Measuring the robustness of a developmental system based on sequential growth rules

Michael Cohen[1], Mark Miodownik[2] and Buzz Baum[1]

[1]MRC-LMCB, University College London, UK
[2]Material Research Group, King's College London, UK
m.cohen@ucl.ac.uk

Abstract

Understanding how complex structures emerge from localised interactions in a robust way is essential to unraveling the mechanisms that underlie developmental processes in both biological and artificial systems. This study investigates the effects of genome complexity on robustness using a simple, evolved developmental system in which cellular automata (CA) rules are applied in sequence in order to generate a 1D pattern of cells. The system employs a 1D two state CA with 128 distinct nearest neighbour update rules. Each developmental run is initiated with a single cell. The cell update rules adopted by every cell at each time-step are allowed to change sequentially at different times according to the instructions contained in a 'genome'. In order to generate a set of productive developmental programs for this analysis, a genetic algorithm was used to select for individuals whose cell states, after a fixed number of time steps, match a set of pre-defined target patterns. This was repeated for genomes of different sizes. The robustness of evolved and randomized CA patterns were compared by systematically applying single cell state perturbations during pattern development. This analysis revealed that in these evolved systems genome size has a positive effect on robustness by freeing the system to generate patterns using a relatively unbiased set of rules, which have very different individual properties. In contrast, smaller genomes are frequently forced to rely on complex patterning rules to generate complex patterns, which amplify damage and hence reduce their robustness. In addition, pattern size (the number of cells) was found to be a major factor in the measured robustness in this system. This is because the cumulative damage induced by developmental perturbations does not scale with pattern size. As a result, increasing pattern size reduces the percentage damage following perturbations and improves overall robustness. In conclusion, we have shown that pattern robustness is an additive effect of the ability of individual rules to propagate and heal defects resulting from environmental perturbation in this simple CA system, and is potentially increased by increasing pattern size and genome size. These results have implications for our understanding of robustness in biological and artificial systems.

Introduction

Both natural and artificial developmental systems are known to generate physical forms that are self-regulating and as such are highly robust to perturbations of many kinds including artificial wounding or cell removal (Wolpert, 2002; Kumar and Bentley, 2003). Understanding how complex structures emerge from localised interactions in a robust way is essential to unraveling the mechanisms that underly developmental processes. In biological and artificial developmental systems, development is often governed by cellular interactions. Fundamentally different processes may occur in sequence at different developmental stages in order to generate overall pattern and form (Wolpert, 2002). This paper explores how using different numbers of rules in sequence during the development of a cellular automata (CA) pattern affects the overall robustness of the patterning process.

Robustness to cell perturbation (or 'wounding') and self-regulation of developed patterns or 3D forms has previously been observed as an emergent property of evolved developmental CA systems (Andersen et al., 2006; Miller, 2004; Basanta et al., ress; Devert et al., 2007; Federici and Downing, 2006; Grajdeanu and Kumar, 2006; Streichert et al., 2003). However, because these systems are complex it is unclear precisely how the implicit developmental rules in these systems lead to a robust developmental program. In order to create a simple system in which the processes underlying the evolution of developmental robustness could be simply and rapidly analysed in detail, we developed a 1D model, in which CA rules are applied in series. CA rules are known to produce characteristic patterns relating to their dynamical properties and overall system stability (Wolfram, 2002) but it is not immediately apparent how such properties may contribute to the effect of cell perturbations during their development. In particular, no previous study has established how using CA rules in temporal sequence may effect the overall system stability and hence the robustness of patterning. Using this system, we explored the roles of evolution and genome complexity on developmental robustness. In each case, we compared results from evolved genomes with those obtained using an equivalent set of random genomes.

Method

The experiment uses a 1D two-state CA of the type defined by Wolfram (Wolfram, 2002). This system consists of a line of *cells* in one of two states; black or white. (The lines are

effectively infinite to avoid edge effects.) In this paper a black square is referred to as a cell a white square represents an empty space. At each time-step in the running of the CA, each location is updated according to a set of conditions dependent only on its previous state and the state of its two adjacent neighbours. The complete set of conditions defines an update rule, which operates on all cells in the system at any one time step. Here, a sub-set of 128 rules are used which exclude those rules whereby a cell can emerge from an empty neighborhood of cells. These are labelled according to Wolfram's numbering scheme and comprise the even numbers between 0 and 255.

The CA are developed for 51 time-steps at which time the 1D pattern generated is referred to as the 'end-state pattern'. In this system the rules are allowed to vary over different time periods, as shown in figure 1, where in this case 6 distinct rules are implemented in series. The particular rule applied to every cell at each time-step is contained in the 'genome for each individual run of a CA. The whole population in any single experiment has the same genome length, or number of genes, n. The specific case illustrated by figure 1 is represented by the n=11 genome: 10 50 174 242 230 122, 9 15 24 32 45. Here the first six numbers represent the set of rules (R1-R6). The remaining five numbers represent the transitions times (T1-T5) at which the rules change. The transition times are constrained to occur in evenly distributed fractions of the total 51 time-steps. For example. in the n=11 case shown, the 5 transition times occur in bins of 10 time-steps. Where the CA patterns are directed by artificial evolution, the fitness function (defined subsequently) is applied at time-step 51, where the end-state pattern of cells is compared to a pre-defined target pattern (shown in grey).

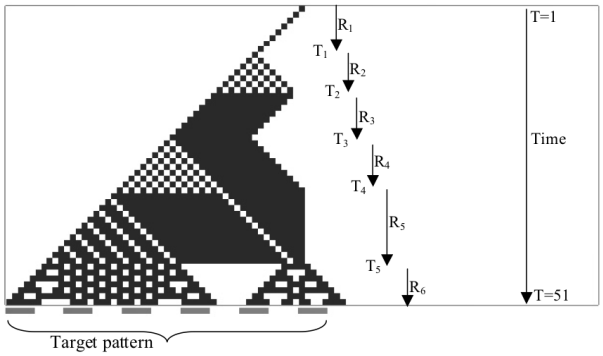

Figure 1: A screen shot of an individual CA run. The end-state pattern at time-step 51 is developed according to the cell update rules. Six rules (R1 to R6) are applied to the system over six different time periods; the transition points of which are labelled T1-T5. The light grey pattern below the box shows the target pattern, P1, towards which the system may be evolved.

Evolving patterns

To test the behaviour of the system under specific types of directed patterning the CA were evolved using a Genetic Algorithm (GA) (Davis, 1991; Mitchell, 1998). This was applied as follows. A population of size N=500 individual genomes was created and these were each developed in accordance with the CA program. Genes were initially seeded by a random number generator. The rule defining genes were selected randomly from the complete set and the time values were randomized within the time period constraint as previously described. A fitness function scored each individual according to the similarity of their end-state pattern, at time-step 51, with a pre-defined target pattern. The target patterns used are shown in figure 2. These were selected to test the effects of varying pattern regularity, symmetry and breadth of distribution. The first six patterns, P1-P6, are the same size, 30 cells, to enable direct comparison, whilst patterns P7 and P8 are 60 cells in size to control for the effects of pattern size.

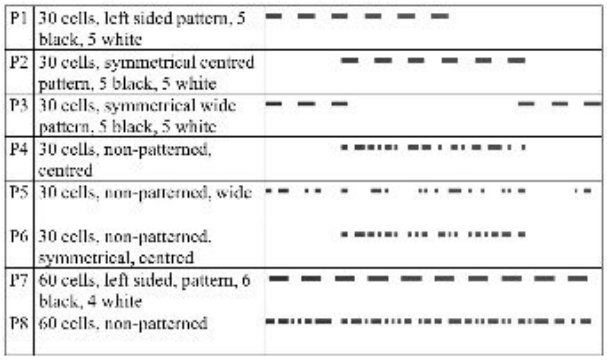

Figure 2: Target patterns selected to test for pattern regularity, symmetry, distribution and size.

The fitness function sums the number of cells that differ in their location between the target pattern and end-state pattern of the developed CA. This is equivalent to the 'Hamming distance' between the two bitwise pattern encodings (Hamming, 1950). Thus the most 'fit' individuals have the lowest 'fitness score' and a perfect correlation scores zero. Tournament selection was used to determine which individuals pass to the next generation; whereby, two individuals are randomly chosen and the fitter individual selected. Crossover was not found to benefit the GA and was not used. The genomes of the next generation were mutated by randomly selecting either new CA rules from the complete set or transition times from within the constraints previously described. The mutation rate, per genome, used at each genome size, n=3,11 and 23, were; 0.6, 0.8, and 1.0 respectively.

The process of selection and mutation leads to a new generation after which the whole process is repeated. Throughout the experiment a fixed population of N=500 was used

and the system was evolved for 1000 generations for target patterns P1 to P8 as well as for an extended 5000 generations for pattern P1 (this set of data is referred to in the results as P1+). Ten evolutionary runs were carried out for every genome size and target pattern. These parameters were all optimised prior to the experiment and were found to be sufficient to achieve stable average fitness scores of low variance. Genomes of sizes 3,11 and 23 were used to compare the effects of genome complexity in this system.

Robustness Testing

Evolved solutions and unevolved, randomly generated genomes were tested for their robustness to cell perturbations. Each single (black) cell was systematically perturbed (cell state changed to white), one at a time, during the pattern development. The emergent end-state pattern after each cell perturbation was compared with that of the unperturbed CA (see figure 3). The damage caused by each cell perturbation was measured in terms of the Hamming distance between the perturbed pattern and the original end-state pattern. This difference was then expressed as a percentage of the original pattern size (the total number of black cells in the end-state pattern). The overall developmental robustness of a particular individual was regarded as being inversely proportional to the averaged percentage damage caused by all developmental cell perturbations. Mean data from 750 randomized genomes of each genome size was compared with mean data from the 10 evolutionary runs at each target pattern.

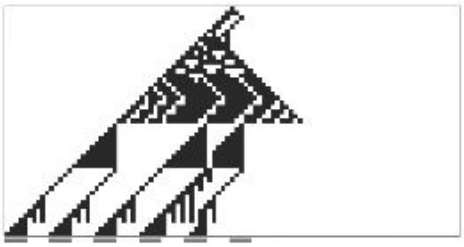

(a) Original n=23 solution consisting of 25 (black) cells.

(b) A single cell perturbation (from black to white) causes a shift in the end-state pattern such that 10 (black) cells are in a different location. Equivalent to a damage score of 40% of the final pattern size.

Figure 3: Measuring the effects of cell perturbations.

Results

A set of CAs with genomes of different sizes (n=3, 11 and 23) were evolved under a genetic algorithm by selecting for their ability to match a set of 8 pre-defined target patterns. The genomes contained instructions for the transient update of the CA rules. In this section the evolved solutions are investigated with regard to their relative success in matching target patterns, the developmental methods adopted to try meet those target patterns and their robustness to developmental cell perturbations.

Pattern Characteristics

Examples of the evolved solutions are shown in figure 4. There were variations in the 10 solutions obtained at each evolutionary run and the subset shown here are intended to illustrate some of the generic differences between the target pattern types and genome sizes. Most immediately striking is the difference in the developmental profiles (that is all the cells at each time-step leading up to the end-state pattern) among the different genome sizes. The n=3 solutions have very distinct profiles characterised by the two different rules applied to meet the target pattern. In contrast the n=23 developmental profiles share a common feature of branching or segmentation at the transition between the 12 rules comprising their genome. There is a complexity of patterning that arises as a result of these rule transitions. The n=11 solutions reflect an intermediate case. It is immediately apparent that the n=11 and n=23 genomes are good at matching the more regularly spaced target patterns but bad at matching a highly distributed random target pattern such as at P5. For the larger patterns, P7 and P8, all individuals of the 3 genome sizes rely on rules that cause an expansion or growth in the number of cells present, as might be expected.

Whilst the target patterns P1-P6 all consisted of 30 cells, the evolved end-state patterns varied in size between 8 and 35 cells. Among randomly generated genomes there was also a significant variation in pattern size. In data obtained from 750 random genomes of each genome size, the average end-state pattern size for n=3, 11 and 23 was 14, 7 and 3 cells respectively. Although the average size was seemingly, relatively low, significantly larger patterns of over 60 cells were also generated by the random samples. The size of the end-state patterns was found to have a significant effect on the robustness of the CA, as is shown later in these results.

Fitness of Evolved Solutions

The GA is designed to identify solutions that match the target pattern. This was shown to be the case, since for all genome sizes the GA yielded patterns with an improved fitness score. The average scores obtained for each of the genome sizes are shown in figure 5. Average scores are given for the champion individuals at the first and last generations. Overall the larger genomes show slightly less fit

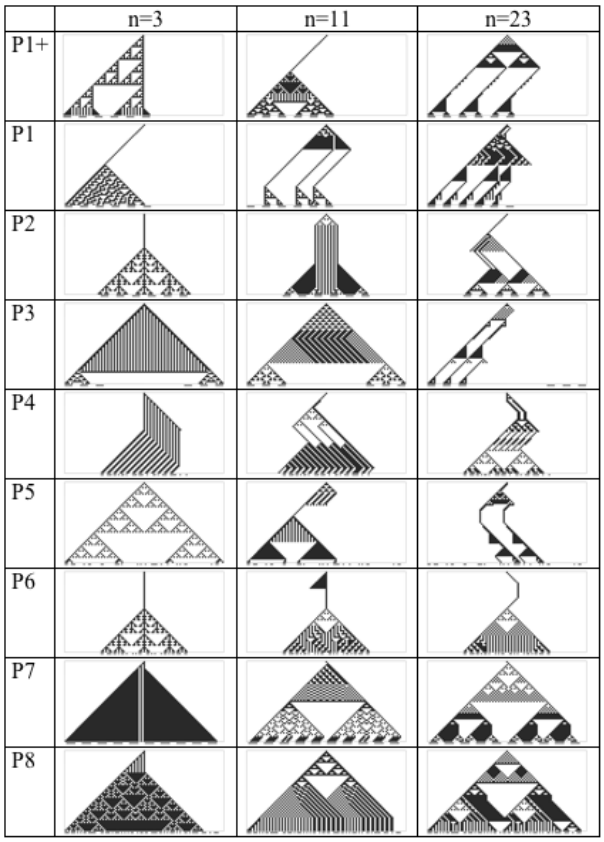

Figure 4: Examples of evolved champion solutions obtained at the last generation of evolutionary runs carried out for each genome size at each target pattern. The pattern at each time step is shown with developmental time represented in the vertical axis.

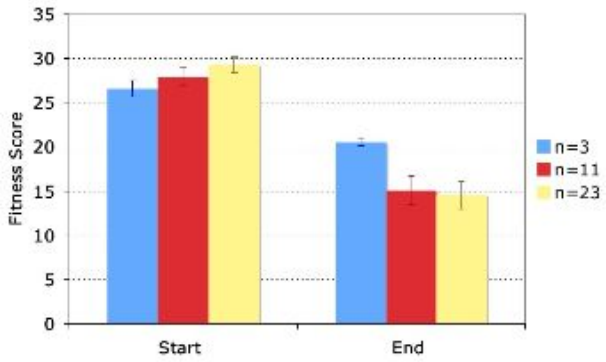

Figure 5: The average champion scores attained by each genome size. The data compares the lowest fitness scores from the first generation (labelled 'Start') and the last generation (labelled 'End'), averaged over all target patterns for all 10 evolutionary runs. Error bars show the 95 percent confidence intervals for the mean values.

(higher) scores at the start of the evolutionary runs, thus indicating that a random population is less likely to match the target patterns. After evolution the n=11 and n=23 genomes achieve very similar average scores both significantly fitter than for the n=3 case.

There were identifiable differences between the target patterns. The n=11 and n=23 genomes consistently outperformed the n=3 genome except in the case of one target pattern, P6. In general for the two larger genome sizes the regularly spaced target patterns P1, P2 and P7 achieve the fittest relative scores. Where more complex arrangements of cells were encountered these systems did less well in matching the end-state patterns.

In order to further qualify the relative evolvability at each genome size the fitness scores obtained by evolution were compared with those of a randomly generated population of 500,000. This is the equivalent number of individuals that are searched by the GA evolving a fixed population of 500 individuals over 1000 generations. The n=3 evolved solutions never out-performed the random search solutions. In contrast, for the n=11 and n=23 genomes all of the evolved solutions outperformed the random search.

Robustness to developmental cell perturbation

To analyse the effects of genome size on developmental robustness in this system cell perturbations were made to both evolved and unevolved individuals (see method for details). Figure 6 shows a plot of this data. Here, the average percentage damage score has been plotted against the size of the end-state patterns. For the random genomes each individual data point is plotted together with a trend line indicating the population mean and associated confidence intervals for this value. For the evolved solutions, data points are plotted showing the mean value obtained over the 10 evolutionary runs with associated confidence intervals.

The evolved solutions for the n=23 genome all sit on the same trend line as for the random genomes and the range of random data in this case is much more constrained than for the n=3 and n=11 genomes. In contrast, for the n=3 and n=11 genomes the distribution of the random data is larger than for n=23. For some target patterns the mean robustness of the evolved solutions patterns is different to the mean random data of equivalent size. The n=3 evolved solutions for target patterns P1, P2, P3, P5 and P6 all have a mean robustness that is significantly lower than for the random data (evolved individuals show higher percentage damage scores within an equivalent pattern size range). For n=11, the solutions at target patterns P1 and P2 are significantly less robust than the average data. This would suggest that evolution towards these specific target patterns has repeatedly selected for combinations of rules and transition times that are less robust than the average random sample. Part of this loss of robustness may be attributed to the fact that these individuals sometimes show a sustained period without pattern

growth that is inherently weak to any perturbation; as can be observed in the examples shown in figure 4 where a single cell is maintained over a number of time-steps before any larger pattern finally develops. A perturbation during this early period without growth will remove the entire pattern. In contrast, the n=23 solutions consistently employ periods of growth and patterning throughout the pattern development. Another factor underlying the loss of robustness of some of the evolved solutions may be a selection for individual rules that are inherently sensitive to perturbations. This will be analysed further in the discussion section.

For all three genome sizes the predominate factor determining robustness is the end-state pattern size itself. To further investigate the effects of end-state pattern size as well as genome size, the mean trend lines from the randomized data are plotted together in figure 7. The curves from each genome size all follow the same trend and there is no significant variation in robustness. Thus it can be concluded that the use of a greater number of rules does not translate into a change in robustness in this system.

The real data are contrasted with curves that represent the effects of altering the state of 2, 4 and 6 cells in the end-state pattern; that is, a theoretical plot in which for each cell perturbation the end-state pattern is altered by a fixed amount. The curves obtained from the randomized genomes all follow a trend very similar to that of a fixed 4 cell perturbation. Only for very low pattern sizes, below approximately 10 cells, do the curves align more closely with a fixed absolute damage of 2 cells. This would suggest that regardless of the size of the pattern generated (and thus the average rate of growth of black cells) the average, absolute damage caused by cell state perturbations remains fairly constant over a wide number of randomized genomes. It is important to note that this is an average quantity. The effect of a cell perturbation early in development, where there are fewer cells, causes significantly more absolute damage than one very late in development (where there are likely to be many more cells). What is suggested here is that, averaged over developmental time, the absolute damage caused by a perturbation is largely independent of the ultimate pattern size. The effects of a cell perturbation do not scale in accordance with the rate of pattern growth and end-state pattern size, as might be expected. Hence, the percentage damage caused by a single perturbation rapidly decreases with increasing pattern size as the curves here demonstrate.

The results have shown that for all three genome sizes, there is very similar trend between the average robustness of randomly generated CA and their end-state pattern size. For evolved CA, the average robustness is shown to differ among solutions obtained at the different target patterns. Whilst the variation in evolved robustness can principally be explained by differences in evolved pattern size, some evolved solutions show a lower than average robustness than was obtained for CA derived from random genomes of equiva-

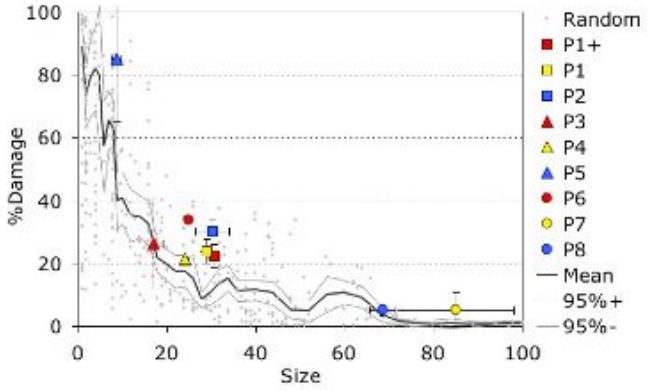

(a) n=3

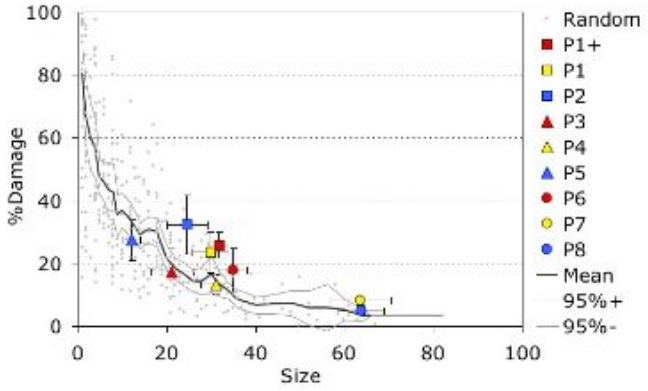

(b) n=11

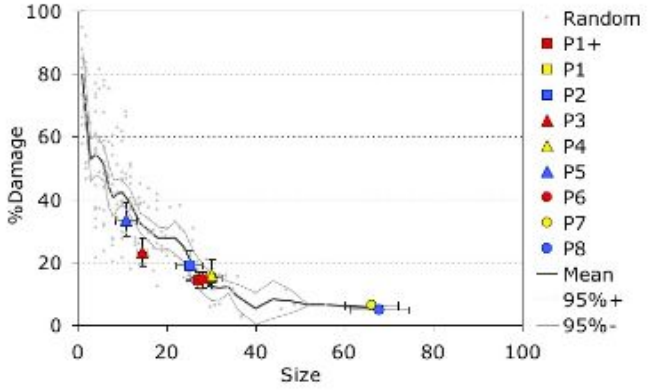

(c) n=23

Figure 6: A plot of end-state pattern size against cell perturbation damage expressed as a percentage of size. Data was obtained from 750 randomly generated genomes of each genome size. The trend line shown the mean of this data with associated 95 percent confidence intervals (derived from data bins across ranges of sizes). For the evolved solutions, mean values and confidence intervals derived over 10 evolutionary runs are shown.

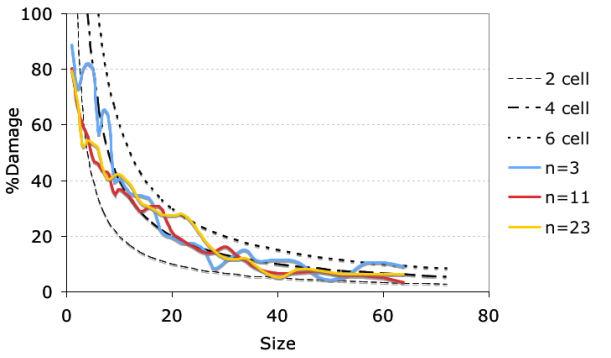

Figure 7: The mean trends of end-state pattern size against the average cell perturbation damage expressed as a percentage of the original pattern size. The data was derived from 750 randomly generated genomes of each genome size. This is contrasted with model curves representing a fixed absolute damage, at all pattern sizes, of 2, 4 and 6 cells.

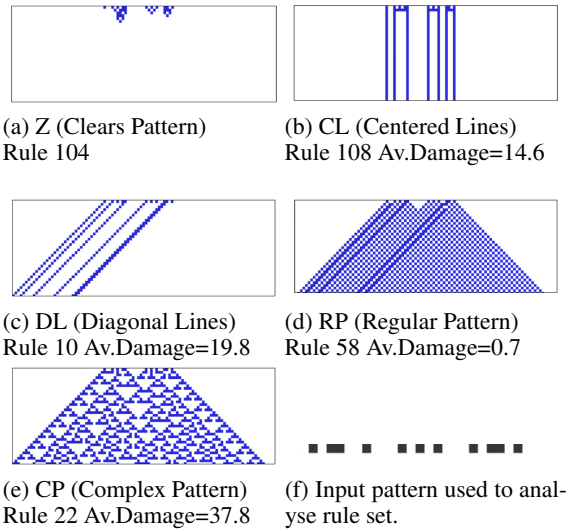

Figure 8: Rules classified according the defined criteria. Shown here are examples of each rule 'type'. The rule number is quoted along with average percentage damage score for that particular rule when each cell was systematically perturbed.

lent pattern size. Therefore, it can be inferred that in order to match targets the evolutionary algorithm is repeatedly selecting for particular combinations of rules that degrade overall robustness in these particular cases.

In order to better understand the effects of individual CA rules on robustness in this system, the rules were categorized and analysed in isolation. Figure 8 demonstrates how the individual rules were categorized. The figures show the behaviour of each rule after input from an arbitrary pattern comprising 11 cells in 9 discrete blocks at time-step one. Each was run for only 40 time-steps to account for the additional 'width' of the input pattern. This 'input' pattern was selected to illustrate the behaviour of the rules at some time into the development of a pattern, as distinct from seeding by a single cell.

A measure of the end-state pattern size and the average percentage damage caused by cell perturbations was made for each individual pattern. These are plotted in figure 9. This shows how the regular patterning (RP) rules are significantly more robust to cell perturbation than the complex patterning (CP) rules, regardless of the pattern size. The emergence of a regular pattern of growth from the irregular input pattern indicates that the system has a stable attractor state that is largely insensitive to initial conditions. Thus perturbing the system later in development has a similarly low effect on the emergent pattern. There is a self organization inherent in these types of rules. For the complex patterns the system is more sensitive to the initial conditions and forms complex pathways in the development of the pattern, with subsequent interactions when pathways intertwine; this results in the nested triangles characteristic of a complex pattern developmental profile. In this case information about previous cell states is transmitted throughout the CA in such a way that cell perturbations have an escalating effect on the emergent patterns at subsequent time-steps. The DL and CL rules that produce substantially less pattern growth show a perturbation response that scales very sharply with pattern size.

The mean trend line gives an indication of the average damage at each size for all the individual rules. It is interesting to note that when contrasted with the curves shown in figure 7 the mean trend among the individual rules closely follows the mean trend for the randomized genomes. This suggests that the average robustness of each of the combinatorial rule systems is essentially the same as the average robustness of the individual rules themselves. This reinforces the finding that the genome size has no intrinsic effect on the average robustness. In addition, it appears that the approximation towards a constant absolute damage (of approximately 4 cells), that was noted previously, can be attributed to a combinatorial effect of the different types of rules. Individually the different types of rules have quite distinct relationships in regards to pattern size and robustness. However the trend line shows that their aggregated relationship closely mimics that of a system with a fixed average response to perturbations in regard to pattern size.

It should be noted that the classification scheme adopted here is not concrete and there are a few rules that generate pattern that appear to be on the border between these types of classification. There is a correspondence with Wolfram's classification system for this type of CA, such that CL and DL are Class 2, RP Class 1 and 2, and CP Class 3. Rules that fall between between RP and CP are Class 4 systems (Wol-

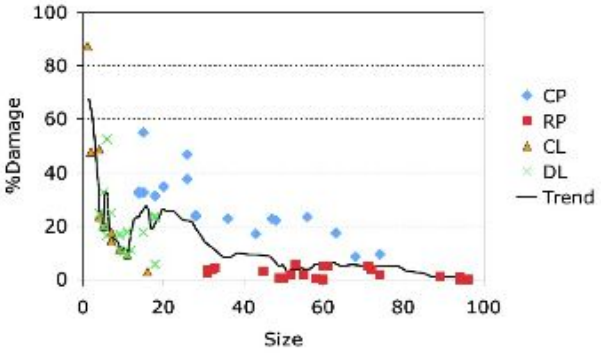

Figure 9: The robustness of individual CA rules of each classification type. The average percentage damage caused by cell perturbations is plotted against the end-state pattern size at time-step 40. The trend line shows a rolling mean average of all the data.

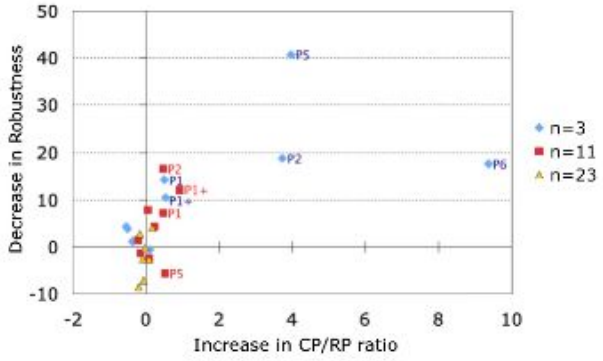

Figure 10: The effect of complex patterning rules on the robustness of evolved solutions. The x-axis shows the average difference in the evolved CP/RP ratio with the rule-set ratio. The y-axis shows the average difference between the evolved robustness scores (expressed as an average percentage damage due to cell perturbation) and the mean robustness of random data of equivalent pattern size (from figure 6). Data points located in the upper right quadrant reveal a correlation between complex patterning and a loss of robustness.

fram, 2002). This system of classification has convergence with other definitions relating to the dynamical properties of CA (Wuensche and Lesser, 1992). The principle distinction made here is that the RP rules are more dynamically stable than the CP rules.

To further investigate why the evolved solutions showed differences in their robustness in comparison to the randomized data, an analysis was carried out with regard to the proportion of rules adopted by the evolved genomes. For each of the evolved solutions the ratio of CP rules to RP rules was determined. The increase in this ratio, as compared with the actual rule set was then calculated. This value is plotted in figure 10 against the increase in the average perturbation damage score obtained by evolved solutions as compared to the mean randomized data of equivalent size (as illustrated in figure 6).

This analysis reveals that where CP rules have been used in high proportion, there is, in most cases, an equivalent decrease in robustness (increase in the percentage damage caused by cell perturbations). Therefore, it seems that in general a loss of robustness can be explained by an increased uptake in CP rules, which are required in order to match certain target patterns and are thus selected for by the GA. This generalization is true in all but one example, where for the n=11 genome at target pattern P5 the evolved solutions are seemingly more robust than average whereas the CP/RP ratio is higher than for the rule set itself. This may be attributable to the very small end-state pattern size that was adopted by these solutions, making them more robust than equivalently sized random patterns, even though a significant amount of their development was undertaken by complex growth rules.

Discussion

In summary, this analysis has demonstrated that there is no intrinsic emergent robustness as a result of increasing the number of sequential 'rules' in a CA system but there is a potential loss of robustness associated with evolved rule biases in smaller genomes. On average the two larger genomes were shown to evolve better (more fit) solutions than the smaller genome. The evolvability of the larger genome sizes is related to the size of the parameter space that they may select from. The greater complexity of the genome provides the means for complex adjustments in the patterning of cells that is not present in the individual rules themselves. Thus the n=3 genomes and to some extent the n=11 genome were more reliant on the use of specific rules for the generation of particular patterns and it was shown that when complex rules were used their robustness was degraded. Though the n=23 solutions were not inherently more robust to cell perturbations than the n=3 or n=11 genomes, they did not deviate from a random distribution in their selection of rules and so showed higher levels of robustness when the smaller genomes were forced to do so in order to achieve the required patterning.

Robustness, here, was explicitly defined as a percentage change in the phenotypic patterns as this was considered to provided the most informative comparison between different evolved solutions. It was shown that, on average, the propagation of a perturbation through development only effects a limited number of cells in the end-state pattern. This means that it is predominantly the size of a developed pat-

tern that contributes to its overall robustness such that the percentage impact of cell perturbations is reduced as size increases. In biological systems there may be a corresponding relationship between organism size and robustness such that larger organisms, containing a greater number of cells, may show less phenotypic response to both developmental and genetic perturbation. Research into the evolutionary adaption of size highlights the physiological or environmental constraints acting on an organism (LaBarbera, 1989). It may be that there is an underlying selective pressure to increase organism size for overall robustness.

There was no evidence for emergent robustness as a product of the GA itself. Adding stochasticity or noise to the CA development, by introducing cell death, may cause the system to evolve more robust solutions. In this scenario it may be that 'fit' solutions that can withstand developmental noise are more likely to be repeatedly selected for during evolution.

The particular CA used here update rules that operate at every site with relatively complex asymmetrical configurations. Rather than cell growth, it more closely represents a collection of established cells making internal decisions about their differentiation between states. Future work may explore how a similar system may be reconfigured to better represent more realistic cellular growth rules.

Conclusion

This study has provided a measure for the developmental robustness of evolved CA patterns in a simple one dimensional system. It has shown that there is no robustness intrinsically associated with using additional rules. However, increasing the complexity of a genome has a beneficial effect on robustness simply because it frees the system to generate patterns using a relatively unbiased set of rules.

For randomized genomes of each genome size individual cell perturbations, on average, produced approximately the same amount of absolute damage to the emergent patterns. This was shown to be equivalent to approximately 4 cells. Robustness, here, was explicitly defined as a percentage change in the phenotypic patterns. Hence, there was a strong correlation between robustness and pattern size.

It was revealed that the robustness of randomized genomes could be attributed to the aggregate effect of selecting from the complete rule set. Different types of rules demonstrated very distinct relationships between robustness and pattern size. However, a trend line showing the mean robustness of each of the individual rules approximated the trend one would expect given a fixed amount of absolute damage, regardless of overall pattern size. Thus the average robustness of the randomized genomes could be simply interpreted as a reflection of the average robustness of the individual rules.

In the analysis of individual rules it was shown that rules generating complex patterns were sensitive to cell perturbations. By contrast, the regular patterns in this system act as a stable attractor in which the state of cells enter a single homogenous state or a predictable cycle that is insensitive to changes in their input. Where genomes of the evolved CA patterns adopted a large proportion of 'complex' patterning rules, their robustness is shown to be reduced in comparison to the average robustness of patterns of equivalent size.

Acknowledgements

This work was supported by CoMPLEX, UCL.

References

Andersen, T., Newman, R., and Otter, T. (Fall 2006). Development of virtual embryos with emergent self-repair. In *American Association for Artificial Intelligence - Symposium on developmental systems*, pages 16–23.

Basanta, D., Miodownik, M. A., and Buzz, B. (in press). The evolution of robust homeostasis and the stem cell niche in artificial multicellular organisms. *PLoS Computational Biology*.

Davis, L., editor (1991). *Handbook of Genetic Algorithms*. Van Nostrand Reinhold.

Devert, A., Bredeche, N., and Schoenauer, M. (2007). Robust multi-cellular developmental design. In *Genetic and Evolutionary Computation – GECCO-2007, Part I*, pages 982–989.

Federici, D. and Downing, K. (2006). Evolution and development of a multicellular organism: scalability, resilience, and neutral complexification. *Artif Life*, 12(3):381–409.

Grajdeanu, A. and Kumar, S. (2006). A novel developmental system for the study of evolutionary design. In *Developmental Systems; AAAI Fall Symposium*.

Hamming, Richard, W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, 26(2):147–160.

Kumar, S. and Bentley, P. J., editors (2003). *On Growth, Form and Computers*. Elsevier Academic Press.

LaBarbera, M. (1989). Analysing body size as a factor in ecology and evolution. *Annual review of ecology and systematics*, 20(97-117).

Miller, Julian, F. (2004). Evolving a self-repairing, self-regulating, french flag organism. In et. al., K. D., editor, *Genetic and Evolutionary Computation – GECCO-2004, Part I*, volume 3102, pages 129–139.

Mitchell, M. (1998). *An Introduction to Genetic Algorithms*. Masachusetts Institute of Technology.

Streichert, F., Speith, C., Ulmer, H., and Zell, A. (2003). Evolving the ability of limited growth and self-repair for artificial embryos. In Banzhaf, W., editor, *ECAL*, volume 2801, pages 289–298.

Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media Inc.

Wolpert, L. (2002). *Principles of Development (2nd Ed.)*. Oxford University Press.

Wuensche, A. and Lesser, M. (1992). *The global dynamics of cellular automata*. Addison Wesley publishing company.