

Article

On Hybrid Creativity

Andy Lomas 

Department of Computing, Goldsmiths College, University of London, London SE14 6NW, UK;
andylomas@gmail.com

Received: 2 May 2018; Accepted: 5 July 2018; Published: 9 July 2018



Abstract: This article reviews the development of the author’s computational art practice, where the computer is used both as a device that provides the medium for generation of art (‘computer as art’) as well as acting actively as an assistant in the process of creating art (‘computer as artist’s assistant’), helping explore the space of possibilities afforded by generative systems. Drawing analogies with Kasparov’s Advanced Chess and the deliberate development of unstable aircraft using fly-by-wire technology, the article argues for a collaborative relationship with the computer that can free the artist to more fearlessly engage with the challenges of working with emergent systems that exhibit complex unpredictable behavior. The article also describes ‘Species Explorer’, the system the author has created in response to these challenges to assist exploration of the possibilities afforded by parametrically driven generative systems. This system provides a framework to allow the user to use a number of different techniques to explore new parameter combinations, including genetic algorithms, and machine learning methods. As the system learns the artist’s preferences the relationship with the computer can be considered to change from one of assistance to collaboration.

Keywords: art; computer; evolutionary design; machine learning; computationally assisted design

1. Introduction

How are we to work creatively with generative systems that computationally create results? 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, this 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 interesting behavior: selecting and influencing rather than dictating results.

Generative systems are typically based on algorithmic processes that are controlled by a number of parameters. 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 (Digital Art Museum 2009a) to more recent work by artists such as William Latham (Todd and Latham 1992), Yoichiro Kawaguchi (Digital Art Museum 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 generally complex and non-linear, with effects such as sensitive dependence on initial conditions. This makes working with such systems particularly challenging: both fascinating and potentially frustrating.

With a small number of dimensions, such as up to 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, or in the animation and visual effects industry is commonly called a wedge sheet. 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 1 shows one of the charts the author created when working on his 'Aggregation' series, exploring the effect that two different parameters had on the forms created. In this example the author was taking 8 samples in each dimension. With two parameters only 64 samples were needed to complete this chart. However, as the number of parameters goes up the number of samples needed to explore different sets of combinations using this method increases rapidly. Three parameters would require 512 samples. Four parameters would need 4096 samples. If we had a system with 10 parameters then just over a billion samples would be needed. 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. Even if enough samples can be taken, how to visualize and understand the space becomes a significantly difficult problem and concepts such as finding the nearest neighbors to any point in the parameter space become increasingly meaningless (Marimont and Shapiro 1979).




Figure 1. Phase space plot from the Aggregation series (Lomas 2005).

One approach is to simply limit the number of parameters, but this can be at the expense of overly limiting the type of system that we are willing to work with. If we are working with richly emergent system these problems are often further compounded. A direct consequence of emergence

is that the parameters 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. Indeed, the most interesting results, such as those in the right hand columns of Figure 1 above, are often near regions of instability or at transitions between behaviors. In particular, in many systems the most interesting emergent behavior occurs close to the boundary between regularity and chaos (Kauffman 1996). The shape of this type of boundary can be extremely complex, a classic example being the ornately fractal shaped boundary between regularity and chaos in the Mandelbrot set (Douady and Hubbard 1984).

This raises the idea of working with the machine not merely as the medium but as an active collaborator in the process of exploration and discovery. Can computational methods be used to allow exploration of generative systems like these in ways that would not be otherwise possible? The computer becomes an active part of the process of discovery, not just as the medium used to create artefacts.

One analogy worth exploring is that of Advanced Chess: a form of the game introduced by Garry Kasparov where each human player can use a computer to assist them to explore possible moves (Kasparov 2017). In particular, 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). These 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.

How are we to express an artistic opinion while working with a complex space of possibilities? The space of possibilities may be rich, but it can be a challenging territory to explore. How should we explore the space of possibilities to find the most interesting results when dealing with more and more parameters? There is a danger of becoming overwhelmed by too many controls. After we have tried out some initial parameter values, what values to try next? This problem, of how to choose new parameter values based on the limited number of parameter combinations we've sampled so far, is one that a computational algorithm may be able to actively help with.

We can think about the possibility for different relationships with the computer. Can using them actively in the creative process allow more fearless engagement with difficult 'unruly' generative systems, exploring spaces with more parameters and complex inter-dependence between them than would otherwise be possible? To try to address these questions, the author has created his own system, called 'Species Explorer', that provides a range of different methods to help the user select parameter combinations to try when working with generative systems.

2. Working with a Complex Space of Possibility

Most of the author's recent work involves exploration of morphogenesis: creating forms generatively through simulation of growth. The aim is to create systems that have the potential for a rich variety of different types of three-dimensional structure and form. All these come emergently from low level rules such as forces between cells, how cells split and connect to their neighbors, and how food needed for growth is created and shared between cells. The work can be considered as explorations of artificial life: inspired by, rather than trying to copy, biology. Actively exploring whether different rules create familiar or deeply alien structures.

The aim is to create systems that have the potential for a rich range of possibilities. Typically, each system has a quite large number of parameters, any of which could influence the development of the forms in potentially interesting ways. The simplest systems that the author has created as part of his Cellular Forms series (Figure 2) has 12 parameters, with more recent variations, such as Hybrid Forms (Lomas 2015) and Vase Forms (Lomas 2017), having 30–40 parameters.

A big question is how to explore a landscape with this number of parameters with a creative intent. As was previously discussed, structured sampling taking a fixed number of samples independently in each dimension would result in an overwhelmingly large number of samples to deal with. Even if we could take the required number of samples, how are we to visualize or otherwise make sense of the results? On the other hand, if we have a smaller number of less structured samples it is very difficult to make sense of what the data means and decide on new parameter combinations to try.

This is an area where computational methods, such as genetic algorithms or machine learning techniques, may be useful. In effect: allow the computer to help us make the best use of the data that has been acquired so far to decide where to next sample.

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). A number of systems that use evolutionary selection for design are described in (Bentley 1999).




Figure 2. Examples of a range of morphologies from the author's Cellular Forms (Lomas 2014).

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 can be seen as a bad strategy for actively finding novel solutions which may be in areas of the landscape that have had very few samples so far.

Another issue worth considering is that for creative work there is often a need for different phases of exploration. Initially we may be actively experimenting: trying to get a feel for the capabilities of the medium we are working with. Once we have done some initial experiments we may want to continue to explore broadly, but with a general focus on regions that seem to have promise. Once we have found some particularly interesting results we may wish to further refine them into presentable artefacts for exhibition, or want to switch to actively looking for novel results that are significantly different to those we have found so far. In other words, the intent of a process of exploration changes over time. If a computer is assisting us we may want it to work in different ways depending what type of result we are currently searching for.

One analogy is with journeying into an unexplored landscape looking for the particularly interesting locations. When we first go into the landscape we may go to a few random places, exploring along a single path. Once we have got an idea of the different types of terrain, we may want to go back and explore some of the most interesting locations we have found in greater detail. If we get bored with the places we have seen so far, we may want to see if we can find new distinctive types of terrain different to any of those we have seen previously.

There are at least four different phases of exploration worth considering separately: Initial Exploration, Secondary Exploration, Refined Focus, and Looking for Novelty.

2.1. Initial Exploration

Initial exploration is where the user is trying to get an initial basic understanding of how a system works and what sort of result may be possible. At this stage they typically have little or no idea what places in space may be interesting, and are effectively just randomly trying out parameter values to test if the system works as expected. They start to get a flavor for what types of behavior may be achieved, and are interested in hints of what may be possible rather than detailed exploration.

2.2. Secondary Exploration

At this stage the user has done some initial exploration. They want to use the information gathered so far to help guide the search: to be steered into broad areas of the parameter space that appear to be potentially fruitful, and away from regions that have been found to yield invalid results or are otherwise undesirable. However, the search should still be a broad one, avoiding what would be considered in optimization as becoming trapped in a local maximum by over-refining too early and in the process missing potentially even more interesting results.

2.3. Refined Focus

When the user has found some results that appear particularly interesting they may want to focus on those for further refinement. This is effectively a focused exploration within a small range of parameter values, such as to create final artistic exhibitable artefacts.

2.4. Looking for Novelty

However, once the user has found and refined some particularly interesting results, they may want to switch to looking for novelty. Are there additional rich seams in the landscape that have not yet been discovered? Can the space be searched for results unlike those seen previously, which may in turn take the exploration in fruitful new directions?

At this stage it is sensible to want to make good use of all the data that has been acquired so far, but without being overly biased towards regions that have been explored in detail. The user is willing to get results that are less 'fit' if they hold the promise for something genuinely new.

3. Evolutionary Methods and Machine Learning

As previously discussed, evolutionary methods are a commonly suggested approach for creatively working with generative systems. For the different phases of exploration, the author considers that they are particularly effective when refining previously found promising solutions, are also useful for secondary exploration, but because of their strong bias towards areas that have already been highly sampled are generally very poor for looking for novelty. With evolutionary techniques the general approach to find novel solutions is to dramatically increase mutation rates, or to simply start with a completely new population with the hope that a fresh search may explore a new path and find novel solutions. However, this is at the expense of throwing away hard-earned data about what has already been discovered in exploring the landscape of possibilities.

In more recent years 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.

Machine learning is a potentially interesting technique when looking for novelty. In particular, it should be possible to use machine learning as a method to predict fitness at arbitrarily chosen new points in the parameter space based on all the data that has been collected so far, and conduct a search explicitly biased in favor of new samples that are a long distance in parameter space from previous samples.

4. Species Explorer

In response to these issues, the author has developed a program called Species Explorer to assist the process of generating parameter values to be used with generative systems (Figure 3). Developed out of necessity, the specific need 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). This program provides a framework for various methods to be used to assist exploring the landscape of possibilities.

The software is designed to be used on computers running Windows and Linux, but everything has been written in an operating system agnostic manner that should facilitate support for other operating systems. It is implemented in Python together with Qt, using the PySide Qt bindings (The Qt Company, Oslo, Norway). This has allowed rapid development and experimentation.

The software provides an interface for the user to rate and categorize the results from running a generative system with different parameter values. Various methods, including random sampling, evolutionary search techniques and machine learning, can be used to generate new parameter values to try out. 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. The user can then rate and categorize the results of these new samples, and the computer in turn suggests parameter values for new places in the landscape of possibilities to try.

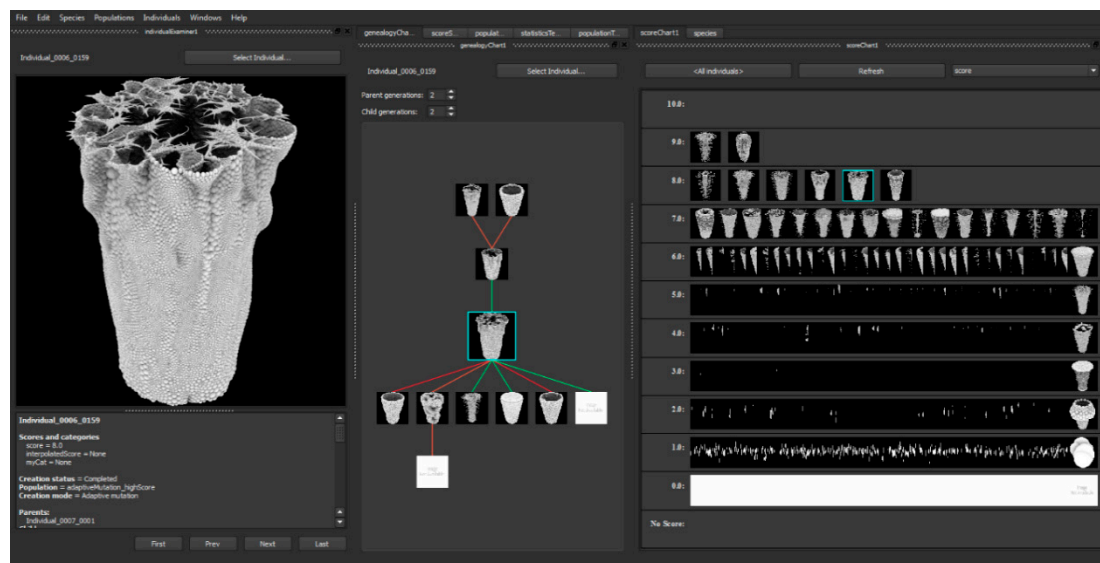


Figure 3. Species Explorer user interface.

The software allows the user to select from a number of different ‘creation methods’, each of which use different techniques for selection of new parameter values based on the data gathered so far. This 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 their own custom ways for how samples are chosen.

Currently the following creation methods are provided. For more technical detail see (Lomas 2016).

- Uniform random selection of parameter value
- User specified parameter values
- Wedge Test (parameter values with regular spacing)
- Standard Mutation
- Adaptive Mutation
- Cross-breeding
- Cross-breeding confined to neighbors
- Selection using lazy machine learning

These provide a flexible palette of different methods for generation of new parameter values. In the author’s experience using the system, he has found that different techniques are appropriate depending on his current intent. For the initial stages of exploration the author generally uses simple random parameter selection. Once he has evaluated some results, scoring them using rating values in a range from 0 to 10, he then uses evolutionary methods such as cross breeding or estimated scores using machine learning to do ‘secondary exploration’ as described above in Section 2.2. When he finds results that he considers particularly interesting and he wants to refine further, adaptive mutation (where mutation rates are varied depending on the distance to the closest neighbors in parameter space) and cross-breeding confined to the neighbors both work well.

However, when searching for novelty the author prefers machine learning techniques over evolutionary ones. These methods use lazy machine learning to estimate scores at new positions in the parameter space. New individuals are chosen based on these estimate values using a Monte Carlo method that estimates the score at a number of candidate points, and chooses one of the candidates

with a probability proportional to the estimated values. This means that parameter combinations which are expected to have high score values will be preferentially selected.

The user can also provide custom 'score expressions' to be used when selecting parents for evolutionary techniques or for the value to be estimated using machine learning. For instance a score expression can be used to raise a score to a power to bias the selection of parents even more towards those with higher score values, to combine the effects of different score values (such as if the user has rated results both on an overall 'score' and on a rating for how 'hairy' structures appear), or to deliberately bias the selection of new samples to less explored regions of parameter space by including the distance to the closest existing sample point in the score expression.

As well as providing numeric ratings, the user can group samples into categories, such as forms that look like 'brains', 'broccoli', or 'corals'. Using score expressions, the user can either restrict breeding to sets of parents in specific categories, or use the machine learning creation methods to find new parameter values that are predicted to have a high probability of being in a given category.

The author has used Species Explorer for all the work he has created in recent years, including his Cellular Forms and more recent related series (Figure 4). The key intent with the system is for the computer to act as an active assistant, helping guide users as they explore a system to discover its potential capabilities and making the best use of all the input the user has made. The user should be able to steer the search with a creative intent, refining particularly interesting results, with the computer assisting them in exploring the space for novel rich behavior.




Figure 4. Morphogenetic Creation exhibition at Watermans, 2016. All works in the exhibition were created using Species Explorer to discover parameter values for the forms presented.

This also raises the question of how the relationship between the user and the computer can change over time. The use of genetic algorithm and machine learning methods means that the system adapts through use, effectively learning about the user's preferences. The author believes that a consequence of this is that the relationship with the computer can also change over time, from one where the computer is very much in the role of purely an assistant where all creative choice is explicitly

controlled by the user, to one where the computer can be seen as more of a collaborator suggesting parameter combinations that have a high likelihood of producing interesting possibilities.

Currently all aesthetic judgements when rating or categorizing results are explicitly made by the user, but one potentially interesting direction for future enhancement would be to use the system as a platform to provide training data for machine learning to model how the user is rating and categorizing. This could allow a system of interactive training, with the user rating and categorizing some initial samples, the machine making predictions about score and categories for new samples, and the user correcting the predictions if they are wrong. Using such a system we could explore whether we can reach a point where the computer can successfully predict how the user will rate and categorize.

5. Conclusions

Working with generative systems gives the potential for rich possibilities, but also presents many challenges, particularly when trying to work with a creative intent. It is natural to look towards working collaboratively with a computer as an inherent part of the process, but what is the relationship that we want to develop with the computer? Is there a way that humans and computers can work together to their mutual strengths? In the author's experience the computer can become an active assistant in the process of discovery as well as being a medium to work with, enabling creative exploration with systems that the author previously found overwhelming (due to the systems having large numbers of parameters any of which could affect the results in difficult to predict but potentially interesting ways). The process changes to one that feels like a productive active collaboration, with the computer freeing the author from anxiety when creating new systems and adding parameters that he believes may have the potential to generate unexpectedly interesting results.

In particular, as an artist it can be important to have a relationship that gives at least the plausible illusion of being relevant and provides for interaction that feels rewarding, where decisions by the artist are steering the work in ways that match their intent. The nature of using a system that adaptively changes based on the user's input is that the relationship with the computer can change over time, from one where the computer acts purely as a technical assistant to one where the computer can be seen as a collaborator in the process of creation.

Funding: This research received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

References

- Bellman, Richard. 1961. *Adaptive Control Processes: A Guided Tour*. Princeton: Princeton University Press.
- Bentley, Peter J. 1999. *Evolutionary Design by Computers*. San Francisco: Morgan Kaufmann, ISBN 978-155860605X.
- Conway, John H. 1970. The game of life. *Scientific American* 223: 4.
- Digital Art Museum. 2009a. Paul Brown. Available online: <http://dam.org/artists/phase-one/paul-brown> (accessed on 20 June 2018).
- Digital Art Museum. 2009b. Yoichiro Kawaguchi. Available online: <http://dam.org/artists/phase-one/yoichiro-kawaguchi> (accessed on 20 June 2018).
- Dawkins, Richard. 1986. *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design*. New York: WW Norton & Company, ISBN 978-0141026169.
- Donoho, David 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; Available online: <https://pdfs.semanticscholar.org/63c6/8278418b69f60b4814fae8dd15b1b1854295.pdf> (accessed on 6 July 2018).
- Douady, Adrien, and John Hamal Hubbard. 1984. *Etude Dynamique des Polynômes Complexes*. Prépublications Mathématiques d'Orsay. Orsay: Université de Paris-Sud.
- Franco, Francesca. 2017. *Generative Systems Art: The Work of Ernest Edmonds*. Abingdon: Routledge, ISBN 978-1472436009.
- Hanna, Sean. 2007. Inductive machine learning of optimal modular structures: Estimating solutions using support vector machines. *AI EDAM* 21: 351–66. [CrossRef]

- Ikeda, Ryoji. 2018. Ryoji Ikeda. Available online: <http://www.ryojiikeda.com/> (accessed on 20 June 2018).
- Kasparov, Garry. 2017. *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*. London: John Murray, ISBN 978-1473653504.
- Kauffman, Stuart. 1996. *At Home in the Universe: The Search for Laws of Self-Organization and Complexity: The Search for Laws of Self-Organisation and Complexity*. London: Penguin, ISBN 978-0140174144.
- Lomas, Andy. 2005. Aggregation: Complexity out of Simplicity. In *ACM SIGGRAPH 2005 Sketches*. New York: ACM, Available online: http://www.andylomas.com/sketch_0087_final.pdf (accessed on 10 March 2018).
- Lomas, Andy. 2007. Flow. Available online: <http://www.andylomas.com/flow.html> (accessed on 10 March 2018).
- Lomas, Andy. 2014. Cellular Forms: An Artistic Exploration of Morphogenesis. Paper presented at AISB-50, London, UK, April 1–4; Available online: http://www.andylomas.com/extra/andylomas_paper_cellular_forms_aisb50.pdf (accessed on 10 March 2018).
- Lomas, Andy. 2015. Hybrid Forms. Available online: <http://www.andylomas.com/hybridForms.html> (accessed on 26 April 2018).
- Lomas, Andy. 2016. Species Explorer: An interface for artistic exploration of multi-dimensional parameter spaces. Paper presented at EVA London 2016 Conference, Electronic Workshops in Computing (eWiC), London, UK, July 12–14; London: BCS.
- Lomas, Andy. 2017. Vase Forms. Available online: <http://www.andylomas.com/vaseForms.html> (accessed on 26 April 2018).
- Marimont, R. B., and M. B. Shapiro. 1979. Nearest Neighbour Searches and the Curse of Dimensionality. *IMA Journal of Applied Mathematics* 24: 59–70. [CrossRef]
- Mohr, Manfred, and Margrit Rosen. 2014. *Der Algorithmus des Manfred Mohr: Texte 1963–1979*. Oakland: Spectator Books, ISBN 978-3944669168.
- Oberhauser, Matthias, Sky Sartorius, Thomas Gmeiner, and Kristina Shea. 2015. Computational Design Synthesis of Aircraft Configurations with Shape Grammars. In *Design Computing and Cognition'14*. Cham: Springer International Publishing, pp. 21–39.
- Nake, Frieder. 2005. Computer art: A personal recollection. Paper presented at the 5th Conference on Creativity & Cognition, London, UK, April 12–15.
- Reas, Casey. 2018. Home Page of Casey REAS. Available online: <http://caesuras.net/> (accessed on 20 June 2018).
- Stein, Gunter. 2003. Respect the unstable. *IEEE Control Systems* 23: 12–25. [CrossRef]
- Sutherland, Major J. P. 1968. *Fly-by-Wire Flight Control Systems*. Dayton: Air Force Flight Dynamics Lab Wright-Patterson AFB.
- Todd, Stephen, and William Latham. 1992. *Evolutionary Art and Computers*. London: Academic Press, ISBN 978-0124371859.
- Turing, Alan M. 1952. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 237: 37–72. [CrossRef]
- Umetani, Nobuyuki, Takeo Igarashi, and Nilroy J. Mitra. 2012. Guided exploration of physically valid shapes for furniture design. *ACM Transactions on Graphics* 31: 86–1. [CrossRef]



© 2018 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).