SCIENCE ROBOTICS | RESEARCH ARTICLE

MANUFACTURING

A framework for robotic excavation and dry stone construction using on-site materials

Ryan Luke Johns[1,2]*, **Martin Wermelinger**[1], **Ruben Mascaro**[3,4], **Dominic Jud**[1], **Ilmar Hurkxkens**[2], **Lauren Vasey**[2], **Margarita Chli**[3,4], **Fabio Gramazio**[2], **Matthias Kohler**[2], **Marco Hutter**[1]

Automated building processes that enable efficient in situ resource utilization can facilitate construction in remote locations while simultaneously offering a carbon-reducing alternative to commonplace building practices. Toward these ends, we present a robotic construction pipeline that is capable of planning and building freeform stone walls and landscapes from highly heterogeneous local materials using a robotic excavator equipped with a shovel and gripper. Our system learns from real and simulated data to facilitate the online detection and segmentation of stone instances in spatial maps, enabling robotic grasping and textured 3D scanning of individual stones and rubble elements. Given a limited inventory of these digitized stones, our geometric planning algorithm uses a combination of constrained registration and signed-distance-field classification to determine how these should be positioned toward the formation of stable and explicitly shaped structures. We present a holistic approach for the robotic manipulation of complex objects toward dry stone construction and use the same hardware and mapping to facilitate autonomous terrain-shaping on a single construction site. Our process is demonstrated with the construction of a freestanding stone wall (10 meters by 1.7 meters by 4 meters) and a permanent retaining wall (65.5 meters by 1.8 meters by 6 meters) that is integrated with robotically contoured terraces (665 square meters). The work illustrates the potential of autonomous heavy construction vehicles to build adaptively with highly irregular, abundant, and sustainable materials that require little to no transportation and preprocessing.

INTRODUCTION

The construction industry is of central importance to human productivity but will require substantive innovation to meet growing demands and to reduce its high environmental impact. Construction is among the world's largest economic sectors, representing 13% of global gross domestic product (*1*). By extension, it is a vital component in virtually every other industry: Construction is responsible for the production of transportation and agricultural infrastructure, manufacturing facilities, dams, power plants, and residences. Global deficiencies in housing supply in the face of population growth and inadequate civil infrastructure (*2*, *3*) are among the causes for an expected surge in demand for building, a demand that is unlikely to be met without substantial technological advances: Labor productivity growth in the building sector has averaged only 1% per year in the past two decades, in contrast to 2.8% across the global economy and 3.6% in manufacturing (*1*). The coming decades will also bring new challenges and opportunities to the industry, such as the building of coastal defenses to combat the effects of climate change (*4*) and the need for efficient extraterrestrial infrastructure to support space exploration (*5*). Among the most critical of these industry challenges is the immediate imperative to reduce the industry's own ecological footprint: Construction is a major contributor to global greenhouse emissions, resulting primarily from the manufacturing of building materials such as cement, steel, plastics, aluminum, brick, and glass (*6*). Cement (used for concrete, blocks, mortar, and plaster) alone accounts for 8 to 9% of global anthropogenic CO_2 emissions and 2 to 3% of global energy (*6*). At the end of its life, concrete is also a major contributor to construction and demolition waste (CDW): Global CDW exceeds 3 billion tons annually (*7*), and concrete is 70% of the United States' contribution to that figure (*8*).

Robotics offers an opportunity to increase labor productivity in construction while also mitigating the concomitant concerns of safety and fatigue in a notoriously dangerous, unstructured, and physically taxing workplace environment (*9*). Beyond improvements to economic and production efficiency in building, robots also provide an opportunity to reduce the industry's environmental footprint by enabling the direct use of local, low-energy, natural, and recycled materials as alternatives to concrete. In situ resource utilization (*10*) has long been proposed as a necessity for facilitating extraplanetary construction: In space, resources are scarce, transportation is expensive and energy intensive, and there is an extreme dearth of skilled tradespeople; thus, it is prudent to develop automated systems that can make use of the resources at hand. Confronted with the climate crisis and a stagnating construction industry, these are also terrestrially resonant problems, and the approach of minimizing energy-intensive manufacturing and transportation by building intelligently with hyperlocal materials can also be applied here on Earth. Building with regionally sourced and minimally processed materials such as stone and soil has been demonstrated to reduce the carbon footprint of built structures by upward of 50% (*11*), but these methods are difficult to adopt in developed areas where mass-manufactured building supplies are more economical.

Historically, building with indigenous materials was commonplace: Stone masonry, for example, has been used in construction for millennia but fell into decline with the advent of inexpensive canal and railway transportation, rising labor costs, and industrially manufactured alternatives that required less-skillful processing on

---

[1]Robotic Systems Lab, ETH Zurich, Zurich, Switzerland. [2]Gramazio Kohler Research, ETH Zurich, Zurich, Switzerland. [3]Vision for Robotics Lab, ETH Zurich, Zurich, Switzerland. [4]Vision for Robotics Lab, University of Cyprus, Nicosia, Cyprus.
*Corresponding author. Email: rjohns@ethz.ch

the construction site (*12*, *13*). As stone's popularity as a building material diminished, so too did the expertise and regional quarries that made the building method widely accessible and affordable (*12*, *14*). By equipping robots with the capacity to work with irregular local materials in situ, such sustainable historic building practices can be reimagined as economically viable alternatives in the context of contemporary construction.

Toward these ends, we present a 12-ton robotic excavation platform capable of building large-scale stone walls and landscapes from locally available, unmodified natural and reclaimed materials (Fig. 1). Our comprehensive pipeline allows for the construction of arbitrarily curved and explicitly defined wall shapes from highly heterogeneous stones, using light detection and ranging (LiDAR) mapping, learned image segmentation, grasp heuristics, and classification to isolate, digitize, and manipulate stone instances in unstructured environments. We present an online planning algorithm that determines a limited-horizon building sequence, leveraging geometric attributes to quickly seed candidate stone poses that are further aligned using a torque-and-penetration-constrained variant of iterative-closest-point (ICP) registration. The resulting candidate placements were classified using a three-channel signed distance field (SDF) representation that encapsulates the local context and desired-wall-geometry around each unique solution, and the selected stone placements were executed with a grasp-and-trajectory plan that considers collision constraints at both ends of the pick-and-place sequence. The presented construction and landscaping solutions were executed in a closed loop, where the accumulated LiDAR map was used for ICP refinement of the as-built structure in stone construction or to measure displacement between desired and actual terrain in excavation. This single platform provides the ability to work with a broad array of in situ granular materials, ranging in scale from clay mud (intermediate axial length $d_I < 0.004$) to coarse boulders ($d_I \approx 2$ m) (*15*). Our work facilitates the full-scale robotic construction of permanent and performative dry stone walls, leveraging multimodal tool use (of a gripper and shovel) to complete both infrastructure-scale landscaping and building tasks on a single job site (Movie 1).

Robotic construction with raw and reclaimed materials

The automation of construction tasks gave rise to some of the earliest robotic systems (*16*), and the first full-scale single-task construction robots (STCRs) were introduced in Japan in the late 1970s (*17*). Whereas STCRs were limited to highly specific and repetitive tasks, the early 21st century brought the application of industrial robots to formally complex, nonstandard, and mass-customized architectural assemblies (*18*, *19*). Early work in this domain explored the aesthetic potential of robotic tools to materialize digital complexity, for example, by creating varied arrangements of bricks (*20*), dimensional lumber (*21*), or custom components (*22*), but remained dependent on standardized, industrially manufactured building materials.

With improved contextual awareness through vision and depth sensing, architectural robotics has expanded in the past decade to include construction from materials with natural geometric heterogeneity. Beyond manufacturing artificial variation from digital models, these processes adapt to existing variations in the environment to facilitate fabrication from readily available and sustainable materials. This approach has been applied to timber structures, where subassembly details or the overall design adapt to the "material morphospace" (*23*) of available, naturally nonstandard trunks, branches, and tree forks (*24*–*27*). Others have leveraged three-dimensional (3D) scanning to reduce waste in manually constructed walls built from digitally machined stones (*28*, *29*). Such applications have demonstrated the potential of nonstandard materials in

Fig. 1. HEAP is an autonomous excavator for construction with in situ resources. HEAP is a modified 12-ton walking excavator equipped with multiple cabin- and arm-mounted LiDAR sensors and cameras, facilitating the construction of the pictured 65-m-long retaining wall and terraced landscape from unmodified local materials. (**A**) Arm-mounted Livox Mid-70 sensor, (**B**) rooftop perception box with multiple LiDAR sensors and camera, and (**C**) lower-cabin camera.

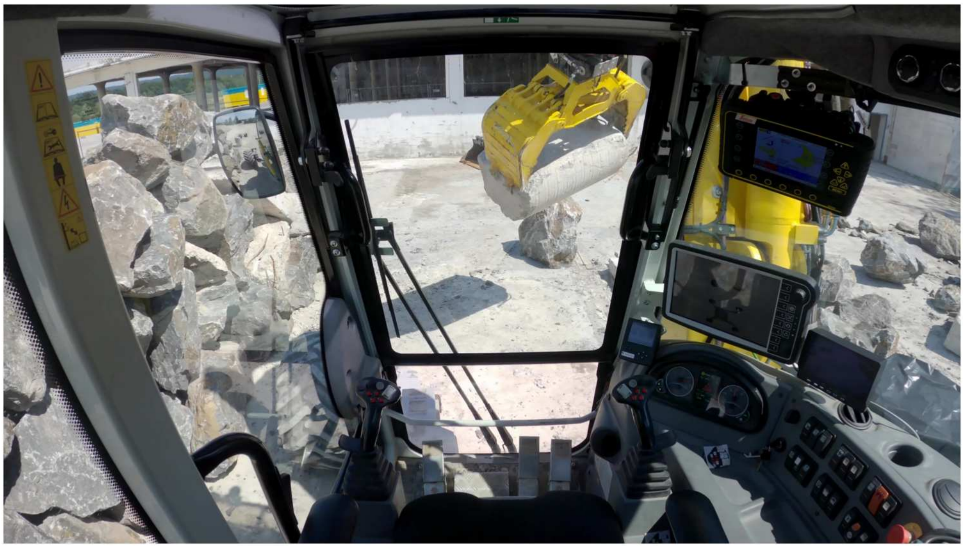

Movie 1. A framework for robotic excavation and dry stone construction using on-site materials.

construction but typically require substantial intervention in the loop: Scanning is executed in a separate stage, and with the exception of Wu and Kilian's work (*27*), the layout is fixed after the initial optimization for piece-by-piece machining and assembly. Our process instead digitizes objects automatically during the construction process and updates the building plan continually, thus allowing for reduced error accumulation by accounting for settling and ensuring robustness if a single building element is damaged or goes missing.

Because the machining of stone introduces additional process complexity and can substantially increase the carbon footprint of masonry construction (*30*), there is motivation to develop methods that can autonomously plan structures using entirely unmodified stone. Previous work in this domain has used simulated annealing to search for stone poses in a 2D scene (*31*) or a combination of physics simulation and heuristics to find stone poses for AR-assisted (*32*–*34*) or robotic (*35*) assembly. Furrer *et al.* (*36*) used gradient descent and a support polygon–based cost function to create stacked towers of up to four stones, whereas Liu *et al.* (*35*) achieved stacks up to six stones, and single-layer walls up to four courses, with a combination of hierarchical heuristics adapted from dry stone masonry guidebooks: Both studies built physical structures at the desktop scale in an open loop using prescanned stones that were localized for grasping with a manipulator-mounted RGB-D sensor. Deep Q-learning has been applied for learning stone assembly policies from physics simulations in 2D scenes (*37*) and as a model-free method that uses 2.5D representations of the scene-top and stone-underside to position stones with fixed orientations toward the construction of double-faced walls in simulation (*38*). One limitation of the latter approach is the lack of consideration of the top surface of the candidate stones: Stones with outward-sloping top surfaces jeopardize the stability of subsequent stones (*39*, *40*), severely limiting the heights that can be achieved with the method.

Our research improves on these studies by targeting the in situ construction of arbitrarily shaped double-faced dry stone walls built from large-scale boulders and concrete demolition debris (we use "stones" henceforth to describe both of these elements interchangeably). Building on our earlier work in geometric planning, stone segmentation, and grasping (*41*–*43*), we present a synthesized system capable of autonomously detecting, grasping, scanning, and reorienting stones in the wild, allowing our geometric planner to optimize for stable and geometrically registered solutions that are unconstrained in orientation. We constructed permanent robotic stone walls, reaching heights up to 6 m, and used these demonstrators to perform an environmental impact assessment for the building method. We also present the preliminary use of an excavator-mounted force-torque sensor to evaluate structural stability at the end of construction, building on small-scale precedents in robotic structural testing (*44*, *45*).

Mobile robotic platforms for autonomous construction

To facilitate adaptive construction tasks, robotic systems must be highly mobile and capable of navigating in dynamic, unstructured environments. Small-scale robots have demonstrated the use of available resources, such as blocks (*46*) or compliant, "amorphous" materials (*47*), to modify their environment for increased navigability (*48*, *49*). A variety of mobile-based manipulators have been developed for in situ additive manufacturing with urethane foams (*50*), clay (*51*), concrete (*52*–*54*), and plaster (*55*) and for the assembly of brick and block walls (*56*–*59*). Toward the manipulation of irregular materials such as stone, specialized hardware has been developed for assistive rock clearing in agriculture (*60*) and teleoperated leveling and assembly of rubble breakwaters (*61*, *62*).

For heavy construction tasks, a range of tracked and wheeled excavators have been modified for autonomous and semi-autonomous operation in the past 25 years (*9*, *63*–*65*). Our work makes use of hydraulic excavator for an autonomous purpose (HEAP) (*66*), a customized robotic excavator previously deployed for trenching (*67*), tree harvesting (*68*), and teleoperation (*69*). A brief system overview is provided in section S2. For an overview of existing autonomous construction platforms, we point the reader to reviews covering topics of STCRs (*17*), multirobot systems (*70*), and on-site construction robots (*71*).

RESULTS

Robotic stone wall assembly

The autonomous construction of stone walls required locating and digitizing discrete building materials and a combination of assembly, grasp, and motion planning to aggregate those toward some globally defined form. To accomplish this, we first mapped and segmented the environment to discover stone instances and then used the segmented point cloud representations of those instances for grasp planning. Stones were individually picked up, laser-scanned in the excavator's gripper, and returned to the ground for storage. In this manner, we built a limited inventory until it contained 20 to 40 stones. Those were used by the geometric planner, which found a feasible subset that could be placed in stable and well-aligned poses on the current wall. Each stone was then regrasped for placement in the wall, with possible intermediate reorientations as necessary. Once placed, the stone pose was refined by laser scanning and updated in the geometric planner. This entire process was repeated each time after all scanned stones were placed or after the planner was no longer able to find solutions with the available inventory. An overview of the dry stone assembly process is provided in Fig. 2, and a detailed description is given in Materials and Methods. We demonstrated the process with two full-scale walls built from more than 1000 elements with an average mass of 997 ± 438 kg.

Freestanding wall

The freestanding wall was robotically constructed from 109 unique elements, with a mixture of waste concrete (24%) and gneiss boulders. The structure measured 10 m by 1.7 m by 4 m, with a tapering thickness from a symmetric batter or rake angle of 3° from the vertical plane (where 1.7 m was the mean); this is notably steeper than the 5.7° to 11.3° rake angle range typical of dry stone masonry walls (Fig. 3, A and B) (*40*). Constructed during a 2-week period, the structure was built at a rate of 17.4 min per stone placement, with direct placements taking an average of 11.8 min and placements requiring intermediate reorientations taking 21.8 min each, inclusive of regrasp attempts. Upon refining the as-placed stone poses with HEAP's laser map, we evaluated the median positional error of each settled stone with respect to the planned pose as 0.108 m. The mean pose error $\mathscr{E}_{\text{pose}}(\mathbf{O}_{planned}, \mathbf{O}_{placed})$ was (0.128 ± 0.088, 0.218 ± 0.172), where $\mathscr{E}_{pose}$ is defined in Materials and Methods (Eq. 12).

Retaining wall

The largest structure built using our method is a permanent, 65.5-m-long retaining wall consisting of 938 unique elements, with a mixture of reclaimed concrete (7%), mixed erratics (3%), and gneiss boulders (90%; Fig. 3, C to F). Designed with variable height and incline (section S3), it has a mean thickness of 1.8 m and an apex height of 6 m. A conjoined parapet (freestanding portion) continues 1.5 m above the retained earth, with a fixed rake angle of 11.3°. The structure was built at a rate of 16.25 ± 6.88 min per stone placement, inclusive of grasping, driving, and refinement, measured from the time between subsequent placements for 300 stones over 21 nonconsecutive working days. We refined the stone poses of the final structure with a high-density point cloud produced by a surveying laser scanner (Leica RTC360) to assess the median positional error between planned and placed poses as 0.115 m and the mean pose error as (0.135 ± 0.089, 0.089 ± 0.111).

Robotic landscaping

HEAP (*66*) is additionally capable of autonomously excavating freeform embankments with high precision (*72*), given a designed target geometry defined as a 2.5D height map. Our design process generally begins with a map captured by surveying equipment or by aerial or excavator-mounted LiDAR. Incorporating considerations like building codes, slope stability, and drainage, the design model can be generated by signed distance functions that operate directly on the height map, using operators based on input points, paths, and areas with parametric connections to specified locations, dimensions, slopes, and volumes (*73*). The excavation planner is initialized using this designed height map, together with a design line that is aligned with the general shape of the embankment, tangential to the digging direction (Fig. 4). In our excavation work, this height map was used to determine the starting location for each digging or dumping cycle, where the planner was run in an iterative fashion until the entire site was excavated to a certain error threshold (receiving an updated environmental LiDAR map for every new planning step). Rather than being position or velocity defined, the digging strategy of HEAP uses a force-torque defined digging trajectory to adapt to different soil types and interaction forces, allowing us to circumvent the inherently difficult task of estimating or predicting soil properties (*67*). Thus, the shovel could dig freely in fine-grained materials (sand, silt, and clay) while adapting to densely compacted or frozen regions, pebbles, and intermittent fine cobbles. Because the size of the landscape was much larger than the workspace of the machine, we facilitated automated repositioning with a driving path that was offset from the design line, where the ultimate position of the machine was found by projecting the excavation or dump points orthogonally onto the driving path. Because the design contained horizontal, hard edges (from steps), the bucket was implicitly aligned to those edges, thus conveniently enabling excavation with the necessary precision. For motions through the air, such as approaching the excavation or dump points, an SDF-informed trajectory planner was used to avoid unwanted collisions with the ground.

Contoured landscape

The robotically constructed retaining wall was built in the context of a public park (Fig. 5A), with the adjacent landscape excavated by HEAP to form four contoured terraces that facilitate water management, erosion control, and pedestrian circulation: The recreational terraces accommodated height disparities on site, providing two 6% access ramps that run along the wall. On each slope, two intermediate berms (0 to 1.2 m) decelerated water runoff with a 10% backsloping geometry, preventing excessive erosion and collecting fine sediment to provide space and humidity for spontaneous plant growth. A rock trench ran along the bottom of each slope, collecting water for on-site storage.

About 1 to 1.5 m of depth was autonomously excavated to produce the precisely rippled terrace surfaces (Fig. 5, B to D). For the largest terraces where the reach of HEAP proved inadequate, the embankment was split into two levels: where the upper part was first excavated with the machine partially driving on the slope using the actively controlled chassis to maintain a level cabin. This process used the excavation workflow developed in our previous work

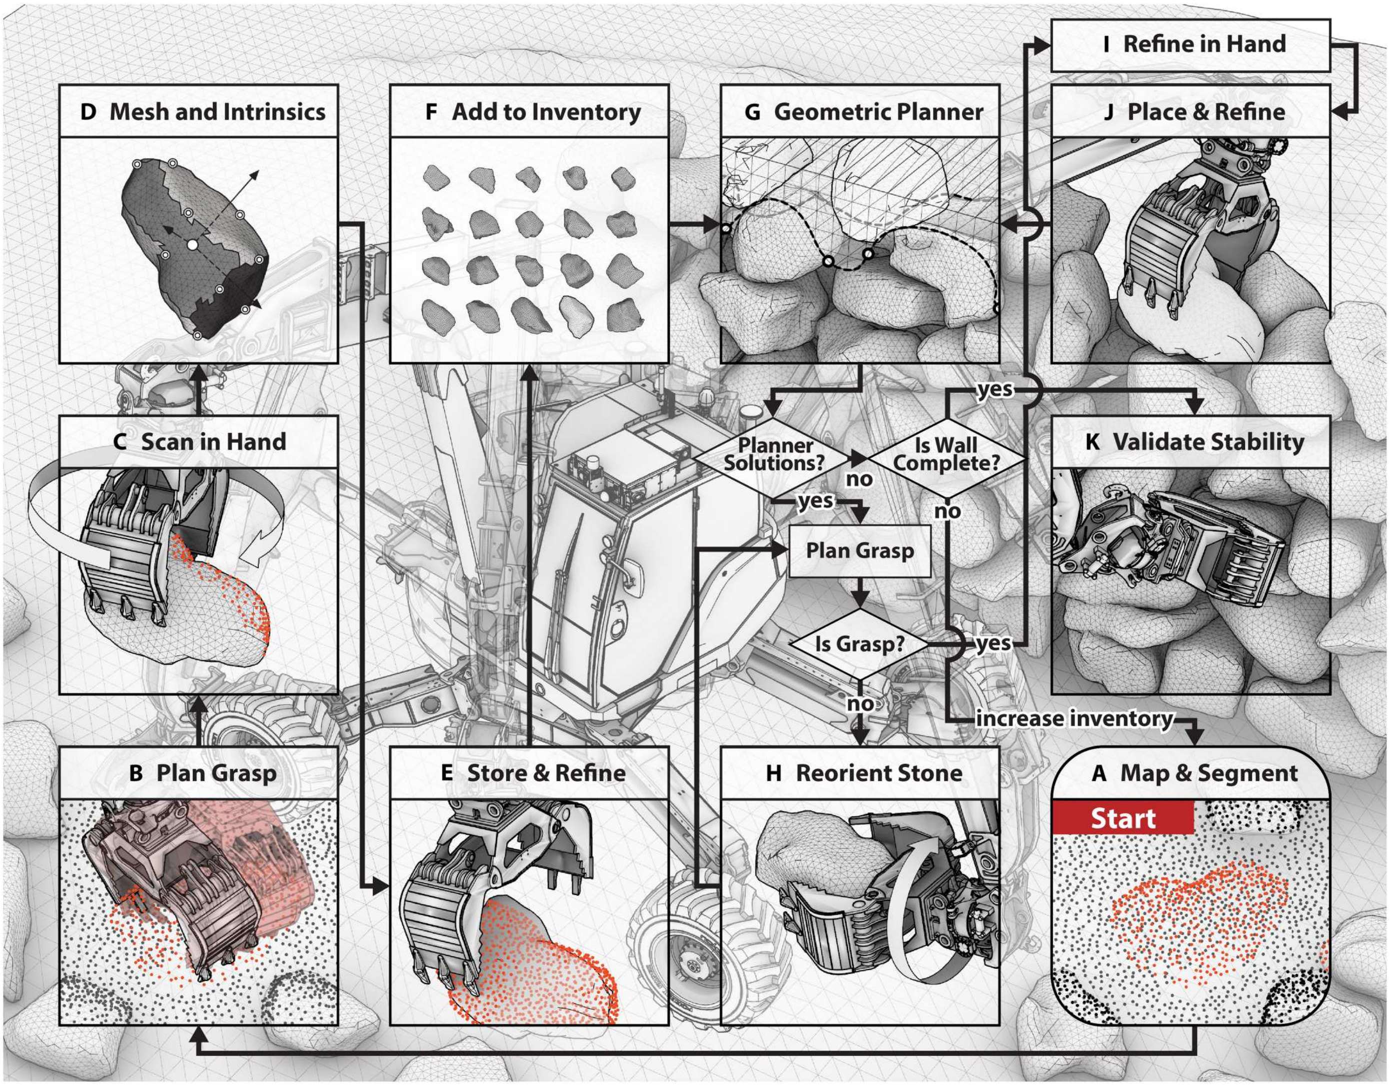

Fig. 2. Autonomous dry stone construction sequence. The construction process began with stones in the environment and a point cloud map produced using the cabin- and boom-mounted LiDAR sensors. Individual stone instances were segmented from this map (**A**), and an initial grasp was planned for each of these segments (**B**). Once grasped, each stone was fully scanned in hand (**C**), and the intrinsic properties of the stone were computed from the resultant mesh (**D**). The stone was then returned to the ground for storage, with its pose refined by LiDAR map upon release (**E**). An inventory (**F**) was accumulated in this manner, which was then used by the geometric planner (**G**) to compute poses within the target wall. Each solution was grasped directly for placement or with an intermediate reorientation (**H**) when necessary. The poses of held stones were refined in hand (**I**) using the cabin-mounted LiDAR, placed in the wall, and refined again (**J**). Stability probing (**K**) could be done either at the end of construction or in batches after each solution set was placed.

(*72*), with the exceptions being that the initial survey was completed by a tripod-mounted surveying laser scanner (Leica RTC360) rather than by drone and that the construction plan was entirely subtractive (it did not require the additional deposition of soil to build up the terrain). Thus, the excavated earth could be simply dumped to the left side of the machine for later removal (fig. S1). The four terraces have a combined area of 665 m^2 and were excavated autonomously over 8 nonconsecutive days, with an average active soil removal rate of 4.3 min/m^3. The finished surface has an average grade of 38.5 ± 18.6%, ranging from 2 to 65%. Taking the distances between the surveyed point cloud of the final excavation and the input design mesh, we observed an absolute average distance error of 0.023 ± 0.020 m and a root mean square error (RMSE) of 0.030 m.

DISCUSSION

The robotic construction processes developed in this work facilitate large-scale material-aware landscaping and construction tasks with high geometric complexity. Our assembly pipeline combines capabilities in stone segmentation and manipulation (Fig. 6) with online 3D scanning (Fig. 7) to accumulate an inventory of unique stone models that can be used for geometric planning (Fig. 8). The synthesis of these elements allowed for the completion of physical structures, which could be used to understand both the efficacy of

Fig. 3. Robotically constructed stone walls. (**A** and **B**) A 10 m–by–1.7 m–by–4 m freestanding wall built with 109 automatically digitized stones and reclaimed concrete elements. (**C** to **F**) A 65.5 m–by–1.8 m–by–6 m retaining wall built to comply with global geometric constraints. (C) Front elevation of retaining wall with partial overlay of mesh models from digital twin. (D and E) Overview of the built structure with variable height and incline. (F) Iteratively refined 3D digital twin model displayed in the geometric planning software.

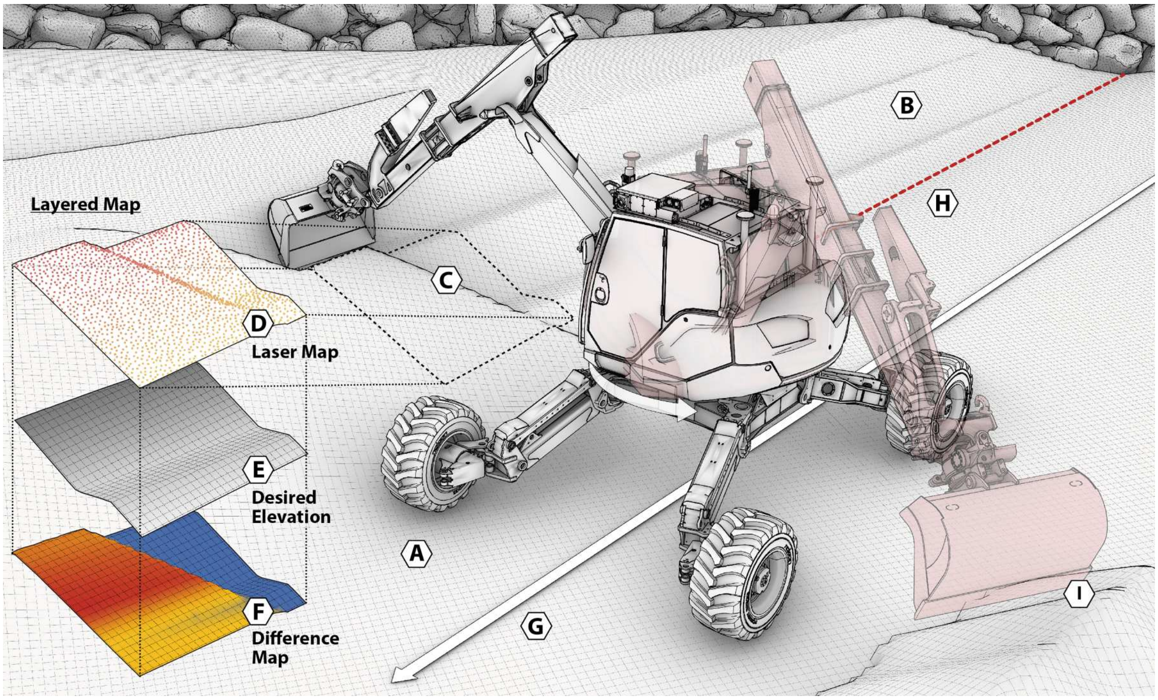

Fig. 4. Autonomous excavation procedure. HEAP moved gradually along the terrain, bringing the untouched landscape (**A**) toward the target geometry (**B**) with angled cuts (**C**) that reduced the effects of actively dug soil disturbing the already-surfaced areas. The machine perceived the environment by LiDAR and assembled laser (**D**) and elevation maps from the data. By subtracting the designed desired elevation map (**E**) from the laser map, a difference map (**F**) was constructed. This difference map was used to determine the next excavation point, to which the machine approached along the driving path (**G**), to work orthogonally on the embankment. The driving path was offset from the design line (**H**), which was provided along with the desired elevation map. After filling the bucket in an excavation cycle, the material was dumped to the side by rotating the cabin (**I**), and the process was repeated.

the proposed building method and the potential to improve and expand it toward new applications in the future.

Construction speed

Combining the average duration of the scanning and placing routines, our process currently requires 21 min per stone placement. Given the size of our stones, this equates to 1 m^3 or 2.8 t of placed material per hour. Considering the mean wall thickness and void ratio of our demonstrators, this is equivalent to 1.1 m^2 of finished wall surface per hour. Excluding the preparatory stone scanning, we have demonstrated an average placement rate up to 12.2 min per stone for 28 consecutively placed stones in a single day. For comparison, our field observations revealed experienced excavator operators building single-layer walls against earth with an average rate of 11 min per stone placement: Hence, we are currently outperformed by human excavator operators in placement speed. Such operators, however, typically require string and paint references with which to register their construction and often a second or third person outside the machine to provide guidance and to insert small supporting stones, gravel, and soil by hand and shovel. In contrast, our process can build complex nonplanar global surface geometries without physical reference markers, does not require a skilled driver or small supporting stones, and provides a full digital twin of the built structure for better accountability and future reuse. Because of the large scale of the materials, our robotic process is notably faster than the manual per-person building rate of 2 t/day for a trained stonemason's construction of two-layer cobble walls (*40*). However, these dextrous workers will likely have a higher daily output in quantity of stones, amounting to more than 87 individual stones, if we assume that each is within the 23-kg upper-lifting limit prescribed by the National Institute for Occupational Safety and Health (NIOSH) (*74*).

Multiscalar construction

The constructed freestanding and retaining walls exhibit 42.2 and 40.2% void ratios, respectively, considering stone that occupies the enclosed mesh volume of the designed wall (excluding any stone mass that protrudes beyond the wall surface). These void ratios are near the coarse end of the 20 to 40% spectrum considered typical for manually constructed dry stone walls (*40*, *75*). Such traditional walls, however, generally include stones that have been reshaped for an improved fit and contain a multitude of much smaller "hearting" stones that are carefully laid to fill the core of the wall and "wedging" stones that fill and support the underside of larger stones (*40*); together, this stone debris accounts for 5 to 20% of the wall volume (*76*). Our current process does not allow for modification of the stones, nor does the large-scale gripper of HEAP facilitate construction with smaller cobble-scale stones (*15*). Although the current level of porosity is unproblematic for a variety of utility structures and can even be desirable as supplemental habitat for local flora and fauna (*77*) or for water flow in breakwaters, reducing this void ratio can increase a stone gravity wall's resistance to incident forces (*40*) and reduce the permeability of structures intended to enclose indoor spaces. Potential approaches to filling the gaps

Fig. 5. Robotic landscaping. (**A**) Overview of park with integrated robotic terracing and stone retaining wall. (**B**) HEAP positioning to modify the lower slope with a grading bucket. (**C**) Grading bucket exit path while cutting the middle slope. (**D**) Intermediate berms on the sloped landscape for erosion control, displaying first signs of ruderal plant growth.

between stones could include mortar printing (*52*), rammed-earth or clay construction (*78*), and jammed tensile-reinforced gravel printing (*79*) or, in an extraterrestrial context, regolith sintering (*80*) or binder-supplemented microscale assembly (*81*). Toward the more traditional insertion of small filling stones, we consider the application of a compound manipulator for macro-micro assembly (fig. S2), with the understanding that parallelization with multiple machines would be necessary for practical implementation (planning and placing a small stone requires a similar execution time as placing a single large boulder).

High-level planning

The present workflow requires the global target geometry of stone walls and landscape features to be manually specified by a designer. This designer must, for example, correctly size the goal wall to meet structural requirements or other site-specific demands. Although this preemptive planning is practical for the majority of scenarios that must integrate with other preplanned building phases (and where material is being delivered on the fly), it is undesirable in remote or relief scenarios where materials are limited to the direct vicinity of the construction site and where the utility of the built structure is more important than its specific shape. In such applications, high-level planning routines should be capable of surveying the available resources and automatically modifying the overall construction plan accordingly to suit the size distribution and quantity of stones or of adjusting the target terrain to accommodate local features while steering toward high-level goals, such as navigability (*82*, *83*).

Autonomy in construction

Construction sites are highly dynamic and safety-critical environments, providing numerous challenges to the future of full construction autonomy. The freestanding wall prototype was built on a planar concrete surface in a facility protected from outside disturbances, and construction could be completed in that context without a human operator in the machine. The retaining wall and terraced landscape, however, were built as permanent installations in an active construction site: Additional manually operated machines and human workers frequently required access to the area for activities such as material handling and storage, earthmoving, and the routing of drainage or electrical services. Weather conditions and excavation works also introduced mud and erosion challenges to the process. For these reasons, the retaining wall was constructed predominantly with a human operator inside the

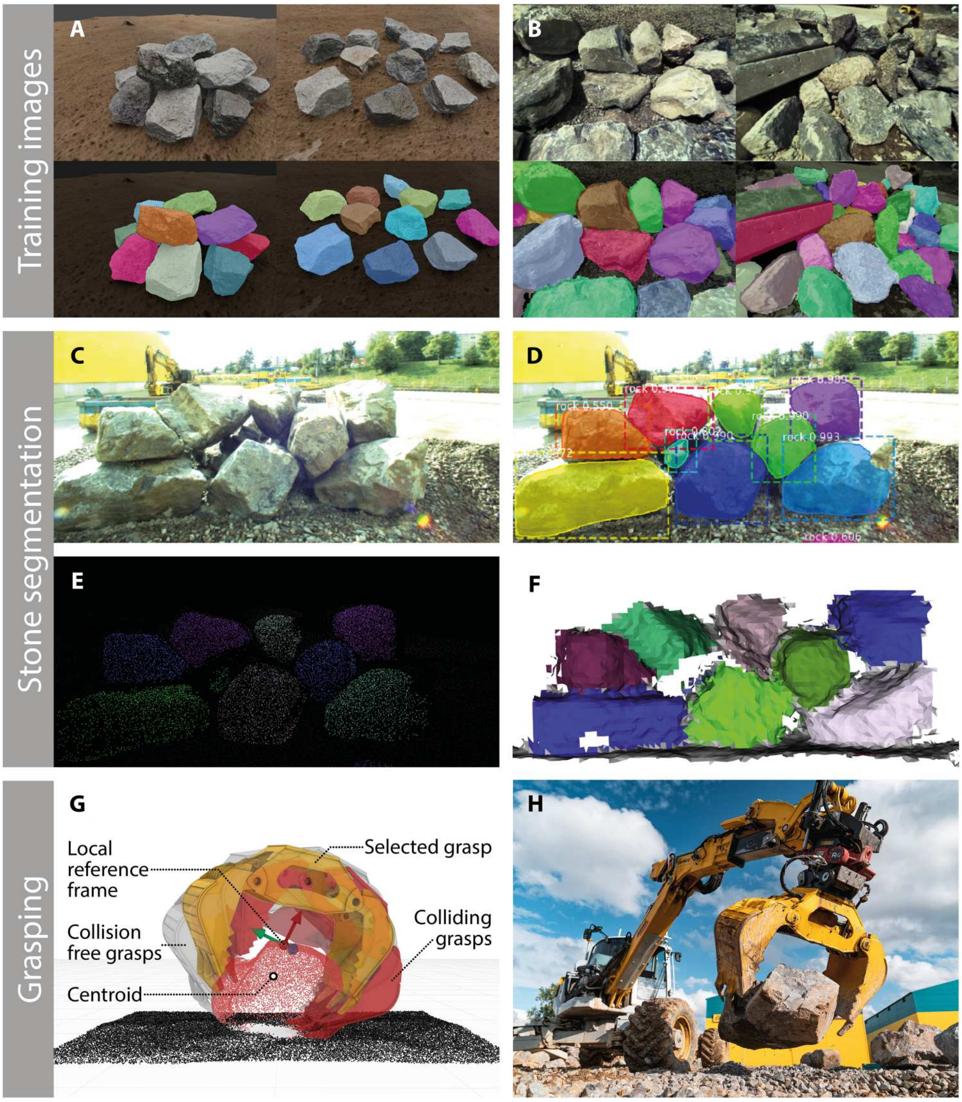

Fig. 6. Stone segmentation and grasp planning. The stone segmentation pipeline was trained with a combination of automatically labeled stone piles generated in simulation (**A**) and manually labeled real images (**B**). RGB input images (**C**) were masked with identified stone instances (**D**). The LiDAR map was then projected onto the image frame and segmented using the RGB masks (**E** and **F**). The grasp planner determined a feasible grasp on the segmented cloud (**G**) that was then executed by HEAP (**H**).

excavator cabin to provide oversight and avoid accidents: Picking, placing, and scanning were executed autonomously, whereas driving between these operations was done manually. In both cases, high-level task sequencing and the handling of unexpected situations were managed by a human supervisor. Unexpected situations can include stones that fall entirely off the wall or the need to remove stones with problematic settling (the freestanding and retaining walls had 6 and about 25 such events, respectively). In these cases, the process is interrupted by the supervisor, the problematic stone is moved away, and the planning-and-building process is restarted by first relocalizing each stone using the LiDAR map. To remove these requirements, future work can implement automation of this overarching management, together with providing a high-level understanding of holistic site goals beyond a single task, thereby allowing robotic devices to safely recognize, collaborate with, or cede access to adjacent robots and operations.

In our digging and assembly experiments, the machine relied on a Global Navigation Satellite System (GNSS) for site localization: On several occasions, this reliance resulted in undesirable down time because of connection failures to the RTK base station. These concerns have been subsequently mitigated with a graph-based multisensor fusion method (*84*) that also incorporates LiDAR for pose estimation. Future works might also include incorporating texture- or LiDAR-based methods that provide local adjustments to the global end-effector pose to improve alignment in grasping and placement.

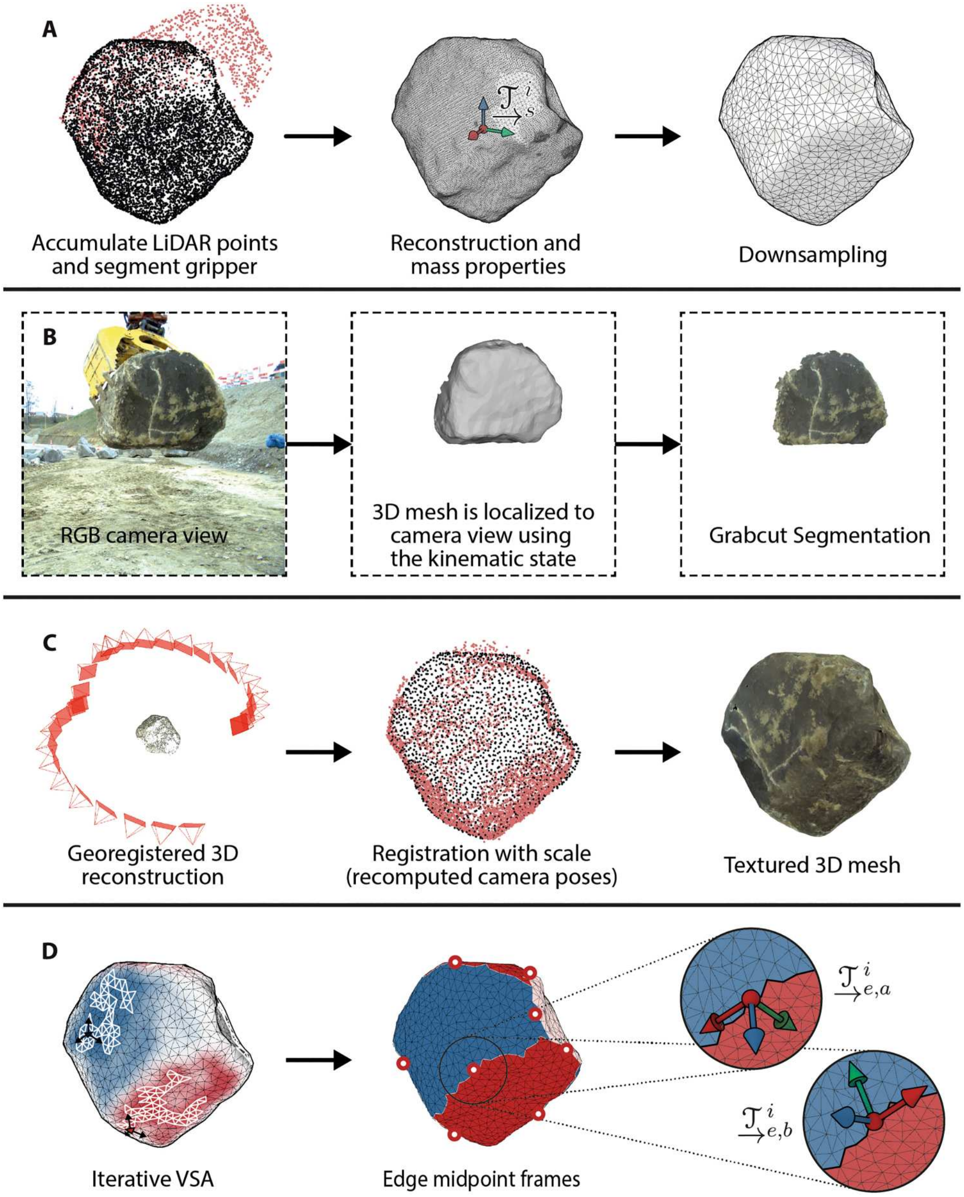

Fig. 7. Stone digitization. (**A**) The stone mesh and mass properties were generated from a cloud accumulated while the stone was held in the gripper and rotated. (**B**) The salient texture region was extracted from the RGB camera images by first generating a corresponding image of the stone mesh in the camera frame using the kinematic state and then applying that image as a seed for grabcut segmentation. (**C**) For improved alignment for texturing, we recomputed camera poses using a sparse reconstruction, aligned the 3D reconstruction to the base mesh, and projected the segmented images from these updated camera poses onto the mesh to create a textured model. (**D**) Using the downsampled mesh, we iteratively applied the seed-and-flood VSA method to identify meta-faces and edge frames.

Outlook

Our use of local and reclaimed materials for the robotic assembly of dry-stacked structures has been applied toward the construction of freestanding and retaining walls, but the process could also be adapted to the development of coastal defenses, transportation, agricultural, or extraterrestrial infrastructure and as an alternative building method in sustainable disaster recovery (*85*, *86*).

Dry stone structures are not only materially efficient to produce but are fully reversible and reusable at their end of life: Our process extends this circular capacity by producing a process-inherent 3D database of the materials available within the wall, useful for future planning. Robotic automation of stone masonry construction has the potential to restore the widespread feasibility of a task that otherwise requires expensive and time-consuming expertise on site, allowing for the use of nontoxic, low-energy, local, and natural materials that can reflect regional vernacular traditions and provide an improved sense of place. Similarly, the building process can support local economies of material acquisition (*12*) and encourage sustainable urban mining or upcycling of disused structures. We have used these walls as the basis for a cradle–to–end of construction evaluation of their associated greenhouse gas emissions in section S4, showing that our process demonstrates a conservative 41%

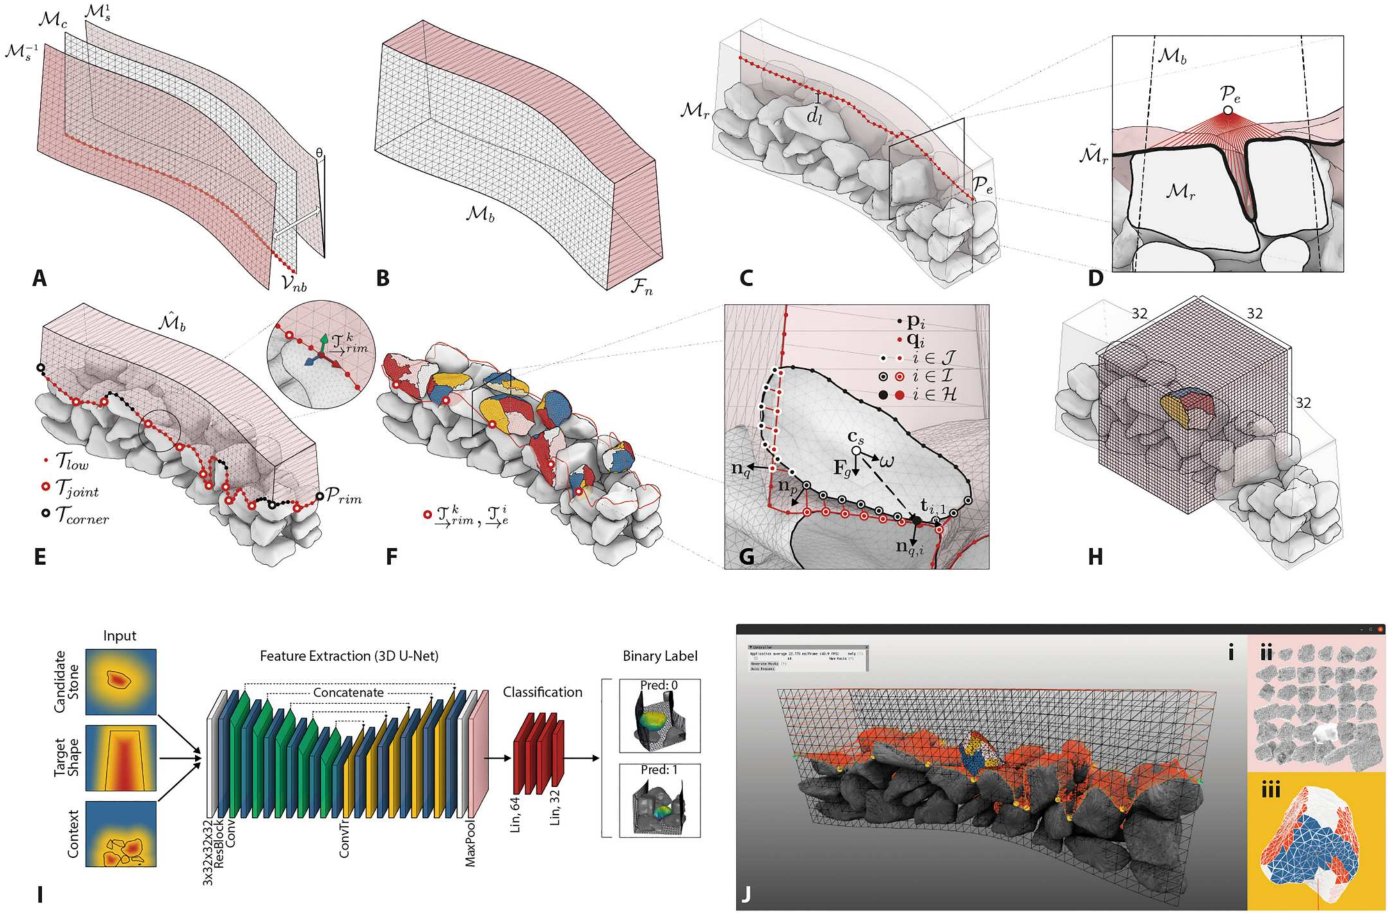

Fig. 8. Stone placement planning. The construction objective was specified by a single designed surface that was offset with a specified rake angle θ (**A**) to produce a closed volume (**B**). At each planning iteration, rays were cast in all directions from a curve $\mathscr{P}_e$ lofted above the center line of the as-built wall (**C**) to approximate the top surface of the existing ground and placed stones (**D**). The intersection of this surface with the designed bounding volume defined the rim path $\mathscr{P}_{\text{rim}}$, which was filtered on the basis of masonry-informed heuristics (**E**) to provide the basis for the initial seed frames of candidate stone placements (**F**). These initial poses were further refined for alignment and stability (**G**). Each of the registered and stable candidate solutions were then encapsulated in an offset and voxelized bounding box centered on the stone origin (**H**), and an SDF was used to classify the placeability of the solution (**I**). A custom software interface (**J**) provided interactive visualizations of the wall state, search area, and solutions (i); the available inventory (ii); and individual stone models (iii).

reduction in embodied carbon equivalents when compared with a similarly performing reinforced concrete wall. With the promise of reducing embodied energy through such adaptive assembly methods, future robotics research can facilitate new models for multiscalar manufacturing and construction, unlocking the latent potential of nonstandard and reclaimed materials toward sustainable products, buildings, and infrastructure.

MATERIALS AND METHODS

Mapping and segmentation

To accumulate an inventory of digitized building materials, we must first localize and segment candidate building elements that are freely distributed on the construction site. By generating a precise map that includes the location and shape of nearby stones, we can plan grasp configurations and arm motions that allow for additional in-hand analysis of these objects.

Our mapping pipeline builds on the LaserSLAM (*87*) framework for incremental pose-graph optimization, with extensions that consider odometry constraints from the fusion of GNSS and inertial measurement unit (IMU) readings and scan-matching constraints generated by the ICP registration (*88*) of incoming LiDAR scans to a subset of previously acquired scans. These scans are filtered using the kinematic state of the machine (to remove self-seen points from the excavator′s arm and legs) and integrated into a volumetric map that is built around the current robot location, uncorrupted by its moving parts. This map is leveraged for landscaping, stone detection, and grasping and is also shared with the geometric planner, where it is meshed to provide collision geometries of the ground that can be regularly updated as necessary (such as with the retaining wall, where surrounding earthwork and backfilling operations were conducted during construction).

Although our previous work also used LiDAR-based mapping to facilitate autonomous grasping (*42*), its reliance on geometric segmentation limited applications to semicontrolled environments with approximately flat ground. To handle more challenging, unstaged, and naturally occurring settings, including arbitrary piles of stones and rough terrain, we augmented the LiDAR sensors with a cabin-mounted camera, providing visual cues that have proven to be more suitable for instance segmentation in cluttered environments.

The improvements to instance-aware mapping presented here leverage state-of-the-art, learning-based visual object recognition techniques to process the live camera feed: We used a custom Mask R-CNN architecture (*89*) trained for the task of stone detection and segmentation on RGB images. To this end, we generated a training dataset containing automatically masked, rendered images of photorealistic 3D stone models that were randomly aggregated in simulation (Fig. 6A) (*90*). After training with synthetic images, the neural network was fine-tuned on a small set of manually labeled images captured on site with a handheld camera (Fig. 6B). For each recorded trajectory, we generated a 3D reconstruction using COLMAP (*91*, *92*) and labeled the resulting dense point cloud using the Supervisely (*93*) annotation tool. Last, 3D labels were projected back to the original images using COLMAP's estimated camera poses, and the resulting 2D instance segmentation maps were refined by means of a conditional random field. At test time, for each stone recognized in the current image, the neural network predicted a 2D segmentation mask that was propagated from the image plane to the 3D map and fused with instance predictions from previous views to achieve an accurate and consistent 3D instance-level segmentation (Fig. 6, C to F).

To propagate instance segmentation information from 2D to 3D, scans acquired between two input images were projected onto the current image plane using LiDAR odometry estimates and the precalibrated camera-LiDAR transformation. The generalized hidden point removal algorithm (*94*) was then applied to discard projected 3D points that were not visible to the camera. This determination of view-specific point visibility was crucial, given that the camera and the LiDARs were mounted some distance apart. Last, the 2D segmentation masks predicted by Mask R-CNN were used to segment the projected point cloud, and the resulting 3D segments were associated with existing stone instances already present in the map (that were detected in previous views). This was achieved by searching, for each stone instance in the map, for the best-fitting 3D segment predicted in the current view. Any remaining unmatched segments were used to initialize a new stone instance in the map.

As newly segmented raw point clouds were integrated into the map, the 3D segmentation was continually regularized in a separate thread by means of an efficient, graph-based label diffusion approach. This method allowed instance labels to be easily propagated between neighboring 3D voxels if the underlying surface was smooth, thus enforcing 3D segmentation boundaries to match actual geometric boundaries. For more specific technical details regarding the label diffusion approach and online instance-level mapping pipeline, we refer the reader to (*95*).

Grasp planning

The goal of the grasp planner was to find viable grasp configurations on the detected object instances that were collision-free and have contact wrenches that span the object's centroid (force closure; Fig. 6, G and H). We generated grasp hypotheses for a two-finger gripper model by sampling surface points on the point cloud representation of the object instance (*42*). For each sample point, a local reference frame was assigned with its x axis pointing along the normal and its z axis in the direction of minimal curvature. A convex polyhedral gripper collision model was aligned with the reference frame and moved along the normal direction with several opening angles until the palm or jaws were in contact with the sampling point cloud. These many sampled grasp hypotheses were first evaluated for validity (force closure) by a grasp classifier using a four-layer convolutional neural network (CNN) presented in (*96*). The valid grasps were further ranked by a task-specific heuristic cost, which resulted in a physical grasp success rate of 82.2%, where failures most commonly occurred because of stone movement during grasp closure, either from imprecise grasp pose execution, slight grasp asymmetries that caused the stone to shift before both sides of the gripper made contact, or wedge-shaped stones that were pushed away by the gripper's closing force. In the case of a slipped grasp or dropped stone, the grasp planner started afresh. At present, these grasp failures are detected manually, but this could be achieved autonomously in the future by comparing the actual gripper closing angle with the specified closing angle for a given grasp. A more detailed overview of grasp planning from segmented object instances can be found in (*42*, *43*).

Stone digitization

Scanning

After the initial grasp that was planned on a partially obscured point cloud, we generated a full mesh of each held object by scanning it in hand. Using the cabin-mounted LiDAR, we accumulated points in the end-effector coordinate frame while spinning the gripper and stone at a constant velocity of 0.18 rad/s and filtering out points belonging to the gripper, achieving full coverage in a single grasp. To remove undesired artifacts (such as from falling dust), we applied a statistical outlier removal filter and isolated the largest cloud after Euclidean cluster extraction (r = 0.2 m) (*97*) for meshing with Poisson surface reconstruction (*98*). Each mesh was automatically validated and repaired if necessary (*99*) and then used to compute the mass properties (*100*) and to produce a downsampled copy. Because the initial reconstruction output overly dense meshes, with about 28,000 faces and 14,000 vertices/m^2 (±40%), downsampling allowed for improved loading, rendering, and planning performance. These downsampled copies have an average resolution of 523 ± 11 faces/m^2 and 261 ± 4 vertices/m^2. Each stone was assigned an intrinsic reference frame $\underset{\rightarrow}{T}{}_{s}^{i}$ that originates at its centroid and is aligned to its principal axes (computed with principal components analysis on the original mesh vertices; Fig. 7A) (*41*). The full scanning sequence took 4.7 min per stone, averaged across 110 field-digitized boulders.

Texturing

To obtain a textured mesh, we simultaneously recorded RGB images from a cabin-mounted camera during the scanning process. Using the known excavator and reconstructed stone geometry in each view, the images were masked to remove the background and excavator using GrabCut (*101*), initialized with the center of the visible stone mesh (Fig. 7B). To correct for misalignments of the camera pose with respect to the held stone, we reconstructed the camera poses with a georegistered sparse 3D reconstruction and registered and scaled this sparse cloud to the LiDAR-produced point cloud using ICP registration, applying the same transformation to the camera poses. The images at these recomputed camera poses were then projected and blended onto the mesh using MVS texturing (Fig. 7C) (*102*). The texturing pipeline is currently used for documentation and to assist human recognition and troubleshooting with specific stones. In the future, it could facilitate feature-based pose estimation or enable planning for visual patterns and color distributions on wall faces.

Face approximation

To assist the assembly planner, we precomputed additional intrinsic geometric properties and stored them in the database alongside each stone object. Edges and meta-faces (approximately planar regions on the stone) were computed by iteratively applying the variational shape approximation method (VSA) (*41*, *103*): Beginning with a number of randomly seeded proxy faces ($k_s = 4$), we repeated the VSA process of distortion-minimizing flooding and proxy fitting until each region had an $\mathscr{L}^{2,1}$ error metric below a specified threshold. If the maximal $\mathscr{L}^{2,1}$ error condition was not met after 10 iterations, then the region with the most distortion was split (incrementing k_s), and the process was repeated until there was some quantity k_s of meta-faces that all fell within the specified $\mathscr{L}^{2,1}$ error constraint. A meta-face adjacency matrix facilitated the identification of stone edges, and each edge midpoint was assigned two local reference frames $\underset{\rightarrow}{T}^{i}_{e,a}, \underset{\rightarrow}{T}^{i}_{e,b}$ in $\underset{\rightarrow}{T}^{i}_{s}$ whose normals were aligned with the mean area-weighted normal of the mesh faces in each adjacent meta-face region and whose x axis was aligned to the stone edge (Fig. 7D). Fast point feature histogram (FPFH) descriptors (*104*) with a radius of 0.2 m were precomputed at each of these edge midpoints using points populated on the mesh surface and normals estimated from the mesh.

Geometric planning

The geometric planner considered the inventory of scanned stones and attempted to find a collection of candidate placements $c_i \in \mathscr{C}$ in the target wall, where each candidate c_i consisted of the tuple (s_i, $\mathbf{O}_i$, v_i) containing the unique stone identifier s_i, the planned pose $\mathbf{O}_i$ of the stone in the world reference frame, and an assigned fitness value v_i.

Each target wall geometry was defined by an arbitrary open triangular mesh that represented the center of the wall $\mathscr{M}_c = (\mathscr{V}_c, \mathscr{F}_c)$, offset in opposing directions with a specified distance and rake angle [*40*] to produce the outer wall surfaces $\mathscr{M}_s^{-1}$ and $\mathscr{M}_s^{1}$ (if desired, then $\mathscr{M}_s^{-1}$ and $\mathscr{M}_s^{1}$ can also be set explicitly) (Fig. 8A). These surfaces were joined into a closed bounding volume $\mathscr{M}_b = (\mathscr{V}_s^{-1} \cup \mathscr{V}_s^{1},\ \mathscr{F}_s^{-1} \cup \mathscr{F}_s^{1} \cup \mathscr{F}_n)$, where $\mathscr{F}_n$ consists of faces derived from the trivial triangulation of the corresponding naked border vertices in $\mathscr{V}_s^{-1}$ and $\mathscr{V}_s^{1}$ (Fig. 8B).

Candidate seeding

Each search for placeable candidates $\mathscr{C}$ began with a superset of candidates $\mathscr{S} \supset \mathscr{C}$, seeded within the boundary volume $\mathscr{M}_b$, that were gradually refined and reduced toward $\mathscr{C}$. Rather than randomly selecting stones and poses within the intractably large search space of $\mathscr{M}_b$, we leveraged the intrinsic geometric properties of the stones to quickly produce a fixed number of partially registered candidates. Taking the mesh $\mathscr{M}_r$ representing the rigid objects in the scene (a concatenation of the meshes for the ground and any previously placed stones) with vertices $\mathscr{V}_r$ and the naked boundary vertices $\mathscr{V}_{nb} \subset \mathscr{V}_c$ of the lowest edge in $\mathscr{M}_c$, we produced a collection of points $\mathscr{P}_e$ projected onto $\mathscr{M}_c$ and lifted slightly above the center of the built structure ($d_l = 0.2$ m; Fig. 8C)

$$\mathscr{P}_e = \operatorname{proj}_{\mathscr{M}_c}\left(\mathbf{v}^i_{nb,x}, \mathbf{v}^i_{nb,y}, \max\left\{ z \,\middle|\, \begin{matrix} (x,y,z) \in \mathscr{V}_r \quad \wedge \\ \mid \mathbf{v}^i_{nb,x} - x \mid < 1 \quad \wedge \\ \mid \mathbf{v}^i_{nb,y} - y \mid < 1 \end{matrix} \right\} + d_l \right)$$

$$\forall \mathbf{v}^i_{nb} \in \mathscr{V}_{nb} \tag{1}$$

From each point in $\mathscr{P}_e$, we cast rays in all directions and remeshed the set of points where these rays strike $\mathscr{M}_r$ using Poisson reconstruction to produce a single simplified surface $\tilde{\mathscr{M}}_r$ that approximated the top of the existing structure and lofted over small irregularities near the wall boundaries (Fig. 8D). The subspace of $\mathscr{M}_b$ above $\tilde{\mathscr{M}}_r$ represents the boundary of the active search volume $\hat{\mathscr{M}}_b$, and the loop of counterclockwise-ordered vertices at the intersection of $\tilde{\mathscr{M}}_r$ and $\mathscr{M}_b$ defines the rim path $\mathscr{P}_{\text{rim}}$, the upper perimeter edge of the in-progress wall. The collection of frames $\underset{\rightarrow}{T}^{k}_{rim} \in \mathscr{T}_{\text{rim}}$ were positioned at each point in $\mathbf{p}^k_{\text{rim}} \in \mathscr{P}_{\text{rim}}$, with normals aligned to those of $\mathscr{M}_b$ and with x axes tangent to $\mathscr{P}_{\text{rim}}$. Candidate seed poses were thus obtained by registering each stone edge frame $\underset{\rightarrow}{T}^{i}_{e,a}$, $\underset{\rightarrow}{T}^{i}_{e,b}$ with each rim frame $\underset{\rightarrow}{T}^{k}_{rim}$.

To improve the bonding and stability of a stone wall, traditional masons suggest placing each stone above the joint between two underlying stones, such that there are few vertical running joints between layers (section S5) (*39*, *40*). It is also typical to maintain an even height during construction and to build inward from the wall corners (*105*). We incorporated these concepts as a heuristic filter to our perimeter frames $\mathscr{T}_{\text{rim}}$ by first identifying a set of key frames $\mathscr{T}_{\text{key}} = ((\mathscr{T}_{\text{corner}} \cup \mathscr{T}_{\text{joint}}) \cap \mathscr{T}_{\text{low}}) \subset \mathscr{T}_{\text{rim}}$ where

$$\mathscr{T}_{\text{joint}} = \underset{\rightarrow}{T}^{k}_{rim} \forall\, k, \text{such that (s.t.)} \left\{ k \mid o_{\mathbf{p}^k_{\text{rim}}} \neq o_{\mathbf{p}^{k-1}_{\text{rim}}}, \mathbf{p}^k_{\text{rim}} \in \mathscr{P}_{\text{rim}} \right\} \tag{2}$$

$$\begin{aligned} &\mathscr{T}_{\text{low}} = \underset{\rightarrow}{T}^{k}_{rim} \forall k, \\ &\text{s.t.} \left\{ k \mid \mathbf{p}^k_{\text{rim},z} \le z_{\min} + b \left\lceil r_b \left\lceil \frac{z_{\max} - z_{\min}}{b} \right\rceil \right\rceil, \mathbf{p}^k_{\text{rim}} \in \mathscr{P}_{\text{rim}} \right\} \end{aligned} \tag{3}$$

$$\begin{aligned} \text{with } z_{\min} &= \min\{z \mid (x,y,z) \in \mathscr{P}_{\text{rim}}\} \\ \text{and } z_{\max} &= \max\{z \mid (x,y,z) \in \mathscr{P}_{\text{rim}}\} \end{aligned} \tag{4}$$

$$\begin{aligned} &\mathscr{T}_{\text{corner}} = \underset{\rightarrow}{T}^{k}_{rim} \forall k, \\ &\text{s.t.} \left\{ k \mid \left(\frac{\mathbf{p}^k_{\text{rim}} - \mathbf{p}^{k-1}_{\text{rim}}}{\|\mathbf{p}^k_{\text{rim}} - \mathbf{p}^{k-1}_{\text{rim}}\|} \right) \cdot \left(\frac{\mathbf{p}^{k+1}_{\text{rim}} - \mathbf{p}^k_{\text{rim}}}{\|\mathbf{p}^{k+1}_{\text{rim}} - \mathbf{p}^k_{\text{rim}}\|} \right) < 0.5, \mathbf{p}^k_{\text{rim}} \in \mathscr{P}_{\text{rim}} \right\} \end{aligned} \tag{5}$$

where $o_{\mathbf{p}^k_{\text{rim}}}$ is the index of the closest object in the combined set of rigid objects $\mathscr{M}_b$ to the kth rim point $\mathbf{p}^k_{\text{rim}}$. Toward building with horizontal layers, we divided the rim path points into bands with height $b = 0.3$ m and prioritized filling in the lower areas by including only the bottom half of these bands ($r_b = 0.5$; Fig. 8E). We then finalized a set of seed frames $\mathscr{T}_{\text{seed}} \subset \mathscr{T}_{\text{rim}}$ to include all perimeter frames within a given radius of a key frame ($r_{\text{key}} = 0.75$ m)

$$\mathscr{T}_{\text{seed}} = \underset{\rightarrow}{T}{}^{k}_{rim} \forall k,$$

$$\text{s.t.} \left\{ k \mid \arg\min_{j} \|\mathbf{p}^{k}_{\text{rim}} - \mathbf{p}^{j}_{\text{rim}}\| \leq r_{\text{key}}, \mathbf{p}^{k}_{\text{rim}} \in \mathscr{P}_{\text{rim}}, \underset{\rightarrow}{T}{}^{j}_{rim} \in \mathscr{T}_{\text{key}} \right\} \tag{6}$$

After computing FPFH descriptors at each frame in $\mathscr{T}_{\text{seed}}$ (*106*), we found the k-closest correspondences between each stone edge and each frame in $\mathscr{T}_{\text{seed}}$ ($k \approx 10$). From these correspondences, we prefiltered any likely unstable stone poses that have a low depth-to-height ratio (*41*) and kept the top n correspondences as set $\mathscr{S}$ for further refinement ($n \approx 2 \times 10^4$). This approach yields a relatively fast set of approximately feasible solutions (Fig. 8F), with the entire process of reconstructing $\mathscr{P}_e$, $\tilde{\mathscr{M}}_r$, $\hat{\mathscr{M}}_b$, and $\mathscr{P}_{\text{rim}}$; filtering; and computing descriptors and correspondences typically taking under 10 s (on a laptop running Ubuntu 20.04 with an Intel i7-8750H 4.1-GHz processor and 64 GB of RAM).

Torque and intersection constrained registration

The candidate placements seeded in $\mathscr{S}$ can exhibit infeasible intersections with the existing context (the ground and partially built wall), poor geometric alignment, and instability. Thus, we required some approach to refine these toward a new pose that is both stable and well aligned to the desired wall surface. Here, we adapted the approach of ICP registration, typically used for aligning point sets from observations of corresponding objects or scenes, for the task of rigidly transforming a stone such that it was well mated with its surroundings. In our previous work (*41*), we applied this strategy by using conventional point-to-plane ICP registration until convergence and then iterating a rigid-body physics simulation for a maximum number of time steps or until equilibrium was reached. The downsides of this method are that the registration and physics simulation are separate processes that can pull the candidate stone in opposite directions. In addition, the ICP registration typically results in many intersections between the stone and the surrounding rigid objects, which can lead to instability in the physics solver; we would thus project the stone up from any intersections after ICP registration, degrading the geometric alignment to the local context. To overcome these issues, we introduced a torque and intersection constrained registration (TICR) approach that combines contact-induced rotation with ICP registration in a single step, allowing us to compute contact points only as needed for registration rather than relying on a separate representation (such as convex decomposition) for collision detection.

We took the set of sample points $\mathscr{P}_s = \{\mathbf{p}_i, \cdots, \mathbf{p}_{|\mathscr{N}|}\}$, populated with a specified density ($D_p \approx 1000$ points/m^2) on each candidate stone with area A_s, where $\mathscr{N}$ is the set of sequential indices {1, ⋯, $\lceil D_p A_s \rceil$}. Given an equally dense set of points $\mathscr{Q}_b$ on our search volume $\hat{\mathscr{M}}_b$, we considered foot points $\{\mathbf{q}_i, \cdots, \mathbf{q}_{|\mathscr{N}|}\} \in \mathscr{Q}_b$, where $\mathbf{q}_i$ is the closest point on the stationary search volume for each point $\mathbf{p}_i$. We obtained surface normals at each point from the underlying mesh geometries, where $\mathbf{n}_{p,i}$ and $\mathbf{n}_{q,i}$ are the normals at points $\mathbf{p}_i$ and $\mathbf{q}_i$, respectively. Note that because the normals $\mathbf{n}_{q,i}$ were derived from the outward-facing closed mesh $\hat{\mathscr{M}}_b$, those normals that correspond to the ground or already placed rocks were reversed. We considered the subsets of point indices $\mathscr{J} \subseteq \mathscr{N}$ for all points where $\mathbf{n}_{p,i} \cdot \mathbf{n}_{q,i} \geq 0$ and $\|\mathbf{p}_i - \mathbf{q}_i\|^2 \leq d_t$ for some distance threshold d_t; $\mathscr{I} \subseteq \mathscr{J}$ for indices where $\mathbf{q}_i$ is on a portion of our search volume that belongs to a rigid object (the ground or previously placed stones); and $\mathscr{H} \subseteq \mathscr{I}$ for a 0.1-m voxel subsampling of the remaining points $\|\mathbf{p}_i - \mathbf{q}_i\|^2 \leq d_c$ within a contact distance threshold d_c (Fig. 8G).

Within each single iteration, we first determined whether the object is in static equilibrium by minimizing the sum of contact forces ($\mathbf{f}_i$) and torques subject to friction constraints using quadratic programming

$$\begin{aligned} \min_{\mathbf{x}} \|\sum_i \mathbf{f}_i + \mathbf{F}_g\|^2 + \|\sum_i (\mathbf{p}_i - \mathbf{c}_s) \times \mathbf{f}_i\|^2, \quad \forall i \in \mathscr{H} \\ \text{s.t.} \quad | \mathbf{f}_i \cdot \mathbf{t}_{i,j} | \leq \mu(-\mathbf{n}_{q,i} \cdot \mathbf{f}_i), \quad j = 1, 2 \\ -\mathbf{n}_{q,i} \cdot \mathbf{f}_i \geq 0 \end{aligned} \tag{7}$$

where $\mathbf{x}$ is the stacked vector of contact-forces $[\mathbf{f}_{i,\, x}, \mathbf{f}_{i,\, y}, \mathbf{f}_{i,\, z}, \cdots, \mathbf{f}_{|\mathscr{H}|,\, x}, \mathbf{f}_{|\mathscr{H}|,\, y}, \mathbf{f}_{|\mathscr{H}|,\, z}]^\top$ and where $\mathbf{F}_g = m_s\mathbf{g}$, for stone mass m_s and gravitational acceleration $\mathbf{g}$. We use $\mathbf{c}_s$ for the stone centroid, μ for the coefficient of friction, and $\mathbf{t}_{i,j}$ as an arbitrary pair of orthonormal vectors perpendicular to the surface normal $\mathbf{n}_{q,i}$.

We then continued to register the stone by computing its transformation with the point-to-plane objective function for ICP (*88*), applying a zero penetration constraint (*107*) to rigid objects, an incrementally applied partial penetration constraint to the target wall surface ($d_p \approx 0.1$ m), and increased weight (ψ_i) to foot points on rigid objects. When the stone was in contact but not in equilibrium, we also constrained its rotation to align with the contact-induced rotation axis $\boldsymbol{\omega}$, derived from the sum of residual torques of Eq. 7

$$\begin{aligned} \min_{\hat{\mathbf{T}}} \quad & \sum_i [(\mathbf{T}(\mathbf{p}_i) - \mathbf{q}_i) \cdot \psi_i \mathbf{n}_{q,i}]^2, && \forall_i \in \mathscr{J} \\ \text{s.t.} \quad & \mathbf{n}_{q,i} \cdot (\mathbf{T}(\mathbf{p}_i) - \mathbf{q}_i) < 0, \psi_i = 1, && \forall_i \in \mathscr{I} \\ & \mathbf{n}_{q,i} \cdot (\mathbf{T}(\mathbf{p}_i) - \mathbf{q}_i) < d_p, \psi_i = 0.5, && \forall_i \in \mathscr{J} \backslash \mathscr{I} \\ & a \cdot (\boldsymbol{\omega}/\|\boldsymbol{\omega}\|) > \phi \Rightarrow \|\boldsymbol{\omega}\| > 0, && \phi \geq 0 \\ & \|\mathbf{a} \times \boldsymbol{\omega}\| \Rightarrow \|\boldsymbol{\omega}\| > 0 \end{aligned} \tag{8}$$

where the operator $\mathbf{T}(\mathbf{p}_i)$ applies a homogeneous transformation to point $\mathbf{p}_i$ using the linearized 4 x 4 transformation matrix

$$\hat{\mathbf{T}} = \begin{bmatrix} 1 & -\gamma & \beta & t_x \\ \gamma & 1 & -\alpha & t_y \\ -\beta & \alpha & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{9}$$

with rotation component $\mathbf{a}$ = (α, β, γ). Because the rotation angle is determined by the magnitude of $\mathbf{a}$, we forced a small rotation by increasing the value of ϕ when the object is not in equilibrium to escape local minima where the stone is aligned geometrically but unstable ($\phi = 0.1$). We then allowed convergence by reducing this value by half at each impulse where $\boldsymbol{\omega}_{\text{current}} \cdot \boldsymbol{\omega}_{\text{previous}} < 0$.

For each candidate solution, we incrementally applied TICR up to a maximum number of iterations ($n_{\text{iter}} \approx 50$) or until convergence. Convergence was accepted when the pose error between current and previous candidates $\mathscr{E}_{\text{pose}}(\mathbf{O}_{\text{current}}, \mathbf{O}_{\text{previous}}) < (1 \times 10^{-4}, 1 \times 10^{-4})$. After registration, these solutions were sorted by the foot-point distance RMSE, and the top candidates ($n_{\text{top}} \approx 5 \times 10^3$) were run through a rigid-body physics engine (*108*). Because of its consideration of penetration and torques (although not yet sliding), TICR-registered solutions were 1.6 times more likely to reach equilibrium in the physics simulation compared with conventional ICP, a figure that we obtained by comparing the simulation survival rate of 10,000 TICR-registered and ICP-registered seed candidates on the same partially built wall (24.6 and 15.0%, respectively). This entire sequence of registration and simulation was iterated twice per execution of the geometric planner (section S6).

Classification

Because our method attempts to avoid the error accumulation and high logistical and computational costs associated with precomputing the full structure, it is essential that individually placed stones are not detrimental to the feasibility of downstream solutions. This can be achieved by applying heuristics from traditional masonry guidebooks (*35*, *41*), such as having the top surface of stones slope toward the interior of the wall and avoiding solutions that leave too little space for adjacent stones by extending close to the opposite side of the wall (*40*). In practice, however, these heuristics can become complicated by special cases, for example, by throughstones that span both sides or others that fall fully within the center of the wall. The surrounding context is also often as important as the candidate stone in determining the validity of the solution. Mitigating these challenges, we used an SDF-based classifier to determine whether an otherwise registered and stable solution is a placeable one. We trained our 3D U-net–based (*109*–*111*) model on hand-labeled three-channel SDFs from a database of candidate solutions in automatically generated walls (fig. S3). For each candidate c_i, we created a bounding region centered on $\mathbf{O}_i$ whose nonuniform dimensions were determined such that there was one meter of clearance on all sides of the axis-aligned bounding box of c_i. This box was subdivided into a 32^3-point voxel grid, and from each point in the grid, we computed three separate signed distances: one to the surface of the candidate stone c_i, one to the boundary volume $\mathcal{M}_b$, and another to the existing rigid objects $\mathcal{M}_r$ (Fig. 8, H and I). The model was trained on 9000 subjectively labeled SDFs and obtained 87% accuracy with the labeled test set of previously unseen stones.

Subset selection

After refinement and filtering, each execution of the solver yields distributed solutions that consider the full inventory across the extended search space of the wall: A single stone can have several valid poses along the length of the wall, and a single position in the wall can have multiple valid solutions with different stones. To reduce the overhead of running the solver for each placed stone, we maximized the number of high-quality solutions that can be placed simultaneously from a given solution set.

Given a solution set $\mathcal{S}$, we looked for the subset $\mathcal{C} \subseteq \mathcal{S}$ to maximize fitness

$$\begin{aligned} \max_{k\{c_1,\cdots,c_k\}} \quad & \sum_{i}^{k} v_i \qquad \forall (s_i, \mathbf{O}_i, v_i) \in \mathcal{S} \\ s.t. \quad & \neg(c_a \mid c_b) \quad \forall a, b \in \{c_1, \cdots, c_k\} \end{aligned} \tag{10}$$

where the incompatibility between two candidates $c_a \mid c_b$ is defined as

$$c_a \neq c_b \wedge (s_a = s_b \vee c_a \cap c_b \neq \emptyset) \tag{11}$$

and where $c_a \cap c_b \neq \emptyset$ indicates a physical intersection of the two candidate meshes computed using the fast winding number (section S7) (*112*). This subset selection process allows for the placement of many solutions in longer walls (around four stones per solver iteration with 15-m walls). Given the relatively constant duration of the full solver (≈4.8 min), this process typically saves more than 75% of computing time per placed stone compared with only accepting a singular top candidate.

Pick and place

Given $\mathcal{C}$, the final set of compatible candidate solutions, we planned feasible grasps and motion trajectories that allowed these solutions to be placed in the in-progress structure. Grasp planning for placement is similar to the initial grasping of segmented object instances, as outlined in the grasp planning section above. In this case, however, we were able to augment the point cloud used for sampling grasp hypotheses with the meshes of already scanned and localized stones. Picking for placement additionally requires consideration of collision constraints at both ends of the sequence: We first generated grasp hypotheses on the desired stone by transforming its mesh to the target placement pose and sampling noncolliding grasps. These grasp hypotheses were then reprojected to the current stone pose (for picking), evaluated for collision and force closure, and ranked by task-specific heuristic costs (*43*).

If no direct pick-and-place sequence was feasible, then we reoriented the stone to a more suitable intermediate pose using a strategy that exploits the particular range of motion of the excavator arm. We first selected a preferred placement grasp that was unconstrained by the picking orientation and context. Then, we planned for an intermediate pose by letting the stone settle in a physics simulation from several starting orientations: The settled pose that allowed for a grasp closest to the preferred one was selected as a stable intermediate placement pose. A reorientation grasp was then planned that lies on the mid-plane between the actual ground plane and the projected ground plane of the intermediate pose. By revolving the stone about the wrist's infinite yaw axis while it was held in the reorientation grasp, the stone was brought into the intermediate pose, and grasp planning for placement was applied again.

To avoid unwanted collisions during placement, we computed the direction of a 1-m linear motion trajectory to be tracked during object placement. This desired placement direction ($\mathbf{n}_{\text{place}}$) was determined by casting rays in all directions from the stone centroid at the desired pose $\mathbf{O}_i$. Rays that did not hit the ground or surrounding objects were averaged to obtain the placement direction with the highest margin to obstacles (*41*).

Refinement

At each physical interaction with a stone (grasping, placement on the ground, and placement in the wall), unpredictable settling occurs that causes the stone to translate or rotate slightly with respect to its planned reference frame. To adapt to these perturbances, we refined the pose of the stone using ICP at each step, registering a point cloud sampled on the reconstructed stone mesh to the LiDAR map accumulated in the neighborhood of the approximately localized stone (*41*, *43*). Upon releasing a stone on the ground or in the wall, both cabin- and arm-mounted LiDAR sensors were used for mapping its surrounding area, where the reach of the excavator facilitated refinement even when stones were placed on the backside of the structure. Points belonging to the ground or to previously refined stones were segmented out of the laser map before ICP registration, and the newly refined stone pose was updated in both the inventory database and geometric planner for accurate grasping and simulation in subsequent operations.

Stone poses were similarly refined in hand during placement: The held stone was slowly rotated while being transferred to the wall, accumulating scan points as with the scanning operation. The final placement trajectory was modified to account for any detected grasp deviations, or the stone was returned to the ground for regrasping if the shifted stone could no longer be placed because of gripper collisions or kinematic constraints.

Stability probing

Stone walls must withstand a variety of load cases, including lateral loads from soil retention, self-weight due to gravity, and dynamic loads such as human climbing and wind. According to masonry guidelines (*113*, *114*), the possible failure modes of slipping and overturning must be considered globally, but also locally, on an individual stone or subset of stones. Global failure modes can be mitigated in the design stage using established formulas to verify the resistance to overturning and sliding given the overall geometry and the assumed density of the wall (*113*, *114*). Such global methods, however, do not allow for the stability assessment of individual elements: Thus, we introduce a physical testing method to alleviate safety concerns from local failures in robotically assembled stone walls.

Our probing routine used a custom-built force-torque sensor based on strain gauges (section S8) to verify that placed stones were load-bearing and locally stable and could resist sliding under applied target threshold forces. Identified stones that slip or rock at low threshold forces could be further stabilized with manually placed wedging stones or mechanical fastening. An overview of the method is provided in section S9.

Construction materials

Our walls were constructed from a combination of gneiss boulders extracted from a quarry in southern Germany (35 km from the building site), erratics unearthed at nearby construction sites, and concrete debris from recently razed buildings within the metropolitan area of Zurich. A detailed outline of the stones used in construction is defined in section S10, and the full archive of 3D models has been made available (*115*).

Pose error

For assessing the placed accuracy relative to the planned pose and to identify convergence in registration, we defined the pose error between two poses $\mathbf{O}_i$ and $\mathbf{O}_j$ as

$$\mathscr{E}_{\text{pose}}(\mathbf{O}_i, \mathbf{O}_j) = (\mathscr{E}_{\text{pos}}, \mathscr{E}_{\text{orient}}) \tag{12}$$

where each pose is composed of a translation $t \in \mathbb{R}^3$ and rotation $\mathbf{R} \in SO(3)$ and where the position and orientation errors are defined respectively as

$$\mathscr{E}_{\text{pos}} = \|t_i - t_j\|^2 \text{ and } \mathscr{E}_{\text{orient}} = \sin\theta \tag{13}$$

where we obtain θ from the axis-angle representation of $\mathbf{R}_{\text{diff}}$, with axis $\mathbf{w}$

$$\mathbf{w}\sin\theta = \frac{1}{2}\begin{pmatrix} r_{32} - r_{23} \\ r_{13} - r_{31} \\ r_{21} - r_{12} \end{pmatrix} \text{ and } \mathbf{R}_{\text{diff}} = \mathbf{R}_i\mathbf{R}_j^{\mathrm{T}} := \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \tag{14}$$

Supplementary Materials

This PDF file includes:
Figs. S1 to S16
Sections S1 to S10
Tables S1 to S4
Algorithm S1
References (*118–130*)

REFERENCES AND NOTES

1. F. Barbosa, J. Woetzel, J. Mischke, *Reinventing Construction: A Route of Higher Productivity* (Tech. Rep., McKinsey Global Institute, 2017).
2. American Society of Civil Engineers, ASCE's 2021 American Infrastructure Report Card | GPA: C– (2017); https://infrastructurereportcard.org/.
3. G. Habert, S. A. Miller, V. M. John, J. L. Provis, A. Favier, A. Horvath, K. L. Scrivener, Environmental impacts and decarbonization strategies in the cement and concrete industries. *Nat. Rev. Earth Environ.* **1**, 559–573 (2020).
4. S. N. Jonkman, M. M. Hillen, R. J. Nicholls, W. Kanning, M. Ledden, Costs of adapting coastal defences to sea-level rise—New estimates and their implications. *J. Coast. Res.* **290**, 1212–1226 (2013).
5. S. Lim, V. L. Prabhu, M. Anand, L. A. Taylor, Extra-terrestrial construction processes—Advancements, opportunities and challenges. *Adv. Space Res.* **60**, 1413–1429 (2017).
6. P. J. M. Monteiro, S. A. Miller, A. Horvath, Towards sustainable concrete. *Nat. Mater.* **16**, 698–699 (2017).
7. A. Akhtar, A. K. Sarmah, Construction and demolition waste generation and properties of recycled aggregate concrete: A global perspective. *J. Clean. Prod.* **186**, 262–281 (2018).
8. U.S. Environmental Protection Agency, "*Construction and Demolition Debris Generation in the United States, 2015*" (Tech. Rep., U.S. Environmental Protection Agency, Office of Resource Conservation and Recovery, 2018).
9. L. Zhang, J. Zhao, P. Long, L. Wang, L. Qian, F. Lu, X. Song, D. Manocha, An autonomous excavator system for material loading tasks. *Sci. Robot.* **6**, eabc3164 (2021).
10. K. Fairchild, W. W. Mendell, Eds., *Report of the In Situ Resources Utilization Workshop* (NASA Conference Publication 3017, National Aeronautics and Space Administration, Scientific and Technical Information Division, 1988), p. 102.
11. J. C. Morel, A. Mesbah, M. Oggero, P. Walker, Building houses with local materials: Means to drastically reduce the environmental impact of construction. *Build. Environ.* **36**, 1119–1126 (2001).
12. A. Klemm, D. Wiggins, in *12 - Sustainability of Natural Stone as a Construction Material in Sustainability of Construction Materials*, J. M. Khatib, Ed., (Woodhead Publishing Series in Civil and Structural Engineering, Woodhead Publishing, ed. 2, 2016), pp. 283–308.
13. J. Dunbar-Nasmith, Why traditional building materials? in Conference Proceedings, I. Maxwell, N. Ross, Eds. (Edinburgh, 1997).
14. N. Crishna, P. F. G. Banfill, S. Goodsir, Embodied energy and CO_2 in UK dimension *stone. Conservation and Recycling* **55**, 1265–1273 (2011).
15. T. C. Blair, J. G. McPherson, Grain-size and textural classification of coarse sedimentary particles. *J. Sediment. Res.* **69**, 6–19 (1999).
16. G. P. Taylor, *An Automatic Block-Setting Crane: Meccano Model Controlled by a Robot Unit*, The Meccano Magazine (1938), p. 172.
17. T. Bock, T. Linner, *Construction Robots: Elementary Technologies and Single-Task Construction Robots* (Cambridge Univ. Press, 2017), vol. 3.
18. M. Kohler, F. Gramazio, J. Willmann, *The Robotic Touch: How Robots Change Architecture* (Park Books, 2014).
19. B. Kolarevic, Ed., *Architecture in the Digital Age: Design and Manufacturing* (Taylor & Francis, 2004).
20. T. Bonswetch, D. Kobel, F. Gramazio, M. Kohler, The Informed Wall: applying additive digital fabrication techniques on architecture, Synthetic Landscapes in *Proceedings of the 25th Annual Conference of the Association for Computer-Aided Design in Architecture*, pp. 489–495 (CUMINCAD, 2006).
21. S. Oesterle, Cultural performance in robotic timber construction, in *reForm( ) - Building a Better Tomorrow: Proceedings of the 29th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA 09)*, T. de Estree Sterk, R. Loveridge, D. Pancoast, Eds. (ACADIA, Chicago, IL, 2009), pp. 194–200.
22. W. McGee, J. Feringa, A. Sondergaard, *Processes for an Architecture of Volume, in Rob Arch 2012*, S. Brell-Çokcan, J. Braumann, Eds. (Springer, 2013), pp. 62–71.
23. R. L. Johns, N. Foley. Bandsawn bands, in *Robotic Fabrication in Architecture, Art and Design 2014*, W. McGee, M. Ponce de Leon, Eds. (Springer International Publishing, 2014), pp. 17–32.

24. Y. Wang, *Hook Park biomass boiler house*, in *Advancing Wood Architecture : A Computational Approach*, A. Menges, T. Schwinn, O. D. Krieg, Eds. (Routledge, 2016), pp. 169–180.
25. M. Self, E. Vercruysse, Infinite variations, radical strategies, in *Fabricate 2017*, A. Menges, B. Sheil, R. Glynn, M. Skavara, Eds. (UCL Press, 2017), pp. 30–35.
26. M. Larsson, H. Yoshida, T. Igarashi, Human-in-the-loop fabrication of 3D surfaces with natural tree branches, in *Proceedings of the ACM Symposium on Computational Fabrication- SCF '19* [Association for Computing Machinery (ACM) Press, 2019], pp. 1–12.
27. K. Wu, A. Kilian, *Designing Natural Wood Log Structures with Stochastic Assembly and Deep Learning*, in Robotic Fabrication in Architecture, Art and Design 2018, J. Willmann, P. Block, M. Hutter, K. Byrne, T. Schork, Eds. (Springer International Publishing, 2019), pp. 16–30.
28. J. de Boer, K. Klinger, J. Vermillion, B. Greenberg, G. Hittler, K. Perry, Smart Scrap (2007); https://web.archive.org/web/20210813045203/http://i-m-a-d-e.org/?p=2201.
29. B. Clifford, W. McGee, M. Muhonen, Recovering cannibalism in architecture with a return to cyclopean masonry. *Nexus Netw. J.* **20**, 583–604 (2018).
30. D. Ioannidou, S. Zerbi, G. Habert, When more is better—Comparative LCA of wall systems with stone. *Build. Environ.* **82**, 628–639 (2014).
31. M. Lambert, P. Kennedy, Using artificial intelligence to build with unprocessed rock. *Key Eng. Mater.* **517**, 939–945 (2012).
32. S. A. Nielsen, A. Dancu, Fusing design and construction as speculative articulations for the built environment, in *Future of Architectural Research*, A. Aksamija, J. Haymaker, A. Aminmansour, Eds. (Perkins+Will, 2015), pp. 65–73.
33. B. Wibranek, O. Tessmann, Digital Rubble Compression-Only Structures with Irregular Rock and 3D Printed Connectors, in *Proceedings of IASS Annual Symposia, IASS 2019* [International Association for Shell and Spatial Structures (IASS), 2019] 6, pp. 1–8.
34. A. Settimi, Q. Wang, E. Andò, J. Gamerro, K. Beyer, Y. Weinand, Projector-based augmented stacking framework for irregularly shaped objects. *Constr. Robot.* **7**, 159–175 (2023).
35. Y. Liu, J. Choi, N. Napp, Planning for robotic dry stacking with irregular stones, in *Field and Service Robotics*, G. Ishigami, K. Yoshida, Eds. (Springer Proceedings in Advanced Robotics, Springer, 2021), pp. 321–335.
36. F. Furrer, M. Wermelinger, H. Yoshida, F. Gramazio, M. Kohler, R. Siegwart, M. Hutter, Autonomous robotic stone stacking with online next best object target pose planning, in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2017), pp. 2350–2356.
37. Y. Liu, S. M. Shamsi, L. Fang, C. Chen, N. Napp, Deep Q-learning for dry stacking irregular objects, *in 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2018), pp. 1569–1576.
38. A. Menezes, P. Vicente, A. Bernardino, R. Ventura, From Rocks to Walls: a Model-free Reinforcement Learning Approach to Dry Stacking with Irregular Rocks, *in 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (CVPRW) (IEEE, 2021), pp. 2057–2065.
39. J. Vivian, *Building Stone Walls* (Storey Publishing, 1976).
40. Environmental Action Foundation, Ed., *Dry Stone Walls: Basics, Construction, Significance* (Scheidegger & Spiess, 2019).
41. R. L. Johns, M. Wermelinger, R. Mascaro, D. Jud, F. Gramazio, M. Kohler, M. Chli, M. Hutter, Autonomous dry stone. *Constr. Robot.* **4**, 127–140 (2020).
42. R. Mascaro, M. Wermelinger, M. Hutter, M. Chli, Towards automating construction tasks: Large-scale object mapping, segmentation, and manipulation. *J. Field Robot.* **38**, 22007 (2021).
43. M. Wermelinger, R. L. Johns, F. Gramazio, M. Kohler, M. Hutter, Grasping and object reorientation for autonomous construction of stone structures. *IEEE Robot. Autom. Lett.* **6**, 5105–5112 (2021).
44. C. Calvo Barentin, T. Van Mele, P. Block, Robotically controlled scale-model testing of masonry vault collapse. *Meccanica* **53**, 1917–1929 (2018).
45. N. Montagne, C. Douthe, O. Baverel, C. Fivet, Experimental investigation of beams under coupled bending and torsion, in *Proceedings of IASS Annual Symposia* (IASS, 2019), issue 11, pp. 1–8.
46. S. Magnenat, R. Philippsen, F. Mondada, Autonomous construction using scarce resources in unknown environments. *Autonom. Robot.* **33**, 467–485 (2012).
47. N. Napp, O. R. Rappoli, J. M. Wu, R. Nagpal, Materials and mechanisms for amorphous robotic construction, in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2012), pp. 4879–4885.
48. R. Fujisawa, N. Nagaya, S. Okazaki, R. Sato, Y. Ikemoto, S. Dobata, Active modification of the environment by a robot with construction abilities. *ROBOMECH J.* **2**, 9 (2015).
49. M. Saboia, V. Thangavelu, N. Napp, Autonomous multi-material construction with a heterogeneous robot team. *Robot. Autonom. Syst.* **121**, 103239 (2019).
50. S. J. Keating, J. C. Leland, L. Cai, N. Oxman, Toward site-specific and self-sufficient robotic fabrication on architectural scales. *Sci. Robot.* **2**, eaam8986 (2017).
51. D. Jenny, H. Mayer, P. Aejmelaeus-Lindström, F. Gramazio, M. Kohler, A pedagogy of digital materiality: Integrated design and robotic fabrication projects of the master of advanced studies in architecture and digital fabrication. *Arch., Struct. Const.* **2**, 649–660 (2022).
52. X. Zhang, M. Li, J. H. Lim, Y. Weng, Y. W. D. Tay, H. Pham, Q.-C. Pham, Large-scale 3D printing by a team of mobile robots. *Autom. Const.* **95**, 98–106 (2018).
53. Constructions-3D, Maxi Printer, https://en.constructions-3d.com/la-maxi-printer.
54. CyBe Construction, CyBe Robot Crawler, https://cybe.eu/3d-concrete-printing/printers/cybe-robot-crawler/.
55. J. Pankert, G. Valsecchi, D. Baret, J. Zehnder, L. L. Pietrasik, M. Bjelonic, M. Hutter, Design and motion planning for a reconfigurable robotic base. *IEEE Robot. Autom. Lett.* **7**, 9012–9019 (2022).
56. V. Helm, S. Ercan, F. Gramazio, M. Kohler, Mobile robotic fabrication on construction sites: DimRob, in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2012), pp. 4335–4341.
57. M. Silver, Peripatetic fabrication: Architecture, co-robotics and machine vision, in *Fabricate 2014*, F. Gramazio, M. Kohler, S. Langenberg, Eds. (UCL Press, 2014), pp. 243–247.
58. S. Peters, R. Belden, SAM, the robotic bricklayer, in *SMART/Dynamics of Masonry* (2014), pp. 10–14.
59. FBR, FBR | Industrial Automation Technology (2019); https://fbr.com.au/.
60. TerraClear, TerraClear | Rock Picking Implement; https://terraclear.com/.
61. C. H. Kim, M. K. Lee, Development of armour stone covering robots for breakwater construction, *in 2014 14th International Conference on Control, Automation and Systems* (ICCAS, 2014), pp. 1652–1656.
62. T. S. Kim, I. S. Jang, C. J. Shin, M. K. Lee, Underwater construction robot for rubble leveling on the seabed for port construction, *in 2014 14th International Conference on Control, Automation and Systems* (ICCAS, 2014), pp. 1657–1661.
63. A. Stentz, J. Bares, S. Singh, P. Rowe, A robotic excavator for autonomous truck loading, in Proceedings, in *1998 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Innovations in Theory, Practice and Applications, IEEE, 1998), vol. 3, pp. 1885–1893.
64. Q. Ha, M. Santos, Q. Nguyen, D. Rye, H. Durrant-Whyte, Robotic excavation in construction automation. *IEEE Robot. Autom. Mag.* **9**, 20–28 (2002).
65. T. Emter, C. Frese, A. Zube, J. Petereit, Algorithm toolbox for autonomous mobile robotic systems. *ATZoffhighway worldw.* **10**, 48–53 (2017).
66. D. Jud, S. Kerscher, M. Wermelinger, E. Jelavic, P. Egli, P. Leemann, G. Hottiger, M. Hutter, HEAP—The autonomous walking excavator. *Autom. Constr.* **129**, 103783 (2021).
67. D. Jud, P. Leemann, S. Kerscher, M. Hutter, Autonomous free-form trenching using a walking excavator. *IEEE Robot. Autom. Lett.* **4**, 3208–3215 (2019).
68. E. Jelavic, D. Jud, P. Egli, M. Hutter, *Towards Autonomous Robotic Precision Harvesting: Mapping, Localization, Planning and Control for a Legged Tree Harvester, Field Robotics* (2021).
69. M. Hutter, T. Braungardt, F. Grigis, G. Hottiger, D. Jud, M. Katz, P. Leemann, P. Nemetz, J. Peschel, J. Preisig, N. Sollich, M. Voellmy, M. Zimmermann, S. Zimmermann, IBEX—A tele-operation and training device for walking excavators, in *2016 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)* (IEEE, 2016), pp. 48–53.
70. K. H. Petersen, N. Napp, R. Stuart-Smith, D. Rus, M. Kovac, A review of collective robotic *construction. Robotics* **4**, eaau8479 (2019).
71. N. Melenbrink, J. Werfel, A. Menges, On-site autonomous construction robots: Towards unsupervised building. *Autom. Constr.* **119**, 103312 (2020).
72. D. Jud, I. Hurkxkens, C. Girot, M. Hutter, Robotic embankment. *Constr. Robot.* **5**, 101–113 (2021).
73. I. Hurkxkens, "Robotic landscapes: Topological approaches to terrain, design, and fabrication," thesis, ETH Zurich (2020).
74. T. Waters, V. Putz-Anderson, A. Garg, *Applications Manual for the Revised NIOSH Lifting Equation, Publication No. 94–110 in National Institute for Occupational Safety and Health, DHHS (NIOSH)* (U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, September 2021 edition, 1994).
75. C. Mundell, "Large scale testing of drystone retaining structures," thesis, Department of Architecture & Civil Engineering, University of Bath (2009).
76. B. Villemus, "Etude des murs de soutènement en maconnerie de pierres sèches," thesis, Civil Engineering, L'institut National Des Sciences Appliquees De Lyon, Lyon, France (2004).
77. F. Witschi, Flora, Fauna, in *Dry Stone Walls: Basics, Construction, Significance*, Environmental Action Foundation, Ed. (Scheidegger & Spiess, ed. 1, 2019), pp. 329–378.
78. K. Dörfler, S. Ernst, L. Piškorec, J. Willmann, V. Helm, F. Gramazio, M. Kohler, Remote material deposition: Exploration of reciprocal digital and material computational capacities, in *What's the Matter? Materiality and Materialism at the Age of Computation* (European Network of Heads of Schools of Architecture, 2014), pp. 361–377.

79. P. Aejmelaeus-Lindström, G. Rusenova, A. Mirjan, J. Medina Ibáñez, F. Gramazio, M. Kohler, Rock print pavilion: Robotically fabricating architecture from rock and string. *Sci. Robot.* **4**, 97–113 (2020).
80. A. Meurisse, A. Makaya, C. Willsch, M. Sperl, Solar 3D printing of lunar regolith. *Acta Astronaut.* **152**, 800–810 (2018).
81. E. Eckman, M. A. Peck, N. Napp. Lunar Infrastructure via Microscale Regolith Assembly, *in AIAA Scitech 2021 Forum (American Institute of Aeronautics and Astronautics*, 2021).
82. N. Napp, R. Nagpal, Robotic construction of arbitrary shapes with amorphous materials, in *2014 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2014), pp. 438–444.
83. V. Thangavelu, M. S. da Silva, J. Choi, N. Napp, Autonomous modification of unstructured environments with found material, in *2020 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2020), pp. 7798–7804.
84. J. Nubert, S. Khattak, M. Hutter, Graph-based multi-sensor fusion for consistent localization of autonomous construction robots. arXiv:2203.01389 [cs.RO] (2022).
85. E. Zea Escamilla, G. Habert, Global or local construction materials for post-disaster reconstruction? Sustainability assessment of twenty post-disaster shelter designs. *Build. Environ.* **92**, 692–702 (2015).
86. F. Faleschini, M. A. Zanini, L. Hofer, P. Zampieri, C. Pellegrino, Sustainable management of demolition waste in post-quake recovery processes: The Italian experience. *Int. J. Dis. Risk Reduct.* **24**, 172–182 (2017).
87. R. Dubé, A. Gawel, H. Sommer, J. Nieto, R. Siegwart, C. Cadena, An online multi-robot SLAM system for 3D LiDARs, *in 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2017), pp. 1004–1011.
88. Y. Chen, G. Medioni, Object modeling by registration of multiple range images, in *1991 IEEE International Conference on Robotics and Automation Proceedings* (IEEE, 1991), vol. 3, pp. 2724–2729.
89. K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask R-CNN, *in 2017 IEEE International Conference on Computer Vision (ICCV)* (IEEE, 2017), pp. 2980–2988.
90. Blender Foundation, Blender—A 3D modelling and rendering package; https://blender.org/.
91. J. L. Schönberger, J.-M. Frahm, Structure-from-motion revisited, *in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, 2016), pp. 4104–4113.
92. J. L. Schönberger, E. Zheng, J.-M. Frahm, M. Pollefeys, in *Pixelwise View Selection for Unstructured Multi-View Stereo, in Computer Vision–ECCV 2016*, B. Leibe, J. Matas, N. Sebe, M. Welling, Eds. (Lecture Notes in Computer Science, Springer International Publishing, 2016), pp. 501–518.
93. Supervisely: Unified OS for computer vision; https://supervisely.com/.
94. S. Katz, A. Tal, On the visibility of point clouds, *in 2015 IEEE International Conference on Computer Vision (ICCV)* (IEEE, 2015), pp. 1350–1358.
95. R. Mascaro, L. Teixeira, M. Chli, Volumetric instance-level semantic mapping via multi-view 2D-to-3D label diffusion. *IEEE Robot. Autom. Lett.* **7**, 3531–3538 (2022).
96. A. ten Pas, M. Gualtieri, K. Saenko, R. Platt, Grasp pose detection in point clouds. *Int. J. Robot. Res.* **36**, 1455–1473 (2017).
97. R. B. Rusu, "Semantic 3D object maps for everyday manipulation in human living environments," thesis, Technischen Universität München (2010).
98. M. Kazhdan, M. Bolitho, H. Hoppe, Poisson surface reconstruction, in *Proceedings of the Fourth Eurographics Symposium on Geometry Processing, SGP '06* (Eurographics Association, 2006), pp. 61–70.
99. M. Attene, A lightweight approach to repairing digitized polygon meshes. *The Visual Computer* **26**, 1393–1406 (2010).
100. M. Bacher, E. Whiting, B. Bickel, O. Sorkine-Hornung, Mass properties of triangulated solids and their derivatives, supplemental material for 'Spin-It: Optimizing moment of inertia for spinnable objects. *ACM Trans. Graph.* **33**, 1–10 (2014).
101. C. Rother, V. Kolmogorov, A. Blake, "GrabCut": Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* **23**, 309–314 (2004).
102. M. Waechter, N. Moehrle, M. Goesele, Let there be color! Large-scale texturing of 3D reconstructions, *in Computer Vision—ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars, Eds. (Lecture Notes in Computer Science, Springer International Publishing, 2014), pp. 836–850.
103. D. Cohen-Steiner, P. Alliez, M. Desbrun, Variational shape approximation. *ACM Trans. Graph.* **23**, 905–914 (2004).
104. R. B. Rusu, N. Blodow, M. Beetz, Fast point feature histograms (FPFH) for 3D registration, in *2009 IEEE International Conference on Robotics and Automation* (IEEE, 2009), pp. 3212–3217.
105. I. Cramb, *The Art of the Stonemason* (Alan C. Hood, 1992).
106. R. B. Rusu, S. Cousins, *3D is here: Point Cloud Library (PCL), in 2011 IEEE International Conference on Robotics and Automation*, 1–4 (IEEE, 2011).
107. S. Flöry, "Constrained matching of point clouds and surfaces," thesis, Vienna University of Technology, Vienna (2010).
108. E. Coumans, Bullet physics simulation, in *ACM SIGGRAPH 2015 Courses, SIGGRAPH '15* (ACM, 2015), p. 1.
109. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, O. Ronneberger, 3D U-Net: Learning dense volumetric segmentation from sparse annotation, in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016*, S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, W. Wells, Eds. (Lecture Notes in Computer Science, Springer International Publishing, 2016), pp. 424–432.
110. A. Wolny, L. Cerrone, A. Vijayan, R. Tofanelli, A. V. Barro, M. Louveaux, C. Wenzl, S. Strauss, D. Wilson-Sánchez, R. Lymbouridou, S. S. Steigleder, C. Pape, A. Bailoni, S. Duran-Nebreda, G. W. Bassel, J. U. Lohmann, M. Tsiantis, F. A. Hamprecht, K. Schneitz, A. Maizel, A. Kreshuk, Accurate and versatile 3D segmentation of plant tissues at cellular resolution. *eLife* **9**, e57613 (2020).
111. J. Nubert, E.Walther, S. M. K. Khattak, M. Hutter, *Learning-based Localizability Estimation for Robust LiDAR Localization* (2022).
112. G. Barill, N. G. Dickson, R. Schmidt, D. I. W. Levin, A. Jacobson, Fast winding numbers for soups and clouds. *ACM Transactions on Graphics* **37**, 1–12 (2018).
113. P. F. McCombie, J.-C. Morel, D. Garnier, J.-C. Morel, D. Garnier, *Drystone Retaining Walls: Design, Construction and Assessment* (CRC Press, 2015).
114. L. Alejano, M. Veiga, J. Taboada, M. Díez-Farto, Stability of granite drystone masonry retaining walls: I. Analytical design. *Géotechnique* **62**, 1013–1025 (2012).
115. R. L. Johns, M.Wermelinger, R. Mascaro, D. Jud, L. Vasey, I. Hurkxkens, M. Chli, F. Gramazio, M. Kohler, M. Hutter, *Dataset for the paper "A framework for robotic excavation and dry stone construction using on-site materials"* (2023); 10.5281/zenodo.10038881.
116. R. Mascaro, *Diffuser* (2023); https://github.com/VIS4ROB-lab/diffuser.
117. R. L. Johns, Ashlar: Mesh Properties for Stone Datasets (2023); https://github.com/gramaziokohler/ashlar.
118. E. Jelavic, Y. Berdou, D. Jud, S. Kerscher, M. Hutter, Terrain-Adaptive Planning and Control of Complex Motions for Walking Excavators, *in 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2020), pp. 2684–2691.
119. G. Wernet, C. Bauer, B. Steubing, J. Reinhard, E. Moreno-Ruiz, B. Weidema, The ecoinvent database version 3 (part I): Overview and methodology. *Int. J. Life Cycle Asses.* **21**, 1218–1230 (2016).
120. GOV.UK, "*Greenhouse gas reporting: conversion factors 2021*" (Tech. Rep., Department for Business, Energy & Industrial Strategy, 2021).
121. Y. H. Dong, L. Jaillon, P. Chu, C. S. Poon, Comparing carbon emissions of precast and cast-in-situ construction methods—A case study of high-rise private building. *Construct. Build Mater.* **99**, 39–53 (2015).
122. C. Jones, G. Hammond, "*ICE (Inventory of Carbon & Energy) Database*" (Tech. Rep. 3.0, Circular Ecology, 2019); https://circularecology.com/embodied-carbon-footprint-database.html.
123. M. Holian, J.-H. Pyeon, "*Analyzing the Potential of Hybrid and Electric Off-Road Equipment in Reducing Carbon Emissions from Construction Industries*" (Tech. Rep. 12-78, Mineta Transportation Institute, 2017).
124. H. S. Oh, U. Kim, G. Kang, J. K. Seo, H. R. Choi, Multi-axial force/torque sensor calibration method based on deep-learning. *IEEE Sensors J.* **18**, 5485–5496 (2018).
125. H. H. Le, J.-C. Morel, D. Garnier, P. F. McCombie, A review of methods for modelling drystone retaining walls, in *Proceedings of the Institution of Civil Engineers-Geotechnical Engineering* (Thomas Telford, 2014), pp. 262–269.
126. B. Zhao, J. Wang, 3D quantitative shape analysis on form, roundness, and compactness with µCT. *Powder Technol.* **291**, 29 (2016).
127. D. C. Froehlich, Mass angle of repose of open-graded rock riprap. *J. Irrig. Drain. Eng.* **137**, 454–461 (2011).
128. Y. Liu, N. Napp, Angle of repose in robotic construction (2022).
129. T. Zingg, *Beitrag zur Schotteranalyse*, Doctoral Thesis, ETH Zurich (1935).
130. H. Wadell, Volume, shape, and roundness of rock particles. *J. Geol.* **40**, 443–451 (1932).

Acknowledgments: We thank C. Girot for assistance and oversight with formulating the approach for robotic landscaping as well as V. Buff, V. Pajovic, and M. Kuhlmann-Joergensen for their assistance with operations in the retaining wall construction. **Funding:** This research was supported by the Swiss National Science Foundation through the National Centre of Competence in Digital Fabrication (NCCR dfab) and the European Space Agency (ESA). Testing area and building materials for the demonstrators were provided by Eberhard Bau AG. Funding for the load testing experiment was provided by the ETH Zurich Foundation Partnership Council for Sustainable and Digital Construction, together with Siemens and Geberit (grant number 2021-FS-211). **Author contributions:** R.L.J., M.W., R.M., D.J., M.C., F.G., M.K., and M.H. formulated the approach for mobile robotic assembly with found objects. D.J., I.H., M.H., F.G., and M.K. formulated the approach for robotic landscaping. R.M. developed the stone instance detection.

M.W. developed the grasp planner and stone picking and placement routines. M.W., R.M., D.J., and R.L.J. developed the stone digitization pipeline. R.L.J. developed the geometric planner. L.V. and M.W. developed the wall probing routine. D.J. developed the autonomous excavation controller for HEAP. I.H. integrated landscape design input and feedback for the excavation process. R.L.J. and M.W. developed the freestanding wall experiment. R.L.J. oversaw construction of the retaining wall, working together with M.W. and L.V. for operations. D.J. and I. H. oversaw the autonomous terrace excavation. R.L.J. drafted the text and figures, with relevant section contributions from M.W., R.M., D.J., I.H., and L.V. M.C., F.G., M.K., and M.H. advised the research, contributed to experiment design, and refined this manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions of this manuscript are included in the main text or the Supplementary Materials. Code is provided for diffusion-based instance segmentation (*116*) and stone form properties (*117*). The 3D stone models, the labeled dataset of candidate placements, and the candidate classification model (*115*) are also available.

Submitted 18 January 2023
Accepted 27 October 2023
Published 22 November 2023
10.1126/scirobotics.abp9758

Science Robotics

A framework for robotic excavation and dry stone construction using on-site materials

Ryan Luke Johns, Martin Wermelinger, Ruben Mascaro, Dominic Jud, Ilmar Hurkxkens, Lauren Vasey, Margarita Chli,
Fabio Gramazio, Matthias Kohler, and Marco Hutter

Sci. Robot. **8** (84), eabp9758. DOI: 10.1126/scirobotics.abp9758

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.