

The dynamical regime and its importance for evolvability, task performance and generalization

Jan Prosi^{1,2*}, Sina Khajehabdollahi^{1*}, Emmanouil Giannakakis^{1,2}, Georg Martius³ and Anna Levina^{1,2,4}

¹ Department of Computer Science, University of Tübingen, Tübingen, Germany

² Max Planck Institute for Biological Cybernetics, Tübingen, Germany

³ Max Planck Institute for Intelligent Systems, Tübingen, Germany

⁴ Bernstein Center for Computational Neuroscience Tübingen, Tübingen, Germany

* These authors contributed equally to this work

Abstract

It has long been hypothesized that operating close to the critical state is beneficial for natural and artificial systems. We test this hypothesis by evolving foraging agents controlled by neural networks that can change the system's dynamical regime throughout evolution. Surprisingly, we find that all populations, regardless of their initial regime, evolve to be subcritical in simple tasks and even strongly subcritical populations can reach comparable performance. We hypothesize that the moderately subcritical regime combines the benefits of generalizability and adaptability brought by closeness to criticality with the stability of the dynamics characteristic for subcritical systems. By a resilience analysis, we find that initially critical agents maintain their fitness level even under environmental changes and degrade slowly with increasing perturbation strength. On the other hand, subcritical agents originally evolved to the same fitness, were often rendered utterly inadequate and degraded faster. We conclude that although the subcritical regime is preferable for a simple task, the optimal deviation from criticality depends on the task difficulty: for harder tasks, agents evolve closer to criticality. Furthermore, subcritical populations cannot find the path to decrease their distance to criticality. In summary, our study suggests that initializing models near criticality is important to find an optimal and flexible solution.

Introduction

Operating close to the critical point at a phase transition between order (subcritical) and disorder (supercritical) has long been associated with optimal performance of complex systems. Several biological systems, such as gene regulatory networks (Balleza et al., 2008; Rämö et al., 2006), neural networks (Tkacik et al., 2015; Schneidman et al., 2006), collectively behaving cells (Halley et al., 2009; De Palo et al.,

2017) or swarms (Cavagna et al., 2010; Chaté and Muñoz, 2014) have been shown to operate close to a critical point. Criticality has been associated with an ability to solve complex tasks (Villegas et al., 2016), flexibility towards changes in the environment and good evolvability (Aldana et al., 2007) in complex living systems (Kauffman, 1993). All these properties provide an adaptive advantage in natural environments, leading to the assumption that evolutionary dynamics push living systems close to the critical regime.

On the other hand, it has been suggested that the ubiquitous presence of noise in nature pushes living systems into a more robust subcritical regime. For example, in an evolutionary model of Random Boolean Networks (RBNs) decreasing the system size, making the task less complex, or introducing noise to the system pushes the optimal regime further into the subcritical range (Villegas et al., 2016). Similarly Rämö et al. (2007) observed, that information propagation is maximized in critical RBNs, however the optimal regime shifts slightly into the subcritical regime under the presence of noise. A related phenomenon has been observed on neuromorphic chips, where simpler tasks lead to optimal behavior in the subcritical regime, whereas harder tasks require progressively more critical dynamics (Cramer et al., 2020). Finally, for some applications, the combination of systems at different distances to criticality is shown to lead to optimal results (Zierenberg et al., 2020). The supercritical state has been universally observed to perform poorly (Villegas et al., 2016; Kauffman, 1993).

The benefits of criticality for the evolvability of living systems have been associated with the genotype-phenotype coupling. Specifically, it has been shown (De Jong, 2006) that a tight genotype-phenotype coupling leads to optimal evolvability. Due to this coupling, the dynamical regime has an impact on the properties of the fitness landscape. In an RBN model, the super- and subcritical regimes were shown to disturb the genotype-phenotype coupling (Kauffman, 1993) and lead to either very rugged or overly flat fitness landscapes. A rugged fitness landscape means that the evolutionary dynamics are just a random search – inefficient in high dimensions (Kauffman and Levin, 1987), whereas a

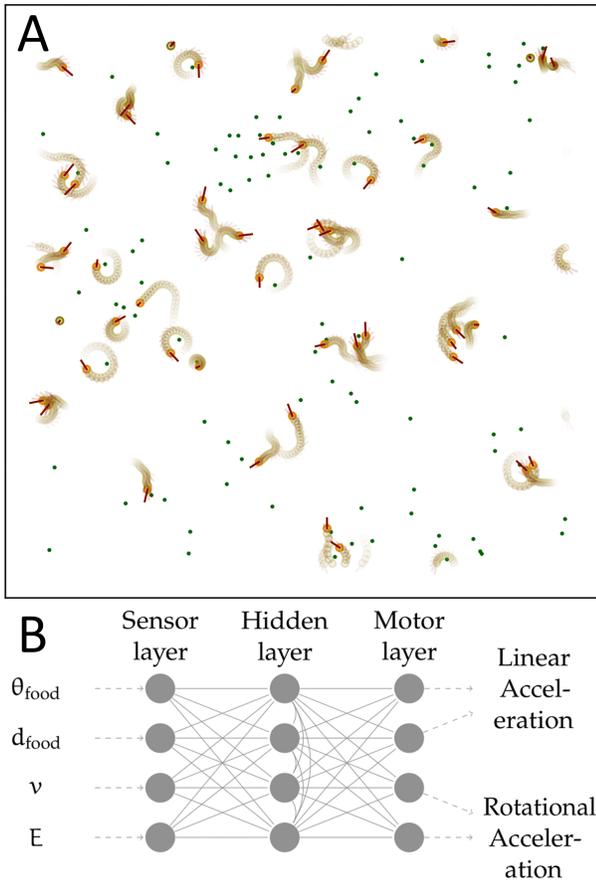


Figure 1: Snapshot in time of population dynamics and schematic representation of the control network. **A:** Environment with 50 organisms (red circles with trails) foraging for food (green dots). **B:** A network with 12 neurons. Four sensory, four hidden and four motor neurons. All potentially allowed edges are displayed. The exact topology and edge weights are subject to evolution by the EA.

very flat landscape strongly dampens the optimization process. Both phenomena lead to a complexity catastrophe, where an increase of system size leads to a failure to discover satisfying solutions with evolutionary search. Critical RBNs result in an intermediately rugged fitness landscape which allows for efficient hill climbing search and is less prone to the complexity catastrophe.

We study how the dynamical regime of populations of evolving organisms influences their ability to solve a task. Our investigation is conducted in a simple foraging game of scalable difficulty, where organisms can gain energy by eating food particles and spend energy when moving. This system allows us to analyze the changes in the dynamical regimes during the evolution. In addition, we propose a potential answer to the question of which dynamical regime demonstrates the best performance and stability with respect

to changes in the environment.

Methods

We investigate a 2D environment where organisms controlled by individual neural networks forage for food. Each organism gains energy by eating food particles and consumes energy by moving. The organisms eat the food particles as soon as they run over them. We can increase the difficulty of the task by requiring the organism’s velocity to be below a certain threshold when running over food in order to be able to consume it. The fitness of an organism is determined by its average energy throughout its lifetime, and an evolutionary algorithm (EA) optimizes the network that controls the organism to maximize its fitness.

Organism

The organisms in our model are controlled by an Ising neural network (INN) that has been previously used by Aguilera and Bedia (2017) as well as Khajehabdollahi and Witkowski (2020). The Ising network consists of N neurons that can be in one of two states $s_i \in \{-1, 1\}$, $i = 1, \dots, N$. There are sensory neurons that only receive input from the sensors and motor neurons that control the agent. All other neurons are hidden units. The connectivity of the network is described by the weight matrix $J \in [-2, 2]^{N \times N}$ and the adjacency matrix $A \in \{0, 1\}^{N \times N}$, as shown in Figure 1B. The adjacency matrix is such that no direct connections between sensor and motor neurons exist. Following the Ising model, each network state has an associated energy

$$e(s_1, \dots, s_N) = - \sum_{\{i,j|A_{ij}=1\}} J_{i,j} s_i s_j \quad (1)$$

where $\{i, j \mid A_{ij} = 1\}$ is the set of connections. The network stochastically minimizes the energy by following Glauber dynamics: At each network iteration, all non-sensor neurons are updated in a random order and the state of neuron i changes from s_i to $-s_i$ with probability:

$$p_i = \frac{1}{1 + e^{\beta \cdot \Delta e_i}}, \quad (2)$$

$$\Delta e_i = e(s_1, \dots, s_i, \dots, s_N) - e(s_1, \dots, -s_i, \dots, s_N),$$

where β is the inverse temperature of the network ($\beta = 1/(T \cdot k_B)$, k_B is the Boltzmann constant which we set to one and omit for simplicity) and Δe_i is the change in the energy of the network that is caused by the spin-flip of the i_{th} neuron (changing its state s_i to $-s_i$). The energy change Δe_i is determined by the connectivity matrix J and the states of neighboring neurons. A negative energy change leads to a greater likelihood of a flip. The parameter β controls the likelihood of energetically unfavorable flips. A larger β leads to deterministic network behavior dominated by the connectivity, whereas a smaller β leads to more random behavior. In principle, many iterations would be required to

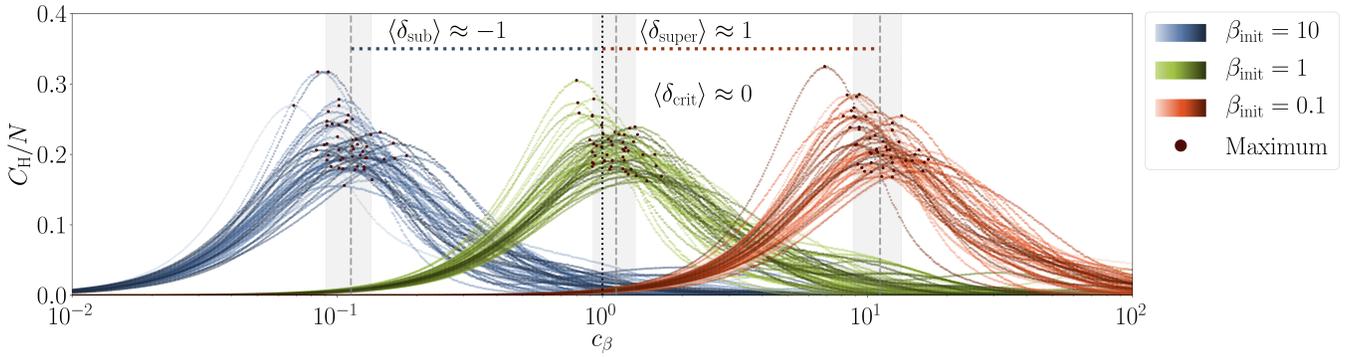


Figure 2: The dynamical regime of a network can be calculated by scaling the inverse temperature with a factor c_β to measure its heat capacity and measuring its distance to the corresponding peak. Heat capacity (Eq. 3) of the Ising networks (Figure 1) for 50 initially subcritical ($\beta_{\text{init}} = 10$, blue), critical ($\beta_{\text{init}} = 1$, green), and supercritical ($\beta_{\text{init}} = 0.1$, red) organisms as a function of c_β . For each organism it reaches the maximum (marked by a dot) at individual values $c_\beta = c_\beta^{\text{crit}}$. Dynamical regime $\delta = \log(c_\beta^{\text{crit}}) \approx -\log(\beta_{\text{init}})$. The displayed populations are unevolved and the resulting dynamical regimes closely correspond to their respective β_{init} .

converge to equilibrium. For practical reasons, we perform a fixed number of iterations (10) to process a new sensor input and create motor commands.

An organism has four input neurons that receive information about the angle θ_{food} and distance d_{food} from the closest food particle as well as its own velocity v and energy E (this energy is distinct from the Ising energy e). Moreover, each organism has four output neurons that control linear and rotational acceleration (2 neurons each) and N_h hidden neurons (Figure 1B). For each time step in the environment we assign a normalized real value to the sensor neurons according to the environmental input. The hidden and motor sensors can only obtain binary states (-1,1). We equilibrate the hidden and motor neurons for an arbitrary number of iterations using a Metropolis algorithm (Metropolis, 1953) which implements equation 2 and subsequently read the states for the motor neurons. The agent accelerates in case both neurons of a motor unit have positive states, decelerates in case both have negative states whereas the velocity remains unchanged for different neuron states.

For most simulations we take $N_h = 4$, to test the scalability of learning we also consider a network with $N_h = 20$. At the beginning of each simulation, an organism is provided with an amount of initial energy $E_{\text{init}} = 2$. Movement reduces energy and consuming food particles increases it. We consider two versions of this environment: In the *simple task* organisms consume food when passing over it. In the *hard task* organisms have to slow down and almost stop to be able to consume food. Unless stated differently, a simulation lasts for a *lifetime* of $t = 2000$ time steps after which the evolutionary algorithm (EA) is applied, and the task is simple. 50 INN-controlled organisms are placed in a 2D environment with periodic boundaries and ever-respawning food particles, conserved to a value of 100 (Figure 1A).

Evolutionary algorithm

The evolutionary algorithm applied to the INNs consists of a combination of elitism, mutation, and mating. At the end of the 2D simulation described above, the fitness of each organism is defined as their mean energy throughout their lifetime. Subsequently, the 20 fittest organisms are selected for reproduction.

The next generation of organisms is produced by applying a combination of copying, mutation, and mating procedures to the selected individuals. The copying algorithm transfers some of the fittest individuals into the next generation unchanged. The mutation algorithm adds or deletes edges in A (connections not present in Figure 1B cannot be added), re-samples a random edge weight in J from a uniform distribution $\mathcal{U}(-2, 2)$, and perturbs the inverse temperature with multiplicative noise $\beta' = \beta \cdot \Delta\beta$ where $\Delta\beta \sim \mathcal{N}(1, 0.02)$. Finally, the mating algorithm takes a weighted average of the connectivity J and inverse temperature β from two parents to produce an offspring. In most of our simulations, the EA iterates for 4000 generations.

Defining the dynamical regime of an organism

We use the heat capacity from statistical physics to derive a measure of an organism's dynamical regime (sub-, super-, critical). In our finite system, we estimate the putative divergence point by changing the inverse temperature β multiplying it with a scaling constant c_β . This change of temperature influences how likely the state of the neurons will flip (Eq. 2), and thus change the equilibrium distribution of energies E (Eq. 1). Now we search for a c_β that delivers the maximal value of the heat capacity $C_H(c_\beta)$, defined as:

$$C_H(c_\beta) = \frac{1}{T^2} \text{Var}(E) = c_\beta^2 \beta^2 \text{Var}(E). \quad (3)$$

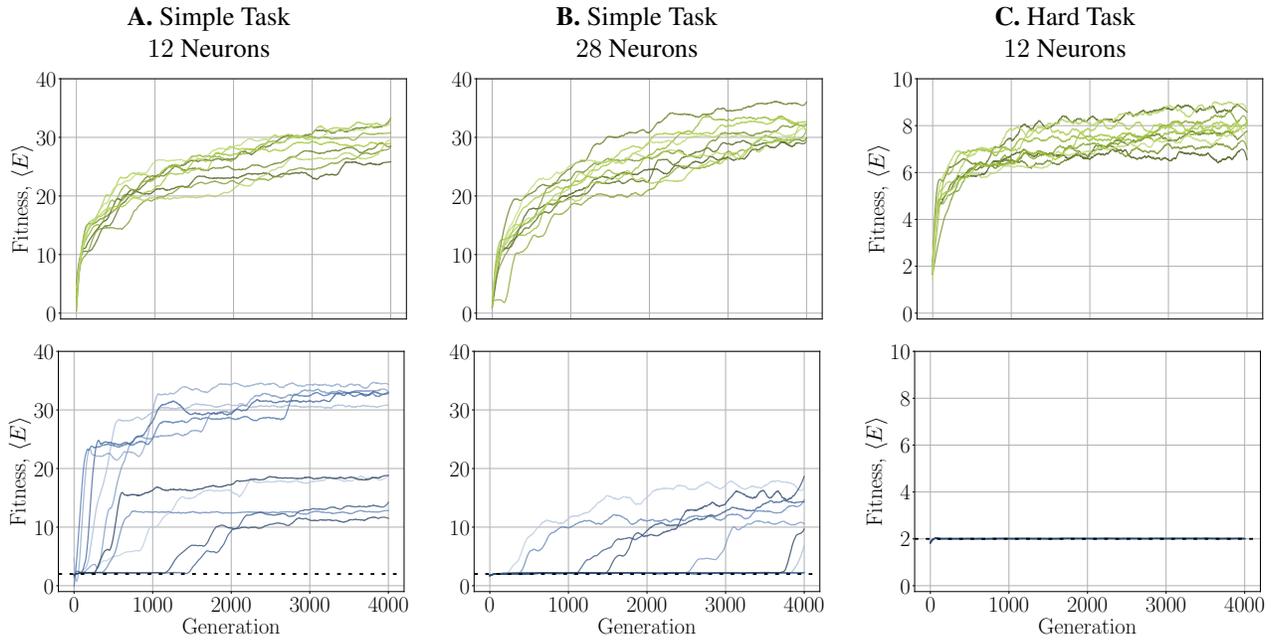


Figure 3: Critically initialized populations can be successfully evolved in different circumstances, whereas for subcritically initiated populations, a harder task or an increased system size can lead to the breakdown of evolutionary dynamics. For each panel 10 initially critical ($\beta = 1$, green, top row) or initially subcritical ($\beta = 0.1$, blue, bottom row) populations evolve for 4000 generations. The dashed line at fitness = 2 in the subcritical panels correspond to the organisms’ initial energy. It can be seen in the 12 neuron, hard task panel (bottom right) that the network is unable to achieve a fitness above its starting value.

Algorithm

```

1: Evolutionary Algorithm  $\triangleright$  Evolves  $J$ ,  $A$ , and  $\beta$ 
2: for generation = 1 to Generations do
3:   Foraging Game in 2D environment
4:   for  $t = 1$  to Time Steps do
5:     Update Sensor Neurons(organism)
6:     Update INN
7:     for iter = 1 to Network Iterations do
8:       for non-sensor neuron in INN do
9:         Potential spin-flip(neuron)
10:      end for
11:    end for
12:    Read Motor Neurons(organism)
13:    Move in 2D environment(organism)
14:  end for
15:  Evolve(population of organisms)
16:  Reset 2D environment
17: end for

```

We define the $c_{\beta}^{\text{crit}} = \operatorname{argmax}_{c_{\beta}} C_H(c_{\beta})$. An analogous procedure was used in (Tkacik et al., 2015).

We define the distance of the network from the critical point by the logarithm of the scaling factor required to bring the network to criticality, $\delta = \log(c_{\beta}^{\text{crit}})$. For unevolved organisms (first generation, Figure 2), the relationship be-

tween β_{init} and δ can be approximated by:

$$\delta \approx -\log \beta_{\text{init}}, \beta_{\text{init}} \in [0.1, 10] \quad (4)$$

On a technical side, for the stability of the numerical procedure for finding the equilibrium distribution to estimate $C_H(c_{\beta})$ we initiate the motor and hidden neurons in the state (s_1, \dots, s_N) delivering minimum of $E(s_1, \dots, s_N)$, while keeping sensory inputs fixed to the observed values. During the evolution, the distance from the critical point (and will) change.

Results

Convergence of evolution

Populations of different initial states follow distinct evolutionary strategies but are all able to solve the standard foraging task. We observe evolution for 4000 generations, concentrating on the populations initiated at the supercritical ($\beta = 10$, $\delta \approx 1$), critical ($\beta = 1$, $\delta \approx 0$) and subcritical ($\beta = 0.1$, $\delta \approx -1$) regimes. Critical populations begin to rapidly gain fitness from the first generation in all EA realizations (Figure 3A). The gradual and stable increase of fitness of the initially critical population suggests that successful hill climbing on the fitness landscape is taking place. In contrast, for subcritical populations, fitness mainly evolves via random jumps, and only

half of the simulations reach the same fitness as the critical populations after 4000 generations (Figure 3A, video: <https://vimeo.com/547613948>). Such fitness dynamics indicate a random search strategy, which often leads to a population getting trapped in local maxima for extended periods of time. Confirming the previous observations by (Khajehabollahi and Witkowski, 2020), we see that supercritical populations, after an initial random period, follow the same path as the critical ones.

For successfully evolvable populations, moderate changes in the complexity of the control network should not destroy the ability of the EA to reach a good fitness.

We test the differences in evolvability for initially critical and subcritical populations by changing the size of the network from 12 to 28 neurons. Same as for smaller networks, the initially critical populations rapidly evolve in all realizations. By generation 4000 they even reach a slightly larger fitness than populations of smaller networks that had evolved for the same amount of time (Figure 3B). At the same time, subcritically-initialized populations do not reach even half of their normal fitness. We observe the same difference between the dynamical regimes when we increase the task's complexity requiring organisms to slow down almost to zero velocity in order to consume food particles (Figure 3C, video: <https://vimeo.com/547615705>). In this harder task, the evolved populations' maximal fitness is expected to be lower than for the simple task. For the initially critical populations, we still observe the same hill climbing dynamics. However, the initially subcritical populations stay at an energy level of exactly two. This signifies that they do not use the originally supplied energy for moving and remain static throughout all 4000 generations, trapped in a local optima.

Overall, we see that although in simple tasks all populations can converge to approximately the same fitness, there exists a significant difference between the initially subcritical and initially critical populations. Namely, a convergence of the EA for critical populations resembles hill climbing. It is stable (all populations follow very similar fitness growth) and behaves similarly for larger networks or more complex tasks. At the same time, for subcritical populations, the EA resembles random search, which is stochastic and fails to find solutions in high-dimensional cases or for more complex tasks.

Evolution of the dynamical regime

Next, we investigate how evolution changes the state of the populations. To do so, we select a wide range of initial dynamical regimes ($\delta \in [-1.5, 1]$) and examine how the dynamics of populations initialized in each of these regimes change throughout evolution. Regardless of their initial dynamical regime, all populations converge to the subcritical regime, however, with different distances from the critical point (Figure 4). We also observe that strongly subcritical

populations ($\delta < -1$) generally achieve lower fitness levels in the simple task and are unable to solve the hard task.

Both near-critical and supercritical populations rapidly change their dynamical regime and by generation 4000 reach an intermediately subcritical regime, whose δ we refer to as δ^* . This is a relatively broad dynamical regime that is intermediately sub-critical but not necessarily deeply sub-critical and is characterized by distinct properties.

Strongly subcritical populations with $\delta \ll \delta^*$ remain at their initial regimes, demonstrating a lack of evolutionary mobility, whereas subcritical populations initialized at different $0 \geq \delta \geq \delta^*$ can still approach δ^* (Figure 4).

Task complexity changes the dynamical regime found in the evolutionary limit. The hard task requires a smaller distance from criticality. We check the evolution of the dynamical regime in both simple and hard tasks (Figure 5). We utilize the observation that all populations with an initial regime $\delta > \delta^*$ converge to similar values. Thus, we consider only initially critical populations. We obtain the distribution of dynamical states by considering 75 independent runs of evolution in both tasks after 4000 generations. The harder task results in a smaller distance from the critical regime (large δ), and the difference is significant ($p < 10^{-13}$, obtained via a two-tailed t-test). The difference is observed throughout the evolution (Figure 5, the main panel based on ten populations).

Initiating the evolution close to the critical regime is important for unknown task complexity. We observe that the dynamical regime never changes towards larger values, but the subcritical convergence point can be at different distances from criticality. Thus, only starting at the critical point guarantees that the optimal dynamical state can be reached by evolution.

Generalizability

For successful biological systems robustness towards environmental change is the paramount feature, therefore, it should be used to distinguish the successfully evolved artificial organisms. We propose a simple measure to investigate how the model behaves outside of its explicit training conditions. Specifically, for a population trained for the organism's lifetime t_{train} we define generalizability as the speed of growth of the average fitness if the organism's lifetime is extended to t_{extend} . Formally:

$$\gamma_t = \frac{\langle E_{t_{\text{train}}} \rangle / t_{\text{train}}}{\langle E_{t_{\text{extend}}} \rangle / t_{\text{extend}}} \quad (5)$$

The *stable generalizability*, $\gamma_t = 1$ corresponds to linear growth whereas sublinear behavior $\gamma_t \ll 1$ indicates possible overfitting to the particular organism's lifetime t_{train} .

We consider initially critical ($\delta \approx 0$) and initially subcritical ($\delta \approx -1$) populations evolved for 4000 generations and then test their performance for an extended lifetime of

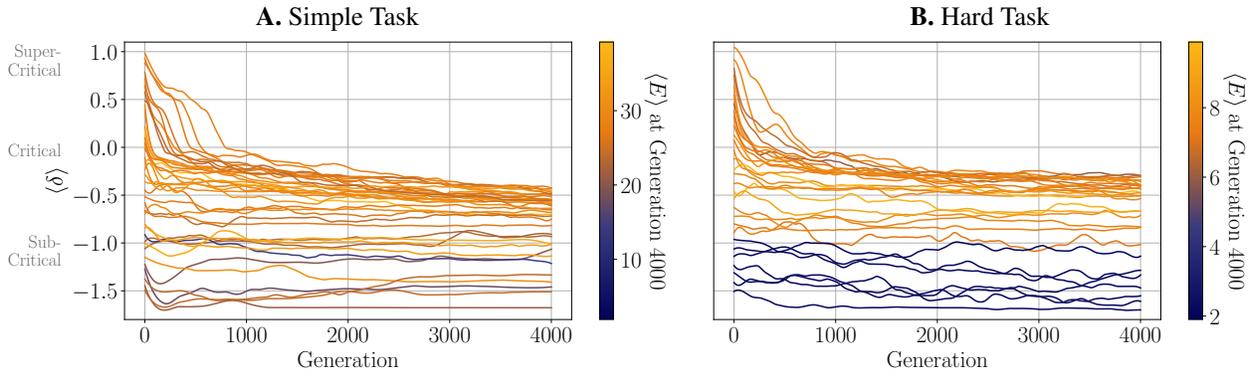


Figure 4: Changes in the distance to criticality over the course of evolution. Populations initiated at various distances to criticality (δ between -1.5 and 1) and evolved on a **A**: simple task and **B**: hard task. The color indicates their fitness at generation 4000. Populations with $\delta < -1$ remain at fitness 2 for the hard task, signifying no evolutionary progress.

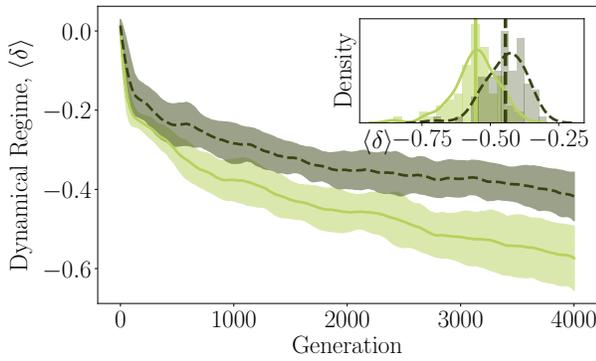


Figure 5: The dynamical regime of the initially critical population in the harder task remains closer to the critical regime, than in the simple task throughout evolution. Main plot: For each condition 10 populations are initiated in the critical regime and evolve for 4000 generations (dark - hard task, light - simple task). Inset: histograms of dynamical regimes at generation 4000 of 75 populations per task, vertical lines denote means, $p < 10^{-13}$. Difference between means in inset and main panel are due to the small sample size for the main panel.

50 000 time steps (instead of the 2000 in training). As reported in previous sections, the critical populations converge to $\delta \approx -0.55$, and they all have a similar fitness after training. Interestingly, when increasing the organisms' lifetime, the fitness of the critical population continues to grow linearly, signifying almost perfect generalizability. About half of the subcritical populations reach the same fitness level. However, the subcritical populations split up into two clusters: one with generalizability close to 0 and another with generalizability close to 1 (Figure 6A). Surprisingly, there is no difference in fitness between these two clusters. When we look more precisely at the individual organisms from the two clusters, we discover that the strategy of the organisms in the

non-generalizable cluster is to increase the velocity permanently until the end of their training lifetime (Figure 6B). However, moving with such a high velocity is not compatible with the energy influx from feeding, and they break down shortly after the end of their training lifetime, this demonstrates that these organisms overfit the training conditions. For all generalizable populations (clusters 1 and 3), the velocity fluctuates around a stable level, and the energy of the organisms grows linearly.

Overall, the initialization in the critical regime results in almost perfect generalizability of evolved populations, whereas initially strongly subcritical populations risk overfit their training conditions.

Effect of genetic perturbations on the fitness

Next, we examine the stability of the evolved organisms to genetic perturbations. We apply genetic perturbations of different magnitudes to the evolved organisms of initially critical $\beta_{\text{init}} = 1$ and subcritical β_{init} populations. We perturb all weights of the connectivity matrix by randomly adding or subtracting a number f_{pert} and then evaluate the fitness of the resulting organism. We find that fitness rapidly declines with perturbation magnitude for both batches of populations, however the subcritical ones decline faster (Figure 7A). We evaluate the fitness decline by the slope of an exponential function fitted to the fitness. For the critical populations it is $\alpha = -2.26$ and for the simple subcritical ones it is $\alpha = -5.03$, which is more than double the decline rate, indicating a much higher sensitivity of subcritical systems to perturbations.

The EA is a source of constant genetic perturbations that are necessary in the beginning of evolution but can become detrimental later. We consider the individual effect of the evolutionary operators (copy, mutate, and mate) on the resulting fitness of the organisms. The variability of fitness for copying simply reflects the natural variability in community fitness rankings and organism behaviour. However,

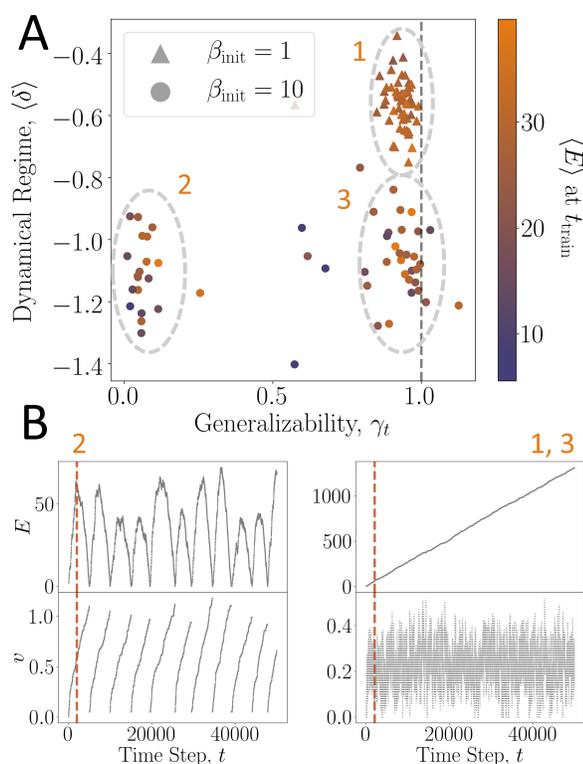


Figure 6: Populations initialized at criticality are always able to find solutions whose behaviors generalize in time beyond their training condition, whereas subcritical populations often overfit to the same conditions. **A**: 54 populations of each type (triangles - initially critical, circles - initially subcritical). After 4000 generations of evolution all critical populations reach a high fitness (indicated by color) and nearly perfect (with one exception) generalizability γ_t (Eq. 5). The initially subcritical organisms split into badly generalizing cluster 2 (19 populations), generalizable cluster 3 (28 populations), and 7 populations that we did not assign to any cluster. This split is independent of their attained fitness. **B**: Energy and velocity as a function of time for representative examples of the organisms from cluster 2 and clusters 1,3 (marked in panel A). Organisms in cluster 2 reach their maximal fitness (sometimes higher than in clusters 1 and 3) at the end of their lives. The dashed orange line denotes the training lifetime of 2000 time steps.

both mating and mutation in fully evolved subcritical populations typically results in a fitness close to 2 – signifying totally unfit organisms (Figure 7B). At the same time, initially critical organisms retain diverse fitness values after mutation and mating, some being close to the optimum. This indicates that the originally critical populations retain their evolvability as opposed to the rigid search performed by strongly subcritical populations.

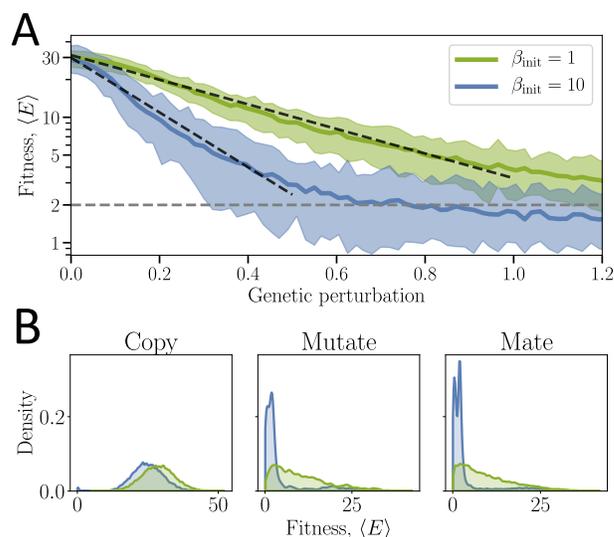


Figure 7: Initially critical populations show larger genotypic stability than initially subcritical ones. **A**: The phenotype (fitness) as a function of the genotypic perturbation (changes in connectivity) for two initial conditions: $\beta_{\text{init}} = 1$ and $\beta_{\text{init}} = 10$. Dashed black lines indicate exponential fit with exponent -2.26 (critical) and -5.03 (subcritical). **B**: The histograms of the fitness values for nearly fully evolved agents (born between generation 3500 and 4000), categorized according to the last evolutionary operator (copy, mutate, or mate) that was applied to them. The $\beta_{\text{init}} = 10$ agents are much less likely to remain fit when their genotype is changed by either mutation or mating.

Discussion

We demonstrate that in various scenarios evolving populations of agents converge to a moderately subcritical state with the resulting deviation from criticality depending on the tasks difficulty. This might appear to be a contradiction to the previous studies, suggesting that operating close to criticality is optimal for natural systems (Mora and Bialek, 2011; Munoz, 2018; Roli et al., 2018). However, a recent body of research showed that for simple tasks, operating at some distance to criticality might be an optimal solution for the sensitivity/stability tradeoff (Hidalgo et al., 2014; Tomen et al., 2014; Villegas et al., 2016; Cramer et al., 2020).

We observe that the distance from criticality affects an agent’s ability to solve complex tasks and to robustly evolve generalizable behavior. Specifically, we observe that slightly subcritical populations are evolvable for different complexities of the control network and task, whereas strongly subcritical populations fail in both cases. Given that these properties are crucial for adaptation in natural environments, we propose that living systems operate in the subcritical regime in close proximity to the critical point. Moreover, we show that the optimal regime moves closer to criticality as we in-

crease the task difficulty, which suggests that the optimal distance from criticality varies. Those findings are confirmed by Cramer et al. (2020) as well as Villegas et al. (2016), who showed that the optimal distance from criticality in the sub-critical regime decreases for higher task complexity or larger system size. We further observe that populations can only become more subcritical during evolution and fail to decrease their distance to criticality even when this would have eventually led to superior behavior. As it is *a priori* unknown which distance from criticality will be optimal when evolving for a new task, starting at the critical point could be the only way for the evolutionary process to descend to the optimal regime. However, in the long run this would require some sub-populations to always maintain closeness to criticality. How this can be achieved for neuronal networks is a subject of vivid research (for a review, see (Zeraati et al., 2021; Buendía et al., 2020; Kinouchi et al., 2020)) and for the embodied Ising agents it remains open for further investigation. Maintaining evolvability in simpler systems was, for instance, observed by switching between different rough energy landscapes (Wang and Dai, 2019). The inhomogeneity of the environment and coevolution can also contribute to the preservation of the critical regime (Hidalgo et al., 2014). Overall, the maintenance of evolvability throughout evolution is an important question beyond the embodied Ising agents studied here.

Our results extend and partly revise the earlier findings of Khajehabdollahi and Witkowski (2020), that reported a superior evolvability of critical populations and an approximate convergence to criticality during evolution. We confirm that the critical regime allows reliable evolvability, additionally, we extend our understanding by considering a set of tasks and architectures in the model. However, our more precise procedure to infer the dynamical regime and the fine sampling of initial conditions uncover additional intricate dynamics. Namely, that critical populations of Ising-agents converge to the subcritical regime, and the distance to criticality depends on the task complexity. We also propose a new way to investigate capabilities of the resulting organisms by defining generalizability and genetic stability measures. Both measures reveal the benefits of staying close to the critical state beyond a simple fitness comparison.

Acknowledgments

This work was supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation. SK and EG thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. GM is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645. We acknowledge the support from the BMBF through the Tübingen AI Center (FKZ: 01IS18039B).

The code can be accessed at https://github.com/LevinaLab/evolution_dynamical_regime.

References

- Aguilera, M. and Bedia, M. G. (2017). Criticality as it could be: organizational invariance as self-organized criticality in embodied agents. In *Artificial Life Conference Proceedings 14*, pages 21–28. MIT Press.
- Aldana, M., Balleza, E., Kauffman, S., and Resendiz, O. (2007). Robustness and evolvability in genetic regulatory networks. *Journal of theoretical biology*, 245(3):433–448.
- Balleza, E., Alvarez-Buylla, E. R., Chaos, A., Kauffman, S., Shmulevich, I., and Aldana, M. (2008). Critical dynamics in genetic regulatory networks: examples from four kingdoms. *PLoS One*, 3(6):e2456.
- Buendía, V., di Santo, S., Bonachela, J. A., and Muñoz, M. A. (2020). Feedback mechanisms for self-organization to the edge of a phase transition. *Frontiers in Physics*, 8:333.
- Cavagna, A., Cimarelli, A., Giardina, I., Parisi, G., Santagati, R., Stefanini, F., and Viale, M. (2010). Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences*, 107(26):11865–11870.
- Chaté, H. and Muñoz, M. A. (2014). Insect swarms go critical. *Physics*, 7:120.
- Cramer, B., Stöckel, D., Kreft, M., Wibrall, M., Schemmel, J., Meier, K., and Priesemann, V. (2020). Control of criticality and computation in spiking neuromorphic networks with plasticity. 11.
- De Jong, K. (2006). Evolutionary computation: A unified approach. bradford book.
- De Palo, G., Yi, D., and Endres, R. G. (2017). A critical-like collective state leads to long-range cell communication in dictyostelium discoideum aggregation. *PLoS biology*, 15(4):e1002602.
- Halley, J. D., Burden, F. R., and Winkler, D. A. (2009). Stem cell decision making and critical-like exploratory networks. *Stem Cell Research*, 2(3):165–177.
- Hidalgo, J., Grilli, J., Suweis, S., Munoz, M. A., Banavar, J. R., and Maritan, A. (2014). Information-based fitness and the emergence of criticality in living systems. *Proceedings of the National Academy of Sciences*, 111(28):10095–10100.
- Kauffman, S. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, 128(1):11–45.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA.
- Khajehabdollahi, S. and Witkowski, O. (2020). Evolution towards criticality in ising neural agents. *Artificial Life*, 26(1):112–129.
- Kinouchi, O., Pazzini, R., and Copelli, M. (2020). Mechanisms of self-organized quasicriticality in neuronal network models. *Frontiers in Physics*, 8:530.
- Metropolis, N. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092. Coauthors: Arianna Rosenbluth, Marshall Rosenbluth, Augusta Teller and Edward Teller.

- Mora, T. and Bialek, W. (2011). Are biological systems poised at criticality? *Journal of Statistical Physics*, 144(2):268–302.
- Munoz, M. A. (2018). Colloquium: Criticality and dynamical scaling in living systems. *Reviews of Modern Physics*, 90(3):031001.
- Rämö, P., Kauffman, S., Kesseli, J., and Yli-Harja, O. (2007). Measures for information propagation in boolean networks. *Physica D: Nonlinear Phenomena*, 227(1):100–104.
- Rämö, P., Kesseli, J., and Yli-Harja, O. (2006). Perturbation avalanches and criticality in gene regulatory networks. *Journal of Theoretical Biology*, 242(1):164–170.
- Roli, A., Villani, M., Filisetti, A., and Serra, R. (2018). Dynamical criticality: overview and open questions. *Journal of Systems Science and Complexity*, 31(3):647–663.
- Schneidman, E., Berry, M. J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012.
- Tkacik, G., Mora, T., Marre, O., Amodei, D., Palmer, S. E., and Bialek, W. (2015). Thermodynamics and signatures of criticality in a network of neurons. *PNAS*, 112(37).
- Tomen, N., Rotermund, D., and Ernst, U. (2014). Marginally subcritical dynamics explain enhanced stimulus discriminability under attention. *Frontiers in systems neuroscience*, 8:151.
- Villegas, P., Ruiz-Franco, J., Hidalgo, J., and Muñoz, M. A. (2016). Intrinsic noise and deviations from criticality in boolean gene-regulatory networks. *Scientific Reports*, 6:34743.
- Wang, S. and Dai, L. (2019). Evolving generalists in switching rugged landscapes. *PLoS computational biology*, 15(10):e1007320.
- Zeraati, R., Priesemann, V., and Levina, A. (2021). Self-organization toward criticality by synaptic plasticity. *Frontiers in Physics*, 9:103.
- Zierenberg, J., Wilting, J., Priesemann, V., and Levina, A. (2020). Tailored ensembles of neural networks optimize sensitivity to stimulus statistics. *Physical Review Research*, 2(1):013115.