Morphogenetic Vase Forms

Andy Lomas[1]

[1]Department of Computing, Goldsmiths, University of London, SE14 6NW, UK
a.lomas@gold.ac.uk

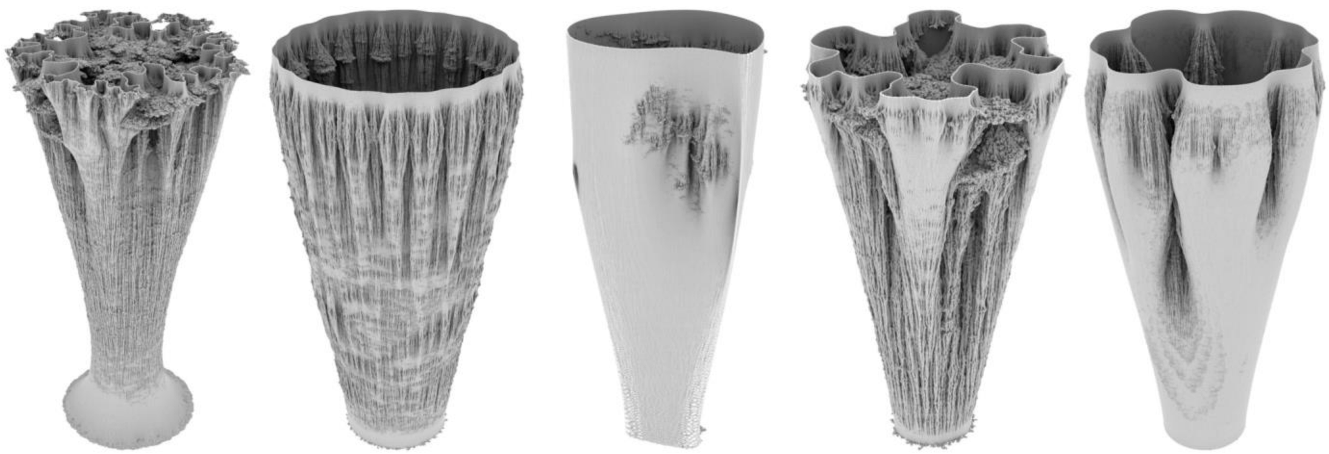

Figure 1: Examples of five Vase Forms

Abstract

This paper describes Vase Forms: a series of art works created using morphogenetic processes. A key motivation for these works was exploration of ways of working creatively with complex generative processes, such as morphogenetic systems, where the desire is to be able to influence the process in creative directions whilst achieving desired properties, such as fabricability using 3D printing, in a manner that retains rich emergence. The paper describes methods used in the creation of these works, including directly affecting morphogenetic processes using constraints and differential growth rates, combined with evolutionary search and machine learning algorithms to explore the space of possibilities afforded by the system. As well as describing the creation of Vase Forms, which have been successfully used to create sculptures, the paper looks at the closely related Mutant Vase Forms: an additional series of artworks created by accident when the system exploited bugs in the rules for the growth system resulting in unexpected but aesthetically interesting structures. These Mutant Vase Forms are not fabricable as physical sculptures with the originally intended methods, but now exist as virtual sculptures in stereoscopic installations.

Introduction

Morphogenesis and Generative Systems

Morphogenesis is a theme that has been explored by a number of artists. In 1951 Richard Hamilton curated an exhibition at the Institute of Contemporary Art (Massey, 1996) of work by a number of artists inspired by D'Arcy Thompson's 'On Growth and Form' (Thompson, 1917). In more recent years, growth has been a subject explored by computational artists including Yoichiro Kawaguchi's 'GROWTH Model' (Kawaguchi, 1983), William Latham's evolved forms (Todd and Latham, 1992) and Daniel Brown's series of digitally generated flowers (Brown, 2018). Interest by artists in morphogenesis, and D'Arcy Thompson specifically, has been sufficient for the University of Dundee to receive support from Art Fund to create a collection of artwork dedicated to this subject (University of Dundee, 2011).

One question raised by generative systems, such as those that use simulation of morphogenesis, is that of how are we to work creatively with them? In particular, how should we work with systems deliberately designed to encourage emergence: complex systems where results are intrinsically difficult to predict? There is a strong analogy with plant breeding, where we are working with a medium that is naturally rich. Through experimentation and experience we can develop insights into what is possible and how to influence plants to develop in ways that give desired properties. We need to discover the potentialities of the system we are working with, as well as the limits of its capabilities. Which features can be independently influenced, and which are co-dependent? Whether art, design or architecture, working in this manner involves changing our relationship with the computer. Traditional top-down design methods are no longer appropriate. We need to be open to a process of exploration. Participating in a search for rich interesting behavior: selecting and influencing rather than dictating results.

Generative systems are typically based on algorithmic processes that are parametrically controlled. Given a set of parameter values the process is run to create an output. Classic examples include Conway's Game of Life (Conway, 1970) and reaction diffusion equations (Turing, 1952). Generative systems have been used by a number of artists, from pioneering early work by Algorists such as Manfred Mohr (Mohr and Rosen, 2014), Frieder Nake (Nake, 2005), Ernest Edmonds (Franco, 2017) and Paul Brown (DAM, 2009a), to more recent work by artists such as William Latham (Todd and Latham, 1992), Yoichiro Kawaguchi (DAM, 2009b), Casey Reas (Reas, 2018), and Ryoji Ikeda (Ikeda, 2018).

The most interesting systems are generally those that create emergent results: genuinely unexpectedly rich behavior that cannot be simply predicted from the constituent parts. For these systems the relationship between the input parameters and the output is often complex and non-linear, with effects such as sensitive dependence on initial conditions. This can make working with such systems particularly challenging.

Creative Exploration of Parameter Space

One problem is that of how to work with systems with large numbers of parameters. With a small number, such as two or three parameters, the space of results can be relatively easily explored by simply varying individual parameter values and plotting the effects of different combinations. One common technique is to create charts where all the parameters are sampled independently at regularly spaced values and results are plotted to show the results. What scientists would call a phase space plot. This method of parameter exploration can be effective and was used by the author for earlier work such as for his 'Aggregation' (Lomas, 2005) and 'Flow' (Lomas, 2007) series (Figure 2).

As the number of parameters increase, the number of samples needed to explore different sets of combinations using this type of method increases rapidly. This problem is commonly called the 'Curse of Dimensionality' (Bellman, 1961) (Donoho, 2000), where the number of samples that need to be taken increases exponentially with the number of parameters. One approach is to simply limit the number of parameters, but this can be at the expense of overly restricting the range of behavior the system is capable of. If we are working with richly emergent systems these problems are often further compounded: a direct consequence of complexity is that parameters that drive the system often work in difficult to comprehend, unintuitive ways. Effects are typically non-linear, often with sudden tipping points as the system goes from one type of behavior to another. In particular, in many systems the most interesting emergent behavior occurs close to the boundary between regularity and chaos (Kauffman, 1996).

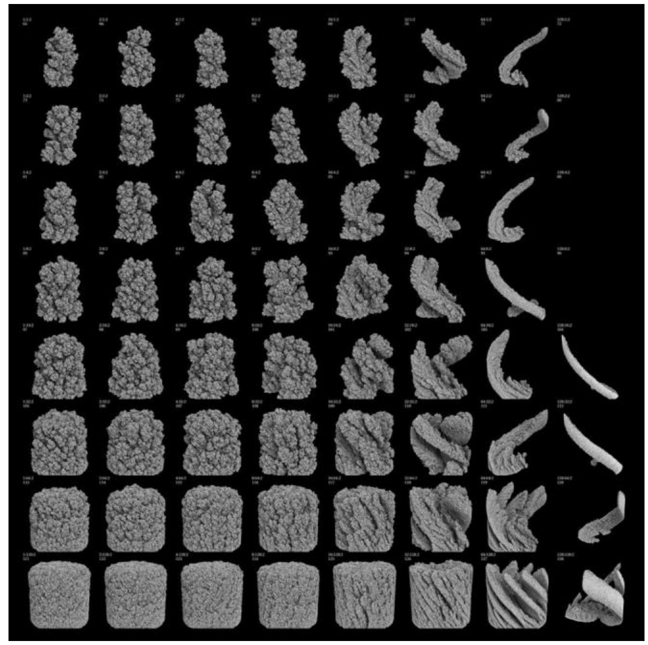

Figure 2: Phase Space plot from the Aggregation Series

This raises the idea of working with computers not merely as a medium to generate artwork but as active collaborators in the process of exploration and discovery. The use of tools to help the process of exploration can materially change both the creative process and the complexity of systems that we can effectively work with.

One analogy is that of Advanced Chess: a form of the game where each human player can actively use a computer to assist them to explore possible moves during games (Kasparov, 2017). Computer chess programs are generally very good at quickly detecting whether a proposed move will have catastrophic results. The effect of allowing a human player to test potential moves with a computer assistant is to make the game blunder-free. By removing the stress of making easily punished mistakes the human in the collaboration is freed to approach the game in a much more actively experimental way.

Another potentially rich analogy is with fly-by-wire systems in aircraft (Sutherland, 1968), that allow designs of aircraft to be created which are inherently unstable but can perform complex maneuvers beyond the performance envelope of conventional aircraft (Stein, 2003). These include designs that would be difficult, or even impossible, for a human pilot to directly control. Through the use of digital fly-by-wire technology, where the pilot uses their controls to indicate their intent but all the data is passed through a computer before being fed to actuators on the control surfaces, such aircraft can be flown safely.

A number of authors have proposed using evolutionary methods to allow artists and designers to explore systems with large numbers of parameters. Examples include Dawkins' Biomorphs (Dawkins, 1986) and Mutator (Todd and Latham, 1992). More recent examples, that use collaborative evolutionary interfaces for creation of images and forms, include Picbreeder (Secretan et al, 2008) and Endless Forms (Clune and Lipson, 2011). A number of systems that use interactive evolutionary computation for art and design are described in (Bentley, 1999) and (Takagi, 2001).

As demonstrated by natural processes, evolutionary methods can be effective even with extremely large numbers of parameters. One problem, though, can be that these methods generally lead to exploring a small number of paths within the space of available possibilities. The nature of these types of methods are to bias the search towards the most successful areas of the parameter space that have already been highly sampled. New samples are taken by mutation or cross-breeding of the gene codes from previous samples that are deemed fittest according to a specified fitness function. This means that previously highly sampled areas are likely to be even more highly sampled in the future as long as they contain 'fit' individuals. This is a good strategy for exploiting the best results that have been previously found, but is potentially a bad strategy for actively finding novel solutions which may be in areas of the landscape that have been very sparsely sampled.

In recent years there have been a number of studies into methods to keep diversity when working with evolutionary techniques, such as Novelty Search with Local Competition (Lehman and Stanley, 2011) and MAP-Elites (Mouret and Clune, 2015). These methods generally require the defining of a domain specific feature vector that represents the behavior of the system to enable a meaningful measure of the distance between individuals in behavior space to be calculated. The creation of such a function to represent behavior is often not easy (Lehman and Stanley, 2008).

A number of authors have proposed using machine learning techniques to assist human designers. In general these are for domain specific applications, such as for architectural space frame structures (Hanna, 2007), structurally valid furniture (Umetani et al. 2012) or aircraft designs (Oberhauser et al. 2015). In these systems, machine learning is typically used to learn about specific properties of the system. This is then used to provide interactive feedback for the user about whether an object designed by them is likely to have desired properties, such as being structurally feasible, without having to do computationally prohibitive tasks such as evaluation of structural strength using finite element analysis.

One thing that needs consideration is that creative work with generative systems often needs different phases of exploration, with the intent of the artist or designer changing over time. Initially they may be actively experimenting: trying to get a feel for the capabilities of the system they are working with. Once they have done some initial experiments they may want to continue to explore broadly, but with a general focus on regions that seem to have promise. When some particularly interesting results have been found they may wish to further refine them into presentable artefacts, or want to switch to actively looking for novel results that are significantly different to those they have found so far. These considerations mean that if a computer is being used to assist them explore the space of possibilities they may want it to work in different ways depending on their current intent.

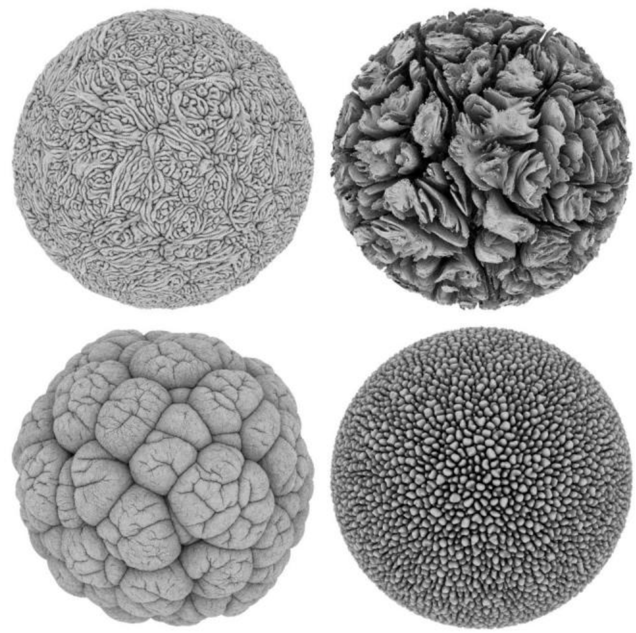

Figure 3: Examples of Cellular Forms

Motivation

The author is a practicing computational artist, whose work explores how complex organic forms can be created through digital simulation of morphogenetic processes. Inspired by Alan Turing's use of simple equations to create rich self-organizing patterns (Turing, 1952), the author's work focuses on creating simplified models of growth at the level of individual cells and exploring the emergent forms that can be created for these low-level rules (Lomas, 2014) (Figure 3). To explore the space of possibilities the author uses a hybrid system that combines several techniques, including evolutionary design search methods and lazy machine learning, to discover and fine-tune parameter combinations that appear to create particularly interesting results (Lomas, 2016).

The motivation behind the work described in this paper was how the cellular growth system used for works such as Cellular Forms (Lomas, 2014) could be modified to direct it towards the more specific goal of creating three dimensional structures that could be turned into physical sculptures. In particular, it was desired that the forms could be physically realized using computationally controlled additive fabrication techniques such as 3D printing using fused filament fabrication, or at larger scales using robots that deposit sequential layers of molten material. It was also desired that the sculptures created could potentially be suitable for use as the supports for tables, so should have flat bases and tops and potentially be strong enough to be load bearing.

One of the main restrictions with fused filament fabrication is that every layer to be printed needs to be supported by the previous layers. This means that overhang areas in the structures have to be within maximum size and angular limits or additional support structures have to be printed and removed after fabrication. The aim in this work was to explore how the structural needs for physical fabrication could be combined with aesthetic goals when working with a morphogenetic system to create forms that, though completely synthetic, exhibit complex detailed organic structures such as are typically found in natural forms.

Methods

Cellular Growth System

The generative system for this work was based on the model of growth using cellular division that the author created for his Cellular Forms work (Lomas, 2014). This system uses a simple particle spring model (Reeves, 1983), with each cell represented by particles with links to a number of other connected cells. The topology of connection between cells means that they form a sheet-like structure embedded in three-dimensional space, with the sheet free to fold into complex geometrical shapes. Interactions between connected cells are implemented using spring-like rules. The key elements of the system are:

- Rules for the generation of food, with rates for the spontaneous generation of food by all cells, and for the generation of food by simulated photosynthesis by firing light rays at cells and generating food in each cell proportional to the number of rays that hits it.
- Rules for whether cells are ‘greedy’ and directly accumulate the food they generate themselves or are ‘cooperative’ and share the food they generate with the other cells they are connected to.
- A threshold for how much food a cell needs to accumulate before it is selected for dividing into two cells.
- Rules for the direction that a cell splits in, which can be influenced by factors such as the curvature of the sheet of cells, the local direction of most tension in the surface, or in a randomly selected direction.
- Forces that try to maintain a constant rest-length separation between linked cells.
- Forces that tend to make the sheet of cells bulge outwards if that local area of the sheet is in compression (cells packed together closer than the rest-length).
- A relaxation rule that tends to move cells towards the average of their neighbors.

The system aims to be sufficiently simple to enable the simulation of morphogenesis to be implemented using massively parallel processing on consumer level graphics processing units, with simulations that can scale to tens of millions of cells and tens of thousands of simulation steps using conventional PC hardware. The code for the simulation engine is implemented in C++ and CUDA. For more details of the algorithms and implementation see (Lomas, 2014).

Additional Constraints and Influences

For this work the cellular division system previously described was augmented with a number of additional constraints and influences to steer the growth system to create structures with desired properties.

As previously described, the aim was to create forms that could be fabricated as sculptures that could also potentially be used as supports for tables. With this in mind, a decision was made to create forms with the topology of a tubular structure with open edges at the top and bottom. The initial configuration of cells was a simple cylinder (Figure 4).

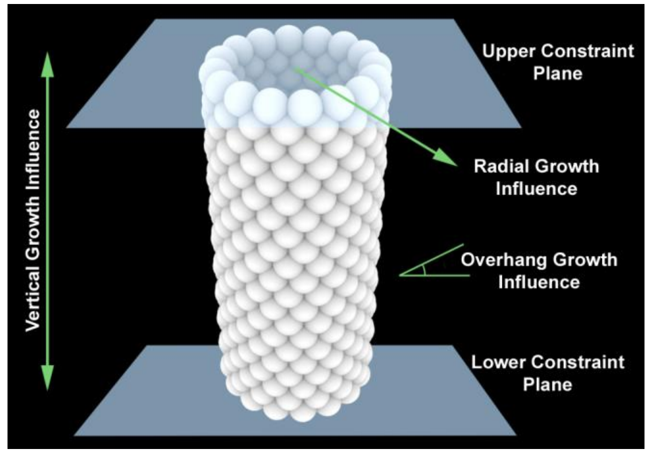

Figure 4: Initial configuration, constraints and influences

To make the top and bottom of the structure stay flat, so that the form should stand on the ground and to potentially support a table top, the cells along the open edges were constrained to horizontal planes so that these cells were only allowed to move in two dimensions within those planes. The constraining planes were maintained at a constant separation distance from each other. When cells divide, a test was run to check whether the resulting cells were in one of the open edges of the sheet. If so, they were added to the sets of constrained cells.

Since the horizontal constraining planes were kept at a constant separation distance, as the cells in the sheet grow there could be a tendency to over crowd the space between the planes if the rest-length between cells was kept constant. To prevent this from happening a rule was added to the system which adaptively reduced the rest-length between cells by a constant factor each time step. The value of this factor was one of the parameters used to drive the system.

Finally, a number of differential growth rates were added to the system which affect the rate of cell division in different parts of the structure. An analogy can be made with controlling the growth of plants by the selective application of nutrient in certain areas and a growth retardant in others. Differential growth rates were implemented by modifying the rate at which food is generated and accumulated in cells, hence controlling the rate at which they divide. Three different factors were allowed to affect growth rates:

Vertical Growth Influence. This allowed the position of the cell along the vertical axis to affect the growth rate, with cells closer to the top of the structure growing faster than cells lower down. The aim was to encourage the formation of generally vase-like forms that are wider at the top than the bottom.

Radial Growth Influence. To stop structures growing too large horizontally, the radial distance from the central vertical axis was also allowed to affect growth rates, retarding the growth of cells further away from this axis.

Overhang Growth Influence. As previously described, one of the main limitations of using fused filament fabrication is that structures can only have limited overhangs if they are to be made without the generation of additional support structures. To try to naturally encourage the growth of forms within overhang limits the local angle between the sheet of the cells and the vertical axis was also allowed to affect the growth rate, reducing the rate of growth if it exceeded a threshold angle.

This set of influences using differential growth rates was selected to work with the requirements of this specific system, but this approach should be suitable to be generalized to give creative influence over a number of other similarly morphogentically based generative systems.

Parameter Selection

With the addition of these constraints and influences the simulation system for generating Vase Forms had 29 parameters, each of which could be set independently. As described in the introduction, having a system with this number of parameters raises challenges of how to explore the space of possibilities of the system in order to find parameter combinations that produce desired results.

In response to these issues, the author has developed a program called Species Explorer (Figure 5) to assist the process of generating parameter values to be used with generative systems. The initial requirement for such a system came from the number of parameters that the author found he needed when he was developing the simulation engine for his Cellular Forms work (Lomas, 2014), but is designed to work in general with systems driven by a fixed number of parameters, and provides a framework for various methods to be used to assist in exploring the landscape of possibilities.

The software provides an interface for the user to specify the programs that need to be run to generate each individual. Once a set of parameter values has been chosen the system writes out a ‘creation script’ (Linux shell script, Windows batch file or Python script) that can be executed on the computer to run the generative system with the specified values. Once an individual has been generated the user can then use the interface to rate and categorize the results.

The software allows the user to select from a variety of ‘creation methods’ each of which use different techniques to generate sets of new parameter values to try. Examples of the creation methods the user can select from include:

- Simple random selection of parameter values from uniform distributions within a specified range.
- Evolutionary search methods using mutation and cross-breeding between the parameters used for previous individuals.
- A ‘fitness landscape’ method where parameter values are selected using lazy machine learning to estimate how the user would rate and categorize individuals at new coordinates in parameter space. The system implements two different options for lazy machine learning: nearest k-neighbors and interpolation using radial basis functions.

Using the interface, different creation methods can be used for each generation and fitness functions (such as for use with evolutionary search or lazy machine learning) can be customized using a simple Python based expression syntax. One feature of these expressions is that they can include the distance in parameter space to the nearest previous sample that has already been taken, which allows a simple implementation of methods to maintain diversity in the genotype space.

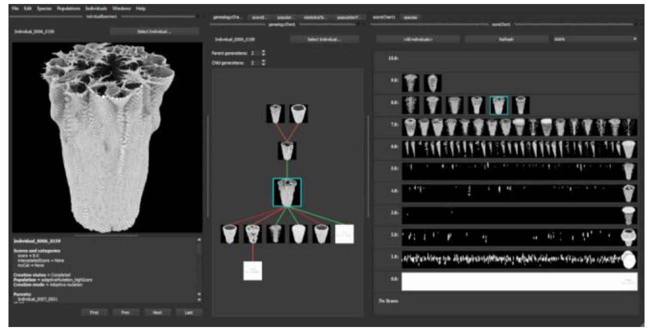

Figure 5: Species Explorer user interface

The use of a variety of creation methods provides the flexibility to allow the user to explore the space of possibilities in different ways depending on their intent (such as focused refinement based on some previous samples, or an active exploration for potentially novel results). The software also provides a framework for plugins to implement new creation methods, so the user can specify custom ways for how parameters for new individuals are chosen. For more technical detail about Species Explorer see (Lomas, 2016).

Results

Mutant Vase Forms

The initial results from the system were genuinely unexpected: instead of creating structures that were likely to be fabricable using 3D printing, the system would often create forms with finely detailed approximately horizontally-oriented branching structures (Figure 6). This was the exact opposite of the type of structure the author was hoping to create by differentially adjusting growth rates to slow down cell division in regions with high overhang angles.

Analyzing the results led to the realization that these unexpected results were due to bugs in how the system had

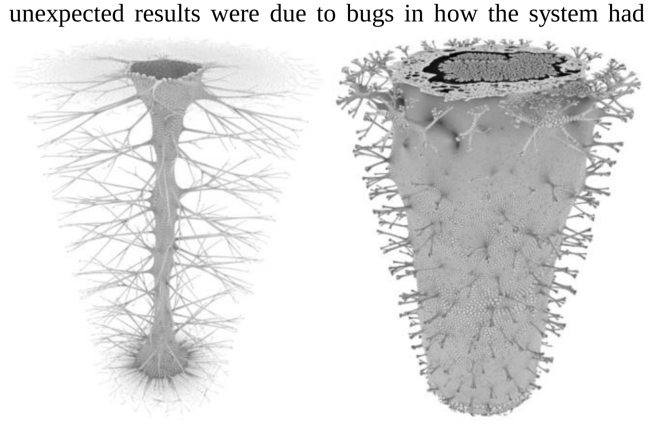

Figure 6: Mutant Vase Forms

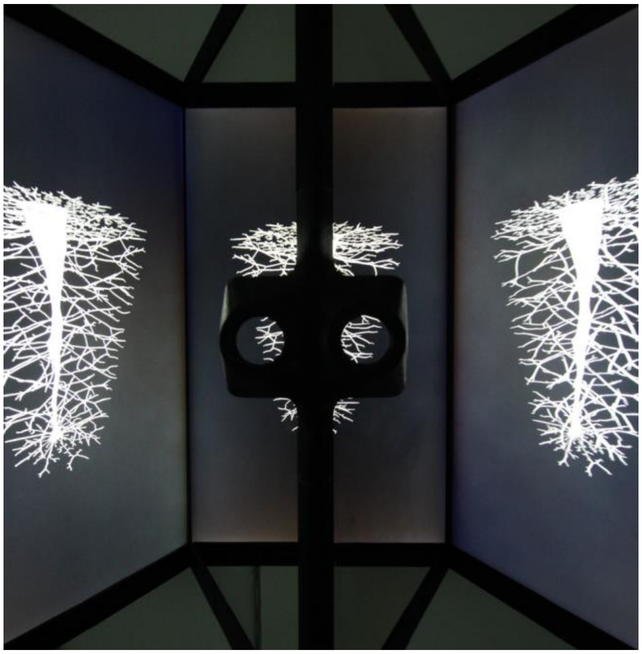

Figure 7: Mutant Vase Forms Stereoscopic Installation

been implemented:

- As previously described, during the growth simulation the rest-length between cells was reduced by a constant factor each time step. However, in the code for generating food by light rays hitting cells, the spheres used to represent the cells were being left at their initial size instead of using the modified rest-length to adjust their radius. The result was that the cells were effectively a lot larger than they should have been, meaning that cells that stuck out further than their neighbors were struck by all the light rays and became the places that all the food from photosynthesis was generated.
- The code for differentially adjusting growth rates was initially implemented by affecting the size of time steps, which was also affecting the physics simulations governing how the cells move. This meant that areas of low growth were also areas where structures became 'frozen', so if a region developed an overhang this feature would become geometrically fixed position-wise in space.

The effect of these two bugs was to accidentally create a recipe for generating horizontally oriented structures with growth focused at the tips. Though this wasn't what was originally intended, the author considered the results to be surprisingly aesthetic, particularly when they create structures with multiple filigree branches. The results can also be seen as genuinely emergent: they are the consequence of bugs in a system which has sufficient complexity so that changes to the rules can lead to rich unexpected consequences.

Though the resulting forms were not suitable for 3D printing using fused filament fabrication, the author considered that the results were sufficiently interesting to make them into a series of artworks in their own right: the 'Mutant Vase Forms'. These have been exhibited using stereoscopic installations so that they can be experienced as animated three-dimensional forms (Figure 7) (Lomas, 2017).

Vase Forms

After investigating the reasons behind the initial unexpected results, the simulation code was modified to fix the bugs previously described. This resulted in the generation of forms with more expected properties: generally vase like structures, broader nearer the top, narrow at the base, and with flat regions at both the top and the base (Figure 1). Though the forms often have overhangs, these are typically within angular limits or are of a sufficiently small size to allow 3D printing using fused filament fabrication methods without the need for additional support structures.

The generated Vase Forms appear to exhibit an interesting range of morphologies, with structures reminiscent of coral and plant-like forms. Structures often have complex ridges and folds, which as well as being aesthetic have the potential to have useful structural performances. There are also often structures that have surprised the author, such as the spontaneous generation of canopy-like structures at certain height ranges, which are probably the results of the differential growth rates. Many forms exhibit a variety of different surface patterns, from regions where the surfaces are relatively smooth, to other sections where the sheet of cells is folded into complex structures.

The author has printed a number of Vase Forms using an Ultimaker 2+ 3D printer, with conventional PLA filament. For exhibition the works have been presented as a combination of 40cm high final sculptures (which needed to be printed in two parts due to the build volume restrictions of the Ultimaker 2+) together with series of smaller 20cm high maquettes that show a number of different stages of development of each form from its initial configuration of a small number of cells in a cylindrical shape to the final structure with several million cells (Figure 8). These 'developmental series' can be seen as echoing the models illustrating embryo development that are commonly seen in natural history museums. Typical final forms used for fabrication have between 5 million and 20 million cells.

As well as the sculptures 3D printed using fused filament fabrication, the author has been able to fabricate a larger 60cm high form in polyamide using selective laser sintering. This is a process that requires less constraints on the shapes that can be fabricated due to overhang limits or the need for support structures.

The author has also used computer rendered image files from data at different timesteps during the growth simulation to create animations showing the process of forms growing by cell division (Lomas, 2018). In a number of exhibitions these animations have been shown together with the physical sculptures (Figure 9), giving another view into the story of how simulation of morphogenesis was used to create the forms.

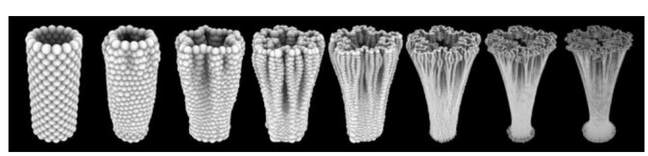

Figure 8: Vase Forms Developmental Series

Discussion

The creation of Vase Forms, and the accidental creation of the closely related Mutant Vase Forms, has been an exploration in trying to take an unruly, but potentially interesting, system of generating structure using morphogenesis and steer it in directions with both aesthetic and functional goals. The combination of using constraints and specific modifications to the growth rules, together with the use of evolutionary and machine learning methods to discover parameter combinations that give desired results, appears well suited to working with morphogenetic systems.

The work can be seen as a case study of engaging with rich emergent systems, emulating the way that nature works through evolution and natural selection but with a design intent. The use of influences that affect growth rates in different parts of the structure were designed for the specific needs of the Vase Forms, but should be generalizable to give a degree of creative influence over other similar morphogenetic systems.

The discovery of parameter combinations to create the final exhibited forms used a hybrid set of methods including interactive evolutionary computing and machine learning. One thing that needs to be considered for creative tools is the different needs of artists or designers that are committed to an extended process of exploration, but whose intention changes during the course of the development of a work, and more casual users, such as gallery visitors or visitors to a website, who are probably only going to engage with a system for a limited amount of time and for whom a simple interface that offers a single mode of interaction is probably more suited. For an artist or designer developing their own work, issues of user fatigue and perceived loss of control can be important, but we can also assume an extended commitment over time. Having a range of customizable tools that allow the user to direct the exploration with different intents can be important.

Conclusion

The aim with this work was to create sculptural forms that could be fabricated without the need for extensive support structures, while avoiding overly constraining the system and losing the potential for rich emergence. This appears to have been successful, creating vase-like structures that have a surprisingly natural appearance even though they have been completely synthetically generated and fabricated. The project can be seen as having genuinely emergent results, particularly with the Mutant Vase Forms where bugs in the algorithm used to generate structures resulted in completely unexpected, but aesthetically interesting, consequences.

The Vase Forms have been well received, including an invitation to feature in a special exhibit at the Victoria and Albert Museum for the 2018 London Design Festival (London Design Festival, 2018). The work has also been shown together their closely related siblings, the Mutant Vase Forms, in a number of exhibitions, including ‘bubble, bulge, bleb’ (LifeSpace, 2017), an exhibition celebrating the centenary of

Figure 9: Vase forms exhibit for the London Design Festival at the V&A Museum, London

the publication of D’Arcy Thompson’s ‘On Growth and Form’.

With the advent of new techniques, such as multi-material fabrication and the control of structure at microscopic levels, morphogenetic systems that can produce rich continually varying complex patterns and forms have the potential to contribute novel solutions beyond those that can be created by conventional assembly out of discrete parts. The question of how to appropriately steer morphogenetic systems, allowing humans to creatively engage with them, whilst keeping rich emergent behavior deserves further study. Working with such systems using a combination of ‘hard constraints’ and ‘soft influences’, such as differentially influencing growth rates, appears to be fruitful.

Acknowledgments

I would like to acknowledge the help of Alisa Andrasek and Daghan Cam at UCL The Bartlett School of Architecture, with whom discussions about creation of structures suitable to act as table supports was an initial inspiration for this work. I would also like to acknowledge the support from Materialise in fabricating a Vase Form using selective laser sintering.

References

Bellman, R. (1961). *Adaptive Control Processes: A Guided Tour*. Princeton: Princeton University Press.

Bentley, P. J. (1999). *Evolutionary Design by Computers*. San Francisco: Morgan Kaufmann.

Brown, D. (2018). *Daniel Brown’s*. Available at http://danielbrowns.com/

Conway, J. H. (1970). The game of life. *Scientific American* 223: 4.

Clune, J., and Lipson, H. (2011). Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In *ECAL 2011*, pages 141-148.

DAM, (2009a). *Digital Art Museum: Paul Brown*. Available online: http://dam.org/artists/phase-one/paul-brown

DAM, (2009b). *Digital Art Museum: Yoichiro Kawaguchi*. Available online: http://dam.org/artists/phase-one/yoichiro-kawaguchi

Dawkins, R. (1986). *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design*. New York: WW Norton & Company.

Donoho, D. L. (2000). High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality. Paper presented at *AMS Math Challenges Lecture*, Los Angeles, CA, USA, August 6–11.

Franco, F. (2017). *Generative Systems Art: The Work of Ernest Edmonds.* Abingdon: Routledge.

Hanna, S. (2007). Inductive machine learning of optimal modular structures: Estimating solutions using support vector machines. *AI EDAM 21*: 351–66.

Ikeda, R. (2018). *Ryoji Ikeda*. Available online: http://www.ryojiikeda.com/

Kasparov, G. (2017). *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*. London: John Murray.

Kauffman, S. (1996). *At Home in the Universe: The Search for Laws of Self-Organization and Complexity.* London: Penguin.

Kawaguchi, Y. (1983). *Yoichiro Kawagucho: Growth*. Available online: http://coppergiloth.net/siggraph/art.php?id=71

Lehman, J., and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE 2008*, pages 329-336.

Lehman, J., and Stanley, K. O. (2011). Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211-218. ACM.

LifeSpace, (2017). *bubble, bulge, bleb*. Available online: http://lifespace.dundee.ac.uk/exhibition/bubble-bulge-bleb

Lomas, A. (2005). Aggregation: Complexity out of Simplicity. In ACM *SIGGRAPH 2005* Sketches. New York: ACM.

Lomas, A. (2007). *Flow*. Available online: http://www.andylomas.com/flow.html

Lomas, A. (2014). Cellular Forms: An Artistic Exploration of Morphogenesis. Paper presented at *AISB-50*, London, UK.

Lomas, A. (2016). Species Explorer: An interface for artistic exploration of multi-dimensional parameter spaces. In *EVA London 2016*, Electronic Workshops in Computing (eWiC). London: BCS.

Lomas, A. (2017). *Mutant Vase Forms: New Growth and Light Rays.* Available online: https://vimeo.com/321525568

Lomas, A. (2018). *Vase Forms.* Available online: https://vimeo.com/321534705

London Design Festival (2018). *Vase Forms.* Available online: https://www.londondesignfestival.com/event/vase-forms-0

Massey, A. (1996). *The Independent Group: Modernism and Mass Culture in Britain, 1949-59*, pages 42-45. Manchester University Press.

Mohr, M. and Rosen, M. (2014). *Der Algorithmus des Manfred Mohr: Texte 1963–1979*. Oakland: Spectator Books.

Mouret, J. B., and Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

Oberhauser, M., Sky S., Thomas G. and Kristina S. (2015). Computational Design Synthesis of Aircraft Configurations with Shape Grammars. In *Design Computing and Cognition’14*, pages 21–39. Springer, Cham.

Nake, F. (2005). Computer art: A personal recollection. In *Proceedings of the 5th Conference on Creativity & Cognition*, pages 54–62. New York: ACM.

Reas, C. (2018). *Home Page of Casey REAS*. Available online: http://reas.com/

Reeves, W. T. (1983). “Particle Systems: A Technique for Modelling a Class of Fuzzy Object”, *ACM Transactions on Graphics, Volume 2 Issue 2*, April 1983: 91-108.

Secretan, J., Beato, N., D Ambrosio, D. B., Rodriguez, A., Campbell, A., and Stanley, K. O. (2008). Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1759-1768. ACM.

Stein, G. (2003). Respect the unstable. *IEEE Control Systems 23*: 12–25.

Sutherland, J. P. (1968). *Fly-by-Wire Flight Control Systems*. Dayton: Air Force Flight Dynamics Lab Wright-Patterson AFB.

Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation. In *Proceedings of the IEEE*, 89(9): 1275-1296.

Thompson, D. W. (1917). *On Growth and Form*. Cambridge University Press.

Todd, S. and Latham, W. (1992). *Evolutionary Art and Computers*. London: Academic Press.

Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 237: 37–72.

Umetani, N., Igarashi, T. and Mitra, N. J. (2012). Guided exploration of physically valid shapes for furniture design. *ACM Transactions on Graphics 31*: 86–1.

University of Dundee, (2011). *D’Arcy Thompson Zoology Museum Art Collection*. Available online: https://www.dundee.ac.uk/museum/collections/zoology/art