# Science Robotics

# Supplementary Materials for

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

Ryan Luke Johns et al.

Corresponding author: Ryan Luke Johns, rjohns@ethz.ch

*Sci. Robot.* **8**, eabp9758 (2023) DOI: 10.1126/scirobotics.abp9758

#### This PDF file includes:

Figs. S1 to S16 Sections S1 to S10 Tables S1 to S4 Algorithm S1 References (*118–130*)

# **Supplementary Materials**



Fig. S1. Automated terrace excavation. Digging (A) and dumping (B) are executed within reach of the cabin turn joint.



**Fig. S2. Macro micro manipulation.** Initial study of the use of a compound manipulator for filling voids with small stones. Refined stone mesh overlaid to show correlation with 3d model.



**Fig. S3. Candidate solution dataset.** (**A-F**) Samples of the automatically generated wall states used to produce the candidate dataset. (**G-J**) Examples from the hand-labelling interface, which directly remeshes the raw 3-channel SDFs using marching cubes. Candidate solutions are shown in white, together with existing context stones (grey) and the wireframe of the wall bounds. Accepted solutions (**G-H**) are well aligned to the target surface. Unacceptable solutions (**I-J**) may be near the desired surface, but with outward sloping regions that would be problematic for the stability of downstream solutions.



**Fig. S4. Freestanding wall stone refinement** (**A**) Front and (**B**) side elevations of the 3D wall model with stone meshes, and the excavator-generated laser map used for stone pose refinement.



**Fig. S5.** As-built landscape. Laser scan for one of four completed terraces. (A) Stepped geometry with dip angle shading (B) Absolute surface error to designed target mesh.



**Fig. S6. Earth platform for elevated construction.** (**A**) For a more stable reach to areas above 4 m, and due to constraints in site-logistics that prevented continuous construction from the back side of the wall, the soil that was excavated to form the landscape was used as a temporary platform for wall assembly. (**B**) Near the end of construction, the platform was removed by a large manually-operated excavator. (**C-F**) Construction stages: (**C**) Foundation-level assembly, (**D-E**) elevated platform assembly, (**F**) platform removed—construction is completed from the back side of the wall.



**Fig. S7. Sample stone grasps.** (A-B) Picking for placement. (C-D) Transferring. (E-H) Placement in the wall.

# S1 Nomenclature

- $\mathcal{S}$  set of candidate stone placements
- $\mathcal{C}$  set of mutually compatible candidate stone placements to be placed
- c candidate stone placement
- v candidate fitness
- s stone identifier
- O stone pose
- $\mathbf{c}_s$  stone centroid
- $k_s$  quantity of stone meta-faces
- $m_s$  stone mass in kilograms
- $A_s$  stone surface area in square meters
- $\mathcal{M}_{(\cdot)}\,$  mesh, consisting of a set of vertices  $\mathcal{V}_{(\cdot)}$  and faces  $\mathcal{F}_{(\cdot)}$
- $\mathcal{P}, \mathcal{Q}$  sets of points in  $\mathbb{R}^3$ 
  - $\mathcal{T}$  set of coordinate frames
- $\underline{\mathfrak{T}}$  coordinate frame with position in  $\mathbb{R}^3$  and orientation in SO(3)
- d distance in meters
- f contact force
- g gravitational acceleration
- $\mu$  friction coefficient
- $\omega$  angular velocity
- $\psi$  foot point weight
- $\mathcal{E}$  error
- $\left\lceil \cdot \right\rceil$  ceiling
- $|\cdot|$  l<sub>1</sub> norm or set cardinality
- $||\cdot||$   $l_2$  norm

#### S2 HEAP

These experiments are developed using HEAP (Hydraulic Excavator for an Autonomous Purpose), a modified 12-ton Menzi Muck M545 excavator equipped with force-controllable hydraulic cylinders with integrated pressure transducers and servo valves (*66*, *67*). The base excavator system is specifically engineered for work on unevent errain, with an undercarriage supported by highly maneuverable, wheeled outriggers—a morphology recognized in the construction trades as a "walking" or "spider" excavator. Our incorporation of pressure- and IMU-feedback allows the use of these wheeled-legs for automated load distribution and chassis balancing on the irregular ground conditions typical to construction environments. In this work, we generally separate leg and arm movements: digging, picking, and placing trajectories are executed while the base is mostly stationary, while driving occurs between these actions. However, previous work has also explored the potential of whole-body motions that combine all five limbs toward stepping and balancing maneuvers in especially complex terrain (*118*).

Excavator localization is achieved using a Leica iCON iXE3 with two cabin-mounted GNSS receivers, and GNSS RTK corrections that are received over the internet from permanently installed base stations. Full 6-dof pose estimation is made possible with the addition of cabinand chassis-mounted SBG Ellipse2-A IMUs. A combination of cylinder-mounted draw-wire encoders and link-mounted IMUs provide the kinematic joint states. The addition of a Rototilt R4 beyond the standard shovel pitch joint provides roll, pitch, and infinite yaw to the end-ofarm shovel or gripper. A total of 23 individually-controllable axes allow for the manipulation of objects weighing up to 3000 kg, with a vertical reach up to 9 m.

Exteroceptive sensing is made possible with two roof-mounted LiDAR sensors (Velodyne Puck VLP-16, Livox Mid-70). An additional arm-mounted Livox Mid-70 provides an overhead view, and the ability to reach over and accumulate points on the back side of objects. A forward-facing RGB camera (Flir Blackfly) is mounted at the base of the cabin for capturing stone texture

information, and a roof-mounted camera is is employed for stone segmentation.

The load testing process utilizes a custom 6-axis force-torque sensor mounted between the end effector and Rototilt baseplate. For more information on this sensor, refer to supplement S8.

## S3 Retaining wall geometry

The built retaining wall was designed with variable height and incline, accommodating pedestrian ramps behind the structure, and a sloped foundation that allows water to drain into a culvert buried near the center of the wall. The height at any point along the wall is given by h(x):

$$h(x) = f(x) - g(x) \tag{15}$$

where the piecewise linear functions f(x) (Eq. 16) and g(x) (Eq. 17) provide the relative elevations for the top of the wall and the foundation, respectively (Fig. S8).

$$f(x) = \begin{cases} 0.032x + 6 & \text{if } -37.5 \le x < 0\\ -0.073x + 6 & \text{if } 0 \le x \le 28 \end{cases} g(x) = \begin{cases} -0.044x + 0.5 & \text{if } -37.5 \le x < -9\\ -0.2x - 0.9 & \text{if } -9 \le x < -4.5\\ 0.16x + 0.07 & \text{if } -4.5 \le x \le 28\\ (17) \end{cases}$$

The rake angle of the wall is separated into two regions: Above the crease line k(x)(Eq. 18), it is a fixed constant of 11.3°. Below k(x), there is a variable rake angle  $\theta(x)$ , also provided as a function of the x axis position (Eq. 19):

$$k(x) = \begin{cases} 0.071x + 4.3 & \text{if } -37.5 \le x < 0\\ -0.056x + 4.3 & \text{if } 0 \le x \le 28 \end{cases}$$
(18)

$$\theta(x) = -0.0004x^3 - 0.015x^2 - 0.25x + 39 \tag{19}$$



**Fig. S8. Retaining wall height and incline.** Front elevation of the robotically-constructed retaining wall (above) with section cuts at the indicated x-axis values (below).

#### S4 Sustainability assessment

**Embodied carbon of robotic stone walls** In order to benchmark the sustainability aspects of robotic dry stone masonry, we estimate the environmental impact by considering (i) the cradle-to-gate energy of the raw materials, (ii) the transport of material to the site, and (iii) the diesel used by the excavator during construction.

(i) Our built structures consist of a combination of quarry stones, erratics unearthed in nearby construction sites, and waste concrete. We compute the embodied carbon associated with the quarry extraction process in table S1. For upcycled erratics and concrete debris, we consider only the impact of transportation.

**Table S1.** Life Cycle Inventory Assessment—embodied carbon equivalent ( $CO_2e$ ) of gneiss boulders extracted in southern Germany with an average mass of 1 t. Quarry-reported values (per tonne) were converted using a density of 2.8 t/m<sup>3</sup>. Method from Ioannidou et al. (*30*)

Resource, unit	Quarry (units/m <sup>3</sup> )	kgCO <sub>2</sub> e/unit	Ecoinvent Process*	kgCO <sub>2</sub> e/m <sup>3</sup>	
Blasting, kg	$3.5 \times 10^{-1}$	4.31	Blasting (RER)	1.51	
Diesel, MJ	$6.75 \times 10^{1}$	$9.1 \times 10^{-2}$	Diesel, burned in building machine (GLO)	6.14	
Electricity, kWh	$6.4 \times 10^{-1}$	$3.9 \times 10^{-2}$	Electricity, medium voltage, at grid (CH)	$2.5 \times 10^{-2}$	
Drill bits, kg	$2.8 \times 10^{-1}$	$3.8 \times 10^{-1}$	Steel, low-alloyed, at plant (CH)	$1 \times 10^{-1}$	
Water, m <sup>3</sup>	$1.4 \times 10^{-3}$	$8.1 \times 10^{-2}$	Tap water, at user (CH)	$1.13 \times 10^{-4}$	+
				7.78	

#### \*IPCC2021 GWP100a (119)

(ii) All materials are delivered to the construction site by diesel trucks (>17 t HGV) dispatched within 40 km. For delivery, we consider the 100 % laden conversion factors for a rigid axle HGV of 0.12 kgCO<sub>2</sub>e/tkm (*120*), or 0.34 kgCO<sub>2</sub>e/m<sup>3</sup>km considering a material density of 2.8 t/m<sup>3</sup>. For the unladen return trip, we distribute the conversion factor of 0.77 kgCO<sub>2</sub>e/km across a typical single delivery payload of 7 m<sup>3</sup> to yield a contribution of 0.11 kgCO<sub>2</sub>e/m<sup>3</sup>km. Combining the impact of the delivery and return transport yields a conversion factor of

$$0.446 \,\mathrm{kgCO_2 e/m^3 km}$$
 (20)

(iii) HEAP's diesel consumption during wall construction is within the range of 6.2-9.5 L/m<sup>3</sup>

of placed stone. The upper bound is used in our calculations, and represents approximate fuel consumption over a three month period. This estimate factors in a considerable amount of idling, early stage troubleshooting, and peripheral tasks—and should thus be considered as a generous estimate for the process. The lower bound represents the consumption rate in the last stages of construction and can be considered as a more realistic value for future deployment of the process, with the potential to be improved further with more process-integrated engine speed and idling control. Multiplying a conversion factor of 36.1 MJ/L (*120*) with the diesel conversion factor from table S1 yields a carbon equivalent conversion factor of

$$3.29 \,\mathrm{kgCO_2 e/L/m^3}$$
 (21)

resulting in a carbon equivalent contribution of  $20.4-31.2 \text{ kgCO}_2\text{e/m}^3$  for our process and machine. Fuel usage of the robotic excavator is thus the largest contributor to the overall carbon cost of our built structures.

Combining the conversion factors 20 and 21 with the total from table S1 provides a generalized equation f(p, d, l) for the carbon impact of our building process per cubic meter of placed material (kgCO<sub>2</sub>e/m<sup>3</sup>)—given the fractional percentage p of quarry-extracted stones in the structure, the average one-way delivery distance d of materials to the site, and the typical excavator diesel consumption l in liters per placed cubic meter

$$f(p, d, l) = 7.78p + 0.446d + 3.29l$$
<sup>(22)</sup>

**Evaluation of prototypes** The freestanding wall is built from a mixture of reclaimed concrete and quarried gneiss boulders (p = 0.76) with a mean transport distance d = 27 km. Considering our upper bound fuel consumption, we have

$$f(0.76, 27, 9.5) = 49.2 \,\mathrm{kgCO_2 e/m^3}$$
 (23)

The retaining wall is built from a mixture of reclaimed concrete, erratics unearthed on nearby construction sites, and quarried gneiss (p = 0.9) with a mean transport distance d = 35 km

$$f(0.9, 35, 9.5) = 53.9 \,\mathrm{kgCO_2e/m^3}$$
 (24)

**Comparison to concrete** We can estimate the comparative carbon contribution of our building method by considering an equivalent structure built with reinforced concrete. For example, we compare our stone retaining wall to a reinforced concrete wall with a thickness of 30 cm, and single-face surface area equal to that of our built structure  $(313 \text{ m}^2)$  with a total volume of 94 m<sup>3</sup>. Our stone structure is substantially thicker (average 1.8 m), but with a 40 % void ratio and and a placed material volume of 339 m<sup>3</sup> (computed from the scanned mesh models of the placed stones). By comparing the volume of these equivalent structures, we obtain the material performance equivalence

$$1 \text{ m}^3$$
 reinforced concrete  $\equiv 3.6 \text{ m}^3$  dry stone (25)

and the equivalent embodied carbon ratio  $r(e_d, e_c)$  of our method as

$$r(e_d, e_c) = 3.6 * \frac{e_d}{e_c}$$
 (26)

Where  $e_d$  and  $e_c$  are the embodied carbon equivalents per cubic meter of dry stone and concrete, respectively.

The embodied carbon of cast-in-situ reinforced concrete is dependent on a combination of regional factors and the composition of the concrete mix, with one recent estimate as high as 770 kgCO<sub>2</sub>e/m<sup>3</sup> for cast-in-place walls in an urban environment with a high carbon electrical mix, inclusive of formwork preparation (*121*). Using this figure as  $e_c$  and taking  $e_d$  from equation 24 would indicate, for example, that our method contains 25.2 % of the embodied carbon

for an equivalent structure in reinforced concrete. For a more conservative comparison, we consider only the raw material embodied carbon of  $330 \text{ kgCO}_2\text{e/m}^3$  for 32/40 MPa concrete (122) as  $e_c$ , indicating that our method would contribute 58.8 % of the CO<sub>2</sub> of an equivalent structure in concrete. For reference, substituting our more realistic fuel consumption of 6.2 L in equation 24 would reduce this value to 46.9 %, using non-quarry materials sourced within 35 km would reduce it to 39.3 %, and using stones found directly on site would reduce it further still to 22.2 %.

At present, the largest CO<sub>2</sub> contributor is the excavator's diesel consumption—accounting for 47–100 % of process emissions, where the upper bound again represents theoretical construction on a site with abundant boulders that require no additional emissions for extraction and transportation. As such, the process has the potential to see drastic improvements in carbon efficiency if migrated to an electrified ex cavator. Electrification of hydraulic construction machinery has been demonstrated to reduce greenhouse gas emissions by 79 % based on the electrical grid of New York State (*123*). Considering that the emissions factor of the Swiss electrical grid is 15 % of the value reported for New York (*119*), this could represent a further reduction of 46–97 % in total process emissions in our context.

#### S5 Lateral Bonding

Double-faced stone walls are stabilized in part by having sufficient lateral bonding between the two outer wall faces. In this supplement, we outline the key characteristics that describe this lateral bonding, in order to provide benchmarks of our own structures—such that they can be compared with traditional methods and future work.

Throughstones in traditional masonry In traditional dry stone masonry, lateral bonding is partially achieved by the inclusion of throughstones (stones which span both sides of the wall). While our literature review has not revealed a suggested ratio of throughstones, we can estimate this based on the recommended use of one throughstone per linear meter of manually constructed walls below 1.5 m in height (40, p. 285): Given a section of wall length (L = 1 m), the batter percentage in decimal form ( $b_p = 0.1$ ), the coping or apex thickness ( $w_c = 0.3$ ), height ( $h \le 1.5$  m), stone density ( $\rho_s = 2800 \text{ kg/m}^3$ ), void ratio ( $r_v = 0.33$ ), and average stone mass ( $m_s = 23 \text{ kg}$ ) (74), we can estimate the quantity of stones ( $k_s$ ) in a typical wall section using equation 27.

$$k_{s} = \frac{L(w_{c}h + h^{2}b_{p})(1 - r_{v})\varrho_{s}}{m_{s}}$$
(27)

Manually constructed walls with heights of 1-1.5 m should thus contain upward of 33-55 stones per linear meter. With those ranges in mind, throughstones can be relatively infrequent, representing only about 1.8 to 3% of total stones.

**Throughstones in our method** For evaluating our method along these lines, we define throughstones as stones that are within some threshold distance  $d_T$  to both opposing wall faces ( $d_T = 5\%$  of the wall thickness). While we do not explicitly seed throughstone candidates, they typically occur when a candidate stone is dimensioned such that it has foot points on both opposing wall faces. Considering our candidate dataset of approximately nine thousand automaticallyplaced stones, such throughstones naturally represent 1.1 % of solutions before classification. However, these throughstones are 1.6 times more likely to be accepted compared to nonthroughstones in this hand-labelled dataset (43.9 % vs. 28.2 %), and thus consist of 1.8 % of accepted solutions. Note that because the dataset is produced from automatically generated solutions without any intermediate classification, there are many more candidate solutions in lower-wall layers where the target wall geometry is thicker, and thus the throughstone distribution of the dataset might not match that of a fully-built wall.

We provide two values for each of our built structures, considering that masonry guidelines generally distinguish between throughstones (fully surrounded by other stones) and capstones (top layer stones which also typically span the full wall width): the freestanding wall has 3 stones which span both sides of the wall (2.75%), of which 2 of these are below the top layer (1.83%). The retaining wall has 26 stones which span both sides of the wall (2.8%), of which 18 of these are below the top layer (1.9%). While these throughstone distributions are similar to those suggested in masonry guidelines, future work would be required to explicitly inform their spatial distribution as suggested in the literature, or to specifically seed candidates in areas where their longest axial dimension corresponds to the wall thickness.

It is important to note that traditional throughstone guidelines are targeted toward conventional, manually-constructed walls, with a distinct inner core of smaller hearting stones that separates the outer-two wall-faces or wythes. Throughstones cannot always be present (for example, where there are insufficient stones whose longest axial length corresponds to the wall thickness), and it is generally acceptable in these cases to create stable structures through "solid masonry, i.e. without a distinct core of backing stones, and in a manner that bonds every course continuously to the next, not just longitudinally, as usual, but also traversely" (40, p. 286). By this definition (and through the exclusion of small filling stones), our work resembles such solid masonry structures, and we encourage this bidirectional bonding by seeding candidate solutions at stone intersections, and increasing the weight of stone-based foot points during registration.

**Header ratio** An additional factor toward ensuring lateral stability is the ratio of headers (stones whose long axial dimension runs perpendicular to the wall face) to stretchers (stones whose long axial dimension runs parallel to the wall face), where German standards for wellcrafted stone walls suggests a minimum ratio of 1:2 (33.3 % headers) (40, p. 281). Broadly defined, the difference between a header and a stretcher can be determined solely based on the aspect ratio of the placed stones: by measuring the stone width ( $d_X$ ) along the horizontal wall tangent, and the stone depth ( $d_Y$ ) along the reference-wall's closest XY-projected surface normal, a header is simply any placed stone where

$$\frac{d_X}{d_Y} < 1.$$

By this definition, our freestanding wall and retaining wall consist of 43.1 % and 57.2 % headers, respectively. This distribution is highly sensitive to stone properties such as elongation, and to the wall orientation, batter angle, and thickness. For example, by redefining  $d_Y$  as being measured along the true surface normal of the reference wall, the inclined retaining wall would instead contain only 48.9 % headers. For rough comparison, the S-curved structure produced in our previous work (41) contains only 30 % headers by either definition, and has no throughstones.

**Bed Depth** Beyond their own width to depth ratio, it is also important that headers extend far enough with respect to the total wall thickness, in order to sufficiently bond with the stones on the opposing wythe. As such, we use the portion of  $d_Y$  inside the wall to describe the bed depth, or the distance a given stone extends into the wall (Fig. S9). Table S2 lists the average bed depth

Set	Average bed depth (%)	Upper-third bed depth cutoff (%)
Candidate Dataset	57	62.2
Freestanding Wall	55.9	63
Retaining Wall	55	62
S-curved wall (41)	51	60

Table S2. Bed depth as a percentage of wall thickness.

of the stones in our built structures and candidate dataset, as a percentage of the wall thickness at each measurement location. An increased average bed depth percentage implies that rocks extend further toward the opposing wall face, and thus there is more bonding between wythes. To relate to the suggested header ratio described above, we also list the bed depth percentage above which 33.3 % of stones extend.



Fig. S9. Bed depth. The bed depth percentage is measured in relationship to the wall thickness, where the bed depth is the portion of  $d_Y$  that falls fully inside the wall.

**Contacts** To evaluate the bonding of dry stone structures, we can also consider the average number of stones that are in contact with each stone (masonry guidelines commonly suggest resting each stone on at least two others). In Figure S10, we outline the average stone contacts for the walls presented in this work, and in our previous work. To reduce the wall-size bias

coming from edge-stones with fewer contacts, we exclude the outer-edge and foundation stones for this analysis. Additionally, given that there is some error in the stone refinement, and that nearby contacts can also improve the robustness of walls as objects shift and settle, we plot the average number of contacting stones at multiple contact thresholds, between 0 and 0.15 m.



**Fig. S10.** Average stone contacts. Left) Average number of unique stone instances that are within the contact threshold of any given stone, in any direction. Right) Average number of unique 'supporting' stones that are within the specified contact threshold of any point on the stone surface that has a downward-pointing surface normal.

## S6 Candidate seed registration

For the geometric planning of stone placements, it is desirable to have seed poses that closely resemble their stable and registered states. To aid in understanding the relationship between the seed candidates and the final accepted poses, we describe the translations during the registration and simulation stages in figure S11. As our approach emphasises fast candidate seeding using few geometric attributes, intersections and/or large gaps with the existing context are common for these seed poses. We observe typical translations of  $(0.41 \pm 0.32)$  m between the seed pose and the final accepted stable pose, and rotations of  $(46.1 \pm 15.1)^{\circ}$ . The majority of iterations and motion occur in the first round of TICR iterations, with the successive stages of simulation and registration contributing comparatively small pose refinements.



**Fig. S11. Candidate registration stages.** Distribution of steps-to-convergence and per-axis translation for 1,000 seeded candidate stone poses. Following the initial TICR pass, we observe reduced translation and and iteration steps for subsequent rigid-body simulation and TICR passes.

#### S7 Subset selection

In order to find a mutually-compatible subset of suitable stone solutions ( $C \subseteq S$ ), we first rearrange the candidates  $c \in S$  into a set of sets S composed of ordered subsets  $\{\mathcal{R}_i, \dots, \mathcal{R}_I\}$ unique to each individual stone, where I is the number of discrete stone identifiers among candidates in S. Notably, we augment each set  $\mathcal{R}_i$  with an additional null-pose candidate  $(s_i, \emptyset, v)$ , with fitness v = 0, and subsequently sort these candidates such that  $v_j > v_{j+1}$ . This formulation of S allows us to reframe the selection of our optimal non-colliding subset C: from each  $\mathcal{R}_i$ , we must select exactly one candidate, while maximizing the sum of selected candidate fitnesses. Among the available solutions for each stone id from S, we have the additional option of choosing a zero-weighted non-placement. Thus, we finally obtain  $\mathcal{C} = \text{SelectNext}(S, \emptyset)$ , as defined in algorithm S1.

Algorithm S1 SelectNext (S,  $C_{partial}$ ) 1:  $C_{best} \leftarrow \emptyset$  $\triangleright C_{best}$  to be populated with our best selection 2:  $i \leftarrow |\mathcal{C}_{partial}|$ 3: for each  $c \in \mathcal{R}_i$  do if CollisionFree ( $c, \mathcal{C}_{partial}$ ) then 4:  $\mathcal{C}_{select} \leftarrow \mathcal{C}_{partial} \cup \{c\}$ 5: if  $|\mathcal{C}_{select}| < |S|$  and not <code>Timeout</code> then 6:  $C_{select} \leftarrow \texttt{SelectNext}(S, C_{select})$ 7:  $\triangleright$  recurse end if 8: if Score ( $C_{select}$ ) > Score ( $C_{best}$ ) then 9:  $\mathcal{C}_{best} \leftarrow \mathcal{C}_{select}$ 10: 11: end if end if 12: 13: end for 14: return  $C_{best}$ 

Where the subfunctions are defined as follows:

Score  $(\mathcal{X})$ : Returns the sum of fitness values  $v_i$  for the set of candidate solutions  $c_i \in \mathcal{X}$ , and  $-\infty$  where  $\mathcal{X} = \emptyset$ . CollisionFree  $(x, \mathcal{X})$ : Checks the candidate placement x for compatibility with each member in the set  $\mathcal{X}$  of previously selected candidates, and returns True if there are no collisions. Collisions are looked up in a precomputed  $|\mathcal{S}| \times |\mathcal{S}|$  matrix, where trivial collisions are assigned to candidates sharing the same stone identifier  $s_i$  (the same stone cannot be put in the wall twice). Otherwise, non-collisions are assigned to any pair where either candidate represents a non-placement. For all remaining candidate pairings, we first check for axis-aligned bounding box (AABB) overlap, and subsequently compute the collision state between candidate meshes using the fast winding number (112) where AABB collisions exist.

Timeout: Returns True if a predesignated time limit has elapsed. In cases where many hundreds of candidates are distributed among many stones, a full graph search demands excessive computation time. Considering that we begin our search with the highest-value candidates, we simply accept the best performing combination achieved thus far in cases where the timeout threshold has been surpassed.

## S8 Rokubi Mega

For monitoring the forces during the structural stability assessment, a custom 6-axis Bota Systems force-torque sensor is used (see Fig. S12). The sensor is mounted between the Rototilt baseplate and the tool (shovel or gripper) and measures the forces and torques transmitted through it. This allows the measurement of all force and torque components directly at the end-effector, instead of relying on inaccurate estimates through cylinder pressure. The measuring range of the strain gauge based sensor is listed in Table S3. As an hydraulic excavator can exert forces considerably higher than the measuring range, especially under compression, the sensor has mechanical end-stops to prevent damage to the sensing element.



Fig. S12. Custom 6-axis force-torque sensor from Bota Systems

In order to physically calibrate the sensor signal, a reference signal has to be generated.

Table S3.Measuring ranges.

Axis	Range	Unit
Fx	25000	Ν
Fy	25000	Ν
Fz	50000	Ν
Mx	35000	Nm
My	35000	Nm
Mz	50000	Nm

As Oh et al. suggest (124), for this task the excavator arm itself can be used to collect the data. A reference weight is added to the end-effector, and the sensor data is collected using gravitational effects on that weight. The load applied is calculated analytically from the shape of the trajectory, and linear regression is used to obtain the calibration. This methods allows for recording a rich excitement of single-axis and combined loadings on the sensor that correspond to the exerted forces and torques during use.

# **S9** Stability probing

Given the heterogeneity of building elements, there is no established practice for ensuring the local stability and safety of robotically assembled stone structures. As such, we developed a probing routine using a custom force-torque sensor in order to assess the local stability of individual stones at the end of construction. We apply a force  $\mathbf{f}_p$  at the contact point made by approaching each stone centroid  $\mathbf{c}_s$  perpendicular to the wall bounds  $\mathcal{M}_s^{1,-1}$ , increasing the magnitude of  $\mathbf{f}_p$  linearly with the distance from the top of the wall (125). For the built retaining wall, we applied maximal testing forces of 18 kN to a subset of placed stones. The testing setup is presented in Figure S13, where the orientation of the closed gripper is also parameterized to minimize eccentric gripper-to-stone collisions except at the specified contact point.



Fig. S13. Robotic stability probing (A-C) Parameterization of our FTS stability testing method: (A) Stone centroid  $c_s$  (B) contact point determined from the approach vector and stone centroid (C) gripper contact pose, oriented to avoid unintended collisions. (D-E) Probing a surface stone with the custom force-torque sensor mounted between the gripper and rototilt.

This probing pipeline was first deployed to evaluate the stability of elements within the retaining wall, at the end of construction. As the first permanent and publicly accessible structure built using our method, these measurements supplemented the responsible civil engineer's required inspection, and helped give an understanding of the final wall's stability.

Our method best checks for a single-stone's resistance to sliding (the primary load case for soil retaining walls), rather than overturning due to eccentric torques. For an inverse understanding of the structural mechanics, we are first interested in being able to classify the stones according to whether they slipped or not, and for the case of slippage, to extract the applied force at which slipping commenced. While future work will correlate these findings toward an evaluation of the stability of such stone structures, our initial findings demonstrate the potential of this method to provide data that could reveal if a given stone slipped under the specified load case. Figure S14 shows plots of stable and unstable behavior under the applied loads. While at present this classification is manual, and requires visual monitoring of the gripper (the data is rendered inconclusive if the gripper slips off the stone during probing), these labels could be automated in future experiments.



Fig. S14. Force torque sensor readings Sample readings from the force torque sensor, indicating the x,y,z components and magnitude of  $\mathbf{f}_p$ . Plots (A-C) correspond to stones that did not slip during the applied load, while (D) illustrates a stone that slipped under an applied load case around 10 kN.

#### S10 Geometric properties of nonstandard construction materials

**Building materials overview** Using the combined dataset of scanned stones from both the freestanding and retaining wall sites, the typical object oriented bounding box dimensions of our building elements are  $d_L = (1.23 \pm 0.28)$  m,  $d_I = (0.90 \pm 0.17)$  m, and  $d_S = (0.67 \pm 0.16)$  m. With intermediate axial lengths between 0.38 and 1.66 m, these stones are classified in the range of fine to coarse boulders on the modified Udden-Wentworth grain-size scale (15). They have a mean surface area of  $(3.1 \pm 0.9)$  m<sup>2</sup>, a mean volume of  $(0.356 \pm 0.156)$  m<sup>3</sup>, and mean mass of (997 ± 438) kg. For these mass estimates and in our simulations, we assume a general material density of 2800 kg/m<sup>3</sup>. From weighing eight scanned stones on an industrial scale, this estimated density has a maximal error of 12 % (0.04 ± 0.07) for our gneiss and erratics. Concrete debris has a typical density closer to 2400 kg/m<sup>3</sup>, however as of yet we neither automatically classify the material, or consider this difference in our simulations.

**Shape properties** The assembly of dry stone structures from nonstandard materials becomes more challenging with increased geometric heterogeneity—intuitively it is easier to build a wall using uniform bricks or ashlars than it is with randomly-sized and irregularly-shaped rubble. To facilitate an understanding of the materials used in our work, and to aid in future related research, we outline their properties in Table S4. We additionally plot the distribution of these properties in Figure S15, including the coefficients (a,b,c) for the gaussian function

$$g(x) = ae^{-\frac{(x-b)^2}{2c^2}}$$

and the coefficient of d etermination  $R^2$ , where r elevant. This information c an b e u sed, for example, toward the procedural generation of larger model datasets with distributions that correlate to those found in the field (*126*). We additionally provide the full dataset of 1,100 scanned 3D models obtained during our construction process (*115*). Although traditional form metrics such as elongation, flatness, and roundness have been demonstrated to influence the packing density (126) and angle of repose (127, 128) of aggregated stones, existing measures are not well suited to describe the ability of materials to be easily stacked in construction. Flatness and elongation, for example, are more descriptive of the object bounding box, and do not express the geometry contained within (for example, they do not differentiate between a sphere and a cube). The 3D rectangularity measure is useful in describing how an object fills its bounding box, but this ratio can be highly influenced by slight protrusions in otherwise flat, rectangular objects. Similarly, a flat disk can be easily stacked, but exhibits low rectangularity due to the circular faces.

In order to provide a more descriptive metric for stone construction, we introduce 'ashlarness' in Table S4 to describe the resemblance of an arbitrary mesh to a stone that has been reshaped specifically for use in masonry walls (an a shlar). This descriptor leverages the metafaces (proxy regions) produced by iterative variational shape approximation (*103*), and prioritizes stones that have a pair of broad, parallel, and similarly sized meta-faces. Ashlarness values fall within the range of 0 to 1, where a thin wafer-like shape represents the upper bound. Although further work is required to quantify the relationship between this new metric and the ease of robotic construction with irregular objects, a visual analysis of the ashlarness value of stones within our dataset indicates an improved correspondence between this metric and the stackability of stones (Fig.S16) (*115*).

**Arbitrary materials** Although our robotic dry stone process facilitates construction with geometrically heterogeneous building elements, we assume that the building material is stone or concrete with an approximately equivalent density and performance under compressive loads. Irregularities among these materials, however, can include protruding steel reinforcement in building debris, fragile cleavage planes in sedimentary or foliated metamorphic rocks, or thin



**Fig. S15. Distributions of stone properties.** Probability distributions of the properties outlined in table S4 for the combined dataset of scanned stones.



**Fig. S16.** Ashlarness samples from the scanned dataset. Stone mesh instances with indicated ashlarness values. The corresponding flatness, e longation, 3 D r ectangularity, and convexity measures are also provided for reference.

**Table S4. Stone form factors.** Form factors for the dataset of scanned stones and concrete debris used in construction. For each object,  $d_S, d_I, d_L$  are the dimensions of the minimum object-oriented bounding box (OOBB), where  $d_S < d_I < d_L$ .

Index	Formula	Range	Dataset mean $\pm$ SD	Description
Volume (V)		$0 \text{ to } \infty$	$(0.356 \pm 0.156) \mathrm{m^3}$	Volume computed from the 3D mesh (100)
Mass	2800 V	$0$ to $\infty$	$(997 \pm 438)  \text{kg}$	Assumed density of 2.8 t/m <sup>3</sup>
Length $(d_L)$		$0$ to $\infty$	$(1.23 \pm 0.28)\mathrm{m}$	Longest axis of OOBB
Width $(d_I)$		$0$ to $\infty$	$(0.90 \pm 0.17)\mathrm{m}$	Intermediate axis of OOBB
Height $(d_S)$		$0$ to $\infty$	$(0.67 \pm 0.16)\mathrm{m}$	Shortest axis of OOBB
Elongation	$d_I/d_L$	0-1	$0.75\pm0.15$	Primary aspect ratio (129)
Flatness	$d_S/d_I$	0-1	$0.75\pm0.15$	Secondary aspect ratio (129)
Convexity	$V/V_c$	0-1	$0.87\pm0.04$	Ratio of object volume to the volume of the convex hull ( $V_c$ ) (126)
Sphericity	$\frac{\pi^{\frac{1}{3}}(6V)^{\frac{2}{3}}}{A_s}$	0-1	$0.81\pm0.05$	Ratio of the surface area of a sphere with volume $V$ to the surface area of the stone $A_s \ (130)$
Inscribed Radius $(R_{MIS})$		$0$ to $\infty$	$0.28 \pm 0.06$	Radius of the maximal inscribed sphere
Roundness	$\frac{1}{R_{MIS}} \left( \frac{\sum \left( \frac{A_n}{\sqrt{\kappa_n^G}} \right)}{\sum (A_n)} \right)$	0-1	$0.62 \pm 0.04$	Ratio of the mean corner radius $R_G$ to the radius of the maximal inscribed sphere $R_{MIS}$ (130). $R_G$ is computed using the gaussian curvature ( $K_n^G$ ) of all vertices with a curvature greater than that of the maximal inscribed sphere, weighted by the areas of each vertex neighborhood ( $A_n$ ) (126)
Rectangularity	$V/V_{oobb}$	0-1	$0.47\pm0.07$	Ratio of object volume to volume of OOBB ( $V_{oobb}$ )
Ashlarness	$\arg\max_x \frac{1}{3} \Big( \frac{\min\{A_i,A_x\}}{\max\{A_i,A_x\}} + \frac{A_i + A_x}{A_s} + (\mathbf{n_i} \cdot \mathbf{n_x})^2 \Big)$	0-1	0.67 ± 0.09	Approximate resemblance to a dressed stone, computed from the variational shape approximation (VSA) with a maximal $\mathcal{L}^{2,1}$ error metric (103) of 0.26. $A_i$ and $\mathbf{n}_i$ are respectively the area and proxy normal of the meta-face whose area weighted proxy normal is most aligned with the normal of the largest face of the OOBB. $A_x$ and $\mathbf{n}_x$ are the area and proxy normal of the meta-face that provides the highest resulting ashlarness value, where $\mathbf{n}_i \cdot \mathbf{n}_x < 0$
Meta-faces		2-25	9.2 ± 4.7	The minimum number of approximate face regions (meta-faces) computed using iterative VSA with a maximal $\mathcal{L}^{2,1}$ error metric of 0.26. We use an arbitrary fixed upper bound of 25 meta-faces

regions that can easily shear under load. Such features are not prevalent in our dataset, however our model notably does not yet account for these material-based failure modes, and treats objects as entirely rigid and uniform in simulation. The ability to recognize such material attributes, or to plan and adapt to a wider range of arbitrary material distributions is a largely unsolved problem in robotic construction (70), and future work should incorporate methods of estimating anomalous object properties beyond geometry to ensure process robustness in diverse environments.