



The evolution of neural-network movement strategies in fluctuating environments

Luis Fernando Valdés Villarreal Boullosa

S3019632

September 2020



rijksuniversiteit
 groningen



Supervised by:

Main supervisor

Prof. Dr. Franz J. Weissing

f.j.weissing@rug.nl

Model implementation

Dr. Hanno Hildenbrandt

h.hildenbrandt@rug.nl

Faculty of Science and Engineering

GELIFES — Groningen Institute for Evolutionary Life Sciences

Nijenborgh 7

9747 AG Groningen, The Netherlands

Table of Contents

1 Introduction.....	1
1.1 The importance of movement for survival.....	1
1.2 Evolution in fluctuating environments.....	1
1.3 Movement strategies.....	1
1.4 Network models for the evolution of plasticity.....	2
1.5 Aims and approach.....	2
1.6 Structure of the report.....	3
2 Methods.....	4
2.1 Model purpose.....	4
2.2 The landscape and its dynamics.....	4
2.3 Resource landscape.....	5
2.4 Neural networks.....	7
2.5 Detailed model description.....	12
2.6 Model output.....	18
2.7 Model implementation.....	18
2.8 Simulations.....	19
3 Results.....	19
3.1 A typical evolved reaction norm.....	20
3.2 Reaction norms evolved in all architectures.....	21
3.3 Movement strategies coexist with sessile strategies.....	24
3.4 Uninformed individuals tend to stay where they are when environmental change is very slow. .	25
4 Discussion.....	28
4.1 The importance of ecological inheritance.....	28
4.2 Coexistence of several strategies.....	28
4.3 Neural networks and plasticity.....	28
4.4 Response to fluctuating environments.....	29
4.5 Movement and decision making.....	29
4.6 Future developments, pros and cons of theoretical models and assumptions.....	29
5 Acknowledgments.....	30
6 Bibliography.....	30

Abstract

Movement is an important survival mechanism among species living in changing environments. This is all the more important when environmental change can be predicted from present or past information. Here I view condition-dependent movement as a special case of phenotypic plasticity. Previous research predicts that phenotypic plasticity will only evolve when the right conditions are met: when environmental change is sufficiently fast to be noticeable within the lifetime of an individual, when the direction and/or degree of change can to a certain extent be predicted on the basis of some cues, and when the costs of plasticity are not too high. In the present work, I investigate whether, and to what extent, these predictions apply to movement. This is not self-evident, because standard theory only considers fluctuations in time, while condition-dependent movement takes place under temporally and spatially changing conditions. In my model, the movement decisions are taken by evolvable neural networks. Each individual harbours such a network, which processes local information on resource densities to determine the direction and/or step length of movement. The networks are transmitted from parents to their offspring (subject to some mutation). Individuals that take the most adequate movement decision can gather more resources and produce more offspring. This way, the movement-determining neural networks evolve over the generations. For simplicity, I neglected local resource depletion in my simulations. Still my simulations led to quite complicated eco-evolutionary dynamics. Moreover, I regularly observed polymorphisms, such as the dynamical coexistence of sessile and motile individuals. Although time did not permit a full analysis of the model, I can conclude that the evolution of condition-dependent movement is strongly affected by the pattern of ecological heterogeneity and the mechanisms underlying movement decisions.

1 Introduction

1.1 The importance of movement for survival

Movement, either through active or passive processes, is a universal aspect of life on Earth. From the microscopic movements of bacteria to the great animal migrations or the large-scale dispersal of plant seeds or pollen, movement allows organisms to find nutrients, mating partners, or favourable living conditions, and to avoid competition and predation. Even viruses differ in their ability to surf aerosols and droplets. Movement is essential when an organism is either not born in the ideal conditions, or when the conditions deteriorate as a result of its own metabolism, the metabolism of others, or extrinsically-caused variations in nutrients, temperature, light availability or toxic substances.

1.2 Evolution in fluctuating environments

Fluctuating environments pose a particular challenge to organisms. From tidal systems, to day and night cycles, to seasons, repeating weather events and recurrent predators, fluctuating systems are everywhere. Populations living today must have evolved the ability to cope with changes in their environments. There are different ways to cope with fluctuating conditions (so-called ‘modes of adaptation’), such as bet-hedging, phenotypic plasticity, or adaptive tracking. Previous theoretical work (e.g. Botero et al. 2015) has shown that (at least) two extrinsic factors are crucial for the evolution of an ‘adequate’ mode of adaptation: the rate and predictability of environmental change. If the fluctuations occur in the span of many generations, populations will track the changing environmental variable. Conversely, when fluctuations happen within the lifespan of the individuals, plasticity or bet-hedging will be the adaptive responses.

1.3 Movement strategies

Because movement requires time and effort, good decisions are paramount. What is a good decision, what is enough information and how quick or long-lasting a decision has to be is completely context-dependent. The principles behind these strategies are not well understood. Though efforts have been made from both modelling and statistical perspectives, the assumptions on the mechanisms underlying movement decisions, if any, have been overly simplistic. Many models that explore the interaction between behaviour, ecology and evolution focus on environmental complexity; these models tend to simplify the genetic and behavioural components to such an extent that they are quite unrealistic. For

example, most attempts to infer mechanism from ecological patterns (Foley et al. 2015), assume *a priori* that animal movement can be characterized by Brownian motion or Levy walks (Benhamou 2007; Smouse et al. 2010). Adding complexity at the genetic and behavioural layers, we allow for degrees of freedom that may challenge our assumptions about these systems (de Jager et al. 2014). Evolving Artificial Neural Networks provide just one such a way.

1.4 Network models for the evolution of plasticity

Condition-dependent movement is a special case of phenotypic plasticity (Price et al. 2003). Plasticity is typically described by so-called ‘norms of reaction’ (Woltereck 1909), and models for the evolution of plasticity typically assume that some parameters determining the shape of the reaction norms are the target of selection. This is problematic for two reasons: the range of possible reaction norms is constrained *a priori* by the modeller (thus not allowing the evolution of ‘surprising’ movement patterns), and the assumed straightforward relationship between reaction norms and underlying genetics precludes complex eco-evolutionary dynamics that often result from the intricate genotype-phenotype mappings of more realistic mechanistic models. For this reason, I will here adopt a ‘regulatory network’ approach, where movement decisions are taken by evolving neural networks. This approach comprises the traditional reaction norm approach, because the evolution of such reaction norms is formally equivalent to the evolution of neural networks with a very simple structure (no processing layer). Networks with a more complicated structure have more evolutionary degrees of freedom and can produce much more complex reaction norms.

1.5 Aims and approach

The main focus of this project was to design and construct the model itself. The resulting model allows for extensive research and experiments beyond the scope of the current manuscript. The main goal of the model was to understand the evolution of movement strategies in resource landscapes that vary in their predictability in space and time. These strategies were to be implemented mechanistically by Artificial Neural Networks (ANNs) that take inputs from the current environment and output a magnitude and direction of movement. We contemplated several architectures and node functions for the ANNs, to get a glimpse at the influence of the architecture itself on the evolved movement strategies and the possible coexistence of different movement strategies (phenotypes) or node-weight values (genotypes). In this way, ANNs provide a mechanistic and multi-dimensional genotype-to-

phenotype mapping that allows for more evolutionary degrees of freedom than traditional one-locus or even multi-locus approaches to movement.

The simulations consist of an Individual-Based-Model of a non-interacting population of individuals distributed in a continuous one-resource landscape. For simplicity, the individuals on the population do not interact and cannot sense each other. Individuals sense resources in their local environment (a point), and in the more complex model, around their local environment (four more points). Based on the sensory input, a step size is calculated through the genetically encoded weights of the network. In the simplest model, the angle of the movement is calculated at random, whereas in the complex model it is a second output of the network. We implemented non-overlapping, discrete generations. During their lifetimes, the individuals gather as much resources as available in their position. After a set number of time steps, asexual reproduction takes place as a random lottery where the probability of reproducing is proportional to the relative resources gathered. At reproduction, individual connection weights have a probability of random mutation. I was interested in how evolution shapes movement strategies, particularly when a different set of strategies is possible.

1.6 Structure of the report

Given that information about the future state of the environment would be beneficial, organisms evolving in scenarios where future states are more correlated with current environments would tend to evolve consistent movement reaction norms. In the opposite case, where future states are inherently unpredictable, the reaction norms would tend to be random in nature. Also, if the predictability of the future states of the environment decreases as the distance from the present location increases, organisms would tend to move when in a low quality spot, and stay still when in a high quality spot. That is, they would evolve a kind of giving up density.

In this work, I ran simulations of individuals in a fluctuating, spatially complex landscape of one resource. The individuals consumed resources without depleting them from the environment, and their movement was determined by an artificial neural network with a predetermined architecture and evolving weights. I implemented several scenarios that differed in the network architecture, and the magnitude and frequency of environmental change. Finally, I analyzed the evolutionary outcomes after 100,000 simulated generations.

2 Methods

2.1 Model purpose

We implement a mechanistic, individual-based model for the evolution of movement behaviour in a spatially explicit context. Using artificial neural networks (ANNs), we aim to construct models of intermediate complexity to test how the evolution of movement strategies depends on the regime of environmental spatio-temporal predictability and scale. The model consists of a landscape with a heterogeneous and continuous resource distribution, and a population of individuals that move on the landscape based on genetically encoded ANNs and environmental cues.

2.2 The landscape and its dynamics

The environment was simulated as Perlin noise (Perlin 1985). This allowed us to create computationally efficient, reproducible, continuous landscapes that can be repurposed for different simulations. The landscape consisted of a distribution of one resource, where the resource had values between 0 and 1. To introduce a layer of complexity, we decided to allow for discontinuities by ‘clamping’ the landscape. By this I mean that all resource concentrations below 0.4 and above 0.7 were set to zero (Figure 1-B). This way, climbing the resource peaks became less straightforward. Furthermore, the distribution of resources can change over time. In order to simulate a kind of predictability, I implemented magnitude and frequency of change as the two variables that determine how the environment changes through time. The predictability of the next cue depends then both on the distance to that point, the magnitude of the change (colours in Figure 3-B) and the frequency of the change. In Figure 3-B is represented that predictability as spatio-temporal correlations in random directions when the landscape changes in the next time step. The effect is then buffered when the frequency is lower.




Figure 1: Example of a non-clamped (A) and a non-clamped resource environment.

2.3 Resource landscape

Perlin noise is a common way to produce semi-random landscapes in visual applications like video-games and computer-generated visual content (Perlin 1985). It works by assigning a grid of semi-random vectors (Figure 2-1: $\vec{P}_A, \vec{P}_B, \vec{P}_C, \vec{P}_D$) and calculating weighted values within each grid cell, proportional to the distance (Figure 2-2: $\vec{A}, \vec{B}, \vec{C}, \vec{D}$) to the vertices which the vector was assigned to. Although in the current work we used a three-dimensional landscape where one dimension was time, for simplicity here I will explain how it works in two dimensions. In principle, this process is expandable to n dimensions.




Figure 2: Calculation of a simplified two-dimensional Perlin noise landscape. 1) The vectors are randomly distributed over a 2x2 grid. 2) At any given point (black dot), the distance to the grid corners is calculated and after interpolation, 3) a single value results from the dot products of the vectors.

We then calculate the dot products $\delta_{\vec{A}}, \delta_{\vec{B}}, \delta_{\vec{C}}, \delta_{\vec{D}}$ between the random vectors at each grid point (Figure 2-2: A, B, C, D) as follows:

$$\delta_{\vec{A}} = \vec{A} \circ \vec{P}_A; \delta_{\vec{B}} = \vec{A} \circ \vec{P}_B; \delta_{\vec{C}} = \vec{A} \circ \vec{P}_C; \delta_{\vec{D}} = \vec{A} \circ \vec{P}_D$$

Then a linear interpolation is applied to each row by weighing the dot products that correspond to the top points in the grid (α) and bottom points of the grid (β):

$$\alpha = \delta_{\vec{A}} \times F(1 - \vec{A}_x) + \delta_{\vec{B}} \times F(\vec{A}_x)$$

$$\beta = \delta_{\vec{C}} \times F(1 - \vec{C}_x) + \delta_{\vec{D}} \times F(\vec{C}_x)$$

where F is a smoothing function for a weight w , that in this case is the standardized distance to a particular point of the grid:

$$F(w) = 6w^5 - 15w^4 + 10w^3$$

We then finally interpolate between α and β :

$$V = \alpha \times F(1 - \vec{A}_y) + F(\vec{A}_y) \times \beta$$

The result is a periodic landscape that is computationally efficient, as at any time point we only need to calculate the resource value of the points where the individuals are, as opposed to simulating the whole landscape. This meant we had a wrapped landscape in the shape of a torus (doughnut), which makes for more intuitive visualization and it allowed us to get rid of edge effects.

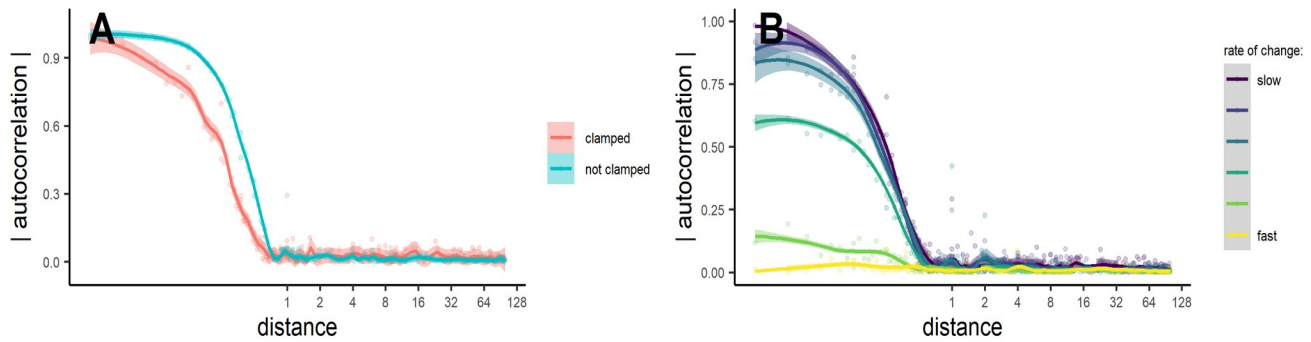


Figure 3. Autocorrelogram of the environmental resources in a static landscape (A) and a dynamic landscape (B). A) Correlation values between the resources of any two places at the point in time. The gray line represents a completely continuous landscape. In red is the clamped landscape we used for the longer simulations, in which there are larger areas without any resource. B) Correlation values between the resources of any one place at the time of sensing (t) and any other place at the time of moving ($t + 1$). These values can be cautiously understood as the environmental predictability, which decreases with spatial distance, discontinuities (i.e. clamping) as well as with faster rates of environmental change.

2.4 Neural networks

Artificial Neural Networks (ANNs) were originally inspired by the functioning of biological neurons. They are a step towards a more mechanistic approach to modelling behaviour. Many simplifying assumptions, such as assuming non-mechanistic node functions (i.e. arc tangent function), no development, are unavoidable at the present time. However, it is necessary to partition behaviour if we are to understand the function of neural architecture and behaviour complexity in tackling behaviour problems.

The ANNs used in the present work are simple forward-feeding networks (Figure 4-B) with nodes organized by layers (Enquist and Ghirlanda 2005). Every node of every layer takes some input from either the environment (Figure 4-A) or the previous layer, and computes an output. The output from the very last layer is translated then into the behaviour. There may be any number of intermediate layers in between the sensory and behaviour layers (Figure 4-B). In the same way, the sensory and behaviour layers can be a single one.

In this work we fix the network architecture within a population, while allowing the weights and biases to evolve. This allowed us to write efficient code while still allowing for a diversity of strategies to evolve. In Figure 4 we detail the different layers of the most complex architecture. By implementing

simpler architectures (Figure 6), we effectively compare a proper network with simple linear mechanisms of movement.

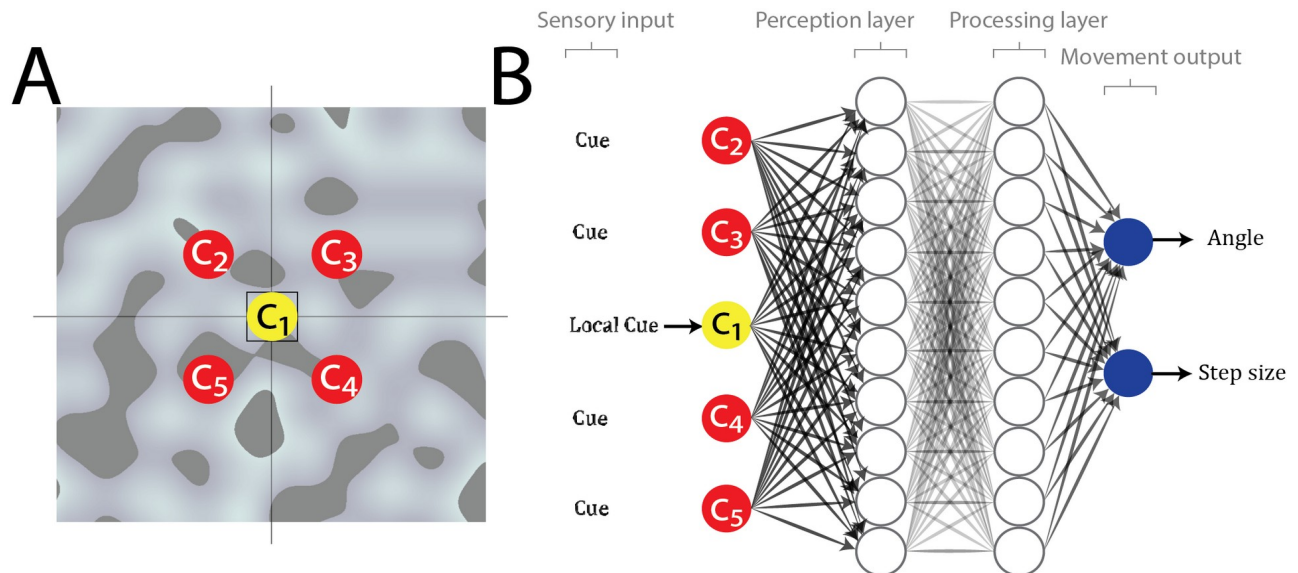


Figure 4. Example of a network with several environmental inputs. A) The individual takes the resource cue from the landscape, at the point where it stands, as well as from four surrounding points. B) A three-layered network with five environmental inputs, which are the resource concentration values at the location of the individual and in four points in the vicinity. The perception layer has ten nodes, as does the processing layer. Finally, the network output is two values: the magnitude and the angle of the movement that the individual is to take on the next time step.

The connection between two nodes of adjacent layers has two attributes, a weight w_i and a bias w_b , through which each input z_i is computed into a y node activity:

$$y = \left(\sum_i w_i * z_i \right) + w_b$$

The node activity is then passed through a node function $f(y)$, where f can be any function. We used one of three functions in this study: a rectified linear activation function (rtlu), a unipolar activation function or a hiperbolic tangent function (Figure 5).

In Figure 6, the six network architectures are shown. They differ in their sensory input, the number of layers, number of nodes and whether the output is exclusively the magnitude of the movement step or the angle of movement as well.

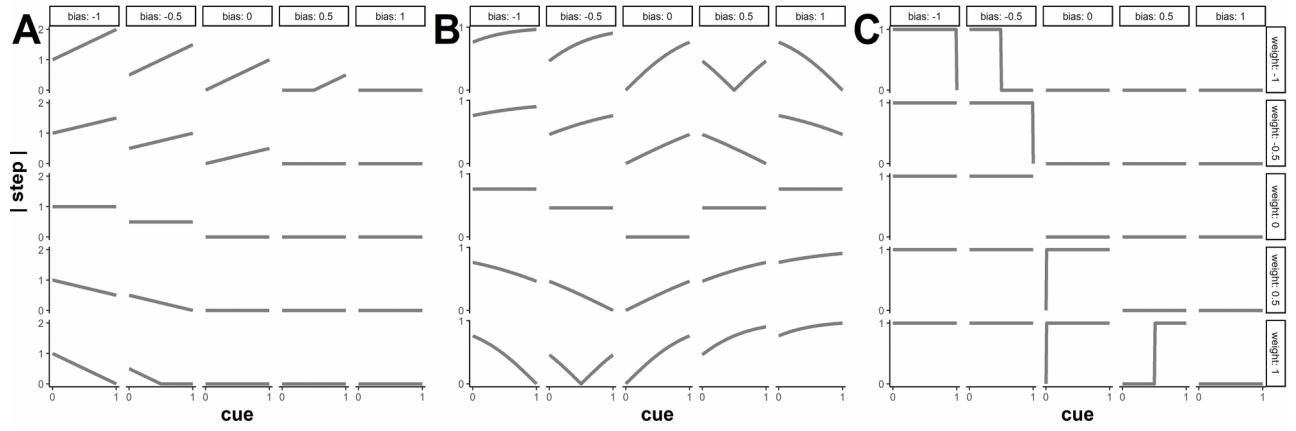


Figure 5. Node functions. A) Rectified linear function (rtlu). B) Hyperbolic tangent function. C) Unit step node function.




Figure 6. The six fixed ANN architectures used throughout the study. A,B,D are networks where there was only one environmental input, the local resource concentration, and only one output: the magnitude of the steps taken. In these, the angle of movement was drawn from a random distribution. D is a network that resembles A,B,D, with the exception that the network determines the angle of movement as well. E,F are networks with five environmental inputs and both the angle and magnitude of movement are determined by the network.

2.5 Detailed model description

2.5.1 State variables and scales

Environment

The environment consists of an infinite virtual landscape, in which a resource is heterogeneously and continuously distributed. The resource distribution consists a randomly generated layer that varies in the degree to which the resource changes through space (spatial autocorrelation) and time (temporal autocorrelation). Heterogeneity in both space and time will be simulated as a Perlin noise cube with different values of autocorrelation.

The landscape is predictable in both space and time. Consider a point (x, y, t) on this landscape, where the value of resources is known. The distance n in space over which the resource values of a point $(x + n, y + n, t)$ can be estimated with some threshold accuracy is the spatial predictability of the landscape. This is implemented in the model as spatial autocorrelation range, and landscapes with high autocorrelation range (ACR) are more predictable than those with low ACR.

The time z over which the resource value of the point $(x, y, t + z)$ can be estimated from the point (x, y, t) is the temporal predictability of the landscape. This is implemented as the timescale of change (R) of resources, and landscapes that change slowly are more predictable than those in which fluctuations are rapid.

Population and individuals

The population consists of non-overlapping generations of individuals that sense, move, and forage on the landscape. The individuals are structured as follows:

Sensory mechanism: Each individual has a neural network with one input that samples the resource value at its current location of the individual. This input is processed through three hidden layers (five nodes each) to generate two outputs: an angle and a distance of movement. The weights of the network are genetically encoded and can mutate during reproduction.

Individual state: Each individual has three internal state variables: the spatial position given by coordinates (x, y) , and stored resources.

Table 1. Model parameters and their values.

Parameter	Default value
General	
Population size	10,000
Number of generations	100,000
Movement decisions per individual	100
Movement	
Movement cost	0
Resource distribution and change	
Distribution function	Perlin noise
Spatial autocorrelation	Variable
Magnitude and frequency of change	Variable
Clamp	All resource concentrations below 0.4 and above 0.7 are set equal to zero
Reproduction	
Dispersal radius	0
Mutation rate per locus	0.001
Cauchy parameter of mutational step size distribution	0.1
Neural Network	
ANN architecture (Number of nodes per layer)	Variable
Max number of nodes and weights	20 nodes in three layers
Node function	Variable
Network type	Feed-forward

2.5.2 Process overview and Scheduling

Within a single time step the following operations take place in sequence:

1. Individuals gather cues at (x_t, y_t) and move to (x_t', y_t') ;
2. The landscape is updated;
3. Individuals take up the resources at their location (x_t', y_t') .

This procedure is repeated for all time steps within a generation.

A generation (= agent lifetime) is comprised of 100 time-steps (Table 1). At the end of each generation, reproduction will take place in the following sequence:

1. Relative fitness calculation;
2. Offspring generation;

3. Offspring mutation;
4. Offspring spatial placing.


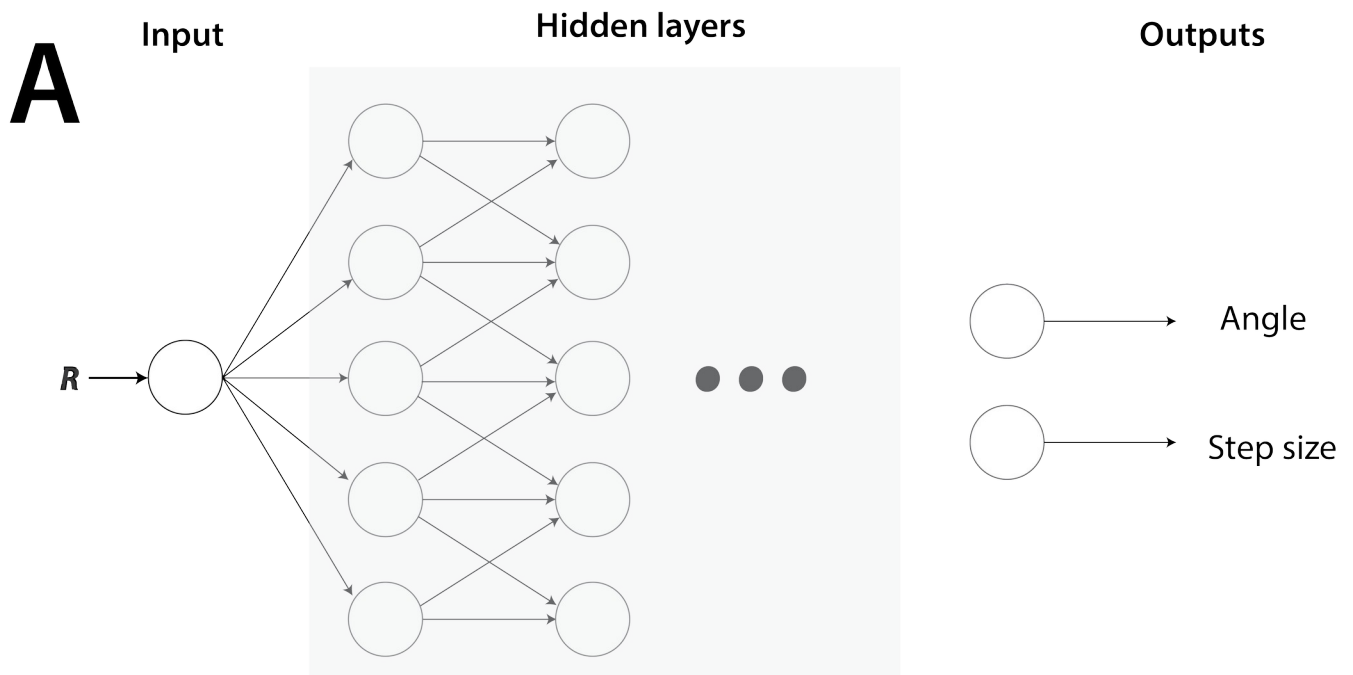


Figure 7. (A) Schematic representation of the ANN. (B) Process scheduling.

2.5.3 Design concepts

Fitness

Individual relative-fitness will be calculated as the stored resources (gathered resources minus costs) in a lifetime divided by the total resources stored by the population.

Adaptation

Populations adapt and improve their fitness by selection on mutation of the ANN weights. This leads to the evolution of functional behavioural strategies.

Sensing

Individuals receive either one or five information inputs, according to their architecture, during each sampling event, consisting of the resource value at the current location.

Interactions

Individuals are not able to sense each other and do not compete for resources. Due to the fixed population size, they do compete for relative fitness share and offspring.

Stochasticity

The stochasticity of the model is produced on the levels of the evolutionary and behavioural time scale. Within one generation, stochasticity arises from random location initialization within the sprouting radius, and spatio-temporal resource dynamics. On an evolutionary time scale across generations, stochasticity is produced by the interplay of mutation, ANN architecture and genetic drift.

Observation

The graphical representation of the model allows detailed observation of movement behaviour and spatiotemporal resource dynamics. Population fitness and neural complexity can be returned as text output. Finally, it is possible to export whole populations as binary objects in order to use a subset of the individuals in further simulations to analyze their movement patterns.

2.5.4 Initialization

Individuals of the population are initialized with a random location, a resource value of zero, an ANN with all weights set to zero.

2.5.5 Submodels

Spatio-temporal resource dynamics

The resource distribution is created through Perlin noise. There are a maximum and minimum values outside which the resources are set to zero. The noise is defined by a frequency, analogous to a measure of spatial autocorrelation. For the dynamically changing noise, a time variable z is added. The relation between z and time steps t determines the rate of environmental change relative to the behavioural time

scale of individuals. Optionally, a clamp can be added to the resource distribution, such that landscape values below a certain threshold are set to zero. This creates a landscape with clumped resources.

Movement

Input and processing

Individuals at time t sample their current location and enter the sample as input to their ANN.

This input is processed through the forward-feeding ANN. The nodes are biased and use a node function to transform node activity to output to consecutive nodes. Nodes in the upper layers use either a rectified linear activation function (rtlu), a hiperbolic tangent function (tanh), or a unit step function (unipolar), while the output node that produces step size has an identity node function, and the output node that produces the angle has a sine node function.

Output and function

The ANN returns a step size β and either a random or ANN-determined movement angle α , depending on the specific architecture. Individuals move to their next position at time $t+1$ as specified by the angle and step size using the formula:

$$\Delta x = \beta * \alpha ; \Delta y = \beta - \Delta x$$

After each sampling event, the ANN updates the individual's angle and step size as described above.

Foraging

Individuals gather the resource value of their current location.

Fitness calculation

The fitness of each individual (always ≥ 0) is based on the resource collected during one generation, and the energetic costs of movement and sampling are set to zero. The relative fitness of an individual is calculated as individual resource stored by total population resources stored. This value gives the chance of each offspring being assigned to this individual.

Reproduction

Reproduction occurs asexually and population sizes are fixed. The relative fitness values are used to create a discrete probability distribution assigning offspring to individuals. Reproduction is clonal regarding all state variables except for location, which is determined randomly for the offspring.

Mutation

The weights of the ANN of each individual are subject to mutation before the new generation starts. Mutation changes a single weight to a new value distributed in a Cauchy distribution of a certain shape (Mutation shape, table 1) around the former value. Mutations occur at a certain rate for each weight and generation (Mutation rate, see table 1). The same mutation rate and shape is also applied to mutate the sampling frequency of individuals.

2.6 Model output

For every simulation, we store the data of only a fraction of the generations due to storage considerations. The data is stored in four ways for data analysis. 1) We save the weight data of all the individuals of a fraction of the generations as binary files in order to do ecological simulation experiments with the evolved populations. The populations can then be ‘revived’ to test in different environments with a different script (runPrintReactionNorms).

The 2) weight data and the 3) cue-response data are stored for a fraction of the generations as text files to be plotted. This is done through a separate script (runPrint) that takes the binary data as input and outputs the weight data as text, as well as the cue-response data can be called through a batch file that is outputted from the evolutionary simulation. Finally, the reaction norms of the fittest individuals are saved as text files by running a script (runRNPrint).

The weight and cue data can then be processed and plotted in an R script (codeProcessCueData).

We implemented a visual representation of the model, where the resource values and the individuals are placed. This allows for interactive, real-time manipulation of the environment while observing the individuals react to it. It acts like an ecological experiment that we used to visualize the evolved movement strategies and how well they coped under environments different to the one they evolved in. There is no reproduction happening, and different aspects of their behaviour are highlighted by color (speed of movement, accumulated resources). Further, individuals can be selected and followed, and their weights are shown on the screen, as well as their inputs and outputs.

2.7 Model implementation

The model is written in C++ and makes use of CUDA cores of modern GPU optimization in OpenGL. Parallel processing is also implemented via multi-threading. The implementation of both the perlin noise landscape and the neural networks allows for such parallel processing.

2.8 Simulations

A typical simulation consists of 10,000 individuals randomly distributed in a Perlin resource landscape. Each individual has an ANN architecture with weight values fixed at birth, and accumulated resources that reset at birth. Every time step, an individual eats all of the available resources from the environment, without depleting them, and then computes the input(s) and moves according to the output(s). The next generation is sampled with replacement from the parent population, and the ANN architecture and weights are inherited with a chance of small mutations. After 100,000 generations, the simulation ends.

Each evolutionary simulation is set under a set of parameters. For the present work, we kept the spatial patterns stable (resources stay between 0.4 and 0.7, with values below and above turning into a value of 0), lacunarity and octaves stay the same. Because we expected spatio-temporal change to be the main source of selective pressure, we varied the frequency and magnitude of change of resource concentrations. As explained above, this was implemented by simulating a third dimension on the perlin landscape, and the jump in that third dimension is the time. It's important to note that the structure of the Perlin landscape also affects the very structure of spatio-temporal change.

For the evolutionary simulations, each simulation had a static rate of change of the environment. From very stable environments to completely unpredictable. We also explored six different networks. Four of them were 'blind', where individuals could only sense their local resource intake. Two of them had four sampling points around their position, which could inform them of the resource gradient at every time step. We also tested whether only choosing a moving distance, or choosing both the distance and the direction. After the evolution rounds were done, the populations were subjected to all of the combinations of rates of environmental change for one generation. Due to time constraints, here we only present the most important results.

3 Results

The movement behaviour results from an evaluation of the cue or cues at and around the location of the individual in the resource landscape. At the end of each generation, the individuals who collected the most resources are the most likely to reproduce. Movement towards higher resource locations or away from resource deserts allows individuals to collect more resources. The movement itself is the result of sequential operations through the nodes and layers of the artificial neural networks. The population

size, network architecture and the structure of the inputs and outputs were fixed within each population. The resource cue is processed through weights and biases which are subject to probabilistic mutations at the stage of clonal reproduction. I explored a set of simulations where populations varied in architectural complexity, number of input cues sensed and whether the angle was also determined by the network (Figure 10).

Because of the complexity of the behaviours, here I will focus on two main outcomes: the reaction norms in response to the main cue (Figure 8-A) and the fitness of the individuals at the end of the simulation (Figure 8-C).

3.1 A typical evolved reaction norm

The magnitude of the step as an output of the networks in response to a main cue is a reaction norm (Figure 8-A). As these artificial neural networks are deterministic, these lead to a straightforward line in simple cases. In more complex cases (Figure 10-E, F), the magnitude of the step depends on 5 cues. In order to simplify its representation, here I project reaction norms against the main cue, that is, the actual resource that the individuals accumulate. Some deterministic variation in output step around the main values results from variation in hidden axes representing the other cues. Each reaction norm has a fitness value associated with it. This is the result of replicate runs of the individual in the conditions in which it evolved, that is, the frequency and magnitude of the temporal change in the landscape. The history of cues throughout these replicates can be seen in Figure 8-B, while the fitness of sessile individuals can be seen on Figure 9.




Figure 8. Fit individuals tend to move when their current environment is resource-poor. Conversely, they tend to stay when the environment is resource-rich. However, due to the spatial structure, that is, to the fact that resource rich areas are often close in space to areas with zero resources, individuals




Figure 9. Distribution of resources accumulated over the lifetime of non-moving individuals under different environmental regimes. Each shape represents the distribution of a particular replica. The dotted line on the top represents the theoretical maximum, and the black line stands for the mean over all replicas.

3.2 Reaction norms evolved in all architectures

The reaction norms of many individuals in the evolved populations were shaped as expected in most scenarios. Individuals would move when in low-resource locations, and then take smaller steps when in high-resource locations. However, networks that could choose the direction of their steps would also move while standing near the resource peaks, giving a V shaped curve (Figure 10). In most scenarios, reaction norms that were beneficial emerged in the first ten generations, though with great variability in their fitness due to stochasticity (Figure 11). Then, through evolution, the reaction norms were refined further and their variability in fitness decreased. That is, the strategies would be more reliable in a variety of situations within the regime of environmental change.

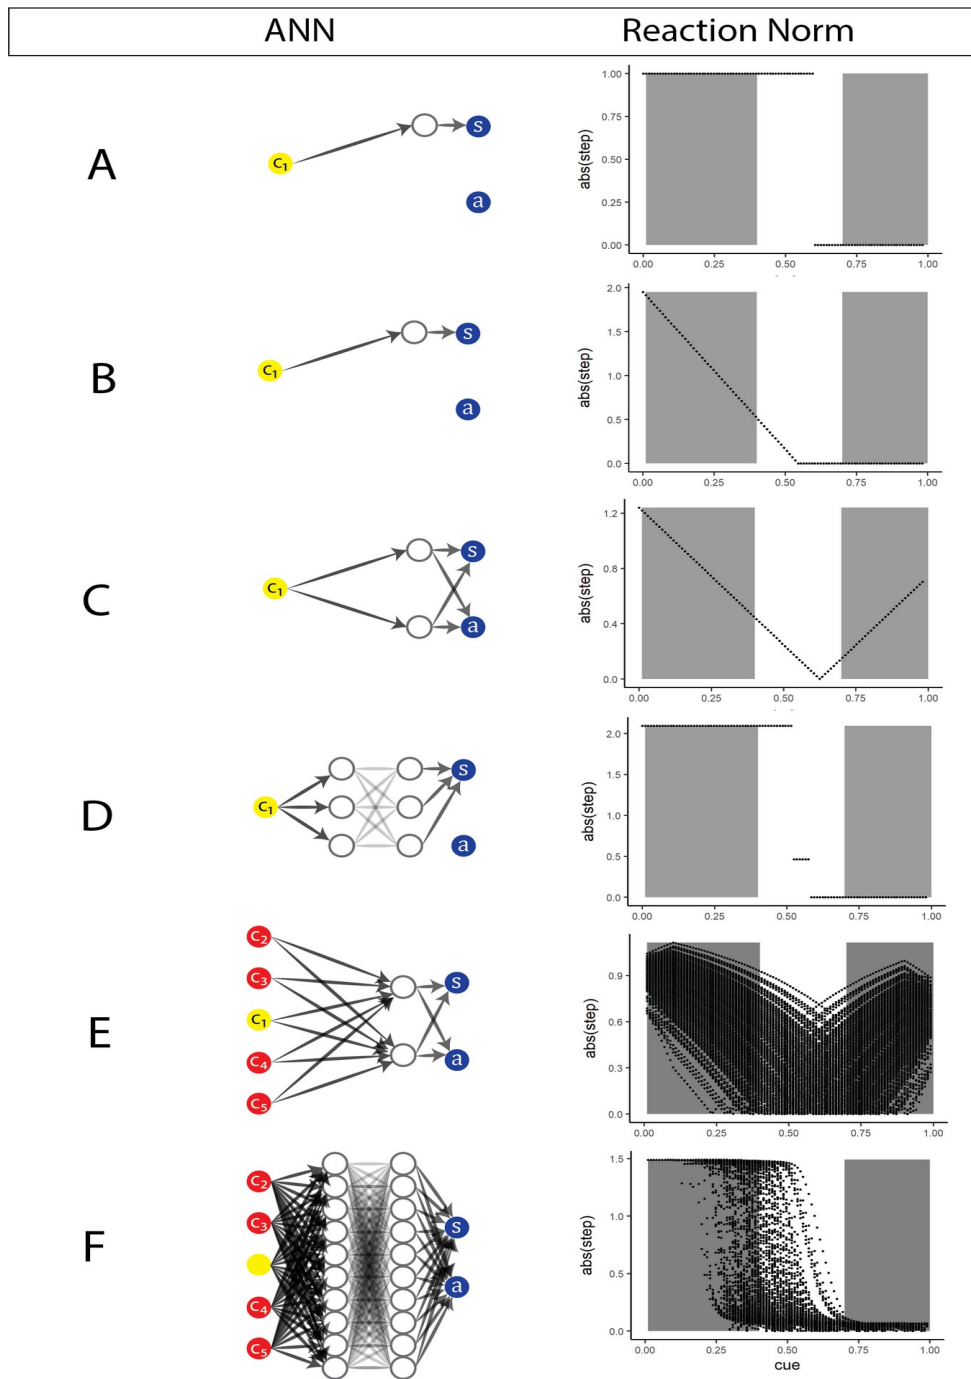


Figure 10. A representation of the ANN architectures and their typical evolved reaction norm. ANN) The network takes in either just a cue (yellow) of the resources at the spot where the individual is, or also cues around the individual (red). A) A network with nodes takes in the cue or cues and outputs a step value. The angle is either an output (C, E, F) or drawn from a random distribution (A, B, D). Reaction norm) The absolute value of the step output by the network in response to the main cue. Since there are more cues in the sighted networks (E, F), only the main cue is represented, and the variability shown in response to the other hidden axes. The gray rectangles are cues that don't exist in the current landscape. Only values of 0 and 0.3 - 0.7 occur.

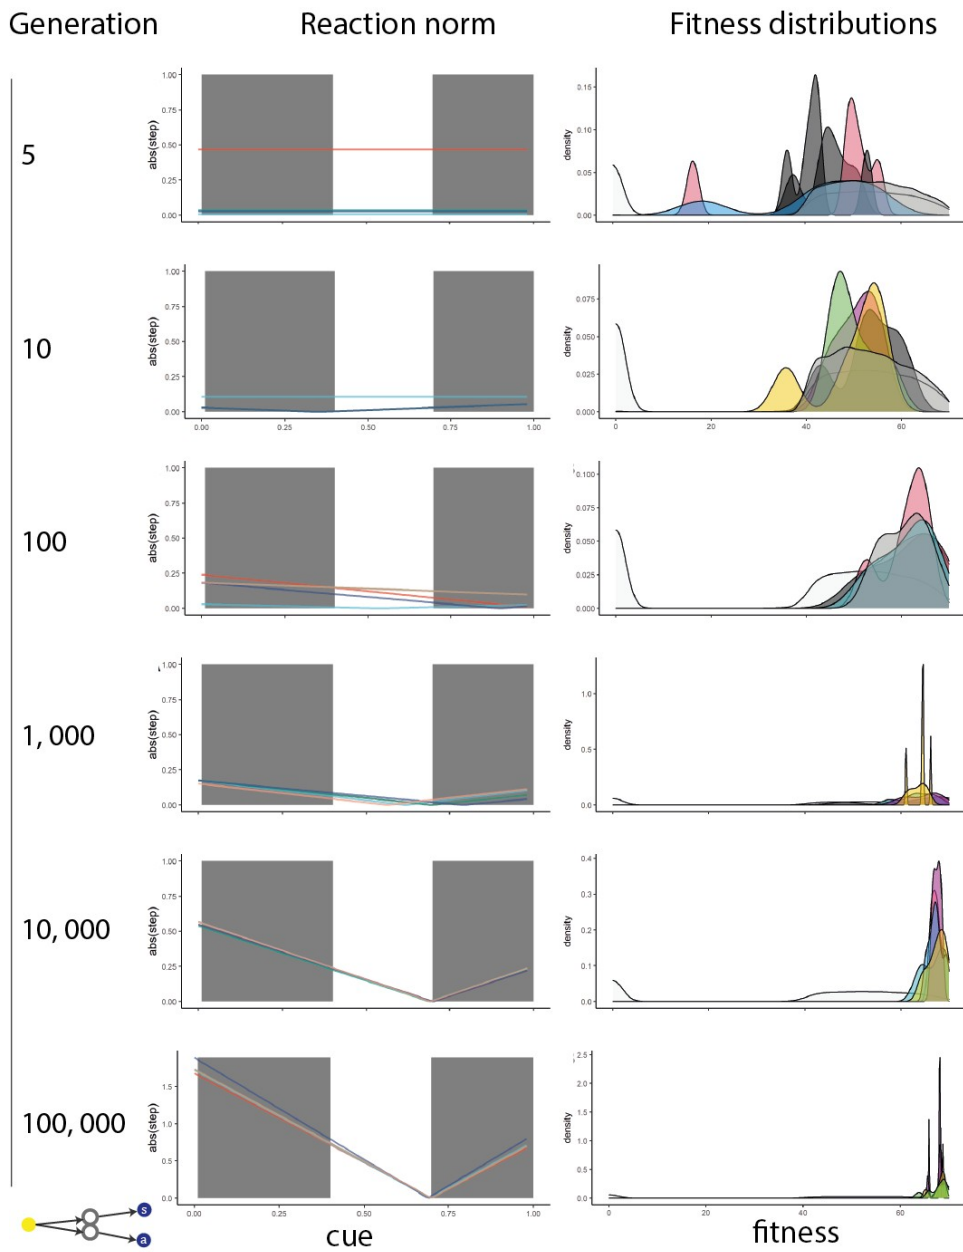


Figure 11. Evolved reaction norms through the generations and their fitness distributions. In a population with an ANN architecture as in Figure 10-C, the first mutations result in flat reaction norms with a wide spread of fitness values. Further selection refines the strategies to the V shape seen from generation 1000 forwards. Each color represents a particular individual in that generation. In white, the expected fitness values for sessile organisms.

3.3 Movement strategies coexist with sessile strategies

In all evolved scenarios, sessile individuals coexisted with motile individuals. Coexistence was not only common, but often some sessile individuals performed as the best 10 % of the population. While on average the motile individuals did much better, the effect of birthplace was strong enough to allow for persistent coexistence.

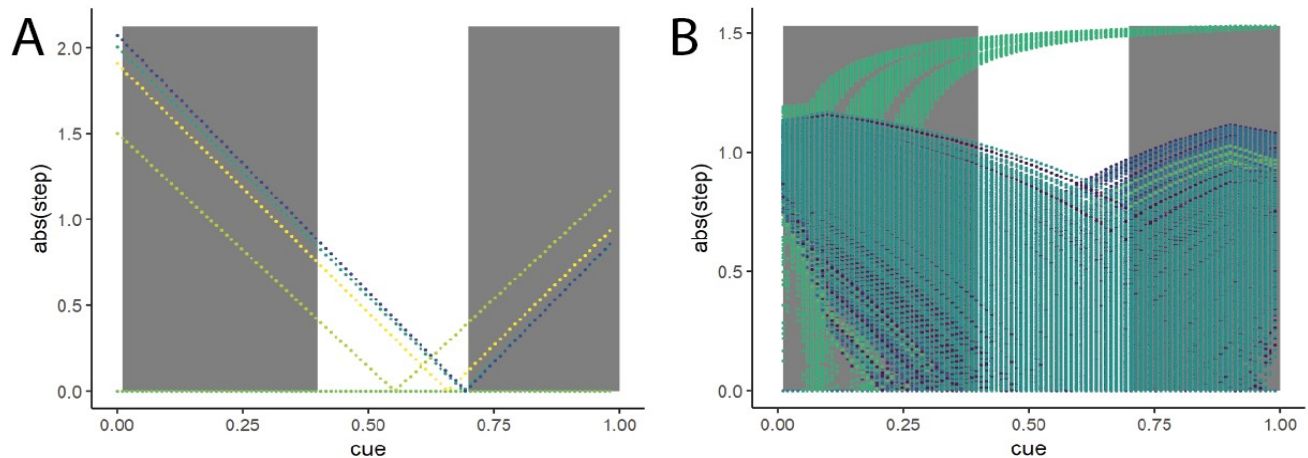


Figure 12. Coexistence of reaction norms. The absolute value of the step output by the network in response to the main cue. Since there are more cues in the sighted networks (B), only the main cue is represented, and the variability shown in response to the other hidden cues. The gray rectangles are cues that don't exist in the current landscape. Only values of 0.0 and 0.3 - 0.7 occur. The colors represent distinct individuals in the evolved populations. Half of the individuals in both scenarios are sessile.

3.4 Uninformed individuals tend to stay where they are when environmental change is very slow

In static resource environments, only populations with ANNs that determined the direction of movement evolved movement behaviour (Figure 13-1 and Figure 13-3), whereas individuals in random-movement-direction populations did not (Figure 13-2). This effect was regardless of the number of nodes or layers (Figure 14).

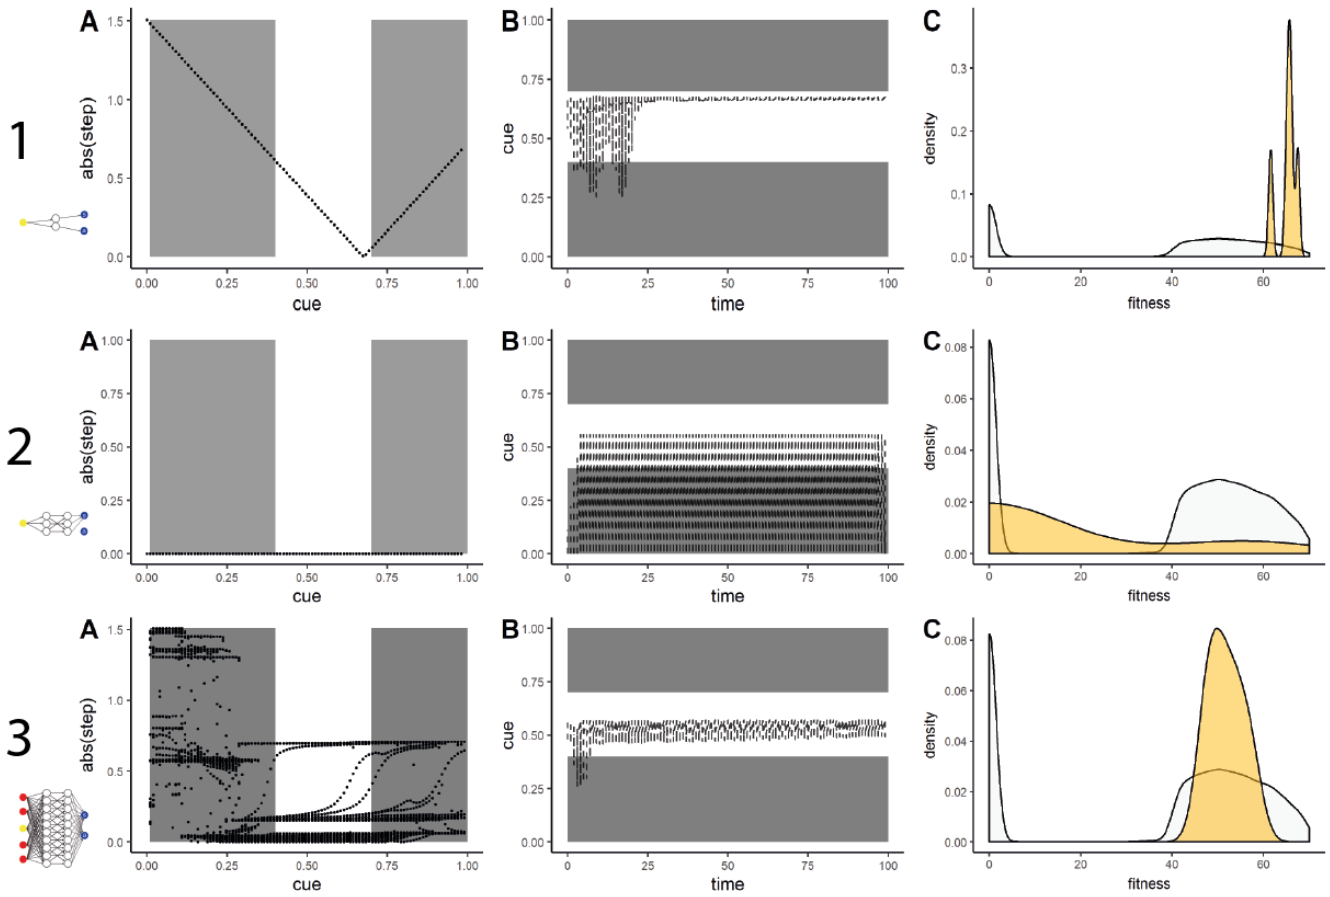


Figure 13. A) Reaction norm, B) cues sensed and C) fitness outcome of an individual from an evolved population with 1) Dual ANN, 2) Smart ANN, 3) Smart Sight ANN.

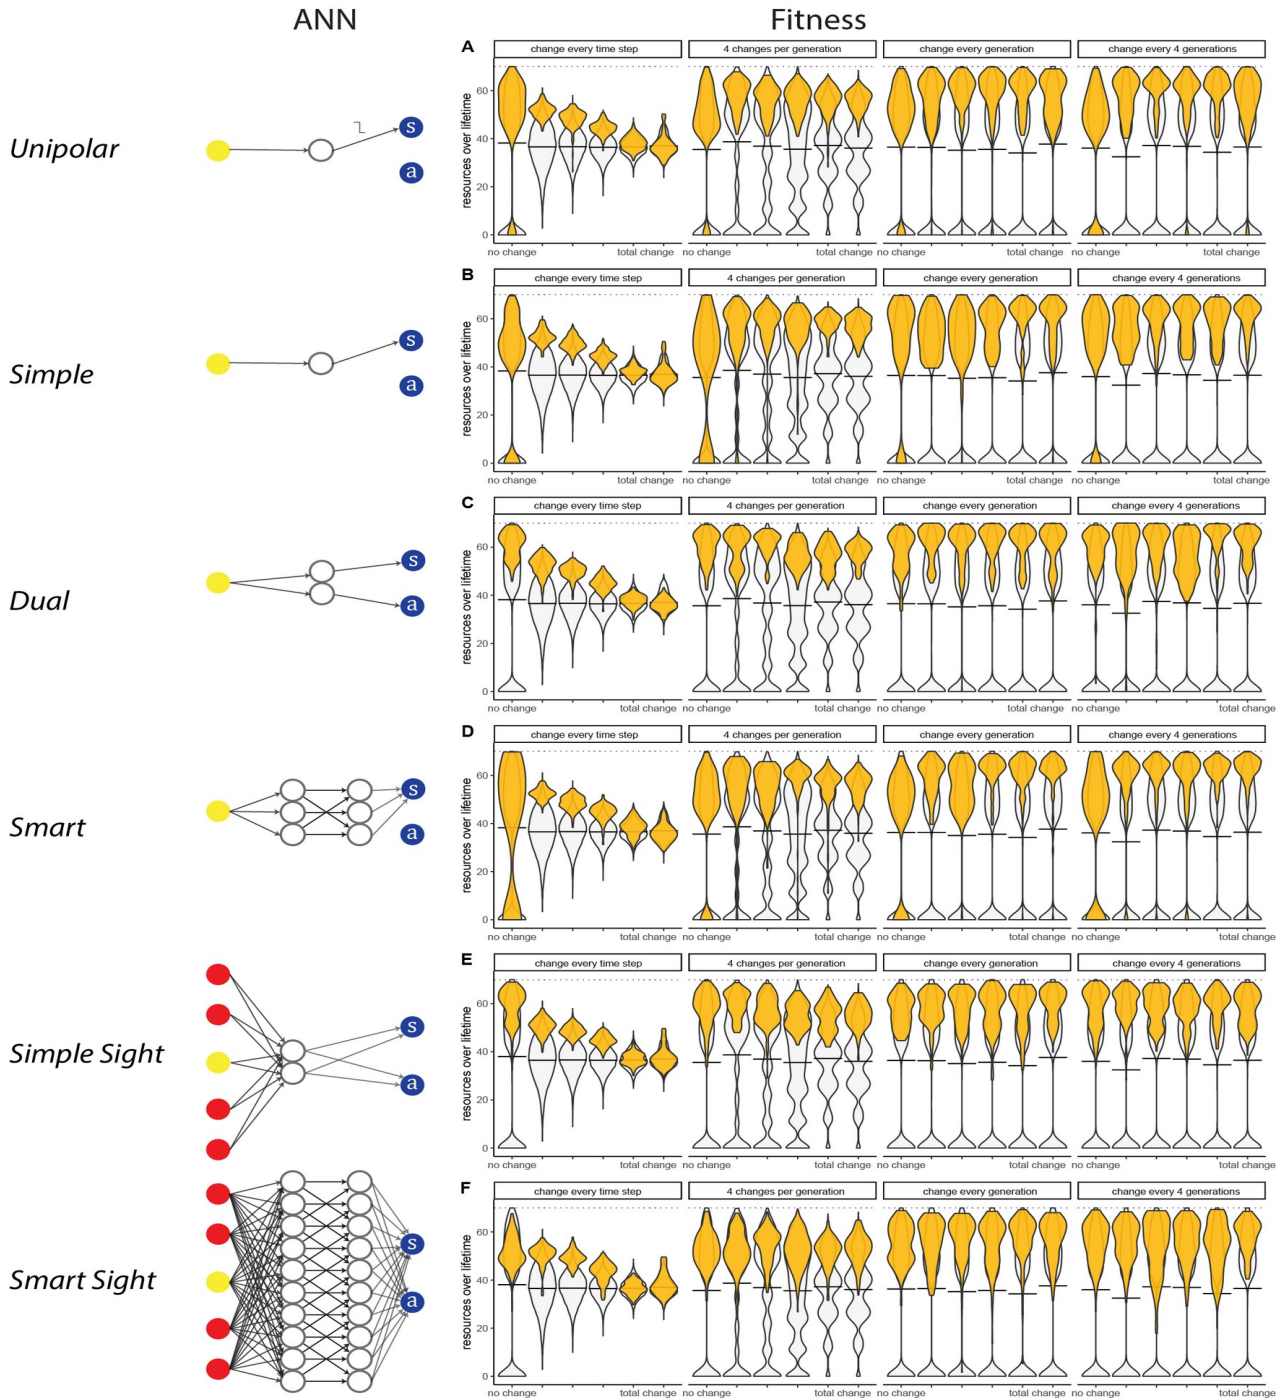


Figure 14: Fitness of the top 10% of the evolved populations. In gray, the distribution of fitness values of non-evolved sessile, control individuals. In yellow, the fitness of the top 10% of individuals of the evolved population. The fitness was calculated for each individual in replicate scenarios in order to cover the span of the effect of the starting point on the fitness.

4 Discussion

4.1 The importance of ecological inheritance

When the environment changes very slowly, there comes a point where sessile and motile individuals perform just as well, after all individuals have established themselves in resource peaks. Through a sort of ecological inheritance, not by modifying the landscape, but by moving to a more suitable environment, individuals effectively alter the selective pressures for movement, as argued by Ghalambor and colleagues (2007). At this point, in fact, mutations that promote movement in resource peaks would be selected against. Evolved reaction norms may then disappear from the population through random drift. This would lead to an eco-evolutionary cycle where movement evolves and drifts by patches in the population. Local stochasticity was acknowledged early on (Chesson 1981; Hubbell 2001) and then pushed to the side, as it presents complications for analytical models (Mullon and Lehmann 2018).

4.2 Coexistence of several strategies

Landscape heterogeneity has been identified before as a promoter of intra-specific trait variation. In related modeling work, landscape patchiness has been shown to promote intra-specific variation, and in turn this variation facilitates species coexistence (Spiegel et al. 2017; Milles et al. 2020). In other frameworks, spatio-temporal heterogeneity, as seen here, has also been shown to promote diversity at several scales (Chesson 1997; Chesson et al. 2004).

Here, I show that indeed coexistence of a continuum of strategies prevailed. Not only because of the more complex genotype-phenotype map, but also due to the inherent spatial stochasticity.

4.3 Neural networks and plasticity

Complex ANNs provided only marginal fitness advantages over simpler ones. Even some architectures approached the functions evolved in the simpler ones. This suggests that the spatial problem was trivial in some conditions, and that without increasing the amount of information perceived by the individuals, or the refinement of the behaviours, could potentially be costly. However, given that decision making systems, as well as the genes involved, are repurposed for multiple behaviours in nature, complex networks may allow for an increased evolvability.

4.4 Response to fluctuating environments

The model is consistent with the notion that the evolutionary responses were predictable, as different reaction norms evolved in the expected manner through different frequencies and magnitudes of environmental change. This is in line with the findings by Botero et al. (2015), though the coexistence of sessile and motile individuals throughout the simulations highlights the relevance of ecological heterogeneity for eco-evolutionary analysis.

Furthermore, it may be that the complexity of the networks has a secondary effect on the dynamics, as the trait becomes more vulnerable to mutations. More complex networks, then, may well evolve faster as a result of the mere amount of loci involved. This was not addressed by Botero et al. and may be an interesting avenue for future research.

4.5 Movement and decision making

An interesting finding was that due to the clamping, which introduced discontinuities in the resource distribution, there were two main strategies, depending on the specifics of the clamping and the architecture of the networks. Whether the individuals could decide on the direction of movement determined how close they would get to the resource peak. Seemingly “smarter” individuals would actually sacrifice the advantage of being on the resource peaks in order to diminish the risk of ending in a resource valley.

4.6 Future developments, pros and cons of theoretical models and assumptions

The purpose of this project was to build a framework to test simple hypotheses about the evolution of movement. Due to time constraints, the focus remained on the analysis of the evolutionary outcomes and not so much in the process itself. A big part of the results is still under analysis, such as the spatially explicit part of the movement strategies, which are available as videos. A thorough examination of the actual eco-evolutionary dynamics is possible with the current framework and may lead to interesting results regarding the interaction of architectural complexity and rapid evolution. Another section would test the hypotheses by Botero et al. () regarding the tipping points in evolutionary strategies and the potential overfitting of strategies that would make populations or subpopulations extinct when environmental conditions change. For example, by testing whether populations already adapted to a regimen of change can cope with a different regimen. There can be several extensions of the current analysis and model, such as introducing more interactions between

organisms beyond competition for reproduction, such as implemented here. Another interesting development would be to allow for the architecture itself to evolve, like it was done in similar studies.

Bet-hedging strategies and developmental effects can be easily implemented. By allowing for an evolving developmental switch, an individual can, either by chance or with information from the environment, choose between two alternative movement strategy networks.

Implementing memory might be an interesting development of the model. This could be done by including either cues from the last local environment or a cumulative average of the past environments as inputs to the network. In this way, individuals might be better equipped to climb resource peaks, particularly when the landscape is large and the environmental variation is low with regards to the movement steps.

Network architecture can be made evolvable. By allowing new nodes disappear or to arise (or nodes to have a chance to be activated or knocked off) as duplicates of others in the same layer, we could test more directly the evolutionary paths of competing neural architectures under different conditions.

Different environmental change regimes could be implemented, such as seasonal change, random disturbance or resource waves.

Multi-resource, risk, and temperature landscapes can be easily implemented by adding more layers to the perlin landscape.

5 Acknowledgments

I would like to thank several people. First, to Franjo for having patiently supervised this work. Second, to Christoph and Pratik for the endless discussions and ideas that they contributed. Third, to Hanno for the amazing model he implemented so I could play with. To my whole research group, for their enthusiasm and contributions. And lastly, but mainly, to my friends, particularly Raphael for helping me see this to an end.

6 Bibliography

Botero, C. A., F. J. Weissing, J. Wright, and D. R. Rubenstein. 2015. Evolutionary tipping points in the capacity to adapt to environmental change. *Proceedings of the National Academy of Sciences* 112:184–189.

- Benhamou, S. 2007. How many animals really do the Lévy walk? *Ecology* 88:1962–1969.
- Botero, C. A., F. J. Weissing, J. Wright, and D. R. Rubenstein. 2015. Evolutionary tipping points in the capacity to adapt to environmental change. *Proceedings of the National Academy of Sciences* 112:184–189.
- Chesson, P. 1997. The roles of harsh and fluctuating conditions in the dynamics of ecological communities. *The American Naturalist* 150:519–553.
- Chesson, P., R. L. E. Gebauer, S. Schwinning, N. Huntly, K. Wiegand, M. S. K. Ernest, A. Sher, et al. 2004. Resource pulses, species interactions, and diversity maintenance in arid and semi-arid environments. *Oecologia* 141:236–53.
- Chesson, P. L. 1981. Models for spatially distributed populations: The effect of within-patch variability. *Theoretical Population Biology* 19:288–325.
- de Jager, M., F. Bartumeus, A. Kölzsch, F. J. Weissing, G. M. Hengeveld, B. a Nolet, P. M. J. Herman, et al. 2014. How superdiffusion gets arrested: ecological encounters explain shift from Lévy to Brownian movement. *Proceedings. Biological sciences / The Royal Society* 281:20132605.
- Enquist, M., and S. Ghirlanda. 2005. *Neural Networks & Animal Behavior*. Princeton University Press.
- Foley, A. M., R. W. Deyoung, D. G. Hewitt, M. W. Hellickson, K. L. Gee, D. B. Wester, M. A. Lockwood, et al. 2015. Purposeful wanderings: Mate search strategies of male white-tailed deer. *Journal of Mammalogy* 96:279–286.
- Hubbell, S. P. 2001. The unified neutral theory of biodiversity and biogeography. *Monographs in Population Biology* (Vol. 32). Princeton Univ Dept of Art &, Princeton, EUA.
- Milles, A., M. Dammhahn, and V. Grimm. 2020. Intraspecific trait variation in personality-related movement behavior promotes coexistence. *Oikos* 1–14.
- Mullon, C., and L. Lehmann. 2018. Eco-evolutionary dynamics in metacommunities: Ecological inheritance, helping within species, and harming between species. *American Naturalist* 192:664–686.
- Perlin, K. 1985. An Image Synthesizer. *SIGGRAPH* 19.
- Price, T. D., A. Qvarnström, and D. E. Irwin. 2003. The role of phenotypic plasticity in driving genetic evolution. *Proceedings of the Royal Society B: Biological Sciences* 270:1433–1440.
- Smouse, P. E., S. Focardi, P. R. Moorcroft, J. G. Kie, J. D. Forester, and J. M. Morales. 2010. Stochastic modelling of animal movement. *Philosophical Transactions of the Royal Society B: Biological Sciences* 365:2201–2211.
- Spiegel, O., S. T. Leu, C. M. Bull, and A. Sih. 2017. What's your move? Movement as a link between personality and spatial dynamics in animal populations. *Ecology Letters* 20:3–18.

Woltereck, R. 1909. Weitere experimentelle Untersuchungen über Artveränderung, speziell über das Wesen quantitativer Artunterschiede bei Daphniden. Verh. D. Tsch. Zool. Ges. 1909:110–172.