

Environmental connectivity influences the origination of adaptive processes

John Shea^{1*}, Sydney Leither^{2,3*}, Max Foreback^{2,3}, Emily Dolson^{2,3}, and Alexander Lalejini¹

¹ School of Computing, Grand Valley State University

² Computer Science and Engineering, Michigan State University

³ Ecology, Evolutionary Biology, and Behavior, Michigan State University

*Equal authorship contribution.

lalejina@gvsu.edu

Abstract

Spatial structure is hypothesized to be an important factor in the origin of life, wherein encapsulated chemical reaction networks came together to form systems capable adaptive complexification via Darwinian evolution. In this work, we use a computational model to investigate how different patterns of environmental connectivity influence the emergence of adaptive processes in simulated systems of self-amplifying networks of interacting chemical reactions (autocatalytic cycles, “ACs”). Specifically, we measured the propensity for adaptive dynamics to emerge in communities with nine distinct patterns of inter-AC interactions, across ten different patterns of environmental connectivity. We found that the pattern of connectivity can dramatically influence the emergence of adaptive processes; however, the effect of any particular spatial pattern varied across systems of ACs. Relative to a well-mixed (fully connected) environment, each spatial structure that we investigated amplified adaptive processes for at least one system of ACs and suppressed adaptive processes for at least one other system. Our findings suggest that there may be no single environment that universally promotes the emergence of adaptive processes in a system of interacting components (e.g., ACs). Instead, the ideal environment for amplifying (or suppressing) adaptive dynamics will depend on the particularities of the system.

Introduction

Prebiotic chemistry can be represented by networks describing possible reactions among the chemicals present at the origin of life. Understanding the origin of life requires that we understand how these chemical reaction networks gave rise to complex biological organisms (Kauffman, 1986; Hordijk et al., 2010; Ruiz-Mirazo et al., 2014). In most cases, biological complexification can be explained by evolution (Lenski et al., 2003). However, traditional evolutionary theory presupposes a population of genetically-encoded self-replicators, which are likely too complex to spontaneously originate from random chemical reactions (Oono, 2012). Indeed, the emergence of genetic replicators and the beginning of Darwinian evolution is commonly defined as life’s origin (Kauffman, 1993). Thus, understanding the origin of life requires illuminating how genetic self-replicators emerged from chemical reaction networks.

Chemical reaction networks often contain autocatalytic cycles, which are groups of chemicals that collectively produce more of themselves via sequences of reactions. One hypothesis for how the origin of life could have occurred is that pre-genetic adaptive processes emerged in chemical reaction networks made up of interacting autocatalytic cycles (Hordijk and Steel, 2004; Xavier et al., 2020). For example, autocatalytic systems that can amplify themselves faster than others by better exploiting available resources or by containing more synergistic sub-components may more rapidly diffuse across an environment. Biases toward more efficient autocatalytic systems could then drive the transition to populations of genetic replicators (Baum et al., 2023). Importantly, this transition is thought to require some form of spatial structure that “encapsulates” its occupants, allowing different autocatalytic systems to exist in different regions. Encapsulation could initially be driven by abiotic factors, such as slow diffusion across a mineral surface submerged in water. Once autocatalytic systems capable of efficiently propagating themselves emerge, encapsulation could transition into full individuation driven by, for example, a member of the system that produces a membrane. This transition to genetic replicators can be thought of as the first egalitarian major transition in individuality, where unrelated units (comprising loosely interacting systems of autocatalytic cycles) unite to form a higher-level individual (Smith and Szathmari, 1997; Queller and Strassmann, 2009).

The prebiotic “chemical soup” of interacting autocatalytic cycles can be seen through an ecological lens wherein an autocatalytic cycle is analogous to a biological species (Peng et al., 2020). As in complex ecologies of species, autocatalytic cycles can exhibit competitive, predator-prey, or mutualistic associations with one another (Peng et al., 2020). Viewing the prebiotic chemical environment as an ecological community allows us to use insights and techniques from ecological theory to understand how chemical systems could have begun to exhibit dynamics consistent with adaptation. Specifically, we can identify scenarios that lead the system to exhibit adaptive dynamics, which we define here as behavior that is better explained by communities acting

as Darwinian individuals than by pure ecological dynamics within those communities (Baum et al., 2023; Leither et al., 2023; Foreback et al., 2023a). While promising preliminary work has tested the feasibility of pre-genetic adaptive dynamics in the laboratory (Vincent et al., 2019; Sokolsky et al., 2024), such experiments are challenging. Thus, computational modeling for rapid hypothesis testing is a valuable complement to conducting laboratory studies. Indeed, recent computational studies support that adaptive processes can emerge from simple “ecologies” of interacting autocatalytic cycles (Foreback et al., 2023a,b; Leither et al., 2023). These studies found that the ability of a system to demonstrate dynamics better explained by evolutionary dynamics than by purely ecological processes is influenced by both 1) abiotic characteristics (e.g., ease of diffusion, rate of disturbance, and rate of seeding new chemicals) (Foreback et al., 2023a) and 2) the particular structure of interactions among autocatalytic cycles (Leither et al., 2023).

Many characteristics of the environment, such as spatial and temporal variability, are also likely to play a critical role in the emergence of adaptive dynamics. In this work, we focus on the environment’s connectivity, which mediates the range of possible interactions among its occupants. Spatial structure is known to influence ecological and evolutionary processes both in natural (Kool et al., 2013) and artificial systems (Dolson and Ofria, 2021; Tomassini, 2005; Plum and Baum, 2022). We use an artificial ecology model to investigate how different patterns of environmental connectivity influence the probability of adaptive dynamics occurring, a property termed “transitionability” (Leither et al., 2023). We hypothesize that the environment’s spatial structure may influence the emergence of pre-genetic adaptive dynamics, and that the environment’s connectivity can be manipulated to amplify or suppress the emergence of adaptive dynamics.

We measured the emergence of adaptive dynamics on ten different patterns of environmental connectivity for each of nine distinct “communities” of interacting autocatalytic cycles. Overall, we found that the pattern of environmental connectivity can dramatically influence the emergence of adaptive dynamics. However, the effect of any particular spatial structure depended on the structure of the community. Each spatial structure that we investigated amplified the emergence of adaptive dynamics for at least one community type and suppressed the emergence of adaptive dynamics for at least one community type.

Background

In community ecology, a metacommunity is a set of distinct communities connected by dispersal. Connectivity within a metacommunity is often determined by spatial distance and other biotic and abiotic phenomena, such as winds, ocean currents, physical barriers, or the migratory patterns of other species (Moritz et al., 2013; Leibold and Chase, 2017). This connectivity controls how organisms disperse across

space. In turn, dispersal patterns determine the range of possible interactions among species, mediate the exchange of individuals (and genes) between populations, and influence where species can move in response to habitat unsuitability (Baguette et al., 2013). For example, populations subjected to a changing environment (e.g., due to climate change, competition with an invasive species, etc.) may have greater extinction risk if dispersal within the metacommunity is difficult (Campos et al., 2020).

The spatial structure of the environment also influences evolutionary processes at both short- and long-term time scales. Environmental connectivity strongly influences selection pressure, the probability of mutation fixation, and the overall rate of evolutionary adaptation (Perfeito et al., 2008; Bailey and Kassen, 2012). For example, fully connected population structures (e.g., well-mixed environments) often allow for rapid selective sweeps, leading to lower amounts of standing genetic variation (Nahum et al., 2015). In contrast, population structures with lower connectivity can slow down selective sweeps, often allowing for greater levels of standing genetic variation (Covert et al., 2014).

Previous work has used graph theoretic approaches to model short-term evolution on a wide range of network structures (i.e., environmental connectivity patterns). Indeed, different connectivity patterns either amplify or suppress the fixation probability of new mutations, directly influencing the overall rate of evolutionary adaptation (Kuo and Carja, 2024; Kuo et al., 2021; Allen et al., 2021; Möller et al., 2019). Graph properties such as the mean degree distribution of graph nodes, variance of degree distribution, and assortative mixing (i.e., preference for nodes to attach to other nodes with a similar degree) predict a spatial structure’s effect on fixation probabilities. For example, structures with low assortativity (e.g., graphs with hub-and-spoke structures) often act as strong amplifiers of adaptive evolution (Kuo et al., 2021). Other studies suggest that rates of adaptation can be controlled, in part, by tuning environmental connectivity patterns (Möller et al., 2019).

At longer time scales, environmental connectivity patterns can drive the evolution and maintenance of biodiversity (Losos and Schluter, 2000; Stein et al., 2014; Hernández-Hernández et al., 2021). For example, allopatric speciation occurs when a population is separated into two reproductively isolated populations (Grant et al., 2000; Gavrillets, 2014; Golestani et al., 2012). This separation is often caused by a physical barrier that isolates two geographic regions from one another (e.g., a mountain range, body of water, etc.). Without such spatial isolation preventing gene flow, populations would more often remain genetically homogeneous (but see (Losos and Schluter, 2000; Covert et al., 2014)). Spatial resource heterogeneity has also been shown to increase diversity and promote the evolution of complex traits (Dolson et al., 2017). Particular regions of spatially heterogeneous environments can also act as evolu-

tionary hotspots where novel adaptive traits are more likely to evolve (Dolson and Ofria, 2017). Likewise, many evolutionary search algorithms impose spatial structure on a population of candidate solutions in order to promote diversity, avoid premature convergence, improve search space exploration, and increase overall problem-solving success (Miralavy and Banzhaf, 2023; Tomassini, 2005).

Methods

We tested the effect of spatial structure on the origination of adaptive processes using an artificial ecology simulation. Specifically, we measured the emergence of adaptive dynamics for nine distinct simulated communities of interacting autocatalytic cycles (ACs) on ten spatial structures. For each distinct community, we compared adaptive outcomes across spatial structures, using a well-mixed (fully connected) control environment as a baseline for expectations in the absence of spatial structure.

Artificial Ecology Model

Our artificial ecology model simulates the growth of N interacting ACs within a metacommunity (Leibold et al., 2004). The metacommunity comprises a network of local sites, each of which may contain different concentrations of each AC. ACs can interact with other ACs within the same site, influencing each other's abundance at that site. An AC's abundance is represented by a continuous value, and changes in any AC's abundance is calculated using a generalized Lotka-Volterra model. That is, the growth of an AC is directly proportional to the abundance of ACs beneficial to it and the scarcity of ACs detrimental to it within its local site. In our experiments, we bounded each AC's abundance (within a local site) to a maximum value of 10,000.

Each AC is defined by its interactions with each other AC, as specified by an interaction network. An interaction network is an N by N weighted and directed matrix where N is the number of unique ACs. All experiments in this work used interaction matrices with $N = 10$ unique ACs. An element x_{ij} within the interaction matrix defines the effect of AC j on AC i . ACs may also have an intrinsic growth rate, which is the value at x_{ii} (i.e., the diagonal). Each interaction x_{ij} is continuous in the range $[-1, 1]$, where positive values are beneficial to growth and negative values are detrimental.

At each simulation time step, ACs may interact within a local site, influencing each others' abundance. Then, ACs may be seeded into a local site at a per-AC seeding rate (0.05 in our experiments), which models rare, long-distance migration events. Finally, ACs may diffuse into adjacent sites, as defined by the environment's spatial connectivity. During diffusion, 5% of each AC at a site diffuses into each neighboring site. We ran all of our experiments for 20,000 time steps. We chose our seeding and diffusion rates based on results from previous studies (Foreback et al., 2023a; Leither et al., 2023) and exploratory experiments.




Figure 1: Example time step in the artificial ecology model, adaptive null model, and assembly null model. In each model, a seeding event has taken place, adding a random AC within the bottom-right site (in pink). The artificial ecology model shows the diffusion of the ACs into adjacent sites. The adaptive model shows a reproduction event in which the site with the highest biomass has reproduced as a whole into another site.

Measuring adaptive dynamics After 20,000 simulation updates, we “stabilize” each site by updating AC abundances for an additional 10,000 time steps. For each stabilization time step, AC abundances are influenced only by local interactions, allowing any recently seeded ACs to grow or decline to a stable abundance level. This stabilization process facilitates grouping communities together for further analysis based on their equilibrium states.

After stabilization, we record the distribution of ACs present at each site in the metacommunity. We define the ACs at a single site as a community. Each community is summarized according to the rank ordering of AC abundance. An AC's abundance must exceed 10.0 to be considered as "present" in this rank ordering. This thresholding is especially important for interaction networks with weak interaction strengths where a post-stabilization abundance may approach zero without reaching it. Additionally, this summarization reduces noise by allowing us to categorize communities according to which ACs dominate, as opposed to the exact abundances present at each site. These ranked communities are then used to detect the frequency at which each community appears in a metacommunity. If two sites have identical AC members present, but each AC's abundance differs slightly between the two sites, we still categorize them as equivalent communities, as long as the rank ordering of abundance levels are the same.

We then compare the final communities found in the simulation to those found in two null models: a *community assembly null model* and an *adaptive null model*. Figure 1 summarizes the differences between the full simulation and the two null models. Each of these two models are initialized in the same way as the full simulation: a metacommunity comprising a network of local sites, each of which

may contain different concentrations of each AC. However, in the community assembly null model, the distribution of ACs present at each site is determined entirely by ecological processes. That is, at each time step, ACs within each site may interact to influence each other's abundance (as determined by the interaction network), and the only way for an AC to enter a site is through a seeding event. The communities that develop are therefore a result of ecological succession and are independent across sites. We record the distribution of communities present in the assembly model's metacommunity after 20,000 updates.

In the adaptive null model, the distribution of ACs present at each site is determined by adaptive processes, where each group of ACs "reproduce" as a single unit and compete with other groups for space in the metacommunity. At each time step, ACs within each site may interact to influence each others' abundance. Then, a group of ACs at a site may reproduce into an adjacent target site (chosen randomly from all adjacent sites). When a group of ACs reproduces into another site, 10% of each AC in the reproducing group is placed (together) into the target site, replacing any previous occupants. In our experiments, the probability of a site reproducing (R_i) is directly proportional to the total abundance of all ACs in the site: R_i equals the total AC abundance at site i divided by the maximum possible AC abundance at a site. After reproduction, new ACs may be seeded into each site (at the per-AC seeding rate). In this way, communities of ACs compete for limited space in the metacommunity, as communities capable of rapid growth out-compete those with slower growth rates. Because communities of ACs reproduce as a unit, any variation in community composition introduced by seeding is heritable. Thus, the combination of this heritable variation and competition for limited space results in pure evolution by natural selection. As in the community assembly null model, we record the distribution of communities present in the adaptive model's metacommunity after 20,000 updates.

For each independent replicate of each experiment, we ran 200 independent community assembly null models and 200 independent adaptive null models. For each null model, we measure the distributions of final communities in order to provide a representative sample of communities that we expect to arise from purely ecological processes or from purely adaptive processes. We then compare the summarized communities recorded from the artificial ecology simulation to those from each null model, measuring "transitionability" as the likelihood ratio of having observed our data under purely ecological versus purely adaptive dynamics:

$$\sum_i^C (\ln(\text{adaptive-prop}_i)) - \sum_i^C (\ln(\text{assembly-prop}_i))$$

where C is the set of communities recorded from the simulation, and adaptive-prop_i and assembly-prop_i are the

proportion of communities equivalent to C_i found in the adaptive and assembly null models, respectively. The adaptive and assembly proportions are adjusted using additive smoothing to avoid division by zero errors. For each site in the simulation, this equation considers the proportion of sites in which the same community was present in each null model. If the community was present at a higher proportion in the adaptive model than the assembly model, the site will contribute positively to the overall transitionability score; the opposite is true for communities that are present more often in the assembly model.

Transitionability scores greater than 0 indicate that the final communities observed in a simulation are more similar to those observed in the adaptive null model than the assembly null model. Conversely, transitionability scores below 0 indicate that these final communities are more similar to those observed in the assembly null model than those in the adaptive null model. Thus, high transitionability scores are evidence that adaptive dynamics have emerged.

Experimental Design

For each of nine AC interaction networks, we compared the emergence of adaptive dynamics (measured as transitionability) across ten environmental connectivity (spatial structure) regimes. We used a fully connected structure (i.e., a well-mixed environment) as a baseline for establishing which regimes amplified versus suppressed adaptive dynamics, as is standard practice (Kuo and Carja, 2024; Möller et al., 2019). For each interaction network, we ran 20 independent replicates of each connectivity regime.

Spatial structures For all experiments, we limited the number of sites in the metacommunity to 100, and the environment's spatial structure determined the connectivity of sites. Connections between sites were undirected. We included spatial structures commonly used in artificial life systems and from studies on the effects of spatial structure on ecological and evolutionary dynamics. For spatial structures generated by stochastic graph generation algorithms, we independently generated 20 structures to be used across experiments (one per replicate). Brief descriptions of the spatial structures are given below:

- **Well-mixed (fully connected):** A fully connected graph where each vertex is connected to all other vertices.
- **Toroidal lattice:** Vertices are organized into a toroidal grid where each vertex is connected to its four neighboring vertices. The vertices in the top and bottom rows and left and right columns are connected, respectively.
- **Linear chain:** Vertices are organized into a linear chain, where each vertex is connected to its two neighbors.
- **Cycle:** A linear chain graph, but the vertices at the two ends of the chain are connected.

- **Wheel:** A single hub vertex is connected to all vertices in a cycle comprising all other vertices in the graph.
- **Star:** A tree with one internal vertex; all other vertices are leaves connected to the single internal vertex.
- **Windmill:** A graph with n size- k cliques that each share a single “hub” vertex. For this work, $n = 10$ and $k = 10$.
- **Comet-kite:** A graph comprising a large, fully connected set of core nodes with random “tails” (Möller et al., 2019). To generate a comet-kite graph, we construct a fully connected core, select a node from the core to attach t initial tail nodes to, and then sequentially attach additional nodes to randomly chosen tail nodes. In this work, we used a core size of 40 nodes, attached 20 initial tail nodes, and added 40 additional tail nodes. These parameters were chosen based on results from (Möller et al., 2019).
- **Random Barabasi-Albert:** A randomly generated, scale-free graph is constructed by sequentially attaching new nodes with m edges, which are preferentially connected to existing nodes with high degree (Albert and Barabási, 2002). Here, $m = 10$.
- **Random Waxman:** A randomly generated graph is constructed by placing nodes uniformly at random in a 2-dimensional space (Waxman, 1988). Each pair of nodes distance d from one another are connected with probability $p = \beta e^{-d/\alpha L}$. Here, we used $\beta = 0.4$ and $\alpha = 0.2$.

We used the networkx library (Hagberg et al., 2008) to generate the star, windmill, cycle, wheel, random Barabasi-Albert, and random Waxman graphs. The graph generation code and all spatial structures used in our experiments (including visualizations of each) are included in our supplemental material (Lalejini et al., 2024).

Interaction networks We tested the effect of differing spatial structure in the context of nine different interaction networks, allowing us to assess interactions between spatial structure and interaction structure. We generated each interaction network with a graph evolution tool that uses a genetic algorithm to construct graphs with specific properties (Leither et al., 2024). Specifically, we generated interaction networks that varied in their connectance and proportion of positive interactions among ACs, as these properties were previously found to be important (Leither et al., 2023). Connectance is calculated via $\frac{E}{N^2}$ where E is the number of edges in the graph, and N is the number of nodes. Positive interaction proportion is the ratio of positively-weighted edges to all established edges. Our interaction networks represent a full factorial combination of the values 0.25, 0.5, and 0.75 for connectance and proportion of positive interactions.

Connectance and positive interaction proportion significantly impact the transitionability of an interaction network

in baseline conditions. By explicitly varying these properties across interaction networks, we ensure a diverse set of possible communities of ACs. Additional properties of the interaction networks may also influence the emergence of adaptive dynamics; however, testing the effects of spatial structure on a greater number of interaction networks falls outside of the scope of this study.

Statistical analyses

For all experiments, we limited comparisons to measurements taken from replicates that shared an interaction network. When comparing distributions of measurements taken from different spatial structure regimes, we performed Kruskal-Wallis tests to screen for statistical differences among independent conditions. For comparisons in which the Kruskal-Wallis test was significant (significance level of 0.05), we performed post-hoc Wilcoxon rank-sum tests to identify pairwise differences, and we corrected for multiple comparisons using the Holm-Bonferroni method.

We used Spearman’s rank correlation method to test for correlations between variables, and we corrected for multiple correlations using the Holm-Bonferroni method. Complete results of all statistical analyses are available in our supplemental material (Lalejini et al., 2024).

Code and Data Availability

Our experiments are implemented in C++ using the Empirical library (Vostinar et al., 2024). Experiment software and data analyses can be found in our supplemental material, which is hosted on GitHub and archived on Zenodo (Lalejini et al., 2024). Our experiment data are archived on the Open Science Framework at (Lalejini, 2024).

Results and Discussion

Environmental connectivity influences the emergence of adaptive dynamics

We investigated the impact of spatial structure in the context of nine different interaction networks that varied in their connectance and proportion of positive interactions among autocatalytic cycles (ACs). For each interaction network, we compared the transitionability scores achieved in each of ten classes of spatial structure. Transitionability score measures the emergence of adaptive dynamics. Communities with transitionability scores greater than zero exhibit behavior that more closely resembles what would be expected from systems undergoing purely adaptive dynamics, and scores less than zero exhibit behavior that more closely resembles expectations under purely ecological dynamics.

Figure 2 shows the transitionability of each community type across each spatial structure. Dashed vertical lines represent the median transitionability achieved (across replicates) in a well-mixed (fully connected) environment. The well-mixed regime represents baseline expectations in the absence of spatial structure. That is, any spatial structures




Figure 2: Transitionability scores achieved on each spatial structure for each community type (interaction network). Each panel corresponds to an interaction network with a particular combination of connectance (rows) and positive interaction proportion (“PIP”, columns). Dashed vertical lines (black) represent the median scores in the well-mixed environment, and solid vertical lines (red) are drawn at 0 transitionability score. Each Kruskal-Wallis test for each interaction network was statistically significant ($p < 0.001$).

that amplify the emergence of adaptive dynamics should result in transitionability scores greater than those under the well-mixed regime, and structures that suppress the emergence of adaptive dynamics should result in scores that fall below those achieved in the well-mixed regime.

As in previous work (Leither et al., 2023), different community structures varied in their baseline transitionability, with some producing scores below zero (e.g., Fig. 2 at PIP 0.25, Conn. 0.5) and others producing scores much greater than zero (e.g., Fig. 2 at PIP 0.75, Conn. 0.5). For each interaction network, we found significant variation in transitionability among spatial structure regimes (Kruskal-Wallis, $p < 0.001$), although for one interaction network (connectance = 0.75, PIP = 0.75) no spatial structures produced transitionability significantly different than well-mixed. For all other interaction networks, we found at least one spatial

structure with significantly different transitionability scores than those achieved in the well-mixed regime (corrected Wilcoxon rank-sum tests, $p < 0.02$). However, which spatial structures differed from well-mixed and the direction of their effects varied by interaction network (Table 1).

The effect of spatial structure varies across community types

The effect of any particular spatial structure varied across interaction networks (Figure 2). For example, the star spatial structure was the strongest amplifier of transitionability for the interaction network with low connectance and low positive interaction proportion ($p < 0.001$). In contrast, the star structure was the strongest suppressor of transitionability for the interaction network with low connectance and high positive interaction proportion ($p < 0.001$).

For each interaction network, we categorized each spatial structure’s effect on transitionability as either an “amplifier”, “suppressor”, or “neither”. If a spatial structure resulted in transitionability scores that differed significantly from those of the well-mixed condition (Holm-Bonferroni corrected Wilcoxon rank-sum tests, $p < 0.05$), we categorized it as either an amplifier or suppressor if scores were greater or lower, respectively, than those from the well-mixed regime. If we failed to detect a significant difference in scores between well-mixed and a particular spatial structure, we categorized that spatial structure as “neither”.

Interaction Network	Spatial structure effects
Conn. 0.25 PIP 0.25	+ Comet-kite, Waxman, Star
	– Lattice, Windmill
	○ Cycle, Linear, Barabasi-Albert, Wheel
Conn. 0.25 PIP 0.50	+ Star , Wheel
	– Comet-kite, Cycle, Linear, Barabasi-Albert , Waxman, Windmill
	○ Lattice
Conn. 0.25 PIP 0.75	+
	– Comet-kite, Cycle, Star , Linear, Lattice, Wheel, Windmill
	○ Barabasi-Albert, Waxman
Conn. 0.50 PIP 0.25	+ Comet-kite, Cycle, Linear, Barabasi-Albert, Waxman, Star , Wheel
	–
	○ Lattice, Windmill
Conn. 0.50 PIP 0.50	+ Comet-kite, Cycle, Linear, Barabasi-Albert, Waxman, Lattice, Wheel , Windmill
	–
	○ Star
Conn. 0.50 PIP 0.75	+ Comet-kite, Cycle, Linear, Barabasi-Albert, Waxman, Star , Lattice, Wheel
	–
	○ Windmill
Conn. 0.75 PIP 0.25	+ Comet-kite, Cycle, Linear, Waxman, Star , Lattice, Wheel , Windmill
	–
	○ Barabasi-Albert
Conn. 0.75 PIP 0.50	+ Cycle, Linear , Star , Lattice, Wheel, Windmill
	– Comet-kite
	○ Barabasi-Albert, Waxman
Conn. 0.75 PIP 0.75	+
	–
	○ Comet-kite, Cycle, Linear, Barabasi-Albert, Waxman, Star , Lattice, Wheel , Windmill

Table 1: Effect of each spatial structure on adaptive processes for each interaction network: amplifier (“+”), suppressor (“–”), or neither (“○”). The structure with the greatest median effect magnitude (positive and negative) for each interaction network is bolded.

Table 1 shows how spatial structures were categorized for each interaction network, and Figure 3 shows the number of times each spatial structure fell into each category across experiments. The overall effect of *every* spatial structure varied by interaction network; each structure acted as either an amplifier, a suppressor, or neither at least once. Moreover, the distribution of how spatial structures were categorized varied across all interaction networks. Only the linear chain and cycle spatial structures behaved similarly to one another across all interaction matrices, which is unsurprising given their similarity in underlying connectivity. The presence of such variation is broadly consistent with evolutionary graph theory literature. For example, star, lattice, windmill, and random Barabasi Albert graphs have each been shown to amplify adaptive evolution in some contexts and suppress it in others (Kuo et al., 2021; Allen et al., 2021, 2020; Kuo




Figure 3: Distribution of effects for each spatial structure across all experiments.

and Carja, 2023; Tkadlec et al., 2020).

Previous artificial life studies investigating the emergence of adaptive dynamics found that a community’s “transitionability” depended on the particular set of abiotic conditions (Foreback et al., 2023a). Indeed, when AC community structures and abiotic conditions were simultaneously optimized for transitionability, the optimal abiotic parameters varied across optimized community structures (Foreback et al., 2023b). Taken together with these previous results, our data suggest that the emergence of adaptive dynamics in any system is sensitive to the broader environmental context in which the system exists. A system that supports adaptive dynamics in one environmental context may not in other environmental contexts.

Different spatial structure properties correlated with the emergence of adaptive dynamics for different interaction networks

The ten spatial structures used in our experiments represent a *small* sample of possible environmental connectivity patterns. To help inform future studies on how to configure connectivity to minimize or maximize the emergence of adaptive dynamics, we screened for spatial structure properties that correlated with transitionability scores. We computed 21 graph properties for each spatial structure. We then identified the top three properties that most strongly correlated (positive or negative) with transitionability score (Holm-Bonferroni corrected Spearman’s rank correlation, $p < 0.05$). We chose efficiently computable properties commonly used in studies investigating the effects of graph structures on evolution (Kuo and Carja, 2024; Hagberg et al., 2008). While the chosen properties are non-exhaustive, they provide a starting point to inform the design of spatial structures in future work. All 21 graph properties are listed and defined in supplemental material (Lalejini et al., 2024); for brevity, we focus our discussion on those that appeared most correlated with transitionability scores.

Interaction network		Top graph properties most strongly correlated with transitionability scores
Conn.	PIP	
0.25	0.25	Num. bridges (0.57), Edge connectivity (-0.51), Node degree variance (0.49)
0.25	0.50	Degree assortativity coeff. (-0.64), Density (-0.44), Radius (-0.41)
0.25	0.75	Median node degree (0.85), Density (0.83), Mean node degree (0.81)
0.50	0.25	Median node degree (-0.68), Edge connectivity (-0.64), Num. bridges (0.58)
0.50	0.50	Median node degree (-0.60), Kemeny constant (0.59), Global efficiency (-0.56)
0.50	0.75	Median node degree (-0.81), Density (-0.67), Mean node degree (-0.65)
0.75	0.25	Median node degree (-0.68), Mean node degree (-0.68), Density (-0.65)
0.75	0.50	Mean node degree (-0.72), Density (-0.66), Average node connectivity (-0.65)
0.75	0.75	Degree assortativity coeff. (-0.37), Radius (-0.35), Diameter (-0.34)

Table 2: Top three statistically significant graph properties most strongly correlated with transitionability score for each interaction network (Holm-Bonferroni corrected Spearman’s rank correlation, $p < 0.05$). Correlation strength is given in parentheses. Positive correlations are in blue, and negative correlations are given in red.

Table 2 shows the graph properties with the three strongest, statistically significant correlations (positive or negative) with transitionability scores for each interaction network. The median degree of all nodes in the spatial structure had the strongest correlation with transitionability score for five out of nine interaction networks (corrected Spearman’s rank correlation, $p < 0.001$); however, the direction of the relationship differed among these interaction networks: four were negative, and one was positive. In general, we found that the directionality of correlations often differed across experiments with different interaction networks.

Assortativity and degree distribution, properties previously found to affect the probability of mutation fixation (Kuo et al., 2021), were represented in the top three correlations for all nine interaction matrices. Consistent with prior work showing that graphs with low assortativity amplify adaptive evolution (Kuo et al., 2021), we found that assortativity was negatively correlated with transitionability; that is, spatial structures with lower assortativity generally resulted in greater transitionability scores.

To identify graph properties that consistently correlated with transitionability score, we counted the number of times each property had a statistically significant correlation of at least moderate strength (magnitude ≥ 0.5). The Kemeny constant had at least a moderate, positive correlation with transitionability score for three interaction networks and was the only graph property to have a consistent relationship with transitionability score (of properties that had significant correlations across more than one experiment). The Kemeny constant measures the time needed to spread across a graph; a lower Kemeny constant indicates a more closely connected graph, whereas a high Kemeny constant indicates a more spread-out graph (Crisostomi et al., 2011). That is, more spread-out spatial structures more often produced greater transitionability scores for at least three interaction networks.

Edge connectivity, which measures the minimum number of edges that must be removed to disconnect a graph, had the greatest total number of significant correlations of at least moderate strength with transitionability score. Edge connectivity negatively correlated with transitionability across five interaction networks and positively correlated with transitionability once. Indeed, nearly all graph properties with at least two moderate correlations had a mix of positive and negative relationships across experiments, further suggesting that the effects of any particular environmental connectivity pattern depend on the structure of interactions.

Conclusion

In this work, we established that the pattern of environmental connectivity can influence the emergence of adaptive dynamics in a simulated system of interacting autocatalytic cycles (Figure 2). However, whether a particular pattern of connectivity amplifies or suppresses adaptive dynamics will likely vary across systems (Figure 3), as we found that spatial structures had different effects for different communities (defined by different interaction matrices). In combination with results from previous studies on the emergence of adaptive dynamics (Foreback et al., 2023a; Leither et al., 2023; Foreback et al., 2023b), our findings suggest that there may be no single environment that universally promotes the origination of adaptive dynamics. Instead, the ideal environment for amplifying (or suppressing) adaptive dynamics will depend on the system. Future work will investigate whether we can predict a spatial structure’s effect given a system’s interaction network.

While our results have established that the pattern of environmental connectivity is an important factor in the origination of adaptive processes, we have yet to disentangle the mechanisms underlying connectivity’s effects. Future work will conduct a case study on selected replicates from our experiments, visualizing how different connectivity patterns influence community diffusion. Additionally, future work will use the graph properties that we identified as most strongly correlated with transitionability scores to design (by hand and using automated graph generation tools) spatial structures that we would expect to maximally amplify or suppress adaptive dynamics.

Acknowledgements

We thank the reviewers, MSU ECODE lab, David Baum, and Eric Smith for feedback on this work. This work was supported in part by the National Science Foundation under Grant No. 2218818 and by GVSU through a Kindschi Research Fellowship to JS. This work was also supported in part through computational resources and services provided by the Institute for Cyber-Enabled Research at MSU.

References

- Albert, R. and Barabási, A.-L. (2002). Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74(1):47–97. Publisher: American Physical Society.
- Allen, B., Sample, C., Jencks, R., Withers, J., Steinhagen, P., Brizuela, L., Kolodny, J., Parke, D., Lippner, G., and Dementieva, Y. A. (2020). Transient amplifiers of selection and reducers of fixation for death-Birth updating on graphs. *PLOS Computational Biology*, 16(1):e1007529.
- Allen, B., Sample, C., Steinhagen, P., Shapiro, J., King, M., Hedspeth, T., and Goncalves, M. (2021). Fixation probabilities in graph-structured populations under weak selection. *PLOS Computational Biology*, 17(2):e1008695.
- Baguette, M., Blanchet, S., Legrand, D., Stevens, V. M., and Turlure, C. (2013). Individual dispersal, landscape connectivity and ecological networks. *Biological Reviews*, 88(2):310–326.
- Bailey, S. F. and Kassen, R. (2012). Spatial Structure of Ecological Opportunity Drives Adaptation in a Bacterium. *The American Naturalist*, 180(2):270–283.
- Baum, D. A., Peng, Z., Dolson, E., Smith, E., Plum, A. M., and Gagrani, P. (2023). The ecology–evolution continuum and the origin of life. *Journal of the Royal Society Interface*, 20(208):20230346.
- Campos, F. S., Lourenço-de Moraes, R., Ruas, D. S., Mira-Mendes, C. V., Franch, M., Llorente, G. A., Solé, M., and Cabral, P. (2020). Searching for Networks: Ecological Connectivity for Amphibians Under Climate Change. *Environmental Management*, 65(1):46–61.
- Covert, A., McFetridge, S., and DeLord, E. (2014). Structured Populations with Limited Resources Exhibit Higher Rates of Complex Function Evolution. pages 129–134. The MIT Press.
- Crisostomi, E., Kirkland, S., and Shorten, R. (2011). A Google-like model of road network dynamics and its application to regulation and control. *International Journal of Control*, 84(3):633–651.
- Dolson, E. and Ofria, C. (2017). Spatial resource heterogeneity creates local hotspots of evolutionary potential. In *Proceedings of the 14th European Conference on Artificial Life ECAL 2017*, pages 122–129, Lyon, France. MIT Press.
- Dolson, E. and Ofria, C. (2021). Digital Evolution for Ecology Research: A Review. *Frontiers in Ecology and Evolution*, 9:18.
- Dolson, E. L., Pérez, S. G., Olson, R. S., and Ofria, C. (2017). Spatial resource heterogeneity increases diversity and evolutionary potential. preprint, Ecology.
- Foreback, M., Leither, S., Baum, D. A., and Dolson, E. (2023a). The role of abiotic parameters in the promotion of egalitarian major evolutionary transitions. In *ALIFE 2023: The 2023 Conference on Artificial Life*. MIT Press.
- Foreback, M., Leither, S., and Dolson, E. (2023b). Using evolutionary computation to find parameters that promote egalitarian major evolutionary transitions. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, pages 135–138.
- Gavrilets, S. (2014). Models of Speciation: Where Are We Now? *Journal of Heredity*, 105(S1):743–755.
- Golestani, A., Gras, R., and Cristescu, M. (2012). Speciation with gene flow in a heterogeneous virtual world: can physical obstacles accelerate speciation? *Proceedings of the Royal Society B: Biological Sciences*, 279(1740):3055–3064.
- Grant, P. R., Grant, B. R., and Petren, K. (2000). The allopatric phase of speciation: the sharp-beaked ground finch (*Geospiza difficilis*) on the Galápagos islands. *Biological Journal of the Linnean Society*, 69(3):287–317.
- Hagberg, A., Swart, P. J., and Schult, D. A. (2008). Exploring network structure, dynamics, and function using NetworkX.
- Hernández-Hernández, T., Miller, E. C., Román-Palacios, C., and Wiens, J. J. (2021). Speciation across the Tree of Life. *Biological Reviews*, 96(4):1205–1242.
- Hordijk, W., Hein, J., and Steel, M. (2010). Autocatalytic sets and the origin of life. *Entropy*, 12(7):1733–1742.
- Hordijk, W. and Steel, M. (2004). Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of theoretical biology*, 227(4):451–461.
- Kauffman, S. A. (1986). Autocatalytic sets of proteins. *Journal of theoretical biology*, 119(1):1–24.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA.
- Kool, J. T., Moilanen, A., and Treml, E. A. (2013). Population connectivity: recent advances and new perspectives. *Landscape Ecology*, 28(2):165–185.
- Kuo, Y. P., Arrieta, C. N., and Carja, O. (2021). A theory of evolutionary dynamics on any complex spatial structure. preprint, Evolutionary Biology.
- Kuo, Y. P. and Carja, O. (2023). Evolutionary graph theory on rugged fitness landscapes. preprint, Evolutionary Biology.
- Kuo, Y. P. and Carja, O. (2024). Evolutionary graph theory beyond pairwise interactions: Higher-order network motifs shape times to fixation in structured populations. *PLOS Computational Biology*, 20(3):e1011905.
- Lalejini, A. (2024). Experiment data (Archived on OSF). 10.17605/OSF.IO/K3D8G. <https://osf.io/k3d8g>.
- Lalejini, A., Shea, J., Leither, S., Foreback, M., and Dolson, E. (2024). Supplemental material (Archived GitHub repository). 10.5281/zenodo.10891182. <https://github.com/amlalejini/alife-2024-spatial-chem-eco>.
- Leibold, M. A. and Chase, J. M. (2017). *Metacommunity Ecology*. Princeton University Press. Publication Title: Metacommunity Ecology, Volume 59.
- Leibold, M. A., Holyoak, M., Mouquet, N., Amarasekare, P., Chase, J. M., Hoopes, M. F., Holt, R. D., Shurin, J. B., Law, R., Tilman, D., et al. (2004). The metacommunity concept: a framework for multi-scale community ecology. *Ecology letters*, 7(7):601–613.

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: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*. MIT Press.

Leither, S., Ragusa, V., and Dolson, E. (2024). Evolving weighted and directed graphs with constrained properties. In *Genetic and Evolutionary Computation Conference (GECCO '24 Companion)*. ACM.

Leski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.

Losos, J. B. and Schluter, D. (2000). Analysis of an evolutionary species–area relationship. *Nature*, 408(6814):847–850.

Miralavy, I. and Banzhaf, W. (2023). Spatial genetic programming. In *European Conference on Genetic Programming (Part of EvoStar)*, pages 260–275. Springer.

Moritz, C., Meynard, C. N., Devictor, V., Guizien, K., Labruno, C., Guarini, J., and Mouquet, N. (2013). Disentangling the role of connectivity, environmental filtering, and spatial structure on metacommunity dynamics. *Oikos*, 122(10):1401–1410.

Möller, M., Hindersin, L., and Traulsen, A. (2019). Exploring and mapping the universe of evolutionary graphs identifies structural properties affecting fixation probability and time. *Communications Biology*, 2(1):137.

Nahum, J. R., Godfrey-Smith, P., Harding, B. N., Marcus, J. H., Carlson-Stevermer, J., and Kerr, B. (2015). A tortoise–hare pattern seen in adapting structured and unstructured populations suggests a rugged fitness landscape in bacteria. *Proceedings of the National Academy of Sciences*, 112(24):7530–7535.

Oono, Y. (2012). *The nonlinear world: conceptual analysis and phenomenology*. Springer Science & Business Media.

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.

Perfeito, L., Pereira, M. I., Campos, P. R., and Gordo, I. (2008). The effect of spatial structure on adaptation in *Escherichia coli*. *Biology Letters*, 4(1):57–59.

Plum, A. M. and Baum, D. A. (2022). ACEs in spaces: Autocatalytic Chemical Ecosystems in Spatial Settings. arXiv:2212.14445 [q-bio].

Queller, D. C. and Strassmann, J. E. (2009). Beyond society: the evolution of organismality. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1533):3143–3155.

Ruiz-Mirazo, K., Briones, C., and de la Escosura, A. (2014). Prebiotic systems chemistry: new perspectives for the origins of life. *Chemical reviews*, 114(1):285–366.

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

Sokolskyi, T., Ganju, P., Montgomery-Taylor, R., and Baum, D. A. (2024). Evidence of Heritability in Prebiotically Realistic Membrane-Bound Systems. *Life*, 14(3):284. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

Stein, A., Gerstner, K., and Kreft, H. (2014). Environmental heterogeneity as a universal driver of species richness across taxa, biomes and spatial scales. *Ecology Letters*, 17(7):866–880.

Tkadlec, J., Pavlogiannis, A., Chatterjee, K., and Nowak, M. A. (2020). Limits on amplifiers of natural selection under death-Birth updating. *PLOS Computational Biology*, 16(1):e1007494.

Tomassini, M. (2005). *Spatially Structured Evolutionary Algorithms*. Natural Computing Series. Springer-Verlag, Berlin/Heidelberg.

Vincent, L., Berg, M., Krismer, M., Saghaei, S. T., Cosby, J., Sankari, T., Vetsigian, K., Cleaves, H. J., and Baum, D. A. (2019). Chemical Ecosystem Selection on Mineral Surfaces Reveals Long-Term Dynamics Consistent with the Spontaneous Emergence of Mutual Catalysis. *Life*, 9(4):80. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

Vostinar, A., Lalejini, A., Ofria, C., Dolson, E., and Moreno, M. A. (2024). Empirical: A scientific software library for research, education, and public engagement. *Journal of Open Source Software*, 9(98):6617.

Waxman, B. (1988). Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622.

Xavier, J. C., Hordijk, W., Kauffman, S., Steel, M., and Martin, W. F. (2020). Autocatalytic chemical networks at the origin of metabolism. *Proceedings of the Royal Society B*, 287(1922):20192377.