



HAL
open science

Dynamics of niche construction in adaptable populations evolving in diverse environments

Eleni Nisioti, Clément Moulin-Frier

► To cite this version:

Eleni Nisioti, Clément Moulin-Frier. Dynamics of niche construction in adaptable populations evolving in diverse environments. *ALIFE 2023 - The 2023 Conference on Artificial Life*, Jul 2023, Sapporo, Japan. 10.1162/isal_a_00668. hal-04356907

HAL Id: hal-04356907

<https://inria.hal.science/hal-04356907>

Submitted on 11 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Dynamics of niche construction in adaptable populations evolving in diverse environments

Eleni Nisioti¹ and Clément Moulin-Frier¹

¹ Flowers Team, Inria and Ensta ParisTech, Bordeaux, France
eleni.nisioti@inria.fr

Abstract

In both natural and artificial studies, evolution is often seen as synonymous to natural selection. Individuals evolve under pressures set by environments that are either reset or do not carry over significant changes from previous generations. Thus, niche construction (NC), the reciprocal process to natural selection where individuals incur inheritable changes to their environment, is ignored. Arguably due to this lack of study, the dynamics of NC are today little understood, especially in real-world settings. In this work, we study NC in simulation environments that consist of multiple, diverse niches and populations that evolve their plasticity, evolvability and niche-constructing behaviors. Our empirical analysis reveals many interesting dynamics, with populations experiencing mass extinctions, arms races and oscillations¹. To understand these behaviors, we analyze the interaction between NC and adaptability and the effect of NC on the population's genomic diversity and dispersal, observing that NC diversifies niches. Our study suggests that complexifying the simulation environments studying NC, by considering multiple and diverse niches, is necessary for understanding its dynamics and can lend testable hypotheses to future studies of both natural and artificial systems.

Introduction

Biological organisms and their environments share a reciprocal relationship: organisms survive and reproduce under selection pressures present in their habitats and environments are modified by their inhabitants, with changes being inherited to the next generation and accumulating with evolutionary time (Odling-Smee et al., 2003). The first process, natural selection, became the cornerstone of evolutionary theory in the early 20th century (Provine, 2001) and, in both natural and artificial life studies, is often seen as synonymous to evolution. The second, niche construction (NC), was characterized as the neglected process in evolution in the early 2000s (Odling-Smee et al., 2003), as evolutionary theory assumed that NC's effect on selection pressures is negligible. Since then, studies of natural populations have

shown that NC often affects selection pressures by helping species protect themselves from environmental uncertainty and accelerates evolution by complexifying environments (Odling-Smee et al., 2003; Schwilk, 2003; Boivin et al., 2016). By now, NC is, at best, the elephant in the room of evolutionary synthesis: despite evidence for its existence, our understanding of it is too limited to enable its study in settings capturing the complexity of the real world.

Recent hypotheses studying major events in evolution point to environmental complexity (Maslin et al., 2015; Potts, 2013; d'Errico et al., 2017). For example, the birth of our own lineage in East Africa occurred amidst large climatic instability that fragmented the landscape into patches of land differing in resource availability and separated by large lakes (Maslin et al., 2014). Other hypotheses further emphasize the ability of natural landmarks such as deserts and oceans to form barriers that isolate populations (Larra-soaña et al., 2013). A similar story is unfolding in artificial life: under Quality-Diversity optimization (Pugh et al., 2016), the search space is divided into behavioral niches to evolve diverse solutions, a paradigm that challenged the dominant approach of training a single agent, as it proved more robust in real-world settings (Cully et al., 2015).

When in a heterogeneous environment, organisms can survive by: a) specializing in a niche to out-compete others, paying the cost that they are out-competed in other niches Nisioti and Moulin-Frier (2022); Grove (2014) b) becoming *phenotypically plastic*: a plastic individual adapts its phenotype to its environment without genetic change and survive in diverse niches within its lifetime. On the downside, plasticity comes with fitness costs, so that plastic individuals are out-competed by specialists in their preferred niche (Grove, 2014; Nisioti and Moulin-Frier, 2022) c) becoming highly evolvable: a high mutation rate enables quick adaptation at an evolutionary scale but comes at the cost of increased deleterious mutations (Lynch et al., 2016; Nisioti and Moulin-Frier, 2022). As both plasticity and evolvability enable adaptation we jointly refer to them as adaptability.

Although limited, our understanding of NC shows that it is influenced both by the adaptability of populations and the

¹We provide an online repo for reproducing our simulations at <https://github.com/eleninisioti/NicheConstructionModel/tree/main>

diversity of environments. For example, processes such as the invention of agriculture (Zohari, 1986) and cultural innovation (Mesoudi and Thornton, 2022) are both influenced by geography and our impressive social learning abilities (Diamond, 1998; Migliano et al., 2020; Migliano and Vinicius, 2022). Explanations for this rely on the inter-play between a population's connectivity and its ability to accumulate solutions: spatial heterogeneity enforces isolation and, therefore, diversity within a population, while adaptability ensures that environmental barriers can be crossed and solutions spread, causing an intensification of NC (Cantor and Whitehead, 2013; Derex and Boyd, 2016; Nisioti et al., 2022).

In this work, we study the evolution of NC in environments divided in diverse niches where populations evolve their adaptability and niche-constructing behavior. Each niche is characterized by its environmental state, which determines how many agents it can fit, and consists of two components: the intrinsic state (which can for example model climatic variations uncontrollable by the agents) and the niche-constructed state, which is the product of niche-constructing individuals inhabiting the niche. The niche-constructed state is also inherited from the previous generation with some decaying applied to capture the fact that ecologically-inherited artifacts cannot persist indefinitely. Modeling NC as a process that changes the capacity of the environment is common (Laland et al., 1999; Krakauer et al., 2009) and finds inspiration in natural behaviors, such as nest-constructing species. Genomes consist of four genes: preferred environmental state (the state at which their fitness is highest), plasticity (the variance in environmental state they can tolerate without a large impact on fitness), evolvability (the mutation rate) and niche construction (the amount by which they change the environmental state of a niche they inhabit). We consider two different mechanisms for selecting which agents will reproduce: a) under *global competition* we select them with a probability proportional to their average fitness across the niches they can survive in and reproduction stops when the environment's capacity is reached b) under *local competition* we consider reproduction in each niche independently: an agent is selected based on its fitness in each niche and reproduction stops when the capacity of that niche is filled. Thus, agents benefit from surviving in multiple niches. Our models of local and global competition have their equivalents both in previous studies in evolutionary computation (with Quality-Diversity algorithms representing local competition and classical algorithms global) and in the study of natural populations (where local adaptation considers that populations of the same species inhabiting different niches experience different natural selection pressures (Leimu and Fischer, 2008)).

An agent may benefit from NC by: a) increasing the capacity of a niche, for example through nest-building, thus, reducing competition for resources b) reducing the capacity of a niche thus making it less desirable to others. This

behavior, which we refer to as negative niche construction, is exemplified in nature by pines, which spread their needles to increase the probability of fire and out-compete not fire-resistant species (Schwilk, 2003) c) bringing the environmental state closer to its preferred niche and further away from others For example, aquatic earthworms can inhabit earth soil only by changing its consistency (Turner, 2000). But NC also comes at a cost: a) increasing the capacity of a niche may invite other species in and increase competition b) decreasing its capacity may make it uninhabitable c) changing the environmental state will require adaptability, which comes at its own cost d) multiple agents co-inhabiting a niche can create high environmental uncertainty that can lead to mass extinctions if adaptability cannot increase enough to respond to it, a case that feels most relevant to our own evolution (Boivin et al., 2016; West, 2017).

Taking into account these complex interactions between environmental heterogeneity, population adaptability and niche construction, we believe that studies of NC should depart from simplified settings with a single niche (Laland et al., 1999; Suzuki and Arita, 2009; Chiba et al., 2020) and non-plastic populations (Laland et al., 1999; Krakauer et al., 2009). As we show in this work, rich models can offer interesting hypotheses about the dynamics of NC, such as that:

- having multiple niches is necessary for avoiding mass extinctions. We observed that, in a single niche, environmental uncertainty induced by NC becomes too high for adaptability to cope with;
- NC promotes adaptability and populations adapt differently depending on the selection mechanism: under global competition NC leads to higher plasticity while under local to higher evolvability;
- the population may niche-construct negatively for prolonged periods of time, making niches uninhabitable. Agents deal with this by evacuating the niche until it becomes inhabitable due to the decay of NC;
- under local competition NC increases genomic and environmental diversity

Related works

Previous studies have primarily modeled NC in two different ways: as an increase in the amount of resources that increases the environment's capacity (Laland et al., 1999; Krakauer et al., 2009) and as a direct increase or decrease in fitness that does not affect capacity (Suzuki and Arita, 2009). We follow the first approach but also allow reductions in capacity. NC studies were at first primarily theoretical, employing differential equations to predict how NC affects evolution (Laland et al., 1999; Krakauer et al., 2009). Such models lent intuitive insights, such that NC can emerge without direct selection (Laland et al., 1999) and that

the ability to monopolize niches enables the emergence of NC (Krakauer et al., 2009) but come with limitations, such as the fact that niches are identical and agents are not plastic.

Suzuki and Arita (2009) study the co-evolution of plasticity and NC in an agent-based model and show that populations alternate between phases of high plasticity and NC, provided that agents niche-construct sequentially. This model differs from ours as the environment consists of a single niche, NC does not modify its capacity and evolvability is constant. Interestingly, when Suzuki and Arita (2009) applied NC in parallel these patterns disappeared due to agents canceling out each other’s behavior. Instead, the patterns in our analysis occur with NC applied in parallel. As we show, this becomes possible due to the presence of multiple niches that stabilize NC. Nisioti and Moulin-Frier (2022) studie the co-evolution of plasticity and evolvability in an environment with multiple niches, but do not consider NC.

Recent advents in deep reinforcement learning (DRL) and neuroevolution have enabled the study of complex behaviors in simulations, such as foraging and tool-use (Pérolat et al., 2017; Hamon et al., 2023; Baker et al., 2020). Chiba et al. (2020) leverage such techniques to study NC in a 2D environment where agents can construct artifacts useful for avoiding predation. Our study can lend insights to this direction, as we can view DRL as the mechanism that underlies the behavioral plasticity considered in our model.

Modeling and methodology

We now separately discuss our model of the environment and genomes, the evolutionary algorithm and the set of metrics we use to monitor evolution.

Modeling the environment

The environment, illustrated in Figure 1, is divided into N niches arranged in a simple latitudinal model: we consider a reference niche at $n = 0$, $N/2$ “northern” niches indicated with positive indexes $n \in (0, N/2]$ and $N/2 - 1$ “southern” niches with negative indexes $n \in (-N/2, 0)$. Each niche is characterized by its environmental state e_n^g which is the sum of two elements:

- an intrinsic state i_n that remains constant with evolutionary time and depends on the location of the niche. Specifically, the state of niche n at generation g is $i_n = i_0 + \epsilon \cdot n$, where i_0 is the state of the reference niche and ϵ is a constant capturing the difference between adjacent niches.
- the niche-constructed state b_n^g , capturing the modifications that agents inhabiting the niche cause. These modifications are carried over generations but are discounted by a factor γ . Formally, the niche-constructed state b_n^g of niche n is equal to $b_n^{g-1} \cdot \gamma$, with $\gamma < 1$, plus the total amount of NC applied by all agents that reproduced in this niche in generation g . We denote this latter amount as $\sum_{k \in \mathcal{K}_n} a_k^g$, where \mathcal{K}_n is the subset of all agents that

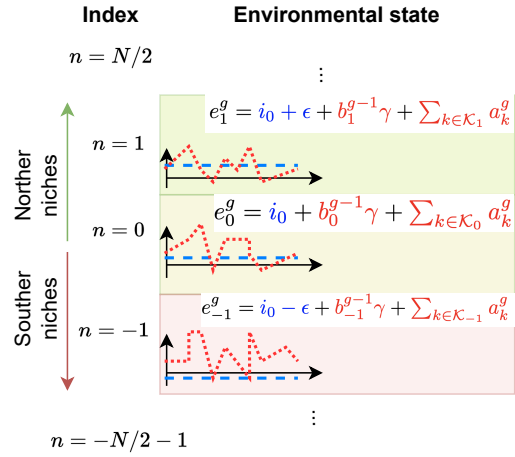


Figure 1: Our model of the environment: each niche n is characterized by its environmental state e_n^g which is the sum of the intrinsic state i_n (in blue) and the niche-constructed state b_n^g (in red). i_n depends on the state of the reference niche (e_0), the offset ϵ and n . b_n^g is the sum of the niche-constructed state at the previous generation discounted by γ plus the sum of the niche-constructing behavior of agents that reproduced in this niche in the current generation.

reproduced in niche n and a_k^g is the amount of niche construction by a single agent (we later explain how agents niche-construct).

Thus, the general equation describing the evolution of niche n with generations g is:

$$b_n^g = b_n^{g-1} \cdot \gamma + \sum_{k \in \mathcal{K}_n} a_k^g$$

$$e_n^g = i_0 + \epsilon \cdot n + b_n^g \quad (1)$$

As we explain later, the environmental state determines the fitness of an agent based on its genome and the capacity of the niche c_n^g as: $c_n^g = e_n^g C_N$, where C_N is the reference niche capacity. Thus, higher environmental states can support larger populations and are termed “high-quality”. To ensure that the maximum population size is independent of the number of niches we define $C_N = C_{\text{ref}}/N$, where C_{ref} is equal to the desirable maximum population size. An assumption of this model is that there is spatial smoothness, i.e. nearby niches are similar. Note that C_{ref} is independent of niche-constructing behavior: by niche-constructing the population can exceed this capacity. We, therefore, further bound the population size by randomly discarding agents if the population exceeds a value K_{max} .

Modeling the genome

We model plasticity through tolerance curves, a tool developed in ecology (Lynch and Gabriel, 1987) and previously

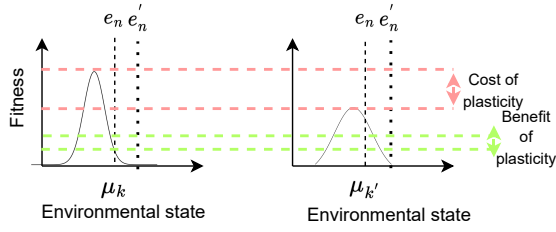


Figure 2: Modeling plasticity as a normal distribution $\mathcal{N}(\mu_k, \sigma_k)$. A non-plastic individual (k) has small σ_k and a high peak at their preferred niche, while a plastic individual (k') has large σ_k and a lower peak at their preferred niche. Fitness in a given niche n is computed as the probability density function of the distribution at the environmental state e_n . This figure also illustrates the cost and benefit of plasticity, assuming that $\mu_k = \mu_{k'}$. If $e_n = \mu_k$ (the actual environmental state is identical to the preferred niche of both individuals) the plastic individual has lower fitness (cost of plasticity). If e_n differs significantly from μ_k (the actual environmental state differs from the preferred one) the plastic individual has higher fitness (benefit of plasticity).

employed in simulation environments (Grove, 2014; Nisioti and Moulin-Frier, 2022). A tolerance curve is a normal distribution with mean μ_k^g , indicating the environmental state of highest fitness for an individual k at generation g , called the preferred state, and standard deviation σ_k^g that captures how quickly the fitness of the genome drops as the environmental state varies from its preferred state. Genomes with large σ_k^g are indicative of plastic individuals (we illustrate the tolerance curves of a plastic and a non-plastic individual in Figure 2).

To model niche construction we introduce a niche-constructing gene $a_k \in [-N/K_{\max}, N/K_{\max}]$ where we bound the amount an agent can niche-construct to ensure that agents inhabiting large environments have limited capabilities. The niche-constructing gene is expressed in the niche an agent reproduces in by modifying its environmental state. Thus, at each generation, niche n is constructed by an amount $\sum_{k \in \mathcal{K}_n} a_k$, where \mathcal{K}_n denotes the subset of agents that reproduced in it. The genome σ_k^g also includes the mutation rate r_k^g . Thus the complete form of a genome is $\sigma_k^g = [\mu_k^g, \sigma_k^g, r_k^g, a_k^g]$ and, upon reproduction, it mutates as:

$$\begin{aligned} \mu_k^{g+1} &= \mu_k^g + \mathcal{N}(0, r_k^g) \\ \sigma_k^{g+1} &= \sigma_k^g + \mathcal{N}(0, r_k^g) \\ a_k^{g+1} &= [c_k^g + \mathcal{N}(0, r_a)]_{-N/K_{\max}}^{N/K_{\max}} \\ r_k^{g+1} &= r_k^g + \mathcal{N}(0, r_k^g) \end{aligned} \quad (2)$$

where $\mathcal{N}(x, y)$ denotes a normal distribution with mean x and variance y and we have highlighted that the niche-

constructing gene is bounded. Note also that, for stability reasons, we employ a different mutation rate for the niche-constructing gene, r_a , that remains constant.

Global and local selection

At the end of a generation agents are selected for crossover based on their fitness. Each chosen individual has two offspring and the next generation consists only of offspring. To compute the fitness of an individual k in generation g we first detect the niches in which it can survive as:

$$n \in \{1, \dots, N\} \mid e_n^g \in [\mu_k^g - 2\sigma_k^g, \mu_k^g + 2\sigma_k^g] \quad (3)$$

and compute its fitness in each one of them as $f_{k,n}^g = pdf(\mu_k^g, \sigma_k^g, e_n^g)$, where pdf denotes the value of the normal probability density function with mean μ_k^g , and variance σ_k^g at location e_n^g . We study two selection mechanisms:

- under global selection all agents in the population are ranked based on their average fitness across the niches they can survive in and reproduce with a probability proportional to it. Agents reproduce until the capacity of the environment is filled.
- under local selection we apply the same criterion but independently for each niche: we detect which agents survive in a niche and reproduce them with a probability proportional to their fitness. Thus, agents that can survive in multiple niches have higher chances of reproduction under this mechanism. Again, agents reproduce in a niche until its capacity is filled.

Thus, at the end of a generation a new population is formed that consists of the offspring of agents from the previous population that were chosen for reproduction, based on their fitness and the environment's capacity. Note that the niche is not inherited from parent to offspring: the offspring will inhabit the niches its genome describes (as in Eq. 3). We present the pseudocode of our algorithm in Algorithm 1 and 2, provided in our online repo due to limited space.

Metrics

In addition to population-wide averages of the genome values and environment-wide averages of the intrinsic and niche-constructed states we monitor the following metrics:

1. $X^g = \sum_k X_k^g$, the number of extinctions. We denote the survival of individual k in niche n at generation g as a binary variable:

$$s_{k,n}^g = (e_{n,g} \in [\mu_k^g - 2\sigma_k^g, \mu_k^g + 2\sigma_k^g]) \quad (4)$$

Thus, an individual goes extinct ($X_k^g = 1$) if $\sum_n s_{k,n}^g$ is zero and survives ($X_k^g = 0$) if $\sum_n s_{k,n}^g$ is positive.

2. V_μ^g , the diversity of the population defined as the standard deviation of the population's preferred state, namely

$$V_\mu^g = \sigma_{\mu^g} \quad (5)$$

This metric captures the genetic diversity of the population computed for the gene of preferred state.

3. D^g , the dispersal of the population, computed as the number of niches over which at least one individual survives for a temporal window of at least w generations. Formally, $D^g = \sum_{n=1}^N d_{n,w}^g$, where $d_{n,w}^g$ denotes the persistence of the population in a given niche for the required time window and is computed as

$$d_{n,w}^g = \begin{cases} 1 & \text{if } \sum_{g'=g-w}^g s_n^{g'} = w \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where s_n^g is indicates the survival of at least one individual in a given niche and is computed as

$$s_n^g = \begin{cases} 1 & \text{if } \sum_k^K s_{k,n}^g > 1 \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

with $s_{k,n}^g$ defined in Eq. (4).

4. H^g , the competition for reproduction within the population. We count, for each niche, the number of agents that survived in it and were fit enough to be chosen for reproduction but did not reproduce because its capacity was reached and, then, sum over all niches.

Results

We now examine the behavior of agents under different settings. First, we compare behaviors under global competition between environments with different number of niches. After, we consider only heterogeneous environments ($N = 100$) niches and we contrast the behavior across two dimensions: a) whether NC takes place or not (where we force all niche-constructing genes to be zero) and b) whether selection is local or global. In all simulations we set the discount factor γ to 0.5, the intrinsic state of the reference niche i_0 to 0.6, ϵ (the difference between adjacent niches) to 0.01, the reference capacity to $C_{\text{ref}} = 1000$, the maximum population size to $K_{\text{max}} = 5000$ and the mutation rate of the niche-constructing gene, r_c , to 0.0003. As we discuss later it would be interesting to study the effect of some of these hyper-parameters. For this study we note that the values of C_{ref} , K_{max} where chosen to limit computational complexity, increasing them further should not qualitatively change our conclusions. Also regarding r_c , setting it to too low a value will disable NC and setting it too high may destabilize NC. We performed ten independent trials and present in plots median values and 95% confidence intervals. When studying niche-constructing behaviors we do not average across trials but present them individually, as differences would average out and conceal information.

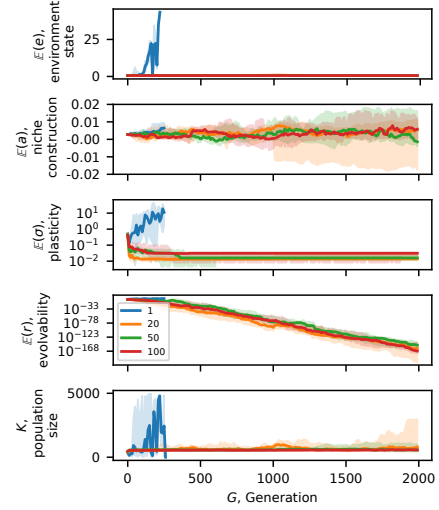


Figure 3: Comparison between a homogeneous ($N = 1$) and a heterogeneous environment ($N = 100$) under global competition

Spatial heterogeneity stabilizes niche construction

In Figure 3 we compare the behavior under global competition for different number of niches ($N \in \{1, 50, 100\}$) for a single trial. We observe that, for $N = 1$, the agents are niche-constructing positively (third row), which pushes the environment state to higher values and more variability (first row). We see that the the population reacts by adapting its preferred niche (second row), increasing its plasticity (fourth row) and keeping its evolvability high (fifth row). Despite that, the population goes extinct around generation 250, after experiencing many booms and busts (sixth row). This behavior was consistent across all trials for $N = 1$, with half of the trials experiencing a mass extinction due to negative NC and half due to positive NC (as in Figure 3).

In the heterogeneous environments, on the other hand, the population survives for the whole simulation. We see that the average NC among agents ($\mathbb{E}(a)$) is similar to that of the homogeneous environment but that this leads to less niche construction in the environment (evident through the environmental state). This is because the population in the heterogeneous environment spreads over multiple niches and often migrates collectively out of a niche (we will take a look at this behavior in the next section), leading to a decrease to the accumulation of NC due to its discounting. In contrast, when $N = 1$ NC accumulates in a single niche. We also observe that NC in $N = 100$ does not increase continuously but experiences oscillations, which suggests that some agents are niche-constructing negatively or reducing their positive NC. As a result, the environment experiences lower oscillations which enables the population to reduce its evolvability. Plasticity remains relatively high but much lower than in the homogeneous environment. We also mea-

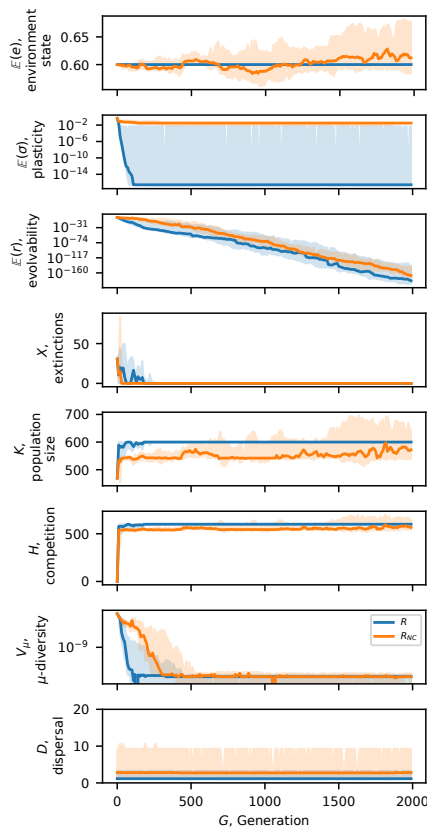


Figure 4: Comparison between populations with (R_{NC}) and without NC (R) under global competition.

sured the number of mass extinctions for environments with different number of niches and observed that even a small number of niches is useful and increasing the number of niches continuously decreases the probability of extinction (10/10 went extinct for $N = 1$, 4/10 for 20, 3/10 for 50 and 2/10 for 100). Thus, the presence of multiple niches was necessary for the survival of the niche-constructing population. Next, we will look into how adaptability interacts with NC to give rise to this behavior. As we hypothesize that the presence of niches matters, we will focus on how agents disperse and niche-construct differently in different niches.

Niche construction promotes adaptability

We now compare a niche-constructing population (denoted as R_{NC}) and a population that cannot niche-construct (R). We examine populations under global selection in Figure 4 and under local selection in Figure 5. We observe that both populations increase their adaptability but do so differently.

Under global selection, the R_{NC} population increases its plasticity but keeps its evolvability low. Dispersal is a bit higher and genomic diversity is initially high but, by the middle of the simulation, reaches the low level of the R population. We also see that R_{NC} has a smaller population size

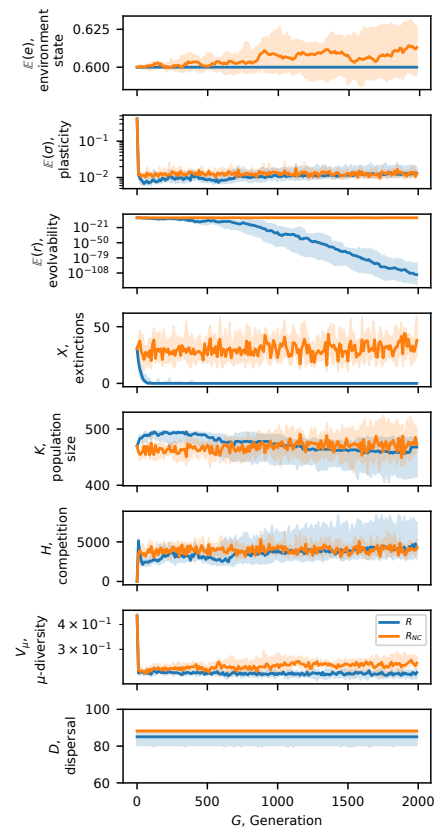


Figure 5: Comparison between populations with (R_{NC}) and without NC (R) under local competition.

early in the simulation. As extinctions are low, this is due to NC reducing the capacity of the environment. Near the end of the simulation we see that R_{NC} reaches and, in some trials, surpasses the size of the R population. This indicates that positive NC has increased the environment's capacity.

Under local selection R_{NC} has very high evolvability but plasticity is at similar level to R . We observe that some extinctions persist throughout evolution but the population size is relatively stable. Why did the niche-constructing population prefer to adapt through evolvability rather than plasticity? As we will see more closely in our analysis of specific trials, agents in a niche-constructing population manage to coordinate their NC behavior within a niche. This would not have been possible under high plasticity, as agents would stochastically niche-construct in different niches and increase environmental uncertainty.

Niche construction can cause an arms race

In Figure 6 we analyze one of the trials for the niche-constructing population under global competition through heatmaps that show the amount of NC at each generation and niche. Populations are initialized with diverse genomes, so that at the beginning they are dispersed in all niches.

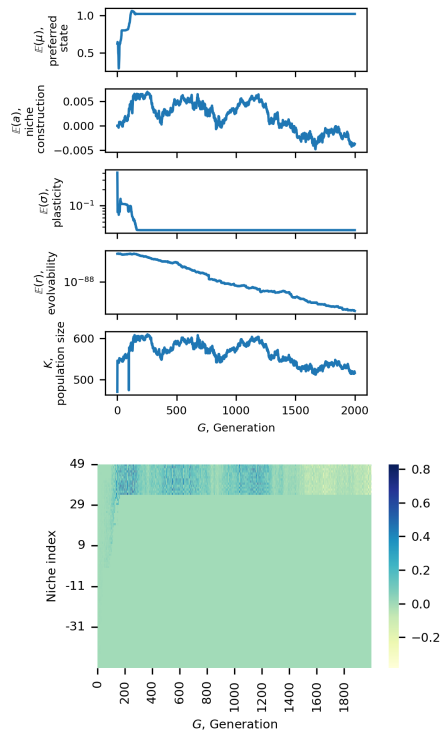


Figure 6: Analyzing one trial for global competition with niche-constructing population where an arms race emerged (Top) evolution of metrics (Bottom) Heatmap with rows corresponding to niches, columns to generations and value/color of cell to the sum of niche construction of all agents in a given niche and generation.

Then, global competition leads to a decrease in plasticity and quickly wipes out the diversity of the population, gathering all agents in the same narrow stripe of niches. Around generation 50, the population starts positively niche-constructing in the middle niche, then starts moving to the north and, once it reaches the end of the environment (around generation 150), stays there and oscillates between phases of positive and negative NC. On the top plot, we see that this population has adapted its preferred state to the last niche and that as soon it reaches it, plasticity stops decreasing and the population size stops increasing. The trial depicted in Figure 6 is representative of 4 out of the 10 trials. In another five trials the same pattern occurred, but in the south instead of the north. In contrast, we observed that populations without NC do not gather at the edge of the environment but stay in a random niche near the center. Thus, NC under global competition leads to a "geographic" arms race due to the following dynamics: once the agents have been gathered to a narrow stripe of niches, they begin niche-constructing, initially randomly. Then, NC will, by random chance become slightly positive or negative. As all agents are competing

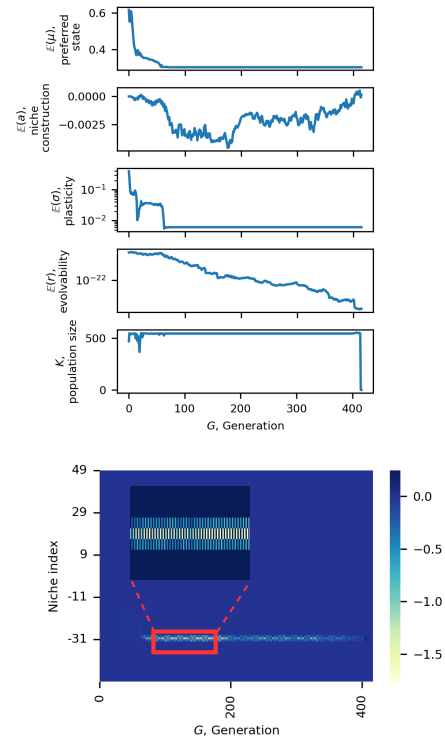


Figure 7: Analyzing one trial for global competition with niche-constructing population where negative niche construction emerged.

in the same niche, this will force them to adapt their preferred state towards the direction that the sum of agents is constructing to. This enables inhabiting the adjacent niche, where the agents will niche-construct towards the same direction. Eventually, the population will reach the end of the environment. Why doesn't the population continue increasing its niche construction to the same direction instead of experiencing cycles? We believe that this is not possible as the mutation rate is low and, therefore, the preferred niche cannot be adapted anymore.

Negative niche construction reduces competition

In Figure 7 we analyze another trial for the niche-constructing population under global competition, where a different behavior emerged: the population niche-constructs negatively, which pushes it to the low-index niches. There we see an interesting pattern: the population is gathered in a narrow stripe of niches where it niche-constructs negatively and with increasing intensity. Then the population leaves this stripe and moves to the adjacent northern and southern niches. This is possible due to the maintained plasticity and does not require an adaptation of the preferred niche. Then, the same behavior happens until the population moves back to the previous stripe and we see many cycles of moving

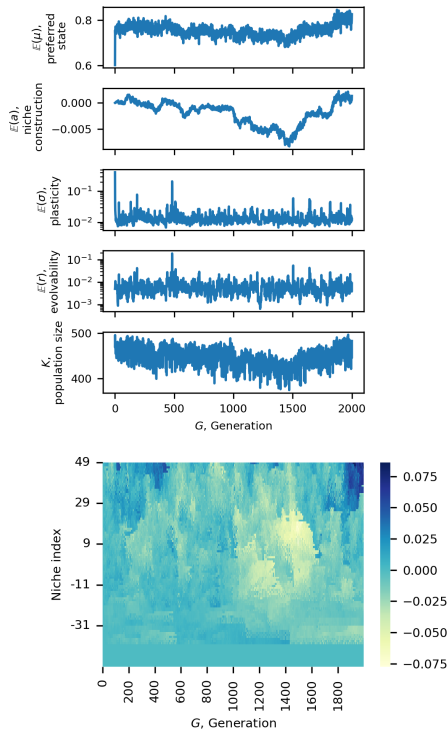


Figure 8: Analyzing a typical trial for local competition with a niche-constructing population.

back and forth. This switching behavior is caused by the fact that, once the population niche constructs negatively in a niche for some time, then this niche becomes uninhabitable, so that the population needs to move out until it becomes inhabitable again. Eventually, some positive NC happens, which leads to an extinction, as evolvability is too low to enable adaptation. This indicates that negative niche construction is not stable in the long-term.

Niche construction diversifies the environment

We now move to local selection, where we saw in Figure 5 that NC leads to a more evolvable population. In Figure 8, we analyze a typical trial where we see that NC is positive and high in northern niches and negative in more southern niches. Thus, the majority lives in the north (this can be inferred from the preferred state $\mathbb{E}(\mu) = 0.8$ surpassing the state of the reference niche $i_0 = 0.6$). The population experiences many changes, with niches often switching between being positively and negatively niche-constructed. We also measured the variance of NC within a niche and found that it is two orders of magnitude smaller than the one under global selection. This suggests that local selection enables agents to coordinate their niche construction within a niche, which diversifies the environment and leads to the higher genomic diversity observed in Figure 5.

An intriguing question is why populations under global selection adapt by increasing their evolvability but keep their plasticity at even lower levels than without NC. We believe that this is due to the interplay between NC and plasticity: very plastic agents could niche-construct in many niches, which would make it difficult to maintain low variance in NC within a niche. By monopolizing a limited number of niches and quickly changing niche through mutations these populations can better coordinate their NC.

Discussion

We have shown that the evolution of niche construction is contingent on the number of niches in the environment and the selection mechanism. We explained this by looking into the interplay between adaptability and NC, in particular their effect on the population’s and environment’s diversity and dispersal patterns. Populations are impressive in their ability to self-regulate themselves: NC remains bounded, even though we do not introduce an explicit cost for it. Populations fail to coordinate their NC and go extinct only when the environment has a single niche.

Our empirical study could be extended in multiple ways. First, we could study the effect of additional hyperparameters. For example, we could allow the intrinsic state to vary with time, following a noisy or periodic signal (Nisioti and Moulin-Frier, 2022). We could also consider that NC is plastic: currently an agent’s gene determines its NC independently of the environmental state while we could imagine scenarios where agents niche-construct differently for different states. Finally, we believe that our empirical observations should be tested in grounded environments with RL agents, such as grid-worlds where a population forages or avoids predation (Hamon et al., 2023; Chiba et al., 2020).

We hope that our work will contribute towards incorporating NC in future studies of collective adaptation. We think this is particularly relevant for the Artificial Intelligence community, which recently focused on agents that can generalize to diverse environments, the importance of environmental complexity and multi-agent dynamics (Reed et al., 2022; Jaderberg et al., 2022; Nisioti et al., 2021; Moulin-Frier, 2022). As we show here, niche construction, which can be seen as a meta-learning mechanism (Constant et al., 2018), can prove promising in complexifying environments and population dynamics and bring us closer to artificial agents with behaviors reminiscent of natural ones.

Acknowledgements This research was partially funded by the French National Research Agency (<https://anr.fr/>, project ECOCURL, Grant ANR-20-CE23-0006). This work also benefited from access to the HPC resources of IDRIS under the allocation 2020-[A0091011996] made by GENCI.

References

- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. (2020). Emergent Tool Use From Multi-Agent Autocurricula. In *International Conference on Learning Representations*.
- Boivin, N. L., Zeder, M. A., Fuller, D. Q., Crowther, A., Larson, G., Erlandson, J. M., Denham, T., and Petraglia, M. D. (2016). Ecological consequences of human niche construction: Examining long-term anthropogenic shaping of global species distributions. *Proceedings of the National Academy of Sciences*, 113(23):6388–6396.
- Cantor, M. and Whitehead, H. (2013). The interplay between social networks and culture: Theoretically and among whales and dolphins. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1618):20120340.
- Chiba, N., Suzuki, R., and Arita, T. (2020). Evolution of Complex Niche-Constructing Behaviors and Ecological Inheritance of Adaptive Structures in a Physically Grounded Environment. *Frontiers in Robotics and AI*, 7:45.
- Constant, A., Ramstead, M. J. D., Veissière, S. P. L., Campbell, J. O., and Friston, K. J. (2018). A variational approach to niche construction. *Journal of The Royal Society Interface*, 15(141):20170685.
- Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, 521(7553):503–507.
- Dere, M. and Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, 113(11):2982–2987.
- d’Errico, F., Banks, W. E., Warren, D. L., Sgubin, G., van Niekerk, K., Henshilwood, C., Daniau, A.-L., and Sánchez Goñi, M. F. (2017). Identifying early modern human ecological niche expansions and associated cultural dynamics in the South African Middle Stone Age. *Proceedings of the National Academy of Sciences*, 114(30):7869–7876.
- Diamond, J. M. (1998). *Guns, Germs, and Steel: the Fates of Human Societies*. W. W. Norton & Co., New York.
- Grove, M. (2014). Evolution and dispersal under climatic instability: A simple evolutionary algorithm. *Adaptive Behavior*, 22(4):235–254.
- Hamon, G., Nisioti, E., and Moulin-Frier, C. (2023). Evolutionary Dynamics of Non-episodic Neuroevolution in Large Multi-agent Environments.
- Jaderberg, M., Mathieu, M., McAleese, N., Bradley-Schmieg, N., Wong, N., Porcel, N., Hughes-Fitt, S., Dalibard, V., and Czarnecki, W. M. (2022). Open-Ended learning leads to Generally Capable Agents. page 54.
- Krakauer, D. C., Page, K. M., and Erwin, D. H. (2009). Diversity, Dilemmas, and Monopolies of Niche Construction. *The American Naturalist*, 173(1):26–40.
- Laland, K. N., Odling-Smee, F. J., and Feldman, M. W. (1999). Evolutionary consequences of niche construction and their implications for ecology. *Proceedings of the National Academy of Sciences*, 96(18):10242–10247.
- Larrasoña, J. C., Roberts, A. P., and Rohling, E. J. (2013). Dynamics of Green Sahara Periods and Their Role in Hominin Evolution. *PLOS ONE*, 8(10):e76514.
- Leimu, R. and Fischer, M. (2008). A Meta-Analysis of Local Adaptation in Plants. *PLoS ONE*, 3(12):e4010.
- Lynch, M., Ackerman, M. S., Gout, J.-F., Long, H., Sung, W., Thomas, W. K., and Foster, P. L. (2016). Genetic drift, selection and the evolution of the mutation rate. *Nature Reviews Genetics*, 17(11):704–714.
- Lynch, M. and Gabriel, W. (1987). Environmental tolerance. *The American Naturalist*, 129(2):283–303.
- Maslin, M. A., Brierley, C. M., Milner, A. M., Shultz, S., Trauth, M. H., and Wilson, K. E. (2014). East African climate pulses and early human evolution. *Quaternary Science Reviews*, 101:1–17.
- Maslin, M. A., Shultz, S., and Trauth, M. H. (2015). A synthesis of the theories and concepts of early human evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1663):20140064.
- Mesoudi, A. and Thornton, A. (2022). What is cumulative cultural evolution? *Proceedings of the Royal Society B: Biological Sciences*, 285(1880):20180712.
- Migliano, A. B., Battiston, F., Viguier, S., Page, A. E., Dyble, M., Schlaepfer, R., Smith, D., Astete, L., Ngales, M., Gomez-Gardenes, J., Latora, V., and Vinicius, L. (2020). Hunter-gatherer multilevel sociality accelerates cumulative cultural evolution. *Science Advances*, 6(9):eaax5913.
- Migliano, A. B. and Vinicius, L. (2022). The origins of human cumulative culture: From the foraging niche to collective intelligence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 377(1843):20200317.
- Moulin-Frier, C. (2022). *The Ecology of Open-Ended Skill Acquisition*. Habilitation à diriger des recherches, Université de Bordeaux (UB).
- Nisioti, E., Jodogne-del Litto, K., and Moulin-Frier, C. (2021). Grounding an Ecological Theory of Artificial Intelligence in Human Evolution. In *NeurIPS 2021 - Conference on Neural Information Processing Systems / Workshop: Ecological Theory of Reinforcement Learning*, virtual event, France.
- Nisioti, E., Mahaut, M., Oudeyer, P.-Y., Momennejad, I., and Moulin-Frier, C. (2022). Social Network Structure Shapes Innovation: Experience-sharing in RL with SAPIENS.
- Nisioti, E. and Moulin-Frier, C. (2022). Plasticity and evolvability under environmental variability: The joint role of fitness-based selection and niche-limited competition. In *GECCO 2022 - The Genetic and Evolutionary Computation Conference*, Boston / Hybrid, United States.
- Odling-Smee, F. J., Laland, K. N., and Feldman, M. W. (2003). *Niche Construction: The Neglected Process in Evolution (MPB-37)*. Princeton University Press.
- Pérolat, J., Leibo, J. Z., Zambaldi, V., Beattie, C., Tuyls, K., and Graepel, T. (2017). A multi-agent reinforcement learning model of common-pool resource appropriation. In Guyon,

- I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Potts, R. (2013). Hominin evolution in settings of strong environmental variability. *Quaternary Science Reviews*, 73:1–13.
- Provine, W. B. (2001). *The Origins of Theoretical Population Genetics: With a New Afterword*. University of Chicago Press, Chicago, IL.
- Pugh, J. K., Soros, L. B., and Stanley, K. O. (2016). Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI*, 3.
- Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S. G., Novikov, A., Barth-Maron, G., Gimenez, M., Sulsky, Y., Kay, J., Springenberg, J. T., Eccles, T., Bruce, J., Razavi, A., Edwards, A., Heess, N., Chen, Y., Hadsell, R., Vinyals, O., Bordbar, M., and de Freitas, N. (2022). A generalist agent.
- Schwilk, D. W. (2003). Flammability Is a Niche Construction Trait: Canopy Architecture Affects Fire Intensity. *The American Naturalist*, 162(6):725–733.
- Suzuki, R. and Arita, T. (2009). Effects of Temporal Locality of Ecological Processes on Coevolution of Learning and Niche Construction.
- Turner, J. S. (2000). The extended organism.
- West, G. (2017). *Scale*.
- Zohari, D. (1986). The Origin and Early Spread of Agriculture in the Old World. In Barigozzi, C., editor, *Developments in Agricultural and Managed Forest Ecology*, volume 16 of *The Origin and Domestication of Cultivated Plants*, pages 3–20. Elsevier.