

# Energy Costs and Neural Complexity Evolution in Changing Environments

Sian Heesom-Green<sup>1</sup>, Jonathan Shock<sup>1,2,3</sup> and Geoff Nitschke<sup>1</sup>

<sup>1</sup>University of Cape Town, South Africa

<sup>2</sup>INRS Montreal, Canada

<sup>3</sup>NiTheCS, Stellenbosch  
South Africa

hsmsia001@myuct.ac.za, jonathan.shock@uct.ac.za, gnitschke@cs.uct.ac.za

## Abstract

The *Cognitive Buffer Hypothesis* (CBH) posits that larger brains evolved to enhance survival in changing conditions. However, larger brains also carry higher energy demands, imposing additional metabolic burdens. Alongside brain size, brain organization plays a key role in cognitive ability and, with suitable architectures, may help mitigate energy challenges. This study evolves *Artificial Neural Networks* (ANNs) used by Reinforcement Learning (RL) agents to investigate how environmental variability and energy costs influence the evolution of neural complexity, defined in terms of ANN size and structure. Results indicate that under energy constraints, increasing seasonality led to smaller ANNs. This challenges CBH and supports the *Expensive Brain Hypothesis* (EBH), as highly seasonal environments reduced net energy intake and thereby constrained brain size. ANN structural complexity primarily emerged as a byproduct of size, where energy costs promoted the evolution of more efficient networks. These results highlight the role of energy constraints in shaping neural complexity, offering *in silico* support for biological theory and energy-efficient robotic design.

## Introduction

The evolution of the brain is a fascinating topic that has been widely investigated, yet much remains to be understood. Many studies have investigated why some animals have evolved larger brains than others (Sayol et al., 2016; Sol, 2009; Michaud et al., 2022). The *Cognitive Buffer Hypothesis* (CBH) suggests that large brains evolved to improve adaptability and enhance survival in changing conditions, such as seasonal environments (Allman et al., 1993; Sol, 2009; Michaud et al., 2022). However, larger brains are typically metabolically costly, and it is not always feasible for organisms to increase their energy intake (Sayol et al., 2016; Michaud et al., 2022; Smaers and Soligo, 2013). The *Expensive Brain Hypothesis* (EBH) highlights this constraint, proposing that an increase in brain size must be met by an increase in net energy intake or reduced energy allocation to other vital organs (Isler and van Schaik, 2009). Despite these limitations, organisms often face the need to adapt to environmental changes (Smaers and Soligo, 2013), raising the question of how changing environments and energy costs impact neural complexity evolution.

Alongside brain size, the organization of the brain plays a crucial role in cognitive ability (Cohen and D’Esposito, 2016). Brain structures that balance functional segregation, where information is processed within specialized neural groups, and integration, where these groups communicate, are key biomarkers for diverse cognitive function (Tononi et al., 1994; Cohen and D’Esposito, 2016). While larger brains demand more energy, which organisms may struggle to obtain, an optimally structured brain could achieve similar cognitive functions with lower energy costs (Smaers and Soligo, 2013). This structural efficiency may be crucial in mitigating the energetic constraints associated with larger brains (Smaers and Soligo, 2013; Oizumi et al., 2014). This raises the question: *Does neural evolution favor larger brains or more efficient wiring under energy constraints?* Hereafter, we use the term neural complexity to refer to the combined influence of brain size and structural organization.

In neuro-evolution, the artificial evolution of *Artificial Neural Networks* (ANNs), provides a valuable tool for studying neural complexity evolution (Miikkulainen, 2025). While both brain size and structural organization contribute to neural complexity, they are often studied independently. For example, previous studies in evolutionary robotics have defined neural complexity in terms of ANN size and investigated whether imposing complexity costs, such as energy penalties, evolves more efficient robot controllers (Nagar et al., 2019; Hallauer and Nitschke, 2020), with less focus on how structural organization in neural controllers might adapt to balance these constraints. Similarly, studies on neural complexity in terms of network structure typically investigate conditions driving evolution towards more complex structures, but with constrained ANN size conditions (Edlund et al., 2011; Joshi et al., 2013; Yaeger and Sporns, 2006; Nitash et al., 2018).

Thus this study investigates neural complexity evolution for ANN size and structure, given testing the CBH remains little investigated using neuro-evolution as an experimental tool. Also, while CBH focuses on how changing environments influence the evolution of brain size, little research has examined its impact on evolving neural structures.




Figure 1: 20x20 grid-world. Cell color indicates food type (edible or poisonous). The agent is the white circle. The legend shows seasonal food-color mappings.

**Research Objective:** Specifically, this study investigates how changing environments and energy costs impact the evolution of neural complexity, defined by both ANN size and structure. To do this, we evolve ANNs of Reinforcement Learning (RL) agents in four environments, ranging from static to highly seasonal (1 to 4 seasons), under two energy regimes (with and without energy costs that scale with ANN size). We thus test predictions of the CBH to examine the relative importance of ANN size versus structural complexity in adaptation, including whether energy costs drive the evolution of more efficient controllers.

## Methods

This section details the task, agent neuro-evolution, energy costs, task performance and neural complexity metrics.

### Task Environment

Agents operate in a 20x20 2D environment (Figure 1), consuming edible foods and avoiding poisonous ones, with edibility signaled by color. Agents expend energy each time step, gain energy from edible foods, and lose energy from poisonous ones (Table 1). At the start of each episode, 10 edible and 10 poisonous food items are randomly placed in the environment, along with the agent’s location. When a resource is consumed, it is replaced with a new item of the same type at a random location. Agents have five discrete actions (move up, down, left, right, or eat) and observe a 9x9 window of nearby cell colors, encoded as 243 RGB values ( $9 \times 9 \times 3$ ), normalized to  $[0, 1]$ . Each episode ends after 100 time steps (Table 1).

**Agent Rewards:** The reward process mirrors the energy dynamics that agents experience over their lifetime.

- -0.01 at each time step (reflects energy expenditure)<sup>1</sup>.
- +1 upon consuming edible food (energy replenishment).

<sup>1</sup>-0.01 is the baseline energy expenditure per step when not scaled by ANN size (see Equation 4).

- -1 upon consuming poisonous food (energy depletion).

**Food Colors and Seasonal Changes:** Environments differ in the number of seasons per episode (1–4), with each season defined by a unique mapping of colors to food edibility (see Figure 1 legend). At the start of each experimental run (i.e., per simulation seed), a predefined set of colors (green, yellow, orange, brown, red, blue, purple, pink) is shuffled: the first four are assigned as edible (for seasons 1–4), and the rest as poisonous. This mapping remains fixed across all episodes and generations within a run but varies between simulation seeds.

Each 100-step episode is evenly divided across seasons, with agents experiencing the same fixed seasonal sequence. For example, a 2-season environment consists of 50 steps of season 1 followed by 50 of season 2. Agents receive no explicit cues about season transitions and must infer them through experience, but these changes follow a fixed sequence in every episode of a run and are therefore predictable rather than stochastic. Additionally, each food item’s final color is generated by sampling each RGB channel within a 0.2 range of its base color.

### Lifetime Learning via Reinforcement Learning

Agents undergo lifetime learning per generation using RL, where during their lifetime, they adapt by adjusting ANN controller connection weights given environment interactions. This lifetime learning mirrors natural adaptive behavior, potentially enhancing agent ability to adapt to environmental changes (Doncieux et al., 2015; Urzelai and Floreano, 2001). RL is guided by the *Proximal Policy Optimization* (PPO) algorithm (Schulman et al., 2017). Each RL agent has an actor and a critic network that share the same ANN topology, which evolves via NEAT (Stanley and Miikkulainen, 2002), but differ in their output layers. The actor network learns the optimal policy, while the critic network estimates the value function.

- **Input layer** (243 nodes): Encodes the current state as a feature vector, representing the RGB color values of each cell in the agent’s 9x9 field of view.
- **Critic output layer** (1 node): State-value function  $V(s)$ , estimating the expected cumulative reward from state  $s$ .
- **Actor output layer** (5 nodes): Represents the policy  $\pi(a|s)$ , defining the probability distribution over the agent’s five possible actions in state  $s$ .

The RL hyper-parameters<sup>2</sup> and the actor’s ANN evolve via NEAT. The critic network is derived by replacing the actor’s output layer with a single state-value node  $V(s)$ , which inherits all incoming output node connections.

<sup>2</sup>RL hyper-parameter details here.

Table 1: Simulation and evolution parameters

Experiment Parameters	
Number of runs	20
Episode length (time-steps)	100
Agent energy usage (per time-step)	-E
Energy gain (edible food)	+1
Energy loss (poisonous food)	-1
Simulation Parameters	
Grid-world size	20x20
Observation / action space size	243 / 5
Neuro-evolution Parameters	
Population size	150
Generations	400
ANN architecture	Feedforward
ANN initial connectivity	Fully connected
Hidden nodes	1
Mutation Parameters	
Weight mutation rate	0.1
Node add / delete probability	0.5 / 0.0
Connection add / delete probability	0.5 / 0.5

### Agent Neuro-evolution

The neuro-evolution process uses the NEAT algorithm to evolve ANN topologies and weights (Stanley and Miikkulainen, 2002), outlined in the following:

1. **Initialization:** 150 RL agents (individuals) are generated with random RL hyper-parameters and a fully connected (initial structural complexity  $N_C = 0$ ; Equation 2) feed-forward ANN with one hidden node. This setup reduces initial network size (fewer connections than no hidden layer). Weights and biases are initialized from  $\mathcal{N}(0, 1)$ .
2. **Lifetime Learning:** Before evaluation, RL agents update their inherited ANN weights over 1000 episodes (100 000 time steps), using their inherited hyper-parameters and the PPO algorithm (Schulman et al., 2017). Importantly, any updates to ANN weights during this learning phase are not inherited by the next generation (Baldwinian Evolution).
3. **Evaluation:** After the learning phase, assess each individual’s fitness based on RL performance (total reward per episode), averaged over 100 episodes (using seeds not seen during learning).
4. **Selection, Reproduction, and Replacement:** The top 10% (15 individuals) are selected based on fitness to produce offspring via crossover and mutation, modifying inherited architectures, initial weights (pre-learning), and RL hyper-parameters. The fittest two are preserved. The remaining 148 are replaced by offspring.

Steps 2-4 are repeated for 400 generations, with each generation using the same fixed training and evaluation environment seed split within a given simulation run. Further details

on the evolutionary parameters<sup>3</sup> used in this neuro-evolution process are provided in Table 1. Evolutionary parameters were optimized via Bayesian hyper-parameter optimization, and connection and node addition and deletion rates were manually tuned over values  $\{0.1-0.5\}$  to balance task performance and complexity variance. Node deletion rate was set to 0, since disconnected nodes could still be removed.

### Task Performance Metric: Net Energy Intake

The task performance metric differs from the fitness function (accumulated rewards over an episode) because it excludes energy expenditure (E), and primarily focuses on net energy intake, measured as the number of edible foods consumed minus the number of poisonous foods consumed during the agent’s lifetime. Excluding energy costs from the task performance calculation enables a better comparison between agents with different energy expenditures, that is, those with energy costs on ANN size versus those without. This still influences the neuro-evolution process and enables us to investigate the impact of higher energy costs for larger ANNs on the evolution of neural complexity.

**Rationale for Energy Budget Design** Agents were not assigned an initial energy budget because energy dynamics did not affect their lifespan or trial duration. This setup was intentionally chosen to avoid biases when comparing performance between the different energy expenditure conditions. In particular, it avoids the issue noted by Hallauer and Nitschke (2020), where energy depletion shortened agent lifespans, giving agents without energy costs more time to perform the task. Since agents in this study could act for the full duration regardless of energy level, assigning an initial energy value was unnecessary.

### Neural Complexity Metrics

**ANN size ( $N_S$ ):** Larger ANNs are often associated with the potential for more complex behaviors due to their increased number of free parameters (Lehman and Stanley, 2011). Therefore, ANN size is commonly defined as the total number of free parameters (connections and non-input nodes) in the network (Nagar et al., 2019; Hallauer and Nitschke, 2020; Lehman and Stanley, 2011; Nitschke and Didi, 2017; Yu, 2010), and is calculated using Equation 1:

$$N_S = (\# \text{ connections}) + (\# \text{ non-input nodes}) \quad (1)$$

**ANN structural complexity ( $N_C$ ):** In neuroscience, two core principles of brain functional organization are *segregation* (specialized processing within groups of neurons) and *integration* (efficient information exchange between these groups) (Sporns, 2013; Cohen and D’Esposito, 2016). A

<sup>3</sup>Full NEAT evolutionary parameter details here.

Table 2: Task Environments (a) and Experiment Sets (b) (*NEC* = *No Energy Costs*; *EC* = *Energy Cost*)

(a) Task Environments		(b) Experiment Sets			
Environment	Seasons	Experiment Set	Energy Expenditure (E)	Experimental Variables	Evaluation Metrics
1 (Static)	1	1 (NEC)	-0.01	Environments: 1-4	Neural Complexity Task Performance
2	1,2				
3	1,2,3	2 (EC)	$-0.01 \times \frac{N_S^{\text{current}}}{N_S^{\text{gen0}}}$	Environments: 1-4	Neural Complexity Task Performance
4	1,2,3,4				

balance between segregation and integration has been shown to support diverse cognitive abilities and rich, flexible dynamics (Tononi et al., 1994; Deco et al., 2015). Neural complexity ( $N_C$ ) is considered high when this balance is achieved, and low when networks are either fully segregated or fully integrated (Tononi et al., 1994). Tononi et al. (1994)’s original  $N_C$  measure, a precursor to Integrated Information Theory (IIT) measures, becomes computationally infeasible for networks larger than 20 nodes, as it requires evaluating every possible bipartition (Toker and Sommer, 2019; Mediano et al., 2018). To address this, we define the  $N_C$  metric for evolved ANN controllers as the ratio between modularity (segregation) and global efficiency (integration), both well-established, scalable graph-theoretic measures for functional segregation and integration (Toker and Sommer, 2019; Palma-Espinosa et al., 2025; Van Diessen et al., 2014; Cohen and D’Esposito, 2016; Capouskova et al., 2022).

**Modularity (M):** Measures segregation (range: [0, 1]) (Cohen and D’Esposito, 2016), by comparing within-module connection density to between-module connections (Newman, 2010). Higher values indicated greater segregation. Modules were identified using the Louvain community detection algorithm (Blondel et al., 2008).

**Global Efficiency (E):** Measures integration (range: [0, 1]) (Capouskova et al., 2022), computed as the average inverse shortest path length between all pairs of nodes (Latora and Marchiori, 2001). Higher values indicate greater integration through more efficient global communication across the ANN.

The neural complexity ratio is computed with Equation 2:

$$N_C = \frac{\min(M(G), E(G))}{\max(M(G), E(G))} \quad (2)$$

Where G is the graph representation of the ANN, analyzed using *NetworkX* (Hagberg et al., 2008). This ratio is sensitive to changes in either component, decreasing when one measure dominates the other. When  $N_C \approx 1$ , segregation and integration are well-balanced (high neural complexity). When  $N_C \approx 0$ , one property dominates, indicating low neural complexity.

## Energy Costs

We consider two energy expenditure scenarios: one without ANN-size-dependent energy costs (NEC) and one where energy costs scale with ANN size (EC).

**NEC (No Energy Costs):** Agents have a constant energy expenditure per time step (Equation 3):

$$E = -0.01 \quad (3)$$

Over a 100 time-step trial, total energy loss is  $-1$ , meaning edible food consumption ( $+1$ ) replenishes energy, while poisonous food consumption ( $-1$ ) depletes it. NEC serves as a baseline to evaluate agent performance without additional energy constraints.

**EC (Energy Costs):** Energy expenditure scales with ANN size (Equation 4):

$$E = -0.01 \times C \quad (4)$$

Where,  $C$  is the ratio of the current ANN size ( $N_S$ ) to its initial size at generation 0 ( $N_S = 254$ ), and where there is a selective pressure imposed against larger networks.

## Experiments

Experiments<sup>4</sup> investigate the impact of energy costs and changing environments on neural complexity evolution.

**The Environment Set:** Consists of four increasingly dynamic environments that differ in the number of seasons an agent experiences over its lifetime (Table 2).

**Experiment sets:** Two experiment sets (Table 2) reflecting different energy scenarios, NEC (No Energy Costs, Equation 3) and EC (Energy Costs, Equation 4), are used to investigate the impact of changing environments and energy costs on neural complexity evolution, where both evaluate neural complexity and task performance metrics of the fittest agent evolved per environment, averaged over 20 runs.

<sup>4</sup><https://github.com/sianmay/ECNCECE>

Comparisons *within* each experiment set examine how changing environments influence neural complexity evolution. This shows if environmental variation alone drives changes in network size and structure, or only when coupled with energy constraints. Comparisons *between* the two experiment sets evaluate how energy costs on ANN size impact neural complexity and task performance versus no energy costs, highlighting trade-offs and their overall effect on agent task performance across environments.

## Results & Discussion

Figure 2 presents neural complexity and task performance results (averaged over 20 runs) for the fittest individuals evolved across environments with increasing seasonality, under both energy scenarios (NEC, EC). To evaluate differences in evolved complexity across environments, Kruskal–Wallis (Kruskal and Wallis, 1952) tests ( $p < 0.05$ ) were conducted and Dunn’s post-hoc tests (Dunn, 1964) with Bonferroni correction for pairwise comparisons (Table 3). Spearman rank correlations (Spearman, 1961) were used to examine associations between seasonality and neural complexity, as well as evolutionary trends over generations (Table 5). Finally, Mann–Whitney U tests (Mann and Whitney, 1947) compared evolved neural complexity and task performance between the NEC and EC scenarios (Table 6).

### No Energy Cost (NEC) Scenario

When we investigated whether changing environments alone impact neural complexity evolution, without considering ANN-size-dependent energy costs (NEC scenario), we found no significant differences in evolved ANN size ( $N_S$ ) or structural complexity ( $N_C$ ) across environments. That is, from the static 1-season environment to the 4-season environment (Kruskal-Wallis,  $p > 0.05$ ). Across all conditions, ANNs increased in both  $N_S$  and  $N_C$  over evolutionary time (Figures 2 (b, d), Table 5). The absence of energy constraints likely allowed unrestricted ANN growth, leading to larger and more complex ANNs regardless of environmental dynamics. These results suggest that environmental dynamics alone do not directly drive the evolution of neural complexity, but instead interact with other selective pressures, such as energy constraints. This result is supported by experimental analysis in evolutionary biology (Fristoe et al., 2017), elucidating that evolutionary transitions in brain size (for example, in an avian global phylogeny case study), resulted in larger brains evolving with equal likelihood in stable and changing environments. This suggests that other environmental and evolutionary factors similarly impact the shaping of neural complexity in both simulation and nature.

### Energy Cost (EC) Scenario

The EC scenario further explores the interplay between changing environments, energy constraints and neural complexity evolution by imposing energy costs on larger ANNs.


### Impact of Changing Environments on ANN size ( $N_S$ ):

When higher energy costs were introduced for larger ANNs, significant differences in evolved ANN sizes ( $N_S$ ) across environments were observed (Kruskal-Wallis,  $p < 0.05$ ,  $\eta^2 = 0.194$ ). Specifically, a notable decrease in  $N_S$  was found in the 4-season environment compared to both the 1-season and 2-season environments (Table 3). Spearman correlation tests (Table 5) revealed that  $N_S$  modestly increased over generations for the environments with 1 to 3 seasons, but this trend disappeared in the 4-season environment. Figure 2b, indicates that unlike other environments, the mean  $N_S$  of the 4-season environment initially grew but then decreased after approximately 150 generations.

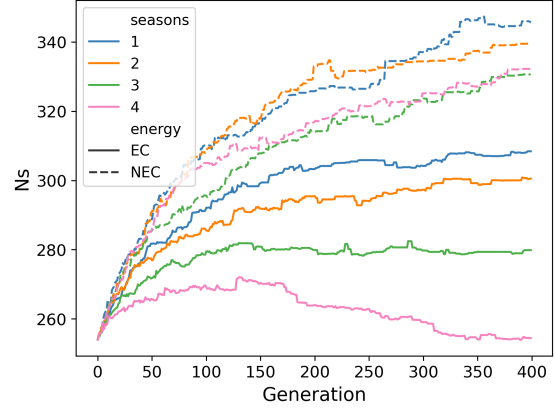
We also confirmed a significant negative association between the number of seasons and  $N_S$  ( $\rho = -0.4322$ ,  $p < 0.05$ , Spearman rank correlation coefficient), indicating that increased environmental change is associated with the evolution of smaller networks, refuting the CBH. One explanation is the impact of increasing seasonal changes on task performance ( $\rho = -0.5389$ ,  $p < 0.05$ , Spearman rank correlation coefficient), suggesting that increased environmental dynamics with concomitantly increased task difficulty, leads to reduced consumption of edible food and consequently lower net energy intake. In the current experiment setup, where agents had no competing energetic costs, the EBH implies that sustaining larger ANNs requires higher net energy intake (higher task performance). Thus, environments enabling higher agent task performance, such as static environments, are more likely to support the evolution of larger ANNs, which is in line with the EBH.

Furthermore, a bootstrap mediation analysis (Preacher and Hayes, 2004) was conducted to test whether task performance (reflecting energy intake) mediates the relationship between environmental variability (number of seasons) and  $N_S$ . The analysis indicated a significant negative indirect effect of seasonal variability on  $N_S$  (Table 4). Specifically, as the number of seasons increased, task performance decreased (e.g., -1.46 for 2 seasons, -2.02 for 3 seasons, and -3.56 for 4 seasons), and higher task performance was associated with larger  $N_S$  (coefficient of 14.50). The indirect effect of -102.09 (95% confidence interval of [-151.91, -60.63]) indicates that decreasing task performance due to environmental variability leads to smaller  $N_S$ . This supports the EBH given that environmental change limits ANN size indirectly by reducing performance, thereby limiting energy intake and the feasibility of evolving larger ANNs.

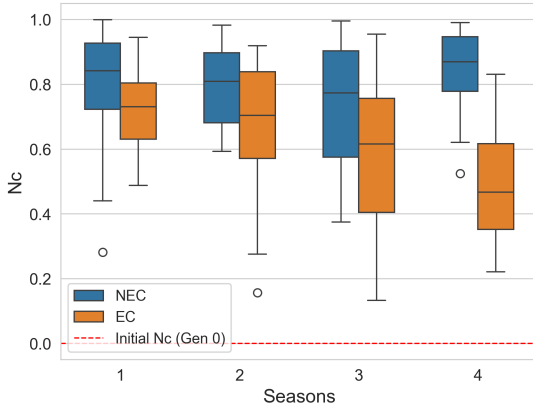
These results align with empirical biological systems research supporting the EBH (Isler and van Schaik, 2009), including studies on amphibians (Luo et al., 2017), eutherians (Van Woerden et al., 2010; Graber, 2017; Van Woerden et al., 2012) and marsupials (Weisbecker et al., 2015), which similarly show that increased environmental seasonality is often associated with smaller brain sizes due to energy limitations.




(a) ANN size ( $N_S$ , Equation 1)




(b) ANN size ( $N_S$ ) over Generations




(c) ANN structural complexity ( $N_C$ , Equation 2)



(d) ANN structural complexity ( $N_C$ ) over Generations



(e) Task performance (net energy intake)



(f) Task performance (net energy intake) over Generations

Figure 2: **Left:** Box plots for neural complexity metrics and task performance. The x-axis represents the environments, labeled by the number of seasons, from the static 1-season environment to the most dynamic 4-season environment. The y-axis displays neural complexity metrics / task performance for the fittest evolved ANNs from the final generation (fitness evaluated after lifetime learning using the same evaluation seeds as during evolution), averaged over 20 runs. **Right:** Neural complexity metrics and task performance over evolutionary time of the current fittest genome, averaged over 20 runs. (*NEC* = *No Energy Costs*; *EC* = *Energy Costs*)

Table 3: Experiment set 2 (EC): Post-hoc Dunn test results (Bonferroni-corrected) for pairwise comparisons of ANN size ( $N_S$ ) and structural complexity ( $N_C$ ) across environments with increasing seasonal changes. Significant p-values ( $p < 0.05$ ) in bold.

Seasons	$N_S$			$N_C$		
	2 seasons	3 seasons	4 seasons	2 seasons	3 seasons	4 seasons
1 season	1.0000	0.1875	<b>0.0028</b>	1.0000	0.3051	<b>0.0040</b>
2 seasons		0.6017	<b>0.0167</b>		1.0000	<b>0.0468</b>
3 seasons			1.0000			0.8836

Table 4: Mediation analysis testing whether task performance mediates the relationship between environmental variability (number of seasons) and ANN size ( $N_S$ ). The indirect effect (ab) is significant, with a 95% confidence interval (CI) of [-151.91, -60.63], indicating that task performance mediates the relationship between environmental variability and  $N_S$ .

a: Effect of n_seasons on Task Performance (2, 3, 4 seasons)	b: Effect of Task Performance on $N_S$	ab: Indirect Effect	95% CI for ab
-1.4585 (2), -2.0205 (3), -3.5610 (4)	14.5020	-102.0939	[-151.91, -60.63]

Table 5: Spearman rank correlation coefficients reflecting the strength and direction of monotonic trends between generation number and neural complexity metrics ( $N_S$  and  $N_C$ ) for each environment and energy scenario. Bold values indicate statistically significant correlations ( $p < 0.05$ ).

Seasons	No Energy Cost (NEC)		Energy Cost (EC)	
	Gen- $N_S$	Gen- $N_C$	Gen- $N_S$	Gen- $N_C$
1	<b>0.5551</b>	<b>0.6127</b>	<b>0.4627</b>	<b>0.6603</b>
2	<b>0.5414</b>	<b>0.6454</b>	<b>0.3288</b>	<b>0.5319</b>
3	<b>0.4531</b>	<b>0.5343</b>	<b>0.1308</b>	<b>0.4675</b>
4	<b>0.6036</b>	<b>0.7521</b>	0.0034	<b>0.4449</b>

However, ascertaining the exact correlation between ANN size and environmental change, in support of the CBH, remains the topic of ongoing research.

**Impact of Changing Environments on ANN structural complexity ( $N_C$ ):** A significant difference was found in evolved ANN structural complexity ( $N_C$ ) across environments (Kruskal-Wallis,  $p < 0.05$ ,  $\eta^2 = 0.168$ ), with the 4-season environment resulting in lower  $N_C$  than both the 1-season and 2-season environments (Table 3). We also confirmed a significant negative association between the number of seasons and  $N_C$  ( $\rho = -0.4062$ ,  $p < 0.05$ , Spearman rank correlation), indicating that increased environmental change is also generally associated with the evolution of less complexly structured networks. However, while ANN size did not significantly increase over generations in the 4-season environment (Figures 2 (b, d), Table 5),  $N_C$  still showed a moderate upward trend ( $\rho = 0.44$ ,  $p < 0.05$ , Spearman rank correlation), suggesting that structural complexity continued to increase concomitant with fitness (Joshi et al., 2013; Edlund et al., 2011; Albantakis et al., 2014).

### Impact of Energy Costs on Neural Complexity and Task Performance

Mann-Whitney U tests ( $p < 0.05$ ) tested for significant differences in task performance and neural complexities ( $N_S$  and  $N_C$ ) between networks evolved with and without energy costs imposed on ANN size (NEC vs. EC scenarios), across each environment in the set. These tests (see Table 6) revealed that, across all environments, imposing energy costs on ANN size led to the evolution of smaller ANNs with comparable task performance.

Structural complexity ( $N_C$ ) was also significantly lower in the EC scenario compared to the NEC scenario, suggesting a potential dependency between  $N_S$  and  $N_C$ . Overall, these results highlight the value of incorporating energy constraints to promote the evolution of more efficient neural architectures. Ongoing work is investigating the effects of varying energy costs to identify thresholds where energy limitations begin to significantly impact task performance.

### Evaluating the Dependency Between $N_S$ and $N_C$

To examine whether structural complexity ( $N_C$ ) was simply a byproduct of network size ( $N_S$ ), 20 additional evolutionary runs were conducted using the same parameters as the previous experiments, but with fitness values assigned randomly from a uniform distribution  $U(0, 1)$ , thereby removing task-based selection pressure. This approach evaluated whether randomly evolved ANNs, under the same evolutionary dynamics, exhibited similar  $N_C$  values for a given  $N_S$  (as those observed in the task-driven experiments).

Given  $N_S$  decreased over generations in some task-driven evolutionary runs, a pattern not observed in the random-fitness runs, the additional random-walk evolutionary runs included an explicit penalty on ANN size to encourage smaller networks and enable a more direct comparison. The resulting scatter plot (Figure 3) presents  $N_S$  versus  $N_C$  across all generations, with random-fitness runs in grey and

Table 6: Statistical comparisons (Mann–Whitney U test) for task performance and neural complexity between networks evolved with no energy costs (NEC) and with energy costs (EC) across environments with increasing seasonal dynamics (1 to 4 seasons). Comparisons are shown for overall task performance, network size ( $N_S$ ), and structural complexity ( $N_C$ ). '==' indicates no significant difference, and '>' indicates significantly greater than (given,  $p < 0.05$ ).

	Task Performance	$N_S$	$N_C$
<b>1 season</b>	NEC == EC (p=0.27)	NEC > EC (p=0.01)	NEC > EC (p=0.03)
<b>2 seasons</b>	NEC == EC (p=0.54)	NEC > EC (p=0.01)	NEC > EC (p=0.04)
<b>3 seasons</b>	NEC == EC (p=0.30)	NEC > EC (p=0.00)	NEC == EC (p=0.06)
<b>4 seasons</b>	NEC == EC (p=0.56)	NEC > EC (p=0.00)	NEC > EC (p=0.00)

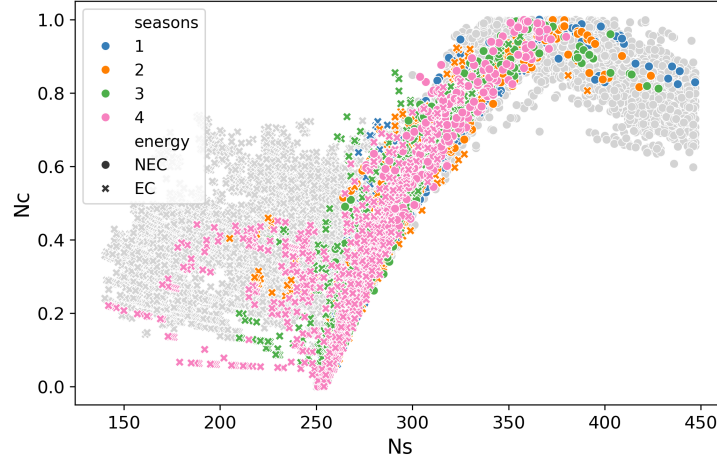


Figure 3: Network size ( $N_S$ ) versus structural complexity ( $N_C$ ) over generations. Grey points and crosses are random-walk evolution without and with an ANN size penalty, respectively. Colored markers show task-based evolution (1-4 season environments). Overlap across conditions suggests  $N_C$  primarily reflects evolved  $N_S$  given mutation rather than task driven selection.

task-driven runs in color. The colored points predominantly fall within the same bounds as the grey points, suggesting that  $N_C$  may have emerged primarily as a consequence of evolving  $N_S$  under the given mutation settings, rather than from task-driven pressures. Structural complexity, therefore, cannot be isolated as a key driver of performance in this context. This is likely a result of the limitations of the feedforward neural controllers (Nolfi and Floreano, 2000) used in this study’s experiments. Ongoing work is testing recurrent connections (memory) in ANNs, to elucidate whether structural complexity plays a critical role in facilitating adaptive and energy efficient solutions across dynamic environments.

## Conclusions

This study investigated how energy costs and changing environments influence the evolution of neural complexity in RL agent ANN controllers. Results indicated that changing environments only impacted neural complexity evolution when energy costs were imposed, with more seasonal environments driving the evolution of smaller networks. Results support the *Expensive Brain Hypothesis* (EBH) over the *Cognitive Buffer Hypothesis* (CBH), within the context of this foraging task, providing *in silico* evidence that or-

ganisms in fluctuating environments may evolve smaller, more energy-efficient brains. Structural complexity increased with fitness, but is hypothesized to have emerged as a byproduct of evolution given current mutation settings. Moreover, imposing energy costs encouraged the evolution of more efficient ANNs, with implications for assisting the design of energy constrained robotic controllers (Nagar et al., 2019), such as those that must adapt to changing robot morphologies and environments (Watson and Nitschke, 2015; Mailer et al., 2021). While these experimental environments allow controlled testing of energy costs and environmental variability, they are highly simplified, so generalization to neural evolution should be made with caution. Overall, this study’s key contribution was its demonstration of the role of energy costs in shaping neural complexity.

## Acknowledgements

Computations were performed using facilities provided by the University of Cape Town’s ICTS High Performance Computing team: `hpc.uct.ac.za`



## References

- Albantakis, L., Hintze, A., Koch, C., Adami, C., and Tononi, G. (2014). Evolution of integrated causal structures in animats exposed to environments of increasing complexity. *PLoS computational biology*, 10(12):e1003966.
- Allman, J., McLaughlin, T., and Hakeem, A. (1993). Brain weight and life-span in primate species. *Proceedings of the National Academy of Sciences*, 90(1):118–122.
- Blondel, V., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.
- Capouskova, K., Zamora-López, G., Kringelbach, M., and Deco, G. (2022). Integration and segregation in the brain as a cognitive flexibility during tasks and rest. *bioRxiv*, pages 2022–10.
- Cohen, J. and D’Esposito, M. (2016). The segregation and integration of distinct brain networks and their relationship to cognition. *Journal of Neuroscience*, 36(48):12083–12094.
- Deco, G., Tononi, G., Boly, M., and Kringelbach, M. (2015). Re-thinking segregation and integration: contributions of whole-brain modelling. *Nature reviews neuroscience*, 16(7):430–439.
- Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. E. (2015). Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2:4.
- Dunn, O. (1964). Multiple comparisons using rank sums. *Technometrics*, 6(3):241–252.
- Edlund, J., Chaumont, N., Hintze, A., Koch, C., Tononi, G., and Adami, C. (2011). Integrated information increases with fitness in the evolution of animats. *PLoS computational biology*, 7(10):e1002236.
- Fristoe, T., Iwaniuk, A., and Botero, C. (2017). Big brains stabilize populations and facilitate colonization of variable habitats in birds. *Nature Ecology & Evolution*, 1(1):1706–1715.
- Graber, M. (2017). *Social and ecological aspects of brain size evolution: a comparative approach*. PhD thesis, University of Zurich.
- Hagberg, A., Swart, P., and Schult, D. (2008). Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory (LANL), Los Alamos, NM (United States).
- Hallauer, S. and Nitschke, G. (2020). Energy and complexity in evolving collective robot bodies and brains. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1–8, Glasgow, United Kingdom. IEEE Press.
- Isler, K. and van Schaik, C. (2009). The expensive brain: a framework for explaining evolutionary changes in brain size. *Journal of human evolution*, 57(4):392–400.
- Joshi, N., Tononi, G., and Koch, C. (2013). The minimal complexity of adapting agents increases with fitness. *PLoS computational biology*, 9(7):e1003111.
- Kruskal, W. and Wallis, A. (1952). Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621.
- Latora, V. and Marchiori, M. (2001). Efficient behavior of small-world networks. *Physical review letters*, 87(19):198701.
- Lehman, J. and Stanley, K. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223.
- Luo, Y., Zhong, M., Huang, Y., Li, F., Liao, W., and Kotrschal, A. (2017). Seasonality and brain size are negatively associated in frogs: evidence for the expensive brain framework. *Scientific reports*, 7(1):16629.
- Mailer, C., Nitschke, G., and Raw, L. (2021). Evolving Gaits for Damage Control in a Hexapod Robot. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 146–153, Lille, France. ACM.
- Mann, H. and Whitney, D. (1947). On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60.
- Mediano, P., Seth, A., and Barrett, A. (2018). Measuring integrated information: Comparison of candidate measures in theory and simulation. *Entropy*, 21(1):17.
- Michaud, M., Toussaint, S., and Gilissen, E. (2022). The impact of environmental factors on the evolution of brain size in carnivorans. *Communications Biology*, 5(1):998.
- Miikkulainen, R. (2025). Neuroevolution insights into biological neural computation. *Science*, 387(6735):eadp7478.
- Nagar, D., Furman, A., and Nitschke, G. (2019). The cost of big brains in groups. In *Proceedings of the 2019 Conference on Artificial Life*, pages 404–411, Newcastle, United Kingdom. MIT Press.
- Newman, M. (2010). *Networks: An introduction*. Oxford University Press, Oxford, UK.
- Nitash, C., Lundrigan, B., Smale, L., and Hintze, A. (2018). The effect of periodic changes in the fitness landscape on brain structure and function. In *Proceedings of the ALIFE 2018: The 2018 Conference on Artificial Life*, pages 469–476, Tokyo, Japan. MIT Press.
- Nitschke, G. and Didi, S. (2017). Evolutionary policy transfer and search methods for boosting behavior quality: Robocup keep-away case study. *Frontiers in Robotics and AI*, 4:62.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press, Cambridge, USA.
- Oizumi, M., Albantakis, L., and Tononi, G. (2014). From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. *PLoS computational biology*, 10(5):e1003588.
- Palma-Espinosa, J., Orellana-Villota, S., Coronel-Oliveros, C., Maidana, J.-P., and Orio, P. (2025). The balance between integration and segregation drives network dynamics maximizing multistability and metastability: J. palma-espinosa et al. *Scientific Reports*, 15(1):18811.

- Preacher, K. and Hayes, A. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior research methods, instruments, & computers*, 36:717–731.
- Sayol, F., Maspons, J., Lapiedra, O., Iwaniuk, A., Székely, T., and Sol, D. (2016). Environmental variation and the evolution of large brains in birds. *Nature communications*, 7(1):13971.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Smaers, J. and Soligo, C. (2013). Brain reorganization, not relative brain size, primarily characterizes anthropoid brain evolution. *Proceedings of the Royal Society B: Biological Sciences*, 280(1759):20130269.
- Sol, D. (2009). Revisiting the cognitive buffer hypothesis for the evolution of large brains. *Biology letters*, 5(1):130–133.
- Spearman, C. (1961). The proof and measurement of association between two things.
- Sporns, O. (2013). Network attributes for segregation and integration in the human brain. *Current opinion in neurobiology*, 23(2):162–171.
- Stanley, K. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127.
- Toker, D. and Sommer, F. (2019). Information integration in large brain networks. *PLoS computational biology*, 15(2):e1006807.
- Tononi, G., Sporns, O., and Edelman, G. (1994). A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences*, 91(11):5033–5037.
- Urzelai, J. and Floreano, D. (2001). Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evolutionary computation*, 9(4):495–524.
- Van Diessen, E., Zweiphenning, W., Jansen, F., Stam, C., Braun, K., and Otte, W. (2014). Brain network organization in focal epilepsy: a systematic review and meta-analysis. *PloS one*, 9(12):e114606.
- Van Woerden, J., Van Schaik, C., and Isler, K. (2010). Effects of seasonality on brain size evolution: evidence from strepsirrhine primates. *The American Naturalist*, 176(6):758–767.
- Van Woerden, J., Willems, E., van Schaik, C., and Isler, K. (2012). Large brains buffer energetic effects of seasonal habitats in catarrhine primates. *Evolution*, 66(1):191–199.
- Watson, J. and Nitschke, G. (2015). Evolving Robust Robot Team Morphologies for Collective Construction. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*, pages 1039–1046, Cape Town, South Africa. IEEE Press.
- Weisbecker, V., Blomberg, S., Goldizen, A., Brown, M., and Fisher, D. (2015). The evolution of relative brain size in marsupials is energetically constrained but not driven by behavioral complexity. *Brain, behavior and evolution*, 85(2):125–135.
- Yaeger, L. and Sporns, O. (2006). Evolution of neural structure and complexity in a computational ecology. In *artificial life X: proceedings of the tenth international conference on the simulation and synthesis of living systems*, pages 330–336. MIT Press (Bradford Books) Cambridge.
- Yu, H. (2010). Network complexity analysis of multilayer feedforward artificial neural networks. *Applications of Neural Networks in High Assurance Systems*, pages 41–55.