

Autonomous Dry Stone

On-Site Planning and Assembly of Stone Walls with a Robotic Excavator

Ryan Luke Johns¹  · Martin Wermelinger²  · Ruben Mascaro³  · Dominic Jud²  · Fabio Gramazio¹  · Matthias Kohler¹  · Margarita Chli³  · Marco Hutter² 

Received: date / Accepted: date

Abstract On-site robotic construction not only has the potential to enable architectural assemblies that exceed the size and complexity practical with laboratory-based prefabrication methods, but also offers the opportunity to leverage context-specific, locally-sourced materials that are inexpensive, abundant, and low in embodied energy. We introduce a process for constructing dry stone walls in situ, facilitated by a customized autonomous hydraulic excavator. Cabin mounted LiDAR sensors provide for terrain mapping, stone localization and digitization, and a planning algorithm determines the placement position of each stone. As the properties of the materials are unknown at the beginning of construction, and because error propagation can hinder the efficacy of pre-planned assemblies with non-uniform components, the structure is planned on-the-fly: the desired position of each stone is computed immediately before it is placed, and any settling or unexpected deviations are accounted for. We present the first result of this geometric- and motion-planning process: a 3-meter tall wall composed of 40 stones with an average weight of 760 kg.

Keywords On Site Robotics · Dry Stone Walls · Adaptive Assembly

Ryan Luke Johns
E-mail: johns@arch.ethz.ch

Martin Wermelinger
E-mail: martiwer@mavt.ethz.ch

¹Gramazio Kohler Research
²Robotic Systems Lab
³Vision for Robotics Lab
ETH Zurich, Switzerland

1 Introduction

Computational planning and fabrication tools have the potential to greatly increase the sustainability of architectural construction, allowing for the use of abundantly available natural and reclaimed materials that are currently too complex or time-intensive for cost-effective widespread application. As compared to many common building materials, minimally-processed rock and locally-recycled demolition debris have extremely low embodied energy (Morel et al. 2001; Alcorn 2003), but also have the potential to express site-specific materiality, revive regional vernacular building methods, and exhibit new forms of “cultural performance” (Oesterle 2009).

This chapter presents an integrated system for constructing double-faced dry stone walls in situ using a customized autonomous mobile hydraulic excavator. We outline a process for mapping the environment, and localizing and digitizing irregular stones. Using this digitized information, we algorithmically determine the position and orientation of stones to align with a designer-indicated goal surface, and conduct grasp- and motion-planning for collision-free placement. We demonstrate the applicability of our method by constructing a 3 m tall wall with a total length of 5 m out of 40 stones with masses between 230-1548 kg.

2 Related Work

A number of recent projects in the domain of architectural digital fabrication have made use of nonstandard input materials combined with robotic construction. Such projects have played a substantial role in re-establishing computational design and fabrication as a practice that can not only materialize digital complexity, but that can understand and adapt to the existing complexity of the material world. Johns



Fig. 1 Stone wall constructed with autonomous hydraulic excavator.

and Anderson (2018) provide an overview of recent projects in this general domain, while a number of notable projects deal more specifically with irregular stone-based assembly: “Smart Scrap” uses automotive digitization tools to inventory partially-dressed limestone for facade planning (de Boer et al. 2007), while “Cyclopean Cannibalism” revives ancient polygonal masonry methods through robotic stone-cutting (Clifford, McGee, and Muhonen 2018). The latter fits parallel-faced polygonal elements within the constraints of demolition concrete and stone offcuts, using 73% of the stock material (selected from a larger set) after dressing. In comparison, the work presented in this chapter makes use of 100% of each stone object, and uses every stone in the available set.

Towards the assembly of such raw materials, the geometric nesting of unmodified parts has been explored in two-dimensions with simulated annealing (Lambert and Kennedy 2012). Others have demonstrated the potential of shape descriptors for matching 3D shapes to existing geometries for the generation of 3D collages without constraints of fabricability (Gal et al. 2007), or for the reassembly of fractured objects without structural constraints (Huang et al. 2006).

The demonstrators presented by Furrer et al. (2017) and Liu, Choi, and Napp (2019) both perform dry stacking of unmodified stones, making use of stationary robotic manipulators and RGB-D scanners for the localization of pre-scanned stones. Both projects consider structural stability using a physics simulator and evaluate a cost function to determine

the validity of each potential pose. The former is focused on the construction of vertical stacks, while the latter can build one-layer thick walls of up to four courses with a 40 % success rate in an open loop fashion (without accounting for the settling of placed stones). In this work, we overcome the design limitations of such 1- and 2- dimensional structures, presenting a planner capable of building a multiple-layer-thick structure following an arbitrarily specified target geometry. Stones are scanned in-situ on the fly, and the as-built wall is monitored in order to refine the placed-stone poses before the planner is executed again.

Several projects have been developed which demonstrate the potential of large-scale autonomous construction with on-site robotics. Generally combining mobile platforms with conventional industrial robots, these compound setups have been predominantly demonstrated with standardized building components (Dörfler et al. 2016) or industrially produced chemical products (Keating et al. 2014). Recent research has demonstrated the potential of robotic systems to construct auxiliary structures in order to achieve and maintain navigability in previously unknown or untraversable terrain, but these experiments have only been demonstrated at a small scale, and do not make use of naturally occurring found construction materials (like stones), but instead use compliant bags (Saboia, Thangavelu, and Napp 2019) or polyurethane foam (Fujisawa et al. 2015) to homogenize the irregularities in their environment.

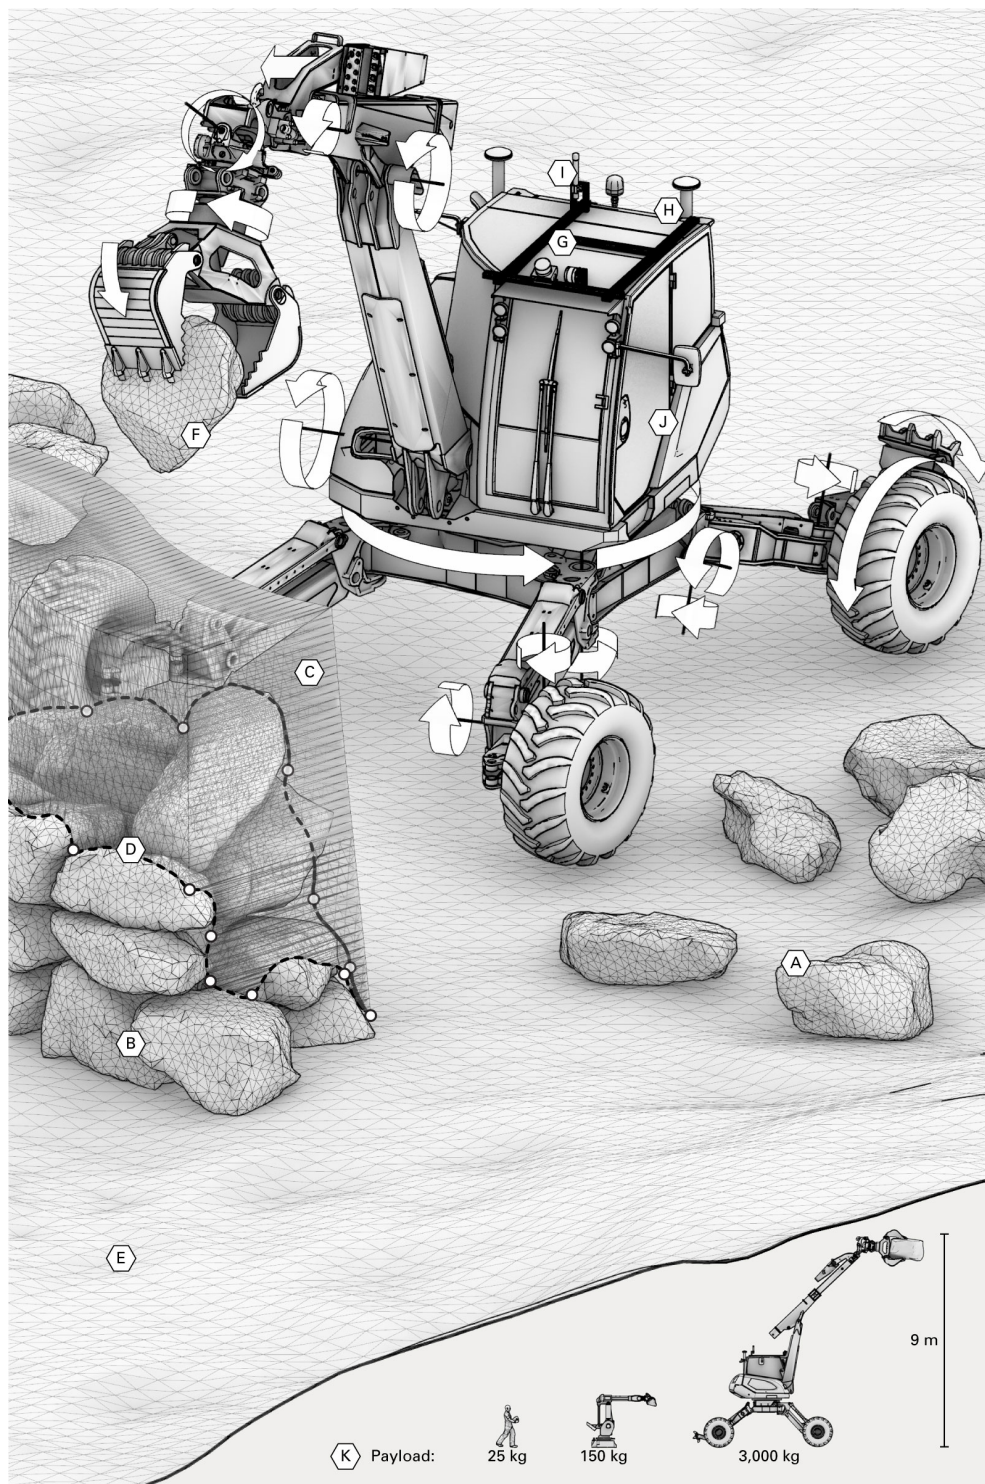


Fig. 2 Setup Diagram:

- A) Available stones
- B) Constructed wall
- C) Search space
- D) Rim path with keypoints
- E) Grid map
- F) Gripped stone
- G) LiDAR sensors
- H) GNSS antennas
- I) Network antenna
- J) HEAP platform with axes indicated
- K) Reference scale and payload comparison with human and industrial robot.

3 Methodology

3.1 Platform

The stone wall was created using HEAP (Hydraulic Excavator for an Autonomous Purpose), a highly customized Menzi Muck M545 12t walking excavator that was developed for

autonomous applications and advanced teleoperation (Figure 2J). The fully-mobile 23-axis excavator is capable of lifting objects to a height of 9 meters, and can freely manipulate items weighing up to 3,000 kg. Adaptable to complex terrain, the machine is equipped with force-controllable hydraulic cylinders that allow for control of ground interaction forces, and cabin- and chassis-mounted Inertial Measure-

ment Units (IMUs) that track machine orientation. Global cabin localization is achieved by a Leica iCON iXE3 with two Global Navigation Satellite System (GNSS) antennas and a receiver (Figure 2H). Real-time kinematic (RTK) corrections for the GNSS signals are received over the network from permanently installed base stations for improved accuracy. Draw wire encoders are used to measure the position and velocity of the hydraulic arm- and grapple-cylinders.

For exteroceptive sensing (scanning of the environment, available stones, and in-progress structure), two Velodyne Puck VLP-16 LiDAR scanners are placed at the front edge of the cabin’s roof to provide 3D scans. The LiDARs are mounted orthogonally to each other, such that scan points accumulate both while driving the robot around and while pivoting the cabin (Figure 2G). For a more detailed description of the system’s sensors and actuators, we refer the reader to (Jud et al. 2019).

All software components of the project are written in C++, and the Robot Operating System (ROS) is used to transfer data over the network between the different software nodes distributed on several computers. The ROS master, managing the connection between processes, is located on the on-board computer of HEAP.

3.2 Stone Localization & Scanning

To reconstruct and localize the stones in the surroundings of the excavator, we use the LiDARs mounted on the cabin roof (Figure 2G). Before the construction process begins, the available stones are loosely distributed on the terrain (Figure 3). This facilitates the detection of single stone instances by first segmenting the ground plane from the LiDAR map, and subsequently performing euclidean clustering on the remaining point cloud to distinguish the single objects. These point clusters are used to compute an initial grasp configuration on the not yet reconstructed stone segments. Once a given stone is picked up using this relatively low-resolution data, a high resolution scan of the full stone is completed while it is held in the gripper. In the following sections, we describe this stone reconstruction process, the LiDAR-based scene mapping, and the object pose refinement based on the reconstruction and scene map information.

Reconstruction For the 3D reconstruction, the end effector is spun with a continuous velocity in front of the LiDAR sensors while holding the stone. With the known geometry of the gripper, we filter out points that belong to the end-effector while accumulating points on the stone. The complete contour of the stone is recorded by concatenating multiple point clouds from different viewing angles, and applying Poisson surface reconstruction (Kazhdan, Bolitho, and Hoppe 2006) to generate the surface mesh from the point cloud data.



Fig. 3 Initial state of stones, loosely distributed on the terrain. LiDAR mapping and RANSAC ground segmentation allows for initial detection and grasping for refined scanning in the gripper.

LiDAR-based Scene Mapping A precise map of the excavator’s complete surrounding (including the as-built structure) is necessary for planning grasp configurations and arm motions, and for facilitating the decision of where and how to position stones. However, simple accumulation of single LiDAR scans is very temporally and spatially limited, as precise knowledge of the sensor’s motion is necessary (and error accumulation leads to drift over time). To create a consistent map over large time- and motion- intervals, we use LaserSLAM (Dubé et al. 2017), a laser-based graph SLAM that considers the odometry constraints provided by the excavator’s state estimation, fusing GNSS measurements with the IMUs, and scan-matching constraints provided by registering a new scan pair using the Iterative Closest Point (ICP) algorithm with a map consisting of the previous scans. Self-see points of the excavator’s arm or legs in the LiDAR scans are filtered to get a 3D point cloud, built locally around the current robot location, without being corrupted by its moving parts (Figure 7A).

ICP Refinement Being able to refine the pose of an already reconstructed stone in the environment (Figure 7B) or in the gripper is an essential capability, as the stone might settle during placement or shift as it is being picked up. For this pose refinement, we first sample a point cloud on the reconstructed mesh of the stone. Using an ICP step, this point cloud is registered to the scene map (if the stone is placed in the environment), or to the accumulated gripper cloud (if it is grasped by the excavator). To improve the localization, we first segment the ground plane in the vicinity of the stone using RANSAC-based 3D plane fitting. Additional filtering is necessary, however, as stones might be cluttered and partially occluded, especially if they are constituents of the in-progress wall. Therefore, the scene map is further segmented

using Euclidean Clustering, and segments belonging to already refined stones or segments too far from the assumed pose are removed before the ICP step. This allows for the reliable pose refinement of objects even if they are partially occluded by the built structure. If the initial orientation of the stone is unknown (e.g. after it has been overturned), the ICP step is performed for multiple initial orientations and the refined transformation with the highest ICP-score is selected.

3.3 Geometric Planning

The iterative selection, positioning and orienting of each stone is determined by a geometric planning software that attempts to construct a double-faced-wall bounded vertically by any user-specified mesh target surface, and below by the LiDAR-scanned elevation map of the existing terrain. Given any number of available stones, the software determines the preferred placement from the available solutions, and the selected stone and the desired transformation are further processed for grasp-planning and physical positioning. Rather than pre-planning the entire wall, the software computes solutions on the fly, stone-by-stone. This is done for two reasons:

First, for sufficiently large constructions, it is impractical to pre-scan, sort, and store hundreds or thousands of stones when the required space and complexity can instead be minimized if they are salvaged, unearthed, or delivered on-demand, as needed. Computing solutions on the fly allows for a more adaptive construction that is not held up if one pre-planned stone goes missing.

Second, the absolute-rigidity of stone makes for significant error propagation that would result in undesirable deviation from the desired geometry, and higher probability of structural collapse in turn. Any minute difference in the scanned geometry can cause unexpected collisions or settling during placement, and the current setup allows for these changes to be automatically accounted-for in the next solution.

Finding a suitable position of even one stone in a relatively small area can be a computationally heavy problem, and the irregular nature of the stones and substrate geometry make it generally difficult to apply stochastic solvers or continuous optimization methods. The software should generally find a valid solution in less time than it takes for the excavator to physically locate, grasp, and place a stone, such that the finding of solutions does not significantly deter progress (at present, the solver takes approximately one minute per stone on a laptop running Ubuntu 18.04 with an Intel i7-8750H 4.1 GHz Processor and 16GB RAM).

Much like the algorithm presented by Liu, Choi, and Napp (2019), we make use of a number of heuristics derived from conventional stone masonry techniques to reduce

the size of the solution space. However, our approach differs significantly in the specification and sequencing of these heuristics: we use both the geometric properties of the stones and the constraints of the designer-specified goal surface to inform an algorithm that begins with fast computational checks, and increases in complexity as the solution space is substantially reduced.

For any given stone placement, the following high level goals are established: determine a structurally stable solution that minimizes the volume directly under the stone in question (Figure 4A) and also minimizes the volume between the exposed stone face and the designated geometry of the wall surface (Figure 4B). To achieve these objectives, we implement the following subroutines:

Stone Attributes At all times, the software holds an inventory of stones that have been scanned. Upon receipt of new information from the scanning node, the stone mesh is immediately cleaned by removing any degenerate faces and disjoint remnants that might remain from the Poisson reconstruction. The mass properties are computed using the method described by Bacher et al. (2014), and each stone is repositioned such that its centroid is at the origin, and its principal axes (computed using PCA on the mesh vertices) are aligned with the global frame. An iterative implementation of the Variational Shape Approximation method (Cohen-Steiner, Alliez, and Desbrun 2004) is used to determine approximate partitioned face-surfaces on the stones given an allowable normal deviation per clustered region of mesh triangles. From these partitioned regions, a face-adjacency matrix is generated, stone edges and edge mid-points are identified, and the mean weighted normal of each region is used to represent the normal direction of each partitioned face-surface (Figure 4C).

Search Space A volumetric search space (Figure 2C) is defined as the region bounded by the designated outer wall surfaces on all sides and above, and bounded below by the ground and the upper surface of the in-progress wall. While this upper surface is in reality generally composed of many disjointed stones, it is approximated as a single continuous surface using a simulated LiDAR scanner on the known mesh data of the ground and placed stones, and meshed using Poisson reconstruction. Thus, both the empty ground and the upper bounds of a half-built wall are treated similarly as the foundation for the next stone. The intersection between this reconstructed “ground” surface and the target wall-face surfaces defines the rim path, which we identify as the upper perimeter edge of the structure at any given state (the contour where the top surface of the as-built structure intersects with the vertical goal wall-surfaces)(Figure 2D).

An even sampling along the rim path serves as a one-dimensional list of possible placement positions, and this list

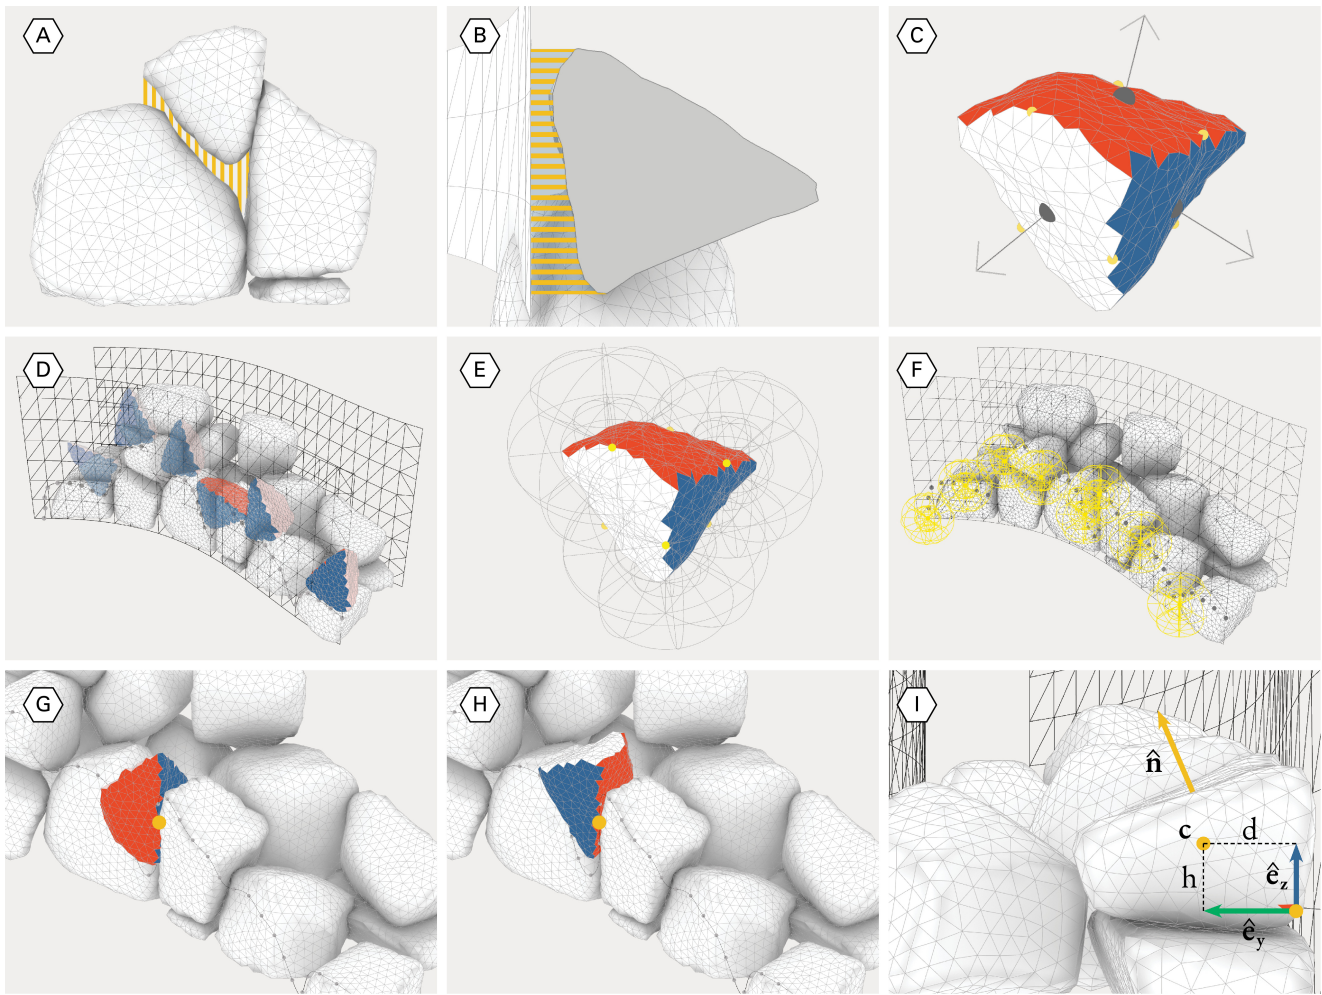


Fig. 4 Stone solver goals and methods. A) Minimize volume under stone, B) Minimize volume between stone face and goal surface, C) Compute faces and edges, D) Sample along rim path, E) Compute descriptors at edge midpoints F) Descriptors at rim path points. G-H) Two possible orientations for each partial match I) Stability and upper normal direction factors.

is further reduced to only include lower points within a given distance of identified keypoints (Figure 2D) along this path: keypoints are defined as any point along the rim path that lies either between two stones, where one stone meets the ground, or at the corners or ends of designated structure. This filtering process ensures that the wall is generally built up evenly, and reduces the likelihood of running joints that weaken the bonding of the wall.

Fit Estimation Given the known edge-midpoints, edge directions, and face-normal-vectors on the stones, and the normal vectors of the goal-wall surface at each point on the rim path, it is trivial to determine the transformations necessary to place any specified stone on this path—such that a given stone edge is tangent to the rim path, and the stone face is aligned with the wall surface. While each combination of path-point and stone-edge could be sampled (Figure 4D), this still represents a substantially large solution-space if the wall is long and there are multiple known stones in the avail-

able inventory. Just as a stone-mason would likely only try placing stones that look like they will fit (instead of brute force sampling every possible position), we first check if a stone is geometrically similar to the region in question using rotationally invariant shape descriptors before attempting any more involved processing (Figures 4E-F, 5). Specifically, we use the FPFH descriptor (Rusu, Blodow, and Beetz 2009) at each edge midpoint and at each rim path vertex coupled with a k-nearest-neighbor (KNN) search to find the k-closest geometric matches at each included point on the rim path ($k \approx 10$). In this demonstrator, the KNN search typically reduced over a million possible combinations of stone edges with rim path points to some thousands. Because of the rotationally invariant nature of such shape descriptors, each match must consider both possible orientations of an edge with the rim path (Figures 4G, 4H).

Stability Heuristics Prior to running any involved physics simulations, we remove potential candidates that are likely



Fig. 5 Elevation detail of constructed wall. Rotationally invariant shape descriptors allow for a tight fit by initially matching available stone surfaces to corresponding negative geometry of the wall.

unstable, or likely to reduce future stability, informed by guidance from stone masonry literature. In dry stone walls, it is generally advised to not place stones in such an orientation that their thinnest dimension defines the depth into the wall—thus acting as a thin veneer that is more likely to fall away from the wall (Environmental Action Foundation 2019). We generalize this requirement as needing to meet a minimum ratio between the horizontal (d) and vertical (h) distance measured between the specified point on the rim path and the stone centroid (Figure 4I):

$$d/h > 0.5 \quad (1)$$

As the construction process uses a double-faced method, the structure benefits from having an inward slope at the top of each placed stone—such that any settling is minimized and supported by stones on the opposing side of the wall (Vivian 1976).¹ For each intended position, we take an even sampling of the exposed stone surface using a raycasting method, and compute the mean top-exposed surface normal ($\hat{\mathbf{n}}$) from the stone normal direction at each hit point. We then verify that $\hat{\mathbf{n}}$ is pointing into the wall and that the stone surface is not excessively steep by

$$\hat{\mathbf{e}}_y \cdot \hat{\mathbf{n}} \geq 0 \quad \text{and} \quad \hat{\mathbf{e}}_z \cdot \hat{\mathbf{n}} > 0.7, \quad (2)$$

¹ Note that this inward slope differs from the inward slope employed by Liu, Choi, and Napp (2019): in our use and in the masonry literature, the inward slope is a sectional property of a double faced wall. Liu et al. use it to mean a lowered center of mass between the end stones of a given bond, as observed across the front elevation (thereby creating a sagging bond in a single-layered structure).

where $\hat{\mathbf{e}}_y$ is the unitized inward-facing surface normal of the target wall surface projected onto the horizontal plane, and $\hat{\mathbf{e}}_z$ is a unit vector pointing against the direction of gravity (Figure 4I).

ICP Refinement Once the pool of potential matches has been reduced by the aforementioned heuristics, each candidate is refined using ICP to better align it with the surrounding stones and goal wall surface.

Physics Following ICP refinement, the remaining solutions are sorted by their ICP score, and (if specified) reduced to a preset maximum number of solutions. These are then simulated for stability with the Bullet physics engine (Coumans 2015), and solutions that do not reach equilibrium within a fixed simulated-world-time period or within a distance moved are ruled out.

Scoring The remaining stones after the physics simulation step are sorted using the combined parameters from figure 4A and B, where the gap volume is calculated with a raycasting technique measuring the displacement of hits measured from local plan- and elevation- projection planes with and without the placed stone (and considering the known stone volume). The best match is then sent for physical placement. A provided Imgui interface (Cornut 2005) allows for alternative high-scoring solutions to be selected manually, if desired.

The geometric planning software makes extensive use of libigl (Jacobson, Panozzo, et al. 2018) for mesh processing, PCL for point cloud operations (Rusu and Cousins 2011), and openframeworks (Lieberman et al. 2005) for visuals (Figure 6).

3.4 Grasp and Placement Execution

The goal of the grasp pose planning is to find viable grasping configurations that allow the excavator to pick a desired stone and place it at the planned location without collisions with the existing wall (Figure 7E). A grasp configuration is defined as 6-DoF pose of the excavator’s gripper, where contact with the desired stone can be performed. In order to find a viable grasp configuration that respects the collision constraints, we sample a large amount of grasp hypotheses on the stone of interest in the map point cloud. Those grasp hypotheses are validated for collisions by intersecting a polyhedral gripper model with the map cloud and searching for inliers. Note that both the grasp configuration and the corresponding placement configuration must be verified to be without collisions (Figure 7D/F). Due to occlusions, the map point cloud may contain holes that can result in undetected collisions. To prevent these situations, the map point cloud

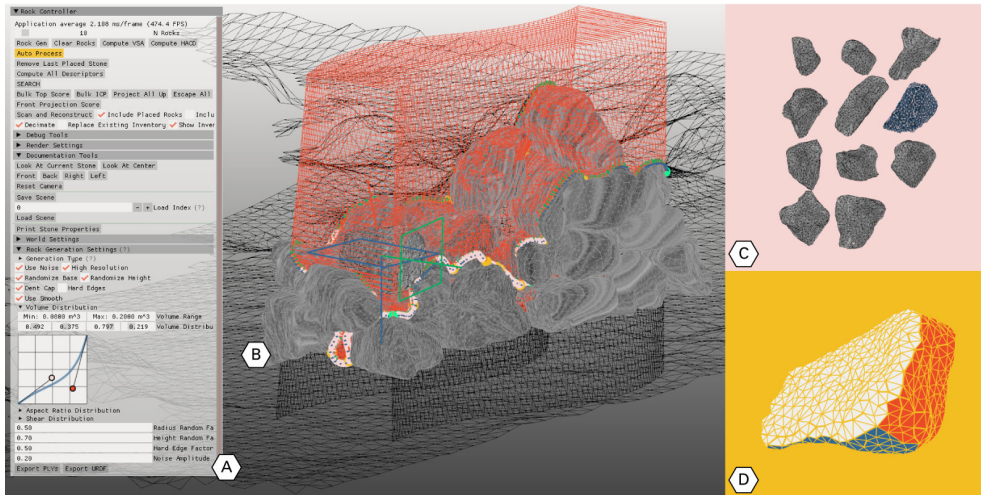


Fig. 6 Screen capture of geometric planning software. A) ImGui Interface, B) Active environment with search space, solutions, and placed stones, C) Available scanned stones in inventory, D) Selected Stone Viewer with VSA faces.



Fig. 7 Grasp planning and placement of a stone to its desired location (red). (A) Collision and localization point cloud, (B) Stones localized in point cloud, (C) Spline motion trajectory, (D, Green) Collision free grasp hypotheses, (F, Red) Grasp hypotheses with collisions, (E) Selected grasp configuration, (G) The collision point cloud is augmented with point clouds of already localized stones.

is augmented with point clouds generated from the meshes of already localized objects (Figure 7G). After ruling out colliding configurations, the remaining grasp poses are evaluated for force closure, and scored by task specific criteria (alignment to ground, grasp encompassment, and distance to stone CoM) to obtain the final grasp.

A motion planner is used to generate a spline trajectory for reaching the desired grasp pose and for moving from the grasp pose to the placement pose (Figure 7C). The approach direction during placement is chosen by ray-casting to provide the largest margin to obstacles. Figure 7 shows the spline trajectory (C, blue) for moving the grasped stone from the ground to its desired placement pose (E, red). Each

time a stone is grasped or released, its pose with respect to the map or end-effector is updated with an ICP refinement step using its previous location as initial guess (Section 3.2).

4 Experiment

The localization, scanning, geometric- and motion- and planning routines developed in this work were used to construct a large-scale (3 m high and 5 m long) test wall using 40 available gneiss stones² with an average mass of 757 kg

² Gneiss is a locally abundant metamorphic rock often used for architectural applications, landscape construction, and retaining walls.




Fig. 8 Digital model of constructed wall with goal surfaces and indicated distance between placed stones and one target surface.



Fig. 9 Side elevation of constructed wall conveying double-faced structure and near-vertical orientation.

the greatest height and number of components of any masonry structure autonomously constructed from irregular stones. It is essential to note, however, that such qualifications are heavily influenced by the geometry of the available material, and one might much more easily construct a many-bonded wall with ashlar-like stones. Despite the relatively large size of the available stones,³ the outer face of the as-built structure was measured to be generally within 10 centimeters of the intended target surface (Figure 8). Figure 9 shows the side elevation of the completed structure, and the front elevation can be seen in Figure 1.

While the geometric planner allows for solutions to be found with any number of stones in the inventory, it predictably finds better correspondences if there are more stones to choose from (at the computational expense of an increased search space). For a much larger construction, one would likely scan and build in batches, keeping tens of digitized stones in the inventory, and refilling it when no solution can be found within some tolerance requirements. For this experiment there was a fixed number of available stones (40), and they were each scanned and added to the digital inventory before the construction of wall commenced. This accelerated construction, and provided a wide variety of stones for the initial planning steps. The planning and execution were performed sequentially, meaning that after each placement the stone pose was refined and updated in the geometric planning before initiating the next update step. Figure 10

(range: 230-1584 kg). To our knowledge, it represents both

³ These are classified as ‘medium-coarse boulders’ on the modified Udden–Wentworth grain-size scale (Blair and McPherson 1999).

shows intermediate steps of the wall construction with two and twenty placed stones.



Fig. 10 Wall construction in progress with two (above) and twenty (below) placed stones out of the 40 used in this experiment.

5 Improvements

A number of improvements are intended which will provide increased process robustness and allow for additional design constraints to influence the solver. While this initial prototype was imagined as a section of a longer wall (Figures 11, 12), future developments will resolve the particularities of terminating details at wall ends and corners: the use of alternating stretchers and headers with returned faces that both strengthen the structure and clearly define the *arris* (edge where surfaces meet) (Cramb 1992).

One major limitation of the setup at present is that we have not yet implemented the autonomous flipping of stones. On the occasion that placement solutions are found which cannot be grasped in the current stone position, an alert is generated that conveys the need for manual intervention: the stone must be flipped manually before the autonomous setup will relocalize the stone and continue construction.

The drawback of using LiDARs for reconstruction and ICP refinement is the need to always move the sensor to accumulate scans into the complete 3D model: this is time consuming, and renders the current sensors ill-equipped for dynamic tracking during stone flipping. This task will be made more feasible with the imminent integration of additional sensors, including RGB cameras. By texturing the reconstruction of the stones (Waechter, Moehrle, and Goesele 2014) it is possible to extract visual features that can be used for recognizing a stone instance or tracking its motion (Marchand, Uchiyama, and Spindler 2016). This will not only better facilitate dynamic tasks such as flipping, but could also be considered as a component in the design process. Provided that a structure contained enough stones to make a pattern legible, it would be possible to incorporate variable distributions of color, scale, or surface shapes to meet additional design goals beyond the global shape of the wall.

Whereas we currently rely on physics simulation for stability assessment, in future work we will consider structural validation by physically probing the as-built construction. The excavator – equipped with pressure sensors and an additional force-torque sensor – would serve as both a construction tool and testing device, using applied loads to verify the stability of the structure and measure deformations.

6 Discussion

This research demonstrates a robotically-enabled process for creating structures from a wide range of input and target geometries: allowing digital design goals to be integrated with existing ancient construction methods for double-faced free-standing dry stone walls built from unprocessed or minimally-processed natural stone. In contrast to cut stone (and even more so to concrete), these exhibit extremely low embodied energy because there is minimal transportation required⁴, and because they do not require energy-intensive cutting or post-processing. While the embodied energy of stone construction varies regionally, post-processing (sawing, etc.) is a major contributor to the overall environmental impact of masonry architecture: despite an increase in the thickness of the built structure when compared to cut stone, building with raw stone can be environmentally advantageous (D. Ioannidou, Zerbi, and Habert 2014).

Rather than modifying individual stones, this method focuses on cleanly fitting the exposed surfaces of stones along the outer two faces of a given wall section, while tolerating necessary gaps between stones within the hidden space of the *poché*. Even in a well-constructed dry stone wall, these gaps can account for 20-40% of the constructed volume (Villemus 2004; Mundell 2009), and are essential in

⁴ For this experiment, the raw stones were delivered by truck from a quarry located 40 km from the test field.



Fig. 11 In progress wall with overlay of potential extension of structure, generated by allowing the solver to continue with additional digital stone models.



Fig. 12 Constructed wall with figure (left) for scale and indicated potential extensions.

enabling the use of completely irregular and indeterminate source materials: a dry stone wall is thus “defined by the spaces between the stones as much as it is by the stones themselves” (Snow 2008). While such gaps can be filled in with smaller ‘chinking’ stones (Vivian 1976) or with mortar, they can also be left as supplemental habitats for flora and fauna to thrive (Witschi 2019).

Dry stone construction has diminished in popularity in the last century, despite the environmental, economic, and aesthetic advantages of building with inexpensive, local, and natural materials: increasing labor costs combined with comparatively inexpensive and simple-to-install mass manufactured building components have rendered irregular stone construction infeasible in many situations (Bätzing 2019; Dimi-

tra Ioannidou et al. 2018). By synthesizing knowledge from manual dry stone masonry handbooks (Vivian 1976; Cramb 1992; Environmental Action Foundation 2019) together with new methods of computational design and robotic construction, this work aspires to reactivate irregular stone construction while enabling new modes of design expression and large-scale fabrication.

7 Conclusion

In this chapter, we introduce a platform and algorithm for the autonomous construction of large scale dry stone walls in situ, with the aid of a customized robotic excavator—addressing the challenge of using unprocessed and locally available construction material. The addition of a high-payload, fully-mobile platform has greatly expanded the ‘machinic morphospace’ (Menges 2013) of robotic construction tools on site, allowing us in turn to better navigate the constrained ‘material morphospace’ of complex and nonstandard found objects (Johns and Foley 2014). The process has potential to greatly reduce the embodied energy of architectural construction in specific contexts, and can be extended to work with a range of source materials, including reclaimed demolition debris. By realizing a double-faced wall on an architectural scale, we demonstrate the applicability of automated stone masonry for future construction processes.

Acknowledgements This research was supported by the Swiss National Science Foundation through the National Centre of Competence in Digital Fabrication (NCCR dfab). Gramazio Kohler Research post-doctoral support: Kathrin Dörfler, Aleksandra Anna Apolinarska, Lauren Vasey.

References

- Alcorn, Andrew (2003). *Embodied Energy and CO₂ Coefficients for NZ Building Materials*. Centre for Building Performance Research.
- Bacher, Moritz et al. (2014). “Mass Properties of Triangulated Solids and Their Derivatives, Supplemental Material for ‘Spin-It: Optimizing Moment of Inertia for Spinnable Objects’”. In: *ACM Transactions on Graphics (TOG)* 33.4, p. 2.
- Bätzing, Werner (2019). “Life in the Alps - Is There a Future?” In: *Dry Stone Walls: Basics, Construction, Significance*. Ed. by Environmental Action Foundation. 1st ed. Zurich, Switzerland: Scheidegger & Spiess, pp. 379–383.
- Blair, Terence C. and John G. McPherson (Jan. 1, 1999). “Grain-Size and Textural Classification of Coarse Sedimentary Particles”. In: *Journal of Sedimentary Research* 69.1, pp. 6–19.
- Clifford, Brandon, Wes McGee, and Mackenzie Muhonen (Dec. 1, 2018). “Recovering Cannibalism in Architecture with a Return to Cyclopean Masonry”. In: *Nexus Network Journal* 20.3, pp. 583–604.
- Cohen-Steiner, David, Pierre Alliez, and Mathieu Desbrun (Aug. 1, 2004). “Variational Shape Approximation”. In: *ACM Transactions on Graphics (TOG)* 23.3, pp. 905–914.
- Cornut, Omar (2005). *Dear ImGui: Bloat-Free Immediate Mode Graphical User Interface for C++ with Minimal Dependencies*. URL: <https://github.com/ocornut/imgui> (visited on 06/11/2020).
- Coumans, Erwin (July 30, 2015). “Bullet Physics Simulation”. In: *ACM SIGGRAPH 2015 Courses*. SIGGRAPH ’15. Los Angeles, California: Association for Computing Machinery, p. 1.
- Cramb, Ian (1992). *The Art of the Stonemason*. Chamersburg, PA: Alan C. Hood & Co. 174 pp.
- De Boer, Jason et al. (2007). *Smart Scrap*. URL: <http://i-m-a-d-e.org/?p=2201> (visited on 03/07/2020).
- Dörfler, Kathrin et al. (2016). “Mobile Robotic Brickwork”. In: *Robotic Fabrication in Architecture, Art and Design 2016*. Ed. by Dagmar Reinhardt, Rob Saunders, and Jane Burry. Cham: Springer International Publishing, pp. 204–217.
- Dubé, Renaud et al. (Sept. 2017). “An Online Multi-Robot SLAM System for 3D LiDARs”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1004–1011.
- Environmental Action Foundation, ed. (2019). *Dry Stone Walls: Basics, Construction, Significance*. Zurich, Switzerland: Scheidegger & Spiess. 472 pp.
- Fujisawa, Ryusuke et al. (Apr. 1, 2015). “Active Modification of the Environment by a Robot with Construction Abilities”. In: *ROBOMECH Journal* 2.1, p. 9.
- Furrer, Fadri et al. (May 2017). “Autonomous Robotic Stone Stacking with Online next Best Object Target Pose Planning”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 2350–2356.
- Gal, Ran et al. (Aug. 4, 2007). “3D Collage: Expressive Non-Realistic Modeling”. In: *Proceedings of the 5th International Symposium on Non-Photorealistic Animation and Rendering*. NPAR ’07. San Diego, California: Association for Computing Machinery, pp. 7–14.
- Huang, Qi-Xing et al. (July 2006). “Reassembling Fractured Objects by Geometric Matching”. In: *ACM Trans. Graph.* 25.3, pp. 569–578.
- Ioannidou, D., S. Zerbi, and G. Habert (Dec. 2014). “When More Is Better – Comparative LCA of Wall Systems

- with Stone”. In: *Building and Environment* 82, pp. 628–639.
- Ioannidou, Dimitra et al. (May 19, 2018). “Where Does the Money Go? Economic Flow Analysis of Construction Projects”. In: *Building Research & Information* 46.4, pp. 348–366.
- Jacobson, Alec, Daniele Panozzo, et al. (2018). *Libigl: A Simple C++ Geometry Processing Library*. URL: <https://libigl.github.io/>.
- Johns, Ryan Luke and Jeffrey S. Anderson (2018). “Interfaces for Adaptive Assembly”. In: *ACADIA // 2018: Recalibration. On Imprecision and Infidelity*. 38th Annual Conference of the Association for Computer Aided Design in Architecture. Mexico City: CUMINCAD, pp. 126–135.
- Johns, Ryan Luke and Nicholas Foley (2014). “Bandsawn Bands”. In: *Robotic Fabrication in Architecture, Art and Design 2014*. Ed. by Wes McGee and Monica Ponce de Leon. Cham: Springer International Publishing, pp. 17–32.
- Jud, D. et al. (Oct. 2019). “Autonomous Free-Form Trenching Using a Walking Excavator”. In: *IEEE Robotics and Automation Letters* 4.4, pp. 3208–3215.
- Kazhdan, Michael, Matthew Bolitho, and Hugues Hoppe (June 26, 2006). “Poisson Surface Reconstruction”. In: *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*. SGP’06. Cagliari, Sardinia, Italy: Eurographics Association, pp. 61–70.
- Keating, Steven et al. (2014). “A Compound Arm Approach to Digital Construction”. In: *Robotic Fabrication in Architecture, Art and Design 2014*. Ed. by Wes McGee and Monica Ponce de Leon. Cham: Springer International Publishing, pp. 99–110.
- Lambert, Malcolm and Paul Kennedy (2012). “Using Artificial Intelligence to Build with Unprocessed Rock”. In: *Key Engineering Materials* 517, pp. 939–945.
- Lieberman, Zachary et al. (2005). *openFrameworks*. URL: <https://openframeworks.cc/> (visited on 06/11/2020).
- Liu, Yifang, Jiwon Choi, and Nils Napp (2019). “Planning for Robotic Dry Stacking with Irregular Stones”. In: *Easy-Chair Preprint no. 1453* 1453.
- Marchand, Eric, Hideaki Uchiyama, and Fabien Spindler (Dec. 2016). “Pose Estimation for Augmented Reality: A Hands-On Survey”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.12, pp. 2633–2651.
- Menges, Achim (2013). “Morphospaces of Robotic Fabrication”. In: *Rob — Arch 2012*. Springer, Vienna, pp. 28–47.
- Morel, J. C et al. (Dec. 1, 2001). “Building Houses with Local Materials: Means to Drastically Reduce the Environmental Impact of Construction”. In: *Building and Environment* 36.10, pp. 1119–1126.
- Mundell, Chris (2009). “Large Scale Testing of Drystone Retaining Structures”. Doctoral Thesis, Department of Architecture & Civil Engineering. University of Bath.
- Oesterle, Silvan (2009). “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)*. Ed. by Tristan de Estree Sterk, Russell Loveridge, and Douglas Pancoast. Chicago, IL: ACADIA, pp. 194–200.
- Rusu, Radu Bogdan, Nico Blodow, and Michael Beetz (May 2009). “Fast Point Feature Histograms (FPFH) for 3D Registration”. In: *2009 IEEE International Conference on Robotics and Automation*. 2009 IEEE International Conference on Robotics and Automation, pp. 3212–3217.
- Rusu, Radu Bogdan and Steve Cousins (May 2011). “3D Is Here: Point Cloud Library (PCL)”. In: *2011 IEEE International Conference on Robotics and Automation*. 2011 IEEE International Conference on Robotics and Automation, pp. 1–4.
- Saboia, Maira, Vivek Thangavelu, and Nils Napp (Nov. 1, 2019). “Autonomous Multi-Material Construction with a Heterogeneous Robot Team”. In: *Robotics and Autonomous Systems* 121, p. 103239.
- Snow, Dan (2008). *Listening to Stone: Hardy Structures, Perilous Follies, and Other Tangles with Nature*. New York, NY: Artisan. 145 pp.
- Villemus, Boris (2004). “Etude des murs de soutènement en maçonnerie de pierres sèches”. Doctoral Thesis, Civil Engineering. Lyon, France: L’institut National Des Sciences Appliquées De Lyon. 254 pp.
- Vivian, John (Jan. 3, 1976). *Building Stone Walls*. Storey Publishing, LLC. 114 pp.
- Waechter, Michael, Nils Moehrl, and Michael Goesele (2014). “Let There Be Color! Large-Scale Texturing of 3D Reconstructions”. In: *Computer Vision – ECCV 2014*. Ed. by David Fleet et al. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 836–850.
- Witschi, Franziska (2019). “Flora, Fauna”. In: *Dry Stone Walls: Basics, Construction, Significance*. Ed. by Environmental Action Foundation. 1st ed. Zurich, Switzerland: Scheidegger & Spiess, pp. 329–378.