# **Evolutionary Fabrication**

# A System of Autonomous Invention

Timothy J Kuehn Union College 807 Union St. Schenectady, NY 12308 kuehnt@garnet.union.edu

## ABSTRACT

Evolutionary algorithms have had success in designing complex objects, ranging from antennae to telescope lenses. However, evolutionary design is limited by the ability of a simulation to accurately represent the physical world. Additionally, evolved designs carry no set of specific instructions describing how to physically create such a design. Our approach, called Evolutionary Fabrication, evolves a process rather than a product. This system of evolution can, in principle, automatically invent and build anything, from soft robots to new toys. We have implemented Evolutionary Fabrication in designing "EvoFab", a machine that consists of four components: A) a genotype for printing objects, consisting of a linear set of instructions sent to a Fab@Home, an open-source 3D printer; B) a way to evaluate printed objects using custom machine vision algorithms; C) a way to automate printing by implementing a custom conveyor belt; D) a way of elaborating upon designs by implementing a genetic algorithm. In the near term, we aim to produce an evolved arch. Current results indicate increased fitness over time. Future improvements are possible through restrictions in extrusion along the Y-axis as well as refining fitness evaluation to be less exploitable.

#### **Categories and Subject Descriptors**

I.2.1 [Artificial Intelligence]: Robotics

#### **General Terms**

Design, Experimentation

#### Keywords

Evolution, Fabrication, Genetic Algorithms, Vision

## 1. INTRODUCTION

Evolutionary algorithms have had success in designing complex objects, ranging from antennae used in NASA's

*GECCO'12 Companion*, July 7–11, 2012, Philadelphia, PA, USA. Copyright 2012 ACM 978-1-4503-1178-6/12/07 ...\$10.00. Space Technology 5 (ST5) mission [5] to astronomical telescope lenses [1]. By evolving the antennae designs using evolutionary algorithms, [5] produced, within one month, two antenna designs suitable for the ST5 mission satellites. [1], in turn, used evolutionary algorithms to automatically design an optical lens, completely from scratch. They found that genetic algorithms were capable of making "humancompetitive" results – lenses that were as good, if not better, than human-designed lenses. Indeed, one of the evolved lenses actually infringed upon a patent of a preexisting lens; another evolved lens improved upon another existing lens patent.

However, evolutionary algorithms have a tendency to exploit their medium. If an evolutionary algorithm is run in a simulation that doesn't exactly replicate a real-world scenario, one may find evolved objects that are suitable for the software in which they are tested but incapable of fulfilling their original intention in practice. [10] found that, when simulating a physical environment, the simulation must be "reasonably accurate"; else, various bugs or rounding errors inevitably occur, leading to undesired solutions that obtain high fitness only through the exploitation of these discrepancies. [2] coined this problem the reality gap: desirable characteristics can be missed and undesirable results can occur when solutions are evolved solely in simulation.

One way to skirt this issue is to evolve solutions in simulation but evaluate their fitness in physical tests. Such a method would prevent any high-fitess traits from evolving if they only have high fitness from exploiting simulations. However, what happens when an evolved design is too difficult to transform into an actual object? This problem is called the **fabrication gap**: evolved designs may be well described, but they often carry no set of specific instructions describing how to get to such an end result [8]. Consider a picture of a soufflé: it may look tasty, but someone with merely a picture would have a difficult time replicating it, due to the hidden process that led to the end result. In fact, [3] have proven that the problem of assembling a final product from blueprints is NP-Complete. Thus, the task of constructing an object from a relatively simple design may not be that difficult for a person to do, but the ease of construction rapidly falls off as the initial design increases in complexity.

[11] provided an elegant solution to the reality gap. Tasked with creating a circuit that could discriminate between 1kHz and 10kHz signals, *without* the benefit of a clock, he used a reprogrammable circuitboard to evolve the *actual hardware* rather than a circuit design. This was very interesting be-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

cause, in simulation, a circuitboard alone would never be able to discriminate between two signals without the aid of a clock. However, [11] found that, when running real-world tests rather than simulated results, the circuitboard developed a pattern that exploited electromagnetic interactions to successfully discriminate between the two signals.

[12] later coined the term "embodied evolution" for an approach similar to that taken by [11]: they showed how an autonomous population of robots could evolve behaviors that would allow them to successfully compete each other, without the aid of a central computer running a genetic algorithm. The main takeaway for our purposes, however, is not that it was decentralized: it's that instead of evolving blueprints that would then later be converted into real products; the robots themselves were directly evolved. This notion of embodied evolution leads us back to our current research.

Our approach, called evolutionary fabrication, bridges the reality and fabrication gaps by combining evolutionary design with a physical process of fabrication. By doing so, we ensure that any evolved product already comes with the build instructions, in addition to ensuring that any way in which the evolutionary algorithm exploits the process will be a replicable exploitation. EvoFab, our machine that implements evolutionary fabrication, works in a three-stage cycle: first an object is printed, then its fitness is evaluated, and then the print space is cleared for the next object to be printed.

This process acts on each member of a population. After a whole population has been printed and evaluated, the best fit members breed a new generation, for which the cycle begins anew. By refining its products over time until it produces a result that fits the design requirements, evolutionary fabrication is, in principle, capable of automatically inventing and building anything, from soft robots to new toys.

In this paper, we will detail the previous work, EvoFab 0.1, and will detail the improvements that went into creating EvoFab 0.2. We will describe the individual mechanisms that, together, make EvoFab. We will show our most currently produced results and discuss current pitfalls of EvoFab 0.2, including ways in which we plan to improve it in the future.

#### 2. PREVIOUS WORK: EVOFAB 0.1

[9] have shown that it is possible to implement a genetic algorithm into the process of fabrication, successfully using a process of evolutionary fabrication to print two-dimensional objects, such as letters. Such an interactive evolutionary algorithm requires a person to judge with their own, subjective eyes the best products of a generation. Additionally, between generations, a person would physically remove the printed objects and reset the platform to be ready for more objects to be printed. By addressing these two primary issues, EvoFab (Figure 1) is now capable of making its own judgment on the level of fitness of printed objects, in addition to being able to fully automate the process autonomously, almost completely removing humans from the loop.

# 3. EVOFAB 0.2: A FULLY AUTOMATED EVO-LUTIONARY FABRICATION SYSTEM

EvoFab is, essentially, a machine that implements a genetic algorithm. EvoFab creates each population via a three-



Figure 1: "EvoFab" consists of a Fab@Home printer, computer vision software to determine fitness, and a conveyor belt, all controlled by an evolutionary algorithm.



Figure 2: A graphical representation of the threestage process: print, evaluate, recycle.

stage process (Figure 2): first, an object is printed; the object is then evaluated; then, the object is moved off the printing platform to begin the process anew. The information for each printed object is stored within a genotype controlled by a genetic algorithm, written in python. What follows is a deconstruction of EvoFab 0.2 into its basic components.

The primary process of fabrication is controlled by Fab@Home, an open-source 3D printer designed by [6] and useful for its low cost and relative ease of use. Fab@Home operates by extruding material through a syringe and depositing it onto a platform, constructing objects layer-by-layer. The carriage that holds the syringe is free to move along the X- and Yaxes. The platform upon which the material is deposited is free to move along the Z-axis.

Fab@Home is currently in its second revision, Model 2 [4]; however, due to currently greater access to the API of the Model 1, we have used the Model 1 as the basis for Evo-Fab. Additionally, previous work attempted to make use of the Model 2's ability to extrude plastic, which would have allowed for magnitudes longer durations of printing without material refills. However, printing with plastic requires the plastic to be melted, and the associated high temperatures to do so necessitate caution and constant vigilance by the user. Because this negates the ability of EvoFab to act autonomously, it was decided that the long-term benefits of other materials (discussed below) outweighed the ability of plastic to print for long periods without refills.

Fab@Home normally builds its products by interfacing through USB with a program that contains the products' blueprints in .STL files. However, our Fab@Home is outfitted with a serial connection in lieu of USB, allowing us to send commands directly to Fab@Home [9]. The utility of this is that it allows the evolutionary algorithm to store genotype encodings as a list of commands. The commands that we use to control the printer's actions are as follows:

- extrude This command causes a small amount of material to be deposited onto the print platform.
- **beginExtrude** This command, rather than send a command directly to the printer, controls the action of the other commands. When activated, all other commands except endExtrude will send their command coupled with an extrude command. Effectively, all other commands say "do this while extruding" when beginExtrude is activated.
- **endExtrude** This command deactivates beginExtrude.
- **goUp** Raises the print platform. The platform starts at its max height and cannot be moved upward until a command has been sent to lower it first. This prevents problems of the platform bumping into the syringe (which would, in effect, break the system). Similarly is **goDown**.
- goLeft Causes the print carriage to move left along the X-axis (in the negative X-axis direction). goLeft and its similar commands all act within certain bounds, outside of which the command ceases to effect movement of the print carriage. For example, the X-axis range may be [-300, 300], so that if the X-axis position is currently -300, goLeft will have no effect on print carriage's movement. Similarly are goRight, goIn, and goOut.

With a printer ready for use, the next step was deciding what material to print with. Previously, [9] chose silicone bath caulk as the material of choice. With new goals, however, come new requirements, and after attempts with plastic (described above) and silicone caulk, we settled on a brand of modeling compound similar to Play-Doh. Silicone caulk is easily extrudable, readily available, and comes in many colors, which is useful in allowing computer vision software to easily differentiate a printed object from its background. However, it is also sticky when first printed, and its cure time of approximately thirty minutes for faster-drying variants is too long to wait between prints. Thus, the material would inevitably stick to the print platform, making automation difficult. Some workarounds were attempted (discussed below), but the Play-Doh variant proved much more usable for our purposes. It has the same benefits of being readily available in many colors and easily extrudable without the drawback of stickiness upon first being extruded. This lack of stickiness comes with its own set of problems (see results), but it has proven to be the best option that has been tried thus far.

When printing objects, original iterations did not have movement in the Z-axis due to restrictions put on the system by the camera (see below). However, more recent revisions to the vision system allowed us to free Z-axis movement. This, in turn, allowed us to set the platform much closer to the syringe tip at the beginning of each fabrication. Previously, the platform was at a constant distance of a few inches from the syringe tip. This tended to cause circular extrusions in which the thread of material would spiral downward. This effect, in turn, caused a high degree of unpredictability in how certain instructions would translate to the print: for instance, a command to extrude in a straight line along the X-axis previously would have resulted in something more similar to a sinusoidal function than a linear thread. Now, however, this has been vastly improved, as the average distance between syringe tip and platform is much smaller.

Using the computer vision libraries openCV, we have developed computer vision software that works in tandem with a camera affixed to the front of the printing platform. In this way, we can reliably control the method of evaluation. The camera that we used for this process is an Ipevo Point 2 View USB Camera, useful for its ability to focus on close-up images. Additionally, the Point 2 View is supported by open-source Ubuntu drivers, making installation and access to a useful API relatively easy. Additionally, the Point 2 View is versatile in its ability to be positioned in various ways, allowing us to attach it directly to the print platform, thus freeing the platform to move along the Z-axis without worry that the platform will move out of visible range of the camera.

To allow for automation, we introduced a conveyor belt into EvoFab. Once an object has been evaluated, it is moved off of the platform and deposited in a disposal container via a conveyor belt that interfaces with the evolutionary algorithm via USB. Through this system, EvoFab can run unattended for approximately thirty minutes before requiring a refilled syringe. Not including refilling, EvoFab can run unattended indefinitely. Thus, the only factor that currently inhibits EvoFab from working completely independently of human attendence is, simply, a large enough syringe.

Initially, when we were still testing plastic printing, we used an Automated Build Platform Kit developed by Maker-Bot Industries for use with their own MakerBot 3D printer. Because the MakerBot also uses plastic, the Automated Build Platform provided a heated base to print onto and would have meshed very well with the printer. However, when switching to silicone caulk, we needed to devise a new conveyor belt, because the Automated Build Platform, as a conveyor belt, would have quickly become too messy to be useable. We therefore devised a new conveyor belt that worked much like a scroll works: instead of depositing products after being printed and evaluated, the "scroll" would roll it up into one side of the scroll, destroying the print but clearing the platform. This solved the short term problem of the objects sticking to the platform; however, it merely delayed the inevitable work required of a human to fix the platform. Once the scroll ran out on one end, it needed to be replaced.

Thus, it did not entirely solve the problem of complete automation. Thus, when we converted to modeling material, it was a natural transition to turn the "scroll" back into a conveyor belt much like the initial Automated Build Platform. There were two primary reasons why we did not go back to the original MakerBot conveyor belt: first, it was simpler to build a larger conveyor belt, given that the Fab@Home used a much larger platform than the size of the conveyor belt. Second, a custom conveyor belt allowed us to control the color of the surface, making the job of the computer vision software that much easier.

In evolutionary algorithms, a different software implementation is required for each type of object one wishes to evolve. Thus, due to the proof-of-concept nature of this research, we implemented software that evaluates exactly one type of object. In choosing what object to evolve, we looked for some shape that is currently not easily produced by a 3D printer. The reason we wanted something that is not easily produceable is that evolutionary algorithms are only truly useful in designing things that have not previously been designed - otherwise, there would be no need to use the evolutionary algorithm in the first place. In determining what a 3D printer cannot easily produce, we focused on the fact that Fab@Home always prints from the bottom-up. In other words, it builds upwards, layer by layer, and cannot construct an object with a portion that is "floating" without any supporting material underneath. Consider, for example, an arch. An arch's supporting columns are easily constructed by Fab@Home, but how will the middle area be created? It cannot deposit material onto mid-air. While a softwarelimited approach may have trouble designing a method of construction for such an arch, an evolutionary algorithm is not restricted to follow any set of rules, allowing it to freely explore all possibilities, thought-of or unthought-of by the algorithm's creators. Thus, the potential for finding a solution greatly increases when using an evolutionary algorithm, making "archness" a good object to evolve as a proof of concept.

In evaluating archness of a printed object, fitness increases proportionally to the percent of overhanging mass that an object contains. Figure 3 shows how such a fitness is evaluated: an image is captured by a camera that views the printing stage. Then, the image is thresholded so that the printed object is white and the background is black. This is made simple by printing in a color negative to that of the background, in this case pink being the negative of green. Then, a bounding box is drawn around the contours of the white image. For all pixels contained within the bounding box, fitness increases for every black pixel that is vertically below a white pixel in its column. Fitness is then divided by total pixels within the bounding box to account for different sized objects, returning the percentage of overhanging mass in the image.

In determining the evolutionary algorithm to use, it was worthwhile to consider the differences between genetic algorithms and random mutation hill climbers. Both genetic algorithms and random mutation hill climbers both create new generations based on the best fit of previous generations, but the way in which they do so is meaningfully different. With genetic algorithms, a child can be created in two different ways:

1. A child can be based off of one mutated parent, in which a combination of the following actions occurs:



Figure 3: Evaluating "archness": Fitness is determined by first, thresholding the image into black and white, and second, drawing a bounding box around the object and calculating the percentage of overhanging mass.

random instructions in the parent's genotype are deleted, new instructions are randomly added, and existing instructions are mutated into different instructions.

2. A child can be based off of *crossover* between two parents: each parent provides a segment of its genotype, and the two segments are combined to make the genotype of the new child

Random mutation hill climbers, in comparison, create children only though method 1. Additionally, the type of hill climber used in this research, called a 1 + N hill climber, differs in another way: each new generation, of size 1 + N, consists of 1 parent and N children all created from that same parent. In other words, all but the one best of each generational population are culled before choosing a parent to use as the basis for the new generation.

Although genetic algorithms may seem more versatile at first glance – and they are, in some ways – they are not, in fact, more efficient than 1 + N hill climbers in many scenarios. In practice, genetic algorithms are only more useful when solutions to the problem at hand are known to be formed from *building blocks* [7] that can be joined together, through crossover, to form higher-order solutions. With the sort of problems that EvoFab is currently trying to solve, there is little to suggest that building blocks exist. Thus, for our purposes, a 1 + N hill climber is satisfactory.

#### 3.1 Results

Over sixteen generations of running EvoFab on a 1 + 4 random mutation hill climber, best fitness per generation increased only three times, as seen in Figure 4. However, average fitness makes quite large jumps between low fitness (around 0.17) to high fitness (around 0.26).

In our runs of EvoFab, we have also found the evolutionary algorithm to consistently exploit the fitness function in an undesirable fashion. The fitness function, as it currently stands, captures a three-dimensional image but treats it as if it were two-dimensional. Because of this, pixels that occur at higher elevations in the captured image are treated as if



Figure 4: While best fitness per generation rarely increases, average fitness jumps quite drastically for better and for worse. One possible cause is too many degrees of freedom.

they are higher along the Z-axis; in reality, they may only be further along the negative Y-axis, which the fitness function treats as nonexistant. Thus, archness is frequently measured to be higher in objects that have an "arm" extending forward on the Y-axis, as seen in the second and third images of Figure 3.

Another undesirable exploitation that occurs when printing is that, if the material is not extruded quickly enough, it will not have time to stick to the platform, causing the print carriage to drag around the thread of material instead. This has led to a certain degree of unpredictability in how a given set of instructions will translate into a printed product. In some cases, this can even cause material that has already been deposited onto the print platform to be dragged along with the hanging thread. This is an example of a mostlynegative exploitation that would go unaccounted for in a simulation-only evolutionary environment. Because of the dramatic effects on overall fitness that this can lead to, we believe that 3-D printers in particular, such as Fab@Home, lend themselves to a method of embodied evolution.

One exploitation that may still turn out to have interesting effects is that, when the platform moves upward along the Z-axis until it is very close to the tip of the syringe, sometimes the syringe tip moves across the object and "picks up" material, moving it to form a new design. This is an example of a potentially positive exploitation that would be nonexistent in a simulation-only evolutionary environment.

### 4. CHALLENGES IN EVOFAB

As seen by the types of evolutions in Figure fig:evaluation, the horizon line of the stage greatly affects measured fitness. When an image is evaluated, pixels higher on the Y-axis are assumed to higher in elevation. However, when the camera captures an image while looking down upon a printed object – thus causing the horizon to shift upward – this measurement becomes skewed. One possible cause for this is that there are too many degrees of freedom in the movement of the carriage. An arch is well suited to be represented in two dimensions; thickness is no concern. Thus, the carriage's extra variability afforded by movement on the Y-axis may very well be the cause of these jumps. Future tests that restrict movement along the Y-axis – so that only movements up, down, left, and right are allowed – will be able to test this hypothesis.

Another possible way in which this can potentially be fixed in future iterations of EvoFab could involve changing the weighting of the fitness function. Currently, an equal weight is applied to all black pixels found beneath white pixels. If, however, we instead weight more heavily black pixels found at higher elevations in the two-dimensional representation, we speculate that this would offset some of the erroneous fitness encountered below the horizon line. We do not want to completely discount the data below the horizon line, because it could still potentially lead to true higher fitness in non-exploitive ways.

With respect to the dragging of threads that have not yet been deposited onto the print platform: while it may seem like increasing extrusion rates would solve this problem, it is not so simple in practice. The problem does not occur constantly, and it is difficult to predict when it will. It is likely due to an uneven dispersion of pressure throughout the syringe, causing some areas to dispense more quickly than others. There is no easy solution, and it presents a problem in that it will always lead to some unreliability in reproducibility of prints. The best way to control for this is careful preparedness in how material is inserted into the syringe. Avoiding air bubbles and using the same method of insertion every time will help control for these kinds of problems. Again, this is another example of a physical-representation that could not be easily represented in simulation.

Regarding the phenomenon of the syringe tip picking up and moving about material, while no positive results have yet been seen from this, it is likely that this exploitation could potentially lead to objects that would not be able to be formed in simulation alone. One could imagine the syringe tip dragging a a dangling thread of material across two pillars, forming a makeshift arch. Hard to anticipate and almost impossible to control manually, this is exactly the kind of thing that makes software simulation so much harder to use effectively.

#### 5. CONCLUSION

We have constructed a system of Evolutionary Fabrication, the first closed-loop cycle of evolving physical objects. We have shown that the system can produce objects of increasing fitness over time and that embodied evolution can uniquely exploit physical realities of various parts of the fabrication process. We have additionally specified areas that currently hamper the system.

Future work will consist of bettering the reliability of Evo-Fab 0.2 through improvements to the accuracy of its fitness function and restrictions of movement along the Y-axis. Additionally, new types of objects will be presented for evolution to test the versatility of the system, along with new fitness functions to tackle those objects.

In *Star Trek: The Next Generation*, machines called replicators are capable of creating any object asked of them: food, toys, clothing, and spare parts for spacecraft repairs are all available at the push of a button. While replicators won't be invented any time soon, this vision of the future is not necessarily impossible. As technology advances, so, too, will implementations of evolutionary fabrication. We envision a future in which a person can ask a machine to create virtually anything for him, and evolutionary fabrication may be the process that gets us there.

#### References

- S. H. Al-Sakran, J. R. Koza, and L. W. Jones. Automated re-invention of a previously patented optical lens system using genetic programming. In M. Keijzer, A. Tettamanzi, P. Collet, J. I. van Hemert, and M. Tomassini, editors, *Proceedings of the 8th European Conference on Genetic Programming*, volume 3447 of *Lecture Notes in Computer Science*, pages 25–37, Lausanne, Switzerland, 2005. Springer.
- [2] N. Jacobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In *Proceedings of the Third European Conference on Advances in Artificial Life*, pages 704–720, London, UK, UK, 1995. Springer-Verlag.
- [3] L. E. Kavraki, J.-C. Latombe, and R. H. Wilson. On the complexity of assembly partitioning. *Information Processing Letters*, 48(5):229–235, 1993.
- [4] J. I. Lipton, D. Cohen, M. Heinz, M. Lobovsky, W. Parad, G. Bernstein, T. Li, J. Quartiere, K. Washington, A.-A. Umaru, R. Masanoff, J. Granstein, J. Whitney, and H. Lipson. Fab@home model 2: Towards ubiquitous personal fabrication devices. In Solid Freeform Fabrication Symposium, pages 70–81, 2009.
- [5] J. D. Lohn, G. S. Hornby, and D. S. Linden. An Evolved Antenna for Deployment on NASA's Space Technology 5 Mission. In U.-M. O'Reilly, R. L. Riolo, T. Yu, and B. Worzel, editors, *Genetic Programming Theory and Practice II.* Kluwer, 2005.

- [6] E. Malone and H. Lipson. Fab@home: The personal desktop fabricator kit. *Rapid Prototyping Journal*, 13(4):245–255, 2007.
- [7] M. Mitchell, S. Forrest, and J. H. Holland. The royal road for genetic algorithms: Fitness landscapes and ga performance. In *Proceedings of the First European Conference on Artificial Life*, pages 245–254. MIT Press, 1991.
- [8] J. Rieffel. Evolutionary Fabrication: the co-evolution of form and formation. PhD thesis, Brandeis University, 2006.
- [9] D. Sayles and J. Rieffel. Evofab: A fully embodied evolutionary fabricator. In G. Tempesti, A. Tyrrell, and J. Miller, editors, Evolvable Systems: From Biology to Hardware, volume 6274 of Lecture Notes in Computer Science, pages 372–380. Springer Berlin / Heidelberg, 2010.
- [10] K. Sims. Evolving virtual creatures. In Proceedings of the 21st annual conference on Computer graphics and interactive techniques, pages 15–22. ACM Press, 1994.
- [11] A. Thompson. An evolved circuit, intrinsic in silicon, entwined with physics. In T. Higuchi, M. Iwata, and L. Weixin, editors, *Proc. 1st Int. Conf. on Evolvable Systems (ICES'96)*, volume 1259 of *LNCS*, pages 390–405. Springer-Verlag, 1997.
- [12] R. A. Watson, S. G. Ficici, and J. B. Pollack. Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In CONGRESS ON EVOLUTIONARY COMPUTATION, pages 335–342. IEEE Press, 1999.