

Evolution of Developmental Strategies in NK Fitness Landscapes

Jacob Ashworth¹, Yujung Lee¹, Jackson Shen¹, Edward Kim¹, Zach Decker¹ and Jason Yoder¹

¹Rose-Hulman Institute of Technology, Terre Haute, IN 47803
ashworjs@rose-hulman.edu

Abstract

Evolution and development are related processes although their relationship is still not well understood. Attempts to explore their relationship are challenged by scales of time and space, but also by the limitations of studies focused on specific constraints of model organisms. To help gain insight into these phenomena, we create an abstract, general model of a developmental process that guides an agent's trajectory through a "tunably rugged" NK fitness landscape. The developmental process is represented by a genotype that is evolved and allows us to investigate periods of exploration and exploitation as they relate to periods of an agent's lifetime and a given landscape's difficulty. Results show that evolution selects for time-sensitive periods of exploration and exploitation, which vary with the difficulty of the landscapes being traversed. Furthermore, our analysis suggests that phenotypic diversity via random exploration present in both early and mid-life can aid the development of superior phenotypes.

Introduction

Learning, development, and evolution work together to produce diverse forms of adaptive organisms. Whereas the relationship between learning and evolution has frequently been explored, development is often overlooked or regarded as a subset of the learning process in many models (Soltoggio et al., 2018). Yet, development is distinct from learning in that it involves structural changes in both the body and brain over time according to a set of processes, though such processes themselves may allow plastic adaptations to environmental cues (West-Eberhard, 2005). Sensitive periods exemplify one such process, and necessitate that development is distinguished from learning with varying degrees of plasticity at different life periods (Knudsen, 2004). Given the limited efforts to explore the relationship between development and evolution, there has been little progress in understanding why certain developmental processes have emerged from evolution as well as how the specific forms observed in nature have arisen (Frankenhuis and Walasek, 2020).

Most models of development have been focused on specific organisms and the various constraints on them. There is an opportunity, however, to develop more abstract models that can explore developmental plasticity generally across

different developmental processes (Belew, 1990; Sommer, 2009; Smart, 2019). One example of such an abstraction is the NK model by Kauffman and Levin (1987), which has been successfully applied to a variety of modeling tasks (Altenberg, 1996; Geard et al., 2002; Fragata et al., 2019; Rhodes and Dowling, 2018). By employing this model as a "tunably rugged" fitness landscape, our work seeks to explore the evolution of developmental strategies in an abstract, but general manner. Our approach is to evolve a representation of different permutations of developmental processes throughout an organisms' full development. This simple developmental "program" is then used to guide an individual from a starting location within an NK model toward, hopefully, locations with higher fitness.

The model and analysis in this work extend previous work in three important directions. First, evolution is shown to produce developmental steps that explicitly explore the landscape in a random manner. Second, evolution drives development to be time-sensitive, with non-uniform patterns of exploration and exploitation comprising the developmental strategy. Third, evolution selects for a more complex developmental strategy marked by multiple transitions between predominantly exploitative and exploratory phases, as the landscape becomes more complex with a greater number of interdependent factors. Furthermore, this work presents a simple yet effective model that can simulate various interactions between evolution and development.

The primary aim of this work is to explore the evolution of development in an abstract model. The rest of this paper is organized to accomplish this goal. First, we summarize related work in in evo-devo models, abstract models in biology, developmental models, and adaptive walks. Next, we provide an overview of the model in this work, including the landscape, agents, actions, and evolution. We then present results from experiments in a gradual sequence of generalization for starting locations, different landscapes, and evolutionary runs. Subsequently, we discuss the major findings and explore them in the context of evolution and development. Lastly, future work identifies major directions to extend upon the results presented in this work.

Related Work

In this section we will discuss areas of relevant literature and their shortcomings. First, we provide a brief background on evo-devo models to offer insight into their particular challenges and limitations. Next, we review the use of abstract models, especially relating to understanding complex phenomena similar to our own domain of interest. Finally, we review developmental models and identify a lack of desirable features which we seek to address.

Evo-Devo Models

Evolutionary developmental biology (evo-devo) has stimulated biological research enormously, both empirically and theoretically (Müller, 2007). For the last two decades, Evo-devo models have been mostly based on evo-devo model organisms such as *Drosophila melanogaster* and *Caenorhabditis elegans* by building on the analysis of those organisms. Although there has been a fruitful expansion in this era of research, the past trend—increasing numbers of organism species such as *Tribolium castaneum* and *Nasonia vitripennis* (Roth and Hartenstein, 2008; Lynch et al., 2006) has led to severe challenges in the scientific methodology and technical difficulties (Sommer, 2009). These new models have driven researchers to build more sophisticated tool kits to investigate the mechanisms of evolutionary change in developmental processes. Developing these involve gene knockout or knockdown, and experimental manipulation which necessitates high complexity and scientific precision. Furthermore, these methods mostly depend on empirical optimizations, which are largely species specific such that protocols cannot be transferred from one organism to another (Sommer, 2009). To overcome these limitations and to capture interactions among complicated phenomena—evolution, development, learning—the models must necessarily become extremely simplified (Belew, 1990) such that an abstract computational model would be advantageous for modeling evo-devo.

Abstract Models

Abstract models have been successfully used to study high level processes, which are otherwise difficult to investigate due to their complexity and scale in time and space. Various abstract models have been developed to explore the relationship between learning and evolution.

One early abstract model developed by Hinton and Nowlan (1987) demonstrated how learning can guide evolution through an idealized simulation model. Similarly, Kauffman and Levin (1987) introduced another abstract model, the NK model, which is a tunably rugged landscape adjustable by two parameters N and K . As this model exhibits how interactions affect dynamics on rugged landscapes, it has been applied to various fields such as learning strategies (Campbell et al., 2020) and ontological development (Panchanathan and Frankenhuis, 2016; Walasek

et al., 2021). NK landscapes are commonly used for examining adaptive walks, which proceed to fitter neighbors resulting from strategies such as ascent, steepest ascent, or minimum ascent (Pitzer and Affenzeller, 2012; Wilke and Martinetz, 1999; Hebborn et al., 2008; Kauffman and Levin, 1987). Extending the NK model, Hebborn et al. (2008) and Wilke and Martinetz (1999) demonstrated significant behavior changes of adaptive walks, while Park et al. (2015) implemented greedy adaptive walks to study haploid asexual population.

Recently, Todd et al. (2020) developed an idealized model of lifetime and evolutionary learning and examined the effect of task-difficulty on the optimal trade-off between learning and evolution. In their model, there are two types of lifetime learning—stochastic hill-climbing and steepest hill-climbing, but the model focuses on learning rather than development and fails to include mechanisms for random exploration, which we believe to be an important component of lifetime adaptation. Overall, NK fitness landscapes have been successfully used to study high level processes in multiple fields and as a result appear viable as a candidate model for studying evo-devo as well.

Developmental Models

There is a need for abstract developmental models which are specific to the field of evo-devo that could address scientific questions in a computational manner. By applying abstract models to developmental, there is an opportunity to gain a deeper understanding of the evolution of development.

Some developmental models have examined the evolutionary selection pressures that produce sensitive periods (Panchanathan and Frankenhuis, 2016; Walasek et al., 2021; Frankenhuis and Walasek, 2020). Unlike a two-stage life history, in which organisms first obtain environmental cues and later develop phenotypes, Walasek et al. (2021) proposed a model where individuals incrementally respond to local environmental conditions with sensitive periods emerging during life. In this model, if an environmental change or migration occurs during its lifetime, then an organism must infer the environmental state and consider environmental change—becoming a complicated inferential task. Due to the complicated inferential nature, the model is highly specific which is a limitation in a developmental context.

Given that most previous models are complex, lack significant mechanisms, and lack tasks with epistatic interactions, we have an opportunity to formulate a novel way of modeling developmental processes. Using this approach, we hope to identify sufficient conditions for the evolution. Furthermore, we aim to explore the relationship between difficulty and different developmental strategies. Ultimately, this paper's goal is to enable work utilizing a computational model to challenge and reformulate existing ideas and produce new hypotheses.

Methods

In this section, we explain the methods used to simulate our model, conduct experiments, and analyze results. First, we will discuss the NK fitness landscape and how it is used within this context. Next, we explain the constraints on the agents' development within the landscape. Afterwards, we describe the process in which the development strategies of the agents are evolved. We then explain the baseline models used for comparison to the evolved strategies. Finally, we describe the universal parameters given to the model across our experiments.

NK Model

In order to model the development of organisms within various environments, we used the NK model as described in (Kauffman and Levin, 1987). Within a NK landscape, the n -value presents the dimensionality of a landscape and the k -value represents the number of epistatic interactions when determining a bitstring's fitness. As k increases its "ruggedness" or the number of local maxima increases. All experiments within this paper are run using landscapes of $n = 15$ and k varying from 0 (smooth) to 14 (rugged). Within the real world, a lower k represents a situation where there are fewer factors in play, and a higher k represents a situation where there are many factors in play. The NK model allows us to adjust the difficulty of the landscape an agent traverses through, allowing us to observe trends within those agents across different model parameters. Agents are placed within this landscape and improve as follows.

Agents and Actions

We define agents to start at a specific location (bitstring sequence) within an NK fitness landscape, this represents the starting point of the agent's developmental strategy. Rather than evolving (i.e. mutating) the starting locations in the landscape, in this work we focus on a consistent starting location and focus on the evolution of a "developmental program." Example target organisms for this model include Salmon and other migrant species (e.g. sea turtles) that are consistently born at a particular location and experience a similar lifelong sequence of challenges.

For our model, the "developmental program" consists of encoding two types of actions over an agent's lifetime: looking and walking. By separating looking actions and walking actions, we get further insight into the patterns of exploration and exploitation that emerge. Taking a look action allows an agent to collect information about the fitness of a bitstring location exactly one bit different from the agent's current location. Alternatively, taking a walk action allows the agent to use its information collected through looking actions to move from its current location to whichever location had the highest recorded fitness. If the highest fitness level of the collected information is lower than the current fitness, then the agent does not move. After a walk action,

the information gathered from previous looks is discarded. Taking combinations of these actions effectively allow an agent to emulate exploratory and exploitative developmental patterns, in a non-Markovian manner.

When an agent explores every possibility around it before walking, i.e. taking all possible look actions before a walk action, we describe it as a steepest hill climb. Conversely, when an agent does not take any looking actions before a walk action, we describe it as a random walk, as it results in the agent blindly moving to a new location without any knowledge of its surrounding landscape. Figure 1 illustrates the actions of an agent on a small ($N=3$) landscape with a genome of length 4.

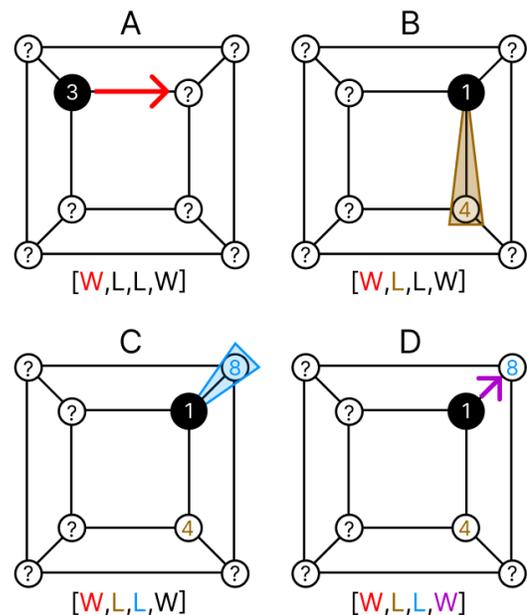


Figure 1: Illustration of an agent performing actions according to its genotype. W's represent walks and L's represent looks. Each graph (A-D) shows the connectivity of an $N=3$ landscape, with the nodes representing locations and their labels representing fitness (with ?s representing unknown fitness). A: The agent's initial location is represented by the black circle. It takes its first action as encoded by its genotype, a walk (W, red). There were no previous look actions so it takes a random walk (in this case to the right). B: Agent's new location is again represented by the black circle, but this time takes a look action (L, brown), resulting in a random adjacent location's fitness value being observed and recorded. C: Agent is still in the same location as in B and takes another look action (L, blue), resulting in another random adjacent location's fitness value being observed and recorded. D: Agent is still in the same location as it was in B and C, but this time takes a walk action (W, purple). The agent moves up and to the right to the highest fitness location recorded (8) and forgets the previous observations.

In this work, we define the lifespan of an agent to be its full developmental period, with a predetermined number of walk and look actions distributed across its lifespan. We take this model, as constrained in the following section, to provide a generalization for the development of organisms. More specifically, we aim to study high level processes guiding a lifetime learning strategy including exploitation and exploration, as represented by combinations of looking and walking steps. While we primarily focus on the evolution of lifetime learning strategies, the model could easily be extended to capture further biological processes, such as sensitive periods.

Evolutionary Algorithm

Within this model, we consider the combination of looking and walking actions over each agent’s lifetime development to be its genotype, that is to say the encoding of a process that guides its developmental trajectory. These genotypes are constrained by having a set number of looking and walking actions, to be distributed over the lifetime of the organism. In other words, we are looking at the times within the agent’s lifespan that it chooses to explore versus exploit.

By implementing these constraints, we create a common ground to analyze the genotypes. The genotypes of the organisms are evolved through a genetic algorithm as follows. First, the organisms within the population are selected based on their fitness at the end of their lifetime (i.e. fitness at the final location in the NK model). The top 50% of the organisms are directly copied to the next generation, and the remaining 50% of the new generation is comprised of mutations of these survivors, where one of the walk steps is randomly moved to a different part of the strategy to modify how look steps are distributed throughout the strategy. The new generation of organisms are then allowed to develop starting from the same starting locations as the previous generation. Through this process, the population of agents evolves their developmental strategy, finding a genotype that generally leads to a higher ending fitness.

In essence, an agent can develop using one of two primary developmental processes: exploration and exploitation. Exploration allows an organism to take risks for the sake of improvement, while exploitation focuses on refinement and efficiency. Within this model, we define exploration as taking multiple consecutive walk actions (resulting in “random walk” steps), while exploitation is taking look actions before a walk action. By allowing agents to evolve the distribution of their looking steps, i.e. choose when in their life to have exploratory and exploitative periods, we can observe the tendencies of these exploratory and exploitative periods to emerge from the evolutionary process.

Baseline Genotypes

We establish several baseline strategies for comparison with the evolved strategies. In order to do so, we will use three

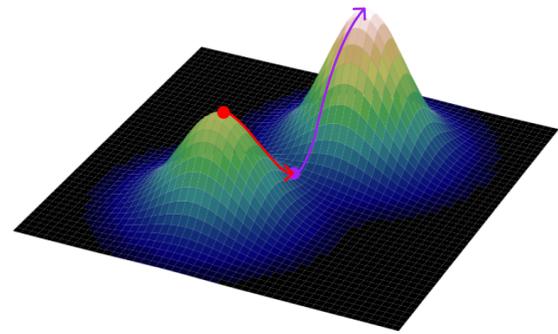


Figure 2: An artistic rendering of an NK landscape to demonstrate the behavior of agents. The starting location of an agent is marked as a red dot. The red path continuing from this depicts a possible path taken with a random walk action. The random walk ends at the red arrow, where the agent then continues its development with a steepest hill climb, reaching a higher optima than its original location. This demonstrates the utility in random walks allowing agents to escape from a local optima.

predetermined types of Walkers for comparison:

A ‘Random Walker’ (RW) takes exclusively walk actions throughout its entire lifetime. This simulates purely random decision making, with no environmental feedback. This walker is not expected to perform well, but will be used to demonstrate the importance of using look actions.

A ‘Steepest Hill Climber’ (SHC) walker looks at every possible adjacent location on the landscape before ever taking a walk action - purely steepest hill climb. This means that a ‘Steepest Hill Climber’ will always take the path of steepest ascent, and creates purely-exploitative behavior.

An ‘Alternating SHC/RW’ walker was designed to combine the strengths of RW and SHC. It accomplishes this by repeatedly taking two SHC steps and then a single RW step. This means it will spend the majority of its developmental process exploiting and ascending to a local optima, but will also have RW steps that allow it to escape local optima it may find itself in, creating potential performance improvements over the normal SHC strategy. Figure 2 provides a simple visualization how an agent at a local optima can use a RW step followed by a SHC step in order to escape a local optima.

Experimental Setup

Across all of our experiments, the parameters of the evolutionary process were held constant. Each evolutionary process consisted of a population of 100 initially randomized strategies simulated for 50 generations. These parameters

were chosen as a result of investigating how long the genetic algorithm generally takes to converge, in our case the algorithm rarely made progress after approximately 45 generations. An important consideration is that on lower- k landscapes the algorithm normally converged much faster, but to allow for fair comparison across different k -values we decided to keep the process consistent.

To make results more easily interpretable, we aimed to develop a more reliable fitness metric for our evolved walkers than the fitness value after a single run of their strategy. Since any strategy with a Random Walk step will vary greatly in its final fitness, there is a chance for a walker to simply get “lucky” and stumble into a high final fitness. In order to combat the inconsistency this would create, the fitness of an individual is determined by the average final fitness across 25 different runs of the strategy. This allows us to ensure that the final evolved strategy must consistently perform well to get a high fitness score.

Results

Relative Performance of Strategies

To verify that our evolved walker is showing improvement over the baseline walkers, we evolved separate walkers on 500 different landscapes for each k -values 0-14 with a consistent $n = 15$ and compared their results to the baseline strategies. Each strategy (evolved and baseline) was 200 steps long, and the evolved strategies were set to have exactly 40 walk steps in total. This 40 walk constraint was introduced to allow for a more direct comparison between different evolved strategies; the specific value of 40 walks was selected because when the number of walks is unconstrained the population regularly converges to having about 20% walk steps within any reasonable length strategy. The NK landscapes were randomly generated, though the same set of landscapes was used to test each individual type of Walker in order to ensure a fair comparison. The results of this experiment can be seen in Figure 3.

There are a few general trends to note from results in Figure 3. First, the random walker shows no sign of progress throughout its entire lifetime, as expected. This shows that without a productive developmental process, we can expect a final fitness of only slightly more than 0. Second, for the more adaptive strategies, we can see that each is able to solve the $k = 0$ landscape entirely and achieve the maximum fitness of 1, but as the difficulty of the landscape increases the fitness reached by each strategy decreases. This decrease in final fitness is sharp when progressing through lower k values, but eventually levels off at high k -values. This is all as expected - a more difficult landscape is inherently much harder to solve, so we would expect every strategy’s performance to degrade as the k -value increases.

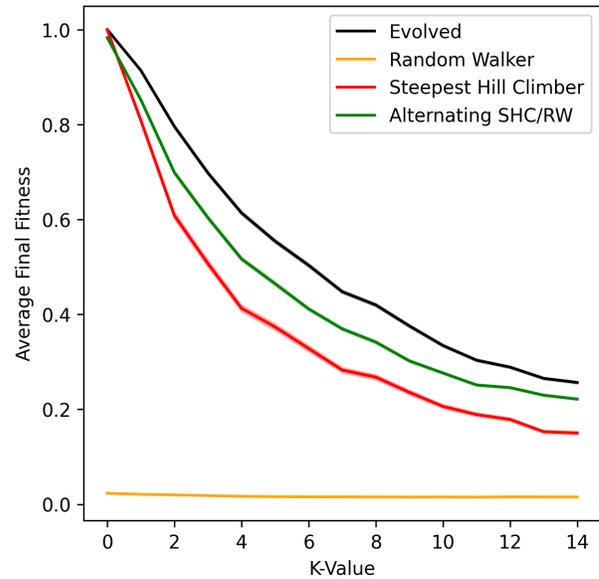


Figure 3: Final fitnesses of strategies across different K values. Average final fitness across 500 different landscapes for k -values from 0 to 14 on a $n = 15$ landscape. Strategy length 200 with 40 walk steps.

Distribution of Look Steps in Evolved Strategies

A more interesting trend appears at k -values greater than zero, where the SHC’s strategy begin to perform worse than the Alternating and Evolved Strategies. This is almost certainly the result of local optima being present on these non-trivial landscapes.

When a purely exploitative strategy like SHC encounters a local maxima, it has no recourse. It will exclusively takes actions that greedily improve its fitness, so once it enters a local maxima it won’t be able to improve anymore, and will then become ‘stuck’ for the remainder of its lifetime. Since local optima become more frequent at higher k -values this means that SHC strategy will increasingly get stuck at an earlier step in its lifetime, causing it to lose performance.

A solution to being stuck at a local optima is taking purely random, exploratory steps as the Alternating and Evolved strategies do. Although random walks are by themselves unproductive, a random walk action followed by subsequent exploitative actions gives a strategy the potential to escape that local optima (see Figure 2 for an illustration). This is why the evolved walker and alternating walker offer improvement over SHC - they are able to escape local optima which results in overall better performance. Although this explains why the SHC strategy performs worse than these two other strategies on more difficult landscapes, the reason for the performance difference between the Evolved and Alternating strategy requires deeper analysis into the structure of these evolved strategies.

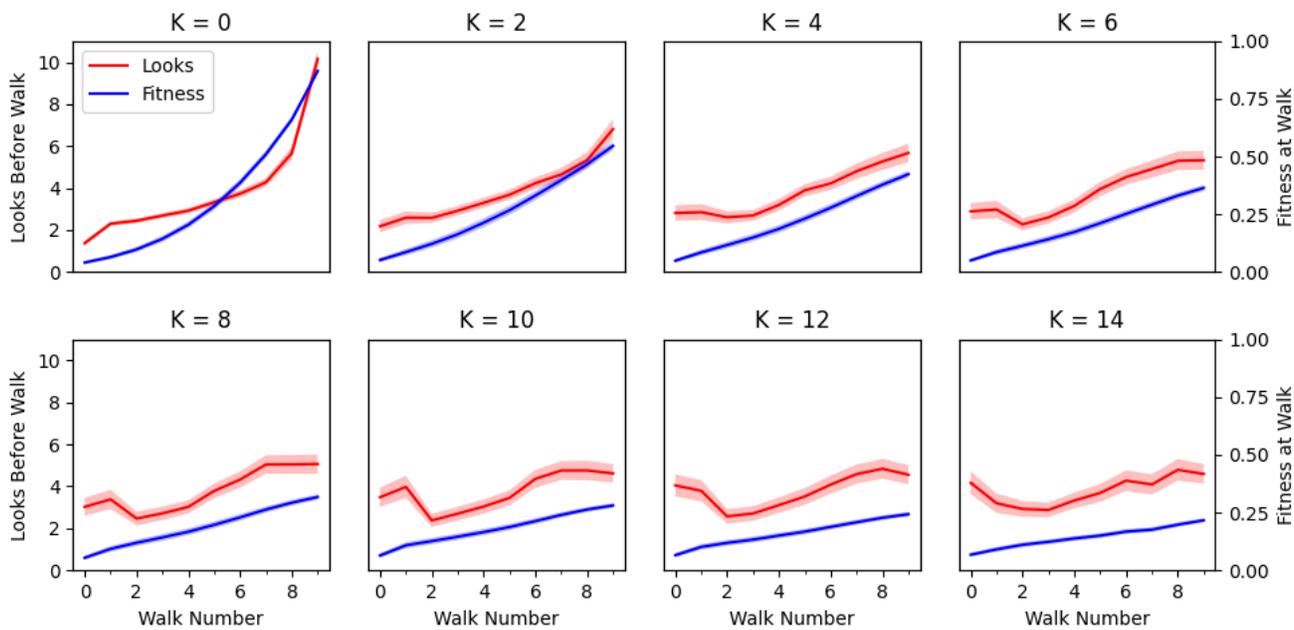


Figure 4: Look before walks and fitness at each walk across K values. Average distribution of Look Steps (in red) and fitness (in blue) of the evolved strategy across 500 landscapes, with error bars (standard deviation). Strategy length at 50 steps with 10 walk steps.

To do this, we will investigate the distribution of look actions in the final evolved strategy in landscapes of certain k -values. We investigated the trends among several sets of parameters and noticed similar structures across most strategy lengths as long as the ratio of walk to total steps remained consistent, so we decided to continue by investigating a shorter length strategy to reduce noise and allow easier data interpretation. These shorter strategies will be 1/4th the length of the previous strategies, with a total length of 50 steps, 10 of which are walks to retain the original 20% walk percentage.

Figure 4 shows the average walk distribution across different landscape difficulties with this more controlled strategy length and reveals several important trends. The first is that across low-difficulty landscapes, we see a preference for early-development exploration (low-look walk steps) and a preference for late-development exploitation (high-look walk steps). This aligns with what has often been observed in developmental processes (Spreng and Turner, 2021). However, once the landscapes become sufficiently difficult (in this case, $k \geq 10$), we see that this trend is significantly weaker, and the level of exploration/exploitation throughout the developmental process doesn't vary nearly as much between early and late development. We believe this is a result of need to escape local optima via use of the exploratory 'Random Walk' steps. At higher k -values we see more frequent local optima, and these exploratory steps be-

come more important at even the late stages of development in order to escape increasing amounts of local optima. The distribution of these purely exploratory random walk steps is discussed in depth in the next section.

Another important trend is how on more difficult landscapes ($k \geq 6$) we see that opening two steps of the strategy frequently have a higher preference for exploitation than the following few developmental steps. A likely explanation for this trend is that Random Walk steps aren't necessarily useful if an individual is not currently located at a local optima, and the chance that an individual is at a local optima before taking any exploitative steps is fairly low. This means that we would expect each walker to take a few exploitative steps before preferring exploration on these high- k landscapes, which is the trend we see emerging.

Distribution of Random Walks in Evolved Strategy

Using the same experimental setup as before, with 50-length strategies that each have 10 walks, we now analyze how the 0-look 'Random Walk' steps are distributed (see Figure 5). The most basic case with $k = 0$ shows no exploratory periods beyond the very first step. This is a result of the specific landscape containing a single global optima and no other local optima. This means that any purely exploitative strategy will succeed on this landscape, so there is never a need for the Random Walk steps.

However, once contribution factor dependence is intro-

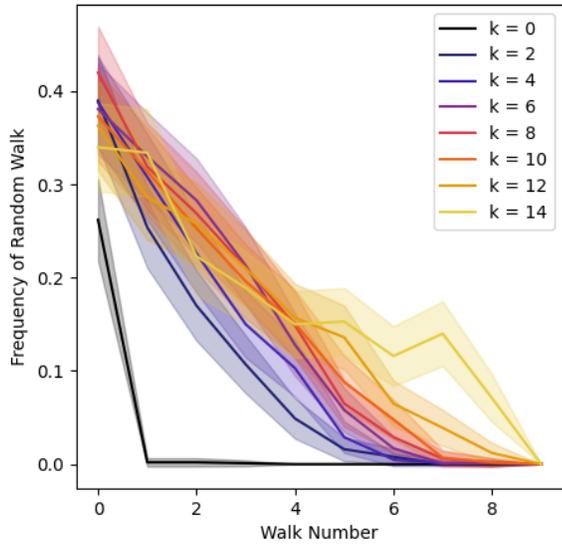


Figure 5: Frequency of Random Walk at Different Walk Steps. Percentage of evolved strategies (1 = 100%) with a Random Walk at each step in their developmental process, with error bars (standard deviation). 500 strategies used at each k -value, and each strategy was evolved on a different landscape to avoid data bias.

duced ($k > 0$), we see the evolutionary process begin selecting for later-in-development exploration steps. These Random Walk steps are present throughout the developmental process for every k -value besides $k = 0$, and as the difficulty continues to increase, we see these exploration actions persist into mid-development, and even into late-development. An important factor to note is that while higher k -valued landscapes do contain more midlife exploration steps, they do not reduce the number of exploration actions in early development. This can be deduced by the fact that for every landscape with k higher than 2 there is a similar chance (± 0.05) for a walk step at the first few steps of development. This shows that while we do see exploration actions persisting later in development on harder landscapes, this does not mean that early development exploration actions are unnecessary, rather this shows the contrary—mid-development exploration periods are more effective when paired with exploration earlier in development.

Another factor that reinforces this fact is how in Figure 3, we see a convergence between the final fitness of the evolved and alternating SHC/RW strategies at high k -values. This is representative of the fact that at higher- k landscapes with more local optima, the strategy to alternate between exploration and exploitation becomes more prominent to escape these optima. Eventually, when the k -value

gets high enough there are so many local optima that the alternating strategy preforms almost equally as well as our evolved strategy at these extremely high k -values. Although this doesn't necessarily mean that the evolved and alternating strategies are using the same method to achieve their fitness, this does tell us that the importance of escaping local optima becomes increasingly important as the k -value increases.

Discussion

By simulating the evolution of developmental strategies in an abstract model, this paper reveals several insights on the relationship between evolution and development that can apply to various organisms and developmental processes. At the highest level, evolution produces developmental steps that explicitly explore the landscape in a random manner. In fact, the presence of such randomly exploring steps is shown to be necessary for all developmental strategies to survive selection, regardless of their location and frequency within the developmental period. It is interesting that evolution drives organisms to ignore environmental cues at times against the promise of a fitness ascent. However, the superiority of populations that develop with random exploration can be attributed to their divergence across a larger search space in the landscape, avoiding convergence at local optima and redirecting themselves to regions of potentially higher fitness. This may help explain the emergence of developmental phenomena such as child rebellion in humans and other animals (Sachser et al., 2018; Chakradhar, 2018; Spear, 2000), where evolution has selected for populations whose adolescents deviate from developmental trajectories fostered by their parents to achieve greater phenotypic and behavioral diversity. Conversely, in the case that only one global optimum exists, represented by the landscape with $k = 0$ (no interdependent factors), evolution indeed produces developmental strategies with no random exploration.

The location of the exploratory steps offers another insight. Evolution drives development to be time-sensitive, with non-uniform patterns of exploration and exploitation comprising the developmental strategy. The populations of evolved strategies in various landscapes all consistently exhibit different degrees of exploitation (and hence those of exploration) at different points in development, notably with large changes at the early and terminal stages of development. For example, the extensive exploitation towards the end of development can be attributed to its direct impact on the final fitness, which is the only fitness function used to evolve the strategies in this model. Similarly, sensitive periods in biological systems demonstrate time-sensitive development and significantly affect organisms towards the end of their development. As the adult form of many organisms are fixed for the rest of their lifetimes with low plasticity, the terminal stage of development is critical for their evolutionary fitness (Spreng and Turner, 2021; Del Maschio et al.,

2018; Brehmer et al., 2014). These are consistent with observations of the uniformly distributed strategy alternating between exploration and exploitation, which was not able to develop higher fitness than the evolved, non-uniform strategies.

Besides the terminal stage, the early and middle stages of development that emerge on different landscapes reveal another perspective. Evolution selects for a more complex developmental strategy marked by multiple transitions between predominantly exploitative and exploratory phases, as the landscape becomes more complex with a greater number of interdependent factors. The early exploitation and middle exploration stages that emerge on high- k landscapes exemplify such phases and their transitions, along with the terminal exploitation stage common to all landscapes. Considering their contrast with the simpler, early exploration stage on low- k landscapes, the multiple transitions can be attributed to a potential mechanism related to the increased number of local optima (and hence the increased risk of suboptimal convergence). The early ascent through exploitation may guide organisms to regions of generally higher fitness, increasing the probability that the following exploration will place the organism near a high-fitness optimum prior to the terminal exploitation. Although it is inconclusive from this work whether this mechanism is truly responsible for the emergence of multiple transitions, it offers a useful direction for future work to verify the presence of such regions or examine completely different hypotheses.

More importantly, it is surprising to note that nature and this highly abstract model share the emergence of a more complex developmental strategy across increasingly complex landscapes. The quantitative development of flatworms exemplifies a simple landscape with few interdependent factors, exhibiting a simple exploitative development increasing or reducing in size based on the nutrition available (Martín-Durán and Egger, 2012). In contrast, the cognitive development of humans has a uniquely high number of interdependent factors due to social influences, evolving multiple distinct stages of development with varying degrees of exploration and exploitation (Thompson, 2021). Along with the aforementioned findings in the value of exploration and the time-sensitivity of development, the emergence of this trend demonstrates the ability of this model to simulate and examine various interactions between evolution and development. Next, we will discuss potential extensions to the model, which would enable future investigations.

Future Work

We see three major directions for future work based on the preliminary results reported on in this paper. First, in our model the only way to escape a local optima is by taking a purely random walk step and getting lucky. A simple extension would be to include a multi-tier look (i.e. change multiple bits in the bitstring) to more intelligently escape optima.

This multi-tier look could also incur an additional cost and potentially offer insight into the underlying selection pressures for sensitive periods. This multi-tier look could allow for greater relative sensitivity to an environment (along with greater cost), which in turn could help identify the usefulness of critical periods in a developmental trajectory.

Second, our model currently only evolves the look and walk actions with a fixed starting location. Often, models employing NK fitness landscapes involve the evolution of the starting location. Extending the model to include starting location would allow more thorough exploration of the interaction between evolution and development.

Finally, in addition to the aforementioned extensions of this model, we see an interesting opportunity to conduct experiments in dynamic fitness landscapes where the fitness values shift over time. This could happen within an agent's lifetime, periodically much like seasonal changes, but also could happen across evolutionary time such that a more flexible developmental strategy might be required to sustain a population. Interesting possibilities could include analysis of the relationship between agent lifetime and the scale of the changes in the environment.

Acknowledgements

We would like to thank the reviewers for constructive feedback which greatly improved the final version of this paper. We would also like to thank Robert J Williamson and Mitchel J Daniel for helpful feedback on an earlier draft of this paper. Finally, we would like to thank Eduardo J Izquierdo and Madhavun Candadai for their conceptual contributions to early explorations of this project.

References

- Altenberg, L. (1996). NK Fitness Landscapes. *Evolution*, pages 1–11.
- Belew, R. K. (1990). Evolution, Learning, and Culture: Computational Metaphors for Adaptive Algorithms. *Complex Systems*, 4:11–49.
- Brehmer, Y., Kalpouzos, G., Wenger, E., and Lövdén, M. (2014). Plasticity of brain and cognition in older adults. *Psychological Research*, 78(6):790–802.
- Campbell, C. M., Izquierdo, E. J., and Goldstone, R. L. (2020). How much to copy from others? The role of partial copying in social learning. In *CogSci*.
- Chakradhar, S. (2018). Animals on the verge: What different species can teach us about human puberty. *Nature Medicine*, 24(2):114–117.
- Del Maschio, N., Sulpizio, S., Gallo, F., Fedeli, D., Weekes, B. S., and Abutalebi, J. (2018). Neuroplasticity across the lifespan and aging effects in bilinguals and monolinguals. *Brain and Cognition*, 125:118–126.
- Fragata, I., Blanckaert, A., Dias Louro, M. A., Liberles, D. A., and Bank, C. (2019). Evolution in the light of fitness landscape theory. *Trends in Ecology & Evolution*, 34(1):69–82.

- Frankenhuis, W. E. and Walasek, N. (2020). Modeling the evolution of sensitive periods. *Developmental Cognitive Neuroscience*, 41:100715.
- Geard, N., Wiles, J., Hallinan, J., Tonkes, B., and Skellett, B. (2002). A comparison of neutral landscapes - NK, NK_p and NK_q. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No.02TH8600)*, volume 1, pages 205–210. IEEE.
- Hebbron, T., Bullock, S., and Cliff, D. (2008). NK α : Non-uniform epistatic interactions in an extended NK model. *Artificial Life*, XI.
- Hinton, G. and Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, 1(3):495–502.
- Kauffman, S. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45.
- Knudsen, E. I. (2004). Sensitive Periods in the Development of the Brain and Behavior. *Journal of Cognitive Neuroscience*, 16(8):1412–1425.
- Lynch, J. A., Brent, A. E., Leaf, D. S., Pultz, M. A., and Desplan, C. (2006). Localized maternal orthodenticle patterns anterior and posterior in the long germ wasp *Nasonia*. *Nature*, 439(7077).
- Martín-Durán, J. and Egger, B. (2012). Developmental diversity in free-living flatworms. *EvoDevo*, 3(1):7.
- Müller, G. B. (2007). Evo–devo: extending the evolutionary synthesis. *Nature Reviews Genetics*, 8(12):943–949.
- Panchanathan, K. and Frankenhuis, W. E. (2016). The evolution of sensitive periods in a model of incremental development. *Proceedings of the Royal Society B: Biological Sciences*, 283(1823):20152439.
- Park, S.-C., Neidhart, J., and Krug, J. (2015). Greedy adaptive walks on a correlated fitness landscape. *Journal of Theoretical Biology*.
- Pitzer, E. and Affenzeller, M. (2012). A comprehensive survey on fitness landscape analysis. In *Studies in Computational Intelligence*, volume 378, pages 161–191.
- Rhodes, M. L. and Dowling, C. (2018). ‘What insights do fitness landscape models provide for theory and practice in public administration?’. *Public Management Review*, 20(7):997–1012.
- Roth, S. and Hartenstein, V. (2008). Development of *Tribolium castaneum*. *Development Genes and Evolution*, 218(3-4):115–118.
- Sachser, N., Hennessy, M. B., and Kaiser, S. (2018). The adaptive shaping of social behavioural phenotypes during adolescence. *Biology Letters*, 14(11):20180536.
- Smart, J. M. (2019). Evolutionary Development: A Universal Perspective. In *Springer Proceedings in Complexity*, pages 23–92.
- Soltoggio, A., Stanley, K. O., and Risi, S. (2018). Born to learn: The inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Networks*, 108:48–67.
- Sommer, R. J. (2009). The future of evo-devo: Model systems and evolutionary theory.
- Spear, L. (2000). Modeling adolescent development and alcohol use in animals. *Alcohol Research and Health*, 24(2).
- Spreng, R. N. and Turner, G. R. (2021). From exploration to exploitation: a shifting mental mode in late life development. *Trends in Cognitive Sciences*, 25(12):1058–1071.
- Thompson, M. J. (2021). Piaget’s stages of cognitive development and erikson’s stages of psychosocial development. In *Child and Adolescent Mental Health: Theory and Practice*.
- Todd, G., Candadai, M., and Izquierdo, E. J. (2020). Interaction Between Evolution and Learning in NK Fitness Landscapes. In *The 2020 Conference on Artificial Life*, pages 761–767, Cambridge, MA. MIT Press.
- Walasek, N., Frankenhuis, W., and Panchanathan, K. (2021). An evolutionary model of sensitive periods when the reliability of cues varies across ontogeny. *Behavioral Ecology*.
- West-Eberhard, M. J. (2005). Developmental plasticity and the origin of species differences. *Proceedings of the National Academy of Sciences*, 102(suppl.1):6543–6549.
- Wilke, C. O. and Martinetz, T. (1999). Adaptive walks on time-dependent fitness landscapes. *The American Physical Society*, 60(2):2154–2159.