Autonomous Multi-Material Construction with a Heterogeneous Robot Team

Article *in* Robotics and Autonomous Systems · September 2019

DOI: 10.1016/j.robot.2019.07.009

3 authors:

Maira Saboia da Silva
California Institute of Technology

Vivek Thangavelu
Cornell University

N. Napp
University of Washington Seattle

Autonomous Multi-Material Construction with a Heterogeneous Robot Team

Maira Saboia, Vivek Thangavelu and Nils Napp

University at Buffalo, Buffalo NY 14260, USA,
{mairasab, vsangara, nnapp}@buffalo.edu

Abstract. We present a construction model that allows robots with different construction capabilities, using materials of different physical properties and sizes, to modify unstructured environments in a distributed system. Building steps are computed reactively so that they can respond to changes in the environment and imperfect assembly. The reactive approach allows robots to coordinate and add material to the same structure. Each robotic agent uses an abstract model of the environment to compute a set of legal construction steps based on its current knowledge of the world, and we show that in this setting more knowledge results in more legal moves. We exploit this capability by letting the system use a variety of materials and choose the most appropriate material given its current knowledge of the state of the structure. We demonstrate the approach by running the system on a variety of terrains and with mixed materials, including both deformable and rigid components.

Keywords: Autonomous Construction, Collective Robotic Construction, Multi-Material, Partial Functions, Local Sensing.

1 Introduction

In collective robotic construction, the goal is for teams of robots to modify their environment in predictable and useful ways [1]. As with autonomous systems in general, one of the main challenges is to ensure reliable operations in unstructured environments. In situations where autonomy is important, the cost of failure is often also high and frequently coincides with irregular terrain and lack of established infrastructure, for example, in disaster response and operation in areas inaccessible to humans, such as extraterrestrial bodies. We present a distributed, autonomous, heterogeneous robot team that can autonomously modify such environments in order to provide mobility. The robots adaptively build structures over unstructured terrain and choose between different material types in order to effectively incorporate preexisting terrain features.

Animals provide many examples of multi-material and adaptable constructions [2]. They modify their environment to provide protection from weather and predators, to regulate moisture and temperature, to ease travel along foraging routes, and to store food. The exact shapes of the final structures adapt to fit into the preexisting environment. For example, no two bird nests are exactly

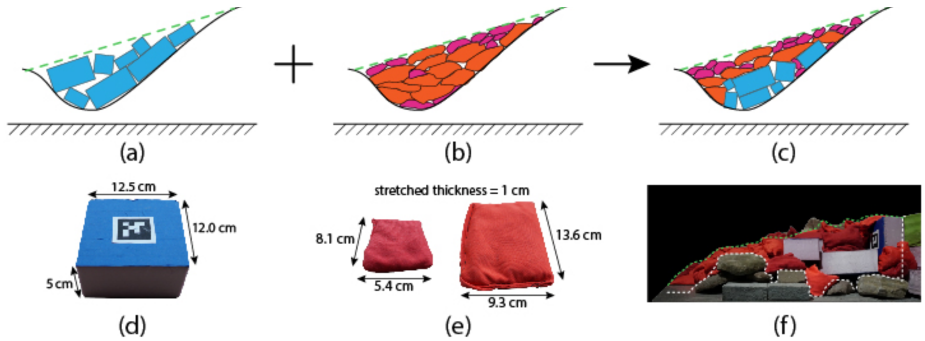

Fig. 1. We present a model and a practical robotic implementation for construction with mixed materials, including rigid foam blocks and compliant bean bags. a) Blocks are relatively large, so construction with them can quickly fill the space. However, due to large discontinuities the structures might not be navigable. b) Using compliant bags results in navigable structures, but can be slow to construct. c) Construction with a mix of materials may leverage the advantages of both. d-e) Blocks and bags used as building materials. f) Cross-sectional view of a multi-material structure built by a robot following the model presented in this paper. The green dotted line show the MARS function; the white dotted line show the initial structure.

alike, but they are a robust solution to their intended function, i.e., keeping eggs warm, hidden, and in place. In biological construction, there are a variety of ways to encode built structures. Here we focus on strategies that are particularly robust and adaptable: function-driven, iterative construction that relies on stigmergy [3]. Stigmergy is a biological phenomenon that has been adopted to robotics, referring to environmentally mediated communication where information about building actions is encoded in the partially built structure. For example, it is known that beavers build in response to the sound of moving water [4]. Since partially built dams restrict flow, building actions of some individuals inform the building actions of others through the environment. Additionally, the cue is directly tied to the function of a dam, which should stop water from moving. The simple way of tying building actions to the function of the resulting structure results in adaptive and robust building behavior. Our system uses the same principle to allow indirect communication between the robots. Each robot decides independently what changes will be made to the structure, and new state after modification might trigger additional building behavior from any of the construction robots. This approach results in a distributed feedback mechanism to achieve a structural function.

The notion of *navigability* introduced in [5] ties the representation of irregular terrains to a robot motion capabilities. It enables a robot to compute the legal building actions given the state and the functional specification of a structure, which is to allow mobility in unstructured environments. In this model, motion is expressed in terms of three scalar parameters, that can provide a conservative lower bound of the motion ability based on the stability and clearance

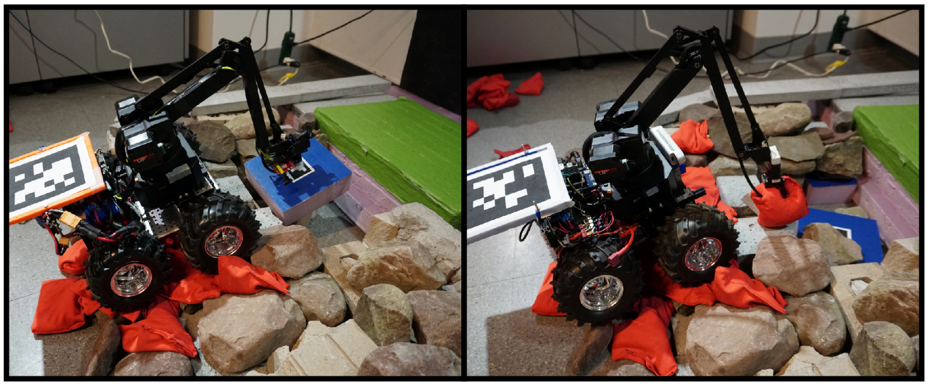

Fig. 2. Robots in action. (left) Foam block handling robot. (right) Bag handling robot.

constraints of the mobile platform. A *deposition model* is used to provides an upper bound on how much the environment changes in response to its actions. Together, the navigability and deposition model allow the robot to adaptively synthesize low-level motion and manipulation plans that are guaranteed to produce navigable structures as long as the environmental changes produced by the robots are consistent with the deposition model at each step. Navigability parameters induce lower bounds on both the allowable deposition models and sensing abilities of a robot. This abstract model for the construction process directly links the motion, sensing, and manipulation abilities of a robot.

1.1 Contributions

The contribution of this paper is two-fold. First, we explicitly model locally sensed data as partial orders, which allows us to easily compute the *maximal* legal modifications a robot is allowed to make while obeying the same guarantees of global minimality as in [6]. Directly computing legal moves based on the available sensing information allows robots much flexibility of construction materials.

Second, we present an autonomous system, for building motion support structures, that uses this abstract construction model to coordinate a team of robots, each with different construction abilities (Fig. 2). We demonstrate that the model allows easy composition of behaviors (Fig. 1): one similar to the behavior in [6] that ensures completion by using small deformable materials and one that uses the new way of computing legal actions to use more volume-efficient rigid materials. The system is both better at building large structures, as it has more freedom to choose materials, and to the best of our knowledge is also the first demonstration of a multi-robot system that adaptively and autonomously decides between a mix of deformable and rigid materials to increase system performance.

1.2 Related Work

Nature-inspired construction systems based on stigmergy have been explored both from an algorithmic perspective [7] and in physical implementation systems [8] since the mid-1990s. The key idea is that the building plan is encoded in rules that agents use when they respond to intermediate build states. The application of these rules, in turn, defines the next possible intermediate states, which can then trigger more rules. This process is thought to govern coordination in swarm construction systems [9], e.g. termite mounds, where no individual agent has an internal estimate of the structure. Work that uses this approach to enable construction often focuses on stigmergy as a process to coordinate agents in a distributed system, e.g. [10,11,12,13,14]. Here we focus on two other properties that many natural construction systems exhibit that stigmergic approaches can provide: adaptability and robustness.

Using this approach is difficult in practice since synthesizing stigmergic rules to build specific structures is challenging, and the contribution of works such as [11,13] is to generate rule sets that can build specific shapes. In order to fully exploit the adaptability, we also use a functional specification to describe the target structures, which is mobility. This has been explored in a theoretical setting [15] and is given as an explanation for animal-built structures that have direct biological functions [3].

Autonomous robotic systems may be broadly categorized based on the number and type of the robots involved, the building material used or the environment they operate on. Most of the prior works fall into the category of distributed (ground and aerial) robots that work on highly processed, specialized building materials and handcrafted rules, operating in a structured environment. Werfel et al. use stigmergy in [11] to build a 3D structure inspired by mound-building termites [12], where a team of identical robots receives a set of low-level rules that collectively produce a specific structure using prefabricated solid bricks. While navigating over the structure, the robots attach bricks in positions that at the same time obey a set of *geometric requirements* and are valid according to the structure plan. Allwright and his collaborators used the same mechanism to conceive the stigmergic block [13]. The robots “mark” the blocks by using NFC (Near Field Communication) to toggle LEDs. The patterns left on the blocks can stimulate construction actions [16]. Although they have a purely reactive model of construction, they suffer from the above mentioned limitations.

In the work of Soleymani et al. [17], a self-contained ground robot builds a protective barrier using compliant pockets. The final shape of the structure is specified via a template and the probability of a certain location to be chosen for deposition is inversely proportional to the *number of pockets* in that location. The work focuses on scalability and demonstrates their approach through experiments and extensive simulation. Another example involves that of Andreen et al. [18], where hundreds of mechanically programmed robots merge foam blocks into structures with pre-defined properties. Both systems are limited to working in structured environments, and the resulting structures are not navigable by the robots. Furthermore, while the former approach clearly demonstrates the use of compliant materials, it does not take advantage of their ability to be easily used on irregular terrain. The work of Fujisawa et al. describes a robot that can modify its environment by depositing foam, which expands and turns rigid [19]. Their work demonstrates the ability to build structures over regular obstacles, but does not provide correctness proofs for the deposition strategy.

Melenbrink et al. show in [14] the ability to build unsupported cantilevers across a gap using force feedback. By looking at how loads transfer to the ground they can anticipate and possibly prevent tipping over by building a structure for counterbalancing. In the work presented by Napp et al. [5], a robot was used to autonomously build a ramp by extruding foam in an unstructured environment. The authors focused on a strategy for adaptive ramp building using a reactive algorithm that iteratively fills *non-navigable* gaps and ledges in the ramp structure. A complementary approach to enabling mobility in unknown environments is presented by Tosun in et al. [20], in which a high-level task planner coordinates modular robots that plan and execute augmentation tasks. This system is autonomous and adaptive, in that it can generate new plans in response to new environments. However, generating task primitives in unstructured terrain is difficult.

To the best of our knowledge, only the works presented in [5] and [14] are able to construct in unstructured environments, however, neither of these present planning and execution in a full 3D state space; Melenbrink et al. [14]. is a 2D simulation, and Napp et al. [5] presents an extrusion of 2D planning and computation into 3 dimensions, and it does not incorporate manipulation and motion plans. In our prior work [6], we adapt the reactive building strategy of [5] to identify *non-navigable* features in a 3D structure, and extend it to incorporate manipulator and motion constraints in the building strategy. The paper contributes a property-driven deposition algorithm to autonomously achieve and maintain navigability conditions over irregular 3-dimensional terrain.

2 Problem Formulation

Navigability [5,6] provides a method to tie robot specific kinematic constraints to the irregular terrain model and give a concise mathematical way to express a set of poses that robots can occupy. After defining navigability, this section describes a method for efficiently computing what changes need to be made in an environment to build a *Perfect Motion Support Structure*, which is defined as a projection operator. We show how this computation is affected by having only limited, partial knowledge of the current structure. Together, these ideas are used in the next section to produce a high-level algorithm for distributed, multi-material construction.

Consider the robot of Fig. 3(b), where $\kappa \in \mathbb{R}^+$ is maximum climbable slope the robot can drive up or down, $\delta \in \mathbb{R}^+$ is the robot body length and $\epsilon \in \mathbb{R}^+$ is maximum discontinuity that a robot can drive over per body length. This discontinuity may present itself in different forms (see examples in Fig. 3(c)), and can be measured empirically or analytically, depending on the robot. This formulation allows for a generalized notion of navigability adaptable to any robot

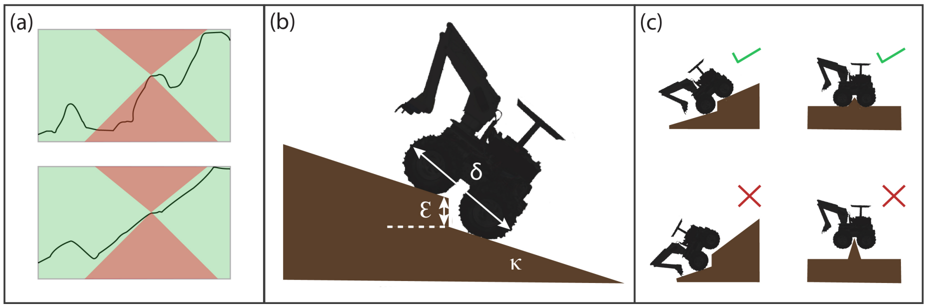

Fig. 3. (a) For Lipschitz continuous functions, there is a double cone (shown in red) whose vertex can be translated along the structure function, such that the structure function remains entirely outside the cone. Top picture is not Lipschitz with the given constant and the bottom structure is. This type of continuity is a way to measure the overall steepness of arbitrary functions. (b) Depicts the robot navigability parameters in the shaded 2D plane space on a upward slope of maximum steepness κ. (c) Examples of valid and invalid robot configurations.

using only three parameters. In our implementation we use: $\kappa = 0.314$, $\delta = 0.30$ and $\epsilon = 0.048$.

We model the environment as continuous functions to allow us to effectively express arbitrary terrains. Consider the construction area Q, as a compact, simply connected, and finite subset of $\mathbb{R}^2$. It is the domain of a bounded, non-negative height function $h : Q \rightarrow \mathbb{R}^+$ which describes a *structure*. h is fully navigable by a robot if

$$|h(\mathbf{p}) - h(\mathbf{q})| \ \leq \epsilon + \kappa|\mathbf{p} - \mathbf{q}|, \ \forall \ \mathbf{p}, \mathbf{q} \in Q \ \text{ s.t. } |\mathbf{q} - \mathbf{p}| < \delta, \tag{1}$$

where $|\cdot|$ represents the Euclidean distance between two points.

Formally, a structure is *navigable* if it is locally (parameter δ) close (parameter ϵ) to K-Lipschitz continuous. The operator $P_\kappa[h]$ projects any structure to the smallest function in $\mathcal{L}_K$, the space of K-Lipschitz functions (Fig. 3(a)) on Q, that is at least as large as h:

$$P_\kappa[h](\mathbf{p}) = \max_{\mathbf{q} \in Q} \{h(\mathbf{q}) - \kappa|\mathbf{q} - \mathbf{p}|\}. \tag{2}$$

In Eq. (2), the value of h at $\mathbf{p}$ of the projected function is sometimes determined by a different location $\mathbf{q}$ in the domain. The set that determines the values of $P_\kappa[h]$ is called the P_κ-*support* of h. For any value $\mathbf{q} \in P_\kappa$-support, $P_\kappa[h](\mathbf{p}) = h(\mathbf{p})$, i.e. $h(\mathbf{p})$ determines its own value in the projection. Additionally, when the max in Eq. (2) is achieved for multiple $\mathbf{q}$, only the one with the largest value of $h(\mathbf{q})$ is included in the P_κ-support. Two functions $h, h' \in Q \rightarrow \mathbb{R}^+$ that agree on their P_κ-support have the same projection $P_\kappa[h](\mathbf{p}) = P_\kappa[h'](\mathbf{p}), \forall \mathbf{p} \in Q$. Further, applying P_κ to a function does not change its P_κ-support, since the values that are exactly K-Lipschitz with other points are determined by the larger of the pair.

The deposition operator, P_κ, computes the Minimal Additive Ramp Structure (MARS), i.e. the minimum volumetric amount of material that needs to be deposited at each point to make a perfect (without discontinuities) motion support structure, which are generalizations of simple ramps. Let h_0 be the original structure, P_κ is used in [5] to prove that a navigable structure can be iteratively built on Q that is bounded above by $P_\kappa[h_0]$. The algorithm works by allowing robots to build approximate MARS that are smaller but still navigable based on the navigability parameters, i.e. the capability of driving over small discontinuities that allows the use of imperfect structures.

Mobile robots are typically equipped with sensors that only provide local/partial information about their environment.

We use *partial functions* to model this behavior [21]. For some subset $Q' \subset Q$ of the building area, a partial function of the structure is defined only in this particular subset, $h' : Q' \to \mathbb{R}^+$. Each different subset of Q defines a different partial function and the set of all partial functions from Q to $\mathbb{R}^+$ is denoted by $Q \multimap \mathbb{R}^+$.

To give a notion of the amount of information that a robot has access to, we compare the domains of partial function and use the subset inclusion order [21] to say that a robot has more information if it can measure subsets of the building domain that are strictly larger. Overlapping subsets that are not contained in each other cannot be compared directly.

3 Abstract Model of Construction

Our abstract model of construction operates by iteratively modifying the existing structure, and thus producing a sequence of substructures, $h_0, h_1, h_2, ...$, in which h_n represents the height function after n depositions. At each step, an agent computes the maximum amount of change that can be made to turn an arbitrary structure into a perfect one i.e. the smallest dominating K-Lipschitz function that can be achieved by adding material.

In order to operate on local sensor information, we compute a local MARS bound by applying the operator $P_\kappa[h]$, Eq. (2) to subsets of Q. The observation subset $Q_r \subseteq Q$ defines where the partial function of the perfect structure $P_\kappa[h]|_{Q_r}$ is computed:

$$P_\kappa[h]|_{Q_r}(\mathbf{p}) = \max_{\mathbf{q} \in Q_r} \{h(\mathbf{q}) - \kappa|\mathbf{q} - \mathbf{p}|\}. \tag{3}$$

When subsets of Q are ordered by set inclusion they induce an order relation between the partial functions of P_κ by

$$P_\kappa[h]|_{Q_r} \leq P_\kappa[h]|_{Q_R} \text{ iff } Q_r \subseteq Q_R \text{ and } P_\kappa[h]|_{Q_r}(\mathbf{p}) \leq P_\kappa[h]|_{Q_R}(\mathbf{p})\ \forall \mathbf{p} \in Q_r.$$

Robots typically lack global knowledge of the current structure state, and we show in Lemma 1 that the MARS bound is a monotonic function of the sensing domain size, as shown in Fig.4. An agent with less information computes a smaller bound and is allowed to build less than an agent with more information. In Theorem 1, we show that modifications based on the local MARS bound are consistent with the global bound.

Lemma 1 (Monotone). *The operator P_κ restricted to subsets of Q is monotone (order-preserving), i.e. $Q_r \subseteq Q_R \subseteq Q$ implies $P_\kappa[h_r]|_{Q_r} \leq P_\kappa[h_R]|_{Q_r}$.*

Proof. Consider $Q_C = Q_R \setminus Q_r$, i.e, the set of elements in Q_R that are not in Q_r. By expanding, rearranging and substituting $P_\kappa[h]|_{Q_r}$ and $P_\kappa[h]|_{Q_C}$ into $P_\kappa[h]|_{Q_R}$, we find that

$$P_\kappa[h]|_{Q_R}(\mathbf{p}) = \max\{P_\kappa[h]|_{Q_r}(\mathbf{p}),\ P_\kappa[h]|_{Q_C}(\mathbf{p})\}. \tag{4}$$

Hence for every $\mathbf{p} \in Q_R$, the value $P_\kappa[h]|_{Q_R}(\mathbf{p})$ is at least $P_\kappa[h]|_{Q_r}(\mathbf{p})$, but is greater when $P_\kappa[h]|_{Q_r}(\mathbf{p}) < P_\kappa[h]|_{Q_C}(\mathbf{p})$. □

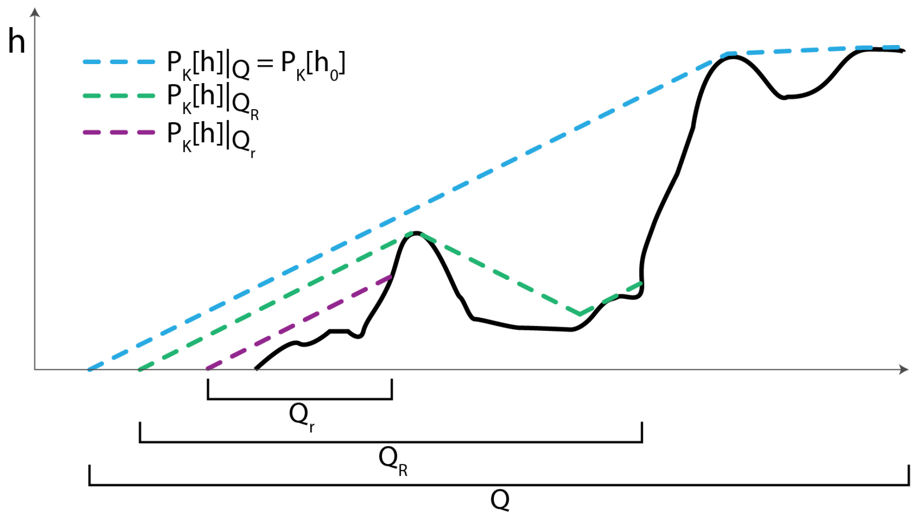

Fig. 4. Illustration of Lemma 1. The different functions are the MARS bounds, $P_\kappa[h|_Q]$, computed on increasingly large subsets. The larger the domain of partial knowledge, the larger the MARS bound.

Theorem 1 (Consistency). *When a robot's iterative depositions are bounded above by P_κ applied to its limiting sensing range Q_r the final global structure is bounded above by $P_\kappa[h_0] = P_\kappa[h]$.*

Proof. Applying P_κ to a structure $h \in Q \to \mathbb{R}^+$, $h' = P_\kappa[h]$, does not change its P_κ-support. Although $P_\kappa[h'](\mathbf{q}) = h'(\mathbf{q})$ everywhere, the max achieved in (3) has a largest point that determines the height of all the locations that are exactly K-Lipschitz with it. This means that $P_\kappa[h_0] = P_\kappa[h_n]$ for all intermediate structures since the global support set does not change. By monotonicity, each local MARS bound produces smaller structures, so they cannot be larger than the global P_κ-projection. Therefore, the global bound holds even if it is only checked locally.□

Theorem 2 (Maximality). *If the robot deposits more than the local MARS bound, i.e. P_κ applied to a partially observed structure $h' \in Q \multimap \mathbb{R}^+$, the final global structure can have points outside the global MARS bound $P_\kappa[h_0](\mathbf{p})$.*

Proof. We need to show that violating any of the MARS bounds computed on partial information can lead to violating the global upper bound. Given some local sensing set $Q_r \subset Q$ that includes points from the P_κ-support of h_0, i.e. the support set of the global structure, then adding ϵ material to point $\mathbf{q}$ that is in the global P_κ-support, will be ϵ outside the global MARS bound. □

This means that for some points, even small changes can affect the global bound. However, without global knowledge they cannot be determined. The P_κ-support for partial structures can include points that are not in the global P_κ-support, so violating the local bounds does not always cause a problem. However, all global P_κ-support points are also P_κ-support points in partial views of the structure, so that staying

Table 1. Building Material Specifications

Type	*Dimensions* ([length×width×height])	*Weight*
Large bags	[12.4cm, 9.6cm, 1.0cm] (±11.4%, ±14.1%, ±10.1%)	71.2g (±17.5%)
Small bags	[6.7cm, 6.05cm, 0.9cm] (±0.25%, ±0.19%, ±0.01%)	34.1g (±2.1%)
Foam block	[10.68cm, 10.34cm, 5.0cm] (±28.2%, ±23.6%, ±0%)	19.9g (±12.3%)

within the local MARS bound preserves the global bounds, and so that it produces the best maximal choice given specific local information. The next section describes how this abstract model of iterative construction maps to a physical implementation.

4 System Implementation

This section describes the physical robot and the details of how the abstract algorithm maps to a physical implementation. The system is used in experiments which are designed to evaluate different material combinations, coordination configurations and terrain conditions.

4.1 Experimental Setup

The robots used for the construction task are identical except in the design of their end effector; one is equipped with a gripper to handle compliant bags and the other with a suction cup to handle foam blocks (as shown in Fig. 2). Each robot is a low-cost mobile manipulator made from off-the-shelf components, capable of maneuvering over irregular terrain.

An AprilTag [22] is mounted on top of the robot for pose estimation. A global 2D occupancy grid map is maintained for motion planning using a single overhead Kinect camera. Depth data are used to get the voxelized representation of the construction area, Q. The perspective of the robots are simulated using the subset of the data from the area surrounding the navigable region. However, because the data are acquired from a overheard perspective, the calculated height function is presumably different from the height function that would be computed if the system implements a local perspective from on-board camera.

An OpenMV M7 Camera, fixed on the end-effector of the arm, is used for foam pickup. As the navigable region changes, new information about the structure is utilized.

4.2 Building Material

The bags used in our experiments are filled with beans and are categorized based on their sizes: large and small bags (see Fig. 1(e)) which are used for coarse and fine depositions, respectively. The foam blocks are cut out from polystyrene insulation foam blocks, see Fig. 1(d). The the mean weight and stretched dimensions [length×width×height] of the various materials are listed in Table 1.

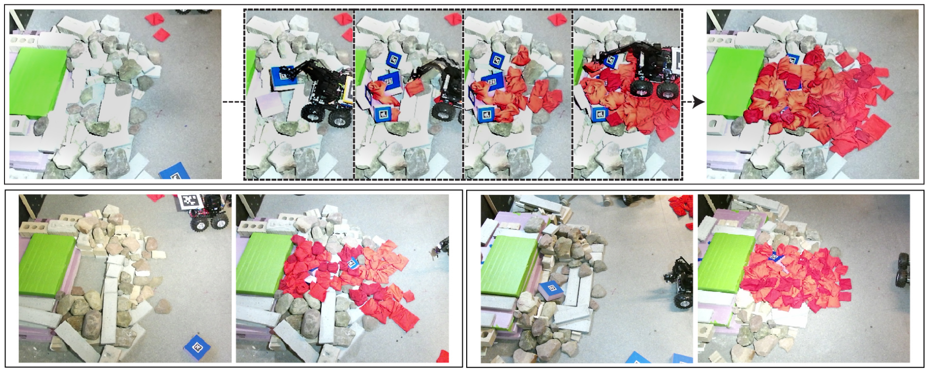

Fig. 5. Overhead view of the construction area. The green block is the target location. (top) The initial terrain is shown on the left and the final structure on the right, with the intermediate states of the structure in between. These states show the robot having to drive over the partially built structure in order to reach deposition sites. (bottom) Two image pairs of initial and final structures for different terrain configurations. The different objects, bags and foam blocks, that the robot can pick up and use are scattered throughout the construction area.

The robot is deployed at a position that is navigable, and from where it has access to an unlimited supply of building materials. The material supply is located at a distance from the construction area such that it is not considered part of the structure. Due to the limited size of the construction area, the robots take turns in modifying the structure, see Sec. 4.4 for details on deciding which robot is *active* and allowed to modify the structure. We present a more detailed description of the robot design and the construction system in [6].

4.3 Discretization

In our system, we work on a discretized version of the workspace. In this discrete representation, we assume the construction grid area G (discretized version of Q) is a finite set of $\Delta \times \Delta$ cells $u = (i, j)$ in $\mathbb{N}^2$ parallel to the construction building area. An individual voxel is given by $v = \{u, l\}$ in $\mathbb{N}^3$, where u is the grid cell and l is its discretized height. A given continuous height function $h(q)$ can be represented as an occupancy grid in the voxel space.

This discretized view of the world will work for the robot, even if the final structure representation is constrained by voxel resolution. The choice of discretization length should be within the discontinuity limit (ϵ) of the robot's navigability constraints i.e $\Delta < \epsilon$, and Δ should be within the minimum possible dimensions of a single deposition.

4.4 Behavior Composition Strategy

The *Behavior Composition Strategy* is a construction approach that dynamically chooses the construction material according to the most appropriate behavior given the current state of the structure. We present three different deposition behaviors using coarse, fortifying, and fine depositions, that are associated with three types of building materials made of foam blocks, large bags, and small bags, respectively. We propose a *priority-based* approach that coordinates the building process so that certain types of building materials are favored, depending on the state of the construction process. Priorities are assigned to each deposition behavior which imposes an order on the robots; the coarse behavior has higher precedence than the fortifying behavior, which in turn has higher precedence than fine behavior. In the remainder of this section, we discuss the framework for material selection and subsequently define the various deposition behaviors.

Due to the limitations of the rover arm and the lack of any explicit physical stability modelling of the depositions, we opt for a more conservative deposition strategy. The robots only plan positioning of the material units center of mass. However, one may design a more general composition and deposition strategy that can take into account the poses of depositions.

We define the *Action Space* $X \subset Q$ of an object m, of height h_m, as the set of valid positions $\mathbf{x}$ where m can be placed on a given structure. Given a position $\mathbf{x} \in X$, the 3D deposition point is $(\mathbf{x}, h(\mathbf{x}) + h_m)$. However, because each type of building material serves a specific purpose, the choice of the most appropriate action follows different selection criteria. These criteria are chosen to ensure that construction steps occur in the desired regions. A criterion vector is assigned to each action in X and the vector is used as the criterion for selection. Consider Y to be the feasible set of criterion vectors in $\mathbb{R}^2$, such that $Y = \{y \in \mathbb{R}^2 :\ y = [c_1(\mathbf{x}), c_2(\mathbf{x})],\ \mathbf{x} \in X\ \}$ is defined differently for each type of material. We use the Pareto frontier [23,24] to eliminate actions that are strictly *dominated* by some other. An action is dominated if there is another action in the action space that has a better score in one criterion and at least the same score in the others. We compute the Pareto frontier as:

$$A(Y) = \{y' \in Y :\ \{y'' \in Y :\ y'' \succ y', y'' \neq y'\ \} = \emptyset\}. \tag{5}$$

$A(Y)$ is a refined action space where all elements are equivalent according to our criteria. A random action is selected among the dominant actions in $A(Y)$.

We define the MARS gap, $\Delta P_\kappa[h_r]|_{Q_r}$, as the difference between the MARS function and the height function h that describes the structure, i.e.

$$\Delta P_\kappa[h]|_{Q_r} = P_\kappa[h]|_{Q_r} - h. \tag{6}$$

In what follows, we will make use of two additional symbols: the target location $\mathbf{t}^*$ is the final goal position that the robot needs to navigate to and $\tilde{\mathbf{t}}$ represents the closest point to the target location that is within the navigable region of the robot, during the construction process.

4.4.1 Coarse Deposition Behavior The purpose of this behavior is to complete the bulk of the structure. It serves to speed up the building process by allowing placement of larger objects, but is not sufficient to reach navigable conditions, see Fig. 1(a) and Sec. 4.4. The robot selects the region in the structure that needs the most material according to the calculated MARS. We utilize large, rigid, cuboid foam blocks that are of the same height, with varying lengths and widths (refer Section 4.1) to implement this behavior.

The problem of fitting irregular objects on an irregular surface is a challenging problem and an area of research by itself [25,26], and beyond the scope of this paper.

Instead, we can generalize the shape of an irregular object by using the cuboid defined by its minimal bounding box. Thus, while using the regular foam blocks simplifies our implementation, it does not restrict the system's use of irregular objects.

Action Space We assume that we can place a rigid foam block in some part of the structure if we can fit its minimal bounding box in the MARS gap. Since we are not modelling the interaction between the material and the structure, or performing any kind of stability analysis, the placement is prone to errors and disturbances. To mitigate such errors in deposition, we define a position as a legal building action for an object if it is possible to fit the given object in any orientation at that position, avoiding making very tight depositions. The action space is the set of legal building actions that are surrounded by other legal building actions, thereby reducing the occurrence of illegal placements (placements above MARS) caused by the large errors in deposition.

Action Selection For this deposition behavior, it is more advantageous if the construction process begins nearest the target location and extends backward towards the initial navigable region. This process delays the coarse modifications from happening near the robots.

Coarse modifications usually do not leave the structure in a navigable state. When these modifications occur near the area where the robots are, the navigable area will shrink, causing the robots to move away from the target and possibly will no longer observe areas of the structure that were already visible. This shrinkage is normal as the navigable area increases and decreases during the construction process, but it is desirable that the shrinkage occurs as late as possible and for the shortest possible duration. Therefore, we use the distance between the deposition location and the target position as one of the selection criteria ($c_1(\mathbf{x})$). We also want to fill in the areas with the biggest MARS gap first, because we can use the larger building materials, and thus, this is another deposition criterion associated with this behavior ($c_2(\mathbf{x})$). We define c_1, c_2 as:

$$c_1(\mathbf{x}) = -|\mathbf{x} - \mathbf{t}^*|,$$
$$c_2(\mathbf{x}) = \Delta P_\kappa[h]|_{Q_r}.$$

4.4.2 Fine Deposition Behavior The *fine deposition behavior* is used to reach navigable conditions. The robot handles small scale depositions which allow for a higher resolution modification to the environment, and therefore is preferably used when a structure is close to being navigable. This behavior is shown in Fig. 1(b-c), where the small, pink bags represent fine depositions.

Action Space: To implement this behavior we use the same building strategy described in [6], in which a robot moves in the direction of its target destination until it finds an obstacle that hinders its progress, i.e. a pair of points $\mathbf{p}, \mathbf{q}$ s.t $|\mathbf{p} - \mathbf{q}| \leq \delta$ that does not satisfy the navigability condition (1), and deposits material in the point of lower height value $\mathbf{x}$, defined by $\mathbf{x} = \text{argmin}\{h(\mathbf{p}), h(\mathbf{q})\}$. The discontinuity value between those points is given by:

$$g(\mathbf{p}, \mathbf{q}) = |h(\mathbf{p}) - h(\mathbf{q})| - \kappa|\mathbf{p} - \mathbf{q}|, \;\; \mathbf{p}, \mathbf{q} \in Q. \tag{7}$$

Given that ϵ is the maximum discontinuity the robot can drive over, when the navigability condition between the points $\mathbf{p}, \mathbf{q}$ does not hold, $g(\mathbf{p}, \mathbf{q}) > \epsilon$. Therefore, ϵ is the lower bound on the amount of material required for the deposition at $\mathbf{x}$.

The bag deposition is modeled by a cone function [5,6], which provides an upper bound on how much the environment changes in response to the deposition actions, while the bottom conforms to the structure. Given an apex position $(\boldsymbol{\phi}, \sigma) \in Q \times \mathbb{R}^+$ and steepness $\kappa_D \in \mathbb{R}^+$, the cone function is defined as:

$$f_{(\boldsymbol{\phi},\sigma)}(\mathbf{x}) = \sigma - \kappa_D|\boldsymbol{\phi} - \mathbf{x}|. \tag{8}$$

However, physical depositions may not be perfect cones. In fact, as long as the deposition, defined by some arbitrary continuous shape function, is bounded above by a cone with a slope greater than κ and a maximum deposition height lesser than or equal to ϵ, the proofs presented in [5] guarantee that no depositions accidentally make intermediate structures larger than $P_\kappa[h]|_{Q_r}$. In our implementation, $\kappa_D > \kappa$ and the deposition height is equal to ϵ and so the small bags always conform to the cone deposition model. Hence, a deposition of a small bag at $\mathbf{x}$ where $h_m = \epsilon$ is a legal action.

```
while Robot not at target location do
    Wait until robot is marked active
    Compute the Minimal Additive Ramp Structure (MARS) bound given
    sensor information
    Use the Behavior Composition strategy to identify next building behavior
    if Robot can perform next building behavior then
        Pick up and deposit the object according to the chosen behavior
    else
        Signal the appropriate robot to take over
    end
end
```

Algorithm 1: Construction Algorithm

Action Selection: During a fine deposition building step, most of the volume of the structure, or at least the part the robot can observe, has already been filled during prior coarse deposition steps. At this stage, the robots need to repair the irregularities in the structure that prevented them from following their path, i.e. the robots want to repair regions that are closer to becoming navigable, but cannot be repaired with large materials. Among these regions, the ones with major irregularities are most likely to hamper the progress of robots, and so the size of the irregularity is used as one of the selection criteria ($c_1(\mathbf{x})$). It is also more advantageous to expand the navigable area as quickly as possible so that the robots can gain access to new unexplored areas, so the other selection criterion is the distance from the deposition point and the navigable area ($c_2(\mathbf{x})$). Therefore, given the action space $X \in Q$, we define c_1, c_2 as:

$$c_1(\mathbf{x}) = g(\tilde{\mathbf{t}}, \mathbf{x}),$$
$$c_2(\mathbf{x}) = -|\mathbf{x} - \tilde{\mathbf{t}}|.$$

4.4.3 Fortifying Deposition Behavior The purpose of this behavior is to increase compaction in the structure for better structural integrity. We utilize large deformable compliant bags (orange bags depicted in Fig. 1(b) and (c)) to implement this behavior. The dimension of the bags used need not conform to the constraints imposed for the bag placement in the *fine deposition behavior*.

Action Space This behavior promotes looking for intermediate volume irregularities. These regions are usually the unfilled space left between foam blocks and other objects, since we avoid making very tight depositions when using the foam blocks. Any cylindrical region in the MARS gap, with a cross-sectional diameter of the maximum bag length and height larger than ϵ, defines a legal building action for this type of material. Unlike foam blocks, the large bags are heavy and do not rotate or turn during placement and therefore require no containment procedure.

Action Selection For large bags, we want to fill areas with the smallest gaps ($c_1(\mathbf{x})$), which are larger than ϵ. Thus, this behavior competes with neither the coarse nor the fine deposition placement areas and hence can occur in parallel without the need for explicit arrangement. Similar to the coarse deposition behavior, this behavior does not leave the structure in a navigable state and thus we want to prioritize depositions that occur are close to the target $\mathbf{t}^*$ ($c_2(\mathbf{x})$):

$$c_1(\mathbf{x}) = -\Delta P_\kappa[h_r]|_{Q_r},$$
$$c_2(\mathbf{x}) = -|\mathbf{x} - \mathbf{t}^*|.$$

4.5 Construction Algorithm

The construction algorithm, shown in Algorithm 1, uses the functional specification of the desired structure, which is navigability, and a composition of building behaviors (Behaviour Composition Strategy) to coordinate the construction process. The proposed algorithm assumes that there is always only one robot working on the structure each time. Thus, the co-ordination between the robots arises directly from the *Behavior Composition Strategy*, which chooses the type of material to be used and in turn the robot that gets to modify the structure and the imposed structure exclusivity, which makes sure that only one robot is working on the structure. Additionally, each robot is able to compute all the legal assembly moves. After a deposition, the *active* robot recomputes the MARS bound to either perform the next construction step or signal the other robot to take over.

The general flow of the algorithm is that large, rigid, and light foam blocks fill large holes in the structure that are accessible by the robot. Then, the fortifying deposition behavior fills the remaining large gaps in the observable structure, such that large bags are placed on and around the previously deposited, loose-footed foam blocks, thus fortifying them. The fine deposition behavior is finally invoked as long as the structure itself is not finished, in order to make it navigable. Hence, the robots coordinate through the structure state.

5 Results and Discussion

We designed experiments to verify the feasibility of using the chosen materials, to evaluate the efficacy of the system over different terrains and to validate the use of a multi-robot system. Each one of the experiments shown in Table 2 represents a different initial structure; they were evaluated in terms of the total time of autonomous

Table 2. Experimental Results

Experiment	*Foams*	*Bags*	*Total time (min)*	*Climbable*	*Mean Foam Volume (cm^3)*	*Mean Bag Volume (cm^3)*
Exp.1 (1)	14	67	82	no	796.8	188.2
Exp.2 (1)	-	49	58	yes	-	127.4
Exp.2 (2)	-	98	109	yes	-	272.0
Exp.3 (1)	-	85	99	yes	-	144.1
Exp.3 (2)	4	61	77	yes	430.5	185.1
Exp.4 (1)	7	119	145	yes	611.6	137.3
Exp.4 (2)	5	120	148	yes	701.5	119.2
Exp.5 (1)	5	97	118	yes	344.383	116.378

operation, whether the structures were navigable, the number and type of materials deposited, as well as the average occupied volume of each type of deposition material. An experimental run is marked as a success when the robots have built a navigable access structure to reach a previously inaccessible target location. In our test environment, the center of the green platform is the target location. Figure 5 shows the initial and final structures for three of the built structures. Additionally, a video[1] containing a complete demonstration of the construction process is available.

5.1 Structural Integrity

In our first analysis, we want to evaluate how the structural integrity of a structure is related to the type of material from which it is composed. We analyze the ability of a structure to support the robots locomotion when composed of a certain type of material, as well as when composed of a mixture of materials.

First, we observe the structural integrity of a structure when constructed on a flat surface, without any irregularities apart from the step to the target platform:

- *Foam blocks*: Due to its lightweight aspect (see Table 1), structures built only with foam blocks lose their integrity immediately after the robots begin to maneuver over them; the blocks are pushed aside and tend to flip when the robot traverses over them. Also, the foam blocks by themselves cannot be used to attain navigable conditions (Fig. 1).
- *Bean bags*: Structures built with only bags are able to maintain structural integrity. The structure deforms moderately during the initial stages of interaction with the robot, causing the material to spread across the boundaries of the structure. More material is added to compensate for this deformation, which ends up leaving the structure heavier and denser, but strong enough to maintain structural integrity during robot locomotion.
- *Foam blocks* and *bean bags*: The ability to maintain structural integrity depends on the arrangement of the building material types; if the foam blocks are loosely packed as part of the outer layers, the friction between the wheel and the block pushes the block away. When the blocks and bags are interleaved, especially if there is a layer of bags covering the blocks, the blocks tend to stand still as they

[1] https://www.youtube.com/playlist?list=PL-hBy4E-KTmi3ZvcO8DlQZHUekdGAL6_f

are fortified by the bags surrounding them. This experiment corresponds to Table 2 Exp.1 run (1). Due to the flat terrain, a large number of foam blocks were used in the structure. Though the fortifying bags were used to interleave block placements and cover them from the top, the blocks in contact with the smoother terrain underwent shifts over the course of the construction process as the robot traverses over the partially built structure. As a result, the robot failed to reach the target platform as the structure could not maintain its integrity.

Next we analyze the structural integrity when construction happens over irregular terrains. To do so, we augment the environment with rocks and rubble to simulate highly unstructured terrains.

- *Foam blocks*: Similar to structures built over flat terrain, when the robots' wheels come into contact with the foam blocks, they are pushed away. However, the surrounding rocks in the environment tend to serve as retaining material. The foam blocks are more robust to shifts than when placed on a flat terrain. However, the foam blocks by themselves cannot attain navigable conditions without the bags (Fig. 1).
- *Bean bags*: The bean bags are very useful on irregular terrain, especially after the robots have traversed over and compressed them into the substrate, because they can fill gaps where our manipulator might not normally be able to place them otherwise. The bags are also useful to hold small rocks in place which, just like the foam blocks, would move and flip when in contact with the wheels. Because of these characteristics, the use of bags increase the overall structural integrity when compared with a structure with only rocks. This experiment corresponds to Table 2 Exp.2 run (1), Exp.2 run (2) and Exp.3 run (1). In Exp.2 run (2), during the construction process, we manually disrupt the structure during the experiment, by adding, removing and/or displacing deposited bags. The construction system then reacts to such disturbances, further demonstrating its adaptability.
- *Foam blocks* and *bean bags*: As long as the foam blocks are used mostly to fill the internal parts of a structure, and the bags are used to fortify the foam blocks and cover them, the structure does not shift due to the weight of a robot as it maneuvers over it. Additionally, the weight (see Table 1) and contact texture of the bean bags provide better friction than the smoother, lighter foam blocks. This experiment corresponds to Exp.3 run (2), Exp.4 run (1) and run (2), and Exp.5 run (1).

5.2 Deformation Analysis

While we exploit the easy deposition and composition of compliant building materials, an unfortunate side effect is that the resulting structures are also deformable. As a result, the structure requires repair during the construction process to compensate for the compaction that occurs during the locomotion of the robots. In this section, we reason about the amount of unmolded structural deformation that occurs during the construction process.

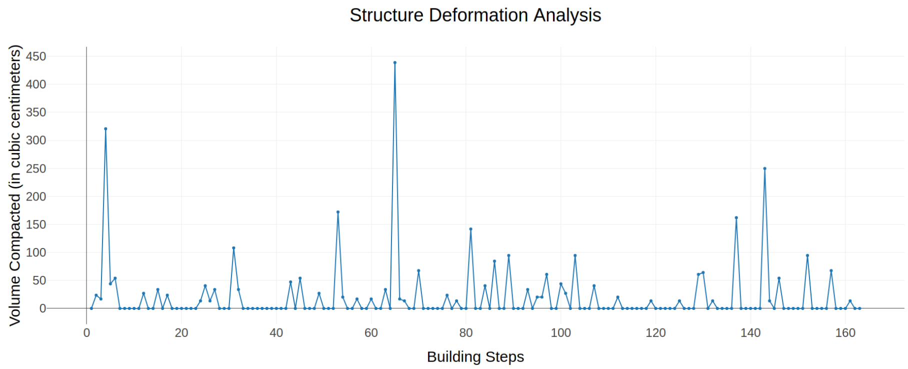

(a) Structure Deformation Plot

(b) Deformation at building step 4.

(c) Deformation at building step 65.

(d) Deformation at building step 143.

Fig. 6. (a) depicts the negative deformation (compaction) in the structure as a function of building steps. In each of (b), (c) and (d): The first two images showcase the structure before and after a deposition at building steps 4, 65 and 143, respectively, and the third image depicts the dense optical flow between the two structure images, in which the various regions of interest are hand labelled, to better visualize the structural changes. The red ellipse depicts negative deformation, blue depicts positive deformation and green depicts the bag deposition.

Figure 6(a) depicts the negative deformation, or compaction, in the structure as a function of building steps for Exp.4, run (2). The compaction is the absolute sum of the negative differences in the structure height functions between every two consecutive building steps. The difference in the number of building steps (164) in Fig. 6(a) and the total number of depositions (125) depicted in Table 2 is due to the fact that the bag robot sometimes fails to grip bags during pickup and hence does not always place a bag successfully during each of its building steps.

The three maximum compactions in Fig. 6(a) occur at building steps at 4, 65 and 143. The maximum compaction occurs at building step 65 and the situation is showcased in Fig. 6(c). The first and second images depict the structure state at building steps 64 and 65, respectively. The third image shows the dense optical flow [27] between the two structure images to better visualize the differences between them, along with hand labelled regions depicting the specific types of changes. The red ellipse depicts the region of the structure the robot moved over and hence, it depicts the negative deformations. The green circle depicts the bag deposition. Fig. 6(b) depicts the situation at building steps 3 and 4. The bag deposition in building step 4 (shown in green) displaces the loosely footed foam block such that a portion of it is pushed down (shown in red) while the remaining is pushed up (shown in blue). It also shifted a stone (shown outside the colored ellipses) during its motion. Fig. 6(d) depicts the situation at building steps 142 and 143. The bag deposition is depicted by the green ellipse. The bags that are placed near the foot of the ramp structure are loosely held together and shift considerably when the rover moves over them. This is captured in the plot since the shifted bags cause negative deformations in the areas they had moved away from.

6 Conclusion

We introduced a concise mathematical model of construction that allows robots to solve a specific piece of a much larger problem of operating in unstructured environments, which is *mobility*. The abstract model of construction uses few parameters that can be computed for most robots, and also allows the use of irregular building materials of different physical properties. We focused on construction for mobility because this is a basic function for a robot to operate autonomously. By implementing our methods on a physical system, we could, examine distributed, multi-material construction with rigid and deformable objects, demonstrating a system of physical robots that coordinate their efforts.

The coordination algorithm is organized around an abstract, geometric model of construction, which is used to show the correctness of the deposition strategy. The main innovation is that by directly computing legal assembly moves as a function of currently available data, agents are allowed to use larger, more efficient materials. Our experiments show that using multiple materials can speed up the process both in terms of the number of depositions and the overall time. This is effective over irregular terrain when the rigid blocks are constrained by preexisting irregular features.

In the system presented, robots with different material handling capabilities cooperate to construct navigable structures on unstructured terrain, exploiting the qualities of a mixture of rigid and deformable building elements. The system can operate completely autonomously for many hours, locating, choosing and depositing different building materials to modify its environment.

We analyze the structural deformation as a function of time to show that environment dependent post-assembly dynamics, such as compaction, are present and can be compensated for by the reactive nature of our approach. We believe this to be a key step in designing and reasoning about construction algorithms that can fully utilize a variety of construction materials.

In future work, we aim to improve the packing efficiency of the foam blocks by utilizing better packing algorithms. This would not only increase the number of foam blocks in the structure and reduce construction time, but would also add more stability to the structure. More immediately, we will focus on a self-contained robotic system with on-board sensors to facilitate SLAM, which bear several practical challenges in unconstrained and unstructured environments. Our long term goal is to extend our robot system to use found materials to build similar access structures and thus move one step closer to construction robots that can operate autonomously in remote locations over long periods of time, for example at extra-terrestrial sites or disaster areas.

Acknowledgements

We would like to thank Prof. Petersen for her helpful feedback, and Walker Gosrich for helping with the development of the robots. We are also grateful for the Science Without Borders Program (CAPES #013584/2013-08), NSF Grant #1846340 and the Computer Science and Engineering Department at the University at Buffalo for supporting this work.

References

1. K. H. Petersen, N. Napp, R. Stuart-Smith, D. Rus, and M. Kovac, "A review of collective robotic construction," *Science Robotics*, vol. 4, no. 28, 2019.
2. M. H. Hansell, *Animal architecture*. Oxford University Press on Demand, 2005.
3. J. S. Turner, *The Extended Organism: The Physiology of Animal-Built Structures*. Harvard University Press, 2002.
4. L. Wilsson, *Observations and experiments on the ethology of the European beaver (Castor fiber L.) : a study in the development of phylogenetically adapted behaviour in a highly specialized mammal*. Uppsala: Almqvist och Wiksell, 1971.
5. N. Napp and R. Nagpal, "Distributed amorphous ramp construction in unstructured environments," *Springer Tracts in Advanced Robotics*, vol. 104, pp. 105–119, 2014.
6. M. Saboia, V. Thangavelu, W. Gosrich, and N. Napp, "Autonomous adaptive modification of unstructured environments," *Proc. Robotics: Science & Systems XIV*, 2018.
7. G. Theraulaz and E. Bonabeau, "Coordination in distributed building," *Science*, vol. 269, no. 5224, p. 686, 1995.
8. R. Beckers, O. E. Holland, and J.-L. Deneubourg, "From local actions to global tasks: Stigmergy and collective robotics," 1994.
9. G. Theraulaz and E. Bonabeau, "A brief history of stigmergy," *Artificial life*, vol. 5, no. 2, pp. 97–116, 1999.
10. Z. Mason, "Programming with stigmergy: Using swarms for construction," in *Proceedings of the Eighth International Conference on Artificial Life*, ser. ICAL 2003. Cambridge, MA, USA: MIT Press, 2003, pp. 371–374.
11. J. Werfel, K. Petersen, and R. Nagpal, "Designing collective behavior in a termite-inspired robot construction team," *Science*, vol. 343, no. 6172, pp. 754–758, 2014.

12. K. Petersen, R. Nagpal, and J. Werfel, "Termes: An autonomous robotic system for three-dimensional collective construction," *Proc. Robotics: Science & Systems VII*, 2011.
13. M. Allwright, N. Bhalla, A. Antoun, H El-Faham, C. Pinciroli, and M. Dorigo, "SRoCS: Leveraging Stigmergy on a Multi-robot Construction Platform for Unknown Environments," *Swarm Intelligence – Proceedings of ANTS 2014 – Ninth International Conference*, vol. 8667, pp. 157–168, 2014.
14. N. Melenbrink, P. Michalatos, P. Kassabian, and J. Werfel, "Using local force measurements to guide construction by distributed climbing robots," *In Proceedings of IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 2017.
15. N. S. Estévez and H. Lipson, "Dynamical blueprints: exploiting levels of system-environment interaction," in *Proceedings of the 9th annual conference on Genetic and evolutionary computation*. ACM, 2007, pp. 238–244.
16. M. Allwright, W. Zhu, and M. Dorigo, "An open-source multi-robot construction system," *HardwareX*, vol. 5, p. e00050, 2019.
17. T. Soleymani, V. Trianni, M. Bonani, F. Mondada, and M. Dorigo, "Bio-inspired construction with mobile robots and compliant pockets," *Robotics and Autonomous Systems*, vol. 74, pp. 340–350, 2015.
18. D. Andreen, P. Jenning, N. Napp, and K. Petersen, "Emergent structures assembled by large swarms of simple robots," 2016.
19. R. Fujisawa, N. Nagaya, S. Okazaki, R. Sato, Y. Ikemoto, and S. Dobata, "Active modification of the environment by a robot with construction abilities," *ROBOMECH Journal*, vol. 2, no. 1, p. 9, Apr 2015.
20. T. Tosun, J. Daudelin, G. Jing, H. Kress-Gazit, M. Campbell, and M. Yim, "Perception-informed autonomous environment augmentation with modular robots," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6818–6824.
21. B. A. Davey and H. A. Priestley, *Introduction to lattices and order*. Cambridge university press, 2002.
22. E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 3400–3407.
23. V. Pareto, "Manuale di economica politica, societa editrice libraria," *Manual of political economy*, vol. 1971, 1906.
24. R. T. Marler and J. S. Arora, "Survey of multi-objective optimization methods for engineering," *Structural and multidisciplinary optimization*, vol. 26, no. 6, pp. 369–395, 2004.
25. F. Furrer, M. Wermelinger, H. Yoshida, F. Gramazio, M. Kohler, R. Siegwart, and M. Hutter, "Autonomous robotic stone stacking with online next best object target pose planning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2350–2356.
26. V. Thangavelu, Y. Liu, M. Saboia Da Silva, and N. Napp, "Dry stacking for automated construction with irregular objects," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018.
27. G. Farnebäck, "Two-frame motion estimation based on polynomial expansion," in *Scandinavian conference on Image analysis*. Springer, 2003, pp. 363–370.