The Two Regimes of Neutral Evolution: Localization on Hubs and Delocalized Diffusion

David Shorten[1,2] and Geoff Nitschke[1,2]

[1] CAIR, CSIR Meraka, Building 43a, CSIR, Meiring Naude Road, Brummeria,
Pretoria, South Africa
[2] Department of Computer Science, University of Cape Town, Rondebosch 7700,
South Africa,
dshorten@cs.uct.ac.za, gnitschke@cs.uct.ac.za

Abstract. It has been argued that much of evolution takes place in the absence of fitness gradients. Such periods of evolution can be analysed by examining the mutational network formed by sequences of equal fitness, that is the neutral network. It has been demonstrated that, in large populations under a high mutation rate, the population distribution over the neutral network and average mutational robustness are given by the principle eigenvector and eigenvalue, respectively, of the network's adjacency matrix. However, little progress has been made towards understanding the manner in which the topology of the neutral network influences the resulting population distribution and robustness. In this work, we build on recent results from spectral graph theory and utilize numerical methods to demonstrate that there exist two regimes of behaviour: convergence on hubs and diffusion over the network. We also derive approximations for the population's behaviour under these regimes. This challenges the widespread assumption that neutral evolution always leads to exploration of the neutral network and elucidates the conditions which result in the evolution of robust organisms.

1 Introduction

When evolution reaches fitness plateaus, the evolutionary dynamics are governed by the topology of the neutral network [36]. In monomorphic populations, where the population size and mutation rate are low, the population performs a random walk on this network [38]. Conversely, in polymorphic populations, at equilibrium, the distribution of the population on the neutral network is given by the principle eigenvector of the network's adjacency matrix [36]. However, there is very little work examining the manner in which the topology of the neutral network influences this population distribution. This paper investigates polymorphic evolution on neutral networks by analyzing the influence of network topology on the principle eigenvector of the neutral network.

When organisms undergo natural evolution, mutation does not act directly on their form, but rather on the genetic code. Similarly, in *Evolutionary Computing* (EC), a *representation* of the problem, upon which mutation can occur, must

be identified. The problem of choosing such a representation, the *representation problem*, has been identified as a critical issue within EC [10], as well as artificial intelligence in general [24].

This necessitates a mapping between the genetic code or representation (genotype) and the organism or resulting problem solution (pheontype): the $G \rightarrow P$ map. the developmental process which translates genetic information into various biological organisms is not well understood [29]. Yet, it has become clear that this mapping is neither one-to-one nor linear [13]. In many organisms and *Ribonucleic Acid* (RNA) folding [9], it has been found that genetic change resulting from mutation is not proportional to phenotypic change [27,29,37]. Moreover, the $G \rightarrow P$ map is highly degenerate, that is many genotypes might encode for an identical phenotype [29].

There exists great variation in the mappings between representations and candidate solutions used in EC. On the one hand, in genetic algorithms, the relationship between representation and solution is often somewhat straightforward [10]. However, within the field of *generative and developmental systems* [7], many highly complicated mappings between representations and evolved forms have been proposed. Such mappings have been applied to a variety of tasks, including robot morphologies and organisms in artificial life studies [33]. Although the properties of individual mappings depend on their definition, some have been shown to be highly degenerate.

Degeneracy introduces the possibility that, when mutated, a genotype will still map to the same phenotype. This implies that the mutation has no effect on fitness and so can be labeled as *neutral*. Kimura [17], along with King and Jukes [18], brought the importance of neutral mutations to the attention of the scientific community through what has come to be known as the *neutral theory of molecular evolution*. This posits that the majority of evolutionary change is the result of the fixation of neutral mutations, as opposed to mutations which confer a selective advantage. Although the level of importance that such genetic drift has on evolution has been controversial [22], it is beyond doubt that certain mutations of certain organisms and structures are selectively neutral [4, 26, 39].

If the genetic code is a string of characters, as opposed to, say, a vector of real numbers, then one can construct networks out of genotypes coding for a given phenotype [36]. Here the vertices represent genotypes, and an edge connects two vertices if there exists a point mutation between their associated genotypes, that is their genetic codes are a hamming distance of one apart. Such *neutral networks* have been studied extensively [1, 4, 26, 36] and it has been shown that, under certain assumptions, these networks permeate sequence space and that any common phenotype can be reached by traveling along them [31].

An important associated concept is that of *mutational robustness* [34]. This refers to the proportion of mutations which leave the phenotype unchanged. The greater the mutational robustness of the genotypes, the larger their neutral networks will be [38]. This has an impact on the *evolvability* of these genotypes, as they can access a greater variety of phenotypes through neutral drift. Moreover,

populations can evolve so as to occupy the most connected parts of the network [36], thus increasing their average robustness.

If one assumes that evolution has reached a fitness plateau, that is that the fitness of all genotypes off the network is lower than that of those on it, then two behavioural regimes emerge. Given a population size M and a mutation rate μ, then if $M\mu \ll 1$ the population is *monomorphic* [3]. Mutations either fix or disappear, that is they either become present in the entire population or none of it. Thus, the entire population is concentrated on a single node of the neutral network. Throughout the neutral epoch the population performs a random walk over the network. On the other hand, if $M\mu \gg 1$, the population is polymorphic and spreads out over the neutral network [38]. Populations of self-replicating RNA, viruses and bacteria are polymorphic , whereas larger organisms are monomorphic [38]. Given the simple dynamics of the monomorphic case, this work focuses exclusively on polymorphic populations.

In their seminal work, van Nimwegen et. al. [36] showed that the equilibrium distribution of a polymorphic population is given by the principle eigenvector of the adjacency matrix of the neutral network and that the average robustness of the population is given by the principle eigenvalue. Despite the insight of this result, little work has been conducted towards determining the manner in which the topology of the neutral network influences the resulting population distribution over the network.

Reeves et. al. [30] were able to derive an upper limit to the principle eigenvalue in terms of the size of the network, by utilising the fact that neutral networks are subgraphs of a hypercube graph. This work, however, said nothing about the effect of other topological features and, moreover, has no implications for the principle eigenvector. Noirel and Simonson [26] were able to show, in simulation, that degree assortativity and the existence of hubs increased the average robustness of populations.

The principle eigenvectors and eigenvalues of graphs are of great importance to a variety of problems [32], principally synchronization phenomena and the spread of epidemics. Since the publication of van Nimwegen et. al.'s seminal paper, there has been significant progress towards describing the behaviour of these two quantities in terms of network topology [5, 14, 19, 28]. To the best of the authors' knowledge, there has been no work published which examines the implications of these results on the neutral evolution of polymorphic populations.

In this paper, we build on the above-mentioned results, both analytically and numerically, in order to elucidate the effect of neutral network topology on the equilibrium distribution and average robustness of polymorphic populations evolving on neutral networks. The principle finding is that there are two distinct behavioural regimes. If the network contains a hub of sufficiently high degree, then the population localizes on this hub. That is the vast majority of the population is found on the hub node and its neighbours. The neutral networks of proteins have been found to have high degree hubs and simulations of neutral evolution have shown that populations converge on these hubs [4, 26, 39]. However, these models also incorporated the stability of the proteins, which acts as

a type of fitness. In most instances, the stability of the sequences correlates with their neutral degree in what is known as a *superfunnel*. Here we show, generally, that this behaviour occurs in the absence of a fitness advantage conferred by hubs so long as the hubs' degree is sufficiently large. Moreover, we demonstrate that, in large networks, the extent of the localization can be much greater than in the smaller networks analyzed in these studies.

This mode of behaviour casts the discussion on the relationship between robustness and evolvability in a new light. Many arguments are based on the assumption that populations spread out over the neutral network. While this is true for monomorphic populations and polymorphic populations evolving on networks that lack high degree hubs, we demonstrate that polymorphic populations evolving on networks with hubs cluster within a very small region of the network. Specifically, there have been two main arguments for how robustness facilitates evolvability. In the first, it is proposed that, when polymporphic populations spread out over the network, the population gains cryptic variation [20]. This variation allows the population to better adapt to changes in the environment. In the second, it is argued that robust genotypes create larger neutral networks. This creates more "stepping off points", and so the population can access more phenotypic variation [37]. Neither of these arguments hold if the population is tightly clustered around a hub. Furthermore, as demonstrated below, when the population clusters around a hub, its average robustness can be substantially higher than the average degree of the network. Thus, this is a case in which robustness and evolvability have a firmly antagonistic relationship.

The second behavioural regime encountered by polymorphic populations is diffusion over the network. We show that such populations are distributed roughly evenly over the network and we derive an expression for the average mutational robustness of the population. This expression shows that the population's robustness is largely determined by two biases: the edge sampling bias caused by mutations and the degree assortativity of the network. Although robustness itself represents a type of mutational bias, one can question the existence and role of higher order mutational biases, that is biases towards biases. Indeed, mutations on the neutral network, that is mutations that lead to viable genotypes, are biased towards higher degree nodes due to the friendship paradox [12]. This effect is named after the phenomenon where, in social networks, the average number of friends of friends is higher than the average number of friends. Moreover, this effect is present in all networks, where the average number of neighbors of neighbors is higher than the average number of neighbors of nodes in the network. The cause of this paradox is that sampling the degrees of neighbors is equivalent to sampling the degrees of nodes at the end of edges, which is biased towards higher degree nodes. The relationship between these two averages can be expressed as: [12]

$$\hat{\lambda} = \langle k \rangle + \frac{\sigma_n^2}{\langle k \rangle} = \frac{\langle k^2 \rangle}{\langle k \rangle} \tag{1}$$

where $\langle k \rangle$ is the average degree (robustness) of genotypes on the neutral network, σ_n^2 is the variance of these degrees and $\hat{\lambda}$ is the average degree of single mutation neighbors. An implication of this result, as demonstrated by van Nimwegen et. al. [36], is that random walks on neutral networks result in an average neutrality equal to $\hat{\lambda}$.

Intuitively, we would not expect populations to converge on an average level of robustness substantially lower than what a random walk provides. Although robust genotypes have a selective advantage in that they produce more viable offspring, if these offspring themselves are not robust it is difficult to see how the population could converge on this lineage. Therefore, the selection of robustness is facilitated by the existence of highly robust nodes whose offspring are also highly robust. This sort of higher order mutational bias is provided by network assortativity, that is, correlation in the degrees of the nodes at the end of edges [23]. By deriving an expression for the average population robustness, we show that it is equal to the mutational sampling bias and rises above or below this figure depending on whether the network has positive or negative degree assortativity.

2 Localization on Hubs

In the context of graph spectra, localization refers to the phenomenon whereby the normalisation weight of an eigenvector ($\sum f_i^2(\lambda)$, where λ is the eigenvalue and $\boldsymbol{f}(\lambda)$ is the eigenvector) is concentrated on a small number of nodes that does not scale with the size of the network [28]. Some authors have suggested using the inverse participation ratio $Y(\lambda)$.

$$Y(\lambda) = \sum_{i=1}^{N} f_i^4(\lambda) \tag{2}$$

as a quantitative measure of localization where, in this case, $\boldsymbol{f}(\lambda)$ is the normalised eigenvector. If, in the limit $N \to \infty$, $Y(\lambda) \sim 1$ then the state is localized. On the other hand, if $Y(\lambda) \to 0$ then the state is delocalized. There are a number of results relating aspects of network topology to localization. Chung et. al. [5] showed that the principle eigenvalue, for a random graph model characterised by a given degree distribution, is given by.

$$\lambda_1 = \begin{cases} \hat{\lambda}, & \hat{\lambda} > \sqrt{k_{max}} \log N \\ \sqrt{k_{max}}, & \sqrt{k_{max}} > \hat{\lambda} \log^2 N \end{cases} \tag{3}$$

where $\hat{\lambda} = \langle k^2 \rangle / \langle k \rangle$ ($\langle k \rangle$ being the average degree and $\langle k^2 \rangle$ being the mean of the squares of the degrees). $\lambda_1 = \hat{\lambda}$ corresponds to the delocalized state and $\lambda_1 = \sqrt{k_{max}}$ corresponds to the localized state.

Goltsev et. al. [14] showed that, for unassortative scale-free networks with degree distribution $P(k) \sim k^{-\gamma}$, the principle eigenstate is localized for $\gamma > \frac{5}{2}$

and delocalized otherwise. The principle eigenvalue is given by $\sqrt{k_{max}}$ and $\hat{\lambda}$ for the localized and delocalized states, respectively.

Martin et. al. [19] demonstrated that for a hub connected to an Erdős-Renyi network, localization occurs when $\sqrt{k_{max}} > \langle q \rangle$ where $\langle q \rangle$ is the average degree of the original Erdős-Renyi network, without the hub. Furthermore, they showed that the eigenvector component on the hub, f_h is given by.

$$f_h = \sqrt{\frac{k_{max} - 2\langle q \rangle}{2k_{max} - 2\langle q \rangle}} \tag{4}$$

Where the average of the components neighbouring the hub, $\langle f_n \rangle$ is given by.

$$\langle f_o \rangle = \frac{f_h}{\sqrt{k_{max} - \langle q \rangle}} \tag{5}$$

and the average of all non-hub components $\langle f_j \rangle$ is.

$$\langle f_j \rangle = \frac{1}{N-1} \frac{f_h}{\sqrt{k_{max} - \langle q \rangle}} \tag{6}$$

Finally, Pastor-Satorras and Castellano [28] have shown that a form of less severe localization can occur on scale-free networks where $\gamma < \frac{5}{2}$.

This then begs the question of whether the neutral networks encountered in natural and artificial evolution meet the topological criteria for localization. Given the wide variety of possible fitness landscapes, it is fair to assume that at least some of them will contain neutral networks with localized principle eigenvectors. However, there is a dearth of mapped-out neutral networks. The authors know of none within EC. Fortunately, some neutral networks of protein [4,26,39] and RNA folding [2] have been mapped. RNA neutral networks would seem to be fairly homogeneous, with narrow degree distributions. However, the neutral networks induced by protein folding contain high-degree hubs. This makes them candidates for localization behaviour.

Moreover, it was reported that localization-like behaviour was observed when evolution was simulated on these neutral networks. However, those simulations incorporated the stability of the proteins, which acts as a type of fitness. The stability of the proteins was strongly correlated with the robustness, that is the highest degree node also had the greatest stability. This has been labeled as the *superfunnel* paradigm [26] and it contains the further assumption that robustness and, by implication, the stability of the sequences decreases with increasing distance from the hub, or 'prototype' sequence. This can be visualized as a funnel, with the bottom placed over the prototype sequence.

The above results concerning the localization of eigenvectors demonstrate that the localization of populations can occur in the absence of the fitness advantage conferred by stability and the anticorrelation of robustness and distance from the prototype sequence. The crucial feature is high degree hubs. Although none of the networks reported on in these studies satisfy the stringent conditions

of Chung et. al. [5], some of those analyzed by Bornberg-Bauer [4] easily satisfy the criteria of Martin et. al. [19] that $\sqrt{k_{max}} > \langle q \rangle$. However, it is unclear how close the topology of the observed networks is to the model of Martin et. al. [19]. Although we do not have access to the full topology of all of the networks studied by these authors, inspection of those for which they presented diagrams leads us to believe that they do conform to this model. Moreover, it is worth bearing in mind that the model of Martin et. al. [19] is considering hubs connected to a *maximally random* network.

2.1 Connected Hubs

The topology of the neutral network of the haemagglutinin protein of the influenza A virus (H3N2) as explored by Wagner [39] appears to conform to a somewhat different topology. Although it easily satisfies the criteria that $\sqrt{k_{max}} > \langle q \rangle$, it would seem to be composed of hubs attached to one another.

In order to shed light on this type of topology, we studied the localization behaviour of a model of random networks, whereby hubs (star networks) were connected by non-preferential attachment. Specifically, a high degree hub of degree m was instantiated by connecting m nodes to a hub node. Further to this, 30 lower degree hubs, of degree $n = 5$ were created. All the hubs were then connected through non-preferential attachment, beginning with the maximum degree hub. Specifically, one low degree hub was connected to the high degree hub and then each subsequent low degree hub was connected to a randomly chosen hub in this connected graph. Figure 1 shows that the localization transition occurs at around $m = 35$.

3 Delocalized Regime

In the delocalized regime, progress on approximating the population's distribution and robustness can be made by assuming that, at equilibrium, for every node in the network, the average population concentration on nodes at a given distance l is equal. That is we utilise a mean-field approximation at a given distance l. This average concentration is the uniform concentration, that is the population size divided by the number of nodes. This is equivalent to assuming that the correlation length for the degrees is low. It has been found that, for most real-world networks, the correlation length is low [21]. Using this assumption we can approximate the proportion of the population which mutates onto a given node, and hence the population distribution and average robustness.

For the cases $l = 2$ and $l = 3$ we make use of the *annealed network approximation* [8], whereby all nodes with a given degree k are approximated as having the same nearest neighbour degree distribution, which is the aggregate distribution over the neighbours of all nodes with degree k. This has the implication that all nodes of degree k have the same average nearest neighbours degree, that is $\bar{k}_{nn}(i) = \bar{k}_{nn}(k_i)$, where i is a node's index and k_i the associated degree. We
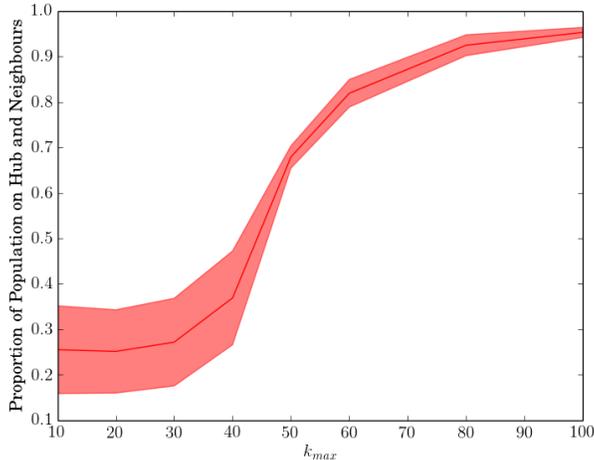
Fig. 1. The proportion of the population found on the largest hub node and its neighbours for our model of connected-hub networks. 30 star networks of degree 5 and a larger star whose degree is plotted on the x axis were connected via non-preferential attachment. The shaded region shows the standard deviation.

also use the approximation:

$$\bar{k}_{nn}(k) \approx \hat{\lambda} + \left(k - \hat{\lambda}\right) r \tag{7}$$

Where $\bar{k}_{nn}(k)$ is the average nearest-neighbour degree of nodes of degree k, $\hat{\lambda} = \langle k^2 \rangle / \langle k \rangle$ ($\langle k \rangle$ being the average degree $\langle k^2 \rangle$ being the average of the squares of the degrees) and r is the assortativity coefficient (the Pearson correlation between the degrees at either end of an edge) [23]. This approximation is derived by considering that r is the root of the coefficient of determination of the linear regression between the degrees of the nodes at either end of an edge.

We introduce the notation λ_1^l to denote the approximation of the principle eigenvalue (population average robustness) based on the assumption of equal average distribution at distance l. Similarly, we use $f_i\left(\lambda_1^l\right)$ to denote the i^{th} component of the principle eigenvector (the proportion of the population having the genotype represented by the i^{th} node), based on the assumption of equal average distribution at distance l.

Some of the below reasoning is based on the particulars of the model of van Nimwegen et. al. [36]. In this model, the population of constant size M resides on a neutral network of size N. The total number of neighbours, neutral and non-neutral, that a given genotype can have is given by U, this limit is determined by the length of the genetic code and the size of the alphabet. Each generation, M genotypes are selected with replacement from the population. These individuals then undergo mutation. With probability k_i/U the individual

remains on the network, where k_i is the degree of the node representing the individual's genotype. If the individual stays on the network, it moves to one of its neighbouring nodes, chosen at random. If it mutates off the network then it is ineligible for selection in the subsequent generation.

3.1 Zero-hop Case

The simplest case is that we assume that the average population concentration at a distance zero from each node is equal, that is we assume that the population is uniformly distributed. The average robustness of the population is therefore, trivially, the average degree. Thus, we have:

$$f_i\left(\lambda_1^0\right) = \frac{1}{N} \tag{8}$$

$$\lambda_1^0 = \langle k \rangle \tag{9}$$

Where $\langle k \rangle$ is the average degree and N is the size of the network.

3.2 One-hop Case

The next case assumes that, at equilibrium, the average population concentration of the neighbours of each node are equal. Therefore, each generation, an average of $k_i M/NU$ individuals mutate onto node i. Normalizing, we arrive at

$$f_i\left(\lambda_1^1\right) = \frac{k_i}{N\langle k \rangle} \tag{10}$$

Multiplying by the robustness (k_i) and summing over all the nodes we arrive at the average robustness of:

$$\lambda_1^1 = \frac{\langle k^2 \rangle}{\langle k \rangle} = \hat{\lambda} \tag{11}$$

3.3 Two-hop Case

If we assume an average uniform population concentration two hops from each node, then, each generation, by the annealed network approximation, the nodes neighbouring node i will receive, on average, $\bar{k}_{nn}(k_i) M/NU$ mutants. This implies that node i will receive $k_i \bar{k}_{nn}(k_i) M/NU^2$ individuals. Substituting in equation (7) and normalizing we arrive at

$$f_i\left(\lambda_1^2\right) = \frac{1}{N\langle k^2 \rangle}\left(k_i\hat{\lambda} + k_i\left(k_i - \hat{\lambda}\right)r\right) \tag{12}$$

Multiplying through by the node's robustness (k_i) and summing over the nodes, we arrive at

$$\lambda_1^2 = \hat{\lambda} + \frac{r\sigma_e^2}{\hat{\lambda}} \tag{13}$$

Where $\sigma_e^2 = \langle k^3 \rangle / \langle k \rangle - \langle k^2 \rangle^2 / \langle k \rangle^2$ is the variance of the node's degrees when sampled by following edges. (13) is equivalent to the approximation derived by Goltsev et. al. [14] through the use of a power iteration.

3.4 Three-hop Case

Our final approximation is based on the assumption that, from any given node, the average population density at nodes three hops away is equal. We consider the node i', a neighbour of i. Each generation, this node will receive an average of $k_{i'} \bar{k}_{nn}(k_{i'}) M/NU^2$ mutants from its neighbours. We then average this over all neighbours i' of node i, that is we want to find

$$I = \frac{M}{NU^2} \left\langle k_{i'}\hat{\lambda} + k_{i'}^2 r - k_{i'}\hat{\lambda}r \right\rangle_{i'} \tag{14}$$

Using the fact that $\left\langle k_{i'}^2 \right\rangle_{i'} \approx \sigma_e^2 + \langle k_{i'} \rangle_{i'}^2$, where the equality is approximate as σ_e^2 is the global variance and not specific to the neighbours of nodes of degree $k_{i'}$, we can arrive at

$$I \approx \frac{M}{NU^2} \left(\hat{\lambda}^2 + \hat{\lambda}\left(k_i - \hat{\lambda} \right)r + \hat{\lambda}\left(k_i - \hat{\lambda} \right)r^2 + \left(k_i - \hat{\lambda} \right)^2 r^3 + \sigma_e^2 r \right) \tag{15}$$

The number of mutants that a node i receives is $k_i I/U$. When we come to normalise this, we find that the total population is

$$P \approx \frac{M}{U^3} \left(\langle k \rangle \hat{\lambda}^2 + \left\langle k \left(k - \hat{\lambda} \right)^2 \right\rangle r^3 + \langle k \rangle \hat{\sigma}_e^2 r \right) \tag{16}$$

The second two terms in the parentheses are much smaller than the first and so, for mathematical expediency, we ignore them. This results in

$$f_i\left(\lambda_1^3\right) \approx \frac{1}{N \langle k^2 \rangle} \left(k_i\hat{\lambda} + k_i\left(k_i - \hat{\lambda} \right)r + k_i\left(k_i - \hat{\lambda} \right)r^2 + \frac{k_i\left(k_i - \hat{\lambda} \right)^2 r^3}{\hat{\lambda}} + \frac{k_i\sigma_e^2 r}{\hat{\lambda}} \right) \tag{17}$$

As previously, we multiply by each node's robustness (k_i) and sum over all nodes to arrive at the approximation for the eigenvalue (population average robustness).

$$\lambda_1^3 \approx \hat{\lambda} + \frac{2r\sigma_e^2}{\hat{\lambda}} + \frac{r^2\sigma_e^2}{\hat{\lambda}} + \frac{r^3 \left(\langle k^4 \rangle - 2\hat{\lambda}\langle k^3 \rangle + \hat{\lambda}^2 \langle k^2 \rangle \right)}{\langle k^2 \rangle \hat{\lambda}} \tag{18}$$

3.5 Numerical Verification

Random networks conforming to the Erdős-Renyi model [11] were created in order to measure the accuracy of the approximations. We needed to test these approximations for various values of the assortativity coefficient (r), however, the expected value of r for Erdős-Renyi networks is 0 [23]. In order to both increase and decrease the assortativity, the rewiring algorithm of van Miegham et. al. [35], was used. This algorithm operates by iteratively picking two edges at

random and observing the degrees of the four nodes at their endpoints. If the goal is to increase assortativity, the two nodes with the highest degrees are connected by an edge and, likewise, the lowest degree nodes are also connected. Moreover, the original two edges are removed from the network. Similarly, if the goal is to decrease assortativity, the maximum degree node is connected to the minimum degree node and the remaining two nodes are also connected. Rewiring does not take place if the desired connectivity arrangement between the four nodes was already present. This iterative process is repeated until the desired value of r is achieved.

We chose 15 values of r on which to test: those between -0.7 and 0.7, inclusive, at intervals of 0.1 . 100 networks, each with $N = 1000$ nodes and average degree of $\langle k \rangle = 10$ were instantiated for each of these values. Figure 3.5 shows the relative error of the derived approximations. It shows that, for positive values of r, the approximation λ_1^3 is more accurate than that of Goltsev et. al. [14] (which is equivalent to λ_1^2).
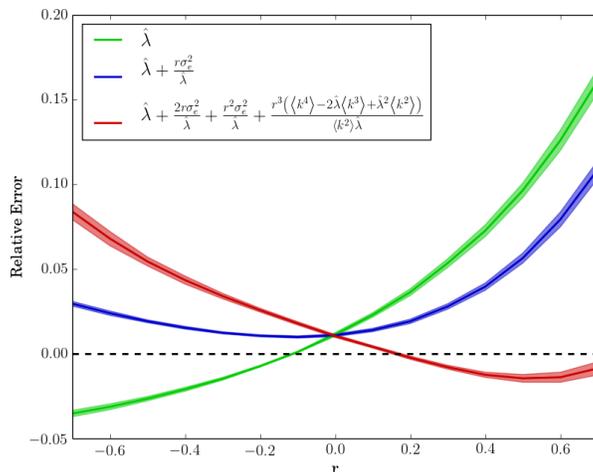
Fig. 2. The relative error of the approximations of the principle eigenvalue measured on Erdős-Renyi networks which have been subjected to a rewiring algorithm [35] in order to display various values of degree assortativity. The shaded region shows the standard deviation.

4 Discussion

In this work, we set out to incorporate and build upon recent results concerning the behaviour of the principle eigenvectors, and associated eigenvalues, of the

adjacency matrices of networks in the context of the study of the dynamics of polymorphic populations evolving on neutral networks.

Much of the discussion surrounding neutral evolution has functioned on the assumption that there is only one regime of polymorphic neutral evolution. In this regime, the population explores much of the network and, in the words of van Nimwegen et. al. [36] "seeks out the most connected areas of the neutral network". This intuition is echoed by many authors, for instance "the population will tend to congregate in regions of a neutral network that have more robust genotypes" [38] and the population will "evolve toward regions denser in neutral genotypes" [1]. Moreover, this is plausible, given that the principle eigenvector (which specifies the population's distribution) is used as a measure of centrality: the *eigenvector centrality* [19].

Our most salient finding is that there are actually two distinct behavioural regimes. Moreover, each of these regimes differ from the above intuition in important ways.

The first regime involves the localization of the population onto a hub and its neighbouring nodes. Here the average mutational robustness of the population is approximated by the square root of the degree of the hub and the number of nodes upon which the bulk of the population resides does not scale with the size of the network.

This has important ramifications for the understanding of neutral evolution. Firstly, under localization, the eigenvector centrality ceases to be a useful measure [19]. This can be illuminated by considering a star network connected to an Erdős-Renyi [11] network by a single edge. Following the arguments outlined in section 2, so long as the degree d of the star's hub is greater than the average degree of the original network $\langle q \rangle$, the population will localize on the star. As a star network has one more edge than vertices (in this case two more, due to the connection to the Erdős-Renyi network), the average degree of this network is approximately two. Therefore, so long as the original Erdős-Renyi network has an average degree greater than two, the majority of the population will be concentrated on a region of the network with below average degree. It is thus concentrated on a region which is neither "most connected" nor "denser in neutral genotypes".

It is useful to query why our intuition failed in this case. After all, the population's average robustness $(\sqrt{k_{max}})$ can be much higher than the network's average degree, even though the population is occupying a region of below average degree. This is due to the fact that such a large proportion of the population is found on the hub itself. Referring again to section 2, we see that, in the case of a star connected to an Erdős-Renyi network, in the limit of large d, around half of the entire population is found on the hub genotype. This concentration is self-reinforcing, as a substantial proportion of the hub's mutants will mutate back onto it, leading to the localization phenomenon [19].

This regime of behaviour has particular relevance to arguments concerning the relationship between robustness and evolvability. These arguments are predicated on the fact that robust genotypes form larger neutral networks [38]. This

then allows for the population to accumulate more cryptic variation as it spreads over the network [20], allowing it to better adapt to changes in its environment. Moreover, it creates more "stepping off points" for the population, allowing it to access more phenotypic variation [37]. However, these arguments fail to take into account the proportion of the network occupied by the population. By definition, localization occurs when the normalization weight of the principle eigenvector is concentrated at a number of nodes that does not scale with the size of the network [28]. Therefore, the cryptic variation in the population and the amount of phenotypic variation accessible to it does not scale with the robustness of the genotypes of which the network is composed. Furthermore, as the population's average robustness is determined solely by the degree of the hub, it is entirely possible to have populations with extremely high average robustness, and very low cryptic variation and access to phenotypic variation. The access to phenotypic variation is particularly poor in this regime, given that such a large proportion of the population is concentrated on the hub and that the hub has such a high proportion of neutral neighbours. It is worth noting that these arguments relating robustness and evolvability are still valid, so long as the robustness of genotypes is homogeneously distributed, in order to avoid localization.

The delocalized case has more in common with the traditional intuition. A principle difference, however, is the level to which concentration on regions of better connected genotypes occurs. Firstly, such a region needs to exist. This requirement will be met in networks with degree assortativity, however, disassortative mixing will result in genotypes with high robustness mutating to those with low robustness, thwarting evolution's attempts at settling on robust nodes. Specifically, by examining equations (10), (12) and (17) we see that, in unassortative networks, the proportion of the population on a given node scales with its degree. This implies that, on relatively homogeneous networks, there will be little difference in the population concentration on various nodes. Furthermore, as shown in (17), disassortative mixing decreases the number of individuals occupying a node in proportion to both the square and cube of its degree. Although we do expect to see a certain degree of concentration of the population in more robust regions of the neutral network in the case that the network exhibits assortative mixing, the severity of this concentration will be relatively mild. Given that, for networks of reasonably high average degree, $\hat{\lambda}$ is substantially larger than r, the latter terms in equation (17) will only play a significant role when k_i is much larger than $\hat{\lambda}$.

We further propose that it is fruitful to think of the delocalized regime of neutral evolution more in terms of a biased sampling process of the genotypes on the network, resulting from mutational biases, as opposed to a population moving between regions of the network. From equations (11), (13) and (18) we can see that, in the absence of assortativity, the population's average robustness is approximated by $\hat{\lambda}$. This is exactly the average robustness which we would expect from performing a random sampling of all possible mutations on the network, as implied by the friendship paradox [12]. Assortative and disassorative mixing by degree will increase or decrease the population's average robustness

above or below this level. Assortativity represents a further mutational bias towards higher or lower degree nodes, dependent on the degree of the node from which the mutation originates.

Given that it is suspected that much of evolution occurs on neutral networks [22] along with the importance of mutational robustness to the survival of organisms and its relationship with evolvability, understanding the impact of the topology of neutral networks on the dynamics of neutral evolution and the resulting robustness of organisms is of great importance. This work has provided insight into these issues in the case of polymorphic populations: large populations evolving at high mutation rate. The directed, neutral, evolution of bio-molecules [6, 15] along with viruses overcoming immunity through neutral evolution [25] fall within this category. These results have potential applicability to these problems. For instance, the neutral evolution of large libraries of molecules [16] will be greatly aided by delocalization, whereas a virus's attempt to escape immunity might be thwarted if its population localizes on a hub.

5 Conclusion

This paper investigated the manner in which neutral network topology influences the resulting population distribution and robustness during neutral evolution at high mutation rates in large populations. In such cases, the population distribution is given by the principle eigenvector of the adjacency matrix of the neutral network and, similarly, the average mutational robustness of the individuals in the population is given by the principle eigenvalue [36]. Hence, we utilized, and built upon, recent results concerning the behaviour of these values from studies concerning the spread of epidemics on networks [14] as well as more general work [19].

It was found that, on homogeneous neutral networks, the population's behaviour could be described in terms of mutational biases. For unassortative neutral networks, it was found that the average mutational robustness was equal to the sampling bias provided by the friendship paradox [12]. Assortative and disassortative mixing by degree raised the robustness above or below this value, respectively. Furthermore, in the process of demonstrating this, we derived a new approximation for the principle eigenvalue of a network in terms of its assortativity and the moments of its degree distribution.

Conversely, for heterogeneous neutral networks with high degree hubs, it was found that the population become concentrated on the nodes around the hub and thus the number of nodes occupied does not scale with the size of the network. Furthermore, the average robustness of the population is given by the square root of the network's maximum degree. These results are particularly relevant to various arguments concerning the relationship between robustness and evolvability [20, 37], which assume that the number of nodes occupied by the population scales with the size of the network.

These results are relevant to the directed evolution of bio-molecules [6, 15], where they can be used to evolve more robust molecules as well as facilitate the

evolution of greater variety. Moreover, they can also further our understanding of the factors that allow viruses to escape immunity along neutral networks [25].

References

1. Aguirre, J., Buldú, J.M., Manrubia, S.C.: Evolutionary dynamics on networks of selectively neutral genotypes: Effects of topology and sequence stability. Physical Review E 80(6), 066112 (2009)
2. Aguirre, J., Buldú, J.M., Stich, M., Manrubia, S.C.: Topological structure of the space of phenotypes: the case of rna neutral networks. PloS one 6(10), e26324 (2011)
3. Bloom, J.D., Lu, Z., Chen, D., Raval, A., Venturelli, O.S., Arnold, F.H.: Evolution favors protein mutational robustness in sufficiently large populations. BMC biology 5(1), 1 (2007)
4. Bornberg-Bauer, E.: How are model protein structures distributed in sequence space? Biophysical Journal 73(5), 2393 (1997)
5. Chung, F., Lu, L., Vu, V.: Spectra of random graphs with given expected degrees. Proceedings of the National Academy of Sciences 100(11), 6313–6318 (2003)
6. Currin, A., Swainston, N., Day, P.J., Kell, D.B.: Synthetic biology for the directed evolution of protein biocatalysts: navigating sequence space intelligently. Chemical Society Reviews 44(5), 1172–1239 (2015)
7. Devert, A.: When and why development is needed: generative and developmental systems. In: Proceedings of the 11th Annual conference on Genetic and evolutionary computation. pp. 1843–1844. ACM (2009)
8. Dorogovtsev, S.N., Goltsev, A.V., Mendes, J.F.: Critical phenomena in complex networks. Reviews of Modern Physics 80(4), 1275 (2008)
9. Draper, D.: The rna-folding problem. Acc. Chem. Res. 25(4), 201–207 (1992)
10. Eiben, A., Smith, J.: Introduction to Evolutionary Computing. Springer, Berlin, Germany (2003)
11. Erdős, P., Renyi, A.: On random graphs i. Publ. Math. Debrecen 6, 290–297 (1959)
12. Feld, S.: Why your friends have more friends than you do. American Journal of Sociology pp. 1464–1477 (1991)
13. Gjuvsland, A., Vik, J., Beard, D., Hunter, P., Omholt, S.: Bridging the genotype-phenotype gap: what does it take? J Physiol. 591(8):, 2055–2066 (2013)
14. Goltsev, A.V., Dorogovtsev, S.N., Oliveira, J., Mendes, J.F.: Localization and spreading of diseases in complex networks. Physical review letters 109(12), 128702 (2012)
15. Jäckel, C., Hilvert, D.: Biocatalysts by evolution. Current opinion in biotechnology 21(6), 753–759 (2010)
16. Kaltenbach, M., Tokuriki, N.: Generation of effective libraries by neutral drift. Directed Evolution Library Creation: Methods and Protocols pp. 69–81 (2014)
17. Kimura, M., et al.: Evolutionary rate at the molecular level. Nature 217(5129), 624–626 (1968)
18. King, J.L., Jukes, T.H.: Non-darwinian evolution. Science 164(3881), 788–798 (1969)
19. Martin, T., Zhang, X., Newman, M.: Localization and centrality in networks. Physical Review E 90(5), 052808 (2014)

20. Masel, J., Trotter, M.V.: Robustness and evolvability. Trends in Genetics 26(9), 406–414 (2010)
21. Mayo, M., Abdelzaher, A., Ghosh, P.: Long-range degree correlations in complex networks. Computational Social Networks 2(1), 1 (2015)
22. Nei, M.: Selectionism and neutralism in molecular evolution. Molecular biology and evolution 22(12), 2318–2342 (2005)
23. Newman, M.: Assortative mixing in networks. Physical review letters 89(20), 208701 (2002)
24. Nilsson, N.J.: The quest for artificial intelligence. Cambridge University Press (2009)
25. van Nimwegen, E.: Influenza escapes immunity along neutral networks. Science 314(5807), 1884–1886 (2006)
26. Noirel, J., Simonson, T.: Neutral evolution of proteins: The superfunnel in sequence space and its relation to mutational robustness. The Journal of chemical physics 129(18), 185104 (2008)
27. Parter, M., Kashtan, N., Alon, U.: Facilitated variation: how evolution learns from past environments to generalize to new environments. PLoS Comput Biol 4(11), e1000206 (2008)
28. Pastor-Satorras, R., Castellano, C.: Distinct types of eigenvector localization in networks. Scientific reports 6 (2016)
29. Pigliucci, M.: Genotype–phenotype mapping and the end of the 'genes as blueprint'metaphor. Philosophical Transactions of the Royal Society of London B: Biological Sciences 365(1540), 557–566 (2010)
30. Reeves, T., Farr, R., Blundell, J., Gallagher, A., Fink, T.: Eigenvalues of neutral networks: Interpolating between hypercubes. Discrete Mathematics 339(4), 1283–1290 (2016)
31. Reidys, C., Stadler, P.F., Schuster, P.: Generic properties of combinatory maps: neutral networks of rna secondary structures. Bulletin of mathematical biology 59(2), 339–397 (1997)
32. Restrepo, J.G., Ott, E., Hunt, B.R.: Approximating the largest eigenvalue of network adjacency matrices. Physical Review E 76(5), 056119 (2007)
33. Stanley, K.O., Miikkulainen, R.: A taxonomy for artificial embryogeny. Artificial Life 9(2), 93–130 (2003)
34. Taverna, D.M., Goldstein, R.A.: Why are proteins so robust to site mutations? Journal of molecular biology 315(3), 479–484 (2002)
35. Van Mieghem, P., Wang, H., Ge, X., Tang, S., Kuipers, F.: Influence of assortativity and degree-preserving rewiring on the spectra of networks. The European Physical Journal B 76(4), 643–652 (2010)
36. Van Nimwegen, E., Crutchfield, J., Huynen, M.: Neutral evolution of mutational robustness. Proceedings of the National Academy of Sciences 96(17), 9716–9720 (1999)
37. Wagner, A.: Robustness and evolvability: a paradox resolved. Proceedings of the Royal Society of London B: Biological Sciences 275(1630), 91–100 (2008)
38. Wagner, A.: The origins of evolutionary innovations: a theory of transformative change in living systems. OUP Oxford (2011)
39. Wagner, A.: A genotype network reveals homoplastic cycles of convergent evolution in influenza a (h3n2) haemagglutinin. Proceedings of the Royal Society of London B: Biological Sciences 281(1786), 20132763 (2014)