Building Like Animals: Using Autonomous Robots to Search, Evaluate and Build

Stanislav Roudavski
Melbourne School of Design, The University of Melbourne
stanislav.roudavski@cantab.net

Abstract

Typical architectural design is reliant on prescriptive top-down decisions that seek efficiency through simplification. Erection of designs produced in this way depends on specially prepared construction sites and access routes. Such building procedures are destructive and cannot gracefully integrate with existing ecosystems. By contrast, this paper is motivated by the desire to assimilate newly built artificial structures with pre-existing dynamic, amorphous, highly variable and fragile sites. As one of the possible approaches to this challenge, this paper discusses the prospects offered by the utilization of intelligent robots. This discussion relies on an empirical study of autonomous robots that outlines the capabilities necessary for such construction.

Architecture of Change

Typical architectural designs rely on static conditions: given briefs, regulated construction sites and industrial materials. They depend on expert predictions and expert control. And yet, there is a growing recognition that the power of design expertise in finding solutions is fundamentally limited. “All buildings are predictions. All predictions are wrong.” [1] Appreciation of this condition has led to the emergence of approaches that seek to embrace and facilitate change, including, for example designing for disassembly and reuse, open buildings, systemic design, service design or holistic approaches to sustainability.

Animal Architecture

The space of this short article does not allow a detailed discussion of animal architecture. However, some of its important characteristics are suggestive as ways of thinking about holistic architectural design and the types of systems that could support it.

Animals construct architecture as a way of extending their control, catching prey at greater distances using spider webs or building mounds to influence interior climatic conditions. Such structures can be seen as extensions of animals’ organisms. They actively modify the surrounding environment, sometimes dramatically. These modifications construct ecological niches for the builders, in competition with others. At the same time, animal-made structures create opportunities for exploitation by other forms of life, sometimes leading to complex symbiotic dependencies. Importantly, animal architecture has to be considered as states within extended evolutionary sequences that also result in particular biomechanics, behaviors, ways to reproduce, procedures for raising progeny, tactics of socialization and so on. Thus, structures produced by animals are locally and temporarily satisfactory solutions that are continuously changing in response to the changing circumstances. As compromises framed by multiple moving targets, these solutions are always sub-optimal. They cannot serve as optimal models for narrow and well-defined tasks, where they can be outcompeted by deterministic artificial systems. However, biologically inspired solutions can be useful in situations where designed systems need to be adaptable, rather than optimal. Evolved life excels at survival. Correspondingly, designed systems with capabilities for long-term autonomy may prove beneficial in situations where human capabilities are limited or direct access by humans is not possible. Beyond extraterrestrial environments and capillary vessels, that are already navigated by autonomous robots, such situations might include environments characterized by large, complex and long-term processes such as those of ecosystems or cities.

Architectural Robotics

Robotics is a growing area of interest in architecture. At the moment, the attention of the field is on industrial robotic arms that are attractive because they are common and comparatively cheap. [2] These devices belong to the subfield of manufacturing automation, or industrial robots. As such, they are intended for environments with existing manufacturing processes, control mechanisms, quality requirements and automation measures. In this context, their purpose is to benefit from financial opportunities that might be available through reduction of waste, improvement of quality, decrease of downtime and so on. Current speculative research in architecture strives to develop software and hardware modifications that can redirect such robots towards operations on bespoke tasks that are common in architectural design.

Autonomous Robots in Dynamic Environments

Beyond automation, robots can be seen as devices that act as programmable replacements for humans but are physically situated intelligent agents within hybrid socio-technical ecologies, capable of more or less independent sensing, analysis, decision taking and action. [3] Here, "[a]n agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators." [4]

The purpose of such agents is not in replacing people but in being better, at least in some aspects, than biological systems, projecting human capabilities into inaccessible environments or enabling otherwise impossible tasks. Robotic agents can achieve these goals not because of their superior intelligence but through a combination of characteristics common to machines: they can operate continuously and for long periods, maintain speed and precision, cope with large volumes of data, etc.

The difference between automation and autonomy can be seen as the contrast between – on one hand – better tools (extension of human arms, tools with “long handles” such as drones, tools moving in bounded regions, not dissimilar to automatic weaving, such as manufacturing robots) and – on the other hand – adaptive, unsupervised agents, such as mobile robots.

Automation is the use of tools for the execution of precise, repetitive actions in well-understood environments. Such environments are treated as controlled, closed worlds and are modelled with confidence, on the presumption that it is possible to build a model that can capture all that is relevant for a set of given operations. Automation robots are common in factories because certain operations on production lines can be specified as such closed worlds. However, some worlds – such as typical architectural environments – are not closed. Where models of such worlds are made, they are complex and incomplete. In such cases, comprehensive preplanning of actions is not feasible or desirable. Instead, sensing and analytic capabilities can be employed for adaptation in response to feedback. Mobile robots operating in open worlds have to understand their environment and thus require intelligence.

In general, the vision of an intelligent independent robot is as yet unfulfilled, especially in the commercial domain. At the moment, most of such robots are employed within small subsets of their possible applications, limiting their potential contribution. However, knowledge accumulated by robotics coupled with the ongoing democratization of coding and mechatronics opens new opportunities as it becomes easier to prototype autonomous systems for new domains.

Related Work

In architecture, interest in autonomy, typically in construction, has a long history. [5] This paper suggests that some of the ideas generated during that journey can be productively revisited with the current tools. Some of the recent related work focuses on the applications of adaptation or emergence, without specifically considering the use of robotic devices. The use of practical robotic experiments to test these approaches is still rare. [6]

In engineering, one existing example includes implementations of bottom-up construction behaviors. [7] Another explores prototypes for intelligent interaction with unknown environments. [8]

More broadly, the ‘autonomic computing’ initiative is of relevance to the present discussion. [9] Its attempts to develop self-maintaining systems have similar goals but fall outside of the scope of this paper.

Experimental Setup

Architectural engagement with the field of robotics, especially with robots as intelligent agents, requires the establishment of new collaborations, protocols, practices and skills. Considering the role of architects in construction, Kieran and Timberlake argue that to retain relevance they have to act as managers of intelligence and “the overseers of the exchange of information”. [10]

Seeking to illuminate the capabilities needed for such overseeing, this research engages in an empirical investigation of robots. It accepts that rich interplay with the surrounding environments and multiple types of co-present agents constitute necessary conditions for intelligence. Given this, the work discussed below investigates which components and methods are required and how they can be coordinated.

The scenario presented in this paper explores construction from objects that can be found on a site, for example rubble left after an earthquake or items collected for recycling. It was presumed that a site containing such materials is initially unknown and partially inaccessible. This type of challenge requires 1) a capability to execute actions in the physical world; 2) an ability to plan actions, such as the use of resources; and 3) a capacity for real-time reasoning based on perception. The experimental setup involves a miniature robotic system that autonomously maps its environment and uses found materials to build shelter-like structures according to high-level criteria specified by humans. Based on a “real but contrived” design, the results of this experiment contribute to the challenge of understanding which capabilities are required for intelligent physical interactions with complex dynamic environments such as architectural sites.

The speculative prototype assumes a landscape-like environment with various degrees of slope. This terrain contains abstracted resources, represented by sticks of three separate lengths, that can be used as building material. At the beginning, the robot does not know the shape of the environment or the distribution of resources. Criteria for selection, sorting and building are supplied by human designers. The robot aims to collect and use material in an efficient manner, saving time and energy.

The setup does not explore all elements that might be needed for the realization of full scale autonomous construction robots. For example, a robotic arm instead of a mobile platform runs a static vision system instead of on-board perception and uses schematic end effectors to simplify aspects that are not central to the investigation. However, an approach that experiments with physical prototypes can highlight important challenges of integration at other scales because some key constraints that limit small experimental systems – size, processing capabilities, power, energy consumption – are similar to those that define implementations at many scales.

Robotic Platform

The project deploys a custom assembly consisting of:

1. Robotic arms with six degrees of freedom (Fig. 1)
2. End effectors: grippers, air blowers, sensors
3. Depth and video cameras
4. Notebook computers and microcontrollers

A robot operating system with custom-written components is used for calibration, control, communication, coordination, perception processing, planning and acting.

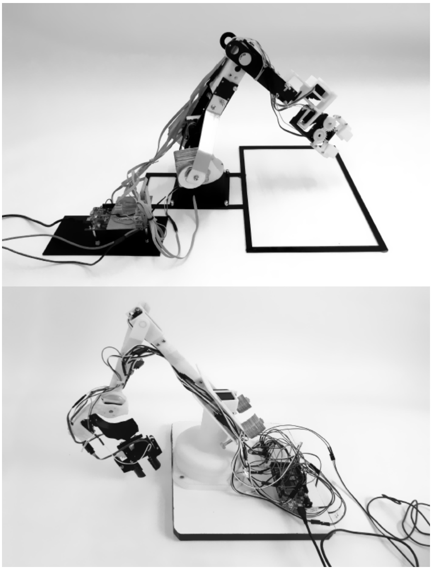

Fig. 1. *Autonomous Architecture*, 2014, two of the several robot arms designed, constructed and operated during the project.

Perception

For perceptions, the project uses depth information and color images supplied by Kinect sensors. These could be positioned in arbitrary locations and calibrated using markers. In real-world scenarios, such information can be provided from multiple sources (satellite, aerial photography, geographical information systems, on-demand surveying (e.g., with drones) on-board sensing, etc.).

Reasoning

The system employs the biologically inspired operational architecture that consists of a behavioral (or skill) layer, a deliberative (or sequencing and planning) layer and an interaction layer. [11, 12] Unlike purely reactive systems (such as boids) that are only able to operate in the present, deliberative agents of this type can consider the past, the present and the future.

Path Finding. To operate in complex environments, mobile robots need to determine their own paths. In the experiments discussed here, these paths have to be generated from the robot's current positions to locations where resources can be picked. These paths can respond to the properties of the terrain such that obstacles (e.g., walls or watered areas) or hard-to-pass regions (e.g., mud) are avoided. Access to resources may require intermediate tasks, such as a construction of access routes by removing material or building bridges. In this project's implementation, three-dimensional evaluations using versions of A* algorithm are used to compare costs of access to the locations mapped by the perceptual system (Fig. 2).

Task planning. In response to a high-level human command or an instruction supplied by an overseeing artificial-intelligence system, a robot needs to plan a sequence of actions. In this case, it:

- Finds an area that is suitable for construction; in this prototype, less sloping areas, as determined through the analysis of the depth-camera data, are deemed to be more suitable;
- Finds an area suitable for exploration;
- Finds resources by analyzing vision data;
- Sorts resources into types;
- Creates a subset of resources within a distance threshold and generates their coordinates (Fig. 2, left);
- Plots paths to these resources and compares retrieval costs (Fig. 2, right);
- Determines the type needed for the current stage of construction and retrieves it; and
- Places the resource into the structure (Fig. 3).

This basic sequence can be made more sophisticated and complex: the site might need to be prepared for construction, for example by removing the resources and placing them into a temporary storage; access routes might need to be prepared to reach blocked regions, etc.

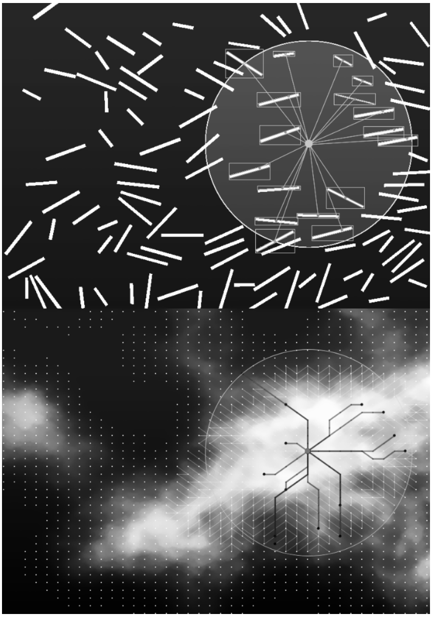

Fig. 2. *Autonomous Architecture*, 2014, detection of resources and comparative path planning.

Actions

Based on the plans described above, the robot performs two main activities: exploration and plan execution. Exploration is done virtually, as an analysis of perceptual data incoming from a vision system that sees all of the scene. In some versions, this is complemented by the sensors mounted on the robot. When an experiment begins, the system starts exploring and building a representation of its environment. The scenario where the perceptual data is supplied by a vision system external to the robot is possible in the real world but unlikely to be sufficient and can to be complemented or replaced with on-board perception that will make map construction more complex and gradual. [13]

When the system receives an order, it uses this representation to build the initial state of the planning problem. The system then performs task planning. When ready, it executes the plan. When execution is completed, the system reverts to the exploration mode that allows it to update its representation of the environment, taking into account any changes.

Exploration. In this project, exploration consists of obtaining and processing of the depth and color information in order to 1) subdivide the site into regions suitable for retrieval of resources and construction, based on heights and slopes and 2) identify resources by the processing of color blobs, filtering of artefacts, calculation of centroids and orientations, estimation of sizes and sorting.

Plan execution. The execution of the plans is done with robotic arms instead of actual mobile robots. This simplification does not undermine the overall value of the experiment because execution is done in a way that is compatible with the use of mobile robots. On the other hand, the use of mobile platforms would introduce new design challenges some of which were considered but are outside of the scope of this paper.

Execution consists of three behaviors: 1) the 'move' behavior takes the robot to a destination resource; 2) the 'take' behavior consists of grasping a resource and requires alignments to be accurately resolved; and 3) the 'build' behavior consists of placing a resource onto a target site in a desired relationship to an already built structure, if any. The form of the resulting structure can be prespecified or generated in response to available resources. While the project has conducted experiments in this area, the detailed discussion of them is also outside the scope of this presentation.

Fig. 3. *Autonomous Architecture*, 2014, retrieval and construction sequence.

Discussion

Using variants of the setup discussed above, the project could experiment with a range of scenarios involving operations with objects and amorphous matter, such as earth and water in a range of speculative architectural proposals.

The experiment confirmed that interdisciplinary teams with engineering and architectural backgrounds can sketch complete systems, designing hardware and software with low-level components and integrating them into workflows that use multiple robotic arms, perceptual devices and custom end-effectors. The use of higher-level tools, such as industrial robotic arms for hardware or ready operating systems such as the Robot Operating System (ROS) can bring obvious benefits in robustness and capacity. However, the use of such systems requires greater prior knowledge, impeding interdisciplinary, participatory collaboration in speculative system design.

Five years since their analysis, this paper agrees with the conclusions of Moussete and Dore that sketching in hardware remains quite challenging despite the recent emergence of new toolkits. [14] Yet, such sketching is necessary for design innovations in the domains where conventional architectural ontologies and epistemologies come into question and reliance on existing types and practices is insufficient or constraining. Thus, this project's experiments highlighted that:

- In unknown environments, a condition that is typical to landscape or urban settings, the behavior of an autonomous system depends on the operation of sensors (and actuators that enable sensing). Collection, analysis and integration of such dynamic and continuous data, as an evolution of site analysis, is a new and interesting challenge in the architectural domain;
- Construction is a sustained and contingent process of many operations. All of these operations can fail in distinctive ways and the system has to cope with these failures without undermining the overall process. In industrial and product design, development of necessary robustness is typically done via prototyping and versioning, an as-yet uncommon workflow in architecture;
- Given that resources found in actual sites are not purpose-specific, they have to be found, recognized, extracted and sorted before use. Fully preprogrammed or purely reactive behaviors are insufficient in such cases because adequate autonomous systems have to evaluate and compare multiple scenarios while avoiding possible impasses.

These findings suggest that the effort needed for the prototyping of robotic systems in architecture can be justified by the benefits that are likely to ensue from the integration of increasingly common industrial robots with the capabilities of artificial intelligence.

Conclusion

This paper presents an empirical study that outlines the capabilities necessary for the tasks of autonomous construction in the physical world. The study used a working prototype to test system integration, practicalities of possible workflows, required knowledge, available building blocks and near-future prospects. The project demonstrated that it is possible to construct a simplified but functional system for design speculation with cheap hardware and commonly accessible software tools. With modifications and additions, the prototypes constructed for the project can support diverse design scenarios and promote otherwise-impossible exchanges with relevant stakeholders.

Acknowledgements

The aforementioned work was developed in collaboration with Denny Oetomo, of Mechanical Engineering at the University of Melbourne, by multidisciplinary teams working on parallel projects. The visual examples above are taken from an experiment that was led by Juan Yang.

References

1. Stewart Brand, *How Buildings Learn: What Happens After They Are Built* (Harmondsworth: Penguin, 1995), 178.
2. Sigrid Brell-Çokcan and Johannes Braumann, "Introduction," in *Rob|Arch 2012: Robotic Fabrication in Architecture, Art, and Design*, ed. Sigrid Brell-Çokcan and Johannes Braumann (Wien: Springer, 2013), 8.
3. Robin Murphy, *Introduction to AI Robotics* (Cambridge, MA: MIT Press, 2000).
4. Stuart J. Russell, Peter Norvig, and Ernest Davis, *Artificial Intelligence: A Modern Approach*, 3rd ed. (Upper Saddle River: Prentice Hall, 2010), 34.
5. Roozbeh Kangari and Rita Gregory, "Feasibility of Automating Military's Environmental Operations," *Automation in Construction* 5, no. 6 (1997).
6. Jan Willmann et al., "Aerial Robotic Construction Towards a New Field of Architectural Research," *International Journal of Architectural Computing* 10, no. 3 (2012).
7. Justin Werfel, Kirstin Petersen, and Radhika Nagpal, "Designing Collective Behavior in a Termite-Inspired Robot Construction Team," *Science* 343, no. 6172 (2014).
8. Stéphane Magnenat, Roland Philippsen, and Francesco Mondada, "Autonomous Construction Using Scarce Resources in Unknown Environments," *Autonomous Robots* 33, no. 4 (2012).
9. Richard Murch, *Autonomic Computing* (Upper Saddle River, NJ: Prentice Hall, 2004).
10. Stephen Kieran and James Timberlake, *Refabricating Architecture: How Manufacturing Methodologies are Poised to Transform Building Construction* (New York: McGraw-Hill, 2004), 22.
11. Murphy, *Introduction to AI Robotics*.
12. Erann Gat, "On Three-Layer Architectures," in *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, ed. David Kortenkamp, Peter R. Bonasso, and Robin Murphy (Cambridge, MA: MIT/AAAI Press, 1997).
13. Magnenat, Philippsen, and Mondada, "Autonomous Construction."
14. Camille Moussette and Fabricio Dore, "Sketching in Hardware and Building Interaction Design: Tools, Toolkits and an Attitude for Interaction Designers" (in Proceedings of Design and Complexity: Design Research Society Conference, ed. by David Durling, et al., DRS: Monteal, CA, July 7–9, 2010).