

Resilient Life: An Exploration of Perturbed Autopoietic Patterns in Conway's Game of Life

Arta Cika^{1*}, Elissa Cohen^{2*}, Germán Kruszewski^{3*†}, Luther Seet^{4*}, Patrick Steinmann^{5*} and Wenqian Yin^{6*}

¹ Dept. of Engineering Science, Oxford University, Oxford, UK

² Dept. of Economics, American University, Washington DC, USA

³ Naver Labs Europe, Grenoble, France ⁴ Non-Affiliated, Singapore

⁵ Biometris, Wageningen University & Research, Wageningen, The Netherlands

⁶ Dept. of Electrical & Electronic Engineering, The University of Hong Kong, Hong Kong, China
germank@gmail.com

Abstract

Complex systems can exhibit autopoiesis—a remarkable capability to reproduce or restore themselves to maintain existence and functionality. We explore the resilience of autopoietic patterns—their ability to recover from shocks or perturbations—in a simplified form in Conway's Game of Life. We subject a large number of autopoietic patterns in the Game of Life to various perturbations, and record their responses using multiple resilience metrics. Our results show that while resilience is rare, we are able to identify structural features improving patterns' resilience. We also draw several parallels between the resilience of patterns in the Game of Life to real-world complex systems. Our work may be useful both for improved searching for resilient patterns in the Game of Life, and for exploring resilience in complex systems.

Introduction

Living systems have been characterized as self-constructing networks of processes that maintain their own boundaries, a concept known as autopoiesis (Varela et al., 1974; Maturana and Varela, 1980). This capacity allows a system, such as the living cell, to persist throughout time, even if all of its individual components have been replaced. Nonetheless, an autopoietic system exposed to external perturbations would not survive for long without *resilience*, the ability to recover from or adapt to shocks or disruptions (Folke et al., 2004).

Here, following the characterization by Beer (2015) of emergent entities in Conway's Game of Life (Gardner, 1970) as autopoietic systems, we present an exploration of their response to perturbations. Conway's Game of Life (GoL) is a cellular automaton which exhibits complex emergent structures, serving as a computational environment to study the behaviour of simple systems when subjected to shocks. In particular, we focus on still lifes, which Beer (2015) identifies as autopoietic structures, together with oscillators and gliders. To study their resilience, we propose a variety of metrics and perturbations to assess whether they are able to maintain their form in response to the shocks, and identify underlying trends in their resilience.

* NB: authors listed in alphabetical order.

† Work done while the author was at Facebook AI.




Figure 1: Effect of additive perturbations to still life *Inflected Clips* (apgcode `xs32_4a9b8b96z259d1d96`), live cells in black. Adding a live cell at green locations yields the original pattern, making the still life resilient to these perturbations. It is not resilient to perturbations at the yellow locations, where adding a live cell does not yield the original pattern.

Conway's Game of Life is, in appearance, quite brittle. Minor perturbations can cause huge collapses. Understanding whether resilience to specific types of perturbations can emerge even in these difficult conditions could shed light on whether resilience is a universal property of computational systems.

Literature Review

Conway's Game of Life

Conway's Game of Life (Gardner, 1970) is an example of a cellular automaton. These are computational models composed of a grid of cells that can each occupy a finite or infinite number of states. In each generation, cells change or maintain their states depending upon (i) the states of their neighbors and (ii) a set of rules. Among cellular automata, Conway's Game of Life has attracted special attention for various reasons. First, despite being based on a mere handful of states and rules, intricate structures and behaviors emerge. These structures appear capable of endless variation and evolution. In this regard, the simulation is remarkably lifelike, and has attracted substantial atten-

tion from those working at the intersection of theoretical computer science and complex systems (Lindgren and Nordahl, 1990; Morita et al., 2002; Margolus, 1984; Adamatzky, 1998; Gots, 2003).

Conway's Game of Life is Turing-complete, meaning that it can be used to perform any computation –another possible hallmark of life (Cook, 2004; Wolfram, 2002; Mitchell, 2005). This may explain why it exhibits a wide range of interesting patterns, potentially including some that “detect” external perturbations and “repair” themselves (Goucher, 2015).

Rules of the Game of Life Cellular automata are defined by the sets of possible states for cells to occupy, the rules which govern evolution of cell states, and the cells' initial states. The Game of Life has a two-dimensional cell arrangement, with only two cell states—“dead” and “alive”—and four rules:

1. If a dead cell has three live neighbors, then it switches states from dead to live.
2. If a live cell has fewer than two neighbors, then it switches states from live to dead.
3. If a live cell has more than three live neighbors, then it switches states from live to dead.
4. If a live cell has two or three live neighbors, then it remains alive.

The GoL uses the Moore cell neighborhood definition (Packard and Wolfram, 1985), in which each cell has eight neighbors. Surprisingly, the four simple rules described above enable Conway's GoL to compute any computable function, as shown by Conway, Berlekamp, and Guy (Austin et al., 1982). An alternative two-rule set also fully defines the GoL (Gots, 2003), but is not commonly used.

Still Lives Within Conway's GoL, patterns with an evolutionary cycle of period one are called still lifes. These patterns' state configurations remain constant across time periods when left on their own, such that all dead cells remain dead, and all living cells remain alive. For example, the “block” pattern is a still life composed of 4 living cells disposed in a square arrangement. Applying GoL rules to compute the next generation yields the exact same pattern because all cells have exactly three neighbours, and thus they remain alive on the next generation, while no dead cell in its boundaries becomes alive because they all have no more than two living neighbours. Still lifes can be considered the residual of an experiment.

Autopoiesis

A system's viability is contingent upon exhibiting autopoietic behavior (Beer, 2015; Maturana and Varela, 1980).

First introduced by Chilean biologists Maturana and Varela (1980) to distinguish living from non-living systems, the term autopoiesis is derived from the Greek words for self (αὐτο- (auto-)) and production (ποίησις (poiesis)) (Mingers, 1991). Such autopoietic systems are comprised of collections of productive units which:

“(i) continuously regenerate and realize the network that produces them, and (ii) constitute the system as a distinguishable unity in the domain in which they exist” (Varela (1997, p. 75), as cited in Beer (2019)).

Previous work on autopoiesis in Conway's GoL has characterized the set of processes that regenerate specific still lifes, oscillators, and gliders, in the absence of perturbations (Beer, 2015). As shown, such systems “realize” their structure in perpetuity when left autonomous and in isolation. Recently, Beer (2019) examined the responses of glider patterns to environmental perturbations. Our analysis, instead, focuses on a large class of patterns, namely, still lifes, characterizing the resilience of each pattern in this class with respect to specific sets of perturbations.

Resilience

Real-world systems constantly experience shocks, perturbations and disruptions. The ability of a system to endure such shocks, while maintaining its form and function, is called resilience (Folke et al., 2004):

Resilience is the capacity of a system to absorb disturbance and reorganize while undergoing change so as to still retain essentially the same function, structure, identity, and feedbacks.

One can view resilience as the natural behavior of a system operating in a state space with attractors (Holling, 1996). The system is initialized somewhere in the state space, and naturally gravitates towards an attractor, such as a fixed point or limit cycle (Strogatz et al., 1994). Natural shocks and perturbations push the system away from the attractor. If the system remains in the original attractor's basin of attraction, the system will return there—otherwise it will find a new attractor. These two responses are commonly referred to as engineering and ecological resilience (Holling, 1996).

The resilience of systems is generally understood in relation to specific system functions (Holling, 1973; Albert et al., 2000). As such, any definition of resilience or robustness must specify both the relevant perturbation and the feature that may persist despite said perturbation (Jen, 2003). Thus, systems do not necessarily possess a general resilience, but instead may possess many specific resiliences to different shocks, disruptions, or other perturbations.

Method

To measure resilience in Conway's GoL, we focused on a specific class of autopoietic patterns. In particular, we restricted our exploration to still lifes. In the absence of perturbations, these patterns are structurally invariant across time, making them the simplest class of such autopoietic patterns. While still lifes can be construed as fixed points in the dynamics of Conway's GoL, it is unclear whether they are stable or unstable. We apply resilience theory on single and multiple basins of attraction (Holling, 1973) as an analogy to explore the behavior of these still lifes when exposed to perturbations.

Quantifying resilience

Studying the resilience of still lifes required us to make two methodological decisions, as follows.

Perturbations For our study, we first identified two distinct types of perturbations to test. The first perturbation adds at least one "live" cell at a position in the Moore neighbourhood of the pattern under examination, while the second perturbation removes (subtracts) at least one "live" cell from the interior of the pattern. As such, we say a structural pattern exhibits *additive resilience* when it returns to the previous life pattern after experiencing the first type of perturbation. If a pattern returns to its previous form after experiencing the second perturbation type, we say it exhibits *subtractive resilience*. We tested the addition of either one (**add one**) or two (**add two**) consecutive living cells to the pattern perimeter, and the subtraction of one (**sub one**) or two (**sub two**) consecutive living cells from within the pattern.

Similarity metrics We next outlined how to measure recovery. While ideally a pattern will always recover its exact original structure after experiencing a shock, it may also need to adapt its structure to return to a stable state. Thus, rather than only considering exact equality to measure resilience, we chose to use multiple similarity measures in order to assess the extent to which the original structure recovered its form. Furthermore, it can take multiple steps before a pattern converges back to a new form, if ever. However, we empirically found that the number of steps required for our metrics to converge would always be well below 100 iterations (more details in the Results Section). Thus, we quantify resilience ρ as the average of the similarities between the shapes resulting from simulating GoL for 100 iterations on each perturbed shape and the original one. More formally, for a given still life s , let $\mathcal{P}(s)$ be the set of all shapes resulting from applying a given type of perturbation \mathcal{P} (for example, $\mathcal{P} = \text{add one}$), σ be a similarity metric, and \mathcal{E}_k be the result of applying Conway's GoL rules for $k = 100$

iterations, then resilience is quantified as:

$$\rho = \frac{1}{|\mathcal{P}(s)|} \sum_{s' \in \mathcal{P}(s)} \sigma(s, \mathcal{E}_{100}(s'))$$

We further identified three variants of similarity to consider over binary vectors s_1 and s_2 representing the original and resulting patterns. Pairs of shapes were converted into same-dimensional binary vectors by taking the bounding box of the largest shape and interpreting living and dead cells as either 1 or 0, respectively.

Equality

$$\sigma(s_1, s_2) = \begin{cases} 1 & \text{if } s_1 = s_2 \\ 0 & \text{otherwise} \end{cases}$$

Inclusion

$$\sigma(s_1, s_2) = \begin{cases} 1 & \text{if } s_{1j} \leq s_{2j} \forall j \\ 0 & \text{otherwise} \end{cases}$$

Cosine

$$\sigma(s_1, s_2) = \frac{s_1 \cdot s_2}{\|s_1\| \|s_2\|}$$

The equality metric indicates whether the still life pattern remains identical in structure post-perturbation. The inclusion metric indicates whether the pattern absorbs, or rather includes, the perturbation into the resulting stable pattern. Lastly, our cosine metric allows us to measure the amount of overlap between the initial static pattern pre-perturbation and the ensuing pattern post-perturbation. Note that because all similarity metrics are bounded between 0 and 1, our resilience measure also lies within the same bounds.

Technical implementation Still life patterns used for our experiments came from the publicly available Catagolue census (Goucher, 2015) having population size between 4 and 42 living cells (for a total of 159 100 patterns). We also used the lifelib library (Goucher, 2017) to simulate Conway's GoL. The full set of experiments took about a day on a single-threaded Python implementation¹.

Structural properties

We also sought to understand what structural characteristics augment a pattern's resilience to the different perturbation types. Do some structural features help protect against one type of perturbation, but not another? To answer these questions we measured five traits for each pattern: density, population size, size of the perimeter, symmetry and number of connected components.

¹We make the code available at <https://gitlab.com/germank/resilient-life>

Density: The fraction of living cells within the bounding box of the pattern.

Population Size: The number of living cells in the pattern.

Size of Perimeter: The number of non-living cells that are at distance 1 (also in the diagonal direction) from a living cell of the pattern.

Symmetry: The similarity (as given by any of the measures above) between the pattern and its rotation or reflection. Here, we also applied the same three measures of similarity defined above. For instance, whether a shape is symmetric with respect to its x-axis is computed can be measured by whether the shape and its mirror image are exactly equivalent (by applying equality), or whether they are just similar (by applying the cosine measure).

Number of connected components: The number of contiguous regions of living cells. Note that every component must be not farther than one single cell apart from another for them to be considered part of the same pattern.

Results

After running our experiments, we ranked each still life by the three self-similarity resilience metrics. We observed the distribution of these three metrics for single-cell perturbations (Figures 2 and 3), and found an exponential decline of resilience at a given rank, suggesting a tendency of the system towards brittleness. The results from the two-cell perturbations (not shown) confirmed this tendency. We found particularly interesting that resilience tended towards a smooth distribution for all resilience metrics, even the equality measure, which had the most stringent criterion. This suggests resilience, as we’ve defined it, is not binary, but rather a graded trait for still lifes in the GoL.

We next examined the highest ranked individual patterns for the different resilience measures within each perturbation type (Table 1). For ease of presentation, resilience values (on a scale of 0 to 1) are presented as percentages (on a scale of 0 to 100). Patterns that ranked highest on one resilience metric within a perturbation category, rarely ranked highest on the other metrics. For example, we found the “inflected clips” pattern (the first pattern depicted in Table 1) ranked first for the *add one* perturbation under the equality measure with an average resilience of 21 percent but ranked 262nd under the inclusion measure and 2027th under the cosine measure. Figure 1 highlights how perturbing cells nearest to the “inflected clips” core allowed the pattern to recover a stable structure, whereas perturbing cells at its periphery induced the pattern to collapse. The points in which the pattern is resilient to perturbations can be explained by the rules. Since the dead cells within the “inflected clips” core are surrounded by sufficiently many live neighbours, perturbations from within the core are eliminated once faced

Pattern	Perturbation	Measure	Resilience	Rank
	add one	equality	21	1
		inclusion	21	262
		cosine	27	2027
	add one	equality	0	53 144
		inclusion	71	1
		cosine	78	1
	sub. one	equality	100	1
		inclusion	100	1
		cosine	100	1
	add two	equality	4	1
		inclusion	4	124
		cosine	11	20 166
	add two	equality	16	1 220
		inclusion	7	1
		cosine	0	1 274
	add two	equality	0	1 220
		inclusion	5	35
		cosine	29	1
	sub. two	equality	6	1
		inclusion	6	1
		cosine	11	86 067
	sub. two	equality	0	76
		inclusion	0	76
		cosine	67	1

Table 1: Patterns that are top-ranked according to at least one of the resilience measures for a given type of perturbation. All resilience values are percentages (i.e. scaled by 100). Ranks are computed out of 159 100 total patterns. In case of ties, the minimum rank is given.

with the overpopulation rule. In other words, since living cells adjacent to this resilient core have only two neighbours, they can withstand the appearance of a third without prompting a reaction. The “tub” pattern (the second pattern depicted in Table 1), in contrast to the “inflected clips” pattern, exhibited no resilience under the equality measure, but obtained the highest ranking under the inclusion and cosine measures with 71 percent and 78 percent resilience scores respectively. The resilience of the “tub” comes from its ability to incorporate or to adapt to the *add one* single cell perturbations. When a cell is added to the four dead corners, the cell can “come alive” and otherwise persist without disturbing the structural integrity of the “tub” pattern. Additionally, the “tub” is resilient to eight other *add one* perturbations, whereby it can evolve into one of the former configurations with a corner filled out. Only the “block” pattern scored a rank of one across all resilience measures in any of the perturbation categories. Specifically, for the *sub one* perturbation, we found that the “block” pattern was fully resilient to removing any of its cells. Since all living cells in the pattern have three neighbours, the removal of any sin-




Figure 2: Distribution of resilience results after *adding* a single cell to the pattern perimeter (*add one*) under the three considered similarity measures.




Figure 3: Distribution of resilience results after *subtracting* a cell from the pattern (*sub one*) under the three considered similarity measures.

gle live cell would automatically result in its regeneration in the subsequent time step. Finally, patterns that maximize resilience metrics for double-cell perturbations showed generally much lower resilience scores, and thus, they are only resilient to few specific positions.

In general, we wanted to understand whether there were some structural features of the patterns that could predict high resilience. For this, we computed the Pearson- r correlations between the structural features described in the previous section and the resilience observations for single cell perturbations. We focused on single cell perturbations, rather than also considering two-cell perturbations because the former showed the highest amount of variance. Results are reported in Table 2, where each percentage value denotes the Pearson correlation between specific resilience measure and structural feature. The number of connected components, the population size, and the pattern’s density come out as the most important predictors of resilience. First, only the number of connected components (*c.c.*) remained a steady positive predictor of resilience, regardless of the resilience metric or perturbation condition. Under most measures of resilience, the number of connected components was also the strongest predictor. Our “inflected clips” example in Figure 1 illustrated how having greater numbers of connected

components may facilitate additive resilience, where disconnected components may create an internal core that is protected from perturbations. The resilience observed in the subtractive case, by contrast, may be a byproduct of the fact that many composite still lifes are often composed of blocks among other shapes. That said, the strength and direction of the relationship between the other predictors and resilience depended on the type of perturbation and resilience metric considered. For example, while high density (*dens.*) positively predicted resilience under the equality metric when performing an additive perturbation, it negatively predicted resilience under the inclusion and cosine metrics. Additionally, under the subtractive perturbations, high density (*dens.*) positively predicted resilience under both the equality and inclusion metrics, but negatively predicted resilience under the cosine metric. Nevertheless, it is important to bear in mind that these are linear correlations, and may not be telling the full story. For example, when analyzing the effect of population size in the additive perturbation under equality we observed that the maximum resilience for each population size tends to grow (with some roughness) until size 32 where it maxes out, but then starts to slowly decline.

Next we analyzed how related our own definitions of resilience were with each other. For this, we computed the




Figure 4: Pearson- r correlation coefficients between different types of perturbations and measures.

correlations between the resilience values obtained from different types of perturbations and different types of measures (Figure 4). We found that the resilience of a pattern to a given type of perturbation does not transfer to other perturbation types when we restrict to measuring resilience through equality (Figure 4c). This supports the idea that systems have specific resiliences to unique disruptions, rather than a general resilience (Carpenter et al., 2001). However, when we adopted a relaxed notion of recovery as given by the cosine measure (Figure 4d), we saw a weak, yet significant, correlation between different types of perturbations. Thus, under this particular definition of recovery, we can see a somewhat more generalized notion of resilience. Conversely, the correlations between our measures also change as a function of the perturbation that we consider. If we restrict to the “add one” perturbation (Figure 4a), the inclusion and the cosine measures tend to highlight the resilience of the same patterns while equality behaves in a more idiosyncratic way. On the other hand, when we focus on the “subtract one” perturbation (Figure 4b) equality and inclusion behave almost in the same way, and differently from cosine.

Last, but not least, we studied the length of the transients to see how many steps these resilient patterns take to recover from a perturbation (up to a maximum of 100). We observed that for both the inclusion and the equality measures, the transient lengths were no longer than 2 steps. For this rea-

son, we restricted the following analysis to just the cosine measure, which displayed more variation. Results are displayed in Figure 5. For small perturbations (Figures 5a and 5b), most transient lengths are concentrated at the lower end of the spectrum. Nonetheless, there is a graded distribution with a decreasing number of patterns having longer transient lengths. Interestingly, when a larger perturbation is applied to the system, the transient lengths increase, as seen in Figures 5c and 5d, while still maintaining a graded distribution.

Discussion

Additive vs. subtractive resilience

We found virtually no still lifes in the GoL which were resilient to both additive and subtractive perturbations. Our work reinforces previous research examining resilience of real-world systems wherein resilience must be considered in specific contexts (Carpenter et al., 2001), and with specific actors and stakeholders in mind (Jones, 2018). That a system exhibits specific resiliences (vs. a general resilience) even in abstract environments such as Conway’s GoL suggests resilience may be a fundamental property of complex systems. Identifying such a property has far-reaching consequences for the study of real-world systems, as it would invalidate a large body of work quantifying “the” resilience of such systems.




Figure 5: Distribution of length of transients when using the cosine measure.

General trends in still life resilience

The Game of Life is highly symmetrical by nature, as all rules function equally in all four cardinal directions. However, symmetry, as measured by the similarity between a shape and its mirror image (both along the x axis and the y axis) did not necessarily yield strong predictions of resilience. Although, there did appear to be a weakly positive relationship between symmetry (as measured by the cosine) and resilience under the equality metric for both perturbation types. This suggests that symmetry is not an crucial factor in a still life’s resilience.

Additionally, a fundamental constraint of the Game of Life is that still life patterns have a maximum density of 0.5 within their boundaries (Capcarrère and Sipper, 2001; Elkies, 1999; Chu et al., 2009; Chu and Stuckey, 2012). This represents a tipping point in the population dynamics of the still life, and follows quite naturally from the rules—as any density above 0.5 creates a cascade of cell deaths, ultimately wiping out all cells in the pattern (Griffeath and Moore, 2003). Nevertheless, we found still lifes with densities approaching 0.5 were more resilient to perturbations than those with sparser densities. It appears that still lifes existing

close to the threshold of their ability to maintain structure—this might be considered their autopoietic limit—are more resilient. This observation was mentioned by Holling (1996), who described how humans operate at body temperatures close to lethal levels—allowing them to mobilize more energy for expanding their living environments or responding to threats. Here we again see an interesting parallel between the abstract dynamics of the Game of Life, and real-life complex systems.

We note here that our density measure is not fully consistent with previous work on the density of still lifes, having used an alternative method for perimeter construction (Elkies, 1999). However, we find support for our method in current autopoiesis discussions regarding what constitutes living system spaces (Villalobos and Razeto-Barry, 2019; Harvey, 2019). As such, we believe the underlying trends are strong enough to suggest a general pattern regarding density, stability and resilience in autopoietic systems.

Conclusions and Future Work

We explored the degree to which still lifes—fixed-points in Conway’s GoL—are resilient to one- and two-celled pertur-

Additive Perturbation					
Equality		Inclusion		Cosine	
Feature	r	Feature	r	Feature	r
pop.	23	c.c.	9	c.c.	17
c.c.	20	dens.	-8	dens.	-10
dens.	19	sx-cos	5	pop.	-7
sx-cos	17	pop.	-4	m-y	3
sy-cos	12	sx	3	sx-cos	2

Subtractive Perturbation					
Equality		Inclusion		Cosine	
Feature	r	Feature	r	Feature	r
c.c.	35	c.c.	36	dens.	-16
pop.	18	pop.	18	c.c.	15
dens.	16	dens.	16	pop.	-14
cw-cos	9	cw-cos	9	cw-cos	-10
ccw-cos	9	ccw-cos	9	ccw	-10

Table 2: Pearson-r correlations (in percentage) between resilience for single-cell perturbations and structural features for the top-5 correlated or anti-correlated features. Key: *pop*=population; *c.c.*=number of connected components; *dens.*=density; *sx/sy*=symmetry on x/y axis; *sx/sy-cos*=symmetry on x/y axis as given by the cosine between the shape and its mirror image; *cw/ccw*=invariance to clockwise or counter clockwise rotation; *cw/ccw-cos*=invariance to clockwise or counter clockwise rotation measured as given by the cosine between the shape and its rotation.

bations. We found certain still lifes served as stronger system *attractors* (Holling, 1996), in that they exhibited higher propensities to return to stable states post-perturbation. Ultimately, we observed a graded, exponentially decaying, distribution of resilience among patterns. This suggests that we may find patterns that are increasingly more resilient to the studied perturbations when examining (exponentially) larger samples. Furthermore, this graded surface may be just what an evolutionary process needs to discover ever more resilient shapes, even though more work characterizing the smoothness of this surface would be needed to support that idea.

While this exponentially decaying trend of resilience in patterns was generally stable across different conditions, the particular definition of resilience given by the used perturbation and the way to measure recovery, was prone to highlight different patterns. This finding supports the idea that no universal notion of resilience exists, but rather, resilience is by nature context conditional.

Our analysis also evidenced some structural features of resilient patterns. Population, the number of connected components, and density served as the strongest general important predictors of resilience; however, for only the number of connected components was the direction of the relationship consistent. Investigations using these structural insights could help lead to the discovery of more resilient patterns in

the future are left for future work. We conjecture that there may exist other not yet discovered patterns which maximize our resilience metrics with respect to the studied perturbations.

While we restricted our analysis to still lifes in this study, future work will extend our exploration to other types of autopoietic patterns in the GoL such as oscillators and gliders (Beer, 2015). Furthermore, the perturbations that we explored were two simple ones; however, they may not arise naturally in the dynamics of the GoL. Thus, a natural extension to our current analysis is to test more “naturalistic” perturbations, such as the impact of a glider hitting a still life. Indeed, some patterns with the name of “eaters” (Griffeath and Moore, 2003) have already been shown to be resilient to these types of perturbations. Thus, a more systematic analysis, such as the one performed in our study could help verify whether our resilience conclusions also hold in more complex scenarios.

Finally, exploring non-resilience is important for its own sake. For instance, through our inclusion and cosine measures of resilience, we found perturbing a still life can result in it re-stabilizing as a transformed still life pattern. Future work will explore how different perturbations may transform still lifes into oscillators, gliders, neither, or some combinations thereof.

To conclude, we have presented a systematic analysis of the resilience of still lifes in Conway’s Game of Life. We have observed that while it is a rare property, increasingly higher levels of resilience are observed. As argued above, this graded landscape may provide evolution with just the right stepping stones that it needs in order to discover ever more resilient forms of organization. We remain cautiously optimistic that, even if rare in a universe which is, in appearance, as brittle as Conway’s Game of Life, resilience can be found.

Acknowledgments

We would like to thank the anonymous reviewers for their feedback, Chris Moore for early discussions, and all the SFI Complex Systems Summer School 2019 staff and participants. A special thanks to John Conway for creating a complex universe that keeps exciting our curiosity.

References

- Adamatzky, A. (1998). Universal dynamical computation in multidimensional excitable lattices. *International Journal of Theoretical Physics*, 37(12):3069–3108.
- Albert, R., Jeong, H., and Barabási, A.-L. (2000). Error and attack tolerance of complex networks. *Nature*, 406(6794):378.
- Austin, A. K., Berlekamp, E. R., Conway, J. H., and Guy, R. K. (1982). *Winning Ways For Your Mathematical Plays*. Academic Press.

- Beer, R. D. (2015). Characterizing autopoiesis in the Game of Life. *Artificial Life*, 21(1):1–19.
- Beer, R. D. (2019). Bittorio revisited: structural coupling in the Game of Life. *Adaptive Behavior*.
- Capcarrère, M. S. and Sipper, M. (2001). Necessary conditions for density classification by cellular automata. *Physical Review E*, 64(3):036113.
- Carpenter, S., Walker, B., Anderies, J. M., and Abel, N. (2001). From metaphor to measurement: Resilience of what to what? *Ecosystems*, 4(8):765–781.
- Chu, G. and Stuckey, P. J. (2012). A complete solution to the maximum density still life problem. *Artificial Intelligence*, 184-185:1–16.
- Chu, G., Stuckey, P. J., and de la Banda, M. G. (2009). Using relaxations in maximum density still life. In Gent, I. P., editor, *Principles and Practice of Constraint Programming*, pages 258–273. Springer.
- Cook, M. (2004). Universality in elementary cellular automata. *Complex systems*, 15(1):1–40.
- Elkies, N. D. (1999). The still-life density problem and its generalizations. *arXiv:math/9905194*.
- Folke, C., Carpenter, S., Walker, B., Scheffer, M., Elmqvist, T., Gunderson, L., and Holling, C. S. (2004). Regime shifts, resilience, and biodiversity in ecosystem management. *Annual Review of Ecology, Evolution, and Systematics*, 35:557–581.
- Gardner, M. (1970). Mathematical games - the fantastic combinations of John Conway's new solitaire game "Life". *Scientific American*, 223:120–123.
- Gotts, N. M. (2003). Self-organized construction in sparse random arrays of Conway's Game of Life. *New Constructions in Cellular Automata*. Oxford University Press, New York, pages 1–53.
- Goucher, A. P. (2015). Catagolue. <https://catagolue.appspot.com/home>. Accessed: 2019-10-25.
- Goucher, A. P. (2017). lifelib. <https://pypi.org/project/lifelib>. Accessed: 2019-10-22.
- Griffeath, D. and Moore, C. (2003). *New constructions in cellular automata*. Oxford University Press.
- Harvey, I. (2019). Habeas corpus: The ins and outs of autopoiesis. *Adaptive Behavior*.
- Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4(1):1–23.
- Holling, C. S. (1996). Engineering resilience versus ecological resilience. *Engineering within ecological constraints*, 31(1996):32.
- Jen, E. (2003). Stable or robust? what's the difference? *Complexity*, 8(3):12–18.
- Jones, L. (2018). Resilience isn't the same for all: Comparing subjective and objective approaches to resilience measurement. *Wiley Interdisciplinary Reviews: Climate Change*, 10(1):e552.
- Lindgren, K. and Nordahl, M. G. (1990). Universal computation in simple one-dimensional cellular automata. *Complex Systems*, 4(3):299–318.
- Margolus, N. (1984). Physics-like models of computation. *Physica D: Nonlinear Phenomena*, 10(1-2):81–95.
- Maturana, H. R. and Varela, F. (1980). *Autopoiesis and cognition: the realization of the living*. D. Reidel.
- Mingers, J. (1991). The cognitive theories of Maturana and Varela. *Systems Practice*, 4:319–338.
- Mitchell, M. (2005). *Computation in Cellular Automata: A Selected Review*, pages 95–140. Wiley.
- Morita, K., Tojima, Y., Imai, K., and Ogiro, T. (2002). Universal computing in reversible and number-conserving two-dimensional cellular spaces. In *Collision-Based Computing*, pages 161–199. Springer.
- Packard, N. H. and Wolfram, S. (1985). Two-dimensional cellular automata. *Journal of Statistical Physics*, 38:901–946.
- Strogatz, S., Friedman, M., Mallinckrodt, A. J., and McKay, S. (1994). Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering. *Computers in Physics*, 8(5):532.
- Varela, F. G., Maturana, H. R., and Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *Biosystems*, 5(4):187–196.
- Varela, F. J. (1997). Patterns of life: Intertwining identity and cognition. *Brain and Cognition*, 34(1):72–87.
- Villalobos, M. and Razeto-Barry, P. (2019). Are living beings extended autopoietic systems? An embodied reply. *Adaptive Behavior*.
- Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.