Morphogenic Shape Grammars for the Design of Engineering Structures

Simon Hickinbotham
School of Physics, Engineering & Technology, University of York
York, UK
simon.hickinbotham@york.ac.uk

Edgar Buchanan
School of Physics, Engineering & Technology, University of York
York, UK
edgar.buchanan@york.ac.uk

Peter Kilpatrick
School of Electronics, Electrical Engineering & Computer Science, Queen's University Belfast
Belfast, UK
P.Kilpatrick@qub.ac.uk

Mark Price
School of Mechanical & Aerospace Engineering, Queen's University Belfast
Belfast, UK
M.Price@qub.ac.uk

Andy M. Tyrrell
School of Physics, Engineering & Technology, University of York
York, UK
andy.tyrrell@york.ac.uk

***Abstract*—Bodies of multicellular organisms are laid out according to *morphogens:* chemical agents which establish a co-ordinate system in the early embryo and use this to decide where body parts should grow. This process offers a mechanism for the automation of the design of engineering constructs via evolutionary search in a similar manner to the way biological evolution has driven the diversity of body forms of life on Earth. There are many ways of encoding such body plans, but the main existing approaches have problems around managing the complexity and stability of the evolutionary search process, particularly when applied to practical engineering design problems. This contribution takes the notion of morphogen chemical gradients and uses it to develop a novel grammar for shape formation. The central idea is to organise and label the spatial sub-regions of a design before making decisions regarding the finished arrangement of structure. This makes it much simpler to explore compositions of the hierarchy of sub-assemblies in a design, and to represent design of shape in an evolvable manner. A worked example of such a morphogenic shape grammar is described, and used in a multi-objective evolutionary search for optimal bridge truss structures over four fitness objectives with two design constraints. The resulting Pareto front shows a wide variety of bridge designs, demonstrating the power of this approach to generate a diverse set of viable options to meet engineering design challenges.**

***Index Terms*—Evolutionary computation, Grammar, Generative AI, Structural Engineering**

I. INTRODUCTION

In biology, morphogenesis is the process by which the growth of an organism from a single cell to maturity is organised. The early stage of morphogenesis has three features [1]: 1) it grows colonies of undifferentiated cells; 2) establishes a co-ordinate system; and 3) labels regions of the growing colony such that they begin to develop into different organs. This process is key to the evolution of body forms of living creatures, which demonstrates a powerful ability to adapt to a vast range of ecological niches. For this reason, morphogenesis is of interest to the field of engineering design, where the problem of building physical objects to solve design problems has many parallels with the concept of filling an ecological niche. The morphogenesis process is akin to the systems architecting phase in design, where the overall configuration of the design solution is established.

Wolpert's pioneering work into mechanisms of morphogenesis [2], [3] led to ambitious research into using direct emulations of the process for design [4], [5]. However, since the number of cells in an organism can number many trillions, the dream of making this work both efficiently and at scale has yet to be realised.

Although there are difficulties in directly emulating morphogenesis for engineering design, alternative approaches also have shortcomings. Compositional Pattern Producing Networks (CPPNs) and related techniques [6], [7] couple evolutionary search with a neural network-based filling of space in response to inputs of the two- or three-dimensional co-ordinates, but it is difficult to couple these approaches with real engineering problems because the representation of shape via superimposition of mathematical functions across the space in which the design sits is too abstract, requiring a majority of contingent positive values in the various functions to generate solid regions of the space. A relatively common approach in the engineering literature is called the "Ground structure" method [8], where a fully-connected lattice of points is pruned in order to find optimal physical forms for loading problems. Evolution within such systems has also been implemented successfully [9]. The main drawback of this approach is that the lattice must be sufficiently coarse-grained to be computation-

---

This work was supported by EPSRC programme Grant Ref EP/V007335/1, "RIED: Re-Imagining Engineering Design", Website: https://riedesign.org/

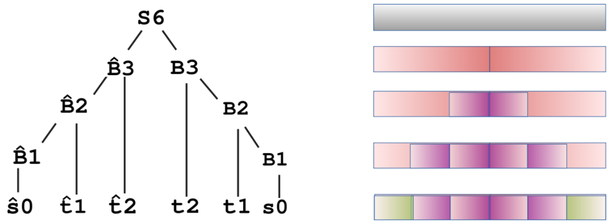

Fig. 1. Conceptual similarities between a production of a morphogen grammar (left) and a hypothetical morphogen concentration gradient (right). The direction of the gradient is indicated by the ‘ˆ’ attribute in the left panel, and by shading in the right panel. The production commences with input attribute 6, giving a start symbol **S6**, which corresponds to an empty space with a vertical morphogen gradient. Each iteration of the production of the grammar (shown in separate rows in the figure) corresponds to a labelling of subregions of the design space. The bottom row illustrates the correspondence between the shape grammar terminals and their position and orientation in the design space.

ally tractable, potentially missing better solutions that could be found within finer grained lattices. Spatial [10] and shape [11]–[13] grammars are primarily used to interpret existing structures, and the relatively few approaches to using shape grammars for design have problems regarding representation of form and linking this to the capacity to search over these representations. [14] use a shape grammar combined with an interactive evolutionary search algorithm to allow users to select through generations of evolved designs, with all the usual problems (user fatigue, slow runtimes, small sample size) associated with human-in-the-loop approaches.

This work is motivated by the observation that the organisation of the growth stages in animals follows a pre-coded “body plan”, that simultaneously increases the size of the body by addition of cells, establishes a co-ordinate system within the body via multiple gradients of chemical agents, and uses those gradients to differentiate the cells into different cell types. There are analogues in this process with the process of generating strings in formal languages. The hierarchy of morphogen agents is analogous to non-terminal symbols in a grammar: both play a role in organising the final structure without being part of that structure. The arrangement of terminal symbols in a sentence of a language organises meaning, which is analogous to the arrangement of the structures in the body of animals organising function.

Existing research into using evolutionary search processes for engineering design have the following issues to address: representation of the design on the genome; composable hierarchy of structure; and managing the complexity of the design space. The novel contributions in this paper offer a route to solving these problems. Firstly, a new class of grammar is developed that allows shape to be represented in linear strings of symbols, thus allowing them to be encoded on a genome. Secondly, since the shape representation is encoded in the grammar, a hierarchy of structure can be explicitly encoded. Thirdly, the complexity of the design problem is managed by exploiting this hierarchy, allowing modularisation of structures and re-casting the nature of the ground structure in an economical manner. Finally, the morphogen model allows fine-tuning of local structure by building and evolving further transformations based on the attributes of the terminal symbols in the grammar.

To demonstrate the application of this grammar to real-world engineering design problems, a new form of shape grammar called a “*Morphogenic Shape Grammar*” (hereafter referred to as a “*morphogen grammar*”) is derived for four different types of bridge truss structure (Warren, Warren-V, Pratt and K trusses). Bridges are linear structures with repeating 3D components that subtly vary, and there exists a range of bridge designs that utilise these. A multi-objective evolutionary search is used to drive the generation of a diverse range of bridge designs using morphogen grammars to encapsulate the form and manage variations in the modules that form the structure.

Supplementary Data: The source code and dataset for this study are available at: doi.org/10.15124/42cf8290-67dd-4a96-b8c2-c51735db4d2c.

II. MORPHOGEN GRAMMARS

Morphogen grammars exploit the ability of attribute grammars [15] to extend Type 2 context-free grammars to pass information around a structure as it is being generated, in a manner analogous to morphogens in a growing organism. This analogy is illustrated in figure 1, where productions are shown as a tree of symbols on the left, and as compositions of shaded gradients on the right, with attributes being passed down the production to indicate the type, magnitude and direction of the morphogen gradients. These grammars are defined by: the *co-ordinate system* κ; the set G of *structure graphs*; the set Σ of *terminal symbols*; the set N of *non-terminal symbols*; the set P of *production rules*; the set A of *attributes*; and S the *start symbol*.

In common with other grammars are the terminal and non-terminal symbols, the production rules and start symbol. The extension to an analogy of morphogenesis is achieved via *attributes* which manage a model of morphogen gradients, and *structure graphs* which are associated with each terminal symbol, and which are represented on the genome via a composite co-ordinate system.

The general strategy employed by morphogen grammars is to delay the generation of form until the final layout of the components of the mature structure is arrived at. This differentiates the approach from that of shape grammars [11], which utilise shape transformation rules at each level of the production. This new approach avoids some of the complications that arise in shape grammars whereby the transformations must accommodate the increase in complexity of the structure and then act upon the more complex structure to form the next stage in production.

Instead of a terminal symbol representing a distinct shape, the terminal initially only accounts for a sub-region of the space in which the structure resides – this subregion is hereafter referred to as a “cell” (although the relative size of the region relative to the complete structure is much bigger

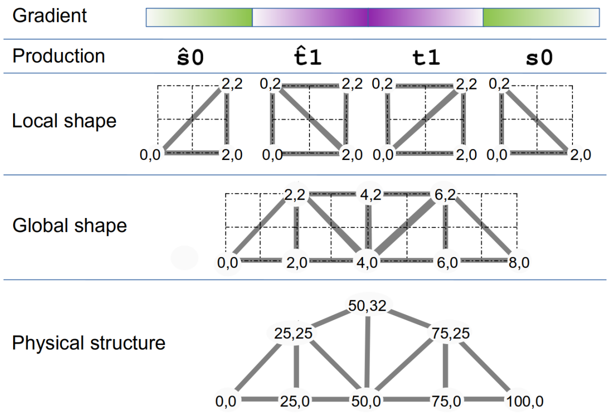

Fig. 2. Gradient analogy and final production for a morphogen grammar. The top row shows the hypothetical morphogen gradient and the second shows the production generated by the grammar as described in figure 1. Each terminal symbol and its attributes correspond to a local shape which is encoded in the grammar (third row). These local shapes are combined to form the global topology by mapping local co-ordinates to global co-ordinates (fourth row), and then the numerical attributes of the terminal symbols are used as a vertical scaling factor when converting the global shape to the physical (or simulated) structure (bottom row).

than that seen in biological organisms). A later processing stage builds a distinct structure in this space, as specified by the type of terminal symbol and its orientation.

The following sections set out the components of the grammar and make reference to a running example for bridge designs. The novel components of the approach are described first, and following this the manner in which the components are linked to established context-free grammars are detailed.

A. Local, Global, and Physical co-ordinate systems

An important property of morphogenesis is the establishment of a co-ordinate system by reference to relative concentrations of chemical agents from cell to cell. To service the ability of the grammar to capture this concept, a three-level co-ordinate system is used. The three levels are *local*, *global*, and *physical* as shown in figure 2.

The local co-ordinate system provides a "cell's eye view" of the structure of individual terminal components. This will be placed within a larger structure as the production is finalised. The local positioning information is mapped to a global Cartesian space by making reference to the cell's relative position in the production. This global space is *discrete* to simplify the process of joining the structures of each cell together at the appropriate place. Finally, the physical co-ordinate space is a mapping of the global space into a continuous space that may be of different scale to the global space, and may be subject to non-linear transformations as specified by the morphogen gradients. For example, the space may be dilated in the 'y' dimension according to the morphogen value of each cell, as shown in the bottom panel of figure 2.

B. Structure graphs

The conceptual gradient analogy along with an example production is shown in figure 2 for a Pratt Truss. The gradient analogy shows the direction and type of local structure as shaded cells. The corresponding production symbols are shown, below which are the local, global and physical representations of the structure.

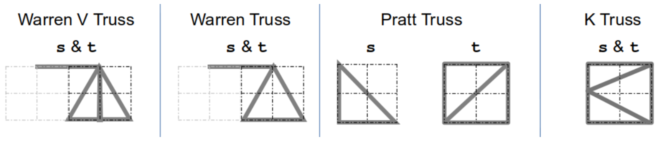

Fig. 3. Structure graphs for four different Truss types that are used for bridge designs. Three of the four truss types have the same structure for terminal symbols 's' and 't', but the Pratt truss has two distinct structures.

Structure in a morphogen grammar is specified by reference to the local co-ordinate system, and specifies connections between points. This is illustrated in figure 3, where terminal structure graphs for four different well-known bridge truss designs [16] are shown, which can be encoded on a genome via indexing the local co-ordinates of each point in the structure.

C. Symbols in morphogen grammars

Productions in morphogen grammars are generated via manipulation of the symbol set using the production rules as in standard context-free grammars. There are three types of symbol: the *start symbol* **`S`**; non-terminal symbols in uppercase, (for example **`B`**); and terminal symbols in lowercase (for example **`t`**,**`s`**).

D. Attributes

Attributes in morphogen grammars are intended to capture the notion of chemical gradients across the developing structure-as-organism and so they are concerned with capturing the magnitude and direction of the concentration as the production rules fire. The magnitude is described using any simple mathematical expression, with the hash symbol '**#**' representing the concentration on the left- and right-hand sides of the production rules.

Since each terminal symbol corresponds with a particular physical structure, the orientation of the structure must also be addressed. The default orientation of a production is to the right, so no attribute is needed to specify this. If the direction of orientation is to the left, this can be specified via a direction attribute, with symbol '^'. This attribute can be used to handle symmetry both when generating productions and when the terminal symbols are mapped into oriented engineering structures. When these symbols are used in generating a production, the symmetry property passes down the hierarchy of symbols.

E. Production rules

Production rules can be used to either generate or parse a language, and in this application, the emphasis is on the generative capability, where the "language" is the arrangement of members of a bridge truss. However, a key difference between morphogen grammars and shape grammars is that the representation of shape does not ascend the parse tree -

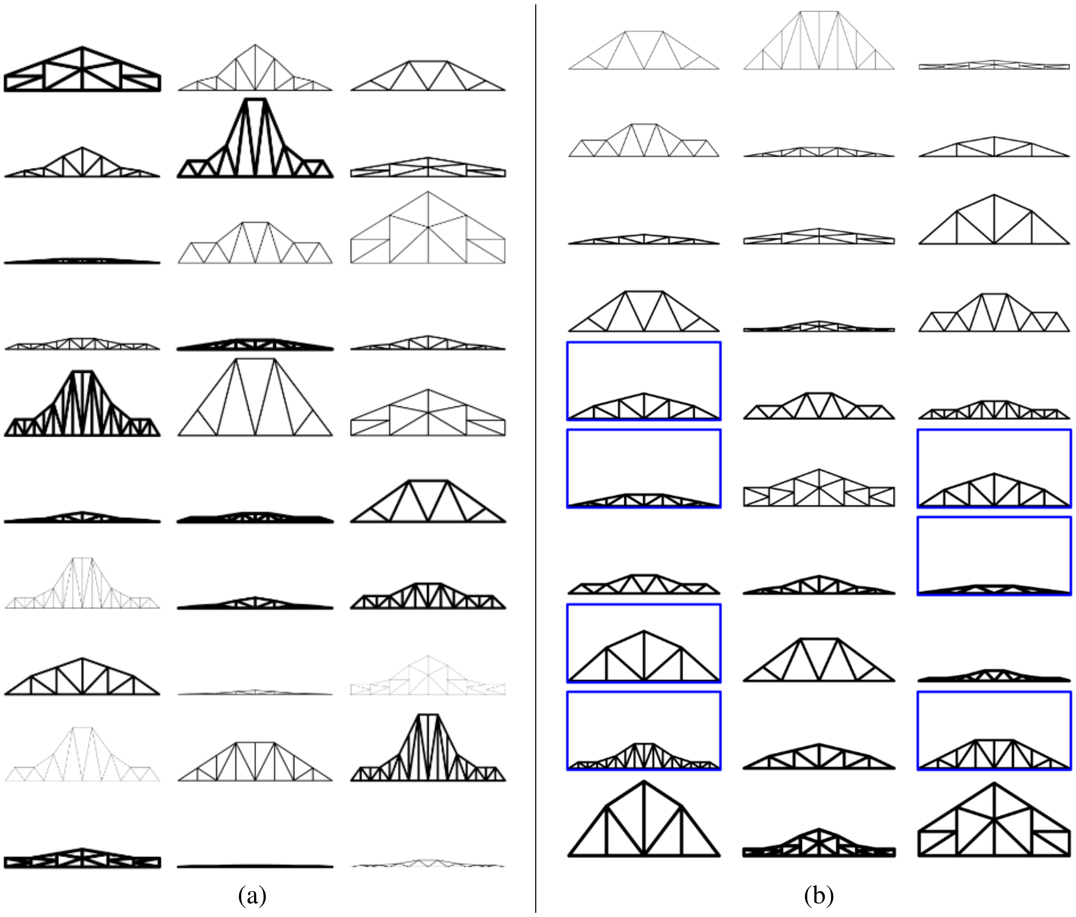

Fig. 4. Example structures from (a) the initial population of bridge trusses and (b) the final Pareto front, showing different height profiles and composition from different structure graphs. Structures satisfying constraints are highlighted with blue boxes.

the only thing that needs to be maintained is the analogue of morphogen concentration, hereafter referred to simply as *concentration*. Thus the production rules are token-based (not shape-based), following standard context-free grammar notation with the addition of attributes to allow contextual information that captures the notion of concentration. This token-based approach makes the encoding of production rules much simpler, allowing them to be represented directly on a genomic representation.

Example: The production rules of the example grammar are:

$$\texttt{S\#} \rightarrow \hat{\texttt{B}}\texttt{\#/2.B\#/2} \quad (1)$$

$$\texttt{B\#} \rightarrow \texttt{t\#-1.B\#-1} \quad (2)$$

$$\texttt{B1} \rightarrow \texttt{s0} \quad (3)$$

A period '`.`' is used to delineate tokens on the right-hand side for clarity. Here, the `#` symbol propagates concentration information through each successive production. For example, rule 1 states that whatever number appears on the left-hand side is halved in the result on the right-hand side, as can be deduced by replacing the '`#`' with a numeric value.

Orientation is determined by the '^' symbol, and states that the orientation should be mirrored in the x-axis for all subsequent productions. By representing gradient information within production rules as a sequence of symbols, it becomes straightforward to encode the production rules directly in a genomic representation, thus allowing the grammar to be embedded in an evolutionary search algorithm.

Example production: Morphogen grammars can be used to generate productions in the same way as other attribute grammars, with the added step of generating the structure from the generated sequence of terminal symbols. Processing begins by specifying an attribute of the start symbol, for example the number 4. The rule shown in equation 1 generates the first production:

$$\texttt{S4} \rightarrow \hat{\texttt{B}}\texttt{2.B2} \quad (4)$$

The input value of 4 on the left-hand side is halved on the right-hand side due to the attribute `#/2`. Each of the symbols on the left-hand side can be processed using the rule in equation 2 to produce a new sequence, and this can be repeated until the token `B1` appears:

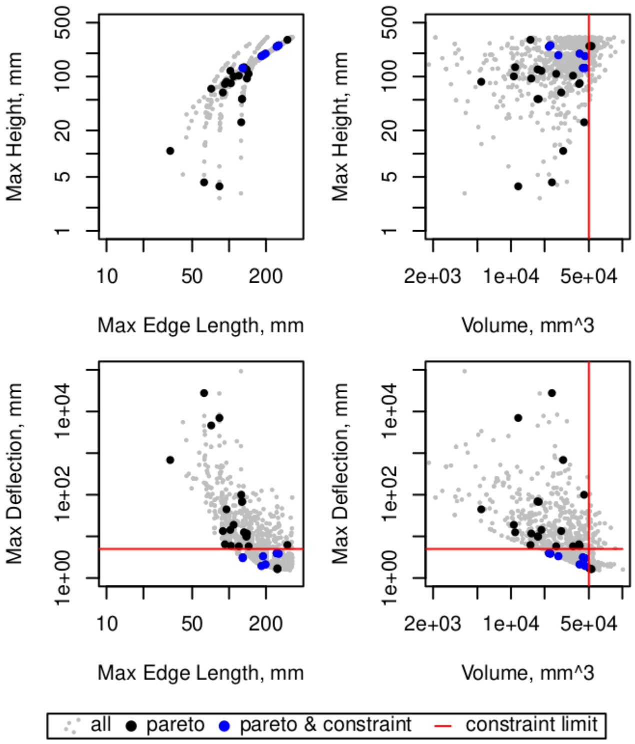

Fig. 5. Fitness of all individuals overlaid with members of the Pareto front from figure 4(b). The four-dimensional fitness landscape is arranged over the four plots, with the maximum edge length and volume shown on the x-axis of the left and right column of plots respectively and the maximum height and deflection shown on the y-axis of the upper and lower rows of plots respectively. This arrangement allows the constraints to be shown as red lines that cross the plots in a consistent manner. The lower left quadrant formed by these lines in the lower right plot shows individuals on the Pareto front that met the constraints.

$$\hat{\mathtt{B}}\mathtt{2.B2} \rightarrow \hat{\mathtt{B}}\mathtt{1}.\hat{\mathtt{t}}\mathtt{1.t1.B1} \tag{5}$$

The final stage in this generation uses equation 3:

$$\hat{\mathtt{B}}\mathtt{1}.\hat{\mathtt{t}}\mathtt{1.t1.B1} \rightarrow \hat{\mathtt{s}}\mathtt{0}.\hat{\mathtt{t}}\mathtt{1.t1.s0} \tag{6}$$

This production is shown in figure 2 and is then used to select the appropriate structure graph for each cell, orient it appropriately, and then assemble the final structure.

III. EXPERIMENTS

In order to demonstrate how the productions of a morphogen grammar can be explored using evolutionary search processes, the preliminary experimental evaluation of this system involves using the NSGA-II multi-objective evolutionary algorithm [17] to discover a Pareto front of bridge designs that can be generated using the morphogen grammar described above. In this initial investigation, four well-known truss designs (Warren, Warren-V, Pratt and K Truss) are encoded into morphogen grammars. These all use the same symbol set, production rules and attributes, but have different mappings of the terminal symbols **`s`** and **`t`** to structure graphs as shown in figure 3.

The evolutionary search has a population of 200 genomes, each of which encodes the following parameters: the period attribute of the start symbol **`S`**, (ranging between 4 and 20); the grammar for truss type (“Warren”, “Warren-V”, “Pratt”, “K”); the cross-sectional area of the members in the structure (between 0.25 and 25 mm^2); the material (aluminium or steel); the scalar on the y-axis height based on local concentration information (between 0 and 1). The evolutionary search algorithm was run for 100 generations using the default settings of the Pymoo library [18].

Each simulated bridge had a fixed support at one end and a roller support at the other and was free to move horizontally (“simply supported” in mechanical engineering terminology). A point load of 17000 Newtons was applied at the centre of the structure. The total span was 500mm. These design decisions were made so that a simulated structure could be fabricated, tested and verified in an engineering workshop in future. This bridge structure was simulated in Calculix [19] as a planar static truss in 3D. Simulation returns forces on the truss members, from which structural properties can be calculated. The fitness measures to be minimised were: the volume of material used; the maximum height; the maximum deflection of any point; the maximum length of any member in the structure.

In addition, the following constraints were used to focus the search on viable structures: the maximum deflection should be less than 5mm; the total volume should be less than 50,000 mm^3. The program was written in Python using the Calculix package for finite element analysis simulation and the Pymoo packages for managing the evolutionary algorithm.

A. *Results*

A subset of the initial population of structures as generated from each of the 200 randomly-initialised genomes is shown in figure 4(a) and a subset of the evolved Pareto front of solutions is shown in figure 4(b) – the latter of these is ordered by volume. Even in the relatively small search space used for this demonstration, it is clear that encoding shape via morphogen grammars offers a means of generating a diverse set of solutions that explore the design space more fully than has been possible previously. This diversity is exploited by the NSGA-II algorithm to yield a range of designs across the four-dimensional fitness landscape.

Structures which meet the constraints on material volume and maximum deflection are highlighted with blue boxes in figure 4(b). Only three of the four available bridge patterns (Pratt, Warren and Warren-V) had configurations that met the specified constraints. The Pratt-based solutions tend to be taller and with an input period of 4 or 6. The Warren-V designs tended to be lower in height and have period from 6 to 8. The Warren truss was the lowest in height, with a period of 4 but with markedly thicker member cross-sectional area. Examples of the K-Truss configuration were found on the Pareto front, but did not pass the constraints, possibly because the inner vertical components of the structure had to be simulated as two members because of the joint halfway up each of these parts of the assembly. K-Trusses have lower risk of buckling than the other structures because the members are generally shorter, but this property was not part of the fitness measure for this experiment. However, all of the structures exploited the facility to use concentration information to transform the physical coordinate space such that an arch shape was achieved.

The evolutionary run evaluated 20,000 designs. The final Pareto front contained 894 individual designs, of which 556 met the two design constraints on volume and deflection. The distribution of these points in the fitness space is shown in figure 5. It can be seen that the constraints on the evolutionary run are quite limiting in terms of the range of designs that are available, but these requirements could be revised if needed for subsequent runs in order to arrive at the final set of design candidates.

IV. Discussion

Grammars are attractive vehicles for design by evolutionary search because of their ability to combine complex concepts in myriad ways that conform to hierarchical rules of composition. However, only a small proportion of published research on grammars pertains to the generation of novel designs (7% of papers reviewed by [20] for example).

The goal of the current endeavour is to generate a highly diverse set of structures that are viable for the application at hand. Completely random structures are unlikely to have utility – instead, a degree of modularity combined with the ability to incorporate existing engineering knowledge into the evolutionary process allows a tractable method to be developed which meets this goal, as shown above.

Although the current experiment is configured for 2D, an extension to 3D is easily attainable by formulating the structure graphs for each terminal in three dimensions. The approach presented here exploits the linear structure of bridge designs - repeated 3D modules in a line - that is not always available to other design problems such as vehicle design. A further extension of the current approach to more general design problems is an active area of research for the authors of this contribution.

The generation of productions for these structures is currently fully deterministic once the attributes of the start symbol are set. It may be possible to use local feedback from the finite element analysis to drive the firing of different production rules in different sub-regions of the structure. A degree of stochasticity may be useful here, which would have the effect of smoothing the fitness landscape as appropriate decision trees evolve.

Having demonstrated this approach on a relatively simple bridge design problem, several avenues of research open up. The simple support model and static load conditions used in the experiment above were relatively idealized, so a more detailed evaluation of the technique on more complex structural design problems, including a detailed analysis of key engineering parameters such as fatigue characteristics and material aging, is the next step in this research. The morphogen grammar has been designed to be fully encodable on the genome, which means that evolutionary approaches can be used to generate alternative production rules and structure graphs to increase the diversity of the solution space even further. It is striking thatv the production rules for the four bridge types presented here were identical - the difference was only in the structure graphs for each bridge as shown in figure 3. The composability of shape that this demonstrates is a key strength of this approach and should be studied further.

Acknowledgment

The authors would like to acknowledge the Institute for Safe Autonomy (ISA)[1].

References

[1] L. C. Biggs and M. L. Mikkola, "Early inductive events in ectodermal appendage morphogenesis," *Semin. Cell Dev. Biol.*, vol. 25-26, pp. 11–21, Jan. 2014.

[2] L. Wolpert, "Positional information and the spatial pattern of cellular differentiation," *J. Theor. Biol.*, vol. 25, no. 1, pp. 1–47, Oct. 1969.

[3] J. Sharpe, "Wolpert's french flag: what's the problem?" *Development*, vol. 146, no. 24, Dec. 2019.

[4] R. Doursat, H. Sayama, and O. Michel, "A review of morphogenetic engineering," *Nat. Comput.*, vol. 12, no. 4, pp. 517–535, Dec. 2013.

[5] R. Doursat and C. Sánchez, "Growing fine-grained multicellular robots," *Soft Robotics*, vol. 1, no. 2, pp. 110–121, Jun. 2014.

[6] N. Cheney, J. Bongard, and H. Lipson, "Evolving soft robots in tight spaces," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '15. New York, NY, USA: Association for Computing Machinery, Jul. 2015, pp. 935–942.

[7] S. L. Thomson, L. K. Le Goff, E. Hart, and E. Buchanan, "Understanding fitness landscapes in morpho-evolution via local optima networks," Feb. 2024. [Online]. Available: http://arxiv.org/abs/2402.07822

[8] T. Zhao, "An implementation of the ground structure method considering buckling and nodal instabilities," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2015. [Online]. Available: http://paulino.princeton.edu/theses/ms/14zhao_msthesis.pdf

[9] P. Hajela and E. Lee, "Genetic algorithms in truss topological optimization," *Int. J. Solids Struct.*, vol. 32, no. 22, pp. 3341–3357, Nov. 1995.

[10] A. McKay, S. Chase, K. Shea, and H. H. Chau, "Spatial grammar implementation: From theory to useable software," *Artif. Intell. Eng. Des. Anal. Manuf.*, vol. 26, no. 2, pp. 143–159, May 2012.

[11] G. Stiny, "Introduction to shape and shape grammars," *Environ. Plann. B Plann. Des.*, vol. 7, no. 3, pp. 343–351, Sep. 1980.

[12] U. Roncoroni, "Programming complex 3D meshes. a generative approach based on shape grammars," in *XXV Generative Art Conference*, 2022.

[13] M. Barros, J. P. Duarte, and B. M. Chaparro, "A grammar-based model for the mass customisation of chairs: Modelling the optimisation part," *Nexus Netw. J.*, vol. 17, no. 3, pp. 875–898, Dec. 2015.

[14] M. O'Neill, J. McDermott, J. M. Swafford, J. Byrne, E. Hemberg, A. Brabazon, E. Shotton, C. McNally, and M. Hemberg, "Evolutionary design using grammatical evolution and shape grammars: designing a shelter," *Int. J. Des. Eng.*, vol. 3, no. 1, p. 4, 2010.

[15] D. E. Knuth, "Semantics of context-free languages," *Math. Systems Theory*, vol. 2, no. 2, pp. 127–145, Jun. 1968.

[16] J. Kulicki, "Highway truss bridges," *Bridge Engineering Handbook*, Jan. 2014.

[17] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002.

[18] J. Blank and K. Deb, "Pymoo: Multi-objective optimization in python," *IEEE Access*, vol. 8, pp. 89 497–89 509, 2020.

[19] Dhondt, "Calculix crunchix user's manual version 2.12," *Munich, Germany*, 2017. [Online]. Available: http://www.dhondt.de/ccx_2.19.pdf

[20] S. M. Haakonsen, A. Rønnquist, and N. Labonnote, "Fifty years of shape grammars: A systematic mapping of its application in engineering and architecture," *Int. J. Archit. Comput.*, vol. 21, no. 1, pp. 5–22, Mar. 2023.

---

[1] Website: https://www.york.ac.uk/safe-autonomy/