Robustness and Adaptability in a Dynamic Neural Field Architecture Subject to Degeneration

João G. Cunha[1], Raymond H. Cuijpers[2] Wolfram Erlhagen[3], and Estela Bicho[1]

[1] Center Algoritmi, University of Minho, 4800-058 Guimaraes, Portugal,
id10787@uminho.pt, estela.bicho@dei.uminho.pt
[2] Human-Technology Interaction, Eindhoven University of Technology, De Zaale, 5612 AZ Eindhoven, The Netherlands, r.h.cuijpers@tue.nl
[3] Center of Mathematics, University of Minho, 4800-058 Guimaraes, Portugal,
wolfram.erlhagen@math.uminho.pt

Abstract. One of the goals of employing biologically inspired models of cognition in artificial systems is to capture their capacity for robustness and adaptation. In biological organisms, this capability is often linked to *degeneracy*, the ability of structurally different elements to support the same function, allowing recovery from damage or loss. This paper investigates the impact of artificially induced *degeneration* on Dynamic Neural Field (DNF) models used for implementing cognitive behaviour in robotic systems. Using a simulated sorting task, we systematically analyse the resilience of DNFs under neuron loss and synaptic disruption, and we examine whether relearning can reorganise degraded architectures to restore correct behaviour. Two experiments are presented: the first evaluates robustness under progressive degeneration, while the second explores the potential of relearning to reorganise the neural connections within the deteriorated architecture. Our results show that DNFs maintain functionality across substantial degradation and can adapt through retraining. These findings provide quantitative evidence that DNFs, through their inherent degeneracy, offer a promising basis for robust and adaptive robotic systems.

Keywords: Dynamic Neural Fields, Biologically-Inspired Robotics, Degeneracy, Robustness, Relearning

1 Introduction

In recent years, there has been growing interest in developing biologically plausible models of cognition for controlling artificial agents [6]. Some of this interest stems from the remarkable adaptability and resilience observed in intelligent biological systems, which can be attributed to one fundamental characteristic known as *degeneracy* [11,19].

Degeneracy is an ubiquitous property in biological systems, from genes to neural networks and evolution itself [3]. It refers to the capacity of structurally different elements to contribute to the same function or outcome, ensuring that if one component is compromised, others can compensate and preserve behaviour [17,29]. A real-life illustration of this phenomenon is observed in individuals with Huntington's disease, where even after losing 10% of brain tissue, they can still perform cognitive tasks at levels

comparable to those without the condition [5]. If robotic systems could emulate this resilience, they could operate reliably in demanding contexts such as space missions, disaster zones, or long-term deployments where repair is impossible.

Building on this perspective, Tononi et al. pioneered the quantitative measurement of degeneracy using an Information Theoretic approach, distinguishing it from redundancy and complexity in simulated artificial neural networks [28]. Subsequent studies have applied these metrics across diverse complex systems such as random boolean networks [10], Hodgkin–Huxley neural networks [16], genetic networks [14], digital system design to obtain fault-tolerant circuits [15], dynamical systems [13], biological networks [7], and to optimize component level network structures [20]. By harnessing degeneracy as a design principle, computational models can achieve robustness and adaptability beyond classical engineering solutions. In robotics, this opens opportunities for fault-tolerant and adaptive control [11, 19].

Dynamic Neural Field (DNF) models have been widely used as bio-inspired control architectures for robotic systems [4]. They provide a computational framework that links neural processing to behaviour through the dynamics of neural populations [23]. Within such populations, individual neurons typically have different but overlapping selectivity, so that many neurons, but not necessarily all, respond to a given stimulus. The resulting activity patterns are shaped not only by feed-forward inputs but also by recurrent excitatory and inhibitory couplings within the neural pool. Such a population representation is thought to have several computational advantages over a "localist" representation [22] (in which the stimulus information is provided by the activity of a single neuron [21, 27]): it reduces the uncertainty associated with the substantial variation of single-neuron activity across repeated presentations of the same stimulus; generalises more easily to similar inputs; and contributes to the brain's ability to resist external perturbations such as damage or failure of individual or small groups of neurons. Moreover, distributed representations are considered to be more biologically plausible, given that electrophysiological recordings have illustrated their extensive use within the cortex [18, 30].

When designing control architectures for artificial systems, a question we continually face is: should entities (e.g., colours or targets) be represented by single neurons or by bumps of population activity? While the representation of discrete entities, such as coloured objects and target containers, can be achieved through single-neuron dynamics, the vulnerabilities of this approach are evident: a single neuron's death could disrupt the entire system. Population-based representations are generally considered more robust to such failures. However, the extent and limits of this robustness remain insufficiently characterised. In particular, it is not yet clear how much degradation DNFs can withstand, which parts of the architecture are most critical, and under what conditions relearning can actually restore lost functionality.

In this paper, we study the effects of artificially induced *degeneration* (defined as the systematic loss of neurons or synaptic connections) on DNF architectures used for robot control. Using a simulated sorting task, we analyse how resilience manifests under different forms of degradation, identifying where the architecture is most vulnerable and how performance deteriorates as degeneration progresses. We also examine

whether relearning can reorganise degraded architectures, allowing lost behaviours to be recovered.

2 Modelling concepts

2.1 Experimental scenario

To address the research questions, two experiments were conducted using the same simulated scenario. In this scenario, a robotic system was designed to execute a simple sorting task, as shown in Fig. 1 (see Supplemental Material for video demonstration). This setup consists of: i) a UR10e robotic arm with an OnRobot RG2 gripper; ii) seven object shapes, each with a distinctive colour (red, green, yellow, blue, orange, indigo, or violet); iii) a camera that detected and processed the hue values of the objects; and iv) a semicircular structure composed of seven uniformly spaced containers (the angular distance between the centres of each container is 4°).

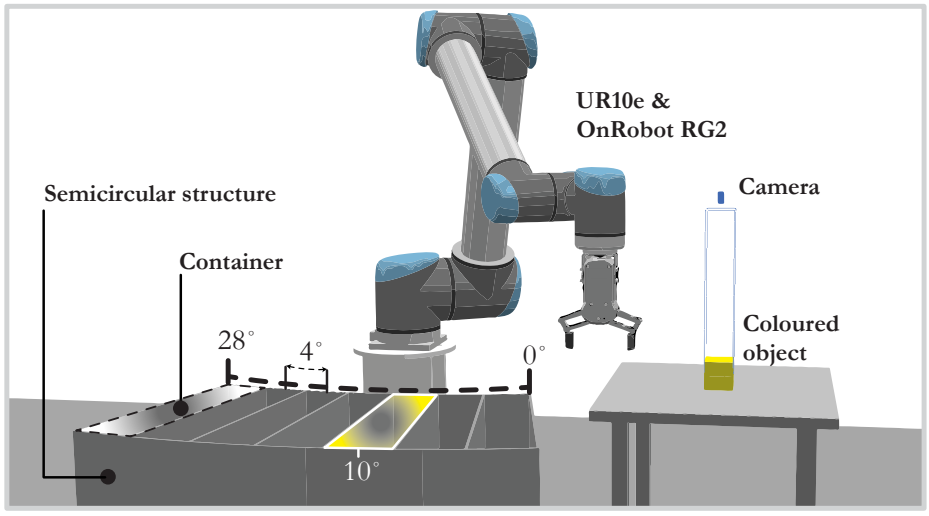

Fig. 1. Experimental setup showing the robotic sorting system, composed of a UR10e robotic arm equipped with an OnRobot RG2 gripper, a camera for hue detection, and a semicircular structure with seven uniformly spaced containers (centres 4° apart).

The robot autonomously picked up each object and, based on its colour, placed it into the corresponding container. This association, *which object goes where*, was learned through a process of learning from demonstration and associative learning. When the robot successfully placed an object inside its respective container, we considered it to have exhibited *correct behaviour*. Conversely, any deviation from this (placement on other containers or no generation of behaviour at all) is considered a *misbehaviour*. This pick-and-place task is well-suited to measure the impact of degeneration because it allows for a clear analysis of correct and incorrect behaviour.

In the first experiment, **inducing degeneration**, the DNF architecture was progressively subjected to several artificial degenerative conditions (including neuron death,

synaptic connection decline, and disruption), until it could no longer generate any behaviour. These computational conditions were inspired by characteristic patterns of human neurodegenerative diseases: communication inefficiency in Huntington's disease [24], gradual synaptic weakening in Alzheimer's disease [26], and selective neuron loss in Parkinson's disease [9]. By tracking how each condition degraded the system's output, we were able to characterise the model's resilience.

In the second experiment, **recovering from degeneration**, the architecture was again exposed to these conditions until the robot started misbehaving. We then investigated the potential for recovery through relearning, exploring the concept of *artificial neuroplasticity*, *i.e.* the ability of a degraded architecture to reorganise its connections and restore previously correct behaviour.

2.2 Model architecture

Controlling the robotic system is a neural architecture comprised of two coupled DNFs: the *perceptual field* and the *output field* (Fig. 2). The perceptual field, u_{per}, receives a transient Gaussian input representing the hue value of the object detected by the vision system. This field is tuned to maintain a self-sustained and self-stabilised bump of activation, thus even after the external stimulus is removed, the bump persists, encoding the hue value of the observed object. This is akin to human perception, where the representation of certain stimuli persists, even if they become absent, until a new stimulus is encountered [8].

The output field, u_{out}, represents the different possible placement angles within the semicircular structure. A bump of activation in this field encodes the target placement angle for the robot to place the object. Unlike the perceptual field, the activation bump in the output field is not self-sustained, *i.e.* it relies on external stimuli provided by the perceptual field to generate a supra-threshold, self-stabilised bump of activation.

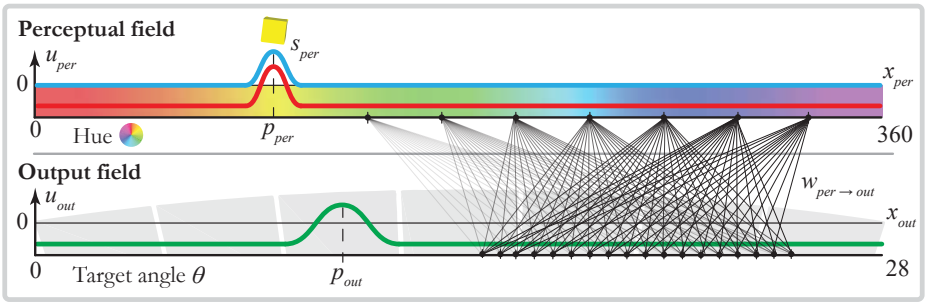

Fig. 2. Control architecture composed of two coupled DNFs: the perceptual field encodes object hue, and the output field encodes the target placement angle. Inter-field connections ($w_{per \to out}$) are learned associatively. A transient stimulus induces a bump in the perceptual field, which drives activation in the output field.

2.3 Model details

In DNFs, neural populations are represented as continuous activation fields, where each point in the field corresponds to the average activation level of a specific group

of neurons. These activation fields evolve based on interactions between neighbouring populations, which can be excitatory or inhibitory. DNFs are capable of emulating a wide range of cognitive processes, including memory, perception, and decision-making, providing a principled framework to study the continuous dynamics that underlie these functions [2, 4, 23].

The specific DNF model employed here follows the formulation originally proposed by Amari to describe pattern formation in neuronal tissue [1]. The population dynamics in each field, $i = \{per, out\}$, are governed by:

$$\tau_i \dot{u}_i(x_i,t) = -u_i(x_i,t) + h_i + s_i(x_i,t) + \int k_{u_i u_i}(x_i - x_i')\, g(u_i(x_i',t))\, dx_i' + q\xi(x_i,t) \quad (1)$$

here $u_i(x_i,t)$ represents the average membrane potential at time t of field i at neural position x_i in a spacial dimension which is a subset of $\mathbb{R}^d$, see implementation details in [4]. $h_i(<0)$ is the resting potential (a measure of the baseline activation level), $s_i(x_i,t)$ represents time-dependent localised input at site x_i from an external source or interconnected fields. The constant τ_i defines the time scale of the field dynamics. The output function $g(u_i)$ dictates that only neurons surpassing a predetermined threshold of activation contribute to the interaction (this function can take the form of, *e.g.*, a Heaviside or a Sigmoid function). The interaction kernel $k_{u_i u_i}(x_i, x_i')$ defines the coupling strength between any two neural positions in the field. ξ is added Gaussian white noise, with strength q (this is used to model the uncertainty seen in actual neurons and also helps in decision making).

All parameters were manually tuned to ensure stable field dynamics and are available at Supplemental Material. In terms of field dimensions, the perceptual field was discretised with spatial step $dx_{per} = 0.5$, a maximum spatial extent of $x_{per}^{max} = 360$, and comprises a total of $N_{per} = x_{per}^{max}/dx_{per} = 720$ neurons. The output field used $dx_{out} = 0.1$, $x_{out}^{max} = 28$, and $N_{out} = 280$ neurons.

External input to the perceptual field is given by a Gaussian stimulus centred in p_{per}, with amplitude a, and width σ, at a certain time t, defined by:

$$s_{per}(x_{per}) = a \exp\left[-\frac{(x_{per} - p_{per})^2}{2\sigma^2}\right] \quad (2)$$

The perceptual field then projects to the output field through weighted inter-field connections:

$$s_{out}(x_{out},t) = \int g(u_{per}(x_{per},t))\, w_{per\to dec}(x_{per}, x_{out}, t)\, dx_{out} \quad (3)$$

where $w_{per\to dec}(x_{per}, x_{out}, t)$ is the inter-field weight, at time t, from neural position x_{per} in the perceptual field to x_{out} in the output field.

Both fields employ a Gaussian interaction kernel:

$$k_{u_i u_i}(x_i - x_i') = c \exp\left[-\frac{(x_i - x_i')^2}{2\sigma^2}\right] - c_{glob} \quad (4)$$

where the amplitude c determines the maximum strength of the connection, σ controls the width or spread of the influence, and c_{glob} represents a global inhibition constant

offset. This choice stems from the need to have, at most, a single activation peak in each field at any given time t.

2.4 How learning occurs

The sorting task, *which object goes where*, was learned through a process of learning from demonstration and associative learning. During learning, the perceptual field (u_{per}) received a stimulus corresponding to the hue value of an object (s_{per}), while the output field (u_{out}) was given a stimulus centred around the corresponding target placement angle (s_{out}^{tar}). This process effectively "showed" the robotic system the association between each object's hue and its correct placement angle within the structure. During this period, the feedforward input from the perceptual to the output field (s_{out}) was disabled, and the output function of the output field is defined as the target output ($g(u_{out}^{tar})$).

The inter-field connection weights $w_{per\to out}$ between u_{per} and u_{out}, were computed using a model based on the Delta Rule for associative learning in DNFs [25]:

$$\tau_w \frac{\delta_w w_{per\to dec}(x_{per},x_{out},t)}{\delta_w t} = g(u_{per}(x_{per},t))g(u_{out}(x_{out},t)) \times (e_{per\to dec}(x_{per},t) - \eta\ w_{per\to dec}(x_{out},x_{out},t)) \quad (5)$$

where the term $-\eta\ w_{per\to dec}(x_{per},x_{out},t)$ is a decay term for avoiding local minima [12]. The error term $e_{per\to dec}(x_{per},t)$ defines the difference between the desired target activity and the current input received by each neuron in the output field:

$$e_{per\to out}(x_{per},t) = g(u_{out}^{tar}(x_{out},t)) - U(x_{out},t) \quad (6)$$

where $U(x_{out},t)$ denotes the total input to each position in the output field:

$$U(x_{out},t) = \int g(u_{per}(x_{per},t))w_{per\to out}(x_{per},x_{out},t)dx_{per} \quad (7)$$

2.5 The degenerative conditions

The artificial degenerative conditions implemented in this study were designed as computational analogues of neurodegenerative phenomena observed in humans. They allow us to systematically investigate how the loss of neurons or synaptic integrity affects the performance of DNF architectures.

A common manifestation of neurodegenerative disease is the progressive death of neurons, as seen, for instance, in Parkinson's disease. To mirror this phenomenon, we devised two distinct computational methods:

- C1 - "Deactivate Pre-Synaptic Neurons": This approach randomly selects a unique neuron in the perceptual field and sets its activation value to zero.
- C2 - "Deactivate Post-Synaptic Neurons": Randomly designate a unique neuron in the output field and, similarly, set its activation value to zero.

Neurodegenerative conditions can also lead to a breakdown in synaptic connections (*e.g.* Huntington's, or Alzheimer's). Our computational strategies for simulating this phenomenon involve three distinct methods:

- C3 - "Deactivate Inter-Field Weights": Here, a unique inter-field weight is randomly selected and assigned a value of zero.
- C4 - "Randomise Inter-Field Weights": This approach introduces randomness to the inter-field weights by selecting a random unique weight and assigning it a value between the initial weight matrix's maximum and minimum values.
- C5 - "Reduce Inter-Field Weights": To mimic the decrease in efficiency of inter-neuronal communication, we systematically diminish the strength of a random unique non-zero weight by 95%.

3 Experiment 1 - Inducing degeneration

The first experiment examined the robustness of DNF architectures *i.e.* the ability to maintain stable and reliable responses, when subjected to repeated instances of neural and synaptic degeneration. All software and implementation details required to reproduce the experiments are publicly available (see Supplemental Material).

3.1 Methodology

For each condition and each colour, we iteratively degraded the condition-specific elements of the architecture. Throughout these experiments, we tracked the centroid of the activation bump in the output field over time, capturing its fluctuations. This allowed us to observe deviations from its original position and assess behavioural stability.

Each experimental trial consisted of two phases: a *setup stage* and an *iterative degeneration stage*. During the setup stage, a Gaussian stimulus centred at position p_{per} was introduced in the perceptual field, and the field dynamics were allowed to stabilise to establish a baseline equilibrium before degeneration. In the subsequent iterative degeneration stage, we applied a certain amount of degeneration and allowed the fields to stabilise accordingly. We repeated this process until all neurons in the output field exhibited sub-threshold activation levels, marking the point of functional failure, when the system could no longer produce meaningful output. In practical robotic terms, this would imply a complete loss of task-relevant behaviour. After each full degradation cycle, all fields were reset to their initial state, and the procedure was repeated for the desired number of trials. The effects of each condition (C1–C5) are illustrated in simulation videos provided in the Supplemental Material.

3.2 Results and analysis

Each condition was tested using inputs corresponding to the seven possible object colours, resulting in 200 trials per condition-colour pair, and a total of 1,400 trials per condition. For each case, we recorded three key metrics: i) the average number of iterations for the bump of activation to disappear; ii) the average number of degenerative iterations for misbehaviour; and iii) the maximum deviation (Table 1).

Disappearance of the activation bump

The architecture exhibited the greatest sensitivity under C1 (deactivation of pre-synaptic neurons), where the activation bump disappeared after only $22.75 \pm 0.40\%$ neurons were affected. In contrast to this, C2 (post-synaptic deactivation) tolerated much higher

Table 1. Impact of different degenerative conditions on the DNF architecture. Metrics include the percentage of affected elements at which the output field's activation bump disappears, the onset of misbehaviour, and the maximum deviation in output angle. Values are reported as mean $\pm$ standard error (SE).

Degenerative Conditions	Affected Elements at Bump Disappearance (% $\pm$ SE)	Affected Elements at Misbehaviour (% $\pm$ SE)	Max. Absolute Deviation ($^\circ \pm$ SE)
C1 – Deactivate pre-synaptic neurons	22.75 ± 0.40	14.77 ± 1.44	7.07 ± 0.80
C2 – Deactivate post-synaptic neurons	89.66 ± 0.32	87.46 ± 1.29	2.15 ± 0.72
C3 – Deactivate inter-field synaptic connections	92.21 ± 0.04	92.21 ± 0.04	1.13 ± 0.09
C4 – Randomise inter-field synaptic connections	87.13 ± 0.12	87.13 ± 0.12	1.58 ± 0.28
C5 – Reduce inter-field synaptic connections by 95%	92.67 ± 0.04	92.67 ± 0.04	1.09 ± 0.07

degeneration, with $89.66 \pm 0.32\%$ neuron loss before the disappearance. The inter-field synaptic conditions (C3–C5) showed similarly high tolerance, with disappearance thresholds above 87%.

This pronounced sensitivity to pre-synaptic degeneration reflects the reliance on the perceptual field's self-sustained activity once the external stimulus is removed. Without sufficient excitatory input, the perceptual representation collapses, preventing activation from propagating to the output field.[4]

Onset of misbehaviour

Misbehaviour appeared earliest under C1, at $14.77 \pm 1.44\%$ affected neurons, confirming that degradation in the perceptual field rapidly disrupts the system's overall function. Under C2, misbehaviour began slightly earlier than bump disappearance ($87.46 \pm 1.29\%$ vs. $89.66 \pm 0.32\%$).

In contrast, when lesioning the inter-field synaptic connections (C3–C5), misbehaviour coincided with bump disappearance, indicating that performance remained correct until connectivity was almost completely lost. This behaviour demonstrates that DNFs are particularly robust to synaptic disruption, even when large fractions of inter-field connections are deactivated, randomised, or severely weakened.

Maximum deviation of activation bump

Condition C1 produced the largest deviation ($7.07 \pm 0.8^\circ$). In a real-world setting, such deviations could lead to frequent misplacements, making the system unreliable in environments requiring high precision. C2 showed a much smaller maximum deviation ($2.15 \pm 0.72^\circ$), while in the inter-field conditions (C3-C5), there were no centroid measurements above or below acceptable limits.

The analysis reveals distinct sensitivities and robustness within the architecture under the different degenerative conditions. The architecture demonstrates a remarkable degree of robustness to degeneration in inter-field synaptic connections (C3, C4, and C5), as evidenced by consistent performance without misbehaviour and limited absolute deviation. The same can be said, to a lesser extent, regarding the deactivation of post-synaptic neurons (C2). On the other hand, the architecture is most vulnerable to the deactivation of pre-synaptic neurons (C1), leading to faster disappearance of the

[4] An alternative design could make the output field self-sustained, but this would mean the robot would maintain internal placement representations even in the absence of perception, which would be undesirable in this task context.

bump, increased misbehaviour, and larger absolute deviations. Nevertheless, despite severe structural loss, the system retains correct behaviour in most conditions, demonstrating quantitative evidence of *degeneracy*.

4 Experiment 2 - Recovering from degeneration

The second experiment investigated whether relearning could serve as a restorative mechanism within a degraded DNF architecture. Understanding relearning in this context is interesting: if a DNF-based controller can recover lost functionality through relearning, system replacement may be avoided, improving operational efficiency and longevity [19].

4.1 Methodology

In this experiment, the architecture was progressively degraded until the robot began to misplace objects. At that point, a relearning procedure grounded in associative learning principles was applied. Deactivated neurons and synaptic connections were excluded from this phase; only functional (”healthy”) neural elements participated in the relearning process. This ensured that the reacquired behavioural patterns resulted solely from reorganisation among functional components rather than reactivation of damaged ones. Thus, we focused on the three conditions involving structural loss: C1 (pre-synaptic neuron deactivation), C2 (post-synaptic neuron deactivation), and C3 (inter-field connection deactivation). Conditions C4 and C5, which merely altered or weakened weights, were excluded since relearning in those cases would ensure recovery.

We began the experiment by introducing incremental degeneration to the architecture, gradually degrading its functionality at specific intervals, depending on the condition. After each round of degeneration, we assess the architecture's performance by observing its ability to exhibit the correct behavioural responses for the different object colours. Degeneration continued as long as the robot executed the task correctly. Once misbehaviour occurred, degeneration ceased, and the relearning phase was initiated.

During relearning, the robot was re-exposed to the correct hue–target associations through repeated demonstrations. Subsequently, we reevaluated the robot's performance. If the robot successfully resumed correct behaviour, the relearning procedure was deemed a success; otherwise, the process was repeated up to a predefined maximum number of demonstrations.

For each trial, we recorded the number of relearning cycles required for successful recovery. This allowed us to understand whether relearning is possible and measure how many demonstrations are needed to restore the desired behavioural dynamics for different levels of degeneration.

4.2 Results and analysis

Deactivation of inter-field synaptic connections

In the deactivate weights condition (C3), the architecture was subjected to a systematic degeneration of inter-field synaptic connections in increments of 0.5% (equivalent to 1008 connections per step). The first signs of behavioural failure appeared at 89.5% degeneration (Fig. 3).

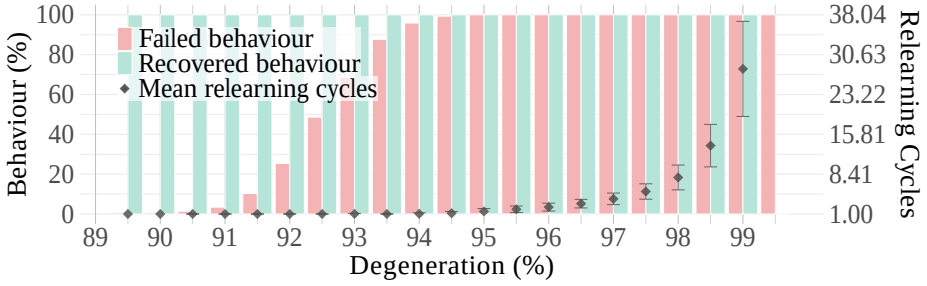

Fig. 3. Impact of deactivating inter-field synaptic weights (C3) on the DNF architecture's ability to recover correct behaviour. Bar plots show the percentage of failed and recovered behaviours across degeneration levels, with diamond markers indicating the mean relearning cycles and standard error.

As degeneration progressed beyond 89.5%, the proportion of fields displaying failed behaviour rose sharply, from 0.68% at 89.5% to 100% at 95% degeneration. The average number of relearning cycles necessary for the recovery of failed behaviours also increased substantially, from one demonstration at 89.5% to an average of 27.98 cycles at 99% degeneration. Additionally, the rising standard error across higher degeneration levels indicates increasing variability in recovery effort, suggesting that the architecture's capacity to reorganise becomes progressively unstable near the upper limits of deterioration. Despite the severe loss of connectivity, full behavioural recovery was achieved through relearning up to 99.5% degeneration. Beyond this threshold, no architecture was able to regain correct performance, indicating a clear point of no return.

Deactivation of post-synaptic neurons

The results from exposing the architecture to the deactivation of post-synaptic neurons condition (C2) tell a slightly different story from the previously analysed condition, as shown in Fig. 4. The first signs of incorrect object placement appeared at 80% degeneration, with a sharp increase in failed behaviour percentage from 0% to 9.36% after just a 0.278% degeneration step (equivalent to one neuron per step). This failure rate remained stable up to 86.12% degeneration, and during this phase, no architectures regained correct behaviour, as indicated by the absence of recovery markers.

Beyond 86.48% degeneration, the percentage of failed behaviour increased exponentially; however, the relearning effort remained constant until 96.56% degeneration. The recovery rate followed an approximately Gaussian distribution, peaking at 52.17% between 92.96% and 93.32% degeneration. This recovery distribution suggests that neuron deactivation leads to a less predictable and more variable recovery process compared to synaptic loss. While some regions of the architecture were able to reorganise or reroute information through remaining neurons, others became irreparably compromised. This uneven response is further reflected in the sharp increase in standard error at higher degeneration levels. At 99.08% degeneration, the standard error rose from 0.088 to 2.368, indicating greater instability in the architecture's ability to recover. As more post-synaptic neurons were deactivated, the architecture struggled to stabilise, and relearning became increasingly irregular and resource-intensive.

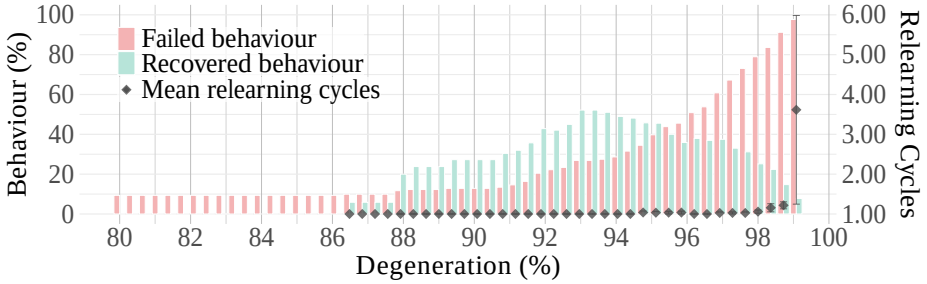

Fig. 4. Impact of deactivating post-synaptic neurons (C2) on behavioural recovery. Bars show the proportion of failed and recovered behaviours; diamond markers represent mean relearning cycles with standard error.

The difference between the deactivation of post-synaptic neurons and the deactivation of synaptic weights stems from the more foundational role the output field has within the architecture. Deactivating synaptic weights reduces connectivity strength but maintains the structural integrity of the neural architecture, allowing for recovery via relearning. In contrast, deactivating post-synaptic neurons directly impairs the architecture's capacity to integrate information, leading to more severe and less recoverable behaviour failures.

Deactivation of pre-synaptic neurons

When the architecture was subjected to the loss of pre-synaptic neurons (C1), its behavioural recovery exhibited distinct dynamics compared to the other types of degeneration. This is to be expected due to the critical role that the pre-synaptic neurons play in transmitting information to the output field (previously discussed in Section 3).

As degeneration increased (discrete steps of 0.1389%; one neuron per step), the behaviour recovery rates decreased rapidly, as shown in Fig. 5. At a degeneration level of 10%, recovery remained relatively high at 87.5%. However, this rate dropped sharply to a minimum of 14.29% at 12.64% degeneration. Interestingly, beyond this point, the recovery rate increased again, reaching a secondary peak of 37.7% around 13.75%, before declining once more. This Gaussian-like pattern (similar to that observed under C2) suggests that the neural architecture was engaging compensatory mechanisms. These mechanisms likely involved strengthening inter-field synaptic connections or recruiting previously unused neurons to compensate for the lost pre-synaptic neurons and restore functionality. Eventually, the recovery rate continues to decline, indicating that the ongoing loss of pre-synaptic neurons overwhelms the architecture's ability to reorganise and maintain behaviour. Another key distinction between the response to C1 and other conditions (C2 and C3) is seen in the evolution of the average number of relearning cycles, which shows considerable variability without a clear pattern.

These results suggest that DNF architectures possess an inherent capability for self-repair, demonstrating that relearning enables the establishment and reorganisation of neural connections within degraded architectures, ultimately restoring the system's initial behavioural patterns.

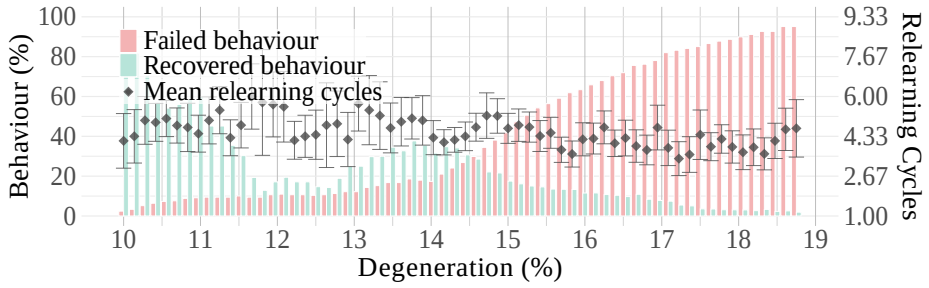

Fig. 5. Impact of deactivating pre-synaptic neurons (C1) on the architecture's capacity to recover behaviour. Bars show the proportion of failed and recovered behaviours; diamond markers represent mean relearning cycles with standard error.

5 Conclusion

Drawing inspiration from biological systems, this work investigated the impact of neural degeneration on Dynamic Neural Field (DNF) architectures, focusing on their robustness and adaptive capacity. We conducted two experiments using a DNF-based control architecture in a simulated robotic sorting task, where objects of different colours were placed into corresponding containers. The architecture was progressively subjected to controlled degeneration, allowing us to determine the thresholds at which the system ceased to exhibit correct behaviour and to investigate whether a degraded system could recover its original functionality through relearning.

Our results provide quantitative evidence of *degeneracy* in DNF architectures. Even under substantial neural lesioning, the models maintained functionality, mirroring the mechanisms observed ubiquitously in biological systems. Furthermore, the experiments demonstrated that DNFs can reorganise and retrain their connections to restore correct behaviour after degeneration, exhibiting a form of *artificial neuroplasticity*. This adaptive capacity distinguishes population-based DNF representations from traditional single-neuron or localist approaches, which typically fail under similar conditions.

These findings suggest that DNF-based architectures offer a promising framework for developing fault-tolerant and self-repairing robotic systems. Their ability to maintain correct behaviour under degradation, and to recover through relearning, positions DNFs as a biologically grounded pathway toward resilient neuromorphic and autonomous systems capable of long-term operation despite hardware decay.

Supplemental Material

Simulation videos (▶ https://youtu.be/) and source code (https://github.com/) supporting this work are available online.

- Main experiment demonstration video ▶ SsyoWnwk_bQ.
- Degeneration condition visualizations ▶: C1 unw0WvjJgUo, C2 PVp4vy-tr1s, C3 j5ee1HWHE0o, C4 LQqGd8gHywA and C5 DKjHoB4Hzw4.
- Code repository : Jgocunha/dynamic-neural-field-degeneration.

Acknowledgements This work was financed by FCT (Portuguese Foundation for Science and Technology) through "I-CATER - Intelligent Robotic Coworker Assistant for Industrial Tasks with an Ergonomics Rationale" Ref. PTDC/EEI-ROB/3488/2021, and through the PhD Fellowship under grant 2022.13948.BD.

References

1. Amari, S.: Dynamics of pattern formation in lateral-inhibition type neural fields. Biological Cybernetics 27, 77–87 (1977), https://link.springer.com/article/10.1007/BF00337259http://link.springer.com/10.1007/BF00337259
2. Coombes, S., beim Graben, P., Potthast, R., Wright, J.: Neural Fields Theory and Applications. Springer Berlin Heidelberg (2014), https://link.springer.com/10.1007/978-3-642-54593-1
3. Edelman, G.M., Gally, J.A.: Degeneracy and complexity in biological systems. Proceedings of the National Academy of Sciences of the United States of America 98, 13763–13768 (11 2001), https://www.pnas.org/doi/10.1073/pnas.231499798
4. Erlhagen, W., Bicho, E.: The dynamic neural field approach to cognitive robotics. Journal of Neural Engineering 3 (2006), https://iopscience.iop.org/article/10.1088/1741-2560/3/3/R02
5. Gray, M., Egan, G., Ando, A., Churchyard, A., Chua, P., Stout, J., Georgiou-Karistianis, N.: Prefrontal activity in huntington's disease reflects cognitive and neuropsychiatric disturbances: The image-hd study. Experimental Neurology 239, 218–228 (1 2013), https://linkinghub.elsevier.com/retrieve/pii/S0014488612004037
6. Hassabis, D., Kumaran, D., Summerfield, C., Botvinick, M.: Neuroscience-inspired artificial intelligence. Neuron 95, 245–258 (7 2017), http://dx.doi.org/10.1016/j.neuron.2017.06.011https://linkinghub.elsevier.com/retrieve/pii/S0896627317305093
7. Hoel, E., Levin, M.: Emergence of informative higher scales in biological systems: A computational toolkit for optimal prediction and control. Communicative and Integrative Biology 13, 108–118 (1 2020), https://pubmed.ncbi.nlm.nih.gov/33014263/
8. Kaminski, J.K., Sullivan, S., Chung, J.M., Ross, I.B., Mamelak, A.N., Rutishauser, U.: Persistently active neurons in human medial frontal and medial temporal lobe support working memory. Nature neuroscience 20, 590 – 601 (2017), https://api.semanticscholar.org/CorpusID:205439887
9. Kashyap, G., Bapat, D., Das, D., Gowaikar, R., Amritkar, R.E., Rangarajan, G., Ravindranath, V., Ambika, G.: Synapse loss and progress of alzheimer's disease -a network model. Scientific Reports 9, 6555 (4 2019), https://www.nature.com/articles/s41598-019-43076-y
10. Kocaoglu, B., Alexander, W.H.: Degeneracy measures in biologically plausible random boolean networks. BMC Bioinformatics 23 (12 2022), https://pubmed.ncbi.nlm.nih.gov/35164672/
11. Krichmar, J.L.: Design principles for biologically inspired cognitive robotics. Biologically Inspired Cognitive Architectures 1, 73–81 (7 2012), https://linkinghub.elsevier.com/retrieve/pii/S2212683X12000084
12. Krogh, A., Hertz, J.A.: A simple weight decay can improve generalization. In: Advances in Neural Information Processing Systems. pp. 950–957. Morgan-Kaufmann (1991), https://proceedings.neurips.cc/paper_files/paper/1991/file/8eefcfdf5990e441f0fb6f3fad709e21-Paper.pdf
13. Li, Y., Dwivedi, G., Huang, W., Kemp, M.L., Yi, Y.: Quantification of degeneracy in biological systems for characterization of functional interactions between modules. Journal of Theoretical Biology 302, 29–38 (6 2012), https://www.sciencedirect.com/science/article/pii/S0022519312000938

14. Macía, J., Solé, R.V., Elena, S.F.: The causes of epistasis in genetic networks. Evolution 66, 586–596 (2 2012), https://academic.oup.com/evolut/article/66/2/586/6851535
15. Man, M., Liu, S., Chang, X., Lu, M.: The biological property of synthetic evolved digital circuits with esd immunity - redundancy or degeneracy? Journal of Bionic Engineering 10, 396–403 (9 2013), http://link.springer.com/10.1016/S1672-6529(13)60235-8
16. Man, M., Zhang, Y., Ma, G., Friston, K., Liu, S.: Quantification of degeneracy in hodgkin-huxley neurons on newman-watts small world network. Journal of Theoretical Biology 402, 62–74 (8 2016), https://pubmed.ncbi.nlm.nih.gov/27155043/
17. Mason, P.H.: Degeneracy: Demystifying and destigmatizing a core concept in systems biology. Complexity 20, 12–21 (1 2015), https://onlinelibrary.wiley.com/doi/10.1002/cplx.21534
18. O'Reilly, R.C.: Six principles for biologically based computational models of cortical cognition. Trends in Cognitive Sciences 2, 455–462 (11 1998), https://linkinghub.elsevier.com/retrieve/pii/S1364661398012418
19. Pfeifer, R., Bongard, J.: How the Body Shapes the Way We Think, vol. s3-XI. The MIT Press (2006), https://mitpress.mit.edu/9780262537421/how-the-body-shapes-the-way-we-think/
20. Randles, M., Lamb, D., Odat, E., Taleb-Bendiab, A.: Distributed redundancy and robustness in complex systems. Journal of Computer and System Sciences 77, 293–304 (3 2011), https://linkinghub.elsevier.com/retrieve/pii/S0022000010000097
21. Rolls, E.T.: Cortical coding. Language, Cognition and Neuroscience 32, 316–329 (3 2017), https://www.tandfonline.com/doi/full/10.1080/23273798.2016.1203443
22. Roy, A.: An extension of the localist representation theory: grandmother cells are also widely used in the brain. Frontiers in Psychology 4 (2013), http://journal.frontiersin.org/article/10.3389/fpsyg.2013.00300/abstract
23. Schöner, G., Spencer, J., Group, D.R.: Dynamic Thinking: A Primer on Dynamic Field Theory. Oxford University Press (11 2015), https://academic.oup.com/book/10360
24. Snowden, J.S.: The neuropsychology of huntington's disease. Archives of clinical neuropsychology : the official journal of the National Academy of Neuropsychologists 32, 876–887 (11 2017), http://www.ncbi.nlm.nih.gov/pubmed/28961886
25. Sousa, E., Erlhagen, W., Ferreira, F., Bicho, E.: Off-line simulation inspires insight: A neurodynamics approach to efficient robot task learning. Neural Networks 72, 123–139 (12 2015), https://linkinghub.elsevier.com/retrieve/pii/S089360801500177X
26. Tansey, M.G., Goldberg, M.S.: Neuroinflammation in parkinson's disease: Its role in neuronal death and implications for therapeutic intervention. Neurobiology of Disease 37, 510–518 (3 2010), https://linkinghub.elsevier.com/retrieve/pii/S096999610900326X
27. Thomas, E., French, R.: Grandmother cells: much ado about nothing. Language, Cognition and Neuroscience 32, 342–349 (3 2017), https://www.tandfonline.com/doi/full/10.1080/23273798.2016.1235279
28. Tononi, G., Sporns, O., Edelman, G.M.: A measure for brain complexity: Relating functional segregation and integration in the nervous system. Proceedings of the National Academy of Sciences 91, 5033–5037 (5 1994), https://pnas.org/doi/full/10.1073/pnas.91.11.5033
29. Whitacre, J.M.: Degeneracy: a link between evolvability, robustness and complexity in biological systems. Theoretical Biology and Medical Modelling 7, 6 (12 2010), https://tbiomed.biomedcentral.com/articles/10.1186/1742-4682-7-6
30. Zhong, S., Choi, J.W., Hashoush, N.G., Babayan, D., Malekmohammadi, M., Pouratian, N., Christopoulos, V.: A neurocomputational theory of action regulation predicts motor behavior in neurotypical individuals and patients with Parkinson's disease. PLoS Comput Biol 18(11), e1010111 (Nov 2022), https://dx.plos.org/10.1371/journal.pcbi.1010111