - S' \cup \text{LL}(S'))$ and

pick S' corresponding to the highest ratio. We then train our HSN to rank the quality of S' without actually calling the costly low-level function $\text{LL}(S')$. Finally, we propose a sampling strategy to further reduce the calls to HSN.

4.3.1 Weighted Average Packing Ratio. In order to compare different choices of S' , we need a metric that measures their similarity to rectangles. To this end, we define the area-averaged packing ratio over all the super-patches in S as follows:

$$\text{pr}(S) \triangleq \frac{\sum_{p \in S} \text{area}(p) \text{pr}(p)}{\sum_{p \in S} |\text{bound}(p)|} \quad \text{pr}(p) \triangleq \frac{\text{area}(p)}{|\text{bound}(p)|},$$

where $\text{area}(p)$ and $\text{bound}(p)$ are the actual area and the bounding box of the super-patch p , respectively. We further observe that the high geometric complexity of the super-patches is due to the interior gaps between them. However, these interior gaps are useless for low-level packing policies because, by the design of our low-level action space of LPN as shown in Fig. 3, we always pack a new patch p_i from the outside of P_{i-1} . Therefore, before inserting $\text{LL}(S')$ to S , we compute the alpha shape [Akkiraju et al. 1995] (Fig. 4) for each new super-patch to fill up the interior gaps and inform the neural networks that interior gaps are useless by design.

4.3.2 Packing Ratio Prediction.

Exhaustively evaluating $\text{pr}(\bullet)$ for all $C_{|S|}^H$ choices of S' would require an intensive amount of calls to the low-level packing policies as well as the costly optimizations (Eq. 1). Instead, we propose a learning-based technique to predict the packing ratio using HSN. Given S' , our HSN uses the same FCN from the low-level policies to encode each patch. The latent codes are then brought through GAT to yield the high-level graph features as in LSN. All the graph features are then brought through a max-pooling layer and a Sigmoid layer to yield the predicted packing ratio. Note the absolute values of packing ratio is less important, since we are only interested in the relative ordering of the potential super-patches. In view of this, we train our HSN via supervised metric learning [Chopra et al. 2005]. During each learning iteration, we randomly sample two (at most) H -patch groups denoted as S' and S'' with groundtruth packing ratios denoted as pr' and pr'' and HSN predicted packing ratio denoted as $\text{HSN}(S')$ and $\text{HSN}(S'')$. We then update HSN via the following margin ranking loss:

$$\mathcal{L} = \max(0, -\text{sgn}(\text{pr}' - \text{pr}'') (\text{HSN}(S') - \text{HSN}(S'')) + \epsilon), \quad (2)$$

where ϵ is the minimal positive margin. We have found that HSN can empirically reach a high prediction accuracy for simple patches, but its accuracy gradually deteriorates as more and more patches are packed into complex-shaped super-patches.

4.3.3 Patch Group Sampling. Even using the HSN to efficiently rank the packing ratio, batch evaluation of $\text{pr}(\bullet)$ for all S' is still time-consuming. To further alleviate the runtime cost, we randomly




Fig. 4. After forming a super-patch (left), we will use alpha shape (right) to fill the gaps and inform the neural networks that these gaps are useless.

sample 400 subsets of patches and predict their packing ratio via a batched HSN evaluation. The top 10 out of 400 best groups are then forwarded to the low-level algorithm to evaluate the groundtruth packing ratio $\text{pr}(\mathcal{S} - \mathcal{S}' \cup \text{LL}(\mathcal{S}'))$, and finally the best of the 10 groups is adopted to form the next super-patch. As outlined in Alg. 1 of the supplementary, this procedure is repeated until the updated packing ratio $\text{pr}(\mathcal{S})$ is not higher than the current value.

4.4 Super-Patch Assembly

After our first stage, the non-tiny patches are grouped into nearly rectangular super-patches. During our second stage, we confidently take the rectangular shape assumption and use the divide-and-conquer algorithm implemented in the Trimesh library [Dawson-Haggerty et al. 2019] to assemble all the super-patches together.

With the initially packed result from bin-packing, we then propose to adjust the joint poses of all the patches, locally squeezing them together via numerical optimization. We denote the bounding box enclosing all the M patches as $\text{bound}(p_1, \dots, p_M)$, and our optimization problem is formulated as:

$$\begin{aligned} \underset{\theta_1, \dots, M, t_1, \dots, M}{\text{argmin}} \quad & |\text{bound}(p_1, \dots, p_M)| \\ \text{s.t.} \quad & [R(\theta_i)p_i + t_i] \cap [R(\theta_j)p_j + t_j] = \emptyset \quad \forall 1 \leq i < j \leq M. \end{aligned} \quad (3)$$

We again use the barrier function technique [Smith and Schaefer 2015] to handle all the collision-free guarantees and ensure that the bounding box is surrounding all the patches. Although this is a joint optimization, it is still efficient to solve, since we only allow rigid motions for all the patches.

Finally, there is another set of tiny patches that are set aside by our filter at the beginning of our pipeline. For these small patches, we adopt a conventional approach to sort them in an area-descending order and then use the scanline algorithm [Hu et al. 2018] to fit them into the gaps and holes of the super-patches. During this stage, we also replace alpha shapes with the original patches to expose the scanline algorithm to potentially useful gaps and holes. These final adjustments are illustrated in Fig. 6.

5 EXPERIMENTS

Network Training. We perform all experiments on a computer with an Intel E5-1650 12-core CPU at 3.60GHz and 32GB RAM. We implement all learning algorithms via Pytorch [Paszke et al. 2017] with the GAT implemented based on [Wang et al. 2019]. For training the LPN, we use DDQN with an experience buffer size of 10^6 transition tuples. We sample roughly 5×10^6 random packing problems with $H = 4$ to populate the experience buffer and we update π^{LPN} using 2×10^4 epochs of stochastic gradient descend (SGD). The same procedure is used for training π^{LSN} . Both learning rates of the LPN and LSN are set to 10^{-4} . We train HSN using a collected dataset of 6×10^4 H -patch subsets with pre-computed groundtruth packing ratios. We update HSN using 500 epochs of SGD with a learning rate of 10^{-3} and a batch size of 256. For each dataset, we use 70% of data for training and the rest for testing.

Runtime Setting. Given the set of input patches, we first sort them in the area-descending order, and then consider the subset of largest patches, whose sum of areas takes up 80% of the area of all patches, which is denoted as the salient subset \mathcal{S}_s . The average area of patches

in the salient subset is denoted as $\bar{a} = \sum_{p \in \mathcal{S}_s} \text{area}(p) / |\mathcal{S}_s|$. Next, we define the tiny patch set as all the patches with an area smaller than $\bar{a}/5$, with other patches classified as non-tiny. Note that the patches are splitted based on their relative sizes instead of absolute sizes, so the final results are insensitive to the threshold values within a reasonable range. The numerical optimizations in Eq. 1 and Eq. 3 are implemented in C++, where we set the maximal allowed iterations and initial step size to 10^3 and 10^{-4} , respectively. In the super-patch assembly, the aspect ratio between the width and height of the texture image domain ranges from 1 to 2. To choose the appropriate aspect ratio, we run the bin-packing algorithms 10 times using different aspect ratios and choose the one with the highest packing ratio. Practical applications would create gutter space around the boundaries of patches to eliminate interpolation artifacts. In our implementations we set the minimal distance between patches to one pixel.

Datasets. We evaluate our method on three datasets of 2D UV patches obtained from UV unwrapping of 3D models using XAtlas. This software can sometimes generate degenerate patches with zero or negative areas, which are removed from our dataset. As illustrated in Fig. 8, Fig. 7, and Fig. 10, our first dataset, “Building” contains 86 man-made building models with mostly sharp features, where the resulting packing problems have 5 to 131 patches. Our second dataset, “Object”, contains 81 3D models representing daily objects with few sharp features, where the packing problems have 9 to 200 patches. Our third dataset, “General”, contains 221 3D models from Thingi10k [Zhou and Jacobson 2016], where the packing problems have 4 to 200 patches. Note that although the size of our datasets is not comparable to the sizes of datasets commonly used for computer vision tasks, our learned packing policy demonstrates pretty good generalization ability, as shown in the results section. This is because the training is conducted on patch-level instead of model level. We train our models on randomly sampled patches from these 3D models, leading to millions of combinatorial packing instances for training.

Baselines. We identify three baselines to compare against. Our first baseline is the exact NFP-based packing approach combining two heuristic methods: maximal packing ratio [Burke et al. 2006] and lowest gravity center [Liu and He 2006], which is denoted as NFP-Heuristic. Given a list of patches, NFP-Heuristic first sorts all patches in the area-descending order and then sequentially packs each patch. For a new patch, NFP-Heuristic considers its 16 rotations and computes NFP for each rotation using Minkowski sum [Bennell and Song 2008] to find collision-free translations. Finally, the pose leading to the highest packing ratio is selected and, if two poses lead to the same packing ratio, we select the one with the lower gravity center position. We compute the NFP using a highly optimized reduced convolution algorithm implemented in CGAL [Fabri and Pion 2009]. Our second baseline is the packer algorithm implemented in the open-source software: XAtlas [Young 2023], which implements aggressive simplification and acceleration techniques, allowing the packing algorithm to scale to problems with hundreds or thousands of patches. For example, it uses voxelized patches instead of piecewise linear ones, so that they could use a scanline algorithm instead of exact NFP computation. Our third baseline re-implements [Sander

et al. 2003] using Python. The major difference between XAtlas and [Sander et al. 2003] lies in their heuristics, where XAtlas maximizes packing ratio and Sander et al. [2003] minimizes wasted area.

Results. For each dataset, we first profile the packing ratio of all the algorithms on the testing problems. As summarized in Tab. 1 (Ours), our algorithm consistently outperforms the other baselines by 5% – 10%. To further justify the generality of our method, we train our networks on the General dataset (Ours[†]) but test it on the other two datasets. In this case, our method suffers from a marginal loss in packing ratio, but still outperforms all the baselines. This result justifies that our method has a reasonable ability of domain transfer and can be ported to pack patches for different classes of 3D models in a zero-shot fashion. More results of our method are visualized in App. C.

Ablation Study. We further analyze aspects of our learning-assisted technique. First, we profile the accuracy of HSN, which is measured by the fraction of patch pairs that are correctly ranked. Our HSN achieves an accuracy of 90.8%, 86.9%, and 84.6% on the Building, Object, and General test sets, respectively. Our trained HSN achieves a high ranking accuracy for the Building dataset and the accuracy degrades for the Object and General datasets, in which the patch shapes are more complex than those in the Building dataset. Next, we highlight the packing ratio of our low-level π^{LPN} and π^{LSN} alone. To this end, we sample a random subset of H patches and use our LL to perform packing. We compare LL with the baselines and summarize the results averaged over 2500 random problems in Tab. 2. It is shown that our DRL-based packing policy still outperforms other baselines for smaller packing problems with H patches, which validates the necessity of using a learned packing policy as the low-level packer. Next, we compare the packing ratio of LL under different horizons H . We train four low-level algorithms with $H = 2, \dots, 5$ and our resulting packing ratios over the 2500 random problems are show in Tab. 3. Our approach performs the worst when $H = 2$, where our low-level policies become myopic, but the quality varies only slightly when $H \geq 2$. Therefore, we choose $H = 4$ for the highest quality.

Finally, to validate the design choices of our pipeline, we analyze several variants of our method. The corresponding results are summarized in Tab. 4. In LPN+LSN, we remove our hierarchical grouping procedure, and train LPN and LSN with a larger horizon $H = 10$ to directly pack a sequence of patches sorted in area-descending order. In NFP+HSN, we replace LL(●) with NFP and area-descending ordering, but still use our HSN for hierarchical grouping. In LPN(c)+LSN+HSN, we train the LPN using continuous action space, where we adopt PPO algorithm [Schulman et al. 2017] to optimize the policy. The learning process of LSN and HSN remains the same as our method. In LPN+LSN+HSN, we use our low- and high-level algorithms but without hole-filling, i.e., we consider all the patches in the high-level algorithm and do not filter out tiny patches. As shown in Tab. 4, LSN+LPN performs extremely poorly. We find that the policy learned with large horizons fails to converge due to the huge search space and extremely high complexity involved in irregular shape packing problems, which validates the necessity of our hierarchical grouping procedure. We

can see that LPN(c)+LSN+HSN leads to better performances than NFP+HSN, verifying the effectiveness of our learning-based packing policy. By the packing ratio comparisons among LPN(c)+LSN+HSN, LPN+LSN+HSN, and ours, we justify the necessity of a discrete action space and the hole-filling strategy.

Computational Cost. Although our method achieves better packing ratios, our computational efficacy is inferior to XAtlas, due to the repeated network evaluation. For the general dataset, the average packing time of XAtlas [Young 2023], [Sander et al. 2003], NFP, and our method are 1.81s, 33.52s, 93.62s, and 37.76s, respectively. The performance breakdown for our packing method on the general dataset is summarized in Tab. 5. It is shown that the computational bottleneck lies in the scanline-based hole filling, which involves nested for loops and is implemented in Python. Our method can be accelerated if we switch to a scanline algorithm implemented in native-C++. We further study the scalability of our method in dealing with large UV packing problems. To this end, we combine patches from several 3D models and use each algorithm to pack all the patches into a single texture. We create a dataset including packing instances with 50, 100, 150, 200, 250, 300 patches, on which we run our method and other baselines. The computational overhead is plotted against the number of patches in Fig. 5. The cost of NFP is much higher than other algorithms due to the superlinear increase of computational complexity in computing the Minkowski sum. By design of our high-level selection policy, our method scales linearly against the number of patches, although our method is slower than XAtlas. In Fig. 9, we show an example in which 784 charts segmented from 6 animal chesses are packed into a single Atlas. The packing time of [Sander et al. 2003], XAtlas, NFP and our method are 180.68s, 4.68s, 2966.63s and 278.87s, respectively. Our method achieves better packing ratio than NFP while requiring significantly fewer computational resources.

User-controlled Aspect Ratio. By default, we search for the optimal aspect ratio of the texture image that maximizes the packing ratio. However, our method can be easily adapted to support a user-specified packing ratio, by forwarding it to the bin-packing procedure. To compare our method with XAtlas, we conduct experiments where we specify the aspect ratio of 1. The results of these experiments are also summarized in Tab. 1 (Ours[★]), where our approach still generates the best results compared to baselines, although with an expected degradation compared to our results with the default setting. Fig. 11 visualizes some of the packing results with fixed and adjustable aspect ratios.

6 CONCLUSIONS AND FUTURE WORK

We present a learning-assisted irregular shape packing algorithm for UV patches. On three datasets with various topology and geometry properties, we achieve 5% – 10% packing ratio improvement over XAtlas, NFP, and [Sander et al. 2003] baseline algorithms. Our algorithm can deal with problem instances with up to hundreds of patches within tolerable computational overhead for offline packing. By optimizing only the rigid transformations for the patches,

Table 1. (Min|Max|Avg) packing ratios of all algorithms on the test datasets; Ours* means our algorithm run with a fixed aspect ratio of 1; † means our algorithm trained on the general dataset.

Test-set	[Sander et al. 2003]	XAtlas	NFP	Ours	Ours*	Ours†
Building	0.470 0.830 0.675	0.525 0.835 0.670	0.499 0.907 0.707	0.683 0.980 0.827	0.683 0.980 0.801	0.683 0.980 0.805
Object	0.290 0.733 0.609	0.385 0.788 0.588	0.439 0.805 0.630	0.377 0.862 0.687	0.377 0.843 0.680	0.377 0.862 0.682
General	0.455 0.883 0.652	0.449 0.886 0.688	0.405 0.886 0.690	0.540 0.937 0.776	0.509 0.937 0.757	-

Table 2. Average packing ratio over 2500 random problems with H patches.

	[Sander et al. 2003]	XAtlas	NFP	Ours (LL)
pr	0.618	0.651	0.582	0.686

Table 3. Packing ratio comparison with different H .

	H = 2	H = 3	H = 4	H = 5
pr	0.741	0.771	0.776	0.770

Table 4. Packing ratio comparison of algorithm variants on General dataset.

Methods	pr
LSN + LPN	0.448
NFP	0.690
NFP + HSN	0.700
LPN(c) + LSN + HSN	0.725
LPN + LSN + HSN	0.744
Ours	0.776

Table 5. The performance breakdown for packing one patch on average.

Procedure	Cost(%)
HSN Evaluation (Sec. 4.3)	10.1%
LPN+LSN Evaluation (Sec. 4.1,4.2)	18.4%
Bin-packing (Sec. 4.4)	3.9%
Joint optimization (Eq. 3)	12.9%
Scanline-based hole filling	52.9%

our approach respects the input UV patch shapes and parameterizations, which can be immediately incorporated into existing UV unwrapping pipelines.

Our method leads to several noteworthy future directions. First, although our method achieves averaged best performance, there can be sub-optimal solution instances as illustrated in Fig. 12. Second, our action space does not inherently allow patches to be put into gaps between other patches, which requires a more flexible action space design. Next, our existing hierarchical grouping approach, while efficient, operates in a greedy manner and may be confined to local optima. An alternative solution is to reframe the grouping process as a sequential decision-making problem that can be solved by advanced policy optimization algorithms. Finally, we are also looking into more efficient policy parameterizing utilizing diffusion models and transformers.

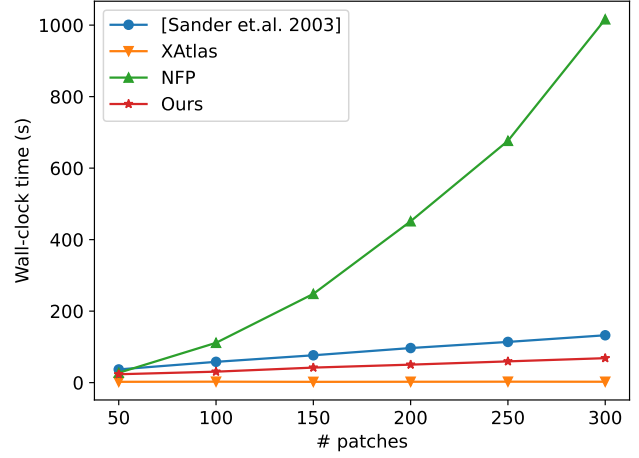


Fig. 5. The average packing time of various algorithms plotted against the number of patches.




Fig. 6. After the first stage, there are gaps between non-tiny patches (top left). We perform a joint optimization to squeeze patches and remove gaps (top right). After hole-filling, we achieve a packing ratio of 76.5% (bottom right). If we do not filter out the tiny patches and forward all patches to HSN, we can only achieve a ratio of 71.9% (bottom left).

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their constructive suggestions and feedback. We also thank Yujie Wang for

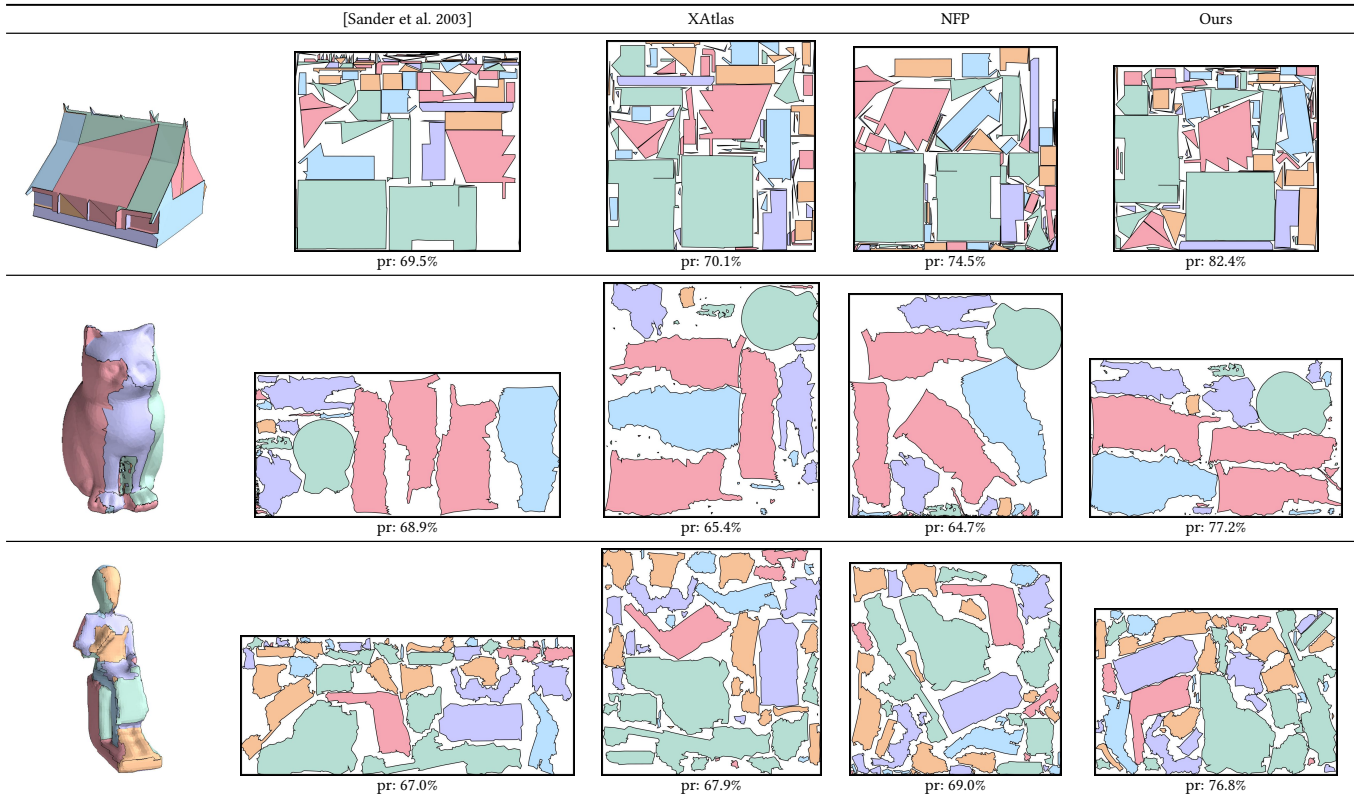


Fig. 7. Three packing instances from the building, object and general dataset.




Fig. 8. A set of 3D model samples from our three datasets.

her valuable discussions and help. Her expertise in solving NP-hard problems using deep neural networks has been instrumental in shaping the direction of this paper, allowing us to refine our methods and focus our efforts in the most promising directions. Additionally, Manyi Li is supported by the Excellent Young Scientists Fund Program (Overseas) of Shandong Province (Grant No.2023HWYQ-034).

REFERENCES

- Nataraj Akkiraju, Herbert Edelsbrunner, Michael Facello, Ping Fu, EP Mucke, and Carlos Varela. 1995. Alpha shapes: definition and software. In *Proceedings of the 1st international computational geometry software workshop*, Vol. 63.
- Marco Attene. 2015. Shapes in a box: Disassembling 3D objects for efficient packing and fabrication. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 64–76.
- Nikhil Bansal, Jos R. Correa, Claire Kenyon, and Maxim Sviridenko. 2006. Bin Packing in Multiple Dimensions: Inapproximability Results and Approximation Schemes. *Math. Oper. Res.* 31, 1 (feb 2006), 31–49. <https://doi.org/10.1287/moor.1050.0168>
- Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. 2021. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research* 290, 2 (2021), 405–421.
- Julia A Bennell and Xiang Song. 2008. A comprehensive and robust procedure for obtaining the nofit polygon using Minkowski sums. *Computers & Operations Research* 35, 1 (2008), 267–281.
- Edmund Burke, Robert Hellier, Graham Kendall, and Glenn Whitwell. 2006. A new bottom-left-fill heuristic algorithm for the two-dimensional irregular packing problem. *Operations Research* 54, 3 (2006), 587–601.




Fig. 9. A large problem instance with 784 charts.




Fig. 10. Three packing instances from the building, object, and general dataset.

Lei Cheng, Liang Deng, and Martin DF Wong. 2005. Floorplanning for 3-D VLSI design. In *Proceedings of the 2005 Asia and South Pacific Design Automation Conference*. 405–411.

Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Vol. 1. IEEE, 539–546.

Teodor Gabriel Crainic, Guido Perboli, and Roberto Tadei. 2009. TS2PACK: A two-level tabu search for the three-dimensional bin packing problem. *European Journal of Operational Research* 195, 3 (2009), 744–760.

Dawson-Haggerty et al. 2019. *trimsh*. <https://trimsh.org/>

Andreas Fabri and Sylvain Pion. 2009. CGAL: The computational geometry algorithms library. In *Proceedings of the 17th ACM SIGSPATIAL international conference on*

advances in geographic information systems. 538–539.

Jie Fang, Yunqing Rao, Xusheng Zhao, and Bing Du. 2023. A Hybrid Reinforcement Learning Algorithm for 2D Irregular Packing Problems. *Mathematics* 11, 2 (2023), 327.

A Miguel Gomes and José F Oliveira. 2006. Solving irregular strip packing problems by hybridising simulated annealing and linear programming. *European Journal of Operational Research* 171, 3 (2006), 811–829.

Ankit Goyal and Jia Deng. 2020. Packit: A virtual environment for geometric planning. In *International Conference on Machine Learning*. PMLR, 3700–3710.

Baosu Guo, Yu Zhang, Jingwen Hu, Jinrui Li, Fenghe Wu, Qingjin Peng, and Quan Zhang. 2022. Two-dimensional irregular packing problems: A review. *Frontiers in Mechanical Engineering* 8 (2022). <https://doi.org/10.3389/fmech.2022.966691>

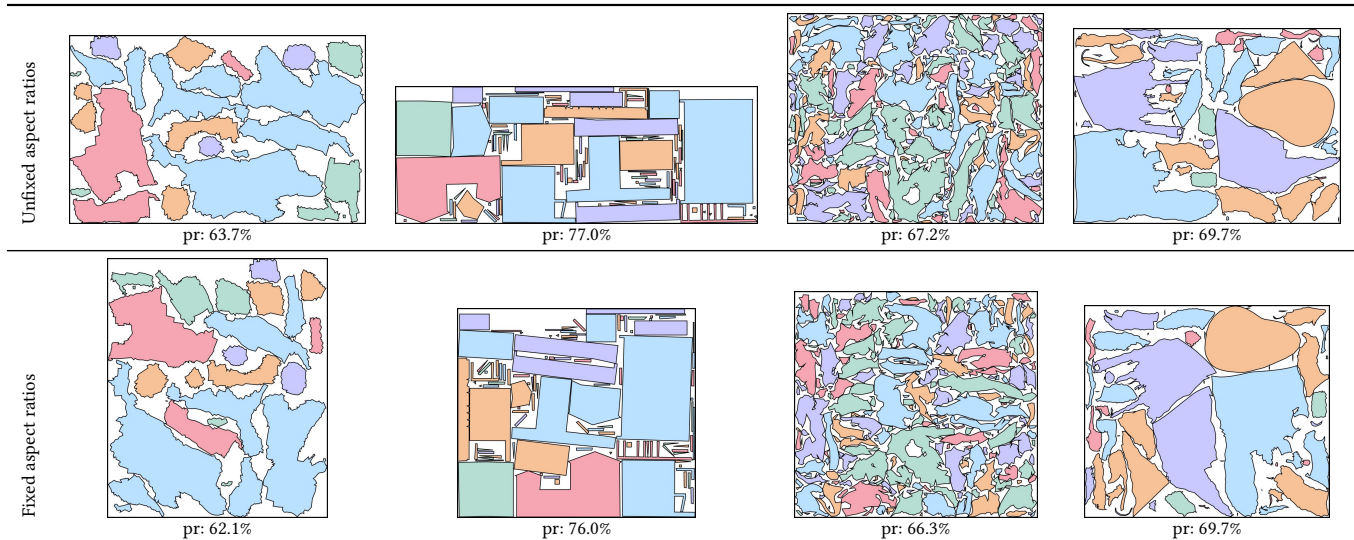


Fig. 11. Packing results with and without fixed aspect ratios.




Fig. 12. Sub-optimal packing instances generated by our pipeline.

- Juris Hartmanis. 1982. Computers and intractability: a guide to the theory of np-completeness (michael r. garey and david s. johnson). *Siam Review* 24, 1 (1982), 90.
- Ruizhen Hu, Juzhan Xu, Bin Chen, Minglun Gong, Hao Zhang, and Hui Huang. 2020. Tap-net: transport-and-pack using reinforcement learning. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.

- Yannan Hu, Sho Fukatsu, Hideki Hashimoto, Shinji Imahori, and Mutsunori Yagiura. 2018. Efficient overlap detection and construction algorithms for the bitmap shape packing problem. *Journal of the Operations Research Society of Japan* 61, 1 (2018), 132–150.
- Yuan Jiang, Zhiguang Cao, and Jie Zhang. 2021. Learning to solve 3-D bin packing problem via deep reinforcement learning and constraint programming. *IEEE transactions on cybernetics* (2021).
- Zhongshi Jiang, Scott Schaefer, and Daniele Panozzo. 2017. Simplicial complex augmentation framework for bijective maps. *ACM Transactions on Graphics* 36, 6 (2017).
- Qingdi Ke, Peng Zhang, Lei Zhang, and Shouxu Song. 2020. Electric vehicle battery disassembly sequence planning based on frame-subgroup structure combined with genetic algorithm. *Frontiers in Mechanical Engineering* 6 (2020), 576642.
- Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Mailliot. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM transactions on graphics (TOG)* 21, 3 (2002), 362–371.
- Max Limper, Nicholas Vining, and Alla Sheffer. 2018. Box cutter: atlas refinement for efficient packing via void elimination. *ACM Trans. Graph.* 37, 4 (2018), 153–1.
- Hao-Yu Liu, Xiao-Ming Fu, Chunyang Ye, Shuangming Chai, and Ligang Liu. 2019. Atlas Refinement with Bounded Packing Efficiency. *ACM Transactions on Graphics (SIGGRAPH)* 38, 4 (2019), 33:1–33:13.
- Hu-yao Liu and Yuan-jun He. 2006. Algorithm for 2D irregular-shaped nesting problem based on the NFP algorithm and lowest-gravity-center principle. *Journal of Zhejiang University-Science A* 7, 4 (2006), 570–576.
- Tobias Nöll and D Strieker. 2011. Efficient packing of arbitrary shaped charts for automatic texture atlas generation. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1309–1317.
- Semih Önüt, Umut R Tuzkaya, and Bilgehan Doğaç. 2008. A particle swarm optimization algorithm for the multiple-level warehouse layout design problem. *Computers & Industrial Engineering* 54, 4 (2008), 783–799.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).
- Roi Poranne, Marco Tarini, Sandro Huber, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Autocuts: simultaneous distortion and cut optimization for UV mapping. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 1–11.
- Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Scalable locally injective mappings. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1.
- Maria Cristina Riff, Xavier Bonnaire, and Bertrand Neveu. 2009. A revision of recent approaches for two-dimensional strip-packing problems. *Engineering Applications of Artificial Intelligence* 22, 4-5 (2009), 823–827.
- Pedro V Sander, John Snyder, Steven J Gortler, and Hugues Hoppe. 2001. Texture mapping progressive meshes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. 409–416.

- Pedro V Sander, Zoë J Wood, Steven Gortler, John Snyder, and Hugues Hoppe. 2003. Multi-chart geometry images. (2003).
- Nico Schertler, Daniele Panozzo, Stefan Gumhold, and Marco Tarini. 2018. Generalized motorcycle graphs for imperfect quad-dominant meshes. *ACM Transactions on Graphics* 37, 4 (2018).
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- Jason Smith and Scott Schaefer. 2015. Bijective parameterization with free boundaries. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–9.
- Olga Sorkine, Daniel Cohen-Or, Rony Goldenthal, and Dani Lischinski. 2002. Bounded-distortion piecewise mesh parameterization. In *IEEE Visualization, 2002. VIS 2002. IEEE*, 355–362.
- Marc Soucy, Guy Godin, and Marc Rioux. 1996. A texture-mapping approach for the compression of colored 3D triangulations. *The Visual Computer* 12 (1996), 503–514.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 30, 1 (2016).
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat* 1050, 20 (2017), 10–48550.
- Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. 2019. Deep Graph Library: A Graph-Centric, Highly-Performant Package for Graph Neural Networks. *arXiv preprint arXiv:1909.01315* (2019).
- Shiyi Wang, Jiong Chen, Xifeng Gao, Hujun Bao, and Jin Huang. 2022. 3D mesh cutting for high quality atlas packing. *Computer Aided Geometric Design* 99 (2022), 102149.
- Shuo Yang, Shuai Song, Shilei Chu, Ran Song, Jiyu Cheng, Yibin Li, and Wei Zhang. 2023. Heuristics Integrated Deep Reinforcement Learning for Online 3D Bin Packing. *IEEE Transactions on Automation Science and Engineering* 0, 0 (2023), 1–12. <https://doi.org/10.1109/TASE.2023.3235742>
- Jonathan Young. 2023. Xatlas. Retrieved May, 2023 from <https://github.com/jpcy/xatlas>
- Chi Zhang, Mao-Feng Xu, Shuangming Chai, and Xiao-Ming Fu. 2020. Robust atlas generation via angle-based segmentation. *Computer Aided Geometric Design* 79 (2020), 101854.
- Hang Zhao, Yang Yu, and Kai Xu. 2021. Learning efficient online 3d bin packing on packing configuration trees. In *International Conference on Learning Representations*. 0–0.
- Hang Zhao, Chenyang Zhu, Xin Xu, Hui Huang, and Kai Xu. 2022. Learning practically feasible policies for online 3D bin packing. *Science China Information Sciences* 65, 1 (2022), 112105.
- Qingnan Zhou and Alec Jacobson. 2016. Thingi10k: A dataset of 10,000 3d-printing models.

A NETWORK ARCHITECTURE

We summarized the detailed architecture of our three network in Tab. 7, Tab. 8, and Tab. 9. The full network pipeline of LPN is defined as:

$$\begin{aligned} (\bar{P}_{i-1}, \bar{p}_i, \dots, \bar{p}_{i+H-1}) &\leftarrow (\text{FCN}(P_{i-1}), \text{FCN}(p_i), \dots, \text{FCN}(p_{i+H-1})) \\ &\triangleq \bar{s}_i \\ Q_{256} &\leftarrow \pi^{\text{LPN}}(\bar{s}_i), \end{aligned}$$

where the output is 256-dimensional Q-values over the action space. The full network pipeline of LSN is defined as:

$$\begin{aligned} \bar{s}_i &\leftarrow \text{FCN}(s_i) \\ (Q_{p_i}, \dots, Q_{p_{i+H-1}}) &\leftarrow \pi^{\text{LSN}} \left[\left(\begin{array}{c} \bar{P}_{i-1} \\ \bar{p}_i \end{array} \right), \dots, \left(\begin{array}{c} \bar{P}_{i-1} \\ \bar{p}_{i+H-1} \end{array} \right) \right], \end{aligned}$$

where the output is the H -dimensional Q-value of each candidate next patch. Finally, the full network pipeline of HSN is defined as:

$$\text{pr}(S') \leftarrow \text{HSN}(\bar{p}_1, \dots, \bar{p}_H),$$

which outputs the predicted weighted average packing ratio.

Table 6. The architecture of FCN.

Layer Type	Input Size	Output Size
Convolution	(batch size, 1, 50, 50)	(batch size, 6, 24, 24)
ReLU	(batch size, 6, 24, 24)	(batch size, 6, 24, 24)
Convolution	(batch size, 6, 24, 24)	(batch size, 12, 12, 12)
ReLU	(batch size, 12, 12, 12)	(batch size, 12, 12, 12)
Convolution	(batch size, 12, 12, 12)	(batch size, 12, 6, 6)
ReLU	(batch size, 12, 6, 6)	(batch size, 12, 6, 6)
Flatten	(batch size, 12, 6, 6)	(batch size, 432)

Table 7. The architecture of LPN.

Layer Type	Input Size	Output Size
Linear	(batch size, $432 \times (1 + H)$)	(batch size, 1024)
ReLU	(batch size, 1024)	(batch size, 1024)
Linear	(batch size, 1024)	(batch size, 512)
ReLU	(batch size, 512)	(batch size, 512)
Linear	(batch size, 512)	(batch size, 256)
ReLU	(batch size, 256)	(batch size, 256)
Linear	(batch size, 256)	(batch size, 256)

Table 8. The architecture of LSN.

Layer Type	Input Size	Output Size
GAT	(batch size, $H, 432 \times 2$)	(batch size, $H, 512$)
ELU	(batch size, $H, 512$)	(batch size, $H, 512$)
GAT	(batch size, $H, 512$)	(batch size, $H, 256$)
ELU	(batch size, $H, 256$)	(batch size, $H, 256$)
Linear	(batch size, $H, 256$)	(batch size, $H, 128$)
ReLU	(batch size, $H, 128$)	(batch size, $H, 128$)
Linear	(batch size, $H, 128$)	(batch size, $H, 64$)
ReLU	(batch size, $H, 64$)	(batch size, $H, 64$)
Linear	(batch size, $H, 64$)	(batch size, $H, 32$)
ReLU	(batch size, $H, 32$)	(batch size, $H, 32$)
Linear	(batch size, $H, 32$)	(batch size, $H, 1$)

Table 9. The architecture of HSN.

Layer Type	Input Size	Output Size
GAT	(batch size, $H, 432$)	(batch size, $H, 256$)
ELU	(batch size, $H, 256$)	(batch size, $H, 256$)
GAT	(batch size, $H, 256$)	(batch size, $H, 64$)
ELU	(batch size, $H, 64$)	(batch size, $H, 64$)
Max Pooling	(batch size, $H, 64$)	(batch size, 64)
ReLU	(batch size, 64)	(batch size, 32)
Linear	(batch size, 32)	(batch size, 1)
Sigmoid	(batch size, 1)	(batch size, 1)

B ALGORITHM PIPELINE

We summarize our training procedure in Alg. 2 and runtime pipeline in Alg. 3.

Algorithm 2 Training π^{LPN} , π^{LSN} , and HSN

Input: A dataset of packing problems $\mathbb{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots\}$
Output: π^{LPN} , π^{LSN} , HSN

- 1: \triangleright Initial training
- 2: **while** not converged **do**
- 3: Sample $\mathcal{S} \in \mathbb{S}$
- 4: Sample *ordered* subset of H patches $\mathcal{S}' \subseteq \mathcal{S}$
- 5: Use MDP of Sec. 4.1 to populate experience buffer
- 6: Update π^{LPN} via DDQN
- 7: **while** not converged **do**
- 8: Sample $\mathcal{S} \in \mathbb{S}$
- 9: Sample *unordered* subset of H patches $\mathcal{S}' \subseteq \mathcal{S}$
- 10: Use MDP of Sec. 4.2 and π^{LPN} to populate experience buffer
- 11: Update π^{LSN} via DDQN
- 12: **while** not converged **do**
- 13: Sample $\mathcal{S} \in \mathbb{S}$
- 14: Sample two *unordered* subsets of H patches $\mathcal{S}', \mathcal{S}'' \subseteq \mathcal{S}$
- 15: Use π^{LPN} and π^{LSN} to compute groundtruth pr' and pr''
- 16: Update HSN using SGD on loss Eq. 2

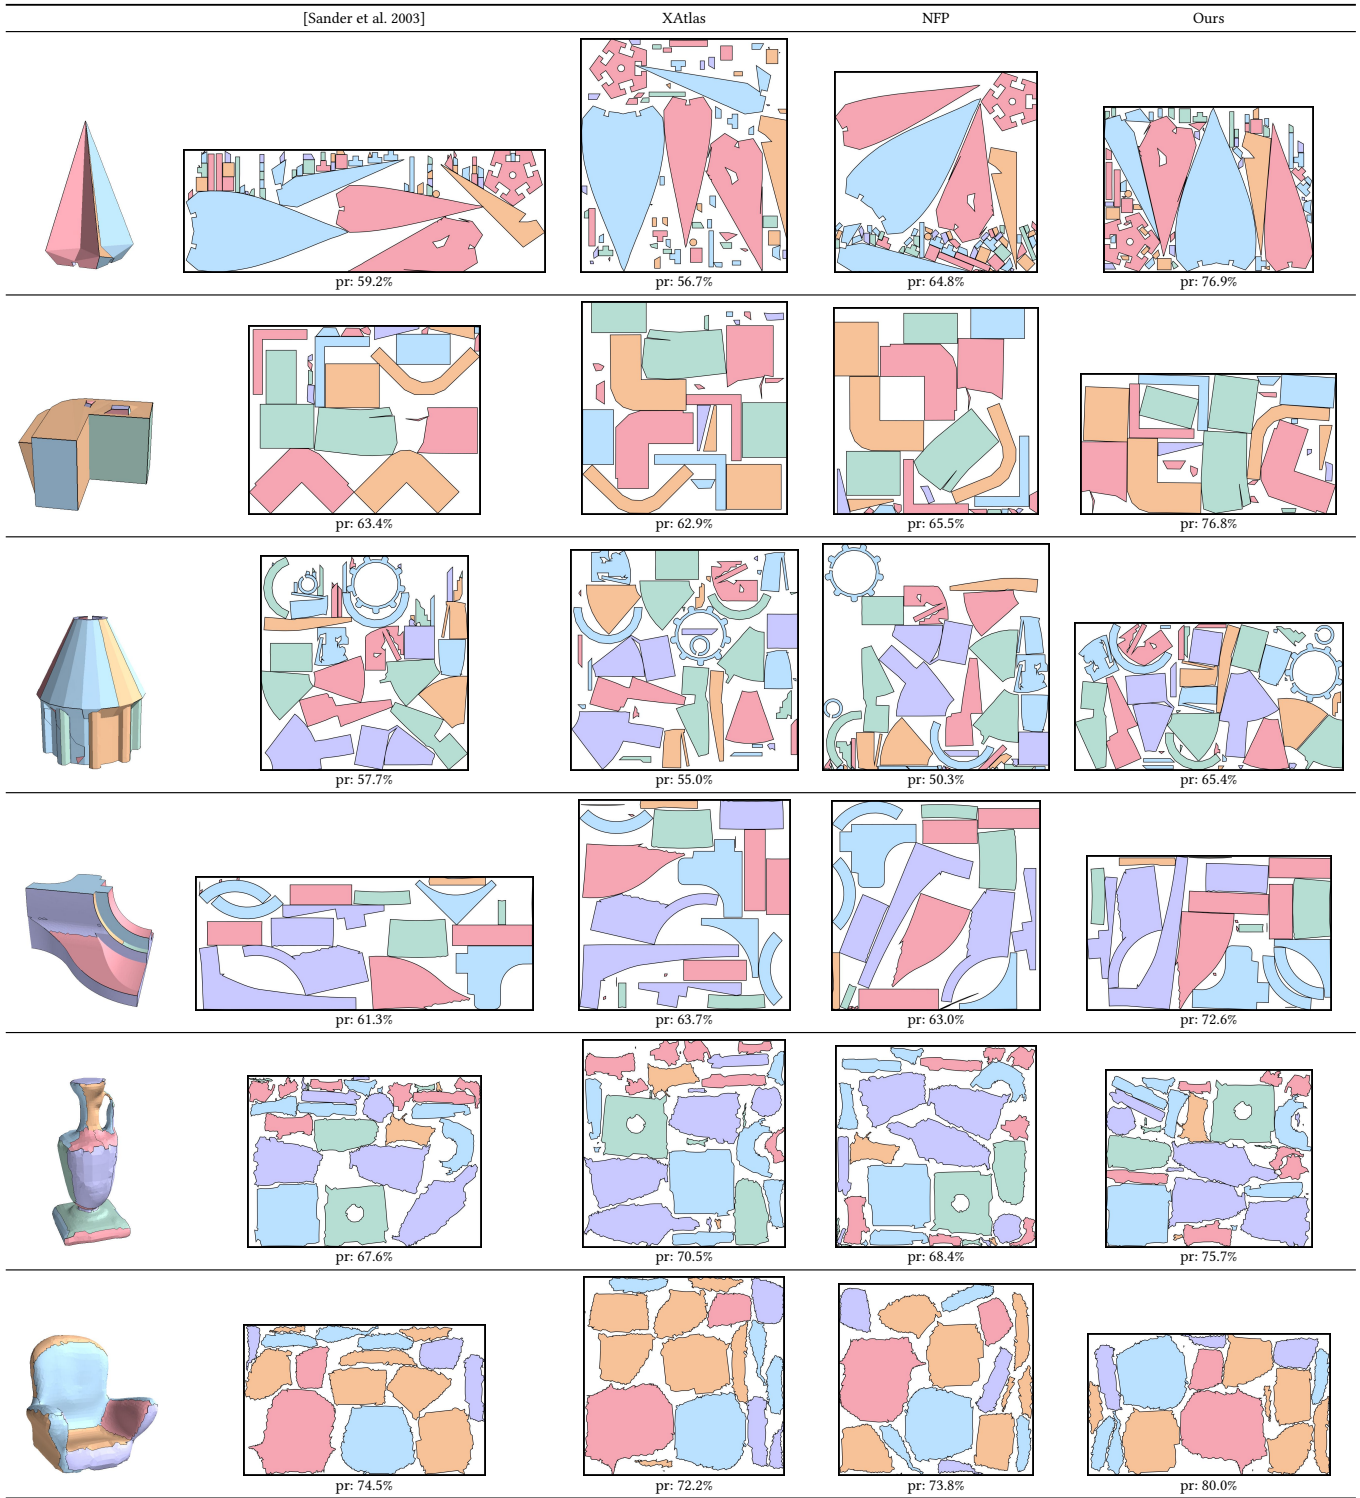
Algorithm 3 Learning-Assisted UV Packing































Input: A packing problem \mathcal{S}
Output: Pose for each $p_i \in \mathcal{S}$































- 1: \triangleright Exclude small patches
- 2: $\mathcal{S}_j \leftarrow \emptyset$
- 3: Calculate average area of salient patches \bar{a}
- 4: **for** $p \in \mathcal{S}$ **do**
- 5: **if** $\text{area}(p) < \bar{a}/5$ **then**
- 6: $\mathcal{S} \leftarrow \mathcal{S} - \{p\}$
- 7: $\mathcal{S}_j \leftarrow \mathcal{S}_j \cup \{p\}$
- 8: \triangleright Main loop
- 9: Compute $\text{pr}(\mathcal{S})$
- 10: **while** not converged **do**
- 11: Sample 400 random subsets of H patches $\mathcal{S}'_{1, \dots, 400}$
- 12: Sort \mathcal{S}'_i in HSN(\mathcal{S}'_i)-descending order
- 13: **for** $i = 1, \dots, 10$ **do**
- 14: Compute groundtruth $\text{pr}_i \leftarrow \text{pr}(\mathcal{S} - \mathcal{S}'_i \cup \{\text{LL}(\mathcal{S}'_i)\})$
- 15: Sort $\mathcal{S}'_1, \dots, \mathcal{S}'_{10}$ in pr_i -descending order
- 16: **if** $\text{pr}_1 > \text{pr}(\mathcal{S})$ **then**
- 17: Compute alpha shape for $\text{LL}(\mathcal{S}'_1)$
- 18: $\mathcal{S} \leftarrow \mathcal{S} - \mathcal{S}'_1 \cup \{\text{LL}(\mathcal{S}'_1)\}$
- 19: $\text{pr}(\mathcal{S}) \leftarrow \text{pr}_1$
- 20: **else** Break
- 21: Bin-packing all patches in \mathcal{S}
- 22: Locally squeeze patches via Eq. 3
- 23: \triangleright Pack small patches
- 24: **for** $p \in \mathcal{S}_j$ in area-descending order **do**
- 25: $\mathcal{S} \leftarrow \text{scanline}(\mathcal{S}, p)$
- 26: Return all poses

C MORE RESULTS

We show more results packed by various baselines.



	[Sander et al. 2003]	XAtlas	NFP	Ours
	 pr: 72.5%	 pr: 63.2%	 pr: 70.8%	 pr: 79.7%
	 pr: 68.0%	 pr: 64.7%	 pr: 64.2%	 pr: 74.3%
	 pr: 65.3%	 pr: 62.9%	 pr: 67.2%	 pr: 78.3%
	 pr: 64.8%	 pr: 63.7%	 pr: 66.5%	 pr: 73.9%
	 pr: 64.8%	 pr: 63.7%	 pr: 66.5%	 pr: 74.9%
	 pr: 66.8%	 pr: 66.1%	 pr: 64.6%	 pr: 76.1%

	[Sander et al. 2003]	XAtlas	NFP	Ours
	 pr: 63.9%	 pr: 64.0%	 pr: 67.5%	 pr: 73.5%
	 pr: 58.2%	 pr: 64.9%	 pr: 65.7%	 pr: 71.7%
	 pr: 71.7%	 pr: 73.4%	 pr: 69.0%	 pr: 79.3%
	 pr: 67.2%	 pr: 69.7%	 pr: 62.5%	 pr: 77.5%
	 pr: 67.7%	 pr: 70.7%	 pr: 65.6%	 pr: 75.9%
	 pr: 67.0%	 pr: 70.3%	 pr: 68.9%	 pr: 79.6%