



Robotic embankment

Free-form autonomous formation in terrain with HEAP

Dominic Jud¹ · Ilmar Hurkxkens² · Christophe Girot² · Marco Hutter¹

Received: 8 January 2021 / Accepted: 5 May 2021
© The Author(s) 2021

Abstract

Automating earth-moving tasks has the potential to resolve labour-shortage, allow for unseen designs and foster sustainability through using on-site materials. In this interdisciplinary project involving robotics and landscape architecture, we combine our previous work on autonomous excavation of free-form shapes, dynamic landscape design and terrain modelling tools into a robotic landscape system. It tightly connects survey, design and fabrication to exchange information in real-time during fabrication. We purposely built a LiDAR survey drone for tight integration. The design environment contains terrain modelling tools to balance cut and fill volumes for material-neutral, on-site construction. Its parametric nature allows it to adapt the geometry to changing site conditions during fabrication. Our autonomous walking excavator is used to create these free-form shapes in natural granular material. We propose an excavation planner for free-form embankments that computes the next excavation location and subsequently the location where the excavated soil should be dumped. This robotic excavation system achieves the world's first autonomous completion of free-form embankments with high accuracy. A 20 m long S-shaped and a two-faced embankment with a corner with roughly 0.03–0.05 m average error were created.

Keywords Autonomous excavation · Landscape design · Free-form embankment · On-site construction · Computational terrain modelling

1 Introduction

The construction sector is one of the largest in the world and vital to every country's economy. However, productivity has not increased significantly over the last decades compared to other sectors. In combination with a rising labour shortage, this sector is in need of a transformation into the digital age through automation to solve these problems. Automating earth-moving tasks will relieve human workers from physically hard, repetitive and dangerous work to focus their skills on more interesting and rewarding tasks (Chui and Mischke

2019). Additionally, these robotic systems will allow for new ways of creating earthworks capable of mediating ecological and mechanical performance in unstructured environments.

This work focuses on highly accurate robotic fabrication of free-form embankments which are seen as templates for larger projects such as road construction, sound barriers, river profiles and protective structures against avalanches, landslides and debris flows. These embankments are built through digital processes that allow novel, material-neutral and terrain adaptive designs. It is achieved by changing the current sequential construction process where first a survey is conducted, a design is created based upon that survey, and the structure is then fabricated according to the design. This one time sequential execution of survey, design and fabrication is changed to a circular feedback scheme where design and fabrication are tightly connected to exchange information both ways.

✉ Dominic Jud
dominic.jud@mavt.ethz.ch
Ilmar Hurkxkens
hurkxkens@arch.ethz.ch

¹ Robotic Systems Lab, ETH Zurich, Leonhardstrasse 21, 8092 Zurich, Switzerland

² Chair of Landscape Architecture, ETH Zurich, Stefano-Francini-Platz 5, 8093 Zurich, Switzerland

1.1 Related work

First, autonomous excavation requires a robotic excavator. One of the first automated excavators was LUCIE developed by the Lancaster University (Bradley and Seward 1998, 1995) more than three decades ago. Their focus lied in control of the excavator for trenching tasks. A decade later, the Autonomous Loading System (ALS) by the Carnegie Mellon University was introduced (Cannon 1999). They were the first to add exteroceptive sensing to an automated machine with a line-laser on the excavator's roof. The focus was mass excavation to load trucks in large open mines. Another decade later, a similar task was again targeted with THOR by the Technical University of Kaiserslautern (Schmidt et al. 2010) exploring behaviour-based control and increasing the applicability through adding construction site navigation (Schmidt and Berns 2015). The last autonomous excavator from academia in this line-up is from the Australian Center for Field Robotics, where they have a long history of industry-driven projects. Their 1.5 tonne excavator by Quang et al. Ha et al. (2002) is used to autonomously dig trenches using a combined notion of force and position. They are the first to abandon the idea of trying to follow a position trajectory closely with the first pass. Instead, they use an iterative approach to converge to the desired profile after a few consecutive digs (Maeda and Rye 2012; Maeda et al. 2014). A trench with 0.025 m accuracy is achieved (Maeda et al. 2015). However, they do not capture the terrain with any sensors but infer the actual elevation only from the bucket edge motion. No dynamic model in the conventional sense is used by Park et al. Park et al. (2017). Instead, the dynamics are learned online with an echo state network and used to track a position trajectory. The network can also compensate for changes during execution, e.g., change in fluid temperature.

The only fully autonomous excavator from industry that has proven itself in a real-world use case so far is AX-162 by Built Robotics. They autonomously dug pipeline trenches across Queensland, Australia Built Robotics 2021.

We use HEAP in this work, the world's first autonomous walking excavator. The development of HEAP from an off-the-shelf walking excavator into an autonomous robot is described in our previous work (Jud et al. 2020). It covers the necessary building blocks, i.e., automation of arm and legs, driving control, state estimation and proprioceptive and exteroceptive sensing.

To work on larger areas, an excavation task planner has to find a series of single dig cycles that complete the excavation task. An example of a global planner is presented by Seo et al. (2011). The excavator moves in a predefined pattern over the excavation site, similar to a milling

machine, and excavates layer by layer until the desired depth is reached. Zhao et al. (2020) show an iterative planner by training a multilayer perceptron to find the next point of attack for excavation. Singh and Simmons (1992) proposed an iterative planner that provides a starting point and the bucket's entire trajectory in the soil. It considers bucket volume, reachability, desired end shape (flat bottom only) and the excavation force. All of the excavation task planners mentioned above are solely intended for mass excavation, loading operations or trenching. In conclusion, no work focuses on the task of precise excavation of more complex forms, except our previous work on free-form trenches (Jud et al. 2019). We proposed an iterative excavation planner that finds the next excavation location by scoring based on the excavation error and the topography. It is extended in this work for 3D embankments.

To explore the creation of precise 3D geometry in loose and granular material, we have to refer to model-scale experiments. Gramazio Kohler Research explored the processing of shapeless sandy materials with a robotic arm equipped with sensors and various end-effectors that allowed for feedback-driven formation processes Gramazio et al. 2021. This work was extended into modelling natural granular material that explored specific digging, shifting and dumping strategies using adaptive feedback systems and dynamic modelling tools by Hurkxkens et al. 2021. Similarly, Bar-Sinai et al. (2019) introduced a protocol combining a theoretical framework and an iterative process for remote groundscaping. However, a schism still exists between models of information (geographic information system), models of design (as intuitive and free-form methods) and models of fabrication. Engineers or landscape designers have to choose between GIS, free-form modelling or computational methods depending on their skill levels. The digital sculpting tools by Westort Westort (1998) are an excellent example of how to integrate the creation of landscape forms with the logic of earth-moving processes. However, it lacks the flexibility of free-form modelling software. In our previous work, we developed computational terrain modelling tools to combine both free-form drawing and computational methods (Hurkxkens and Bernhard 2019). This resulted in the terrain modelling plugin Docofossor (Hurkxkens and Bernhard 2019) for Rhino 3D which was extended to integrate directly into the fabrication environment of HEAP for this work.

1.2 Contribution

The main contribution of this article is the autonomous excavation of free-form embankments in terrain, as shown in Fig. 1. The digital process for robotic creation of embankments includes an initial survey through a drone purposely built for tight integration in this system. A novel digital design process is implemented that allows for a material



Fig. 1 The autonomous walking excavator HEAP can excavate free-form shapes with high accuracy. The image shows the s-curved embankment's progress with the finished surface in the foreground and the unfinished pile to excavate behind it

neutral design and feedback from the construction process. HEAP is used to execute the excavation process. To this end, an excavation planner for 3D embankments was developed. Apart from computing the next excavation point, it can also find a suitable dump location for the soil. In its entirety, the robotic excavation system shows the world's first autonomous completion of 3D embankments with unprecedented accuracy. An 20 m long S-shaped and a straight embankment with 0.03–0.05 m average error were created.

2 Robotic landscape system

As illustrated in Fig. 2, the robotic process is split into three major parts. First, a drone performs an initial survey to gather a three-dimensional map of the construction site. Secondly, design algorithms compute the desired shape of the embankment based on the recorded data. Thirdly, the desired elevation is sent to HEAP, which will execute the necessary excavation and dump actions to turn the design into reality. The fabrication process is running at a feedback rate of 100 Hz to control HEAP. The current elevation map is the feedback from the fabrication to the design. In the future, properties such as soil composition will also be fed back. The feedback rate from fabrication to design can vary according to the project. A fast rate might be helpful to adapt the design often and quickly to changes in the environment. In the other extreme, there can also be no feedback from fabrication to the design and HEAP will try to recreate the initially planned design as well as possible. The drone survey should be used at the beginning to gather a complete map of the environment for the design step. Additionally, the drone can be deployed at discrete times to get a complete overview of the progress and possibly adapt the design. The




Fig. 2 The robotic process starts with an initial drone survey. The map of the construction site is then used in a design environment to produce the desired shape of the embankment. This desired elevation is sent to HEAP, which will realize the plan through excavation and dumping cycles. Information is fed back, which results in a novel, adaptive process

excavator's viewpoint might make it miss certain places due to obstructions.

2.1 Initial survey—landscape mapping drone

The initial survey of the construction site is carried out with a purpose-built drone. The motivation to build an airborne mapping drone, instead of choosing an off-the-shelf product, is, on the one hand, the costs and, on the other hand, tighter integration and flexibility of a self-developed system. Multiple different sensors and their arrangement could be tested out.

There are two possible options for airborne mapping, camera and (LiDAR) based. Where camera-based methods, such as photogrammetry/structure from motion, are

lower cost and require less payload, LiDAR's produce preciser maps, especially in a vertical direction. Additionally, LiDAR's can also capture more complex structures, e.g., trees, bridges and buildings, correctly where photogrammetry will fail. This might not be critical for our use case of mapping landscapes. However, in an architectural context, producing visually appealing maps is a benefit.

Upon the decision of mounting multiple sensors, including a relatively heavy LiDAR on the drone, a DJI M210 RTK drone with 1.72 kg payload was chosen. The real-time kinematic (RTK) enhanced Global Navigation Satellite System (GNSS) positioning will provide a world reference to the map. This is critical, as the later fabrication step with HEAP will also rely on a GNSS-based state estimator. Thus, the drone map can be easily referenced to the excavator and the construction plan in general.

The entire sensor payload is shown at the bottom of Fig. 3. It is built around a light-weight but stiff 3D-printed structure. Three connection points at the top are used to securely mount it to the drone's original damped camera mount points, as shown in the in-flight image at the top of Fig. 3. An Ouster OS1 LiDAR is chosen due to its low weight (380 g) and high-resolution with 64 lines. It is mounted at the front such that the cooling fins are located in the rotor's downwash. Otherwise, the heavy base plate has to be installed for better cooling. Additionally, an RGB camera (Ximea xiC) and a multispectral camera (Micasense RedEdge-M) can be mounted as well. However, they are not used yet in this work. In the future, the multispectral camera will allow augmenting the point cloud gathered by the LiDAR with additional information about the ground material. All of the sensors are connected to an Nvidia TX2 with an Auvideo J120 carrier board. It was chosen for its low weight and low power consumption. The carrier board provides the connections to all sensors at the same time and can also house a large M2 SSD for data storage. The computing unit is also directly connected to the drone through an RS-232 connection to read any information from the drone. The entire sensor payload is powered through the drone's batteries, with the necessary power converters. Thus, a separate heavy battery which would reduce the flight time is not necessary.

The drone is still flown manually in these experiments. However, the RS-232 connection to the drone would allow sending waypoints through the DJI Onboard SDK ROS. On top of commanding the drone, the same interface can also be used to record the state estimate from the drone's onboard state estimation. In our current configuration of the drone, IMU, RTK-GNSS, visual odometry from the onboard downward-facing camera and a barometer are fused. Apart from the drone's estimated pose, the LiDAR and IMU packets are also recorded for later offline mapping.




Fig. 3 The sensor pack is mounted at the front of a DJI M210 RTK, where normally a large camera could be attached. The sensor payload includes a LiDAR (dark blue), RGB camera (black), multispectral camera (red) and IMU (behind the LiDAR) connected to a Jetson TX2 for data gathering and processing. The components are mounted on a stiff 3D-printed frame and powered through the drone with a converter leaving out a heavy additional battery

In an offline processing step, the gathered LiDAR packets are fed into the `lidar_align` package (Taylor and Millane 2018) to find the transformation between the drone's origin and the LiDAR. This step should be carried out for each data recording, as the sensor can be unmounted and LiDAR and drone are not time-synchronized. `lidar_align` will find the transform between the drone and LiDAR as well as the corresponding time offset. Stitching the LiDAR packets together according to the drone's state estimation and the drone-LiDAR calibration will result in a large point cloud as shown in Fig. 4. The test field where the embankment prototype will be created is located in the white rectangle. The 3D point cloud is converted into a 2.5D elevation map using the `grid_map_pcl` package (Jelavic and Jud 2018), which will cluster the points per grid cell. The resulting elevation map is shown in the bottom right corner of Fig. 4.




Fig. 4 The location where experiments are conducted is shown in the large 3D point cloud gathered by the drone and consisting of 17 million points. The test field in the white rectangle is converted into a 2.5D elevation map for use in the design environment

Running a localization and mapping algorithm with the collected LiDAR scan would be an alternative, especially if a LiDAR is deployed on a cheaper, less capable drone, which might not provide an accurate state estimate. Examples for such an alternative would be an (ICP) approach or Cartographer developed by Google (Hess et al. 2016).

2.2 Design environment—docofossor

The design environment consisting of analysis, modelling and validation steps are performed on a 2.5D digital elevation model received from the surveying drone as an elevation map. COMPAS FAB [8] on the Windows side in combination with a `ros_bridge` in the Windows Subsystem for Linux is used to connect the design environment to the local network where messages to the drone and HEAP can be transmitted using Robot Operating System (ROS). The design environment itself runs in Rhino using computational terrain modelling tools based on the terrain modelling plugin Docofossor (Hurkxkens and Bernhard 2019), which was implemented by Hurkxkens (2020) for this interdisciplinary collaboration.

The terrain modelling tools enable dynamic updating of cut and fill operations based on distance functions. To be able to respond to changing site conditions in a dynamic fabrication environment, topographic designs are encoded parametrically. This is achieved by applying primitive shape representations in distance functions to terrain data. This allows for many Boolean operations on a single data set. Because the tools are written in Python, it can be used in the parametric modelling environment of Rhino Grasshopper or in COMPAS. Instead of having a signed distance field that equates to 0 at the surface, the height values determine the distance to 0, e.g. sea level. As such, the distance field is a simple height-map where the values are given by its elevation. The data structure





Fig. 5 2D illustration of the distance functions operating on the found geometry of the terrain

for the distance field now consists of the definition of the elevation map and accompanying z values. By applying a distance function on every point in the grid using an iterator, Boolean operations are calculated simply by shifting the grid-points up or down, while keeping the network topology intact.


Figure 5 shows how the distance functions result in the embankment design. Here we look at the 2D simplification where we make a swale on the left side of the input curve and an embankment on the right. Apart from the input curve, slope angles and two distances (dxB and dxC) are given. Newly designed elevation values are obtained by subtracting or adding the difference (dz) depending on whether a cut or fill is required. The full embankment design is controlled using the amplitude and length of a sine-wave function to achieve the full three-dimensional geometry with varying width and height, see Fig. 6a.

The adapted design line is used to create a material neutral design of an embankment which is visible in Fig. 6b as the blue line. Balancing cut and fill volumes, such that no excess soil is present or additional soil is needed, is therefore integral to this design. The final cut and fill areas are illustrated in Fig. 7. Apart from cut and fill balancing, parameters such as the soil swell factor and maximum achievable slope angles are also considered.

The connection of the design environment to the robotic platform over a network allows for tight integration of the design process before and during fabrication. It can receive an initial map and also get constant feedback on the fabrication progress, as shown in Fig. 2. This feedback can trigger changes in the design on the fly according to local conditions perceived by HEAP. For example, if HEAP discovers soil that can not be used to build an embankment with a particular slope, the design is adapted such that this newly found material can be integrated into the embankment. Or in case cut and fill volumes do not add up towards the end, the design can be changed accordingly. Such feedback can be often, e.g., after each digging cycle, or in the other extreme, never again after the initial design is created.



(a) The curved embankment is created using distance functions on the actual terrain elevation.



(b) The curved embankment is created from a single design line (black) that adapts to the actual terrain elevation (blue). The section view also shows the actual terrain before excavation, illustrating cut and fill regions.

Fig. 6 The robotic embankment design is created on the elevation map gathered with the mapping drone




Fig. 7 Illustration of the earth-movement in cut (red) and fill (blue). The computational design model automatically adapts to balance material volume with a 1.05 swell factor to the found geometry of the terrain

2.3 Fabrication environment—HEAP

The embankment is fabricated with the autonomous walking excavator HEAP developed in our previous work, where we transformed an off-the-shelf walking excavator into a

versatile, mobile manipulator (Jud et al. 2020). The fabrication environment block in Fig. 2 shows the required components for autonomous excavation, namely excavation planning, control, state estimation and mapping. Fig. 8 illustrates the process of a single excavation cycle in the context of the embankment experiment. In the following paragraphs, one excavation cycle is explained in detailed and visualized in Fig. 8. The following section numbers are related to the respective octagons in Fig. 8.

1. Connectivity: HEAP is connected to the design environment through WiFi and receives the desired elevation map and the corresponding design line while feeding back the current elevation map.

2. State Estimation: As both the survey drone and HEAP rely on GNSS-RTK for localization, HEAP is easily referenced to the design through our state estimation approach (Jud et al. 2020) computing the pose of the excavator.

3. Excavation Mapping: Excavation mapping, as shown in our previous work (Jud et al. 2019), generates a 2.5D elevation map of the current topography using LiDARs (Velodyne VLP-16 and Robosense RS-Bpearl) on the excavator's roof. The exteroceptive sensing (LiDARs) is fused with proprioceptive sensing (tracing the bucket edge as it moves through the soil) to quickly adapt the map to changes and increase robustness at the same time. Other possible disturbances, i.e., self-perception, falling soil, are handled as well.

4./5. Terrain-collision free trajectory to Point of Attack: The excavation planner finds the next point of attack for excavation. It is based on our work for free-form trenches (Jud et al. 2017), but further developed in this work to handle 3D embankments, as introduced in Sect. 3.1. Once the point of attack is found by the excavation planner, an end-effector trajectory free of terrain collisions is planned from the current location to either the point of attack or later on also to the dump point. It is an (SQP)-based planner using a (SDF) generated from the elevation map as shown in our previous work (Jud et al. 2019).

5. Single Dig Cycle: We use the soil-independent single dig cycle, as shown in Jud et al. (2017). It is an approach relying on force-control of the excavator's arm. In addition, a state machine can also switch to grading with an inverse kinematic arm controller to achieve a highly accurate and smooth surface (Jud et al. 2019).

6. Dumping: After the bucket is filled with soil, the excavation planner will find a suitable dump point, as described in detail in Sect. 3.2. Again, a terrain-collision free trajectory will guide the bucket to that location.




Fig. 8 HEAP, the autonomous walking excavator, is used to fabricate the embankment. The numbers 1–8 show a single dig cycle, 9 the finished part of the embankment with height lines for better visualization and 10 the unfinished part

7. Driving: HEAP can drive along a predefined path to reach faraway points of attack or dump points. More details on the driving are given in Sect. 3.

8. Active Chassis Balancing: While the excavator drives over undulated terrain, the hip balancing controller takes care of keeping all wheels in ground contact for maximum traction and stability. It relies on force-controlled leg cylinders and a blind balancing algorithm. The hip balancing controller is shut off while the excavator is digging and not driving to allow it to leverage the machine's weight for digging. Additionally, the active chassis plays a crucial role in high-accuracy excavation. It guarantees that the excavator has a stable stand with weight on all wheels such that no rocking motion is possible when the heavy arm moves.

3 Excavation planning For 3D embankments

The excavation planner receives the desired elevation map and the corresponding design line from the design environment. Figure 9 shows the designed embankment from a top view with the height encoded in grey colours together with the design line as the dotted line. The design is only consisting of the swale and the front face without the back face, as the excavator cannot reach the back face. However, it is considered in the design process for the volume calculations.

First, the driving path covering the entire design, shown as the dashed line in Fig. 9, is computed from the design line by offsetting it such that the excavator does not drive over the design, but also such that every point is still reachable

with the arm from the driving path. The excavation point planner, as introduced in the following Sect. 3.1, is then run to find the next point of attack for excavation. The end position of the driving path is found such that the arm will work orthogonally to the design line. Such an end position of the chassis with the bucket over the point of attack is shown with the excavator model depicted in Fig. 9. The reasoning behind the arm's working direction being orthogonal to the design line is illustrated in Fig. 10. The bucket edge will be parallel to the embankment break line without specifically enforcing this in the planner. After a successful excavation cycle, the dump point planner, as shown in Sect. 3.2 is run to find the dump point. In case that the dump point lies outside




Fig. 9 The driving path (dashed line) for a designed embankment is derived from the respective design line (dotted line) provided by the design environment. It is a constant offset line such that all points of the design are reachable from the driving path. The embankment height is encoded in grey colours (black: low, white: high)




Fig. 10 In order to achieve a clear break line between the swale and front face, the bucket has to pass over it with a parallel bucket edge meaning that the working direction of the bucket is orthogonal to the design line

of the reach, the excavator drives along the path until the point is reachable.

Driving on the path is controlled through a pure pursuit implementation in SE(2) (Jelavic 2020). It provides a driving speed and steering angle to the driving controller.

3.1 Excavation point planner

An excavation planner for 2D trenches was introduced in our previous work (Jud et al. 2017). It finds the next point of attack for excavation by scoring based on the distance to the machine and terrain height. For this work on embankments, we expand this work to 3D. The scoring still consists of the terrain height to excavate higher areas before moving to lower areas. However, the machine's distance is not a suitable scoring for large plans, as this will lead to behaviour that depends on where the machine is located or starts. For better regularization, a progression line for scoring is introduced as shown in Fig. 11. With the angled progression line, the machine works from left to right and favours the embankment's upper parts before working at the bottom as soil slides down the embankment.

The mathematical formulation of this approach is as follows. The error matrix $Z_e \in \mathbb{R}^{n \times m}$ between actual elevation $Z_{act} \in \mathbb{R}^{n \times m}$ and desired elevation $Z_{des} \in \mathbb{R}^{n \times m}$ is

$$Z_e = Z_{act} - Z_{des}. \quad (1)$$

We first construct a selection matrix $S \in \mathbb{R}^{n \times m}$ that selects only those cells that have an error higher than a certain threshold $\theta^+ \in \mathbb{R}^+$:

$$S^+(x_i, y_j) = \begin{cases} 1, & \text{if } Z_e(x_i, y_j) > \theta^+ \\ \text{NaN}, & \text{else.} \end{cases} \quad (2)$$




Fig. 11 The planner will choose the next excavation point to work on the biggest blob and neglect the smaller ones (marked blue) for now. This will increase efficiency as the iterative planner will jump less to different blobs. The scoring is computed from the progression line (blue lines) as well as the cell's terrain height. The point with the largest score will be the next point of attack and is encircled in blue. The score is represented through grey values. Black areas do not have a positive excavation error beyond a threshold

The x - and y -position of the point of attack (blue circle in Fig. 11) are the arguments of the maximum

$$\arg \max_{x_i, y_j} S_{lb}(x_i, y_j) S^+(x_i, y_j) (Z_{act}(x_i, y_j) + \mu(\cos(\beta)x_i + \sin(\beta)y_j)). \quad (3)$$

The first term Z_{act} in this equation will cause higher areas to be excavated first. The second term causes the machine to work from left to right along a progression line with the angle β . The two terms, that are depending on z -values and xy -values, respectively, are weighted through the factor μ . μ creates a slope where the machine decides to first remove material on top before moving further down. Fig. 11 shows only a progression line in 2D. However, as one can see from Eq. 3, it is actually a 3D progression plane defined by the two angles $\mu = \tan \alpha$ and β .

This planner is run iteratively without any planning that goes beyond one dig cycle. Thus, it can frequently happen that the machine returns to an already finished part of the embankment to fix some minor errors. Multiple reasons might cause this error, e.g., bad LiDAR measurements, mapping inaccuracies, falling soil. In order to reduce the planner deciding for subsequent point of attacks that are far apart, which reduces efficiency through longer driving distances, a second selection mask is introduced. The mask S_{lb} only keeps the biggest blob in the error map, as illustrated in Fig. 11. OpenCV's (Bradski 2000) *find-Contour* functionality is used to identify the error blobs. Thus, the machine will first work on the biggest pile of soil before taking care of smaller ones. This regularization has improved efficiency as the machine has to drive less when it is consistently working on the same pile. Larger piles are easier to excavate than smaller ones, and it is less likely to be stuck trying to fix some minor error without success.




Fig. 12 After thresholding the error map, error blobs with too small volumes are not considered as candidates for dumping (highlighted in blue). The dump point is then found through scoring with terrain height and distance to the bucket. Grey values represent the scoring, whereas black areas are not considered for dumping

3.2 Dump point planner

A motivation of this work is the ability to realize material neutral designs, which require planning not only for excavation but also for dumping soil. The strategy is to dump soil until a positive error is achieved such that the required accuracy can be produced in a subsequent excavation step. Only dumping soil without a subsequent excavation step will not lead to the desired accuracy. In our work on free-form trenches, the soil was simply dumped to a predefined spot to get it out of the way. This option is still used in this implementation if there is no space available any more on the design to dump soil. This could also be easily adapted to dump soil into a dump truck or processing plant. However, the design has balanced cut and fill volumes, and it is thus the more common case that there is still volume missing to complete the design. A similar approach to the excavation planner is used in that case. First, the error per cell is calculated as in Eq. 1. The selection matrix $S^- \in \mathbb{R}^{n \times m}$ selects those cells that have an error lower than a certain threshold $\theta^- \in \mathbb{R}^+$ and thus are candidates for dumping soil:

$$S^-(x_i, y_j) = \begin{cases} 1, & \text{if } Z_e(x_i, y_j) < \theta^- \\ \text{NaN}, & \text{else.} \end{cases} \quad (4)$$

These candidate cells are scored with two terms. First, the euclidean distance of the bucket to the cell $d(x_i, y_j)$ is computed. Dump points with a lower distance should obviously be preferred as this will lead to short dump cycles and thus increased efficiency. Secondly, the terrain height is also part of the score, as deep holes should be filled first. Figure 12 illustrates the scoring. Similar to Eq. 3,

$$\arg \min_{x_i, y_j} S_{iv}(x_i, y_j) S^-(x_i, y_j) (Z_{\text{act}}(x_i, y_j) + \mu d(x_i, y_j)) \quad (5)$$

finds the x - and y -position of the dump point with the same weighting μ as in Eq. 3, but a different additional selection

mask S_{iv} . This mask removes blobs from the error maps that represent a small error volume $< 0.1 \text{ m}^3$, roughly a fifth of the bucket volume. Fig. 12 shows small volumes that are removed, and in this case, only one large volume is considered. If the planner decided to dump a full bucket of soil in such a spot, it would make it worse than before, possibly requiring multiple excavation steps to clean it up again. A bit of accuracy is thus sacrificed for increased efficiency. Dumping less than the entire bucket is not a feasible approach.

4 Results

The robotic process was put to the test in two different experiments. First, an s-curved embankment was created. It highlights the new circular process of survey, design and fabrication. The design was adapted roughly two thirds through the fabrication process to re-balance cut and fill volumes. Also, the embankment's curved nature would be a tough challenge to build for a human operator on any level. The second experiment shows a two-faced embankment with a corner. It should highlight that this technology is also relevant and useful for current designs and not only for more complex curved shapes.

These experiments are validated through computing error maps as shown in Figs. 16 and 20 as well as an error metric for the entire design where all cell errors are combined in a mean error. The error per cell is the distance of the actual surface measured with the on-board LiDAR's to the desired surface from the design environment in the normal direction of the desired surface. Without using the normal direction of the desired surface, steep slopes would not be reflected correctly in the error metric.

4.1 S-curved embankment

The robotic process shown in Fig. 2 is showcased on an s-curve embankment. The video¹ accompanying this experiment summarizes the entire process from the initial drone flight to the finished embankment. The autonomous excavator operates with Level 4 autonomy Melenbrink et al. (2020) as stated in independent research by Melenbrink et al. (2020), as the optional operator in the cabin does not interact with the machine and only has a supervisory role for this research prototype.

The design can be seen in Fig. 5a and is composed to follow the site's general slope and catch any water runoff in the swale. The embankment slope is set to 33.7° and the total displacement volume to 30 m^3 . Although the design could have been steeper due to the clay-rich soil, the heavy

¹ <https://www.youtube.com/watch?v=Wjq3Nf9rWrM>.



Fig. 13 A top view of the finished embankment prototype shows the curved design and the smooth surface finish




Fig. 14 An elevation map of the s-curved embankment is created from the onboard LiDAR's after the experiment was successfully finished. The colours represent the terrain height

rain during the fabrication period led to a more conservative design. In this experiment, the first precision geometry was created autonomously with our novel robotic process in landscaping.

Figure 13 shows the final top view, Fig. 14 the elevation map at the end and Fig. 15 a time-lapse over the entire fabrication process. These illustrations show that a smooth surface finish, as well as a precise curvature, was achieved. The excavator started by design on the northern corner and progressed to the southern part. A total of 6 h of machine operation was necessary to cut 30.60 m³ and fill 32.20 m³ of soil. A small swell factor of 1.05 was used as some of

the piled-up soil was moved in the past years and thus not compacted over a long time. Nonetheless, digging the swale in front of the embankment went down into hard, compacted and undisturbed soil with small stones. The excavation system successfully handled these soil changes.

Regarding excavation accuracy, there are two unequally challenging parts in this design. The northern part, which is worked on first, is concave with a large radius. It is less challenging to fabricate than the southern part, consisting of a tight convex curvature that is inherently hard to reach with a flat edge shovel. The twofold design was chosen to show on the first part, the accuracy the process can achieve on complex curved surfaces and, at the same time, in the second part, where its limitations lie.

The average error per cell for the first concave part is 0.038 m. The error map in Fig. 16 shows mostly light blue colours representing errors close to zero in this area. This is in the same range as our previous, much more trivial experiments on free-form trenches where 0.027 m and 0.024 m average error was achieved. The overall average error is 0.056 m due to the previously mentioned challenges in the second part. Interestingly, areas with a positive error coincide with areas where cutting was necessary from the initial plan and areas with negative errors coincide with areas where filling was required. This seems logical as the thresholds for the excavation/dumping planner were chosen symmetrically as $\theta^+ = -\theta^-$. In future experiments, $\theta^- = 0$ would be a good choice to force a positive error and possibly a subsequent excavation step for a more equally distributed error over the entire embankment.

Some soil was lost when excavating the swale. It did not end up in the bucket after the bucket closing motion but was instead pushed towards and under the machine. Since the design is balanced regarding cut and fill volumes, the lost volume will be missing at the end of the embankment's fabrication. However, this is not a problem but rather a perfect opportunity to highlight the strength of tightly connected design and fabrication steps. Roughly after two-thirds of the fabrication, the design was altered to re-balance cut and fill volumes. Figure 17 shows the adapted plan. The swale cannot be changed as it is adapted to the overall terrain slope, and its form is functional regarding water runoff. The front face slope should also not be changed as it would otherwise be different from the already built rest of the embankment. However, the ridgeline can be lowered along the front face as illustrated in the section cut in Fig. 17, which removes a large volume (black and white stripes) at the backside of the embankment to compensate for the lost soil. The back face form is not critical and only there to support the front face and can thus be changed as required.



Fig. 15 A time-lapse with six images over the entire fabrication process shows the progress of the s-curved embankment




Fig. 16 The error map of the s-curve embankment shows missing soil (negative error) with yellow colours, excess soil (positive error) with dark blue colour and close to zero error with light blue




Fig. 17 The swale, to catch water runoff, and the front face are functional parts that should not be changed. The embankment volume can still be adapted to compensate an offset in cut and fill volumes by dropping the ridgeline along the front face. This will remove a large volume (black and white stripes) at the backside of the embankment. This is illustrated through a section cut at the point with the biggest correction

4.2 Two-faced embankment with a corner

In a second experiment, a two-faced embankment with a corner was fabricated as illustrated in the plan's rendering in Fig. 18. This experiment did not use the design environment. Instead, the embankment shape was drawn up in Blender and exported as a computer-aided design (CAD)




Fig. 18 The two-faced embankment with a corner is defined by a CAD model designed in Blender

model. The corner was realized by two individual plans for the two faces. The machine switched to the second one as soon as the first one was completed. Instead of using the high-precision servo valves, pilot stage driven main valves were used as this experiment solely needed grading cycles and no force-controlled digging. The reason being that it was performed in loose soil dumped from a truck. Additionally, the grading speed was significantly increased for a more efficient process.

For demonstration purposes, the finished embankment could be torn down by a human operator, and the autonomous excavator would fix the embankment again. The excavation planner would find the areas where soil was missing, identifies the areas with too much soil and move the soil accordingly. The process of destroying the embankment manually and rebuilding it autonomously made for a compelling demonstration. Figure 19 shows the finished two-faced embankment with a corner.

Figure 20 shows a top view of the error map at the end. An average error per cell of 0.032 m was achieved. A smooth surface finish and a clear corner highlight the high precision of this autonomous excavation demonstration. The figure shows positive errors with excessive soil in dark blue colours and negative errors with missing soil in yellow colours. However, the largest parts have only very little error, as illustrated by large patches with light blue colour.



Fig. 19 The image shows the final state of the two-faced embankment experiment. An overall smooth surface finish with a clear and straight corner was achieved




Fig. 20 The top view error map shows missing soil with yellow colours and excessive soil with dark blue colours. Light blue colours are errors around zero. The average error is 0.032 m

5 Conclusion

A unique collaboration between robotics and landscape architecture led to the fully autonomous high-accuracy fabrication of a free-form embankment where the design is material neutral and adaptive to local conditions. This has not been shown before. In detail, the robotic process demonstrated in this article tightly connects drone surveying, design and fabrication to exchange information in real-time during fabrication. The design can adapt immediately to local changes, e.g. incorporate large stones, change the slope according to the soil type, with either feedback from a surveying drone or the excavator used for fabrication. In

order to achieve such tight integration, we purpose-built a surveying drone for gathering a 3D point cloud from a LiDAR and subsequently transformed into an elevation map. RGB and multispectral cameras will add additional information on the soil characteristics in the future. The fabrication is realized with HEAP, the world's first autonomous walking excavator capable of high precision excavation. For this work on 3D embankments, an excavation planner was developed that iteratively plans excavation cycles as well as points to dump the soil. Planning for both excavation and dumping allows for fabricating material neutral designs.

Improvements on efficiency and productivity will be necessary. 6 h of machine operation for a 20 m long embankment moving 30 m³ soil is well below any human operator's performance. Also, larger plans will need the deployment of multiple machines. It is not yet clear how we will solve the collaboration problem. Not only could there be multiple machines of the same type, but also machines of different types working in a collaborative manner as seen on today's large construction sites.

Currently, the surveying drone only feeds information into the design process, but not to HEAP fabricating the embankment. Deploying the drone autonomously during the fabrication step and supplying information also to HEAP, would allow the excavator to make better decisions as some of the blind spots, e.g. areas behind the embankment, can be perceived and correctly handled. This will become even more critical for larger designs.

The experiments brought forth in this article demonstrate how computational design and robotic fabrication can respond dynamically to ongoing changes in natural granular materials. As such, the construction approach with HEAP changes the way we can design and construct with terrain. Apart from construction efficiency, it gives future landscape architects more geometric freedom in the design of landscapes to create innovative grading and drainage strategies and could provide dynamic maintenance along road embankments and rivers where equilibrium in terrain can not naturally be sustained. As such, the potential of autonomous earth-moving tasks is recognized in its ability to mediate between existing and desired states of terrain and how it opens up landscape design to new cyber-physical explorations.

Acknowledgements This work was supported in part by the Swiss National Science Foundation through the National Centre of Competence in Digital Fabrication (NCCR d4b), Hexagon Geosystems, armasuisse Science and Technology and the Landscape Visualization and Modelling Lab (LVML).

Funding Open Access funding provided by ETH Zurich.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Built Robotics and MPC Kinetic in Queensland. <https://www.i-q.net.au/main/autonomous-excavators-on-site>. Accessed: (2020-12-01)
- Bar-Sinai K.L., Shaked T, Sprecher A (2019) Informing Grounds: A Theoretical Framework and Iterative Process for Robotic Groundscaping of Remote Sites. In: Proceedings of the 39th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA), pp. 258–265
- Bradley D, Seward D (1995) Developing real-time autonomous excavation - the LUCIE story., pp. 3028–3033
- Bradley DA, Seward DW (1998) The development, control and operation of an autonomous robotic excavator. *Journal of Intelligent and Robotic Systems* 21(1):73–97
- Bradski G (2000) The OpenCV Library. Dr. Dobb's Journal of Software Tools
- Cannon H (1999) Extended Earthmoving with an Autonomous Excavator. The Robotics Institute PhD
- Chui M, Mischke J (2019) The impact and opportunities of automation in construction. <https://www.mckinsey.com/~media/McKinsey/Industries/Capital%20Projects%20and%20Infrastructure/Our%20Insights/The%20impact%20and%20opportunities%20of%20automation%20in%20construction/The-impact-and-opportunities-of-automation-in-construction.pdf>
- COMPAS FAB: Robotic Fabrication for COMPAS. https://gramaziokohler.github.io/compas_fab/
- Gramazio F, Kohler M, Willmann J Thriving Digital Materiality. In: The Robotic Touch: How Robots Change Architecture. Park Books
- Ha Q, Santos M, Nguyen Q, Rye D, Durrant-Whyte H (2002) Robotic excavation in construction automation. *IEEE Robotics Automation Magazine* 9(1):20–28. <https://doi.org/10.1109/100.993151>
- Hess W, Kohler D, Rapp H, Andor D (2016) Real-time loop closure in 2d lidar slam. In: 2016 IEEE International Conference on Robotics and Automation (ICRA), pp. 1271–1278. <https://doi.org/10.1109/ICRA.2016.7487258>
- Hurkxkens I (2020) Robotic landscapes: Topological approaches to terrain, design, and fabrication. Ph.D. thesis, ETH Zurich, Zurich. <https://doi.org/10.3929/ethz-b-000451100>
- Hurkxkens I, Bernard M (2019) Computational terrain modeling with distance functions for large scale landscape design. *Journal of Digital Landscape Architecture* 2019(4):222–230
- Hurkxkens I, Bernhard M (2019) Docofossor. <https://www.food4rhino.com/app/docofossor>
- Hurkxkens I, Mirjan A, Gramazio F, Kohler M, Girot C Robotic Landscapes: Designing Formation Processes for Large Scale Autonomous Earth Moving. In: C. Gengnagel, O. Baverel, J. Burry, M. Ramsgaard Thomsen, S. Weinzierl (eds.) *Impact: Design With All Senses*, pp. 69–81. Springer International Publishing. https://doi.org/10.1007/978-3-030-29829-6_6
- Jelavic E (2020) Planning and control for car-like vehicles. https://github.com/leggedrobotics/se2_navigation
- Jelavic E, Jud D (2018) Grid map pcl. https://github.com/ANYbotics/grid_map/tree/master/grid_map_pcl
- Jud D, Hottiger G, Leemann P, Hutter M (2017) Planning and control for autonomous excavation. *IEEE Robotics and Automation Letters* 2(4):2151–2158. <https://doi.org/10.1109/LRA.2017.2721551>
- Jud D, Leemann P, Kerscher S, Hutter M (2019) Autonomous free-form trenching using a walking excavator. *IEEE Robotics and Automation Letters* 4(4):3208–3215. <https://doi.org/10.1109/LRA.2019.2925758>
- Jud D, Wermelinger S.K.M., Jelavic E, Egli P, Leemann P, Hottiger G, Hutter M (2020) Heap - the autonomous walking excavator. Unpublished, submitted to *Automation In Construction*
- Maeda GJ, Manchester IR, Rye DC (2015) Combined ilc and disturbance observer for the rejection of near-repetitive disturbances, with application to excavation. *IEEE Transactions on Control Systems Technology* 23(5):1754–1769. <https://doi.org/10.1109/TCST.2014.2382579>
- Maeda G.J., Rye D.C (2012) Learning disturbances in autonomous excavation. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 2599–2605. IEEE
- Maeda G.J., Rye D.C., Singh S.P (2014) Iterative autonomous excavation. In: *Field and service robotics*, pp. 369–382. Springer
- Melenbrink N, Werfel J, Menges A (2020) On-site autonomous construction robots: Towards unsupervised building. *Automation in Construction* 119:103312. <https://doi.org/10.1016/j.autcon.2020.103312>
- Park J, Lee B, Kang S, Kim PY, Kim HJ (2017) Online Learning Control of Hydraulic Excavators Based on Echo-State Networks. *IEEE Transactions on Automation Science and Engineering*
- Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. Standard, Society of Automotive Engineers (2018). https://www.sae.org/standards/content/j3016_201806/
- Schmidt D, Berns K (2015) Construction site navigation for the autonomous excavator thor. In: 2015 6th International Conference on Automation, Robotics and Applications (ICARA), pp. 90–97. <https://doi.org/10.1109/ICARA.2015.7081130>
- Schmidt D, Proetzsch M, Berns K (2010) Simulation and control of an autonomous bucket excavator for landscaping tasks. In: 2010 IEEE International Conference on Robotics and Automation, pp. 5108–5113. IEEE
- Seo J, Lee S, Kim J, Kim SK (2011) Task planner design for an automated excavation system. *Automation in Construction* 20(7):954–966
- Singh S, Simmons R (1992) Task planning for robotic excavation. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* 2:1284–1291. <https://doi.org/10.1109/IROS.1992.594551>
- Taylor Z, Millane A (2018) lidar_align: a simple method for finding the extrinsic calibration between a 3d lidar and a 6-dof pose sensor. https://github.com/ethz-asl/lidar_align
- Westort C (1998) Methods for sculpting digital topographic surfaces. Phd thesis, ETH Zurich
- Zhao J, Long P, Wang L, Qian L, Lu F, Song X, Manocha D, Zhang L (2020) Aes: Autonomous excavator system for real-world and hazardous environments

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.