Planning for Robotic Dry Stacking with Irregular Stones

Yifang Liu, Jiwon Choi, and Nils Napp

1 Introduction

We critically depend on modifying our environment, e.g., by constructing shelters, infrastructure for transportation, water, energy, and waste management, as well as structures that regulate the natural environment such as dams, drainage, and protection against avalanches. Labor and materials are the main cost drivers of the construction industry, which also produces approximately 500 million tons of demolition waste, mostly in the form of concrete [1]. Cement production accounts for (≈%5) of global CO_2 emission [2]. Robotic construction with in situ (found) materials simultaneously addresses primary cost drivers of construction, while mitigating its environmental impact. This idea has been explored in specialized situations. Driven by the need for resource conservation in space, NASA has studied in situ material use for extraterrestrial environments. Launching building materials into space is very costly, yet simple structures–such as berms, walls, and shelters–might be readily built from minimally processed but rearranged materials [3]. Such *utility structures*, i.e., structures that have a specific function, but whose exact shape matters less, are also important on Earth. Examples include erosion barriers for changing coastlines, temporary support structures in disaster sites, or containment structures made from contaminated materials from a nuclear or chemical leak. One particularly well-suited

Y. Liu (✉) · N. Napp
School of Electrical and Computer Engineering, Cornell University,
Ithaca, NY 14853, USA
e-mail: yl892@cornell.edu

N. Napp
e-mail: nnapp@cornell.edu

J. Choi
Department of Computer Science and Engineering, University at Buffalo,
Buffalo, NY 14260, USA
e-mail: jiwoncho@buffalo.edu

© Springer Nature Singapore Pte Ltd. 2021

G. Ishigami and K. Yoshida (eds.), *Field and Service Robotics*, Springer Proceedings
in Advanced Robotics 16, https://doi.org/10.1007/978-981-15-9460-1_23

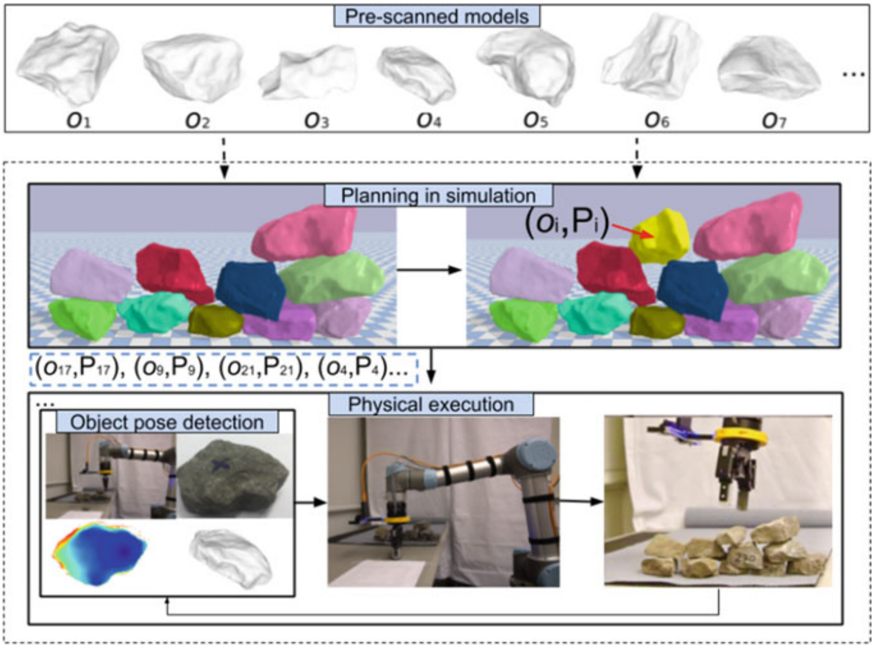

Fig. 1 System overview. All the stone models are pre-scanned. The planning is conducted in a physics simulator, then the assembly sequences are executed open-loop with a robotic arm

construction method for such types of utility structures is dry-stacking stones. This ancient method was practiced by humankind since 3500 B.C. [4], and makes up some of the oldest man-made structures. Theoretically, robots are ideally suited for this work, since robots make work safer and physically less demanding. A lack of understanding of how to pose and solve assembly planning problems of irregular natural material into in situ functional structures, however, is currently hindering robots from performing such useful construction tasks.

We present an algorithmic approach for solving the planning problem of assembling stable structures from a collection of irregularly shaped rigid objects. The application is to enable dry stacking with found, minimally processed rocks. We focus on the problem of high-level placement planning for rocks to build stable structures, and dry-stack structures with tens of rocks, which significantly improves the state of the art. These plans are executed open-loop without additional tactile sensing; however, our results suggest that large-scale dry-stacking robots would benefit from better physical feedback during the construction process. The whole process is shown in Fig. 1. To the best of our knowledge, this is the first system that can automatically dry-stack a wall with 4 courses using natural irregular rocks, which could significantly benefit the large-scale outdoor construction robots.

The rest of the paper is structured as follows: Sect. 2 provides a brief overview of related work; Sect. 3 describes the planning algorithm; results are presented in Sect. 4. Finally, Sect. 5 concludes the paper and discusses future work.

2 Related Work

The fundamental questions in autonomous construction are how to specify, plan, and execute the placement of building elements to achieve a final structure. Approaches [5–8] differ in each of these aspects, and range from determining the assembly order of elements whose position is known in advance [5] to formulating building plans that need to pick the type, shape, and pose of elements to build approximate shapes, or to build structures that fulfill specific functions [9].

There are many examples of multi-robot construction [10]. Prominent examples include 3D construction according to blueprints with climbing [11] and flying [12] robots. These systems rely on uniform, custom-made rigid elements that snap together using magnets, and focus on combinations of pre-processed compilation steps and local runtime control. Another body of work is focused on distributed collective robotic construction of functional structures composed of deformable and amorphous materials [9, 13–16]. In these demonstrations, the deformable nature of the building material compensates for placement inaccuracies and environmental irregularity which simplifies planning.

When building with rigid irregular objects, small surface features substantially affect the friction and stability. Compounding this difficulty, microfracture formation during execution may deform the structure surface. This makes planning and stable placement of irregular objects fundamentally different from building with regular objects that have predictable contact geometry. In our previous work [8, 17, 18], we proposed an architecture for solving the dry-stacking problem, based on heuristics and deep Q-learning to build stable, large-scale structures using physics simulation in 2D.

We leverage guidelines for building a stable structure from instructional books for dry-stacked masonry [19–21]. For example, it is good to place large stones with inward-sloping top surfaces on the corners, and smaller stones in the middle. Such heuristics can provide a structured approach in making assembly decisions, but in their description, much is left to experience and human judgment.

The most closely related work is Furrer et al. [6], who propose a pose search algorithm that considers structural stability using a physics simulator. In addition, they present an autonomous system, using a robot manipulator, for stacking stable vertical towers with irregular stones. The pose search cost function considers support polygon area, kinetic energy, the deviation between thrust line direction and the normal of the support polygon surface, and the length between the new object and the center of mass (CoM) of the previously stacked object. While this paper also uses a physics simulator, we use it primarily to find a finite set of feasible poses, and then apply a layered filtering approach, which we found results in both better and more robust performance.

3 Methods

In this section, we first describe the notation used in this paper; then we elaborate on the planning algorithm for stacking irregular stones. Finally, we provide the object pose detection pipeline used in physical execution.

3.1 Notation

The *world frame* is a 3D coordinate system where the gravity is in the negative z-direction, and the goal is to construct a target structure $T \subset \mathbb{R}^3$, i.e., a subset of the world space that should be filled by selecting and placing elements from a set of objects O. Each object is a connected subset of $\mathbb{R}^3$ with its origin at the CoM.

An assembly $A = (a_1, a_2, ..., a_I)$ is a set of I assembled objects, where each element $a_i = (o_i, P_i)$ is a pair containing an object $o_i \in O$ and its pose $P_i = (p_i, R_i) \in SE(3)$. The position $p_i \in \mathbb{R}^3$ denotes the CoM position of object o_i in the world frame, and $R_i \in SO(3)$ is its orientation. Empty space set is a set $E \subset T$ s.t. every point $e \in E$ can be connected by a straight line from ∂T to ∂A without passing through any other a_i and ∂T, and ∂A denote the boundaries of T and the assembly, respectively. This definition excludes the complicated internal voids created by stacking irregular objects from counting as empty space. The top surface is given by $S = \partial E \cap \partial A$, i.e., the overlapping area of the empty set E and the assembly A, where ∂E denotes the boundaries of empty space E. We define the action space $X(o_i)$ of object o_i to be restricted to have its CoM in E:

$$X(o_i) = \{(p_i, R_i) | p \in E\}. \tag{1}$$

The world is initially assumed to be empty of objects, aside from a support surface at the bottom of T. We want to find an assembly strategy for autonomous agents, i.e., picking a sequence of elements o_i and actions from $X(o_i)$ to build an assembly that occupies the target structure T subject to physical contact, friction, and gravity constraints.

3.2 Structure Planning

This section presents an assembly planning algorithm for irregular objects. Similar to [8], we design a greedy heuristic approach to find the next best pose from a set of feasible poses for a given object. For each object, we use a physics simulator to generate a finite set of feasible *stable* poses, strategically reduce this set, and choose the best available pose. By repeating this sequence, we incrementally build the structure.

Algorithm 1: Feasible Pose Generation

```
Data: o_i: object
Result: Feasible Poses Set
{(x_j, y_j, z_j)} ← discretize position;
for each (x,y,z) in {(x_j, y_j, z_j)} do
    for N_ori ← 0 to N do
        R ← random generate an orientation;
        reset o_i pose to (x, y, z, R);
        while N_contact(o_i) < 3 do
            step physics simulation once;
        end
        pause physics simulation;
        reset o_i linear and angular velocity to 0;
        step physics simulation once;
        while o_i is not stable do
            step physics simulation once;
        end
        if distance between current pose P_i and (x, y, z, R) < Threshold then
            add current pose P_i to Feasible Pose Set X_F;
        else
            continue;
        end
    end
end
```

3.2.1 Feasible Pose Generation

We use a physics simulator to find *physically stable* configurations. We approximate the real-world state with simulation and provide a practical and efficient stability estimate of the system without actually having to physically interact with the external world. This helps us to acquire a good prior estimate for the system.

Since the action space (p and R) is continuous (Eq. 1), we first sample the action space in such a way that the position p is discretized, and each position p corresponds to a set of randomly sampled orientations R. We then make use of a rigid body simulator to find physically stable configurations. The simulator proceeds by first selecting an initial placement (position and orientation) for a given object on the surface of the built structure, and then simulating the forces acting on the object until it settles into a stable pose; see Algorithm 1. Although the number of possible initial placements is large, a substantial amount of them settles down into a small subset of feasible poses. This set of feasible poses for an object o_i denotes $X_F(o_i)$. Even though all the poses in the feasible pose set are stable, many of them are poor choices and result in low overall stability. In the next section, we will discuss how to refine this set by using heuristics gathered from instructional literature for masonry books [19–21].

3.2.2 Action Space Reduction

The refinement of action space is a hierarchical filtering approach, where each filter removes poses that do not meet the minimum requirement for a satisfactory placement according to a specific heuristic. The set of filters used in this work is presented as follows:

- Support polygon area: the area of an object's support polygon. A higher value of support polygons correlates to a stable footing for the object. Similar to the method presented in [6], we robustly find the support polygon from the sparse contacts by updating the simulator 10 steps and collecting all the contacts. Then, Principal Component Analysis (PCA) is used to reduce the 3D contacts to 2D points. Finally, the convex hull of these 2D points is calculated as a support polygon area.
- Normal of support polygon: the normal direction of the support polygon. It measures how much the normal direction deviates from the thrust line direction vector [6].
- Neighbor height: the difference in heights of the object, after its placement, with its left and right neighboring objects. The height of an object is represented by CoM height. This feature helps maintain leveled surfaces in the structure.
- Stone top surface slope: the top surface angle of the object at a given pose. In building a wall, we prefer inward sloping angles to prevent stones from the top layers to fall down from the structure [21, Pg. 49].
- Interlocking: the number of objects in the structure that are in contact with the current object at a given pose. The use of this feature allows for staggered layering and thus helps to prevent vertical stacking in the structure [20, Pg. 19].

Each filter evaluates one of these features. The reductions are applied hierarchically for each object as follows:

- The original feasible pose set denotes X_F.
- At *Filter 1*, only select poses that have an inward sloping top surface angle. The remaining poses after applying this filter denote X_{F1}.
- At *Filter 2*, discard poses with dot product $\|\mathrm{n}_i \cdot \mathrm{v}_i\|$ less than the mean of all poses from X_{F1}, where $\mathrm{n_i}$ represents the normal direction of the contact polygon, and $\mathrm{v_i}$ is the thrust line direction vector. The set of poses after applying this filter denotes X_{F2}.
- At *Filter 3*, remove poses with support polygon area less than the mean of the support polygon area value of all stable poses. The remaining poses after applying this filter denote X_{F3}.

- At *Filter 4*, this is based on the current state of the structure. If it is not a corner placement, we only choose poses whose centroid heights are lower than the average centroid heights of corner stones at the current course. The set of poses after applying this filter denotes X_{F4}.
- At *Filter 5*, remove the poses whose number of interlocking objects are smaller than the mean number of interlocking. The set of poses after applying this filter denotes X_{F5}.

This hierarchical reduction model is carefully designed such that a random pose at each level is more desirable than a random pose drawn at the earlier filtered levels. It is also designed such that no good possible stable poses are removed earlier before reaching the final selection filter. The relation between the various sets of poses is shown in Eq. 2.

$$X_F \supset X_{F1} \supset X_{F2} \supset X_{F3} \supseteq X_{F4} \supset X_{F5}, \tag{2}$$

where X_F is defined in Sect. 3.2.1.

Unlike the pose searching algorithm used in [6], which combines terms similar to Filters 2 and 3, as well as other heuristics into a single scalar cost function and finds poses by gradient descent, the planning algorithm proposed in this paper first considers geometric and physical constraints using a simulator to find a discrete set of feasible actions and further refines this set by using a hierarchical filter based on heuristics gathered from the instructional materials. This approach eliminates the need for tuning the relative weights in a scalar cost function. Without the need for cost-tuning, the algorithm is more adaptable to different stones with various physical properties, such as size, density, and friction. The reason is that with a single scalar function the weights are coupled and the relative importance depends on the set of objects's physical properties. However, in the hierarchical filter, each term is assessed in isolation and thus the method is less sensitive to the change of physical properties in the set of objects.

3.2.3 Proposed Algorithm

Algorithm 2 describes how a structure is constructed. The inputs are the set of available objects along with the target structure to be built. During construction, it builds the structure course by course, and within courses, it first places the corner stones with the inward slope in the two course extrema, as shown in Algorithm 2 Line 2; then it builds the middle area within the course (Lines 3–5). The output is the set of assembly steps.

Algorithm 3 describes the steps to select an object and its pose for the placement. The inputs are the set of remaining objects and their type (corner stone or random stone), since different object types may require different hierarchical filters. The first step is to choose a random object (Line 3) and collect the feasible poses (Sect. 3.2.1) of this object (Line 4). Then, it applies the Hierarchical Filter (Sect. 3.2.2) to reduce the action space at Line 5. If the reduced action space is not empty, we select one pose from it; otherwise, we try this procedure again for a different object until it reaches the maximum number of trials (Lines 6–11).

Algorithm 2: Proposed Assembly Approach

Data: O: object dataset, T: target structure
Result: Assembly steps
while *target area T still has room left to build* **do**
    place corner stones with inward slope in the two course extrema;
    while *current course still has room left to build* **do**
        place stone in the current course ;
    end
end

Algorithm 3: Place Stone

Data: B: set of available objects ($B \subseteq O$), object type
Result: Placed object pose
1 $n \leftarrow 0$;
2 **while** $n \leq$ *Maximum Number of Trials* **do**
3     $b \leftarrow$ randomly choose one object from B;
4     $X_F \leftarrow$ feasible poses set;
5     $X_{final} \leftarrow$ apply Hierarchical Filter to X_F ;
6   **if** $X_{final} \neq \emptyset$ **then**
7   place one of the X_{final} poses;
8   return;
9   **else**
10         $n \leftarrow n + 1$;
11   **end**
12 **end**

3.3 Object Pose Detection

During physical execution, we first detect the object pose in the scene. We start by capturing a set of point cloud data of an object from different views via an RGB-D camera; we then filter out the points that do not belong to the current object by removing the plane points from point cloud data using Point Cloud Library (PCL) [22]; next, we merge the remaining point cloud data. We apply global registration to provide an initial transformation and Iterative Closest Point (ICP) algorithm to further refine the transformation using the Open3D library [23]; finally, we run registration on merged point cloud data and pre-scanned 3D mesh of the object to get the relative pose between them. Similar to the third step above, the registration also contains global registration and ICP. The whole pose detection pipeline is shown in Fig. 2.

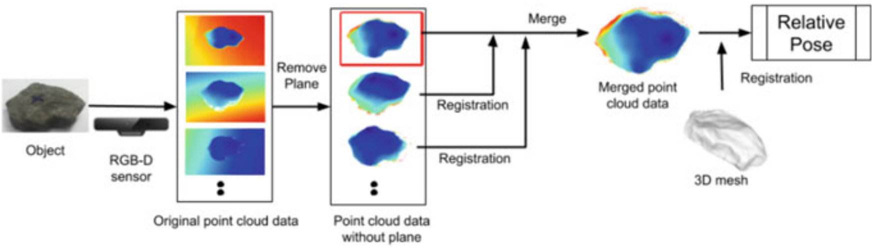

Fig. 2 Object pose detection pipeline showing the sequence of detection, registration, and merging. See Sect. 3.3 for more details

Once the relative pose between the current object and the 3D mesh is detected, we use the manipulator to pick up the stone and apply the same transformation to the end-effector of the manipulator to place the stone as the planned pose.

4 Experiments

In this section, we first describe the experimental setup, then we show the stone towers and stone walls using the proposed planning, as well as the comparison between the pose searching algorithm proposed in [6] and the proposed method.

4.1 Experimental Setup

As shown in Fig. 4, a UR5 manipulator equipped with a ROBOTIQ 2-Finger gripper is used in the manipulation task. An Intel® RealSense™ SR300 RGB-D camera is attached to the UR5 arm for point cloud data acquisition. The MoveIt! [24] package is employed for motion planning. We collected 23 shale stones as irregular objects, which are specifically selected such that they fit the size of the gripper. The average weight of the selected stones is 193 g with a standard deviation of 90 g. The outer bounding box size of the stones are 0.0791 ± 0.0144, 0.0585 ± 0.0086, and 0.04 ± 0.0088 m. During physical execution, objects are manually fed into the pickup area. The gripper grasping position varies depending on the detected object position, but the gripper orientation remains the same. The stone 3D model is acquired with a Matter and Form 3D Desktop Scanner. Figure 3 shows some samples of the stones and their corresponding 3D mesh models. The object pose detection and manipulation parts are implemented using the Robot Operating System.

Fig. 3 Irregular shale stones and their corresponding 3D meshes

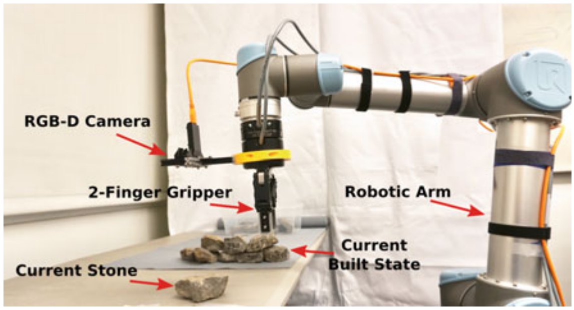

Fig. 4 An overview of the experimental setup

4.2 Results

The autonomous building system is shown in Fig. 1. In this section, we first compare the proposed algorithm with other work in simulation and physical execution; then, we show the physical execution results.

4.2.1 Stone Tower

Our goal was to build a vertical stone tower using the pre-scanned 3D mesh models as a test structure to evaluate planning under stability constraints. We compared the proposed method with the pose searching algorithm proposed in [6]. The comparisons were conducted in PyBullet physics simulation [25].

The pose searching method used in [6] places each object on the top object of the existing stacking using a physics engine. A cost function is introduced to evaluate the "goodness" of each pose, which considers 4 elements: contact area C_i, kinetic energy E_{kin}, the length between the newly placed object pose P_j and the previously placed object pose P_i (denoted as $\mathrm{r}_{P_j P_i}$), and the dot product between normal of the contact polygon and the trust line direction vector $\|\mathrm{n}_i \cdot \mathrm{v}_i\|$. The cost function is defined as

$$f(P_i) = w_1 C_i^{-1} + w_2 E_{kin}(P_i) + w_3 \|\mathrm{r}_{P_j P_i}\| + w_4 \|\mathrm{n}_i \cdot \mathrm{v}_i\| \ , \tag{3}$$

where w_j are tuned for the object set. After assigning the cost to the valid contact pose, gradient descent is used to search the local optimal pose P_i^*.

Since we use a different type of stone from that of [6], and the size of the stones are also different given that we use different arms and grippers, the w_j given by [6] are not optimal for our application. We also modify the last component of Eq. 3 as

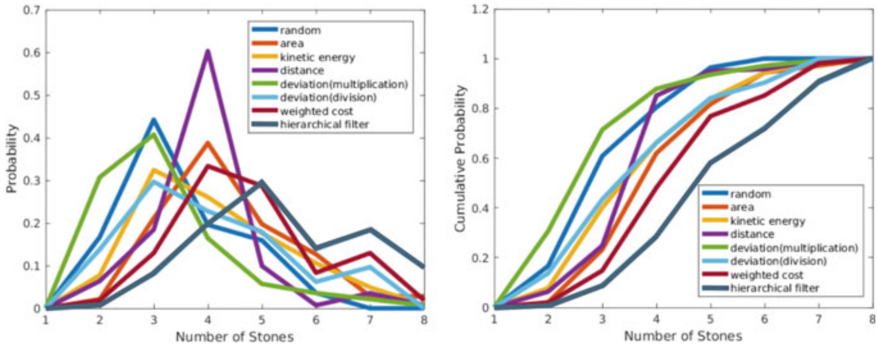

Fig. 5 Vertical tower building results. For each method, we build 150 different vertical stone towers in simulation. The x-axis shows the number of stones each tower has, and on the left figure the y-axis shows the percentage of each height; on the right figure, the y-axis represents cumulative percentage of each height. "random" means that we randomly pick a pose from feasible pose set; "area" represents that the cost function only contains contact polygon area (C) one element, so does "kinetic energy" (E_{kin}), "distance" ($\mathrm{r}_{P_j P_i}$), and "deviation" ($\|\mathrm{n}_i \cdot \mathrm{v}_i\|$). For "deviation", we test both multiplication $\|\mathrm{n}_i \cdot \mathrm{v}_i\|$ and division $\|\mathrm{n}_i \cdot \mathrm{v}_i\|^{-1}$. The "weighted cost" uses the optimized cost function. The proposed "hierarchical filter" significantly outperforms the other methods

$w_4\|\mathrm{n}_i \cdot \mathrm{v}_i\|^{-1}$ given the fact that the larger the dot product is, the smaller the cost should be. At last, we apply Bayesian optimization for Gaussian process modeling called GPyOpt [26] to optimize the weights.

Figure 5 depicts the results of using different cost functions and the proposed hierarchical filter-based algorithm (Algorithm 1). Since the heuristics used in building a vertical tower is different from that of a wall, we modify the filters to fit the task. The filter contains the contact polygon area C, distance $\mathrm{r}_{P_j P_i}$, and the top surface slope. We also evaluate each cost component used in Eq. 3 separately. We can see that the proposed hierarchical filter algorithm has a higher opportunity of building a vertical tower with more than 5 stones compared to all other methods. Table 1 gives the average number of stones each algorithm can build. It also shows that the proposed method can build more stones than other methods. We randomly select 9 towers planned by the proposed method from all of the towers that have a height of at least 6 stones for physical execution. Table 2 shows the building process. Since the first three stones can always be placed successfully, we start the table from the 4th stone. We can see that 4/9 can be built up to 6 stones, and only 1/9 drops at the 4th stone. Compared to the towers built in work [6], which only builds up to 4 stones with a chance of 2/11, our method can build taller stone towers both in simulation and in practice. The reasons for the failures in our execution could be object pose detection error, pickup error, an opening gripper moves an object which is already placed, differences between the object 3D mesh and the real stone, inaccuracies in contact modeling, etc. Figure 6a, b illustrate one tower example of the proposed algorithm in the simulation environment and in practice.

Table 1 Comparison on the average tower height built using different filters

Random	Area	Kinetic energy	Distance	Deviation (m)	Deviation (d)	Weighted cost	Proposed
3.4565	4.3944	4.0845	3.9362	3.2	4.0207	4.75	5.4118

Table 2 Physical execution (d represents stone drops, ✓ represents the successful placements, Ø means that the plan does not contain further steps.)

Stone number	Tower 1	Tower 2	Tower 3	Tower 4	Tower 5	Tower 6	Tower 7	Tower 8	Tower 9
4th	✓	✓	d	✓	✓	✓	✓	✓	✓
5th	✓	✓		d	d	✓	✓	✓	✓
6th	✓	✓				✓	d	✓	d
7th	d	d				Ø		d	

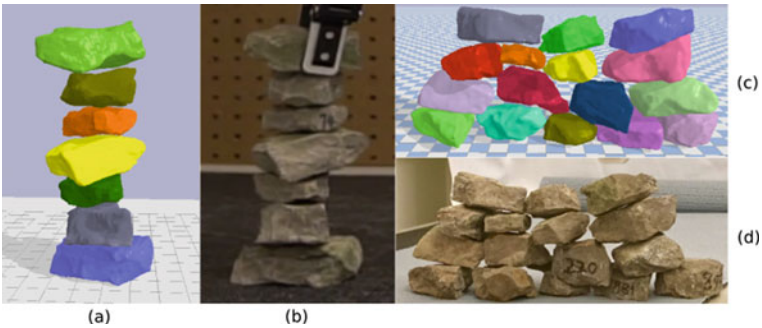

Fig. 6 Vertical tower and wall in simulation and corresponding physical execution results

4.2.2 Stone Wall

In this experiment, the goal is to build a stone wall. The structure is planned in simulation using Algorithm 2, and then the UR5 manipulator places the stones to the planned pose. The planned wall has 4 courses, and each course has 3–5 stones. The execution order is manually calculated, but complies with the assembly order in the simulation if there is no collision during assembly due to the gripper. In other words, the simulated assembly order does not take into account clearances for the fingers or grippers, but if problems exist, reordering stones within a course often fixes potential collisions. As mentioned in the previous section, several things may cause failures during the execution process. We categorize the failure cases into two classes: poor placement and structure collapse. Poor placement contains bad grasping, wrong object pose detection, and stone drops after placement. Structure collapse is the case that after placing the current stone, more than 1 stone falls down. In this experiment, 7 out of 13 walls can be successfully built without collapsing. Table 3 shows the failure

Table 3 Stone wall execution failure rate

	Course 1	Course 2	Course 3	Course 4
Poor placement	0.14	0.18	0.29	0.48
Structure collapse	0	0.04	0.07	0.05

rate during execution. We separate the execution process based on different courses. We can see that as the course increases, the poor placement rate also increases. All the previous minor errors building up to larger errors lead to more failure cases. Figure 6c, d show an example of a planned wall in the simulation and the wall built in the real world using the robotic arm without collapsing. We also compare the proposed algorithm with untrained humans using the same set of stones. We found that humans are much better at executing the generated open-loop plans. However, if the assembly sequences are not provided, our human participants could not plan several steps ahead and retorted to trial-and-error.

5 Conclusion and Future Work

The proposed method is able to plan placements for a set of irregularly shaped rocks and build stable dry-stacked structures. Similar to previous work, we use a rigid body simulation engine in order to find stable poses for rocks. We introduce two primary innovations: first, we use only the physics engine to create a finite set of feasible poses; second, we use a layered refinement architecture that significantly improves performance compared to optimized scalar cost functions in evaluating the quality of feasible poses. We also introduce new filtering terms, which are specific for building walls with interlocking layers, compared to vertical stacks.

We focus on high-level placement planning as it is a central issue in dry stacking. The overall system could be significantly improved with a more specialized and robust execution system, specifically reactively re-planning in the face of errors and unmodeled action outcomes, and in incorporating tactile feedback during placement and pose evaluation. The instructional literature suggests this approach for human builders as well: candidate stones are placed, wiggled, and then either removed or stabilized by wedging small rocks into crevices until the newly placed rock is stable. In any of these situations, having better high-level placement plans that can be executed in an open-loop fashion will be beneficial and this paper represents significant progress in that direction.

Acknowledgements We would like to thank Hironori Yoshida and Dr. Marco Hutter for their valuable input regarding the comparison algorithm, and Jackie Chan for scanning the stones. This work was partially supported by NSF Grant #1846340 and the SMART CoE at UB.

References

1. USEPA: Advancing sustainable materials management: 2014 fact sheet. United States Environmental Protection Agency, Office of Land and Emergency Management, Washington, DC 20460, 22 Nov 2016
2. Green in practice 102—concrete, cement, and CO2
3. Sanders, G.B., Larson, W.E.: Progress made in lunar in situ resource utilization under NASA's exploration technology and development program. J. Aerosp. Eng. **26**(1), 5–17 (2013)
4. Wikipedia Contributors: History of construction—Wikipedia, the free encyclopedia (2019)
5. Chen, I.-M., Burdick, J.W.: Determining task optimal modular robot assembly configurations. In: Proceedings of 1995 IEEE International Conference on Robotics and Automation, vol. 1, pp. 132–137. IEEE (1995)
6. Furrer, F., Wermelinger, M., Yoshida, H., Gramazio, F., Kohler, M., Siegwart, R., Hutter, M.: Autonomous robotic stone stacking with online next best object target pose planning. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 2350–2356. IEEE (2017)
7. Helm, V., Ercan, S., Gramazio, F., Kohler, M.: Mobile robotic fabrication on construction sites: Dimrob. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 4335–4341. IEEE (2012)
8. Thangavelu, V., Liu, Y., Saboia, M., Napp, N.: Dry stacking for automated construction with irregular objects. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 1–9. IEEE (2018)
9. Napp, N., Nagpal, R.: Distributed amorphous ramp construction in unstructured environments. In: Distributed Autonomous Robotic Systems, pp. 105–119. Springer (2014)
10. Petersen, K.H., Napp, N., Stuart-Smith, R., Rus, D., Kovac, M.: A review of collective robotic construction. Sci. Robot. **4**(28), eaau8479 (2019)
11. Petersen, K., Nagpal, R., Werfel, J.: An autonomous robotic system for three-dimensional collective construction. Termes, June 2011
12. Lindsey, Q., Mellinger, D., Kumar, V.: Construction with quadrotor teams. Auton Robot **33**(3), 323–336 (2012)
13. Doerfler, K., Sebastian Ernst, Luka Piškorec, Jan Willmann, Volker Helm, Fabio Gramazio, and Matthias Kohler. Remote material deposition. In *International Conference, COAC, ETSAB, ETSAV*, pages 101–107, 2014
14. Saboia, M., Thangavelu, V., Gosrich, W., Napp, N.: Autonomous adaptive modification of unstructured environments. In: Proceedings of Robotics: Science and Systems, vol. XIV (2018)
15. Soleymani, T., Trianni, V., Bonani, M., Mondada, F., Dorigo, M.: Autonomous construction with compliant building material. In: Intelligent Autonomous Systems, vol. 13, pp. 1371–1388. Springer (2016)
16. Saboia, M., Thangavelu, V., Napp, N.: Autonomous multi-material construction with a heterogeneous robot team. In: Distributed Autonomous Robotic Systems. Springer (2018)
17. Liu, Y., Saboia, M., Thangavelu, V., Napp, N.: Approximate stability analysis for drystacked structures. In: 2019 IEEE International Conference on Robotics and Automation (ICRA). IEEE (2019)
18. Liu, Y., Shamsi, S.M., Fang, L., Chen, C., Napp, N.: Deep q-learning for dry stacking irregular objects. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1569–1576. IEEE (2018)
19. Gardner, K.: Stone Building. The Countryman Press (2017)
20. McRaven, C.: Building Stone Walls, vol. 217. Storey Publishing (1999)
21. Vivian, J.: Building Stone Walls. Storey Publishing (1976)
22. Rusu, R.B., Cousins, S.: 3D is here: point cloud library (PCL). In: IEEE International Conference on Robotics and Automation (ICRA) (2011)
23. Zhou, Q.-Y., Park, J., Koltun, V.: Open3D: a modern library for 3D data processing (2018). arXiv:1801.09847

24. Sucan, I.A., Chitta, S.: Moveit! (2013). http://moveit.ros.org
25. Coumans, E., Bai, Y., Hsu, J.: Pybullet physics engine (2018)
26. GPyOpt: Gpyopt: a bayesian optimization framework in python (2016). http://github.com/SheffieldML/GPyOpt