The Role of Abiotic Parameters in the Promotion of Egalitarian Major Evolutionary Transitions

Max Foreback[1,2,3], Sydney Leither[1,2,3], David A. Baum[4,5] and Emily Dolson [1,2,3]

[1]Computer Science and Engineering, Michigan State University, East Lansing, MI
[2]Ecology, Evolutionary Biology, and Behavior, Michigan State University, East Lansing, MI
[3]BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI
[4]Department of Botany, University of Wisconsin, Madison, WI
[5]Wisconsin Institute for Discovery, University of Wisconsin, Madison, WI
foreba10@msu.edu

Abstract

The problem of identifying conditions that enable major evolutionary transitions, in which distinct units come together to form a new higher level unit, is a complex and difficult topic spanning many disciplines. Here, we approach this problem from the perspective of the origin of life, which allows us to make the simplifying assumption that the lower-level units are not also evolving. This assumption lets us focus on identifying environmental factors that promote egalitarian major transitions in general and the origin of life specifically. To study this question, we build a simple artificial ecology model. We quantify major-transition-like dynamics using a maximum likelihood approach and a set of null models predicting the behavior of our system under various dynamics. Ultimately, we find that, even in a maximally simple artificial ecology model, we are able to observe evidence of community-level selection and thus the beginnings of a major evolutionary transition. The regions of parameter space that promote community-level selection vary based on species interactions but we observe consistent trends.

Introduction

From the unification of prokaryotic cells to form the first eukaryotes to the evolution of eusocial insects, major transitions have played a vital role in the complexification of life on Earth (Smith and Szathmary, 1997; Herron, 2021). A major transition occurs when a new biological unit of organization begins acting as an evolutionary individual. Indeed, the origin of life itself can be thought of as a major transition, in which an ecological community (composed of interacting autocatalytic chemical cycles (Smith and Szathmary, 1997; Queller, 1997)) transitioned into a self-replicating unit. This raises a general question: what factors predispose an ecological community to become an emergent individual rather than merely a sum of its component species? In subsequent major transitions, the formation of a higher-level unit likely involved evolution of the species that combined to form a higher-level unit. However, based on the hypothesis that prebiotic chemical dynamics were purely ecological, such coevolution would have been impossible prior to the origin of life. Thus, the first major transition must have been achieved by only two factors working together: 1) the dynamics of the abiotic environment, and 2) the properties of the (non-evolving) ecological community. Here, we will attempt to understand the abiotic conditions that promote a transition from ecological to adaptive dynamics; for an investigation of the properties of the ecological community, see Leither et al. (2023).

The role of autocatalytic cycles, in which the products of a chemical reaction promote their own formation, has long been emphasized in studies of the origin of life (Kauffman, 1995). It has been shown that these autocatalytic cycles can emerge even at only moderate levels of catalysis (Hordijk et al., 2010) or even without any explicit catalysis as a result of the topology of chemical reaction networks (Blokhuis et al., 2020; Peng et al., 2022). While significant work has focused on individual autocatalytic cycles, less is known about the ecological dynamics of communities of interacting autocatalytic cycles. It has been shown that pairs of autocatalytic cycles can interact similarly to pairs of biological species, via competition, predation/parasitism, mutualism, etc. (Peng et al., 2020). It has also been shown that chemical ecosystems can manifest dynamics similar to ecological succession (Peng et al., 2023).

This perspective suggests that, rather than explicitly simulating complex chemical reaction networks, we might be able to get a handle on the first major transition at the origin of life, by focusing strictly on the ecological dynamics among pairs of autocatalytic cycles. In addition to elucidating the origin of life, such an abstract ecological approach may allow us to draw general conclusions about the conditions necessary to promote major transitions as a whole.

Major transitions can be split into two main categories, egalitarian major transitions, and fraternal major transitions. The difference comes from how the pre-transition units are related - in egalitarian transitions these units are unlike, while in fraternal transitions they are derived from the same parental unit (Queller, 1997). While significant progress has been made in recent years towards understanding fraternal major transitions (Moreno and Ofria, 2019; Goldsby et al.,

2020) the origin of life would likely have been an egalitarian major transition since parent-offspring pairs would not yet have existed.

Studying major transitions is difficult due to their slow speed of occurrence and rarity in nature. Digital systems provide a unique advantage in that experiments are fast, repeatable, inexpensive, and provide perfect information. Investigating major transitions in computational systems has already proven valuable for fraternal transitions (Moreno and Ofria, 2019; Goldsby et al., 2020). Here, we begin the process of doing the same for egalitarian major transitions, specifically those modelled after the origin of life.

Many abiotic conditions are hypothesized to have been important in facilitating the origin of life, most notably the presence of spatial structure and/or compartmentalization (Higgs and Lehman, 2015). Fundamentally, for such a transition to have happened, there must have been some factor that gave some kind of advantage to local communities that were jointly successful. Many properties of a world could create this circumstance, but here we attempt to bring it about with the simplest environment possible. Thus, we conduct this first round of experiments in a simple meta-community model governed by Lotka-Volterra dynamics (Malcai et al., 2002; Leibold and Chase, 2017). In this model, diffusion from locations with high biomass causes communities that are collectively capable of rapid growth to co-propagate to neighboring communities.

Methods

In order to examine egalitarian major transitions, we have constructed an artificial ecological simulation. All code used for this project is open source (Foreback et al., 2023a). This artificial ecology simulates the growth of different species within a world over a certain period of time. While many artificial life simulations define individuals that have discrete life and reproductive cycles, we do not. Our species are defined only by their interactions with each other - one individual is no different from another individual of the same species. Groups of these species make up what we call communities. Each community is a unique combination of species where only the presence or absence of the species is considered, with a species being present if it exists at any level in the community regardless of abundance. This definition results in 2^n possible communities, where n is the number of species in the world.

Interactions among species are controlled by a square, non-symmetric community interaction matrix. This matrix contains values for every possible species interaction and is of size $n \times n$Each species has a row, which represents how other species impact it, and a column, which represents its effects on other species. The diagonal encodes the species' intrinsic growth rate. The matrix is non-symmetric, as interactions do not need to be equivalent in both directions - if species A promotes the growth of species B, species

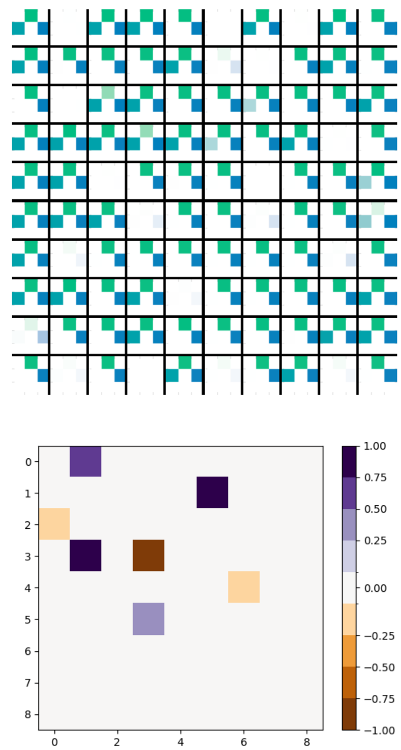

Figure 1: An example world (top) and interaction matrix (bottom). Each position in the world has a population count for each species 0-8. In this world, there are two similar communities that dominate. The first contains species 1, 3, and 5, and the second contains only species 1 and 5. The former community does not completely dominate the world because species 3 represses itself, and will eventually drive itself to extinction if seeded into a cell without other species that support it. The darker the color in the cell, the more abundant that species is.

B could be neutral or even detrimental to the growth of species A. The values within the matrix are encoded as floating point values in the range of $[-1, 1]$. The more positive the value, the more beneficial the interaction for the row species; negative values encode detrimental effects on the row species. A value near zero means that the row species is minimally impacted by the species represented by the column. The growth of any particular species is determined by its intrinsic rate of growth, adjusted based on the local abundance of all other species, weighted by their beneficial or detrimental interactions. Mathematically, the growth of a given species m in a single cell for one time-step is:

$$m_{t+1} = \sum_{i=0}^{n} s_i I_{mi}$$

where s is a vector containing an ordered count of all species in that cell at time t, and I is the interaction matrix.

The world is a square grid of cells that synchronously updates species counts a certain number of times before terminating. Cells on the edge of the world will wrap around to the opposite side (i.e. the world is toroidal). An example visualization of the world can be seen in Figure 1. Each cell tracks the number of each species contained within it and begins completely empty. Species counts are tracked as continuous values, as is typical in generalized Lotka-Volterra models. Every time step, the world updates according to the current communities in it and their compositions. Within a cell, individuals who have more positive interactions acting on them than negative interactions will increase in number over time. Similarly, individuals who have more negative interactions in their cell will decrease in number, possibly going locally extinct if their count reaches zero.

In addition to community interactions, three abiotic parameters can greatly change the developmental trajectory of the world. First is the *diffusion rate*, which ranges from 0 to 1 and dictates how individuals can move through the world. It controls the proportion of individuals in any given cell that will leave that cell and go to one of the neighboring cells. One quarter of this amount will diffuse in each direction. Varying the diffusion rate impacts how quickly species can spread and dominate the world. In addition to diffusion, there is a probability that exactly one individual from a randomly selected species will be added into any given cell. The parameter that controls this probability is called the *seeding rate*, and represents an event like immigration from a global species pool. Varying the seeding rate can also change which species dominate the world. For example, a very high seeding rate will quickly introduce all species to the world, driving out species that are commonly negatively impacted before they have a chance to establish. For the purposes of this paper, we bound seeding within the range $[0, .1]$. Finally, the *clear rate* dictates the probability that all species in a cell will go extinct, and represents ecological

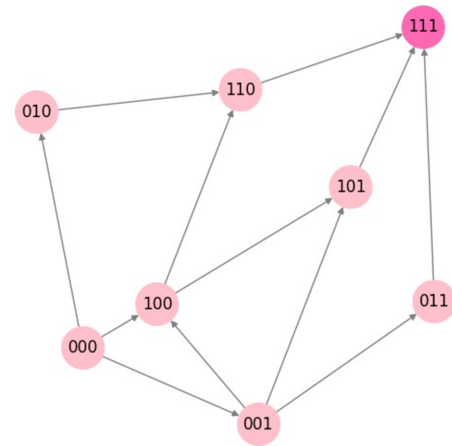

Figure 2: A simple assembly graph for three species. Each community composition is represented by a bit string. A one indicates the species in that position is present, a zero indicates its absence. The sink node here represents all three species coexisting. Since there is no invasion that can destabilize this community, we expect the entire world to converge to it over time.

disturbance. The clear rate is hypothesized to be important because the resulting empty cells tend to be colonized by diffusion of all species that are present in the neighboring cells, which might allow compositions that sustain high total counts to (imperfectly) replicate into cleared neighbors. The clear rate ranges from $[0, 1]$ where 1 means that every cell is cleared every time-step, and 0 means that cells are never cleared. It is these abiotic parameters that we are interested in investigating under different matrix structures, in the hopes of determining whether some parameter ranges promote community level adaptation.

Measuring Adaptation

One approach to identifying major evolutionary transitions is to observe the extent to which fitness is increasing over time across different levels of biological organization (e.g. the "individual" vs. the "community") . Although the theory of multi-level selection suggests that we may see increases at multiple levels, the level at which fitness increases the most can be seen as the level that is acting most like a Darwinian individual (Ratcliff et al., 2015). In the context of our system, if that level is the level of the community, then we have evidence of community-level selection.

A downside to this approach is that it requires being able to quantify fitness, a difficult task, especially when the "individuals" that fitness is acting on change, as they do during a major transition. Our approach is to deploy five different measures of adaptation, each of which could be seen as a fitness proxy. The metrics are calculated for every cell in the world and are as follows:

Biomass
: The total number of individuals present in the cell.

Growth Rate
: How quickly the total biomass in the cell will grow given the current composition and interactions.

Heredity
: The ability of a community to convert formerly empty cells to a composition similar to it.

Invasion Ability
: How quickly a community can traverse an empty world without seeding or clearing.

Resilience
: How resistant the current community is to an invasion.

For the purposes of this work, we focus only on biomass and growth rate as our fitness proxies for two reasons: 1) these are the properties that selection appears to be acting most strongly on in our system, and 2) we have previously found that the interaction matrices we are using in this work result in adaptation to maximize these properties (Foreback et al., 2023b).

Increases in these measures over time indicate that an increase in community fitness is occurring, but this does not necessarily mean that community-level adaptation is occurring. Other ecological dynamics, such as community assembly or succession, could be causing increases in fitness by coincidence. To identify increases in fitness that are most likely the result of community-level adaptive dynamics, we need to compare the expected behavior of the system under ecological vs. adaptive dynamics, as summarized in state-transition graphs describing the possible behavior of a given cell in the world. The graphs are directed, with each node representing a stable community, and each edge representing a transition between stable communities. We define stable communities as those that, when left unperturbed, will not drive any of their member species to extinction. In order for a community composed of N species to be considered stable, when that community is isolated from the stochastic processes of the world (mainly diffusion and seeding) all N species must be able to coexist at some level. If any are driven extinct over time, the community is not stable.

We construct two types of graph for each run of the simulation, one for the case of pure community assembly, and one representing adaptive dynamics optimizing for our measures of fitness. The former is called the community assembly graph, and contains all stable communities that can exist in the current world. We can determine which transitions between stable communities are possible by simulating the effect that possible invasions have on stable communities. If an invasion by a particular species causes a stable community to transition to another, an edge is added from the starting community to the ending community. Nodes that have an out degree of zero represent stable communities that cannot be invaded by any species, and are called sink nodes. Under pure community assembly dynamics, the development of communities would be equivalent to a random walk on this graph. Over a long enough time frame, we would expect the communities in the world to converge only to communities represented by sink nodes. There are scenarios in which this graph will have no sink nodes, i.e. there are no community compositions that are resilient to all invasions. In this case it is possible for a subset of nodes to develop that act collectively as sink nodes, as they only have transitions to each other. An example assembly graph can be seen in Figure 2.

The second type of graph, called a fitness landscape graph, is constructed from the assembly graph. It has the same structure as the community assembly graph, with nodes for stable communities and edges for possible transitions between communities. However, when calculating the fitness landscape graphs, we remove edges that lead from a community with higher fitness to one with lower fitness, as these transitions should be unlikely under purely adaptive dynamics. A separate fitness graph is created for each fitness proxy used, in this case yielding a biomass fitness landscape graph and a growth rate fitness landscape graph. As an example, in the growth rate fitness graph, a community composition that could be invaded and stably transition to another community would only retain its edge if the new community had the same or greater growth rate than the original community. Thus, sink nodes in the fitness landscape graph represent local fitness optima. Note that these fitness landscapes are a slight oversimplification of actual adaptive dynamics (in which deleterious mutations are possible).

We can utilize these two types of graphs – fitness and assembly – to understand what we would expect to happen to the world on average under different dynamics. A random walk along the assembly graph would be our expectation under the scenario in which ecological succession is the main driving force. A random walk along a fitness graph, however, represents a world in which community-level selection on that given fitness proxy was the primary dynamic (i.e. adaptive dynamics). Since both ecological and adaptive dynamics are likely play some role, our goal with this analysis is to determine which factor is dominating the dynamics we observe. To do so, we examine the final state of the world and record the distribution of community compositions. If this set of communities is more likely to result from a random walk on one of the fitness graphs than on the community assembly graph, then we have evidence of adaptive dynamics.

To calculate the probability of arriving at different nodes via random walks on our graphs we utilize the PageRank algorithm and a maximum likelihood statistical approach.

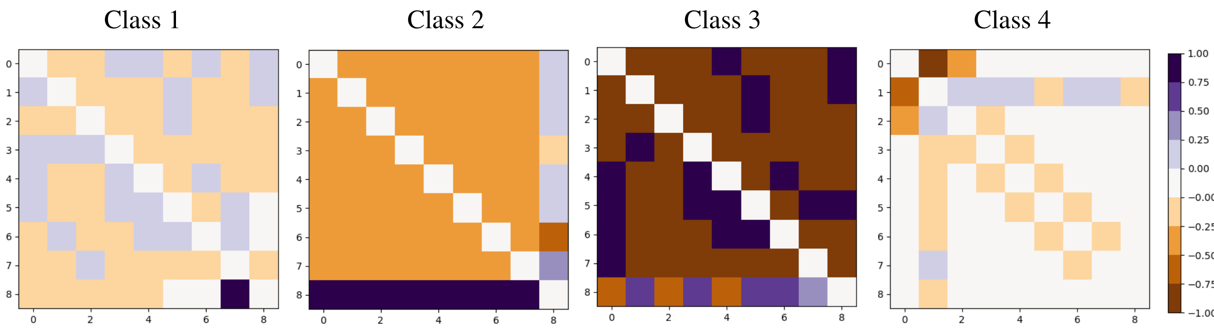

Figure 3: Structures for the four classes of matrices found to promote adaptive dynamics. Note that the genetic algorithm used to generate these structures did not allow for a species to have an intrinsic growth rate, which is why the diagonals are empty.

Each node on the graph, and therefore each stable community that could possibly arise in the world, will be given a PageRank score that reflects the probability a simulated run would end up there over a long period of time. These PageRank values can then be used along with the final world state to calculate the likelihood that we would have observed the resulting set of communities under either assembly dynamics (the community assembly graph PageRank scores) or adaptive dynamics (fitness graph PageRank scores). The likelihood of community composition under community assembly is given by the following equation:

$$L(assembly) = \sum_{i=0}^{n} p_i(PageRank_A(community_i))$$

Similarly for adaptive dynamics:

$$L(adaptive) = \sum_{i=0}^{n} p_i(PageRank_F(community_i))$$

Here n is the number of final communities and p is what proportion that final community was of the whole world. $PageRank_F$ is the given community's PageRank score in the given fitness graph, and $PageRank_A$ is the given community's PageRank score in the assembly graph. To aggregate these score into a single score for each fitness measure, we take the difference between them,

$$AdaptivenessScore = L(adaptive) - L(assembly)$$

We use the difference rather than the ratio, which is more commonly used for likelihoods, for ease of data visualization (the likelihoods tend to create extreme values). This leaves us with two adaptability scores, one for each used fitness proxy. Since any of the fitness proxies could potentially represent the characteristics being acted on by community level selection, a high value of either of the scores is potential evidence of adaptive dynamics.

Exploring adaptive worlds

Choosing which species interaction matrices to explore is an important decision, as the conditions under which adaptive dynamics occur are largely unknown. To constrain our search space to a reasonable size, we explore worlds which have nine species. As a result we have 2^9 possible communities that can form, only some of which are stable. A community interaction matrix for nine species will have 81 entries of continuous values, which makes searching for the combinations that manifest community adaptation using enumeration computationally expensive. When combined with the fact that the three abiotic parameters may change the results of the world under any given matrix, there is a vast parameter space to consider, even when constrained to 9 species.

In order to limit our search to a feasible space, we examine the effects of the abiotic parameters on four different classes of matrices that we previously found to display adaptive dynamics, shown in Figure 3. These matrix classes were discovered via a genetic algorithm based search of the parameter space for conditions that promote adaptive dynamics (Foreback et al., 2023b). While that work identified some parameter ranges that promote adaptive dynamics for these matrices, it did not conduct a systematic analysis on how the abiotic parameters influence the emergence of adaptive dynamics. Here, we perform an extensive search of parameter combinations for all four matrix structures by running simulations for each condition and examining the final set of communities that develop. The examined matrix classes are known to perform well for scores of biomass and growth rate.

Results

Proof-of-concept experiment

To confirm that our adaptiveness metrics are capturing the properties we want, we performed a proof-of-concept experiment with diffusion set to 0. Worlds in which there is no diffusion represent an extreme scenario where species are unable to traverse the world. Consequently, each cell in the meta-community should be governed purely by ecological assembly dynamics, as the only processes affecting species in a cell are seeding and clearing events.

We can also generate worlds governed purely by adaptive

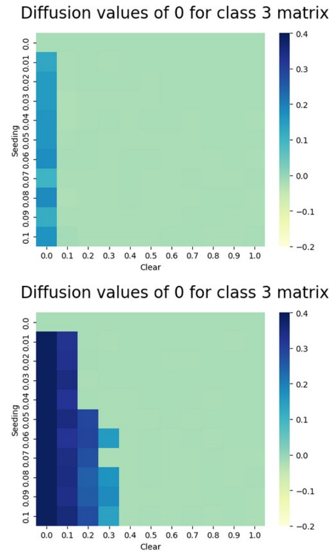

Figure 4: Heatmaps of adaptiveness scores when diffusion is set to 0. Top: A world in which group level reproduction was disabled. Bottom: A world in which group level reproduction was enabled.

dynamics by introducing the possibility for group-level reproduction. Under this condition, once the total biomass of a cell reaches a certain threshold, we seed all of that cell's member species into an adjacent cell. In the absence of diffusion, community-level reproduction is the only way to spread to new cells, meaning that each cell should effectively function as a Darwinian individual.

We tested both of these conditions. While we did see some communities classified as "adaptive" under purely ecological assembly dynamics, we saw dramatically higher adaptiveness under the purely adaptive condition. These results illustrate that, while our metrics are imperfect due to their probabilistic nature, by and large they do capture the dynamics of interest (see Figure 4).

Parameter space exploration

To examine the results of different abiotic parameters on each class of matrix systematically, we run our model over the full range of possible values. We increment diffusion and clearing probability by .1, and seeding by .01, and test all possible combinations for our matrix classes. This setup results in 110 combinations of diffusion and seeding for each increment of clear probability and a total of 1210 combinations for each matrix class. Seeding values of zero are not included in the analysis, as worlds that are never seeded cannot develop communities. For each permutation of parameters, we run five simulations and average their scores, in order to minimize the noise that is inherent in our simulation.

All worlds that develop a set of communities with positive biomass and growth rate scores are potentially undergoing major transitions, as the dynamics of those worlds were more likely to be produced by community level adaptation than by assembly dynamics. However, higher adaptiveness scores are stronger evidence of adaptation. As such, we differentiate between relatively high scores and scores that are only slightly above zero. In general we see two dominant communities under each matrix structure, one that is adaptive (with a score of greater than .1) and one that is non-adaptive (with a score of zero or slightly negative). In some worlds there is a spectrum of compositions that can develop with communities that are similar, but not identical, vying for control. Exactly which community dominates the world depends on the world's abiotic parameters. Below, we discuss the results for each matrix class in turn.

Class 1 Matrices

Class one matrices are characterized by their large numbers of small magnitude weights, both in the positive and negative direction. These small weights have implications for ecological stability in modeled food web interaction networks and have been shown to promote long term community stability McCann et al. (1998). Under conditions of high seeding, low clearing, and diffusion of around .5, class 1 matrices tend to result in the appearance of an adaptive community of four species. Adaptive communities developed most often with a clear probability of zero: of the 110 possible parameter combinations for zero clear, 101 produced an adaptive community. Adaptive communities occurred more rarely for clearing values of .1 and .2, and almost never for higher clearing values. These values of clear probability were also associated with low values of diffusion, typically .1 or .2.

Class 2 Matrices

Matrices of this class produce only one adaptive community due to their strict structure. This adaptive community is composed of species 8 and species 7 working together to dominate all other species. All previous experiments under this class contained the typical high seeding and low clear, but also had low diffusion, less than .4. Our results show that this matrix class is actually extremely resilient to many different values of diffusion and seeding for clear values up to .5. Of the 1210 possible parameter combinations, 600 produced the adaptive community. Of these, 392 occurred for clearing values in the [0, .4] range. As expected, high values of clear probability tended to produce low adaptation in most scenarios; though some worlds with both very high diffusion and seeding were able to occasionally produce an adaptive community in high clearing values. Across all values of clear probability except zero, a diffusion of zero failed to produce any adaptive communities.

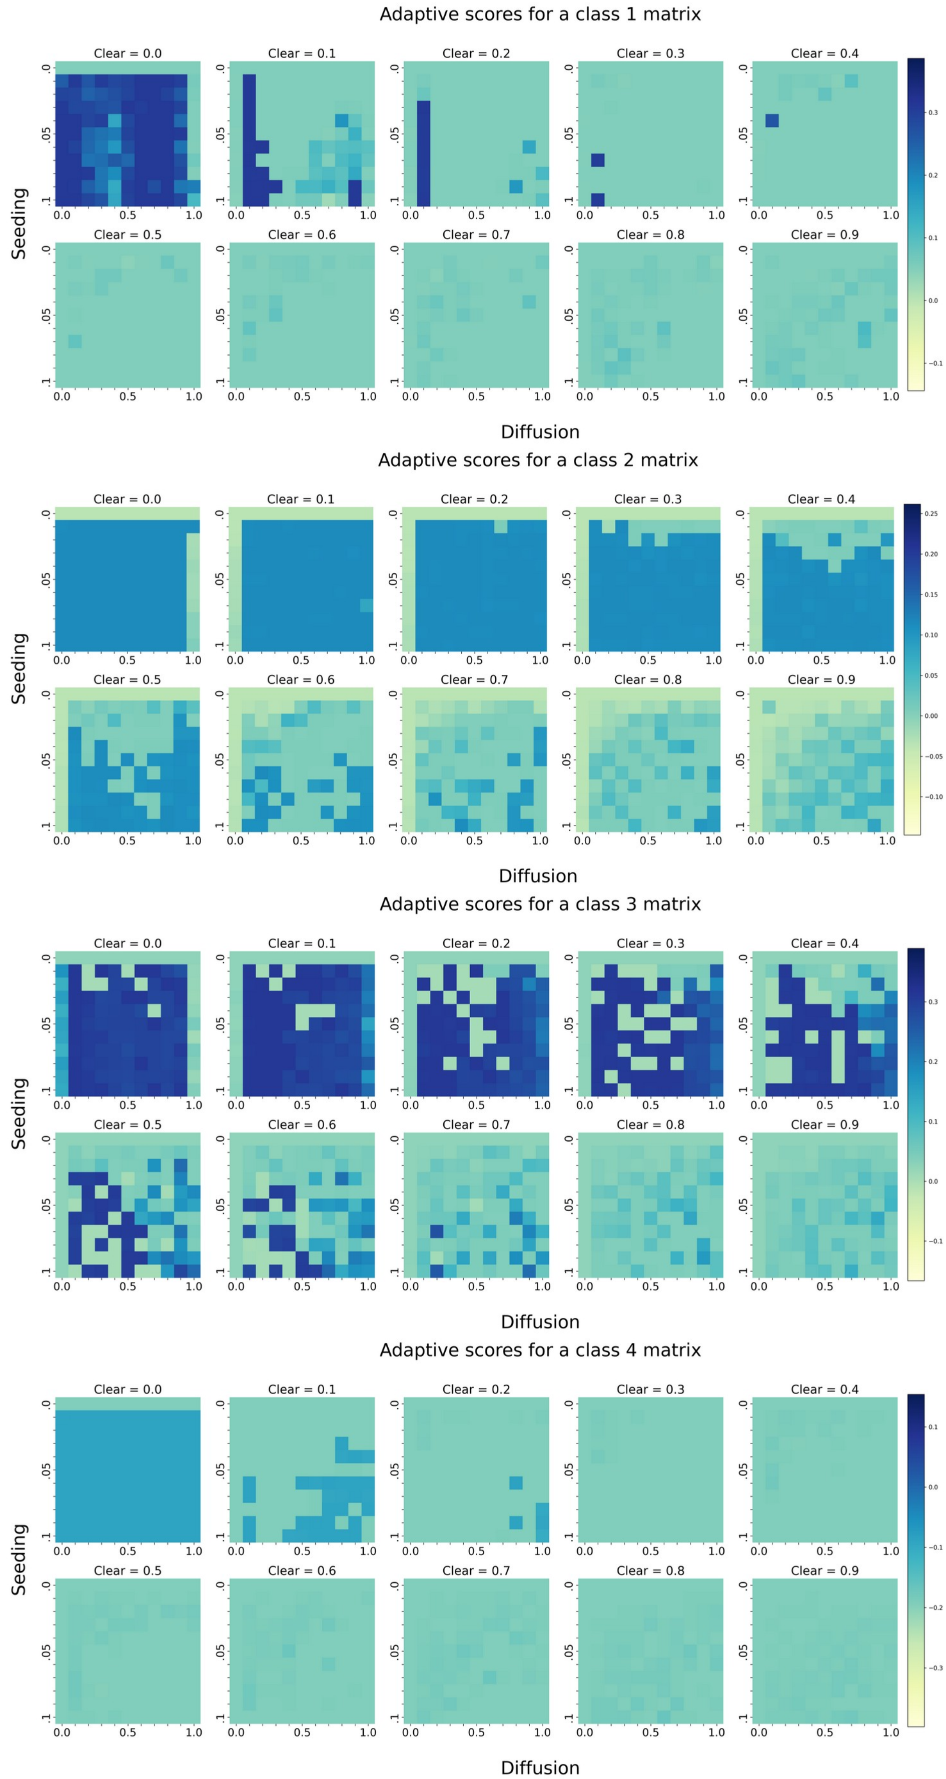

Figure 5: Heatmaps showing the average of 5 runs for all parameter combinations across matrix classes. Darker colors represent higher adaptiveness scores, and more adaptive communities.

Class 3 Matrices

Similar to class 1, class 3 matrices have many weights, with the difference being that class 3 matrices have a mix of strongly positive and strongly negative interactions. These matrices were previously associated with even lower diffusion than class 2 matrices, tending toward values of less than .2 We find that this class is actually capable of promoting two different adaptive communities depending on the abiotic parameters used. One community contains three species and the other contains only two, with the former having a slightly higher adaptive score than the latter.

In our experiments two worlds develop for this matrix class, one in which the 3 species community completely dominates, and one that contains a mix of the two species and three species communities. The former turned out to be far more common, emerging consistently for parameter combinations up to a clear probability of .4. Contrary to previous results, low diffusion was not necessary for the adaptive community to occur, especially with low values of clear. As expected, high clearing values failed to yield adaptive communities, with near zero scores increasing in number as clear probability increased.

Class 4 Matrices

Matrices of class 4 have few connections, that are typically weak, with many species having no interaction at all. Previous work on this class indicated unique results for abiotic parameters, namely high diffusion and seeding values along the range [0, 1](as opposed to the high seeding found in other classes). Our search revealed that this matrix structure was actually very vulnerable to changes in clear. Almost all of the observed adaptive communities develop for a clear probability of zero, with all 110 possible combinations for zero clear probability leading to adaptive communities, while only 29 emerged for a clear probability of .1, and only two more for all clearing values above that. These results indicate that class 4 matrices are extremely dependant on very low clearing values for the development of their adaptive communities.

Discussion

There were some general patterns that emerged across matrix classes. Most notably, a low clear probability was consistently important for facilitating the development of adaptive communities across values of diffusion and seeding. This observation suggests that communities need adequate time to develop within the world before being wiped out by a clearing event. These results also suggest that frequently disturbed environments may be unsuitable for major transitions to occur. Indeed, contrary to our initial hypothesis, a non-zero clear probability does not appear to be necessary for promoting adaptive dynamics. Note, however, that it is possible that other interaction networks exist for which a higher clear probability is necessary to promote adaptive dynamics.

Contrary to previous results, high values of seeding were not necessary for the development of adaptive communities with these matrix structures. As long as the world was not cleared out too often relative to the seeding probability, adaptive communities tended to develop for many different combinations of diffusion and seeding. This finding seems to suggest that, at least with these interaction matrices, diffusion and seeding have less of an impact on the outcome of the world than clear probability.

Conclusion

Some ecological interaction networks when examined in our simple abiotic environment yield phenomena that resemble major-transitions insofar as the dynamics are better explained by community-level adaptation than by purely ecological dynamics. This suggests that we are indeed seeing some level of community-level selection, and that certain regions of parameter space are more likely to produce such community-level selection. Contrary to our preliminary hypotheses, for the matrix classes examined here, adaption was favored by lower values of clear; for example, even a clear probability of zero produced dynamics consistent with adaptation under all matrix structures. Seeding and diffusion had little impact except in a few cases where adaptation emerged at higher clear probabilities.

In future work, we plan to more directly test whether these conditions lead to major transitions. We will do so by 1) allowing the interaction networks to evolve and observing whether they evolve to favor more cooperation, and 2) allowing species to evolve behaviors that harm the individual but help the community. Ultimately, these results inform our expectations about the circumstances under which the origin of life could have occurred, as well as other circumstances under which we would expect egalitarian major transitions to occur.

Acknowledgements

We thank David E. Smith and the Michigan State University ECODE lab for feedback on the ideas in this paper. This material is based upon work supported by the National Science Foundation under Grant No. 2218818. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Blokhuis, A., Lacoste, D., and Nghe, P. (2020). Universal motifs and the diversity of autocatalytic systems. *Proceedings of the National Academy of Sciences*, 117(41):25230–25236.

Foreback, M., Leither, S., and Dolson, E. (2023a). Code and supplemental info for the role of abiotic parameters in the promotion of egalitarian major evolutionary transitions. https://doi.org/10.5281/zenodo.7953995.

Foreback, M., Leither, S., and Dolson, E. (2023b). Using evolutionary computation to find parameters that promote egalitarian major evolutionary transitions. In *Genetic and Evolutionary Computation Conference Companion (GECCO '23 Companion), July 15–19, 2023, Lisbon, Portugal*, New York, NY, USA. ACM.

Goldsby, H., Kerr, B., and Ofria, C. (2020). Major transitions in digital evolution. *Evolution in Action: Past, Present and Future: A Festschrift in Honor of Erik D. Goodman*, pages 333–347.

Herron, M. D. (2021). What are the major transitions? *Biology & Philosophy*, 36(1):2.

Higgs, P. G. and Lehman, N. (2015). The rna world: molecular cooperation at the origins of life. *Nature Reviews Genetics*, 16(1):7–17.

Hordijk, W., Hein, J., and Steel, M. (2010). Autocatalytic sets and the origin of life. *Entropy*, 12(7):1733–1742.

Kauffman, S. A. (1995). *At home in the universe: The search for laws of self-organization and complexity*. Oxford University Press, USA.

Leibold, M. A. and Chase, J. M. (2017). Metacommunity ecology, volume 59. In *Metacommunity Ecology, Volume 59*. Princeton University Press.

Leither, S., Foreback, M., Baum, D. A., and Dolson, E. (2023). Interaction strengths affect whether ecological networks promote the initiation of egalitarian major transitions. In *ALIFE 2023: The 2023 Conference on Artificial Life*. MIT Press.

Malcai, O., Biham, O., Richmond, P., and Solomon, S. (2002). Theoretical analysis and simulations of the generalized lotka-volterra model. *Physical Review E*, 66(3):031102.

McCann, K., Hastings, A., and Huxel, G. R. (1998). Weak trophic interactions and the balance of nature. *Nature*, 395(6704):794–798.

Moreno, M. A. and Ofria, C. (2019). Toward open-ended fraternal transitions in individuality. *Artificial life*, 25(2):117–133.

Peng, Z., Linderoth, J., and Baum, D. A. (2022). The hierarchical organization of autocatalytic reaction networks and its relevance to the origin of life. *PLOS Computational Biology*, 18(9):e1010498.

Peng, Z., Plum, A., and Adam, Z. (2023). Heritable, mutable, and selectable spatial patterns from bistable chemical reaction networks in weakly-coupled reactor arrays.

Peng, Z., Plum, A. M., Gagrani, P., and Baum, D. A. (2020). An ecological framework for the analysis of prebiotic chemical reaction networks. *Journal of theoretical biology*, 507:110451.

Queller, D. C. (1997). Cooperators since life began.

Ratcliff, W. C., Fankhauser, J. D., Rogers, D. W., Greig, D., and Travisano, M. (2015). Origins of multicellular evolvability in snowflake yeast. *Nature communications*, 6(1):6102.

Smith, J. M. and Szathmary, E. (1997). *The major transitions in evolution*. OUP Oxford.