Autonomous Collaborative

Robotic Reconfiguration with Deep Multi-Agent Reinforcement Learning [ACRR+DMARL]

Tyson Hosmer*
Sergio Mutis*
Eric Hughes*
The Bartlett, UCL
Ziming He
Tongji University
Philipp Siedler
Octavian Gheorghiu
Barış Erdinçer
The Bartlett, UCL

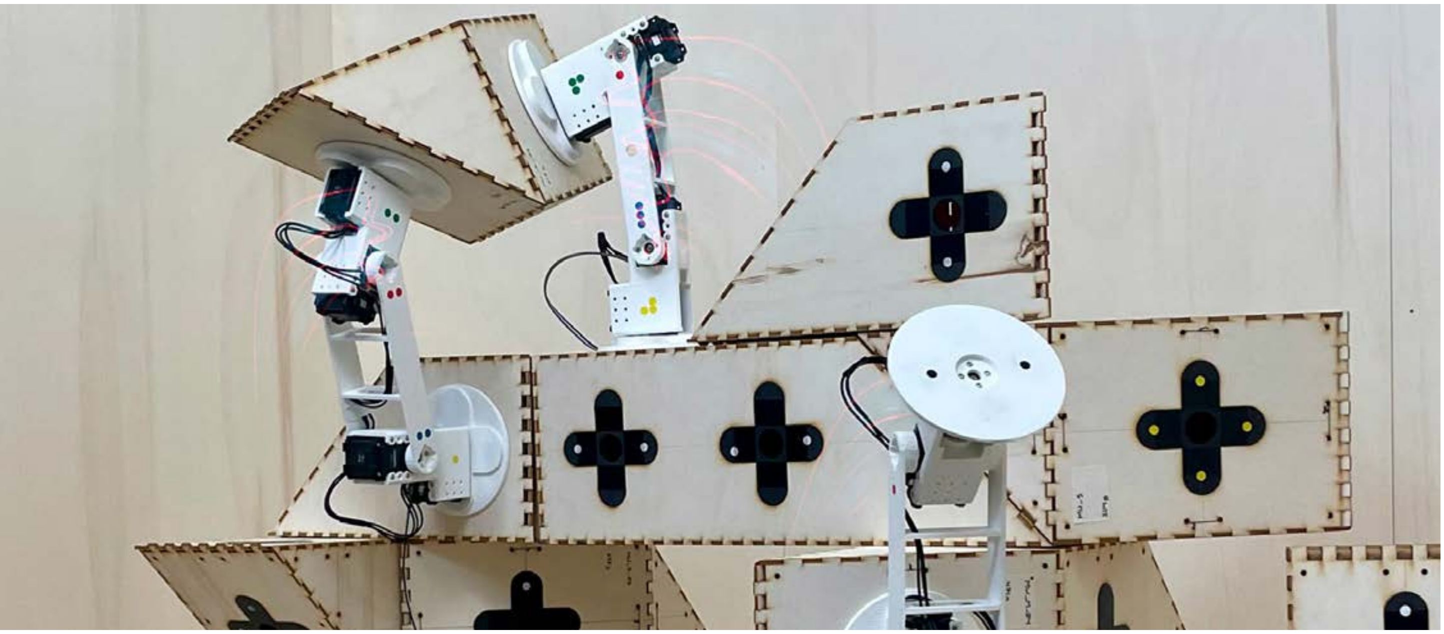

1 Photograph: Collaborative Robots.

ABSTRACT

To address the unprecedented challenges of the global climate and housing crises, requires a radical change in the way we conceive, plan, and construct buildings, from static continuous objects to adaptive eco-systems of reconfigurable parts. Living systems in nature demonstrate extraordinary scalable efficiencies in adaptive construction with simple flexible parts made from sustainable materials. The interdisciplinary field of collective robotic construction (CRC) inspired by natural builders has begun to demonstrate potential for scalable, adaptive, resilient, and low-cost solutions for building construction with simple robots. Yet, to explore the opportunities inspired by natural systems, CRC systems must be developed utilizing artificial intelligence for collaborative and adaptive construction, which has yet to be explored. Autonomous Collaborative Robotic Reconfiguration (ACRR) is a robotic material system with an adaptive lifecycle trained with deep, multi-agent reinforcement learning (DMARL) for collaborative reconfiguration. Autonomous Collaborative Robotic Reconfiguration is implemented through three interrelated components codesigned in relation to each other: 1) a reconfigurable robotic material system; 2) a cyber-physical simulation, sensing, and control system; and 3) a framework for collaborative robotic intelligence with DMARL. The integration of the CRC system with bidirectional cyber-physical control and collaborative intelligence enables ACRR to operate as a scalable and adaptive architectural eco-system. It has the potential not only to transform how we design and build architecture, but to fundamentally change our relationship to the built environment moving from automated toward autonomous construction.

INTRODUCTION

Our global population is estimated to increase to 11.2 billion by the year 2100, requiring us to build 2 billion new homes over the next 80 years. The construction industry creates an estimated 33% of the world's waste, and at least 40% of the world's carbon dioxide emissions (Miller 2021). The construction industry remains one of the least digitized and slowest to adopt disruptive technologies (Agarwal et al. 2016, Loosemore 2015). We continue constructing buildings organized in sheering layers and designed with linear building life cycles eventually ending in demolition (GlobalData n.d. 2018, Ngwepe and Aigbavboa 2015, Brand 1995). To address the unprecedented challenges of the global climate and housing crises requires radically changing the way we conceive, plan, and construct buildings, from static continuous objects to adaptative eco-systems of reconfigurable parts.

Living systems in nature demonstrate extraordinary scalable efficiencies in adaptive construction with simple, flexible parts made from sustainable materials. For example, nomadic ant colonies face extreme pressure to generate foraging routes, moving massive numbers of ants each day, yet through their simple parts and local rules they have shown rapid efficiency in constructing adaptive "living bridges" through the linkage of their bodies (Figure 4). Ants modulate their behavior in response to locally changing environments to adapt to dynamic traffic conditions, recover from damage, and dissemble when underused (Graham et al. 2017).

Inspired by robust natural construction systems, new approaches to construction with teams of robots have become active areas of interdisciplinary research highlighting opportunities for safe, sustainable, and efficient building construction. Collective robotic construction specifically concerns embodied, autonomous, multirobot systems that modify a shared environment according to high-level, user-specified goals integrating architectural design, the construction process, mechanisms, and control (Peterson et al. 2019). These systems typically involve machines that are codesigned with the architectural systems they construct, enabling them to be more adaptive, scalable, and reusable while operating in dynamic environments (Leder et al. 2022, Silver 2017, Lindsey et al. 2011, Kayser et al. 2018, Jenett et al. 2019, Terada and Murata 2008, Napp et al. 2012).

To demonstrate the opportunities inspired by natural systems, CRC systems must be developed with artificial intelligence for collaborative and adaptive construction, which has yet to be explored. Autonomous collaborative

2

3

2 Photograph: Autonomous Collective Robotic Reconfiguration (ACRR), 2.4 m arch configuration.

3 Photograph: Autonomous Collective Robotic Reconfiguration (ACRR), window configuration.

4 Photograph: Ant Bridge (Lutz 2015). 4

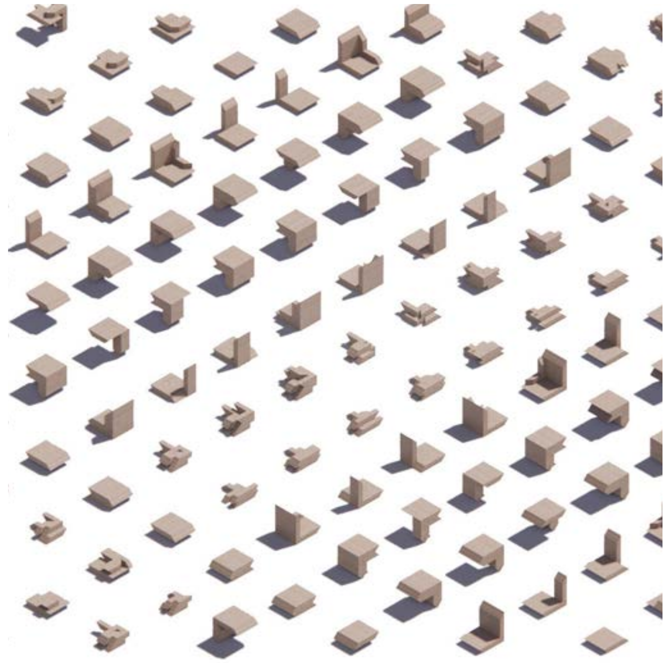

5 Figure: Robotic Construction Configurations. 5

robotic reconfiguration is a robotic material system with an adaptive lifecycle trained with DMARL for collaborative reconfiguration (Figure 1 through 3). Autonomous collaborative robotic reconfiguration is implemented through three interrelated components codesigned in relation to each other: 1) a reconfigurable robotic material system; 2) a cyber-physical simulation, sensing, and control system; and 3) a framework for collaborative robotic intelligence with DMARL. The integration of the CRC system with bidirectional cyber-physical control and collaborative intelligence enables us to project operating as a scalable and adaptive architectural eco-system.

STATE OF THE ART

Automation and Autonomy

Automation comes from the Greek word "automaton", meaning "self-movement," whereas autonomy comes from autonomos meaning "self-law." Autonomous construction in nature is far more rapid, adaptive, and efficient than our built environment. Despite radical advancements in artificial intelligence (AI) and robotics, automation in architecture tends to focus on making incremental improvements to conventional unidirectional processes, while robotic buildings should consider not only design to production, but design-to-production-to-operation chains from a lifecycle perspective relating to the socio-economic and ecological impacts (Bier and Mostafavi 2018; Bier et al. 2018). To move from procedural automation toward autonomy requires developing architecture with the properties of facilitated variation, situated and embodied agency, as well as intelligence (Hosmer and Tigas 2019).

Tibbitts defines self-assembly as a process by which disordered parts build an ordered structure without humans or machines (Tibbits 2017, 2012), while Gershenfeld develops principles for self-assembly around the concept of "digital materials", enabling reversibility and reconfigurability through computational models structuring the combinatorics of discrete parts (Popescu, Mahale and Gershenfeld 2006, Retsin 2019, Retsin and Garcia 2016). Furthermore, we consider the theory of "facilitated variation" in biology suggesting that the intrinsic construction of an organism directly affects its "evolvability (Parter, Kashtan, and Alon 2008, Gerhart and Kirschner 2007, Kirschner 2009)." We extend principles of digital materials and facilitated variation in the design of reconfigurable architecture by embedding effective degrees of freedom and constraint in a cyber-physical simulation model which is assembly aware (Figure 6).

Construction in Nature

Natural builders, such as social insects, exhibit extraordinary levels of efficiency, scalability, adaptability, and robustness in developing complex habitats through forms of collective intelligence and collaborative construction for building nests and living quarters, protection barriers, traps, and mobility scaffolds. This often occurs through the interactions of individuals with little or no global knowledge (Peterson and Nagpal 2017, Peterson et al. 2019, Hansell 2007).

Termites build complex living environments communicating via pheromone deposition with no centralized control,

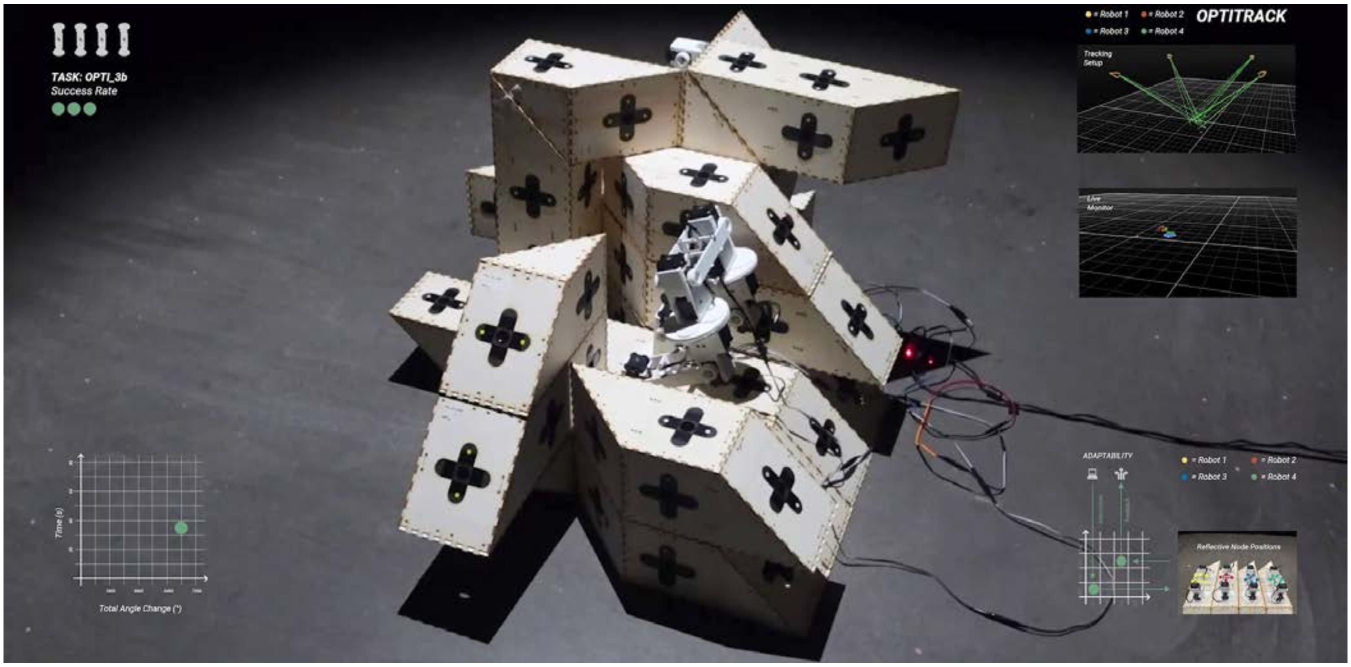

6 Image: Motion Capture Cyber-Physical OptiTrack Sensor Feedback Setup.

6

building stigmergy where individuals demonstrate specific building behaviors in response to existing structures with "stimulating configurations" (Hansell 2007). Termites construct intricate networks of channels for a variety of interrelated functions utilizing soil, saliva, and other organic materials reaching depths of more than 10 meters. Collective construction in nature is adaptive to changing conditions and resilient to localized damage, failure, or the loss of workers (Figure 4). Army ants build "living bridges" by linking their own bodies together to dynamically create adaptive physical structures to improve the efficiency of their foraging paths (Graham et al. 2017, Anderson et al. 2002; Garnier et al. 2013; Reid et al., 2015; Rettenmeyer, 1963; Schneirla, 1972). Researchers have found that efficiency in ant bridges is quantifiable, with costs and benefits not measured in single elements or performances, but revealed in degrees and rates of adaptation to changing environments and shifting goals (Graham et al. 2017).

It is common for natural builders to dynamically create the environment they use to navigate. This strategy can be extended to robotic material systems where active robotic agents and passive material components are codesigned to maximize scalability and adaptability.

Collective Robotic Construction

Industrial implementation of robotics in construction has focused primarily on the automation of conventional prefabrication processes in controlled environments. Recently, models of multi-robot collaboration in advanced research have begun to transition from automated prefabrication tasks (Wagner et al. 2020, Lloret-Fritschi et al. 2018, Chai, Zhang and Yuan 2020) to initial applications of robotics in onsite construction (Petersen, Nagpal, and Werfel 2011; Melenbrink, Werfel, and Menges 2020; Augugliaro et al. 2013; Dakhli et al. 2017; Mechtcherine et al. 2019). Full scale, collaborative, robotic assembly of timber structures has been demonstrated using industrial robotic arms hung in a mobile gantry system (Adel et al. 2018). Complex prefabrication of components with robotic fiber winding has been implemented for construction (Dambrosio et al. 2019). While these systems exhibit high precision automation in controlled prefabrication, they are limited by the scale and degrees of freedom of the robotic arms and gantry systems.

Collective robotic construction systems offer opportunities for scalable adaptive on-site construction, leveraging a range of centralized or decentralized coordination strategies across various codesigned robotic platforms and material systems (Peterson et al. 2019). Construction coordination has been explored through centralized control (Augugliaro et al. 2013), local communication (Jokic et al. 2017), templated control (Saboia et al. 2018), and emergent coordination (Andreen et al. 2016). Principals have been extracted from biological construction and translated into the field of robotics through various algorithmic strategies such as stigmergy (Theraulaz et al. 1995, Napp et al. 2014, Stewart et al. 2006, Soleymani et al. 2015, Allwright et al. 2014, Grushin et al. 2006, Werfel et al. 2007, Martinoli et al. 1999), swarm flight (Zhang et al. 2022, Stuart-Smith

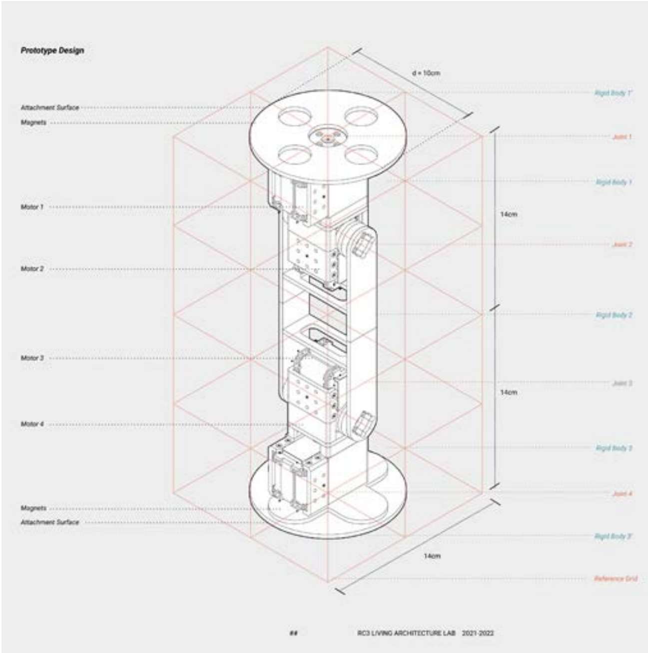

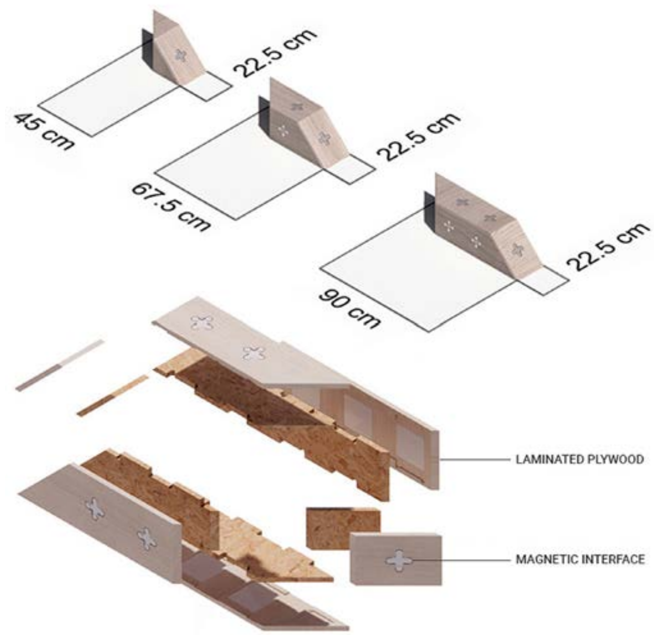

7 Diagram: ACRR Distributed Robot Design.

8 Figure: ACRR Material Unit System.

2016), templating (Stewart and Russell 2006, Soleymani et al. 2015), blind bulldozing (Parker et al. 2003, Parker et al. 2006), reactive and interactive construction (Napp et al. 2014, Estévez and Lipson 2007, Werfel et al. 2007, Veenstra et al. 2015), task allocation (Yun et al. 2011, Meng and Gan 2008, Guo et al. 2009), and specialization (Nitschke et al. 2012). Building element strategies for CRC include predefined elements (Werfel and Nagpal 2006), amorphous materials (Rosman et al. 2018), and continuous elements (Braithwaite et al. 2018), which then translate into a range of mechanisms and material strategies such as robot/brick codesign (Hosmer et al. 2022, Jenett and Cheung 2017, Petersen et al. 2011, Moses et al. 2014, Rosman et al. 2018, Terada and Murata 2004), strut climber codesign (Detweiler et al. 2006, Melenbrink et al. 2017, Melenbrink and Werfel 2019), compliant materials (Stewart and Russell 2006, Soleymani et al. 2015, Napp et al. 2012), amorphous depositions (Napp et al. 2012), and fibers (Felbrich et al. 2017, van de Kamp et al. 2015, Tucker et al. 2022).

Embodied swarm intelligence has been demonstrated through robot reconfiguration (Rubenstein et al. 2012; Petersen and Nagpal 2017) and for small scale CRC with robotic builders navigating over the simple blocks they stack (Petersen and Nagpal 2017, Petersen et al. 2011). The term "relative robot" relates to a codependency between a robotic system and material system for locomotion and assembly through its structured environment (Jenett and Cheung 2017). Examples include the BILL-E robotic robots climbing on the modular lattice structure they assemble, autonomous strut-climbing robots that climb over the trusses they construct, and simple distributed robotic joints leveraging passive timber elements for kinematic chaining and locomotion (Jenett and Cheung 2017, Melenbrink et al. 2017, Melenbrink and Werfel 2019, Leder et al. 2022, Leder et al. 2019). These examples are adaptive and efficient, but face limitations in the scale, materials, and geometries they assemble. Alternatively, aerial additive manufacturing (Aerial-AM) employs high degrees of freedom through teams of aerial robots with coordinated 3D printing, but the resulting structures are not reversible or adaptable (Zhang et al. 2022).

Learning from the above examples and inspiration from natural builders, ACRR is developed as an ecology of active modular robots and passive modular parts with a range of sizes and geometries (Figures 2 and 3). Robots climb on the structures they assemble or navigate on their own (off grid) while learning independent and collaborative behaviors leveraging combinations of multiple robots and differentiated passive parts configuring complex assemblies (Figure 5).

AI for Construction Robotics

Thus far, machine learning (ML) for architectural robotics has focused primarily on solving isolated problems related to fabrication, assembly, or construction engineering rather than holistic CRC strategies. Machine learning has been used for adaptive robotic carving (Brugnaro 2019), automation of scaffold lifting machines (Harichandran 2019), learning robotic behaviors for material manipulation

9 Photograph: Cyber-Physical Simulation and Control System. 9

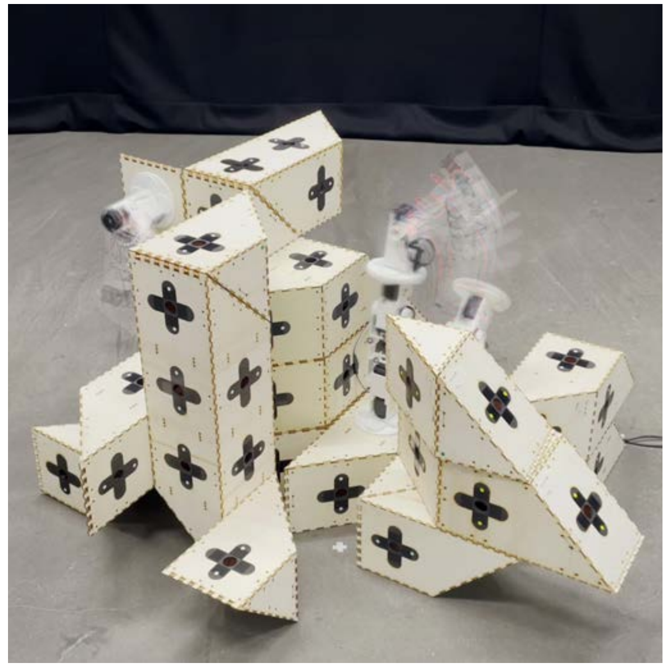

10 Photograph: Motion Capture Robotic Reconfiguration. 10

(Zeng et al. 2019), robotic wire arc additive manufacturing (Dharmawan 2020), and large-scale 3D printing with tower cranes (Parisi et al. 2023). Deep Reinforcement Learning (DRL) has been used with robotic arms in controlled environments for high-precision assembly tasks (Apolinarska et al. 2021, Luo et al. 2019, Luo et al. 2020, Fan et al. 2019, Inoue 2017, Belousov 2022), for solving insertion tasks using Computer-Aided Design (CAD) models (Thomas 2018), and for autonomous block stacking with visual sensing (Felbrich et al. 2022, Zhang et al. 2019). Only a few examples exist for DRL applied to CRC, including a distributed robotic system with a single agent trained to leverage active bending in robotic construction behaviors for the assembly of bamboo structures, and for an autonomous robotic tensegrity system (Łochnicki et al. 2021, Hosmer and Tigas 2019).

Developing autonomous construction through CRC requires the ability to learn intelligent behaviors that are adaptive to complex dynamic environments such as construction sites. Deep reinforcement learning is closely related to the field of optimal control, where one is seeking an optimal policy for controlling a system to optimize objectives (Sutton and Barto 1998). Deep reinforcement learning learns how to control an agent by interacting with the environment through trial and error, making it an optimal strategy for adaptive robotic construction. In this research, we develop a DMARL strategy for collaborative navigation and reconfigurable assembly to demonstrate the potential for adaptive CRC.

METHODS

Autonomous Collaborative Robotic Reconfiguration is implemented via three closely interrelated components: 1) a reversible robotic material system; 2) a cyber-physical simulation, sensing, and control system; and 3) a framework for multi-agent robotic intelligence with deep reinforcement learning. They operate in a cyber-physical feedback loop allowing ACRR to adjust and account for the expected, yet undeterminable, occurrences of real physical environments with intelligent adaptive behaviors.

Reversible Robotic Material System

Inspired by social insects, ACRR's physical system consists of active parts (bespoke distributed robots) and passive parts (biased material units) codesigned to self-assemble and reassemble into complex architectural structures much larger than themselves.

Active parts are simple, modular robots capable of individual locomotion through free crawling or by climbing on the assemblies they manipulate. Each robot, measuring 280x100x100mm and weighing 420g, is composed of three rigid bodies of 3D printed polylactic acid (PLA) and two custom magnetic interfaces. These five elements are articulated by four joints with one degree of angular freedom, actuated by one Dynamixel AX-12A motor each, powered via a U2D2 Power Hub, and controlled with a Raspberry Pi Zero (Figure 7). The simplicity of the robots allows for their easy fabrication and assembly.

Passive parts are building blocks for architectural spaces

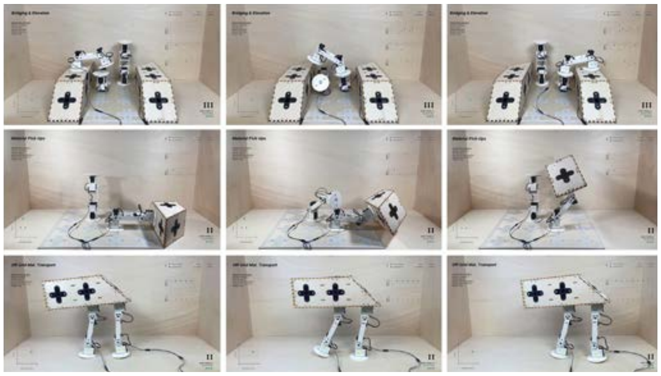

11

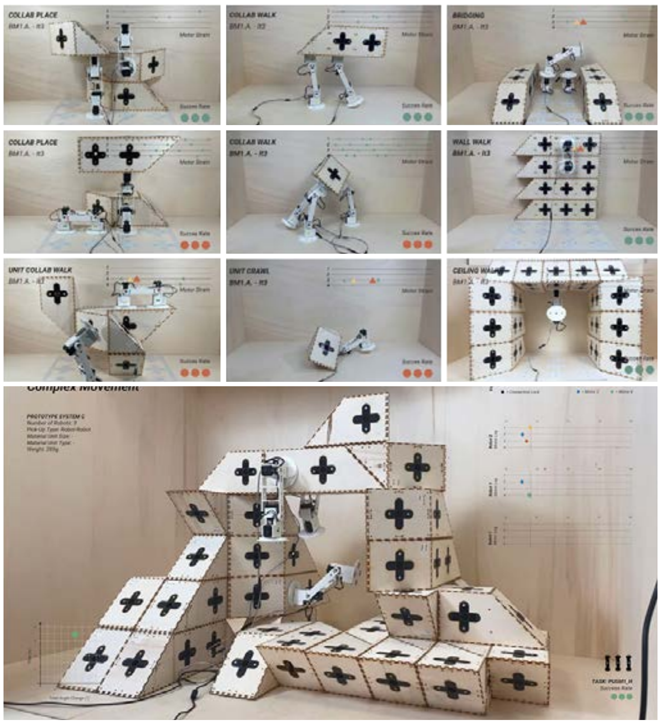

12

RL Setup for Game 1: Simple Locomotion
main goal: translate a robot from a starting position into a goal position.

step threshold (s_{max}) = 500 steps

observations (o)

$o_{1\text{-}4}$ = {motor1_angle, motor_2_angle, motor_3_angle, motor_4_angle}
$o_{5\text{-}7}$ = {robot_x_coordinate, robot_y_coordinate, robot_z_coordinate}
$o_{8\text{-}10}$ = {goal_x_coordinate, goal_y_coordinate, goal_z_coordinate}

actions (a)

$a_{1\text{-}4}$ = {set_motor1_angle, set_motor_2_angle, set_motor_3_angle, set_motor_4_angle}

reward (r)

r = (-|robot_x_coordinate - goal_x_coordinate| + -|robot_y_coordinate - goal_y_coordinate| + -|robot_z_coordinate - goal_z_coordinate|) – t – Σ|Δa|/4

RL Setup for Game 2: Pathfinding
main goal: translate a robot from a starting position into a goal position, avoiding obstacles.

step threshold (s_{max}) = 200 steps

observations (o)

$o_{1\text{-}4}$ = {motor1_angle, motor_2_angle, motor_3_angle, motor_4_angle}
$o_{5\text{-}7}$ = {robot_x_coordinate, robot_y_coordinate, robot_z_coordinate}
$o_{8\text{-}10}$ = {goal_x_coordinate, goal_y_coordinate, goal_z_coordinate}
o_n = {obstacle$_n$_x_coordinate, obstacle$_n$_y_coordinate, obstacle$_n$_z_coordinate}

actions (a)

$a_{1\text{-}3}$ = {advance_x+-1, advance_y+-1, advance_z+-1}
// $a_{1\text{-}3}$ represent best performing behaviors (motor sequences) from previous games //

reward (r)

r = (-|robot_x_coordinate - goal_x_coordinate| + -|robot_y_coordinate - goal_y_coordinate| + -|robot_z_coordinate - goal_z_coordinate|) – t – Σ|Δa|/4

13

11 Collaborative Behaviours: a.) Locomotion: Bridging, b.) Material Pick-Up: Counterweight, c.) Material Translation: Off-grid Bipedal.

12 Photographs: Reversible Robotic Material System Behaviors.

13 Diagram: Deep Reinforcement Learning Setup Game #1 and # 2.

with reversible connections designed in direct relationship with their robotic counterparts. A kit of discrete modular units was developed with different biased geometries, sizes, and materials (Figure 8), serving various architectural roles, and system-specific functions, such as robotic charging stations. The primary timber unit, which comes in three different sizes: 280x140x140mm, 420x140x140mm, and 560x140x140mm, is composed of six CNC-milled laminated plywood surfaces with embedded custom magnetic interfaces and internal laminated plywood reinforcements.

All passive and active parts share a common magnetic interface with four neodymium magnets placed, with alternating polarities, equidistantly along the circumference of a circle with a 4-cm radius. If two interfaces are in close proximity with aligned equivalent polarities (position 0), they repel, and, if aligned with opposite polarities (position 1), they connect. Given that position 0 and position 1 differ by a rotation of 90°, the parts of the system can engage and disengage by rotating their interfaces 90°. Interfaces on passive parts are designed to periodically align along a virtual three-dimensional grid of 140x140x140mm voxels creating a structured environment for the robots to inhabit. Passive parts have between four and twelve of these biased interfaces enabling a plethora of diverse aggregations to be achieved. Simple dynamic interactions between parts enable an array of emergent collaborative behaviors, including robotic translations, material pick-ups, material transportation, temporary support structures, and temporary assemblies of robots and material units into joint body plans (Figure 12).

Cyber-Physical Simulation and Control System

The control system consists of four main elements: a control computer running ACRR's simulation environment, one Raspberry Pi Zero computer per robot, one Dynamixel U2D2 per robot, and four Dynamixel AX-12A motors per robot. The simulator integrates with the bi-directional Dynamixel SDK control protocol to enable the effective communication between Unity, Raspberry Pi, and the Dynamixel hardware (Figure 9). Motor speed and position instructions are calculated within the simulator and sequentially transmitted via WiFi from the control computer to each RaspberryPi, which dispatches the instructions through a USB connection to the U2D2s, converting the data and feeding it into the motors via direct 3Pin TTL and 4Pin RS-485 connectors. In turn, the Dynamixel AX-12A motors read their current positions, speeds, and loads, feeding the data back through the same hardware-software pipeline into the simulation.

ACRR's custom simulation environment is developed in

RL Setup for Game 3: Single-Agent Assembly
main goal: translate a robot from a starting position into a goal position, avoiding obstacles.

step threshold (s_{max}) = 1000 steps

observations (o)
$o_{1\text{-}4}$ = {motor1_angle, motor_2_angle, motor_3_angle, motor_4_angle}
$o_{5\text{-}7}$ = {robot_x_coordinate, robot_y_coordinate, robot_z_coordinate}
$o_{8\text{-}10}$ = {tmu_x_coordinate, tmu_y_coordinate, tmu_z_coordinate}
$o_{11\text{-}13}$ = {goal_tmu_x_coordinate, goal_tmu_y_coordinate, goal_tmu_z_coordinate}
o_n = {mu$_n$_x_coordinate, mu$_n$_y_coordinate, mu$_n$_z_coordinate}

actions (a)
$a_{1\text{-}4}$ = {set_motor1_angle, set_motor_2_angle, set_motor_3_angle, set_motor_4_angle}
a_5 = {navigate_to_location_xyz}
// a_5 represents best performing behaviors (motor sequences) from previous games //

reward (r)
r = (-|tmu_x_coordinate - goal_tmu_x_coordinate | + -|tmu_y_coordinate - goal_tmu_y_coordinate | + -|tmu_z_coordinate - goal_tmu_z_coordinate|) – t – Σ|Δa|/4

RL Setup for Game 4: Collaborative Reconfiguration
main goal: relocate a target material unit (tmu) from a starting into a goal position within an environment of multiple material units (mu) and in collaboration with other robotic agents (ra).

step threshold (s_{max}) = 500 steps

observations (o)
$o_{1\text{-}4}$ = {motor1_angle, motor_2_angle, motor_3_angle, motor_4_angle}
$o_{5\text{-}7}$ = {robot_x_coordinate, robot_y_coordinate, robot_z_coordinate}
$o_{8\text{-}10}$ = {tmu_x_coordinate, tmu_y_coordinate, tmu_z_coordinate}
$o_{11\text{-}13}$ = {goal_tmu_x_coordinate, goal_tmu_y_coordinate, goal_tmu_z_coordinate}
o_n = {mu$_n$_x_coordinate, mu$_n$_y_coordinate, mu$_n$_z_coordinate}
o_i = {ra$_i$_x_coordinate, ra$_i$_y_coordinate, ra$_i$_z_coordinate}

actions (a)
$a_{1\text{-}4}$ = {set_motor1_angle, set_motor_2_angle, set_motor_3_angle, set_motor_4_angle}
a_5 = {navigate_to_location_xyz}
a_6 = {pick_up/place_material}
// $a_{5\text{-}6}$ represents best performing behaviors (motor sequences) from previous games //

reward (r)
r = (-|tmu_x_coordinate - goal_tmu_x_coordinate | + -|tmu_y_coordinate - goal_tmu_y_coordinate | + -|tmu_z_coordinate - goal_tmu_z_coordinate|) – t – Σ|Δa|/4

14 Diagram: Deep Reinforcement Learning Setup Game #3 and #4. 14

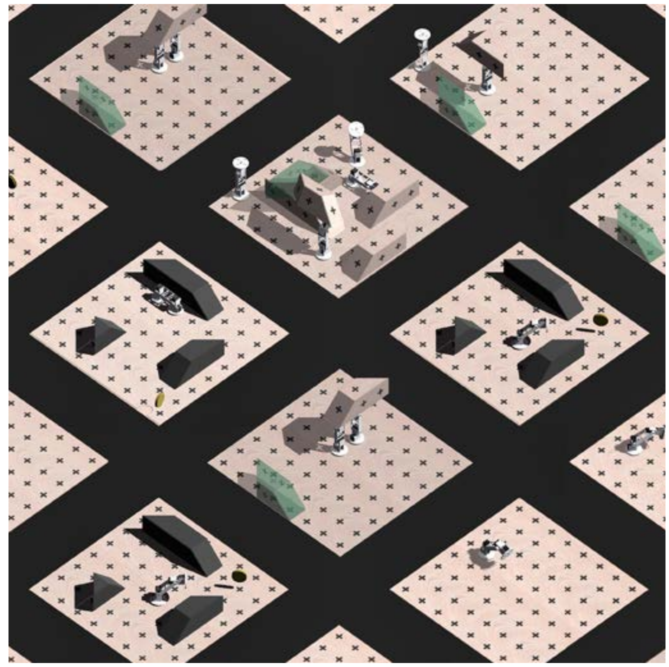

15 Simulation: Multi-Agent Deep Reinforcement Learning Training. 15

Unity3d with an agent-based framework that simulates interactions of robotic agents and passive parts. The simulator is directly linked to the robotic control system, sending instructions, and receiving sensor feedback with a wireless, bi-directional, communication protocol for simultaneous state alignment between the simulation environment and the physical world (Figure 9). Agents trigger encoded actions and action sequences which send goal angle changes to the physical motors and then wait for the motor sensors to send position and load data back indicating success, failure, or critical loads (Figure 10). If there is an obstacle in the real world, and a motor is blocked by it, the system recognizes the location of the obstacle, updates the simulation, and reacts.

For high precision and resilience, the control system was paired with Optitrack motion capture adding three elements into the pipeline: (1) a Motive motion capture session running in the same computer as the Unity 3D simulation; (2) at least four Optitrack cameras; and (3) reflective markers attached to the parts of the robotic material system. As the active and passive parts move in the physical space, the Optitrack cameras register a 2D image of the markers and transmit it to the control computer. This is where Motive triangulates their location to obtain and transmit their 3D positioning to the Simulator via the Optitrack Unity 3D Plugin and the NatNet SDK protocol (Figure 6).

Collaborative Intelligence with Multi-Agent Reinforcement Learning

The autonomous collaborative robotic reconfiguration's Simulation, Sensing and Control System allows for the execution of user-generated and semi-automated action sequences for agents to execute simple locomotion and assembly patterns with the reversible robotic material system. Larger, more complex, scenarios with biased parts, multi-robot collaboration, and reversibility have exponentially larger solution spaces, making real-time human control or pre-coded sequences of individual motors insufficient and unmanageable. To tackle this, our framework leverages DRL to train situated and embodied agents to learn adaptive locomotion and reconfigurable construction behaviors (Sutton and Barto 2018). Given the potential for highly complex sequences of emergent coordinated movement with multiple robots in dynamically changing environments, we extend this strategy to DMARL. Deep Multi-Agent Reinforcement Learning is a sub-field of DRL focused on the study of behavior of multiple agents that coexist in a shared environment. Each agent is motivated by its own rewards, taking actions based on its interests where those interests can be aligned or opposed to other agent’s interests, resulting in complex dynamics (Albrecht and Stone 2017). Our agents are trained to receive the same reward structure, learning when and how to collaborate.

We develop DMARL with self-play through a series of games simulating navigation and reconfiguration tasks of increasing complexity. The games are “played” by humans, automated algorithms, and artificially intelligent agents. Each player, or agent, controls an individual robot seeking

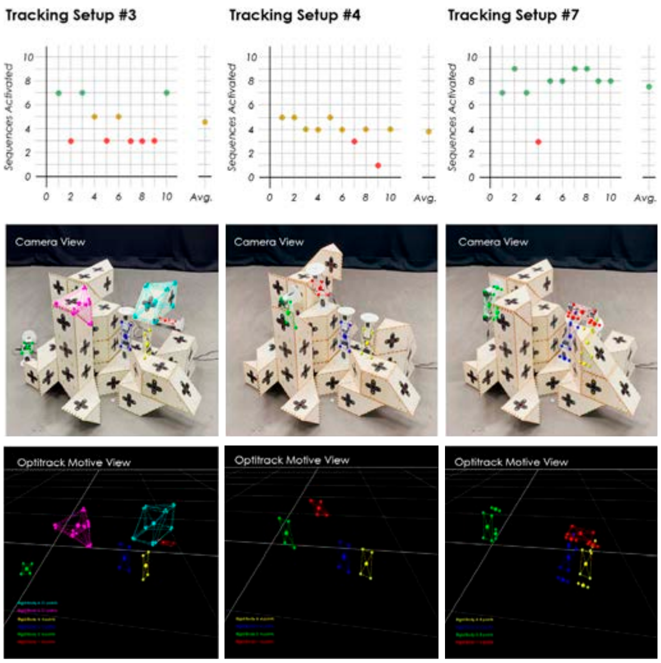

16 Motion-Capture-Based Reconfiguration Studies with OptiTrack Sensors.

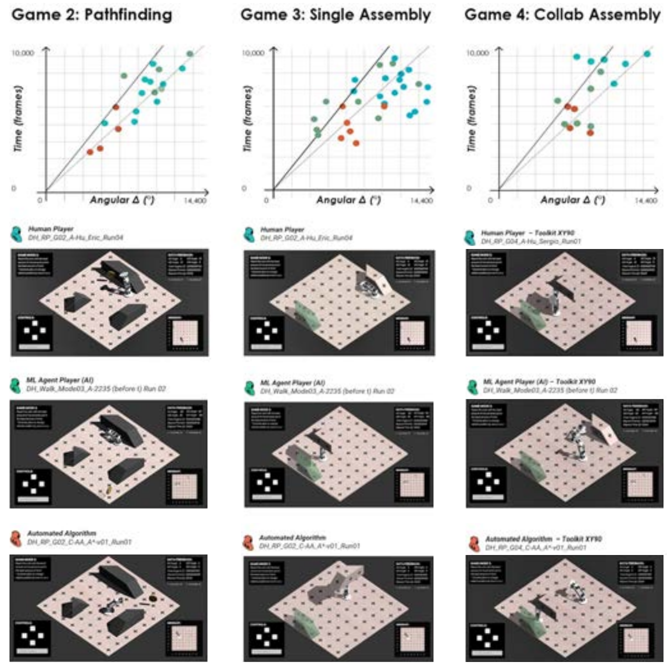

17 Deep Reinforcement Learning Gamified Studies (Human vs AI vs Algorithm).

to complete the global task of the game. Regardless of player type, all agents have access to the same playable actions and receive analogous variable observations of relevant game data. The results of each game are then assessed by speed and energetic efficiency, identifying best performing behaviors and feeding them back into the next game as new playable action sequences. The games are set additively, from simple to complex: Open Navigation, Pathfinding, Single Agent Assembly, and Collaborative Reconfiguration (Figure 13 and 14).

Aritifical intelligence players were trained with the ML agents framework for DRL in Unity 3D (Juliani et al. 2018), using the ACRR simulation environment as a training arena. At each step, a Deep Neural Network (DNN) is given a series of observations (o) and executable actions (a), and depending on its performance towards a goal, it receives a reward (r). The simulation tracks time elapsed (t) and total absolute motor angular change ($\sum|\Delta a|$), which serve in the reward function to promote speed and energetic efficiency, respectively. At the beginning of each game, element positions are randomized to ensure adaptive and scalable behaviors (Figure 13 through 15).

Once trained, the DNNs compete with humans and automated algorithms, developed specifically for each game, to generate ACRR's datasets. High-performing behaviors are sequenced as playable actions in subsequent games, combining the best performing strategies developed by humans, automated algorithms, and AI, in an additive intelligence progression from simple to complex multi-agent spatial reconfiguration.

RESULTS AND DISCUSSION

The ACRR methodology was tested through a case-study project called Diffusive Habitats, which embraces the potential of continuous spatial reconfiguration and non-linear building life cycles, aligned with an innovative distributed ownership model for communal living. Each habitat undergoes constant spatial reconfiguration to adapt to the needs of its changing community and its situated social and environmental conditions.

Reversible Robotic Material System

The system successfully physically demonstrated a comprehensive range of bespoke single-agent and multi-agent behaviors, or action sequences, involving both its active and passive parts (Figures 11 and 12). Collaborative behaviors extended the scope and effectiveness of the execution of essential reconfiguration tasks, including:

- (a) Locomotion: individual robots were able to navigate the structured environment, regardless of gravitational orientation, and translate themselves beyond it by rolling on flat surfaces. Multiple robots effectively collaborated to translate across inclined planes, and over small obstacles.
- (b) Material Pick-Up and Placement: single robots were capable of lifting and placing small and mid-size material units (~460g), while multiple robots effectively coordinated pick-up or counterweight sequences to lift

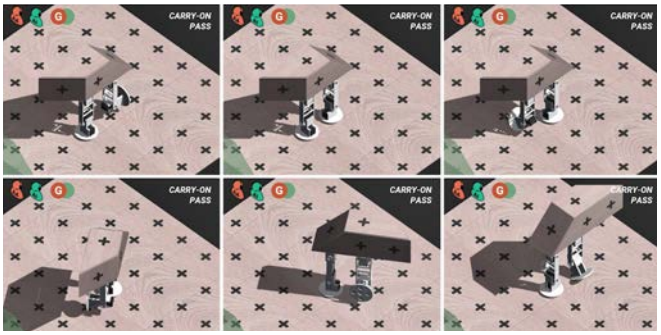

18

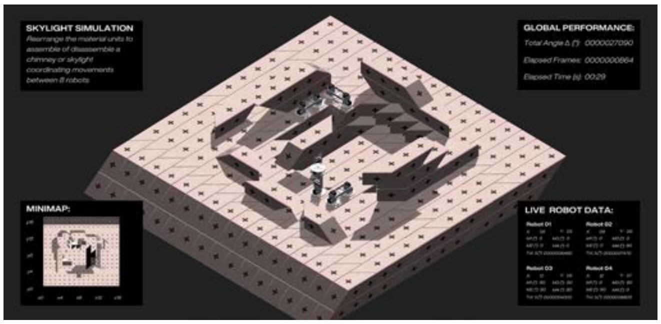

19

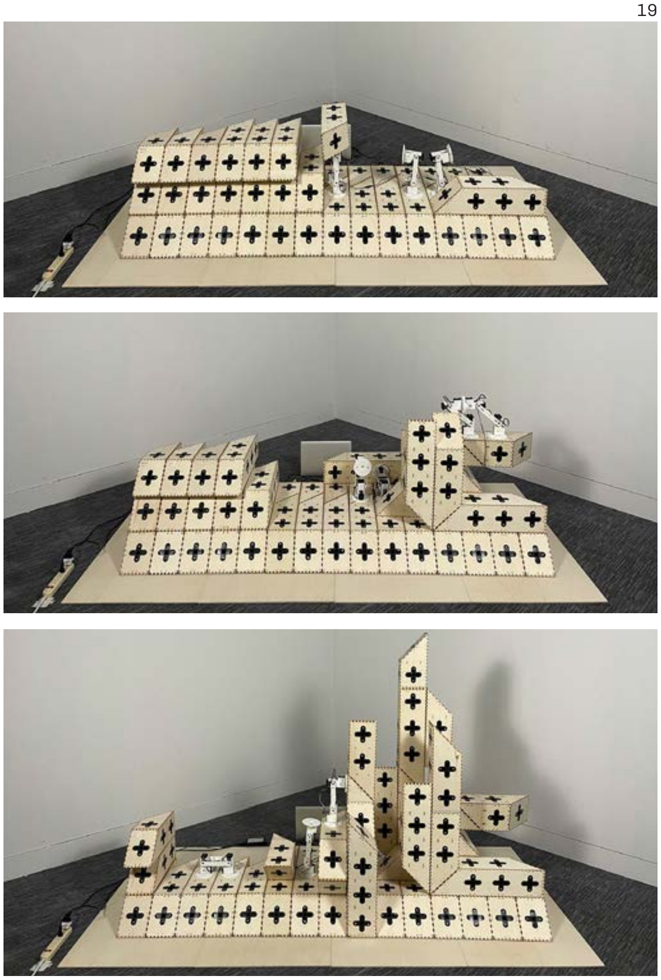

20

18 Simulation: Game 4 Implementation of Multi-Agent Deep Reinforcement Learning for Collaborative Assembly.

19 Simulation: Applied Collaborative RL in Skylight Reconfiguration.

20 Photographs: Collaborative Reconfiguration Sequences.

large material units (around 670g).

- (c) Material Translations: in collaboration, robots effectively passed and translated material ~300% times faster than individual robots, which were unable to move with the units and forced to repeat slow lift, place, and walk-around sequences.

Many of these behaviors involved the clustering of robots and material units into combined body plans, or hybrid material-robotic morphologies, which could surpass basic collaboration. Robots can operate as each other's elevators, bridges, and cranes; robots can combine with material units to move in different clustered arrangements for extra range or stability, such as bipedal or hexapedal clusters; and robots can place material units or themselves as temporary scaffolding for reconfiguration processes (Figure 11 and 12). This type of synergy, having successful analogs in nature (ant bridges), remains an uncharted research thread of great potential in the realm of architectural robotics.

Cyber-Physical Simulation and Control System

Autonomous collaborative robotic reconfiguration's Cyber-Physical Control System was calibrated by executing all the previous physical prototype behaviors from the simulator with motor sensor and Optitrack sensor feedback. The motion capture framework with Optitrack cameras was applied in a short series of experiments, which systematically evaluated different tracking setups, while activating specific sequences based on the positioning data of 4 robots and 18 material units (Figure 16). With an average of ~350% increase in sequence activation, the results validated the value of motion capture as a source of feedback data, which enables higher precision simulations and adaptability to unpredictable physical environments.

Framework for Robotic Intelligence with Deep Reinforcement Learning

We sequentially implemented four training games with human, algorithmic, and AI agent players. Deep neural networks were trained to complete each of the gamified reassembly tasks. Some tasks involved single-agent training, while others used multi-agent training to learn collaborative strategies. Each training session involved approximately five million runs on average, enabling the DNN to learn a series of strategies that maximize the reward and best achieve the objective. Success rate (r) and performance indicators, such as time elapsed (t) and total absolute motor angular change ($\sum |\Delta a|$), varied depending on player type and game, allowing for the identification and adoption of an array of effective strategies and behaviors (Figures 17 and 18).

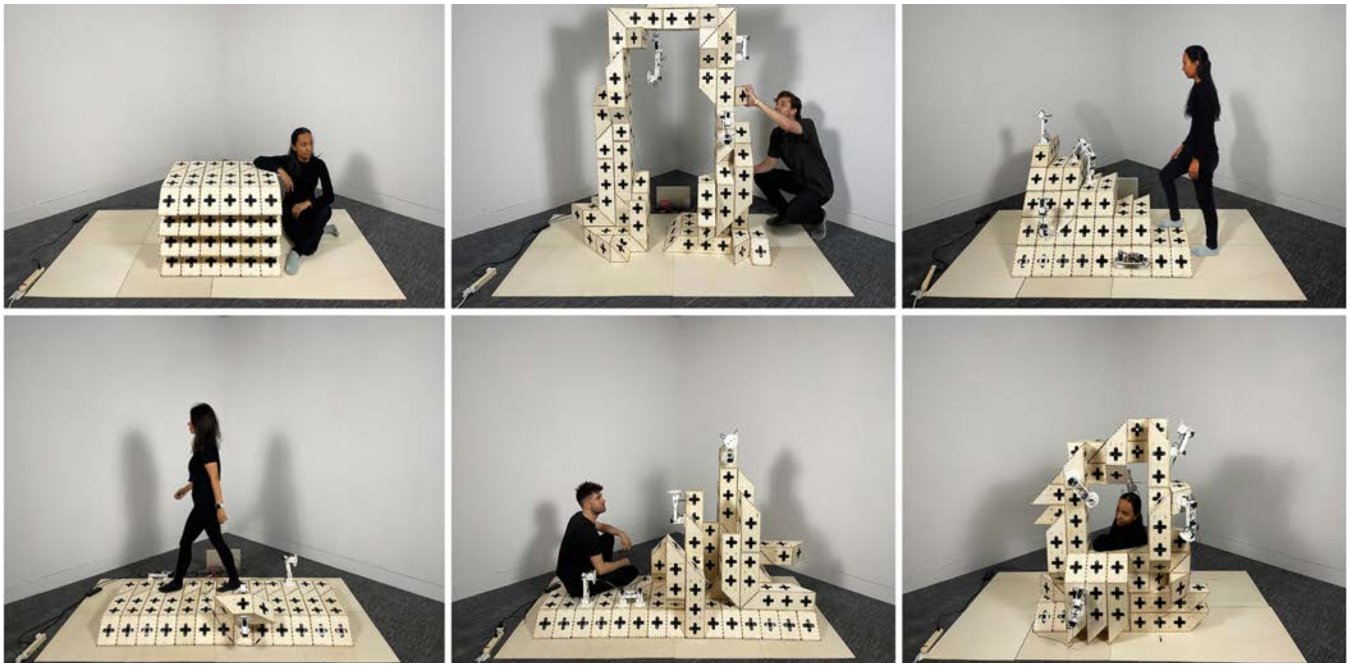

21 Photographes: Cyber-Physical Reconfiguration Studies.

21

Game 1, Open Navigation, exhibited a perfect success rate (r = 100%) from all players in each of its 50 game runs. Two highly efficient walking styles were identified and established for subsequent games: the ‘straight flip’ of the G1-AA-B automated algorithm, with 360° of total absolute angular change (∑|Δa|) and 260 frames elapsed (t) per 140mm of translation, and the ‘tilted turn’ of the G1-NN-K neural network, with 190° of total absolute angular change (∑|Δa|) and 290 frames elapsed (t) per 140mm of translation.

Game 2, Pathfinding, also displayed a perfect success rate from all players across its 60 game runs. Unsurprisingly, the G2-AA-A automated algorithm, based on the A* pathfinding algorithm (Hart 1968), outperformed both its human and AI counterparts in all relevant metrics. Therefore, it was set for the subsequent games as the default playable action for robot pathfinding across the game environment.

Game 3, Simple Assembly, exhibited varying success rates for different player types across its 30 game runs. A series of highly efficient, robot-unit behaviors were identified and sequenced for future use, including material pick-up/placement styles and strategies for single-robot transportation of material units, like the ‘side roll’ transportation strategy identified from the G3-NN-B neural network.

Game 4, Collaborative Reconfiguration, saw AI with the highest success rates across its 30 game runs. Dataset analysis revealed numerous high-performance unique behaviors, all of them collaborative, most of them emergent, ranging from diverse collaborative material passing strategies to collaborative robot-unit translations, such as the emergent ‘carry-on’ pass exhibited by groups of G4-NN-G neural networks (Figure 18).

Implementation of ACRR with DMARL

Next, the fully trained agents were tested for a complex problem of collaborating to create a skylight within a closed aggregation (Figure 19). Six trained robotic agents successfully coordinated the accurate relocation of 23 material units of diverse sizes, transforming a flat roof into a 1.5m skylight in 1min and 54sec. This simulation relied on the combination of AI-generated pass behaviors, pathfinding algorithm sequences, and human-crafted ‘turn’ coordination sequences, demonstrating the value of the mixed approach to building adaptive intelligence.

Finally, the integration of ACRR with Multi-Agent Reinforcement Learning was demonstrated through a series of cyber-physical reconfiguration tasks at 1:1 scale, involving 5 robotic prototypes and 30 to 50 material units (Figure 20). These studies demonstrate ACRR’s capacity to successfully calculate, control, and coordinate multiple robotic agents, in both physical and digital space. From walkable surfaces and arching assemblies (up to 3m tall), to furniture and partitions, these aggregations demonstrate the ability of the system to efficiently adapt itself through autonomous collaborative construction (Figure 21).

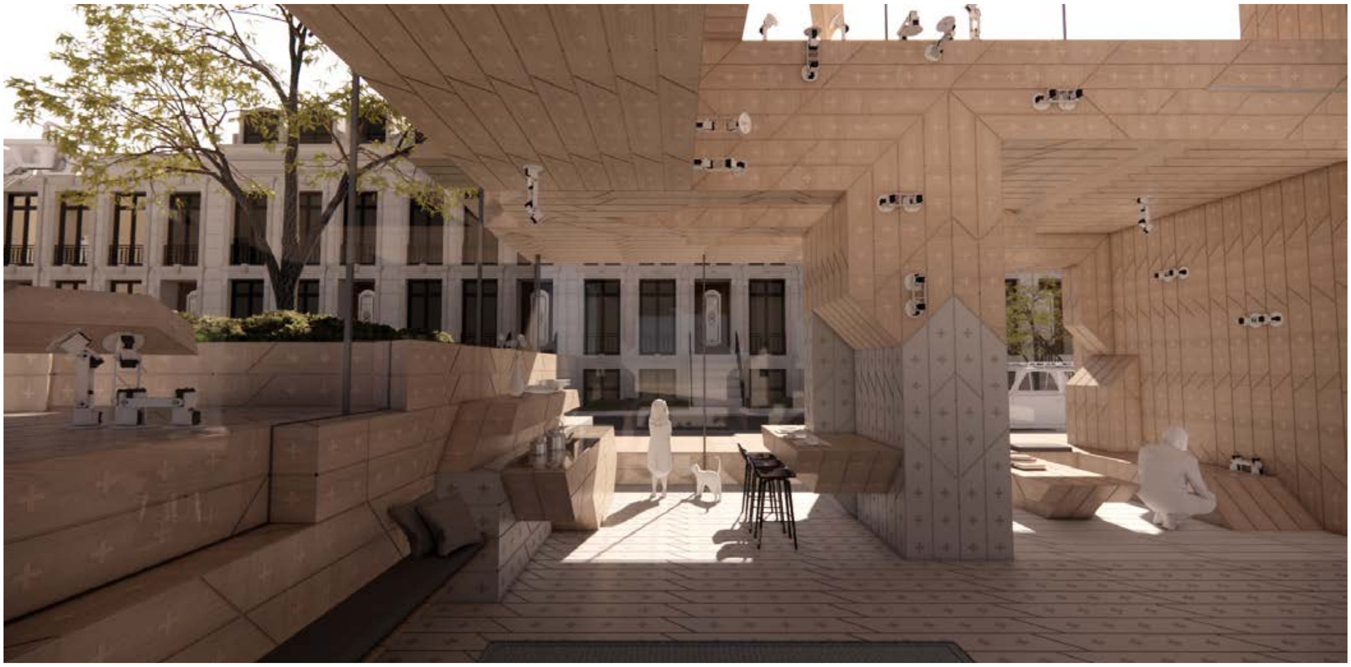

22 Rendering: Speculative Reconfiguration Ecosystems. 22

CONCLUSION

Autonomous collaborative robotic reconfiguration demonstrates a potential for autonomous, reconfigurable, and scalable construction through the collaboration of simple robots. This research challenges the building industry's traditionally linear integration of new technologies through incremental improvements to traditional processes of design, fabrication, and construction. As one of the first implementations of DMARL within CRC, the research demonstrates the potential for embedding adaptive intelligence directly into dynamic construction environments (Figure 22 through 24).

Like natural builders, ACRR has several advantages over traditional processes of construction, as well as existing industrial robotic solutions. Like termites, these adaptive builders move freely without the limitations of gantry systems or wheel-based rovers, enabling scalable solutions while efficiently processing construction tasks in parallel. By assembling reversible passive parts, assemblies can adapt over time, enabling new transformable building typologies. Simple robots with limited degrees of freedom are cheaper to build than large-scale, industrial gantry solutions, while being resilient and adaptive to individual robotic failures. Autonomous collaborative robotic reconfiguration is a collaborative eco-system where active and passive parts form hybrid body plans and exhibit emergent collaborative behaviors that enable the construction of complex structures.

The next steps in our research involve scaling up our robotic system with robust building materials and reversible locking joints, while expanding our incremental learning strategy with deep, multi-agent reinforcement learning to develop broader adaptive intelligence.

There is massive potential for CRC systems to revolutionize the way we build architecture with key challenges in autonomous assembly to overcome, including building performance, resilience, and engagement (Lu 2017). A shift toward autonomy is not just an incremental improvement, but a disruptive technological and cultural change that raises new questions surrounding how we build and interact with an architecture that is self-adaptive (Figure 21).

ACKNOWLEDGEMENTS

Living Architecture Lab, Research Cluster 3 (RC3), Architectural Design, The Bartlett School of Architecture, University College London.
Course Masters: Tyson Hosmer, Octavian Gheorghiu, Philipp Siedler
Machine Learning Tutor: Panagiotis Tigas
Theory Tutor: Jordi Vivaldi Piera
Technical Tutor: Ziming He, Barış Erdinçer
Team Diffusive Habitats (2021-2022): Sergio Mutis, Eric Hughes, Garyfallia Papoutsi, Faizunsha Ibrahim
The Bartlett BPro AD Director: Tyson Hosmer
The Bartlett BPro Director: Professor Frédéric Migayrou

Main Softwares and Libraries used:
Unity3D, ML-Agents by Juliani et al., SpatialSlur Library by David

Reeves, Tensorflow Framework by Martin Abadi et al., McNeel Rhinoceros + Grasshopper, Karamba by Preisinger, Clemens, and Moritz Heimrath, Autodesk Maya, Autodesk 3ds Max, Adobe Photoshop, Adobe InDesign.

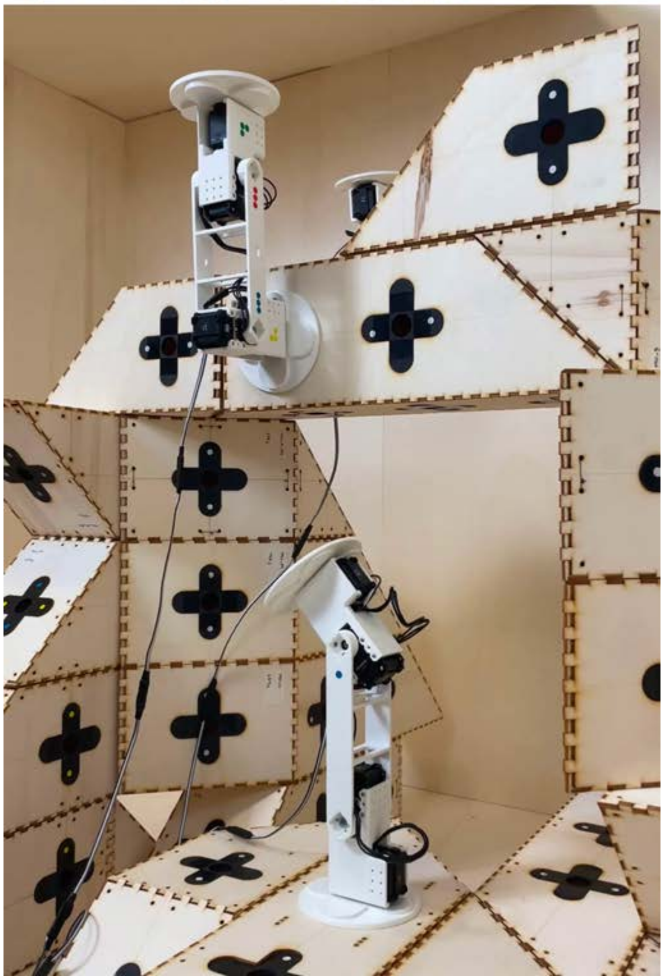

23 Photograph: Autonomous Collective Robotic Reconfiguration (ACRR), Collective Sequence.

24 Photograph: Autonomous Collective Robotic Reconfiguration (ACRR), Collective Sequence.

REFERENCES

Adel, Arash, Andreas Thoma, Matthias Helmreich, Fabio Gramazio, and Matthias Kohler. 2018. “Design of Robotically Fabricated Timber Frame Structures.” In ACADIA 2018: Recalibration: On Imprecision and Infidelity [Proceedings of the 38th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)], Mexico City, Mexico, 18-20 October 2018, edited by P. Anzalone, M. del Signore, and A. J. Wit, 394-403. CUMINCAD.

Agarwal, Rajat, Shankar Chandrasekaran, and Mukund Sridhar. 2016. "Imagining construction’s digital future." McKinsey & Company. Mckinsey Global Institute (MGI) Analysis. https://www.mckinsey.com/capabilities/operations/our-insights/imagining-constructions-digital-future#/

Albrecht, Stefano and Peter Stone. 2017. "Multiagent Learning: Foundations and Recent Trends. Tutorial." IJCAI-17 conference.

Allwright, M., N. Bhalla, H. El-faham, A. Antoun, C. Pinciroli, and M. Dorigo. 2014. “SRoCS: Leveraging stigmergy on a multi-robot construction platform for unknown environments.” In ANTS 2014: Swarm Intelligence, edited by M. Dorigo, M. Birattari, S. Garnier, H. Hamann, M. Montes de Oca, C. Solnon, and T. Stutzle, vol. 8667 of Lecture Notes in Computer Science, 158–169. Springer. doi:10.1007/978-3-319-09952-1_14.

Anderson, C., G. Theraulaz, and J.-L. Deneubourg. 2002. “Self-Assemblages in Insect Societies.” Insectes Sociaux 49 (2): 99–110. doi:10.1007/s00040-002-8286-y.

Andreen, David, Petra Jenning, Nils Napp, and Kirstin Petersen. 2016. “Emergent Structures Assembled by Large Swarms of Simple Robots.” In ACADIA Proceedings. ACADIA.

Apolinarska, Aleksandra Anna, Matteo Pacher, Hui Li, Nicholas Cote, Rafael Pastrana, Fabio Gramazio, and Matthias Kohler. 2021. “Robotic assembly of timber joints using reinforcement learning.” Automation in Construction 125 (May): 103569.

Augugliaro, Federico, Ammar Mirjan, Fabio Gramazio, Matthias Kohler, and Raffaello D’Andrea. 2013. “Building Tensile Structures with Flying Machines.” In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE.

Belousov, Boris, Bastian Wibranek, Jan Schneider, Tim Schneider, Georgia Chalvatzaki, Jan Peters, and Oliver Tessmann. 2022. “Robotic Architectural Assembly with Tactile Skills: Simulation and Optimization.” Automation in Construction 133 (January): 104006. 0926-5805. https://doi.org/10.1016/j.autcon.2021.104006.t

Bier, Henriette, Alexander Liu Cheng, Sina Mostafavi, Ana Anton, and Serban Bodea. 2018. "Robotic Building as Integration of Design-to-Robotic-Production and -Operation." In Robotic Building, edited H. Bier, 97-119. Cham, Switzerland: Springer.

Bier, Henriette and Sina Mostafavi. 2018. “Robotic Building as Physically Built Robotic Environments and Robotically Supported Building Process.” In Robotic Building, edited H. Bier, 253-271. Cham, Switzerland: Springer.

Braithwaite, Adam, Talib Alhinai, Maximilian Haas-Heger, Edward McFarlane, and Mirko Kovač. 2017. “Tensile Web Construction and Perching with Nano Aerial Vehicles.” In Springer Proceedings in Advanced Robotics, 71–88. Cham: Springer International Publishing.

Brand, Stewart. 1995. How Buildings Learn: What Happens after They’re Built. New York: Penguin.

Brugnaro, G. and S. Hanna. 2019. “Adaptive Robotic Carving Training Methods for the Integration of Material Performances in Timber Manufacturing.” In Robotic Fabrication in Architecture, Art and Design 2018, 336–348. Cham: Springer.336–348.

Detweiler, C., M. Vona, K. Kotay, and D. Rus. 2006. “Hierarchical control for self-assembling mobile trusses with passive and active links.” In Proc. ICRA 2006, 1483–1490, May 2006.

Chai, Hua, Liming Zhang, and Philip F. Yuan. 2020. “Advanced Timber Construction Platform Multi-Robot System for Timber Structure Design and Prefabrication.” In Architectural Intelligence, 129–44. Singapore: Springer Nature Singapore.

Dakhli, Zakaria, and Zoubeir Lafhaj. 2017. “Robotic Mechanical Design for Brick-Laying Automation.” Cogent Engineering 4 (1): 1361600.

Dambrosio, Niccolo, Christoph Zechmeister, Serban Bodea, Valentin Koslowski, Marta Gil-Pérez, Bas Rongen, Jan Knippers, and Achim Menges. 2019. “Buga Fibre Pavilion: Towards an architectural application of novel fiber composite building systems.” In ACADIA 2019: Ubiquity and Autonomy. Proceedings of the 39th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA), edited by K. Bieg, D. Briscoe, and C. Odom, 16–29. Austin, TX: ACADIA.

Dharmawan, Audelia G., Yi Xiong, Shaohui Foong, and Gim Song Soh. 2020. “A Model-Based Reinforcement Learning and Correction Framework for Process Control of Robotic Wire Arc Additive Manufacturing.” In 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE. 4030-4036.

Estévez, N. S., and H. Lipson. 2007. "Dynamical blueprints: Exploiting levels of system-environment interaction." In Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation, 238–244. ACM.

Fan, Yongxiang, Jieliang Luo, and Masayoshi Tomizuka. 2019. “A Learning Framework for High Precision Industrial Assembly.” In 2019 International Conference on Robotics and Automation (ICRA). IEEE. 811-817. doi: 10.1109/ICRA.2019.8793659.

Felbrich, B., N. Frueh, M. Prado, S. Saffarian, J. Solly, L. Vasey, J. Knippers, and A. Menges. 2017. “Multi-machine fabrication: An integrative design process utilising an autonomous UAV and Industrial robots for the fabrication of long-span composite structures.” In Acadia 2017 Disciplines & Disruption: Proceedings of the 37th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA), 248–259.

Felbrich, Benjamin, Tim Schork, and Achim Menges. 2022. “Autonomous Robotic Additive Manufacturing through Distributed Model-Free Deep Reinforcement Learning in Computational Design Environments.” Construction Robotics 6 (1): 15–37. doi:10.1007/s41693-022-00069-0.

Garnier, Simon, Tucker Murphy, Matthew Lutz, Edward Hurme, Simon Leblanc, and Iain D. Couzin. 2013. “Stability and Responsiveness in a Self-Organized Living Architecture.” PLoS Computational Biology 9 (3): e1002984.

Gerhart, John, and Marc Kirschner. 2007. “The Theory of Facilitated Variation.” Proceedings of the National Academy of Sciences 104 (suppl 1): 8582–89.

GlobalData. n.d. “Global Construction Outlook to 2022-Q4 2018 Update.” Accessed Sept 10, 2019. https://www.orbisresearch.com/reports/index/global-construction-outlook-to-2022-q4-2018-update.

Graham, Jason M., Albert B. Kao, Dylana A. Wilhelm, and Simon Garnier. 2017. “Optimal Construction of Army Ant Living Bridges.” Journal of Theoretical Biology 435: 184–98.

Grushin, Alexander, and James A. Reggia. 2006. “Stigmergic Self-Assembly of Prespecified Artificial Structures in a Constrained and Continuous Environment.” Integrated Computer-Aided Engineering 13 (4): 289–312.

Guo, Hongliang, Yan Meng, and Yaochu Jin. 2009. “A Cellular Mechanism for Multi-Robot Construction via Evolutionary Multi-Objective Optimization of a Gene Regulatory Network.” Biosystems 98 (3): 193–203.

Hansell, Mike. 2007. Built by Animals: The Natural History of Animal Architecture. OUP Oxford.

Harichandran, A., B. Raphael, and A. Mukherjee. 2019. “Determination of automated construction operations from sensor data using machine learning.” In Proceedings of the 4th International Conference on Civil and Building Engineering Informatics.

Hart, Peter E, Nils J Nilsson, and Bertram Raphael. 1968. “A Formal Basis for the Heuristic Determination of Minimum Cost Paths.” IEEE transactions on systems science and cybernetics 4 (2): 100-107.

Hosmer, Tyson, and Panagiotis Tigas. 2019. “Deep Reinforcement Learning for Autonomous Robotic Tensegrity (ART).” In ACADIA 2019: Ubiquity and Autonomy. Proceedings of the 39th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA), ed. K. Bieg, D. Briscoe, and C. Odom. 16–29. Austin, TX: ACADIA.

Hosmer, T., J. Wang, W. Jiang, and Z. He. 2022. “Integrated Reconfigurable Autonomous Architecture System.” In Proceedings of the 42nd Annual Conference of the Association for Computer Aided Design in Architecture: Distributed Proximities, ACADIA 2022, ACADIA.

Inoue, Tadanobu, Giovanni De Magistris, Asim Munawar, Tsuyoshi Yokoya, and Ryuki Tachibana. 2017. “Deep Reinforcement Learning for High Precision Assembly Tasks.” In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE. 819-825, doi: 10.1109/IROS.2017.8202244.

Jenett, Ben, and Kenneth Cheung. 2017. “Bill-e: Robotic Platform for Locomotion and Manipulation of Lightweight Space Structures.” In 25th AIAA/AHS Adaptive Structures Conference, 1876.

Jenett, Benjamin, Amira Abdel-Rahman, Kenneth Cheung, and Neil Gershenfeld. 2019. “Material–Robot System for Assembly of Discrete Cellular Structures.” IEEE Robotics and Automation Letters 4 (4): 4019–26.

Jokic, S., P. Novikov, J. Shihui, S. Maggs, C. Nan, and D. Sadan. 2017. “Mini Builders.” http://robots.iaac.net/.

Juliani, Arthur, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy et al. 2018. "Unity: A general platform for intelligent agents." arXiv preprint arXiv:1809.02627 (2018).

Kayser, Markus, Levi Cai, Sara Falcone, Christoph Bader, Nassia Inglessis, Barrak Darweesh, and Neri Oxman. 2018. “Design of a Multi-Agent, Fiber Composite Digital Fabrication System.” Science Robotics 3 (22).

Kirschner, Marc. 2009. "Variations in Evolutionary Biology." In The Architecture of Variation, edited by Lars Spuybroek. London: Thames and Hudson. p.26-33.

Leder, Samuel, Ramon Weber, Dylan Wood, Oliver Bucklin, and Achim Menges. 2019. “Distributed robotic timber construction.” In ACADIA 2019: Ubiquity and Autonomy. Proceedings of the 39th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA), ed. K. Bieg, D. Briscoe, and C. Odom. 510–519. Austin, TX: ACADIA.

Leder, S., H. Kim, O. S. Oguz, N. Kubail Kalousdian, V. N. Hartmann, A. Menges, M. Toussaint, and M. Sitti. 2022. “Leveraging Building Material as Part of the In-Plane Robotic Kinematic System for Collective Construction.” Adv. Sci. 9: 2201524. https://doi.org/10.1002/advs.202201524.

Lindsey, Q., D. Mellinger, and V. Kumar. 2011. Robotics: Science and Systems VII. Cambridge, MA: MIT Press.

Lloret-Fritschi, E., F. Scotto, F. Gramazio, M. Kohler, K. Graser, T. Wangler, L. Reiter, R. J. Flatt, and R. Mata-Falcón. 2018, RILEM Int. Conf. on Concrete and Digital Fabrication, Zürich, Switzerland.

Łochnicki, G., N. K. Kalousdian, S. Leder, M. Maierhofer, D. Wood, and A. Menges. 2021. “Co-Designing Material-Robot Construction Behaviors: Teaching distributed robotic systems to intuitively leverage active bending for light-touch assembly of bamboo bundle structures.” In ACADIA – Realignments – towards critical computation [Proceedings of the ACADIA Conference 2021], 11.04 2021.

Loosemore, Martin. 2015. “Construction innovation: fifth generation perspective.” Journal of Management in Engineering 31: 4015012.

Lu, Andong. 2017. “Autonomous Assembly as the Fourth Approach to Generic Construction.” Architectural Design 87(4): 128-33.

Luo, Jianlan, Eugen Solowjow, Chengtao Wen, Juan Aparicio Ojea, Alice M. Agogino, Aviv Tamar, and Pieter Abbeel. 2019. “Reinforcement Learning on Variable Impedance Controller for High-Precision Robotic Assembly.” In 2019 International Conference on Robotics and Automation (ICRA), 3080-3087. IEEE. doi: 10.1109/ICRA.2019.8793506.

Luo, J., and H. Li. 2020. “Dynamic Experience Replay.” In Proceedings of the Conference on Robot Learning in Proceedings of Machine Learning Research, 100:1191-1200. Available from https://proceedings.mlr.press/v100/luo20a.html.

Lutz, Matthew, and Chris Reid. 2015. Living Ant Bridge. https://www.princeton.edu/news/2015/11/30/ants-build-living-bridges-their-bodies-speak-volumes-about-group-intelligence.

Martinoli, A, A.J Ijspeert, and F Mondada. 1999. “Understanding Collective Aggregation Mechanisms: From Probabilistic Modelling to Experiments with Real Robots.” Robotics and Autonomous Systems 29 (1): 51–63.

Mechtcherine, Viktor, Venkatesh Naidu Nerella, Frank Will, Mathias Näther, Jens Otto, and Martin Krause. 2019. “Large-Scale Digital Concrete Construction – CONPrint3D Concept for on-Site, Monolithic 3D-Printing.” Automation in Construction 107 (November): 102933.

Melenbrink, Nathan, and Justin Werfel. 2019. "A Swarm Robot Ecosystem For Autonomous Construction, 2017." In Robotic Building: Architecture in the Age of Automation, edited by Gilles Retsin, Manuel Jimenez, Mollie Claypool, and Vicente Soler, 88-90. München: DETAIL. https://doi.org/10.11129/9783955534257-018.

Melenbrink, Nathan, Justin Werfel, and Achim Menges. 2020. “On-Site Autonomous Construction Robots: Towards Unsupervised Building.” Automation in Construction 119 (November): 103312.

Melenbrink, Nathan, Paul Kassabian, Achim Menges, and Justin Werfel. 2017. "Towards force-aware robot collectives for on-site construction." In Proceedings of the ACADIA Conference, pp. 382-391. 2017.

Meng, Yan, and Jing Gan. 2008. “A Distributed Swarm Intelligence Based Algorithm for a Cooperative Multi-Robot Construction Task.” In 2008 IEEE Swarm Intelligence Symposium. IEEE. 1-6.

Miller, Norman. "The Industry Creating a Third of the World's Waste." BBC Future Planet. BBC Global News Ltd., December 16, 2021. https://environment.ec.europa.eu/topics/waste-and-recycling/construction-and-demolition-waste_en.

Moses, Matthew S., Hans Ma, Kevin C. Wolfe, and Gregory S. Chirikjian. 2014. “An Architecture for Universal Construction via Modular Robotic Components.” Robotics and Autonomous Systems 62 (7): 945–65.

N. S. Estévez, 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.

Napp, Nils, Olive R. Rappoli, Jessica M. Wu, and Radhika Nagpal. 2012. “Materials and Mechanisms for Amorphous Robotic Construction.” In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. Algarve, Portugal: IEEE 4879–4885.

Napp, Nils, and Radhika Nagpal. 2014. “Distributed Amorphous Ramp Construction in Unstructured Environments.” Robotica 32 (2): 279–90.

Nitschke, G.S., M.C. Schut, and A.E. Eiben. 2012. “Evolving Behavioral Specialization in Robot Teams to Solve a Collective Construction Task.” Swarm and Evolutionary Computation 2 (February): 25–38.

Ngwepe, Lusca, and Clinton Aigbavboa. 2015. “A Theoretical Review of Building Life Cycle Stages and Their Related Environmental Impacts.” Journal of Civil Engineering and Environmental Technology 2(13): 7-15.

Parisi, Fabio, Valentino Sangiorgio, Nicola Parisi, Agostino M. Mangini, Maria Pia Fanti, and Jose M. Adam. 2023. “A New Concept for Large Additive Manufacturing in Construction: Tower Crane-Based 3D Printing Controlled by Deep Reinforcement Learning.” Construction Innovation, Vol. ahead-of-print No. ahead-of-print. https://doi.org/10.1108/CI-10-2022-0278

Parker, C.A.C., Hong Zhang, and C.R. Kube. 2003. "Blind Bulldozing: Multiple Robot Nest Construction." In Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453), accessed May 31. IEEE. http://dx.doi.org/10.1109/iros.2003.1248950

Parker, Chris A. C., and Hong Zhang. 2006. “Collective Robotic Site Preparation.” Adaptive Behavior 14 (1): 5–19.

Parter, Merav, Nadav Kashtan, and Uri Alon. 2008. “Facilitated Variation: How Evolution Learns from Past Environments to Generalize to New Environments.” PLoS Computational Biology 4(11): e1000206.

Petersen, Kirstin, Radhika Nagpal, and Justin Werfel. 2011. “TERMES: An Autonomous Robotic System for Three-Dimensional Collective Construction.” In Robotics: Science and Systems VII. Robotics: Science and Systems Foundation. The MIT Press. 257-264.

Petersen, Kirstin H., Nils Napp, Robert Stuart-Smith, Daniela Rus, and Mirko Kovac. 2019. “A Review of Collective Robotic Construction.” Science Robotics 4 (28).

Petersen, Kirstin, and Radhika Nagpal. 2017. “Complex Design by Simple Robots: A Collective Embodied Intelligence Approach to Construction.” Architectural Design 87(4): 44–49.

Popescu, George A., Tushar Mahale, and Neil Gershenfeld. 2006. “Digital Materials for Digital Printing.” In NIP & Digital Fabrication Conference, 58-61.

Reid, Chris R., Matthew J. Lutz, Scott Powell, Albert B. Kao, Iain D. Couzin, and Simon Garnier. 2015. “Army Ants Dynamically Adjust Living Bridges in Response to a Cost–Benefit Trade-Off.” Proceedings of the National Academy of Sciences 112 (49): 15113–18.

Retsin, Gilles, and Manuel Jimenez Garcia. 2016. “Discrete Computational Methods for Robotic Additive Manufacturing: Combinatorial Toolpaths.” In ACADIA 16: Posthuman Frontiers: Data, Designers, and Cognitive Machines; Proceedings of the 36th Annual Conference of the Association for Computer Aided Design in Architecture, 332-341.

Retsin, Gilles. 2019. “Bits and Pieces: Digital Assemblies: From Craft to Automation.” Architectural Design 89(2): 38-45.

Rettenmeyer, Carl W. 1963. "Behavioral Studies of Army Ants: Estudios de Comportamiento de Hormigas Guerreras." The University of Kansas Science Bulletin 44 (9): 281–465.

Rosman, Guy, Changhyun Choi, Mehmet Dogar, John W. Fisher IIIl, and Daniela Rus. 2018. “Task-Specific Sensor Planning for Robotic Assembly Tasks.” In 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE.

Rubenstein, Michael, Christian Ahler, and Radhika Nagpal. 2012. “Kilobot: A Low Cost Scalable Robot System for Collective Behaviors.” In 2012 IEEE International Conference on Robotics and Automation, 3293–98.

Saboia da Silva, Maira, Vivek Thangavelu, Walker Gosrich, and Nils Napp. 2018. “Autonomous Adaptive Modification of Unstructured Environments.” In Robotics: Science and Systems XIV. Robotics: Science and Systems Foundation.

Schneirla, T.C. 1972. Army Ants: A Study in Social Organization. W.H. Freeman & Co Ltd.

Silver, M. 2017. "Fabricate 2014: Negotiating Design & Making." UCL Press, 242.

Soleymani, Touraj, Vito Trianni, Michael Bonani, Francesco Mondada, and Marco Dorigo. 2015. “Bio-Inspired Construction with Mobile Robots and Compliant Pockets.” Robotics and Autonomous Systems 74 (December): 340–50.

Stewart, Robert L., and R. Andrew Russell. 2006. “A Distributed Feedback Mechanism to Regulate Wall Construction by a Robotic Swarm.” Adaptive Behavior 14 (1): 21–51.

Stuart-Smith, Robert. 2016. “Behavioural Production: Autonomous Swarm-Constructed Architecture.” Architectural Design 86 (2): 54–59. doi:10.1002/ad.2024. Note, journal was cited 55 times.

Sutton, Richard S., and Andrew G Barto. 1998. Introduction to Reinforcement Learning. Cambridge, MA: The MIT Press.

Sutton, Richard S., and Andrew G. Barto. 2018. Reinforcement learning: An introduction. MIT press

Terada, Y., and S. Murata. 2004. "Automatic Assembly System for a Large-Scale Module Structure - Hardware Design of Module and Assembler Robot." In 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566). IEEE

Terada, Yuzuru, and Satoshi Murata. 2008. “Automatic Modular Assembly System and Its Distributed Control.” The International Journal of Robotics Research 27 (3–4): 445–62.

Tibbits, Skylar. 2012. “Design to Self-Assembly.” Architectural Design 82(2): 68-73.

Tibbits, Skylar. 2017. “From Automated to Autonomous Assembly.” Architectural Design 87(4): 6-15.

Theraulaz, Guy, and Eric Bonabeau. 1995. “Coordination in Distributed Building.” Science 269 (5224): 686–88.

Thomas, Garrett, Melissa Chien, Aviv Tamar, Juan Aparicio Ojea, and Pieter Abbeel. 2018. “Learning Robotic Assembly from CAD.” In 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE. 3524-3531. doi: 10.1109/ICRA.2018.8460696.

Tucker, Cody, Xiliu Yang, August Lehrecke, Mathias Maierhofer, Rebeca Duque Estrada, and Achim Menges. 2022. “A Collaborative Multi-robot Platform for the Distributed Fabrication of Three-dimensional Fibrous Networks (Spatial Lacing).” In Proceedings of the 7th Annual ACM Symposium on Computational Fabrication (SCF '22), Article 3, 1–18. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3559400.3561995.

Van de Kamp, T., Dörstelmann, M., dos Santos Rolo, T., Baumbach, T., Menges, A., Knippers, J. 2015. "Beetle Elytra as Role Models for Lightweight Building Construction." Entomologie Heute 27: 149–158.

Wagner, Hans Jakob, Martin Alvarez, Ondrej Kyjanek, Zied Bhiri, Matthias Buck, and Achim Menges. 2020. “Flexible and Transportable Robotic Timber Construction Platform – TIM.” Automation in Construction 120 (December): 103400.

Werfel, Justin, Donald Ingber, and Radhika Nagpal. 2007. “Collective Construction of Environmentally-Adaptive Structures.” In 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE.

Werfel, Justin, Donald Ingber, and Radhika Nagpal. 2007. “Collective Construction of Environmentally-Adaptive Structures.” In 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE.

Werfel, J., and R. Nagpal. 2006. “Extended Stigmergy in Collective Construction.” IEEE Intelligent Systems 21 (2): 20–28.

Yun, Seung-kook, Mac Schwager, and Daniela Rus. 2011. “Coordinating Construction of Truss Structures Using Distributed Equal-Mass Partitioning.” In Springer Tracts in Advanced Robotics, 607-623. Berlin, Heidelberg: Springer Berlin Heidelberg.

Zeng, A., S. Song, J. Lee, A. Rodriguez, and T. Funkhouser. 2019. “TossingBot: Learning to Throw Arbitrary Objects with Residual Physics.” In Proceedings of Robotics: Science and Systems (RSS). http://arxiv.org/abs/1903.11239.

Zhang, Ketao, Pisak Chermprayong, Feng Xiao, Dimos Tzoumanikas, Barrie Dams, Sebastian Kay, Basaran Bahadir Kocer, et al. 2022. “Aerial Additive Manufacturing with Multiple Autonomous Robots.” Nature 609 (7928): 709–17. doi:10.1038/s41586-022-04988-4.

Zhang, M., S. Vikram, L. Smith, P. Abbeel, M.J. Johnson, and S. Levine. 2019. “SOLAR: Deep Structured Representations for Model-Based Reinforcement Learning.” http://arxiv.org/pdf/1808.09105v4.

IMAGE CREDITS

Figure 1: © Lutz, Matthew, and Chris Reid. 2015. Living Ant Bridge.

Figures 2 through 22: Images by the authors and Living Architecture Lab, Research Cluster 3 (RC3), The Bartlett School of Architecture, UCL, and team members of Diffusive Habitats (2021-2022).

Tyson Hosmer is an architect, researcher, educator, and founder working at the intersection of design, computation, AI, and robotics. He is an Associate Professor and the Director of the Architectural Design (AD) Master’s program at the Bartlett School of Architecture UCL, where he directs the Living Architecture Lab (Research Cluster 3). Tyson’s research focuses on autonomous reconfigurable architectural systems with artificial intelligence, cognitive agent-based systems, and generative design models with deep learning. His 15 years in practice include working as a Senior Associate Researcher and head of ZHA-Social research group with Zaha Hadid Architects, as well as with Axi:Ome, Asymptote Architecture, Kokkugia, and Cecil Balmond Studio where he was the Research Director for several years.

Sergio Mutis is a researcher within ZHA Code at Zaha Hadid Architects, an architectural robotics tutor at the Bartlett, UCL,

and a graduate visiting tutor at Imperial College London and the Royal College of Art. He holds a B.Arch. from Universidad Javeriana, Colombia, and a M.Arch. from the Bartlett, UCL, UK.

Eric Hughes is an architectural and computational designer who has worked in both Georgia (USA) and New York. Eric holds a B.Arch. from Mississippi State University and a M.Arch. from the Bartlett School of Architecture, UCL.

Ziming He is a senior designer and a software developer focusing on computational design, generative design, architectural visualization and robotics. He is currently working as a Senior Designer in Zaha Hadid Architects and is a Technical Tutor with RC3, Bartlett School of Architecture, UCL.

Octavian Gheorghiu is a London-based architect, design researcher and educator. He holds a Professional Diploma from the Bartlett School of Architecture and a Masters's Degree from the Architectural Association DRL. He is working as an Associate Partner at Foster + Partners, helping to deliver projects at several scales, ranging from urban planning to architecture and product design. As an educator, Octavian is a Lecturer at the Bartlett School of Architecture and a Tutor at the AADRL.

Philipp Siedler is an AI researcher, engineer, game developer, educator, and architect. He graduated from the Architectural Association, Design Research Laboratory (AA DRL) in 2017, and worked for the ZHAI Analytics and Insights research group at Zaha Hadid Architects. Currently, he is an AI Research Engineer at Aleph Alpha, working on a new generation of AI systems beyond supervised learning. He teaches at the University College London (UCL) Living Architecture Lab (RC3) at Bartlett School of Architecture.

Barış Erdinçer is a designer and researcher at Zaha Hadid Architects focusing on developing applications with generative design, machine learning, and immersive interactive environments. He is a researcher and technical tutor with the Living Architecture Lab (RC3) at the Bartlett School of Architecture.