Research Statement: Scalable Autonomous Construction

Samuel Lensgraf

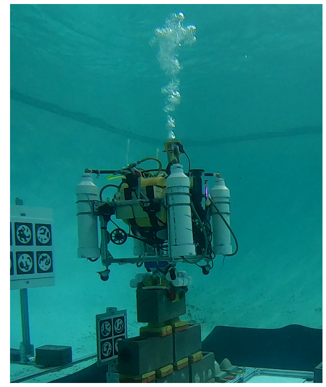

Figure 1: Autonomous underwater construction robot building a wall.

I intend to enable autonomous robots to build large scale structures made of practical materials such as cement, steel, or earth. Achieving this goal requires solving problems in construction planning algorithms, robot hardware design, reliable sensing and control, and the computational design of a system for construction, including blocks, connectors, and grasping strategies. Large scale autonomous construction can allow us to rapidly adapt to changing environmental circumstances by rapidly deploying protective infrastructure that can be moved and re-deployed as needed.

I was awarded an NSF Graduate Research Fellowship Program fellowship which supports my Ph.D. research on autonomous underwater construction. This work led to the first free-floating underwater construction robot [1]–[3]. Figure 1 shows our autonomous underwater construction system. I have developed motion planning algorithms to increase the efficiency of fused-deposition 3D printers [4]–[7]. My first paper in this series was awarded the Best Automation Paper award at the IEEE International Conference on Robotics and Automation (ICRA). This work was the basis of a patent [8].

I am interested in enabling **reliable** autonomous construction and fabrication in harsh environments, such as underwater, in space, or in remote and windy locations on land. In harsh environments, we cannot assume that the robot can ever be controlled with perfect precision: currents or winds may buffet the robot at any moment. We cannot assume it is possible to perceive objects to be manipulated perfectly. Our robot hardware designs and planning algorithms must explicitly take into account ways to **mitigate error**.

I design computational and mechanical techniques which work together to mitigate uncertainty. Mitigating uncertainty is an under-represented perspective in current robotics research, providing fertile ground for interdisciplinary papers and proposals to solve algorithmic, computational, mechanical, and system design challenges. Planning construction sequences amid uncertain placements requires new algorithms for construction planning. Designing construction materials which forgive inevitable placement errors demands new approaches to computational design. Implementing and testing these new algorithms on real robot systems establishes a feedback loop for continuous improvement.

Construction sequencing algorithms Reliably building large scale structures with simple robots demands construction sequencing algorithms that optimize the construction process according to the capabilities of the robots. I have explored planning algorithms that improve 3D printing **speed**, algorithms for construction with buoyant robots that improve **energy efficiency**, and algorithms for the deployment of localization infrastructure and construction to **increase accuracy**. These threads establish a basis on which I intend to build construction and fabrication planning algorithms that allow simple robots to operate at massive scale with high accuracy.

When humans build structures, they partially structure the environment to achieve accuracy despite variations on the site: scaffolds allow higher reach, and guide lines ensure walls are straight. I have directly explored ways our underwater construction robot can maintain and extend its operating area by **structuring its environment** using a combination of automatically deployed infrastructure and a new planning algorithm [3]. Figure 2 shows our construction robot modifying its localization infrastructure. We use a new model of the positioning certainty of visual fiducial markers and a hardware design for reconfigurable visual fiducials.

To enable large scale construction with a buoyant robot, I designed a **convex optimization** approach to allocate buoyancy to transport construction materials. This formulation allows the construction robot to balance two resources: compressed air for buoyancy and battery power. This is the first planning algorithm for construction which balances two finite resources.

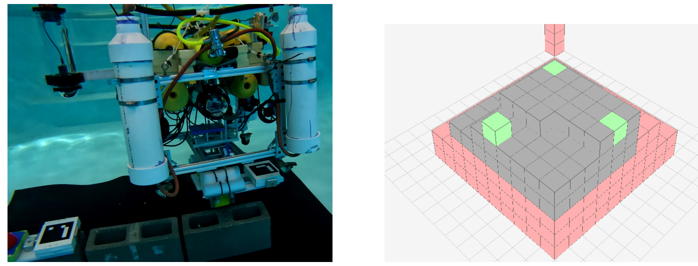

(a) (b)

Figure 2: (a) Robot moving a visual fiducial to extend its workspace. (b) Idealized representation for planning.

I have explored motion planning techniques to improve the efficiency of fused deposition 3D printers. My work in this area created a new type of motion planning for 3D printers which allows them to operate without the layer-by-layer constraint while guaranteeing collision-free printing.

It is my goal to unify these threads of research into **error-aware planning** algorithms which can allow robots to build structures with minimal, autonomously deployed instrumentation of the environment. My work on 3D printing establishes a basis for searching over the space of feasible construction sequences. My work on autonomously deployed and maintained infrastructure validates these ideas in deployment. This type of planning algorithm could enable CNC machines and 3D printers with arbitrary workspaces, sensor networks which can expand and move autonomously using a frame of reference they provide, autonomous construction of modular structures of arbitrary scale, or mining and scrapping robots which structure their environment as they tunnel into it.

Immediate research plan To support my long term vision, I have in mind several immediate research projects. First, develop scaffolding which can be autonomously deployed and extended by an underwater robot. My existing implementation of simultaneous localization infrastructure deployment and construction requires a first layer of blocks to be placed. This project will involve hardware advances, and a new planning problem: to limit error build accumulation as foundations are repeatedly deployed to move the operating area.

Second, formalize and design general algorithms for the problem of error-aware planning. When autonomously deployed objects establish a coordinate frame in which a robot operates, care must be taken to limit drift in that coordinate frame as objects are moved. Given known error distributions on the positions of sensing infrastructure, how can we plan ways to move that infrastructure which maintain accuracy? This can apply to mobile sensor networks and construction robots.

Computational design I will design rigid geometries which help robots achieve manipulation tasks. These rigid bodies will be shaped according to the manipulation strategy used to grasp and manipulate them. Using their geometry alone, they will transform motion into the desired behavior of the system. Designing these rigid bodies will require new simulation techniques and automatic design algorithms. These geometries will complement and extend the planning algorithms discussed in the previous section.

Assembly with passive, error-correcting modules introduces new design constraints on both the structures that are built and the robots that build them. To design construction materials which slide together passively, we must balance a trade-off between ease of assembly and the tightness. As modules fit together more tightly, the likelihood of jamming increases. My work on the kinematics of rigid body chains is a first step towards balancing this trade-off [9]: I defined a **sparse linear program** which captures the directions of maximum flex of the structure using local distance constraints. We can rapidly judge the long range effects of local changes to the geometry of single components by understanding the flexibility of structures. Figure 3 shows an example output of our flexibility algorithm on a large sheet of modules. Tightening local constraints reduces the aggregate flexibility.

Immediate research plan I will develop **differentiable simulation** techniques for designing rigid bodies that passively aid in completing manipulation tasks. Differentiable simulators are an increasingly popular tool for training machine learning systems which interact with the real world. Differentiable simulators in robotics applications have primarily been targeted at training control strategies rather than designing components which rigidly collide and align.

With the needed differentiable simulation framework in place, we can design rigid geometries for various applications. I am particularly interested in automatically co-designing grasping strategies and graspable rigid geometries. Co-designed grasping strategies and rigid geometries could inform the design of construction robots or allow us to design tools for humanoid robots (and humans) with limited dexterity.

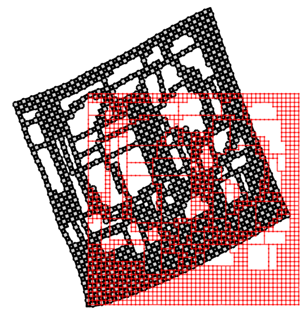

Figure 3: Flexibility of a sheet of loosely connected rigid bodies. Red marks initial positions.

Practical robot system development and deployment Designing and deploying large scale prototypes of new robotic systems serves as inspiration for my research. It is my belief that implementing new systems in the real world is one of the best ways to find surprising new directions of research. To that end, my work focuses on choosing a new goal for a capability for a robot system and then branches off more specialized projects based on the problems I encounter. This approach facilitates both **collaboration and serendipity**.

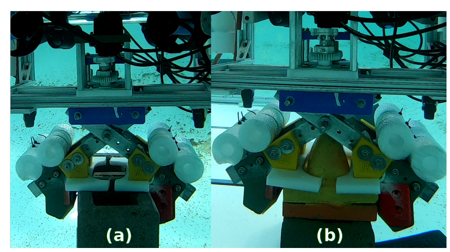

Figure 4: Robot grasping construction materials using a co-designed manipulator.

I have developed the first low cost AUV platform for doing manipulation research [2], [10]. It is also the first buoyant construction robot. This robot serves as a basis for validating the importance of error correcting geometries and grasping strategies. Figure 4 shows a complementing grasping strategy and manipulator design. The grasping strategy corrects error on two dimensions and the manipulator handles another three.

I collaborated with Dartmouth Professor Adithya Pediredla who is experienced in modeling computer vision systems to enable my work on simultaneous infrastructure deployment and construction. Together, we developed the first model of the noise distribution of visual fiducial markers with respect to relative position. Using that model to plan construction sequences required solving a new planning problem: planning for assembly with the additional constraint of sensor coverage and movement. To solve this problem, I collaborated with a graduate student in algorithms. Combining her theoretical knowledge of clustering and my knowledge of the practical deployment problems led us to a solution of the new planning problem.

My work on the autonomous construction robot has inspired new directions in **experimental robotics**. I mentored a masters student through the process of developing and publishing work on automatic exposure control for perceiving visual fiducial markers with high certainty. It is the first work which looks at improving the noise of visual fiducial marker readings using an active exposure control method. This work has been accepted at the International Symposium on Experimental Robotics (ISER) 2023.

Immediate research plan I will develop a walking CNC machine robot. The robot will be equipped with a cutting tool and feet which will allow it to move and extend its workspace. Using the cutting tool, it can carve its own foot holds. This robot could allow the scrapping of arbitrarily sized structures, or mining in locations where it is not feasible to move large infrastructure. This would be the first machine of its kind and would introduce both hardware design challenges and validate the error-aware planning strategies discussed above.

[1] S. Lensgraf, A. Sniffen, E. Honnold, J. Jain, Z. Zitzewitz, W. Wang, A. Quattrini Li, and D. Balkcom, "Droplet: Towards Autonomous Underwater Assembly of Modular Structures," in *RSS*, 2021.

[2] S. Lensgraf, D. Balkcom, and A. Quattrini Li, "Buoyancy enabled autonomous underwater construction with cement blocks," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 5207–5213.

[3] S. Lensgraf, A. Sarkar, A. Pediredla, D. Balkcom, and A. Quattrini Li, *Scalable underwater assembly with reconfigurable visual fiducials*, Oct. 2023. arXiv: 2310.19408.

[4] S. Lensgraf and R. R. Mettu, "Beyond layers: A 3D-aware toolpath algorithm for fused filament fabrication," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden: IEEE, May 2016, pp. 3625–3631, ISBN: 978-1-4673-8026-3.

[5] S. Lensgraf and R. R. Mettu, "An improved toolpath generation algorithm for fused filament fabrication," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, Singapore: IEEE, May 2017, pp. 1181–1187, ISBN: 978-1-5090-4633-1.

[6] S. Lensgraf and R. R. Mettu, "Incorporating Kinematic Properties into Fused Deposition Toolpath Optimization," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid: IEEE, Oct. 2018, pp. 8622–8627, ISBN: 978-1-5386-8094-0.

[7] C. Yoo, S. Lensgraf, R. Fitch, L. M. Clemon, and R. Mettu, "Toward Optimal FDM Toolpath Planning with Monte Carlo Tree Search," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 4037–4043.

[8] S. Lensgraf and R. Mettu, "Method and system for rapid and efficient three-dimensional printing," US10399280B1, Sep. 2019.

[9] S. Lensgraf, K. Itani, Y. Zhang, Z. Sun, Y. Wu, A. Quattrini Li, B. Zhu, E. Whiting, W. Wang, and D. Balkcom, "PuzzleFlex: Kinematic motion of chains with loose joints," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 6730–6737.

[10] S. Lensgraf, A. Sniffen, A. Quattrini Li, and D. Balkcom, "Towards the autonomous underwater construction of cement block structures with free-floating robots," May 2022.