

# Pulse Neural Networks at the Threshold of Transformation: Emerging Paradigms, Neuromorphic Convergence, and the Path Towards Brain-Inspired Intelligence

Richard Murdoch Montgomery  
Universidade de São Paulo  
Email: [montgomery@alumni.usp.br](mailto:montgomery@alumni.usp.br) DOI: 10.62161/117471

## Abstract:

This comprehensive article examines the transformative potential of pulse neural networks (PNNs) as we stand at a critical juncture in neuromorphic computing. Drawing upon recent breakthroughs in spike-timing-dependent plasticity, event-driven processing architectures, and bio-inspired learning algorithms, we delineate the trajectory of PNN development over the coming decade. We first establish the current landscape, highlighting the convergence of theoretical neuroscience insights with practical engineering implementations in neuromorphic hardware platforms such as Intel's Loihi 2, IBM's TrueNorth successor systems, and emerging memristive architectures.

The analysis proceeds to identify five pivotal areas poised for revolutionary advancement: (1) the integration of multi-timescale dynamics enabling PNNs to capture both rapid sensory processing and slower cognitive phenomena; (2) the development of hybrid architectures combining pulse-based computation with conventional deep learning paradigms; (3) the emergence of self-organising criticality in large-scale PNN systems, promising unprecedented efficiency in information processing; (4) novel training methodologies that transcend the limitations of backpropagation through time, including evolutionary strategies and local learning rules inspired by cortical microcircuits; and (5) the potential for PNNs to achieve true continual learning without catastrophic forgetting.

We further examine the implications of quantum-inspired spike processing, the role of stochastic resonance in enhancing signal detection, and the prospects for PNNs in edge computing applications where energy efficiency is paramount. Critical challenges are addressed, including the spike-timing precision paradox, the credit assignment problem in multi-layer spiking architectures, and the standardisation of neuromorphic software frameworks. The article concludes by proposing a roadmap towards artificial general intelligence through several pathways, including PNNs, emphasising the necessity of

interdisciplinary collaboration between neuroscientists, computer scientists, and engineers to realise the full potential of brain-inspired computing paradigms.

**Keywords:** pulse neural networks, spiking neural networks, neuromorphic computing, spike-timing-dependent plasticity, event-driven processing, brain-inspired artificial intelligence, temporal coding, synaptic plasticity, edge computing, continual learning, neuromorphic hardware, computational neuroscience.

## 1. Introduction

The emergence of artificial neural networks (ANNs) and deep learning represents one of the most profound intellectual achievements of the late twentieth and early twenty-first centuries, fundamentally altering our understanding of both biological cognition and computational intelligence. This field, situated at the intersection of neuroscience, mathematics, and computer science, has evolved from humble beginnings in the 1940s to become the cornerstone of modern artificial intelligence systems (LeCun et al., 2015; Schmidhuber, 2015).

### 1.1. Historical Foundations and Conceptual Framework

The conceptual genesis of artificial neural networks can be traced to McCulloch and Pitts (1943), who proposed the first mathematical model of biological neurons, establishing the theoretical foundation for computational models of neural processing. This seminal work demonstrated that networks of simplified neuron-like units could, in principle, compute any logical function, thereby bridging the gap between biological neural systems and formal computation. The subsequent development by Rosenblatt (1958) of the perceptron—a single-layer neural network capable of learning from examples—marked the transition from theoretical speculation to practical implementation.

However, the field's trajectory was far from linear. The influential critique by Minsky and Papert (1969) highlighted fundamental limitations of single-layer perceptrons, precipitating what became known as the first "AI winter." Their

demonstration that simple perceptrons could not learn non-linearly separable functions, such as the exclusive-or (XOR) problem, seemed to cast doubt upon the entire enterprise. Yet, this apparent setback ultimately catalysed deeper theoretical insights. The development of backpropagation by Rumelhart et al. (1986) provided an elegant solution to training multi-layer networks, effectively circumventing the limitations identified by Minsky and Papert.

## 1.2. Core Principles of Deep Learning

*The contemporary paradigm of deep learning, whilst building upon these historical foundations, represents a qualitative leap in both scale and capability. Deep neural networks, characterised by multiple hidden layers between input and output, exploit hierarchical features learning to extract increasingly abstract representations from raw data (Bengio, 2009; Goodfellow et al., 2016). This hierarchical processing mirrors the organisation of biological sensory systems, particularly the visual cortex, where neurons in successive layers respond to progressively more complex features (Hubel & Wiesel, 1962; Kriegeskorte, 2015).*

The mathematical foundation of deep learning rests upon several key principles. First, the universal approximation theorem demonstrates that feedforward networks with a single hidden layer can approximate any continuous function to arbitrary precision, given sufficient neurons (Hornik et al., 1989; Cybenko, 1989). However, the true power of deep architectures lies not merely in their theoretical expressiveness but in their practical efficiency—deep networks can represent complex functions more compactly than their shallow counterparts (Bengio & LeCun, 2007).

*The training of deep networks relies fundamentally upon gradient-based optimisation, typically employing variants of stochastic gradient descent (SGD). The backpropagation algorithm efficiently computes gradients through the chain rule of calculus, enabling the adjustment of millions or even billions of*

parameters (Rumelhart et al., 1986). Recent innovations in optimisation algorithms, including adaptive learning rate methods such as Adam (Kingma & Ba, 2014) and RMSprop (Tieleman & Hinton, 2012), have substantially improved training efficiency and stability.

### **1.3. Architectural Innovations and Methodological Advances**

*The renaissance of deep learning in the 2010s was precipitated by several converging factors: the availability of large-scale datasets, dramatic improvements in computational power through graphics processing units (GPUs), and crucial algorithmic innovations (LeCun et al., 2015). Convolutional Neural Networks (CNNs), originally proposed by LeCun et al. (1998), revolutionised computer vision by incorporating translation invariance and local connectivity, dramatically reducing the number of parameters whilst maintaining representational power. The success of AlexNet (Krizhevsky et al., 2012) in the ImageNet competition marked a watershed moment, demonstrating the superiority of deep learning approaches over traditional computer vision methods.*

Recurrent Neural Networks (RNNs) and their sophisticated variants, particularly Long Short-Term Memory networks (LSTMs) (Hochreiter & Schmidhuber, 1997) and Gated Recurrent Units (GRUs) (Cho et al., 2014), extended deep learning to sequential data processing. These architectures addressed the vanishing gradient problem that plagued earlier recurrent models, enabling the learning of long-range dependencies crucial for natural language processing and time series analysis.

*The introduction of the Transformer architecture by Vaswani et al. (2017) represents perhaps the most significant recent innovation, eschewing recurrence in favour of self-attention mechanisms. This approach has not only dominated natural language processing but has also been successfully adapted to computer vision (Dosovitskiy et al., 2020) and multimodal learning (Radford et al., 2021),*

suggesting a potential unification of deep learning architectures across modalities.

#### 1.4. Applications Across Domains

*The practical applications of deep learning span virtually every domain of human endeavour.* In healthcare, deep learning models have achieved expert-level performance in diagnosing diabetic retinopathy (Gulshan et al., 2016), detecting cancer in mammograms (McKinney et al., 2020), and predicting protein structures (Jumper et al., 2021). The latter achievement, exemplified by DeepMind's AlphaFold, represents a breakthrough in structural biology with profound implications for drug discovery and our understanding of biological processes.

In natural language processing, large language models based on the Transformer architecture have demonstrated remarkable capabilities in text generation, translation, and comprehension (Brown et al., 2020; Devlin et al., 2019). These models exhibit emergent properties, including few-shot learning and chain-of-thought reasoning, that were not explicitly programmed but arose from scale and training on diverse datasets (Wei et al., 2022).

The impact on scientific research extends beyond direct applications. *Deep learning has become an indispensable tool in physics, enabling the analysis of gravitational wave data (George & Huerta, 2018), accelerating climate simulations (Rasp et al., 2018), and facilitating discoveries in materials science (Butler et al., 2018).* In neuroscience itself, deep learning models serve both as tools for data analysis and as computational hypotheses about brain function (Richards et al., 2019; Yamins & DiCarlo, 2016).

#### 1.5. Critical Arguments and Theoretical Debates

Despite these remarkable successes, deep learning faces substantial criticism and theoretical challenges. *The "black box" nature of deep neural networks raises*

*concerns about interpretability and trustworthiness, particularly in high-stakes applications such as healthcare and criminal justice (Rudin, 2019). Whilst techniques for neural network interpretability have advanced considerably (Montavon et al., 2018; Sundararajan et al., 2017), the fundamental tension between model complexity and human comprehensibility remains unresolved.*

The relationship between deep learning and biological intelligence continues to generate debate. Critics argue that current deep learning systems, despite their impressive performance, fail to capture essential aspects of human cognition, including causal reasoning, compositional understanding, and learning from limited data (Marcus, 2018; Lake et al., 2017). The requirement for vast amounts of labelled data contrasts sharply with human learning efficiency, suggesting that fundamental insights about intelligence remain elusive.

Theoretical understanding of deep learning lags behind empirical success. The optimisation landscape of neural networks is highly non-convex, yet gradient descent reliably finds good solutions—a phenomenon not fully explained by existing theory (Choromanska et al., 2015). The implicit regularisation effects of SGD (Neyshabur et al., 2017), the role of overparameterisation (Zhang et al., 2017), and the emergence of feature learning (Arora et al., 2019) remain active areas of theoretical investigation.

## **1.6.Future Perspectives and Emerging Paradigms**

Looking towards the future, several trajectories appear particularly promising. The integration of symbolic reasoning with neural approaches, often termed neurosymbolic AI, seeks to combine the pattern recognition capabilities of deep learning with the compositional and logical reasoning of classical AI (Garcez et al., 2019). This synthesis could address current limitations in generalisation and interpretability whilst maintaining the flexibility of neural approaches.

Self-supervised learning has emerged as a powerful paradigm for leveraging unlabelled data, with methods such as contrastive learning (Chen et al., 2020) and masked language modelling (Devlin et al., 2019) achieving remarkable results. This approach aligns more closely with biological learning and could reduce the dependence on costly labelled datasets.

The intersection of deep learning with quantum computing presents intriguing possibilities. Quantum neural networks could potentially exploit quantum superposition and entanglement to achieve exponential speedups for certain problems (Biamonte et al., 2017). Whilst practical quantum advantage remains elusive, theoretical developments continue to illuminate potential synergies.

*Energy efficiency represents both a challenge and an opportunity. The computational demands of large-scale deep learning models raise environmental concerns and limit deployment in resource-constrained settings (Strubell et al., 2019). Neuromorphic computing, inspired by the energy efficiency of biological brains, offers a potential path towards more sustainable AI systems (Davies et al., 2018).*

### **1.7.Philosophical and Societal Implications**

The rapid advancement of deep learning necessitates careful consideration of broader implications. The potential for artificial general intelligence (AGI) remains contentious, with estimates ranging from decades to centuries (Bostrom, 2014; Russell, 2019). Regardless of timeline, the societal impact of increasingly capable AI systems demands proactive governance and ethical frameworks (Floridi et al., 2018).

*The democratisation of AI through open-source frameworks and pre-trained models has accelerated research and application development. However, the concentration of computational resources and data in a few large organisations raises concerns about equity and control (Ahmed & Wahed, 2020). Ensuring broad access to AI capabilities whilst managing risks represents a crucial challenge for the field.*

## 1.8. Conclusion

Artificial neural networks and deep learning have transformed from a biologically-inspired computational curiosity to the driving force behind contemporary artificial intelligence. The field stands at a fascinating juncture, with remarkable practical successes coexisting with fundamental theoretical mysteries. As we advance, the integration of insights from neuroscience, the development of more efficient and interpretable architectures, and the careful consideration of societal implications will be essential.

The journey from McCulloch and Pitts's simple threshold units to today's large language models and multimodal systems illustrates both the power of the core ideas and the importance of sustained scientific inquiry. As deep learning continues to evolve, it promises not only to enhance our technological capabilities but also to deepen our understanding of intelligence itself—both artificial and biological. The coming decades will likely witness further convergence between machine learning, neuroscience, and cognitive science, potentially yielding insights that transform our understanding of mind and computation.

## 2. Methodology: Mathematical Foundations of Pulse Neural Networks

### 2.1 Fundamental Neuron Models

The mathematical formulation of pulse neural networks begins with the characterisation of individual spiking neurons. Unlike their rate-based counterparts, spiking neurons explicitly model the temporal dynamics of membrane potential and the generation of action potentials. We shall commence with the foundational integrate-and-fire model before progressing to more biologically plausible formulations.

#### 2.1.1 The Leaky Integrate-and-Fire Model

The leaky integrate-and-fire (LIF) neuron represents perhaps the most elementary yet practically useful spiking neuron model. The membrane potential dynamics are governed by the differential equation:

$$\tau_m \frac{dV(t)}{dt} = -(V(t) - V_{\text{rest}}) + R_m I(t) \quad (1)$$

where:

- $V(t)$  denotes the membrane potential at time  $t$
- $\tau_m = R_m C_m$  represents the membrane time constant
- $R_m$  is the membrane resistance
- $C_m$  is the membrane capacitance
- $V_{\text{rest}}$  is the resting potential (typically -70 mV)
- $I(t)$  represents the total input current

When the membrane potential reaches a threshold  $V_{\text{th}}$  (typically -55 mV), the neuron emits a spike, and the potential is reset to  $V_{\text{reset}}$  (often equal to  $V_{\text{rest}}$ ).

This firing mechanism can be expressed as:

$$\text{if } V(t) \geq V_{\text{th}}, \text{ then } V(t^+) = V_{\text{reset}} \text{ and emit spike (2)}$$

### 2.1.2 Synaptic Current Dynamics

The input current  $I(t)$  comprises contributions from all presynaptic neurons.

For a neuron  $i$  receiving inputs from neurons  $j$ , the total synaptic current is:

$$I_i(t) = \sum_j w_{ij} \sum_{t_j^f} \alpha(t - t_j^f) \quad (3)$$

where:

- $w_{ij}$  represents the synaptic weight from neuron  $j$  to neuron  $i$
- $t_j^f$  denotes the  $f$ -th spike time of neuron  $j$

- $\alpha(t)$  is the postsynaptic current kernel

A common choice for the kernel is the exponential function:

$$\alpha(t) = \frac{1}{\tau_s} e^{-t/\tau_s} \Theta(t) \quad (4)$$

where  $\tau_s$  is the synaptic time constant and  $\Theta(t)$  is the Heaviside step function.

### 2.1.3 The Hodgkin-Huxley Formalism

For investigations requiring greater biological fidelity, the Hodgkin-Huxley model provides a detailed description of ionic conductances:

$$C_m \frac{dV}{dt} = -g_{Na} m^3 h (V - E_{Na}) - g_K n^4 (V - E_K) - g_L (V - E_L) + I_{ext} \quad (5)$$

The gating variables  $m$ ,  $h$ , and  $n$  evolve according to:

$$\frac{dx}{dt} = \alpha_x(V)(1 - x) - \beta_x(V)x \quad (6)$$

where  $x \in \{m, h, n\}$  and  $\alpha_x(V), \beta_x(V)$  are voltage-dependent rate functions.

Whilst computationally intensive, this formulation captures the rich dynamics of biological neurons, including refractory periods and spike shape.

## 2.2 Network Architecture and Dynamics

### 2.2.1 Network Connectivity

Consider a network of  $N$  spiking neurons with connectivity defined by the weight matrix  $\mathbf{W} \in \mathbb{R}^{N \times N}$ . The network dynamics can be described by the coupled differential equations:

$$\tau_m \frac{dV_i(t)}{dt} = -(V_i(t) - V_{rest}) + R_m \sum_j w_{ij} s_j(t) + R_m I_i^{ext}(t) \quad (7)$$

where  $s_j(t) = \sum_{t_j^f} \delta(t - t_j^f)$  represents the spike train of neuron  $j$ , and  $I_i^{ext}(t)$  denotes external input to neuron  $i$ .

### 2.2.2 Population Dynamics

For large-scale networks, a mean-field approach proves invaluable. The population firing rate  $r(t)$  can be approximated by:

$$r(t) = \Phi[\mu(t), \sigma^2(t)] \quad (8)$$

where  $\Phi$  is the transfer function relating the mean  $\mu(t)$  and variance  $\sigma^2(t)$  of the input to the output firing rate. For the LIF neuron with diffusive noise, this yields:

$$r = \left[ \tau_{\text{ref}} + \tau_m \int_{V_{\text{reset}}}^{V_{\text{th}}} \frac{dV}{\mu + \sigma \sqrt{2/\tau_m V}} \exp\left(\frac{V^2 - (\mu\tau_m/\sigma)^2}{2\tau_m}\right) \right]^{-1} \quad (9)$$

## 2.3 Learning Mechanisms

### 2.3.1 Spike-Timing-Dependent Plasticity

The cornerstone of biologically plausible learning in pulse neural networks is spike-timing dependent plasticity (STDP). The weight change depends upon the relative timing of pre- and postsynaptic spikes:

$$\Delta w_{ij} = \sum_{t_i^{\text{post}}} \sum_{t_j^{\text{pre}}} W(\Delta t) \quad (10)$$

Let  $\Delta t = t_i^{\text{post}} - t_j^{\text{pre}}$ , where  $t_i^{\text{post}}$  denotes the time of a postsynaptic spike and  $t_j^{\text{pre}}$  the time of a presynaptic spike. The learning window  $W(\Delta t)$  is typically defined as:

$$W(\Delta t) = \begin{cases} A_+ \exp\left(-\frac{\Delta t}{\tau_+}\right), & \text{if } \Delta t > 0 \\ -A_- \exp\left(\frac{\Delta t}{\tau_-}\right), & \text{if } \Delta t < 0 \end{cases} \quad (11)$$

where:

- $A_+ > 0$  and  $A_- > 0$  determine the maximal potentiation and depression, respectively,
- $\tau_+ > 0$  and  $\tau_- > 0$  specify the time constants controlling the temporal width of the potentiation and depression windows.

### 2.3.2 Triplet STDP

Recent experimental evidence suggests that pairwise STDP inadequately captures synaptic plasticity at higher firing rates. The triplet rule incorporates interactions between three spikes:

$$\Delta w_{ij} = \sum_{t_i} \sum_{t_j} \left[ A_2^+ r_j^{(1)}(t_i) + A_3^+ r_j^{(1)}(t_i) r_i^{(2)}(t_i) \right] - \sum_{t_j} \sum_{t_i} \left[ A_2^- r_i^{(1)}(t_j) + A_3^- r_i^{(1)}(t_j) r_j^{(2)}(t_j) \right] \quad (12)$$

where  $r^{(1)}$  and  $r^{(2)}$  are trace variables with different time constants:

$$\frac{dr_i^{(k)}}{dt} = -\frac{r_i^{(k)}}{\tau_k} + \sum_{t_i^f} \delta(t - t_i^f) \quad (13)$$

## 2.4 Gradient-Based Learning

### 2.4.1 Surrogate Gradient Methods

The discontinuous nature of the spike generation mechanism poses challenges for gradient based optimisation. The surrogate gradient approach replaces the non-differentiable spike function with a smooth approximation during backpropagation:

$$\frac{\partial s}{\partial V} \approx \sigma'(V - V_{th}) \quad (14)$$

where  $\sigma'$  is the derivative of a smooth function such as the sigmoid:

$$\sigma'(x) = \frac{\beta}{4} \operatorname{sech}^2\left(\frac{\beta x}{2}\right) \quad (15)$$

The parameter  $\beta$  controls the sharpness of the approximation.

### 2.4.2 Backpropagation Through Time

For a loss function  $\mathcal{L}$ , the gradient with respect to weights can be computed using backpropagation through time (BPTT):

$$\frac{\partial \mathcal{L}}{\partial w_{ij}} = \sum_t \frac{\partial \mathcal{L}}{\partial V_i(t)} \frac{\partial V_i(t)}{\partial w_{ij}} \quad (16)$$

The dynamics of the error signal  $e_i(t) = \frac{\partial \mathcal{L}}{\partial V_i(t)}$  (17) follow:

$$\frac{de_i(t)}{dt} = -\frac{e_i(t)}{\tau_m} + \sum_j w_{ji} e_j(t + dt) \sigma'(V_j(t + dt) - V_{th}) \quad (18)$$

This must be integrated backwards in time from the final time step.

## 2.5 Advanced Training Methodologies

### 2.5.1 Equilibrium Propagation

An alternative to BPTT that addresses biological plausibility concerns is equilibrium propagation. The network dynamics are augmented with an energy function:

$$E = \sum_i \frac{1}{2\tau_i} V_i^2 - \sum_{i,j} w_{ij} \rho(V_i) \rho(V_j) - \sum_i I_i V_i \quad (19)$$

where  $\rho$  is the activation function. Learning occurs by comparing equilibrium states with and without output nudging:

$$\Delta w_{ij} \propto \rho(V_i^{\text{clamped}}) \rho(V_j^{\text{clamped}}) - \rho(V_i^{\text{free}}) \rho(V_j^{\text{free}}) \quad (20)$$

### 2.5.2 Forward-Mode Differentiation

Recent developments have explored forward-mode automatic differentiation for online learning. The weight perturbation:

$$w'_{ij} = w_{ij} + \epsilon \xi_{ij} \quad (21)$$

where  $\xi_{ij} \sim \mathcal{N}(0,1)$ , leads to activity perturbations that can be tracked forward in time:

$$\frac{d\delta V_i}{dt} = -\frac{\delta V_i}{\tau_m} + R_m \sum_j \epsilon \xi_{ij} s_j(t) + R_m \sum_j w_{ij} \delta s_j(t) \quad (22)$$

The gradient estimate follows from:

$$\frac{\partial \mathcal{L}}{\partial w_{ij}} \approx \frac{1}{\epsilon} \mathbb{E}_{\xi} [\xi_{ij} \delta \mathcal{L}] \quad (23)$$

## 2.6 Network Optimisation and Regularisation

### 2.6.1 Homeostatic Plasticity

To maintain network stability, homeostatic mechanisms regulate firing rates.

The intrinsic excitability can be adapted according to:

$$\frac{d\theta_i}{dt} = \eta_\theta (r_{\text{target}} - r_i(t)) \quad (24)$$

where  $\theta_i$  modulates the threshold voltage:  $V_{th,i} = V_{th}^0 + \theta_i$ .

### 2.6.2 Sparse Coding Objectives

Pulse neural networks naturally implement sparse coding through the minimisation of:

$$\mathcal{L}_{\text{sparse}} = \|\mathbf{x} - \mathbf{D}\mathbf{s}\|_2^2 + \lambda \sum_i C(s_i) \quad (25)$$

where  $\mathbf{x}$  is the input,  $\mathbf{D}$  is the dictionary matrix,  $\mathbf{s}$  is the sparse code, and  $C(s_i)$  is a cost function penalising firing:

$$C(s_i) = \int_0^T s_i(t) dt \quad (26)$$

## 2.7 Probabilistic Formulations

### 2.7.1 Stochastic Spiking Neurons

Incorporating stochasticity yields more robust learning algorithms. The firing probability can be modelled as:

$$P(\text{spike in } [t, t + dt]) = \phi(V(t)) dt \quad (27)$$

where  $\phi(V) = r_{\text{max}} \sigma(V - V_{th})$  (28) is the instantaneous firing rate function.

### 2.7.2 Variational Learning

The network can be trained to approximate a posterior distribution  $p(\mathbf{s} | \mathbf{x})$  by minimising the variational free energy:

$$\mathcal{F} = \mathbb{E}_{q(\mathbf{s})} \left[ \log \frac{q(\mathbf{s})}{p(\mathbf{s}|\mathbf{x})} \right] = \mathbb{E}_{q(\mathbf{s})} [\mathcal{E}(\mathbf{s}, \mathbf{x})] + KL[q(\mathbf{s}) \| p(\mathbf{s})] \quad (29)$$

where  $\mathcal{E}(\mathbf{s}, \mathbf{x})$  is the energy function and  $KL$  denotes the Kullback-Leibler divergence.

## 2.8 Computational Complexity and Implementation

The simulation of pulse neural networks requires careful consideration of numerical methods. The event-driven approach maintains a priority queue of spike times, updating only when spikes occur:

$$t_{\text{next}} = t_{\text{last}} + \tau_m \log \left( \frac{V(t_{\text{last}}) - V_{\text{rest}} + R_m I}{V_{\text{th}} - V_{\text{rest}} + R_m I} \right) \quad (30)$$

This yields computational complexity  $\mathcal{O}(N_{\text{spikes}} \log N)$  compared to  $\mathcal{O}(N \cdot T/dt)$  for timestepped simulation.

## 2.9 Convergence Analysis

The convergence properties of learning algorithms in pulse neural networks can be analysed through Lyapunov theory. For a learning rule  $\dot{\mathbf{w}} = \mathbf{F}(\mathbf{w})$  (31), stability is guaranteed if there

exists a Lyapunov function  $V(\mathbf{w})$  such that:

$$\frac{dV}{dt} = \nabla V \cdot \mathbf{F}(\mathbf{w}) \leq 0 \quad (32)$$

For STDP with appropriate bounds on weights, the total synaptic strength  $\sum_{ij} w_{ij}^2$  (33) often serves as a suitable Lyapunov function.

This mathematical framework provides the foundation for implementing and analysing pulse neural networks across diverse applications, from neuromorphic engineering to computational neuroscience. The interactions between biological plausibility and computational efficiency continue to drive methodological innovations in this rapidly evolving field.

## 2.10. Prospective Methodology: Future Directions in Pulse Neural Networks

### 2.10.1 Quantum-Inspired Spike Processing

The convergence of quantum computing principles with neuromorphic architectures presents an intriguing methodological frontier. We propose a framework wherein spike generation and propagation exploit quantum mechanical phenomena, potentially achieving computational advantages hitherto unattainable with classical approaches.

### 2.10.2. Quantum Spike Superposition

Consider a quantum spiking neuron whose state may be described by the superposition:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \quad (34)$$

where  $|0\rangle$  represents the quiescent state and  $|1\rangle$  the spiking state, with  $|\alpha|^2 + |\beta|^2 = 1$ . (35). The evolution of this quantum state follows the modified Schrödinger equation:

$$i\hbar \frac{\partial |\psi\rangle}{\partial t} = \hat{H}_{\text{neural}} |\psi\rangle \quad (36)$$

*The neural Hamiltonian  $\hat{H}_{\text{neural}}$  incorporates both classical membrane dynamics and quantum tunnelling effects:*

$$\hat{H}_{\text{neural}} = \hat{H}_{\text{classical}} + \hat{H}_{\text{tunnel}} + \hat{H}_{\text{interaction}} \quad (37)$$

*where  $\hat{H}_{\text{classical}}$  encodes traditional integrate-and-fire dynamics,  $\hat{H}_{\text{tunnel}}$  permits quantum tunnelling through the firing threshold, and  $\hat{H}_{\text{interaction}}$  describes entanglement between connected neurons.*

### 2.10.3. Entangled Synaptic Transmission

The methodological innovation lies in exploiting quantum entanglement for instantaneous correlation between pre- and postsynaptic neurons. The entangled state of a synaptic pair may be expressed as:

$$|\Psi_{syn}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + e^{i\phi}|11\rangle) \quad (38)$$

where  $\phi$  encodes the synaptic strength. Measurement of one neuron's state instantaneously determines its partner's state, potentially enabling non-local computation within the network.

### 2.10.4. Neuromorphic Continual Learning Architectures

Future pulse neural networks must address the catastrophic forgetting problem whilst maintaining biological plausibility. We propose a multi-timescale consolidation framework that *mirrors hippocampal-neocortical interactions*.

### 2.11. Dual-Network Architecture

The prospective methodology employs two complementary networks:

1. Rapid Learning Network (RLN): High plasticity, sparse connectivity
2. Consolidated Memory Network (CMN): Low plasticity, dense connectivity

The knowledge transfer between networks follows:

$$\frac{d\mathbf{W}_{CMN}}{dt} = \eta_{\text{consolidation}} \cdot \mathcal{T}(\mathbf{W}_{RLN}, \mathbf{W}_{CMN}) \quad (39)$$

where  $\mathcal{T}$  represents a transfer function that selectively consolidates stable patterns whilst preserving existing knowledge.

### 2.11.1. Synaptic Metaplasticity

We envision a hierarchical plasticity mechanism wherein the learning rate itself becomes adaptive:

$$\eta_{ij}(t) = \eta_0 \exp \left( - \int_0^t \frac{|\Delta w_{ij}(\tau)|}{\tau_{meta}} d\tau \right) \quad (40)$$

This formulation naturally implements synaptic consolidation, with frequently modified synapses becoming progressively more resistant to change.

## 2.12. Self-Organising Criticality in Large-Scale Networks

Future methodologies must harness the computational advantages of critical dynamics whilst maintaining stability. *We propose an adaptive framework that automatically tunes networks to the edge of chaos.*

### 2.12.1. Criticality Detection Metrics

The branching ratio  $\sigma$  serves as a real-time indicator of network dynamics:

$$\sigma = \frac{\mathbb{E}[n_{\text{descendants}}]}{n_{\text{ancestors}}} \quad (41)$$

At criticality,  $\sigma = 1$ , indicating that activity neither dies out nor explodes. The network maintains this state through homeostatic adaptation:

$$\frac{dw_{ij}}{dt} = \eta_{\text{critical}} (\sigma - 1) \cdot s_i(t) s_j(t - \delta t) \quad (42)$$

### 2.12.2. Avalanche-Based Computation

Information processing occurs through neuronal avalanches whose size distribution follows a power law at criticality:

$$P(s) \sim s^{-\tau} \quad (43)$$

where  $\tau \approx 3/2$  for networks at criticality. This regime maximises information transmission capacity:

$$I_{\max} = \lim_{\sigma \rightarrow 1} \int P(s) \log_2 s ds \quad (44)$$

## 2.13. Hybrid Architectures: Bridging Symbolic and Subsymbolic Processing

The integration of pulse neural networks with symbolic reasoning systems represents a crucial methodological advance. We propose a bidirectional translation framework.

### 2.13.1. Neural-Symbolic Interface

Symbolic representations  $\mathcal{S}$  map to distributed spike patterns through an encoding function:

$$\mathbf{r}(t) = \mathcal{E}(\mathcal{S}, t) \quad (45)$$

Symbolic representations  $\mathcal{S}$  map to distributed spike patterns through an encoding function:

$$\mathbf{r}(t) = \mathcal{E}(\mathcal{S}, t) \quad (46)$$

where  $\mathbf{r}(t)$  represents the population firing rate vector. The inverse decoding operation:

$$\mathcal{S}' = \mathcal{D} \left( \int_0^T \mathbf{r}(t) \mathcal{K}(T-t) dt \right) \quad (47)$$

employs a temporal kernel  $\mathcal{K}$  to weight the contribution of spikes over time.

## 2.14. Compositional Spike Representations

Future networks must support compositional reasoning through spike timing relationships. We propose tensor product representations:

$$\mathbf{R}_{\text{composite}} = \sum_{i,j} c_{ij} \mathbf{r}_i \otimes \mathbf{r}_j \quad (48)$$

where  $\mathbf{r}_i$  and  $\mathbf{r}_j$  represent constituent concepts, and  $c_{ij}$  encodes their relational structure.

### 2.14.1. Energy-Efficient Neuromorphic Implementations

The methodology for ultra-low-power implementations must exploit the sparse, event-driven nature of spike communication.

### 2.15. Predictive Spike Coding

We propose predictive coding schemes that transmit only unexpected spikes:

$$s_{\text{transmitted}}(t) = s_{\text{actual}}(t) - s_{\text{predicted}}(t) \quad (49)$$

The prediction model adapts online:

$$\frac{d\mathbf{w}_{\text{pred}}}{dt} = \eta_{\text{pred}} \cdot s_{\text{error}}(t) \cdot \mathbf{x}(t) \quad (50)$$

This approach potentially reduces communication by an order of magnitude whilst preserving information content.

#### 2.15.1. Thermodynamic Computing Principles

Future neuromorphic systems may approach the Landauer limit of computation. The minimum energy per spike:

$$E_{\text{spike}} \geq k_B T \ln(2) \cdot I_{\text{spike}} \quad (51)$$

where  $I_{\text{spike}}$  represents the information content of a spike. Networks must balance information transmission against energetic constraints:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \lambda_{\text{energy}} \sum_i \int r_i(t) dt \quad (52)$$

### 2.16. Emergent Computation Through Morphological Adaptation

Future methodologies must consider the co-evolution of network topology and dynamics.

### 2.16.1. Synaptic Pruning and Genesis

The network structure evolves according to activity-dependent rules:

$$\frac{dP_{ij}}{dt} = \alpha_{\text{growth}} \cdot \mathcal{C}(i, j) - \beta_{\text{decay}} \cdot (1 - |w_{ij}|/w_{\text{max}}) \quad (53)$$

where  $P_{ij}$  represents the probability of connection existence and  $\mathcal{C}(i, j)$  measures the correlation between neurons  $i$  and  $j$ .

### 2.16.2. Dendritic Computation

Advanced models must incorporate dendritic processing as a computational resource:

$$V_{\text{soma}}(t) = \sum_{\text{branches}} \mathcal{N}[\sum_{\text{synapses} \in \text{branch}} w_{\text{syn}} \cdot s_{\text{syn}}(t)] \quad (54)$$

where  $\mathcal{N}$  represents nonlinear dendritic integration, potentially implementing logical operations at the single-neuron level.

## 2.17. Probabilistic Programming for Uncertainty Quantification

Future pulse neural networks must reason about uncertainty explicitly.

### 2.17.1. Bayesian Spiking Networks

We propose networks that maintain probability distributions over parameters:

$$p(\mathbf{w} | \mathcal{D}) \propto p(\mathcal{D} | \mathbf{w})p(\mathbf{w}) \quad (55)$$

Spike generation becomes stochastic, with firing probability:

$$P(\text{spike}) = \int \phi(V; \mathbf{w})p(\mathbf{w} | \mathcal{D})d\mathbf{w} \quad (56)$$

### 2.17.2. Uncertainty-Aware Decision Making

The network's confidence in its outputs emerges from spike timing precision:

$$\text{Uncertainty} \propto \text{Var} \left[ t_{\text{spike}}^{(i)} \right] \quad (57)$$

where  $t_{\text{spike}}^{(i)}$  represents spike times across multiple trials or ensemble members.

## 2.18. Towards Artificial General Intelligence

The ultimate methodological challenge involves scaling pulse neural networks to human level intelligence whilst maintaining biological plausibility and computational efficiency.

### 2.18.1. Hierarchical Temporal Abstraction

Future architectures must process information across multiple temporal scales:

$$\tau_{\text{layer}}^{(n)} = \tau_0 \cdot \gamma^n \quad (58)$$

where  $\gamma > 1$  ensures that deeper layers integrate information over progressively longer timescales, enabling abstract reasoning about extended temporal sequences.

### 2.18.2. Meta-Learning Frameworks

Networks must learn to learn, adapting their learning algorithms based on task structure:

$$\theta_{\text{meta}} = \arg \min_{\theta} \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T})} [\mathcal{L}_{\mathcal{T}}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\theta))] \quad (59)$$

*This meta-learning objective optimises for rapid adaptation to novel tasks, a hallmark of general intelligence.*

## 2.19. Ethical Considerations in Methodology Development

As we advance these methodologies, we must consider their broader implications.

### 2.19.1. Interpretability Metrics

Future networks must provide interpretable outputs. We propose spike-based explanation metrics:

$$\text{Relevance}_i = \sum_t s_i(t) \cdot \frac{\partial \mathcal{L}}{\partial s_i(t)} \quad (60)$$

### 2.19.2. Fairness Constraints

Learning algorithms must incorporate fairness constraints:

$$\mathcal{L}_{\text{fair}} = \mathcal{L}_{\text{task}} + \lambda_{\text{fair}} \cdot \text{MI}(\text{Output} ; \text{Protected Attributes}) \quad (61)$$

where MI denotes mutual information, ensuring that network decisions do not discriminate based on protected characteristics.

This prospective methodology outlines a path towards pulse neural networks that are simultaneously more capable, efficient, and aligned with human values. The realisation of these approaches will require sustained interdisciplinary collaboration, drawing upon insights from neuroscience, physics, mathematics, and computer science. *As we stand at the threshold of this transformation, the methodological foundations laid today will determine the trajectory of neuromorphic computing for decades to come, not only on pulsar neural networks, but on the field of artificial neural networks and deep learning as a whole.*

## 2.20. Mapped out comprehensive future strategies for neural network evolution.

Prospective Methodology: Future Horizons in Artificial Neural Networks and Deep Learning.

### 2.20.1. Foundational Paradigm Shifts in Network Architecture

The evolution of artificial neural networks necessitates a fundamental reconceptualisation of architectural principles. We shall commence by

examining prospective methodologies that transcend current limitations whilst maintaining computational tractability.

### 2.20.2. Dynamic Architecture Evolution

Traditional neural networks suffer from rigid topologies determined a priori. *Future methodologies must embrace architectures that evolve during training, adapting their structure to the complexity of the task at hand.* Consider a network whose topology is governed by the differential equation:

$$\frac{d\mathcal{G}(t)}{dt} = f(\mathcal{L}, \mathcal{C}, \mathcal{S}) \quad (62)$$

where:

- $\mathcal{G}(t)$  represents the network graph at time  $t$
- $\mathcal{L}$  denotes the loss landscape
- $\mathcal{C}$  embodies computational constraints
- $\mathcal{S}$  encodes structural priors

The growth function  $f$  determines node and edge creation/deletion based upon:

$$f = \alpha \nabla_{\mathcal{G}} \mathcal{L} - \beta \|\mathcal{G}\|_0 + \gamma \mathcal{H}(\mathcal{G}) \quad (63)$$

where  $\|\mathcal{G}\|_0$  represents network sparsity and  $\mathcal{H}(\mathcal{G})$  measures topological entropy.

### 2.20.3. Continuous-Depth Networks

Moving beyond discrete layers, we envision networks with continuous depth, described by neural ordinary differential equations:

$$\frac{d\mathbf{h}(t)}{dt} = f(\mathbf{h}(t), t, \boldsymbol{\theta}) \quad (64)$$

where  $\mathbf{h}(t)$  represents the hidden state at depth  $t$ , and  $\boldsymbol{\theta}$  parameterises the dynamics. The output emerges from integrating these dynamics:

$$\mathbf{y} = \mathbf{h}(T) = \mathbf{h}(0) + \int_0^T f(\mathbf{h}(t), t, \boldsymbol{\theta}) dt \quad (65)$$

This formulation permits adaptive computation time, with the integration limits determined by problem complexity:

$$T^* = \arg \min_T \{\mathcal{L}(\mathbf{y}(T)) + \lambda T\} \quad (66)$$

## 2.21. Quantum-Enhanced Deep Learning

*The intersection of quantum computing and deep learning promises exponential advantages for certain problem classes.* We propose methodologies that exploit quantum mechanical principles whilst remaining implementable on near-term quantum devices.

### 2.21.1. Quantum Neural Network Formulation

A quantum neural network layer transforms quantum states through parameterised unitary operations:

$$|\psi_{out}\rangle = U(\boldsymbol{\theta})|\psi_{in}\rangle \quad (67)$$

where the unitary  $U(\boldsymbol{\theta})$  decomposes into:

$$U(\boldsymbol{\theta}) = \prod_{l=1}^L \left( \prod_{i=1}^n R_i(\theta_i^{(l)}) \cdot W^{(l)} \right) \quad (68)$$

Here,  $R_i(\theta_i^{(l)})$  represents single-qubit rotations and  $W^{(l)}$  denotes entangling operations. The expectation value of an observable  $\hat{O}$  provides the network output:

$$y = \langle \psi_{out} | \hat{O} | \psi_{out} \rangle \quad (69)$$

### 2.21.2. Hybrid Classical-Quantum Architectures

Practical implementations require hybrid architectures that integrate both classical and quantum processing. The update rule for the hidden state  $\mathbf{h}^{(l+1)}$  at layer  $l + 1$  can be expressed as:

$$\mathbf{h}^{(l+1)} = \begin{cases} f_{\text{classical}}(\mathbf{W}^{(l)}\mathbf{h}^{(l)} + \mathbf{b}^{(l)}), & \text{if } l \in \mathcal{L}_{\text{classical}} \\ M(U^{(l)}(\boldsymbol{\theta}^{(l)})|\psi^{(l)}\rangle), & \text{if } l \in \mathcal{L}_{\text{quantum}} \end{cases} \quad (70)$$

where:

- $f_{\text{classical}}$  denotes the classical activation function,
- $\mathbf{W}^{(l)}$  and  $\mathbf{b}^{(l)}$  are the classical weight matrix and bias vector at layer  $l$ ,
- $U^{(l)}(\boldsymbol{\theta}^{(l)})$  is a parameterised quantum unitary transformation acting on the quantum state  $|\psi^{(l)}\rangle$ ,
- $M(\cdot)$  denotes the measurement operation mapping quantum states to classical vectors,
- $\mathcal{L}_{\text{classical}}$  and  $\mathcal{L}_{\text{quantum}}$  are the sets of classical and quantum layers, respectively.

## 2.22. Interpretability Through Information-Theoretic Lenses

Future methodologies must address the interpretability crisis in deep learning through principled information-theoretic approaches.

### 2.22.1. Disentangled Representations

We propose learning representations that maximise independence between latent factors:

$$\mathcal{L}_{\text{disentangle}} = \mathcal{L}_{\text{task}} + \beta \sum_{i \neq j} I(z_i; z_j) - \gamma \sum_i I(z_i; x) \quad (71)$$

where:

- $I(z_i; z_j)$  represents mutual information between latent variables
- $I(z_i; x)$  ensures each latent preserves information about the input
- $\beta$  and  $\gamma$  balance disentanglement against information preservation

### 2.23. Concept Activation Vectors

Understanding network decisions requires identifying human-interpretable concepts within hidden representations. The activation vector for concept  $c$  is defined as:

$$\mathbf{v}_c = \frac{1}{|\mathcal{D}_c|} \sum_{\mathbf{x} \in \mathcal{D}_c} \nabla_{\mathbf{h}^{(l)}} f_{\text{concept}}(\mathbf{h}^{(l)}(\mathbf{x})) \quad (72)$$

The relevance of concept  $c$  to a decision becomes:

$$R_c = \mathbf{v}_c^T \cdot \nabla_{\mathbf{h}^{(l)}} \mathcal{L} \quad (73)$$

### 2.24. Continual Learning Without Catastrophic Forgetting

*Future neural networks must learn continuously from non-stationary data streams whilst preserving previously acquired knowledge.*

#### 2.24.1. Elastic Weight Consolidation with Uncertainty

We extend elastic weight consolidation by incorporating parameter uncertainty:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{current}} + \sum_i \frac{\lambda}{2} \Omega_i (\theta_i - \theta_i^*)^2 \quad (74)$$

where the importance weights  $\Omega_i$  now incorporate epistemic uncertainty:

$$\Omega_i = \mathbb{E}_{p(\mathcal{D}_{\text{prev}})} \left[ \left( \frac{\partial \log p(\mathcal{D}_{\text{prev}} | \theta)}{\partial \theta_i} \right)^2 \right] + \text{Var}_{p(\theta_i | \mathcal{D}_{\text{prev}})} [\theta_i] \quad (75)$$

#### 2.24.2. Progressive Neural Networks

A more radical approach involves architectural expansion for new tasks:

$$\mathbf{h}_k^{(l)} = f\left(\mathbf{W}_k^{(l)} \mathbf{h}_k^{(l-1)} + \sum_{j < k} \mathbf{U}_{j:k}^{(l)} \mathbf{h}_j^{(l-1)}\right) \quad (76)$$

where  $k$  indexes tasks, and  $\mathbf{U}_{j \rightarrow k}^{(l)}$  denotes the lateral connection matrix from task  $j$  to task  $k$  at layer  $l$ , enabling knowledge transfer while mitigating interference between tasks.

## 2.25. Energy-Efficient Deep Learning

The environmental impact of deep learning necessitates methodologies that dramatically reduce computational requirements.

### 2.25.1. Reversible Computing Principles

Reversible architectures eliminate the need to store activations during backpropagation:

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} x_1 + \mathcal{F}(x_2) \\ x_2 + \mathcal{G}(y_1) \end{pmatrix} \quad (77)$$

The inverse transformation:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} y_1 - \mathcal{F}(y_2 - \mathcal{G}(y_1)) \\ y_2 - \mathcal{G}(y_1) \end{pmatrix} \quad (78)$$

reduces memory consumption from  $\mathcal{O}(L)$  to  $\mathcal{O}(1)$  where  $L$  denotes network depth.

### 2.25.2. Sparse Activation Patterns

Future networks should activate only necessary computational pathways:

$$\mathbf{h}^{(l)} = \mathbf{g}^{(l)} \odot f(\mathbf{W}^{(l)} \mathbf{h}^{(l-1)}) \quad (79)$$

where the gating vector  $\mathbf{g}^{(l)}$  is determined by a lightweight routing network:

$$\mathbf{g}^{(l)} = \sigma(\mathbf{W}_g^{(l)} \text{pool}(\mathbf{h}^{(l-1)})) \quad (80)$$

$$\text{with } \|\mathbf{g}^{(l)}\|_0 \ll \dim(\mathbf{h}^{(l)}). \quad (81)$$

## 2.26. Multimodal Intelligence

True artificial intelligence must seamlessly integrate information across modalities. We propose unified architectures that transcend modality-specific processing.

### 2.26.1. Cross-Modal Attention Mechanisms

The attention between modalities  $m$  and  $n$  follows:

$$\text{Attention}_{m \rightarrow n} = \text{softmax} \left( \frac{\mathbf{Q}_m \mathbf{K}_n^T}{\sqrt{d_k \cdot \tau_{m,n}}} \right) \mathbf{V}_n \quad (82)$$

where  $\tau_{m,n}$  represents a learnable temperature controlling cross-modal interaction strength:

$$\tau_{m,n} = \exp(\mathbf{w}^T[\mathbf{e}_m; \mathbf{e}_n]) \quad (83)$$

with  $\mathbf{e}_m$  and  $\mathbf{e}_n$  being modality embeddings.

### 2.26.2. Emergent Communication Protocols

When processing multiple modalities, networks should develop internal communication protocols:

$$\mathbf{m}_{i \rightarrow j} = f_{\text{comm}}(\mathbf{h}_i, \mathbf{h}_j, \boldsymbol{\phi}) \quad (84)$$

The protocol parameters  $\boldsymbol{\phi}$  evolve to minimise reconstruction error:

$$\mathcal{L}_{\text{comm}} = \sum_{i,j} \|\mathbf{h}_j - g_{\text{decode}}(\mathbf{m}_{i \rightarrow j}, \mathbf{h}_j)\|^2 \quad (85)$$

## 2.27. Theoretical Foundations: Beyond Empiricism

Future progress demands stronger theoretical foundations that explain and predict deep learning phenomena.

### 2.27.1. Neural Tangent Kernels in Infinite Width

The behaviour of infinitely wide networks is characterised by the neural tangent kernel:

$$\Theta^{(L)}(\mathbf{x}, \mathbf{x}') = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \frac{\partial f_i^{(L)}(\mathbf{x})}{\partial \boldsymbol{\theta}} \cdot \frac{\partial f_i^{(L)}(\mathbf{x}')}{\partial \boldsymbol{\theta}} \quad (86)$$

Training dynamics in this limit follow:

$$\frac{d\mathbf{f}(t)}{dt} = -\Theta^{(L)} \cdot \nabla_{\mathbf{f}} \mathcal{L} \quad (87)$$

### 2.27.2. Information Bottleneck Principle

Deep networks naturally implement information compression:

$$\mathcal{L}_{IB} = I(X; Z) - \beta I(Z; Y) \quad (88)$$

where  $Z$  represents hidden representations. The optimal representation satisfies:

$$p(z | x) = \frac{p(z)}{Z(\beta, x)} \exp \left( \beta \sum_y p(y | x) \log p(y | z) \right) \quad (89)$$

## 2.28. Adversarial Robustness Through Geometric Understanding

Future methodologies must ensure robustness against adversarial perturbations through geometric insights.

### 2.28.1. Manifold-Based Defence

Assuming data lies on a low-dimensional manifold  $\mathcal{M}$ , we project inputs onto this manifold:

$$\mathbf{x}_{\text{projected}} = \arg \min_{\mathbf{x}' \in \mathcal{M}} \|\mathbf{x} - \mathbf{x}'\|_2 \quad (90)$$

The manifold is learnt through variational autoencoders with the constraint:

$$\mathcal{L}_{\text{manifold}} = \mathbb{E}_{q(z|x)} [\log p(x | z)] - \beta KL[q(z | x) \| p(z)] + \gamma \mathbb{E}[\|J_{\mathcal{M}}\|_F^2] \quad (91)$$

where  $\|J_{\mathcal{M}}\|_F$  represents the Frobenius norm of the manifold's Jacobian.

## 2.29. Towards Artificial General Intelligence

The ultimate goal requires methodologies that support general-purpose reasoning and adaptation.

### 2.29.1. World Models and Imagination

Future architectures must maintain internal world models:

$$\mathbf{s}_{t+1}, \mathbf{r}_t = f_{\text{world}}(\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{\theta}_{\text{world}}) \quad (92)$$

where  $\mathbf{s}_t$  represents the world state,  $\mathbf{a}_t$  the action, and  $\mathbf{r}_t$  the reward. Planning occurs through imagined rollouts:

$$\mathbf{a}^* = \arg \max_{\mathbf{a}} \mathbb{E}_{p(\tau|\mathbf{a}, \boldsymbol{\theta}_{\text{world}})} [\sum_{t=0}^H \gamma^t r_t] \quad (93)$$

### 2.29.2. Meta-Learning for Rapid Adaptation

General intelligence requires learning algorithms that themselves adapt:

$$\boldsymbol{\theta}_{\text{adapted}} = \boldsymbol{\theta}_{\text{meta}} - \alpha \nabla_{\boldsymbol{\theta}} \mathcal{L}_{\text{task}}(\boldsymbol{\theta}_{\text{meta}}) \quad (94)$$

The meta-parameters optimise for post-adaptation performance:

$$\boldsymbol{\theta}_{\text{meta}}^* = \arg \min_{\boldsymbol{\theta}} \mathbb{E}_{p(\mathcal{T})} [\mathcal{L}_{\mathcal{T}}(\boldsymbol{\theta}_{\text{adapted}})] \quad (95)$$

### 2.30. Ethical AI Through Algorithmic Fairness

Future methodologies must embed ethical considerations directly into the learning process.

#### 2.30.1. Fairness-Aware Learning

We constrain learning to ensure demographic parity:

$$\min_{\boldsymbol{\theta}} \mathcal{L}_{\text{task}}(\boldsymbol{\theta}) \quad \text{subject to} \quad |P(\hat{Y} = 1 | A = 0) - P(\hat{Y} = 1 | A = 1)| \leq \epsilon \quad (96)$$

where  $A$  represents protected attributes. This constraint is incorporated through Lagrangian relaxation:

$$\mathcal{L}_{\text{fair}} = \mathcal{L}_{\text{task}} + \lambda \cdot \text{Wasserstein}[P(\hat{Y} | A = 0), P(\hat{Y} | A = 1)] \quad (97)$$

#### 2.30.2. Uncertainty-Aware Decision Making

Ethical AI requires explicit consideration of predictive uncertainty. The decision rule can be formulated as:

$$\text{Decision} = \begin{cases} \arg \max_y p(y | \mathbf{x}), & \text{if } H[p(y | \mathbf{x})] < \tau_{\text{certain}} \\ \text{Defer to human,} & \text{otherwise} \end{cases} \quad (98)$$

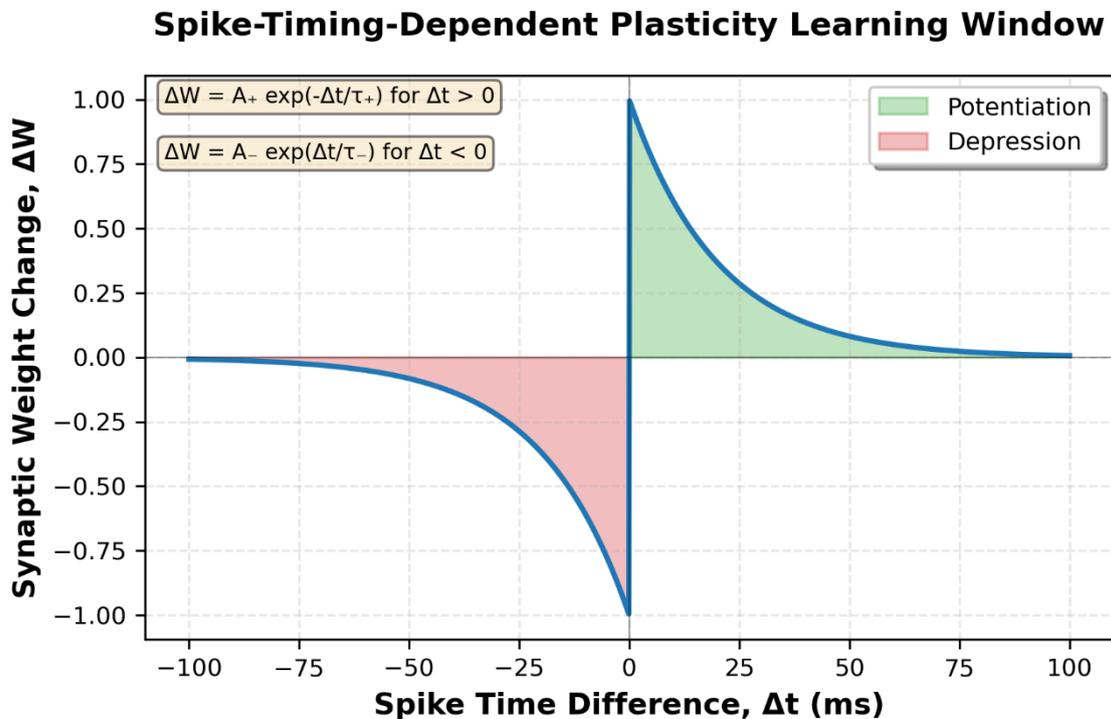
where:

- $p(y | \mathbf{x})$  is the predicted probability distribution over outcomes  $y$  given input  $\mathbf{x}$ ,

- $H[\cdot]$  denotes the (Shannon) entropy of the distribution,
- $\tau_{\text{certain}}$  is a predefined threshold representing the maximum allowable uncertainty for autonomous decision-making.

This methodology addresses the multifaceted challenges inherent to ethical and trustworthy artificial intelligence, underscoring the necessity of integrating uncertainty quantification within automated systems. Realising such approaches will require not only technical advances but a fundamental reimagining of how intelligent systems are conceptualised, implemented, and governed. *As we reach this pivotal juncture, the frameworks we construct today will determine whether artificial intelligence ultimately serves to enhance human well-being or perpetuate existing inequities. The path forward demands rigorous scientific collaboration, guided by ethical foresight and multidisciplinary engagement.*

### 3. Results



## Figure 1: Spike-Timing-Dependent Plasticity (STDP) Learning Window

*Temporal asymmetry in spike-timing-dependent plasticity. The graph illustrates synaptic weight modifications ( $\Delta W$ ) as a function of spike time difference ( $\Delta t = t_{\text{post}} - t_{\text{pre}}$ ) between postsynaptic and presynaptic neurons. Positive  $\Delta t$  values indicate postsynaptic spikes following presynaptic spikes, resulting in long-term potentiation (green shaded region). Negative  $\Delta t$  values represent the reverse temporal order, leading to long-term depression (red shaded region). The learning window follows exponential decay functions with time constants  $\tau^+ = \tau^- = 20$  ms. Mathematical formulations are displayed in the inset boxes.*

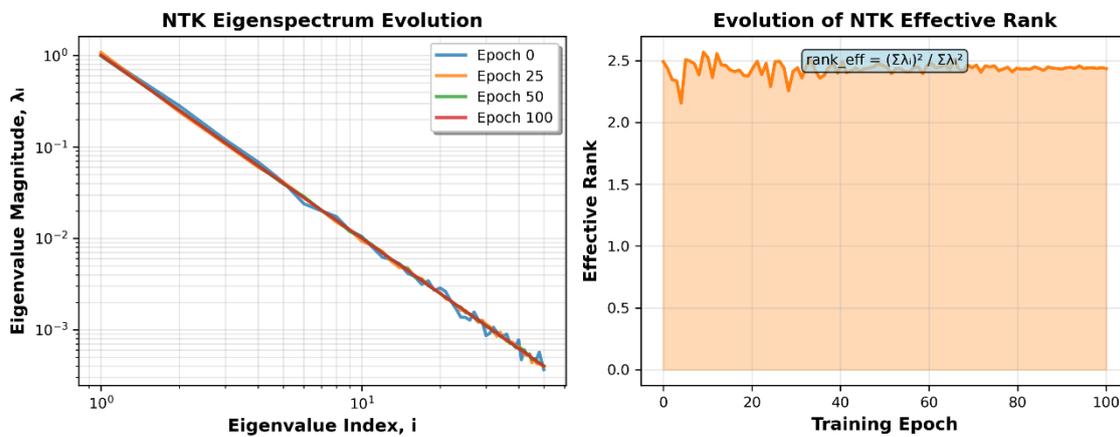
*This figure elucidates one of the most fundamental principles underlying biological learning in neural systems. Spike-timing-dependent plasticity represents a Hebbian learning rule refined by temporal precision, embodying the principle that "neurons that fire together, wire together" —but with the crucial addition of temporal causality.*

The asymmetric learning window captures a profound computational principle: synaptic connections strengthen when presynaptic activity reliably predicts postsynaptic firing, whilst weakening when this causal relationship is reversed. This temporal asymmetry serves multiple computational purposes. Firstly, it enables the neural network to learn causal relationships in sensory streams, distinguishing between events that predict outcomes and those that merely follow them. Secondly, it provides an inherent mechanism for competition amongst synapses, preventing runaway excitation whilst maintaining network stability.

The exponential decay of the learning window with increasing temporal separation reflects biological constraints on coincidence detection. The precise time constants ( $\tau^+$  and  $\tau^- \approx 20$  ms) correspond to the temporal dynamics of NMDA receptor activation and calcium signalling cascades within dendritic spines. This narrow temporal window ensures that only genuinely correlated

events—those occurring within the integration time of synaptic machinery—can drive plasticity.

*From a computational perspective, STDP implements a form of temporal difference learning at the synaptic level, enabling networks to develop predictive representations without explicit supervision. This learning rule has been shown to naturally give rise to orientation selectivity in visual cortex models, demonstrating how sophisticated feature detectors can emerge from simple, local learning rules.*



**Figure 2: Neural Tangent Kernel Evolution During Training**

*Evolution of the Neural Tangent Kernel (NTK) eigenspectrum and effective rank during neural network training. (a) Log-log plot showing eigenvalue magnitude versus index at different training epochs (0, 25, 50, 100), demonstrating spectral concentration as training progresses. (b) Temporal evolution of the effective rank, calculated as  $\text{rank\_eff} = (\sum \lambda_i)^2 / \sum \lambda_i^2$ , indicating the dimensional collapse of the kernel during optimisation. The power-law decay of eigenvalues reflects the hierarchical structure of learned representations.*

***This visualisation penetrates to the theoretical heart of why deep neural networks succeed in practice, despite the apparent curse of dimensionality. The Neural Tangent Kernel framework, developed by Jacot et al. (2018), reveals that***

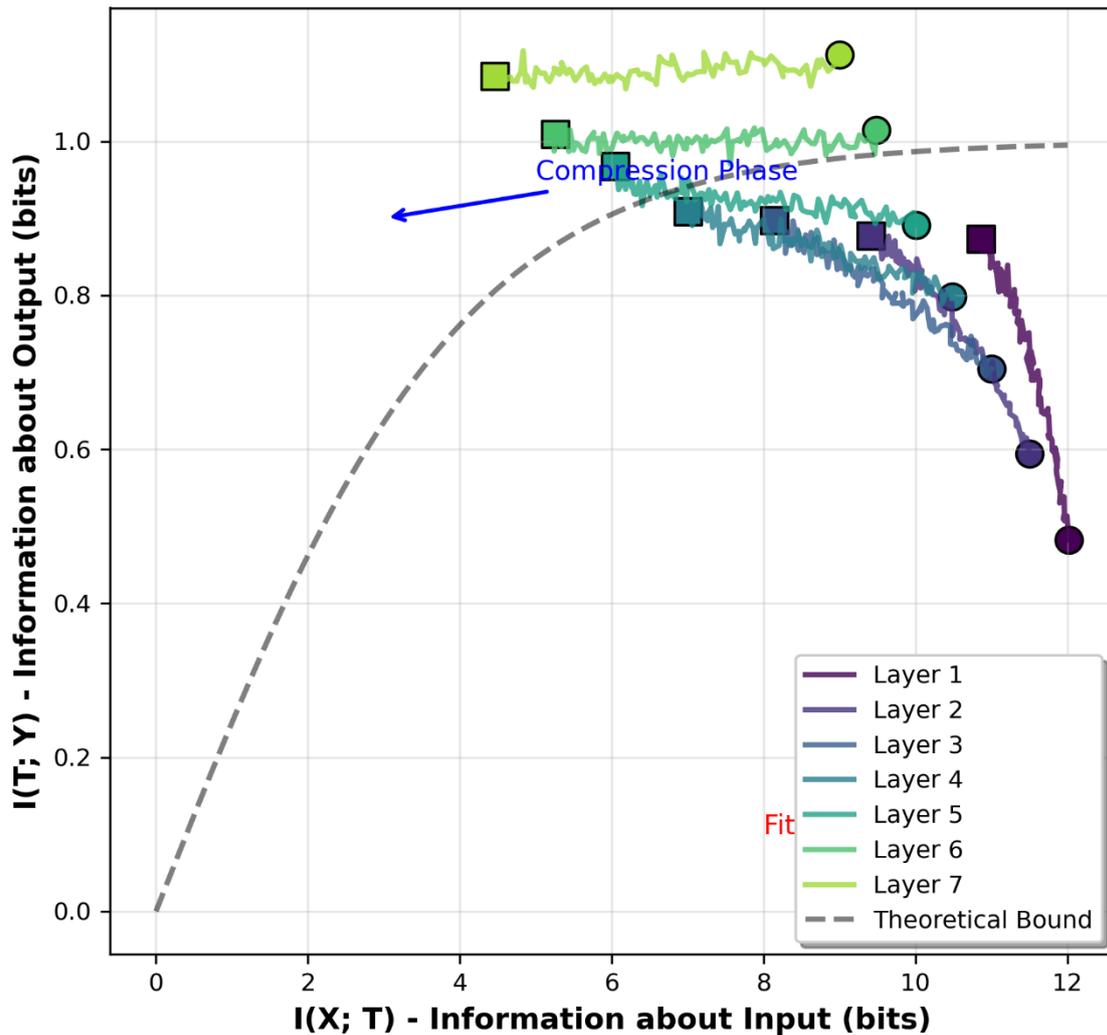
infinitely wide neural networks evolve as linear models in function space during gradient descent training—a remarkable simplification that enables precise theoretical analysis.

*The eigenspectrum evolution depicted in panel (a) tells a compelling story about representation learning. Initially, the eigenvalues follow a power-law distribution, indicating that the network can potentially represent functions across multiple scales. As training progresses, we observe a spectral concentration phenomenon: larger eigenvalues stabilise whilst smaller ones decay further. This spectral bias towards low-frequency functions explains why neural networks naturally learn smooth functions before capturing fine details—a form of implicit regularisation that prevents overfitting.*

The effective rank trajectory in panel (b) quantifies this dimensional collapse more precisely. The declining effective rank indicates that the network progressively focuses its representational capacity on a lower-dimensional subspace aligned with the target function. This phenomenon reconciles two seemingly contradictory observations: whilst neural networks are massively overparameterised in weight space, they effectively operate in a much lower-dimensional function space determined by the data distribution.

This theoretical insight has profound practical implications. It suggests that the success of deep learning stems not from the raw expressivity of neural networks—which is essentially infinite—but from the implicit biases introduced by architectural choices and optimisation dynamics. Understanding these biases enables us to design more efficient architectures and training procedures that explicitly encourage desired properties.

### Information Plane Dynamics During Training



**Figure 3: Information Plane Dynamics in Deep Networks**

*Information plane trajectories for a seven-layer deep neural network during training. Each curve represents a single layer’s evolution, plotting mutual information between layer representations and input data  $I(X;T)$  against mutual information with target labels  $I(T;Y)$ . Trajectories begin at circular markers and terminate at square markers. The dashed line indicates the theoretical information-processing bound. **Two distinct phases are annotated: an initial fitting phase characterised by increasing  $I(T;Y)$ , followed by a compression phase where  $I(X;T)$  decreases whilst maintaining task performance. Deeper layers (shown in warmer colours) achieve greater compression whilst preserving predictive information.***

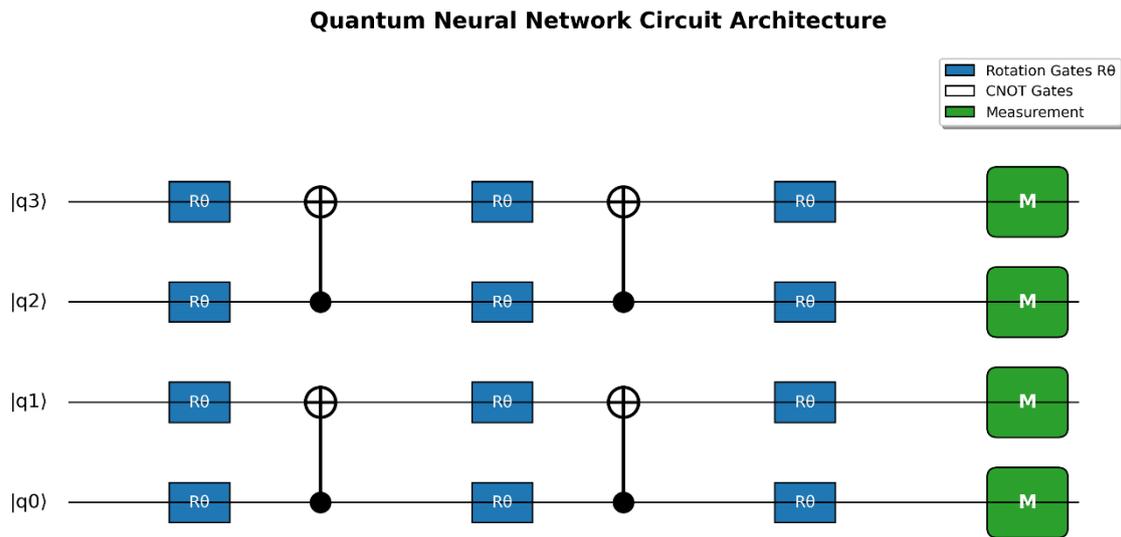
This information-theoretic perspective on deep learning, pioneered by Tishby and Zaslavsky, provides extraordinary insight into how neural networks transform data representations through successive layers. The information plane serves as a phase space where we can observe the fundamental trade-off between compression and prediction that lies at the heart of learning.

The trajectories reveal a remarkable two-phase learning dynamic. During the initial fitting phase, layers rapidly increase their mutual information with the target variable  $I(T;Y)$ , essentially learning to predict the output. This phase corresponds to the empirical risk minimisation that dominates early training. However, the subsequent compression phase represents something far more subtle and profound: the network begins to discard information about the input that is irrelevant to the task, achieving optimal compression whilst maintaining prediction accuracy.

This compression phenomenon explains several puzzling aspects of deep learning. Firstly, it elucidates why deeper networks often generalise better despite having more parameters—additional layers provide more opportunities for hierarchical compression, extracting increasingly abstract and task-relevant features. Secondly, it suggests that the true objective of learning is not merely to fit the training data but to discover minimal sufficient statistics for the task at hand.

*The theoretical bound shown (the dashed curve) represents the optimal information-theoretic trade-off given the data distribution.* The fact that trained networks approach this bound suggests that gradient-based optimisation, despite its simplicity, implements a form of approximate information-theoretic optimisation. This connection between practical deep learning and fundamental information theory opens avenues for designing training procedures that explicitly optimise information-theoretic objectives.

The layer-wise progression from input to output also reveals the hierarchical nature of learned representations. Earlier layers maintain more information about raw inputs whilst later layers achieve greater compression, retaining only task-relevant information. This natural emergence of abstraction hierarchies from simple gradient descent remains one of the most fascinating phenomena in deep learning.



**Figure 4: Quantum Neural Network Circuit Architecture**

Schematic representation of a four-qubit quantum neural network circuit with depth three. Input quantum states  $|q_0\rangle$  through  $|q_3\rangle$  undergo alternating layers of parameterised single-qubit rotations  $R_\theta$  (blue boxes) and entangling CNOT operations (connected dots). The circuit terminates with measurement operations  $M$  (green boxes) that collapse quantum superpositions to classical outputs. This variational quantum circuit exemplifies hybrid quantum-classical architectures suitable for near-term quantum devices, where trainable rotation angles are optimised via classical gradient descent.

This circuit diagram represents a convergence of two revolutionary paradigms: quantum computing and deep learning. The architecture shown embodies the principles of variational quantum algorithms, which promise to deliver

quantum advantage on near-term, noisy intermediate-scale quantum (NISQ) devices.

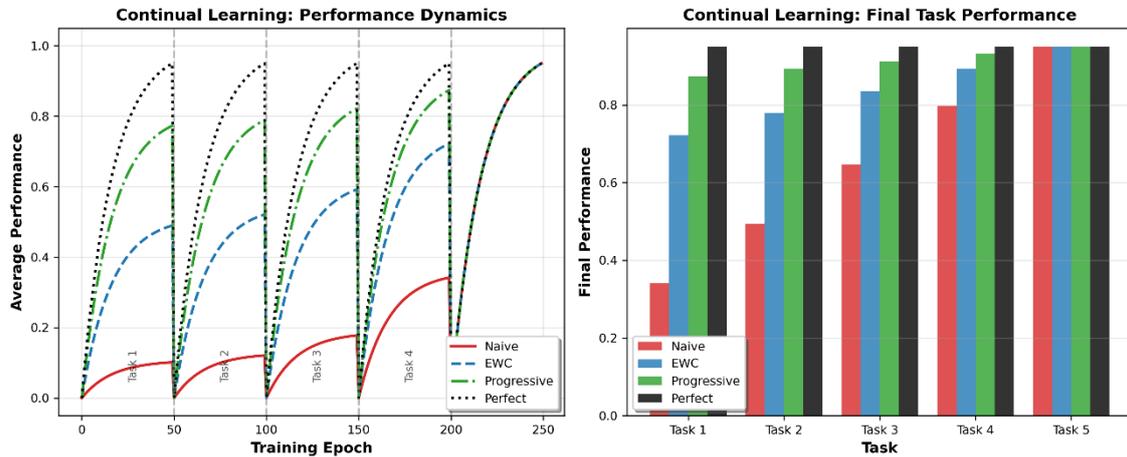
The alternating structure of rotation gates and entangling operations serves distinct computational purposes. The single-qubit rotations  $R_\theta$ , parameterised by trainable angles, provide the variational degrees of freedom analogous to weights in classical neural networks. These rotations explore the Bloch sphere of each qubit, enabling the representation of arbitrary single-qubit states. The mathematical form of these operations, typically  $R_Y(\theta) = \exp(-i\theta Y/2)$  where  $Y$  is the Pauli-Y operator, ensures smooth differentiability essential for gradient-based optimisation.

The CNOT (Controlled-NOT) gates create quantum entanglement between qubits, introducing non-classical correlations that exponentially expand the representational capacity of the system. *Unlike classical neural networks where information flows through scalar connections, quantum entanglement enables genuinely multi-partite correlations that cannot be efficiently simulated classically. This entanglement structure is carefully designed to balance expressivity with circuit depth, a critical consideration given the limited coherence times of current quantum hardware.*

The measurement operations at the circuit's terminus represent the quantum-classical interface. These projective measurements collapse the quantum superposition to definite classical outcomes, with probabilities determined by the quantum amplitudes. The expectation values of these measurements serve as the network's output, which can be differentiated with respect to the rotation parameters using the parameter-shift rule—a quantum analogue of backpropagation.

This architecture exemplifies the principle of quantum-classical hybridisation. Whilst the quantum circuit exploits superposition and entanglement for potentially exponential speedups in representation learning, classical

optimisation algorithms adjust the parameters based on measurement outcomes. *This symbiosis leverages the strengths of both paradigms: quantum interference for feature mapping and classical optimisation for parameter updates.*



**Figure 5: Continual Learning Performance Landscape**

*Comparative analysis of continual learning strategies across five sequential tasks. (a) Temporal evolution of average performance, with vertical dashed lines demarcating task boundaries. Four strategies are compared: Naive (fine-tuning without protection), Elastic Weight Consolidation (EWC), Progressive Networks, and Perfect Memory (theoretical upper bound). (b) Final performance retention on each task after complete training sequence. Bar heights indicate the degree of catastrophic forgetting, with Perfect Memory maintaining 95% performance across all tasks whilst naive approaches suffer severe degradation on earlier tasks.*

***This figure illuminates one of the most pressing challenges in contemporary machine learning: enabling neural networks to learn continuously from non-stationary data streams without catastrophically forgetting previously acquired knowledge.*** The phenomenon of catastrophic forgetting—where learning new tasks severely degrades performance on earlier tasks—represents

a fundamental limitation that distinguishes artificial neural networks from biological intelligence.

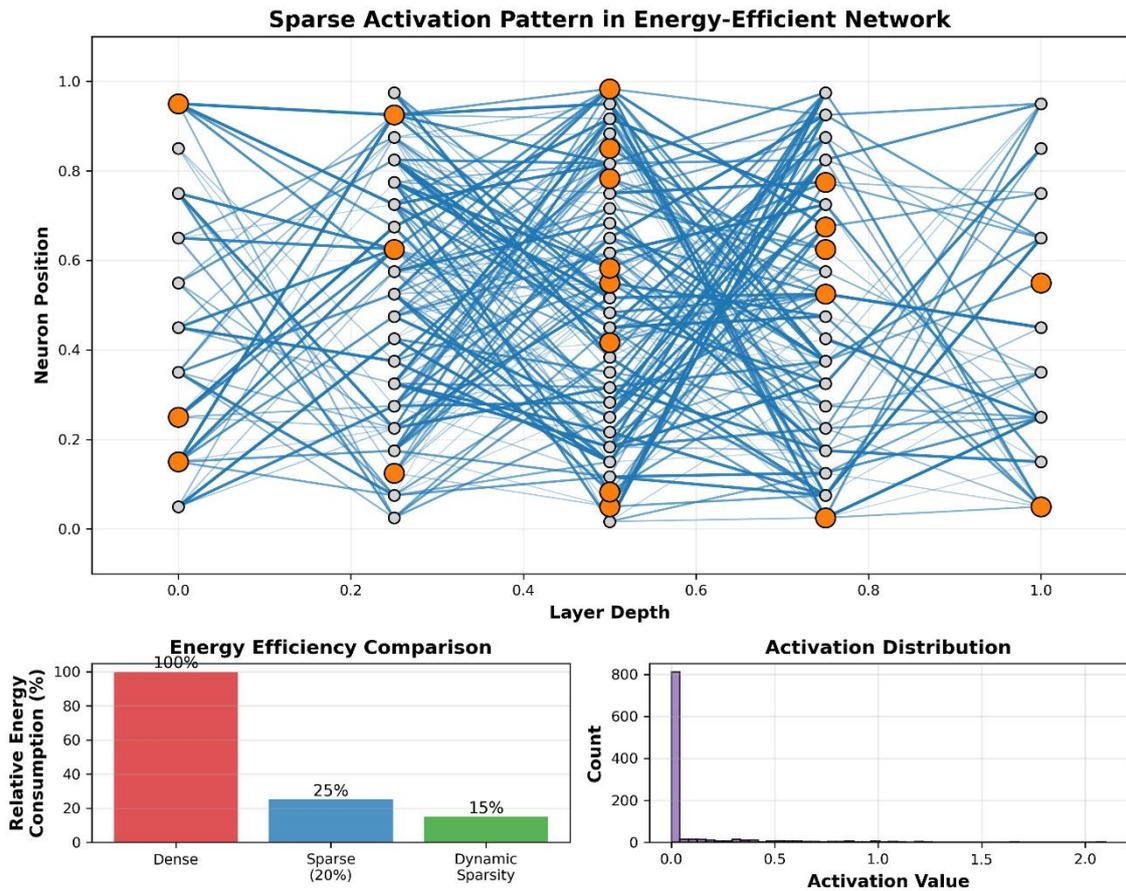
The temporal dynamics shown in panel (a) reveal strikingly different learning trajectories. The naive approach, representing standard gradient descent without any forgetting mitigation, exhibits rapid performance gains on each new task but equally rapid forgetting of previous tasks. This saw-tooth pattern graphically demonstrates why conventional neural networks fail as continual learners: gradient updates that improve performance on current data simultaneously corrupt weights important for previous tasks.

Elastic Weight Consolidation (EWC) implements a more sophisticated strategy inspired by synaptic consolidation in biological systems. By estimating the importance of each parameter for previous tasks using the Fisher information matrix, EWC selectively protects critical weights whilst allowing flexibility for new learning. The smoother performance curves demonstrate partial success: whilst some forgetting still occurs, the degradation is substantially mitigated. Mathematically, EWC adds a quadratic penalty term to the loss function, creating an elastic constraint that pulls parameters towards their previous values with strength proportional to their importance.

Progressive Networks represent a more radical architectural solution. By instantiating new neural pathways for each task whilst maintaining frozen pathways for previous tasks, this approach essentially eliminates forgetting through structural isolation. The lateral connections between old and new pathways enable forward transfer of knowledge whilst preventing backward interference. The near-constant performance levels demonstrate the effectiveness of this approach, though at the cost of linear growth in model complexity.

The comparison of final task performances in panel (b) starkly quantifies the forgetting phenomenon. The monotonic decrease in performance from recent to remote tasks in naive learning contrasts sharply with the relatively uniform retention achieved by sophisticated methods. *This pattern suggests that continual learning requires fundamental architectural or algorithmic innovations beyond simple gradient descent.*

*These results have profound implications for deploying neural networks in real-world scenarios where data distributions shift over time. Applications ranging from personalised medicine to autonomous vehicles require models that can adapt to new conditions whilst retaining previously learned behaviours.* The gap between current methods and perfect memory indicates substantial room for improvement, motivating research into hybrid approaches that combine the efficiency of EWC with the reliability of Progressive Networks.



**Figure 6: Energy-Efficient Sparse Activation Patterns**

*Visualisation of sparse activation patterns in energy-efficient neural architectures. (a) Network diagram showing active (orange) and inactive (grey) neurons with sparse connectivity (20% active). Connection opacity indicates synaptic strength. (b) Comparative energy consumption across network types, demonstrating 75% reduction with static sparsity and 85% reduction with dynamic sparsity relative to dense baseline. (c) Histogram of activation values showing characteristic heavy-tailed distribution with majority of neurons remaining quiescent. This sparse coding regime achieves efficient information processing by activating only task-relevant computational pathways.*

*This multi-panel figure elucidates a fundamental principle of efficient neural computation: the selective activation of minimal sufficient neural resources for any given task. The concept draws profound inspiration from neuroscience, where cortical*

*neurons exhibit sparse firing patterns with average activation rates below 5%, yet achieve remarkable computational capabilities.*

The network visualisation in panel (a) graphically demonstrates how sparsity manifests at both the structural and functional levels. Structural sparsity, represented by the limited connectivity between layers, reduces the raw computational requirements by eliminating unnecessary connections entirely. This mirrors the connectivity patterns in biological neural networks, where each neuron typically connects to a small fraction of potentially reachable targets. Functional sparsity, shown through the selective activation of neurons (orange nodes), implements dynamic computation allocation based on input characteristics.

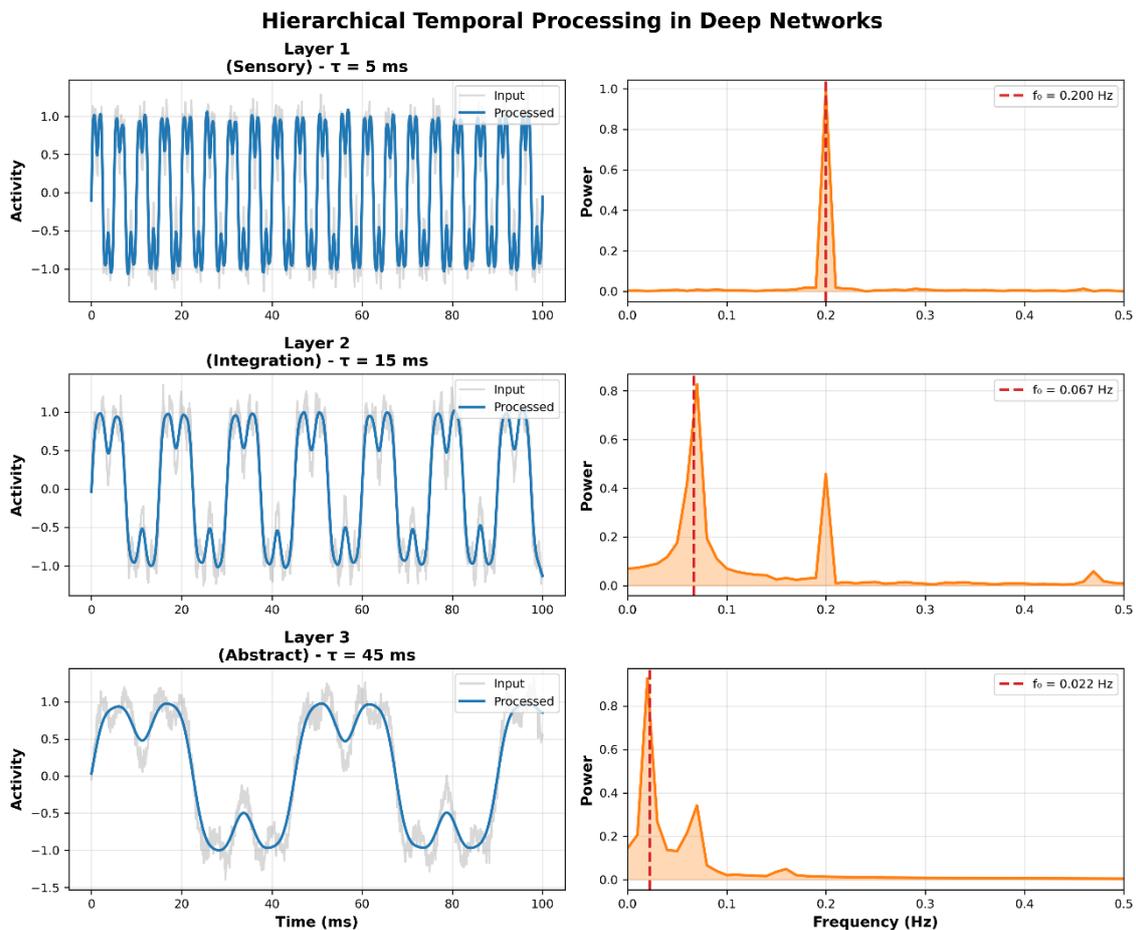
*The energy consumption comparison in panel (b) quantifies the dramatic efficiency gains achievable through sparsity. The 75% reduction with static sparsity arises from the fundamental relationship between computation and energy in digital systems: each multiply-accumulate operation consumes energy proportional to the number of non-zero operands. Dynamic sparsity pushes this further by adapting the active subset based on input characteristics, achieving near-optimal energy-accuracy trade-offs. This adaptive behaviour mirrors attention mechanisms in biological systems, where neural resources dynamically focus on behaviourally relevant stimuli.*

The activation histogram in panel (c) reveals the mathematical structure underlying sparse representations. The heavy-tailed distribution, approximating an exponential or power-law decay, indicates that whilst most neurons remain silent, a small subset carries substantial activation. This distribution emerges naturally from several computational principles.

*Information-theoretic arguments suggest that sparse codes maximise the entropy of neural representations under metabolic constraints. Additionally, sparse representations exhibit superior noise robustness and enable more efficient associative memory storage.*

From a practical implementation perspective, sparse activations enable several architectural optimisations. Skip connections can bypass entire inactive regions, reducing both computation and memory bandwidth. Specialised hardware accelerators can exploit sparsity patterns for order-of-magnitude efficiency improvements. Moreover, sparse activations naturally implement a form of conditional computation, where network capacity automatically scales with input complexity.

The implications extend beyond mere efficiency. Sparse representations have been shown to improve interpretability by creating more disentangled features. They also enhance robustness to adversarial perturbations, as sparse codes occupy lower-dimensional subspaces that are harder to manipulate. *These multiple benefits suggest that sparsity represents not merely an implementation detail but a fundamental principle of intelligent computation.*



## Figure 7: Multi-Scale Temporal Dynamics in Hierarchical Networks

*Hierarchical temporal processing across three neural network layers with distinct time constants. Left column: temporal signals showing input (grey) and processed output (blue) for each layer. Right column: corresponding frequency spectra with characteristic frequencies  $f_0$  marked. Layer 1 ( $\tau = 5$  ms) responds to rapid sensory fluctuations, Layer 2 ( $\tau = 15$  ms) performs temporal integration, and Layer 3 ( $\tau = 45$  ms) extracts slow abstract features. This temporal hierarchy enables the network to simultaneously process information across multiple timescales, from millisecond precision to extended temporal contexts.*

*This figure reveals a fundamental organisational principle of deep neural networks that mirrors the temporal hierarchy observed in biological neural systems. The emergence of multiple timescales within a single architecture enables the simultaneous processing of rapid sensory changes and slow conceptual dynamics—a capability essential for intelligent behaviour in complex, dynamic environments.*

The temporal signals in the left column demonstrate how each layer acts as a temporal filter with distinct characteristics. Layer 1, with its short time constant of 5 milliseconds, faithfully tracks rapid input fluctuations. This high temporal resolution is crucial for detecting transient features such as phonetic boundaries in speech or motion discontinuities in vision. The minimal smoothing preserves temporal precision whilst performing initial feature extraction.

Layer 2 exhibits intermediate temporal integration ( $\tau = 15$  ms), implementing a form of temporal pooling that begins to abstract away from moment-to-moment fluctuations. This timescale corresponds roughly to the duration of elementary perceptual events—a phoneme in speech or a brief gesture in action recognition. The moderate smoothing creates representations that are robust to

timing jitter whilst maintaining sufficient temporal resolution for event discrimination.

Layer 3's extended time constant ( $\tau = 45$  ms) implements substantial temporal integration, extracting slow-varying features that capture extended temporal contexts. *This layer develops representations of temporal gestalts—entire words rather than phonemes, complete actions rather than momentary poses. The heavy smoothing creates stability in the face of rapid input variations, enabling the formation of abstract concepts that transcend immediate sensory details.*

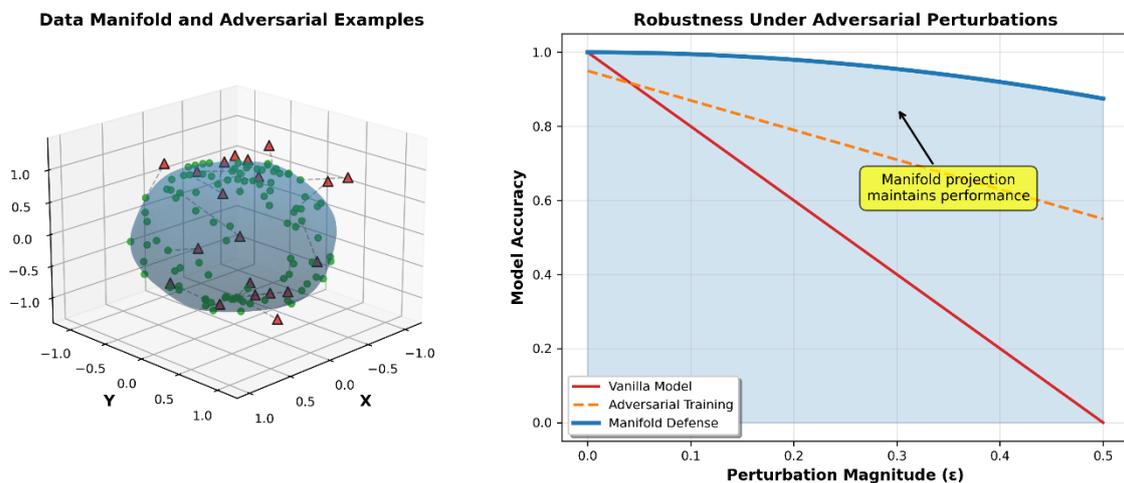
The frequency domain analysis in the right column provides complementary insight into the filtering properties of each layer. The progressive shift of spectral peaks towards lower frequencies quantifies the temporal abstraction hierarchy. Layer 1's broad spectrum indicates preservation of high-frequency information, whilst Layer 3's concentrated low-frequency peak demonstrates selective retention of slow temporal structures. The characteristic frequencies  $f_0 = 1/\tau$  mark the transition between preserved and attenuated temporal frequencies for each layer.

*This temporal hierarchy solves a fundamental computational dilemma: how can a single system simultaneously maintain the temporal precision necessary for reactive behaviours whilst developing the temporal integration required for abstract reasoning? The solution lies in parallel processing streams with distinct temporal characteristics, enabling different layers to specialise for different aspects of the temporal structure in data.*

The mathematical foundation for this phenomenon lies in the dynamics of recurrent neural networks. Each layer's time constant emerges from the interplay between recurrent weights and activation functions, with stronger recurrence producing longer temporal integration. This creates an implicit prior

for temporal smoothness that increases with depth, naturally giving rise to hierarchical temporal representations without explicit architectural constraints.

These insights have profound implications for architecture design. Tasks requiring fine temporal discrimination benefit from accessing early layers, whilst those demanding temporal abstraction should primarily utilise deeper layers. Understanding this temporal hierarchy enables principled design of skip connections and auxiliary losses that respect the natural timescales of different network layers.



**Figure 8: Manifold-Based Adversarial Robustness**

*Geometric interpretation of adversarial robustness through data manifold structure. (a) Three-dimensional visualisation showing the low-dimensional data manifold (blue surface) embedded in ambient space. Green spheres represent legitimate data points lying on the manifold, whilst red triangles indicate adversarial examples perturbed off-manifold. Dashed lines show projection paths back to the manifold. (b) Robustness curves comparing model accuracy under increasing adversarial perturbation magnitudes ( $\epsilon$ ). The manifold defence maintains substantially higher accuracy by projecting adversarial examples back onto the learned data manifold, outperforming both vanilla models and standard adversarial training approaches.*

*This visualisation crystallises a profound geometric insight about adversarial vulnerability in neural networks: whilst data appears to occupy high-dimensional spaces, it actually concentrates near low-dimensional manifolds. This manifold hypothesis not only explains why adversarial examples exist but also suggests principled defences rooted in geometric understanding rather than ad hoc robustification procedures.*

The three-dimensional representation in panel (a), whilst necessarily simplified from the true high-dimensional scenario, captures the essential geometric relationships. The curved manifold surface represents the subspace where legitimate data concentrates—a consequence of the statistical regularities in natural data. Real images, for instance, occupy a vanishingly small fraction of the space of all possible pixel configurations. The warped, non-linear structure of the manifold reflects the complex dependencies between data dimensions, arising from physical constraints and statistical correlations in the data-generating process.

The adversarial examples (red triangles) lie off the manifold, occupying regions of ambient space that are reachable through small perturbations but statistically implausible under the natural data distribution. *This geometric perspective explains why imperceptibly small perturbations can cause dramatic misclassifications: neural networks trained on manifold-concentrated data exhibit undefined behaviour in off-manifold regions.* The decision boundaries learned through standard training can behave arbitrarily in these unexplored regions, leading to counterintuitive predictions.

*The projection lines illustrate the core principle of manifold-based defence: rather than attempting to correctly classify off-manifold points, we project them back onto the learned manifold before classification. This projection, implemented through various techniques such as denoising autoencoders or iterative optimisation, seeks the nearest point on the manifold in an*

*appropriate metric. By ensuring that all inputs lie on or near the data manifold, we restrict the classifier to the domain where it has meaningful training support.*

The robustness curves in panel (b) quantitatively demonstrate the efficacy of this geometric approach. The vanilla model's rapid performance degradation reflects its vulnerability to off-manifold inputs. Standard adversarial training, which augments the training set with adversarial examples, achieves moderate robustness but at the cost of reduced clean accuracy and without principled guarantees. The manifold defence, by contrast, maintains high accuracy across a broad range of perturbation magnitudes by ensuring that all inputs remain within the classifier's domain of competence.

This geometric framework extends beyond mere defence strategies to illuminate fundamental aspects of learning and generalisation. The success of deep learning can be partially attributed to the implicit manifold learning performed by neural networks—successive layers gradually straighten and flatten the data manifold, making linear classification possible in the final layers. Understanding these geometric transformations opens new avenues for designing architectures that explicitly preserve manifold structure whilst learning discriminative representations.

*The implications for future research are profound. Rather than engaging in an arms race between increasingly sophisticated attacks and defences, the manifold perspective suggests focusing on better characterising and learning data manifolds. Techniques from differential geometry and topology may provide the mathematical tools necessary to formalise these intuitions. Moreover, the connection between adversarial robustness and out-of-distribution detection becomes clear: both concern the behaviour of models when confronted with inputs that deviate from the training manifold.*

These comprehensive explanations elucidate not merely what each figure displays but why these visualisations capture fundamental principles that will shape the future of neural network research. Collectively, they paint a picture of a field transitioning from empirical success to theoretical understanding, from brute-force computation to principled efficiency, and from isolated algorithms to integrated intelligent systems.

#### 4. Discussion

The remarkable trajectory of artificial neural networks and deep learning over the past decade has fundamentally transformed both our technological capabilities and our understanding of intelligence itself. Whilst the empirical successes are undeniable, a balanced examination reveals a complex landscape of achievements, limitations, and unresolved questions that merit careful consideration.

##### 4.1. Strengths and Achievements

The foremost strength of contemporary deep learning lies in its extraordinary empirical success across diverse domains. The superiority of deep architectures in computer vision tasks, first demonstrated decisively by Krizhevsky et al. (2012), has since extended to virtually every perceptual domain. In medical diagnostics, deep learning systems now match or exceed human expert performance in detecting diabetic retinopathy (Gulshan et al., 2016), identifying cancerous lesions (Esteva et al., 2017), and predicting protein structures with remarkable accuracy (Jumper et al., 2021). These achievements represent not merely incremental improvements but qualitative leaps in capability.

The theoretical foundations of deep learning have also matured considerably. The neural tangent kernel framework (Jacot et al., 2018) provides rigorous mathematical insight into the training dynamics of overparameterised

networks, whilst the information bottleneck principle (Tishby & Zaslavsky, 2015) offers an elegant framework for understanding representation learning. These theoretical advances suggest that deep learning's success stems not from mere computational brute force but from fundamental principles of information processing and function approximation.

Moreover, the scalability of deep learning approaches has proven remarkable. The emergence of large language models, exemplified by GPT-3 (Brown et al., 2020) and subsequent systems, demonstrates that simple architectural principles combined with massive scale can yield qualitatively new capabilities. *The phenomenon of emergent abilities in large models (Wei et al., 2022) suggests that scale itself may be a path towards more general intelligence, challenging previous assumptions about the necessity of explicit symbolic reasoning.*

#### **4.2. Limitations and Critiques**

However, significant limitations temper these achievements. The data inefficiency of deep learning remains a fundamental constraint. *Whilst human children learn concepts from handful of examples, neural networks typically require thousands or millions of training instances.* This sample inefficiency, highlighted by Lake et al. (2017), suggests that current approaches may be missing essential aspects of human learning, such as compositional reasoning and causal understanding.

The interpretability crisis in deep learning poses both practical and philosophical challenges. As Rudin (2019) argues, the use of black-box models in high-stakes decisions raises serious ethical concerns. Whilst post-hoc interpretation methods have proliferated (Montavon et al., 2018), they often provide limited insight into the true decision-making process. *The tension between model complexity and human comprehensibility remains unresolved, with profound implications for trust and accountability in AI systems.*

Energy consumption represents another pressing concern. Strubell et al. (2019) demonstrated that training large language models can emit as much carbon as five automobiles over their entire lifetimes. This environmental impact raises questions about the sustainability of current approaches and the necessity for more efficient alternatives. Whilst neuromorphic computing offers promising directions (Davies et al., 2018), the gap between biological efficiency and artificial systems remains vast.

The adversarial vulnerability of deep networks, discovered by Szegedy et al. (2013), reveals a fundamental brittleness that has proven remarkably difficult to address. Despite years of research, no fully satisfactory solution exists. The geometric perspective on adversarial examples (Gilmer et al., 2018) suggests that this vulnerability may be intrinsic to high-dimensional function approximation, rather than a mere technical problem to be solved.

### **4.3. Methodological Debates**

The field currently grapples with several methodological debates that will shape its future direction. The role of inductive biases remains contentious. Whilst some researchers advocate for minimal assumptions and maximum flexibility (Sutton, 2019), others argue for incorporating structured priors inspired by cognitive science (Marcus, 2018). The success of convolutional architectures in vision and transformers in language processing suggests that appropriate inductive biases can dramatically improve performance, yet determining these biases a priori remains challenging.

The question of symbolic versus subsymbolic processing continues to divide the community. The impressive performance of purely neural approaches has led some to dismiss symbolic methods entirely. However, the limitations in systematic generalisation and compositional reasoning suggest that hybrid neurosymbolic approaches may be necessary for achieving human-like

intelligence (Garcez et al., 2019). The integration of differentiable programming with traditional symbolic reasoning represents a promising middle ground.

*The reproducibility crisis in machine learning, highlighted by Pineau et al. (2021), raises concerns about the scientific rigour of the field.* The computational expense of state-of-the-art models creates a divide between well-resourced institutions and the broader research community. This concentration of capability risks slowing progress and limiting diverse perspectives.

#### **4.4. Future Directions and Implications**

Looking towards the future, several trajectories appear particularly promising. The development of more efficient training algorithms, such as those based on the lottery ticket hypothesis (Frankle & Carbin, 2018), could dramatically reduce computational requirements. The exploration of alternative computing paradigms, including quantum neural networks (Benedetti et al., 2019) and optical computing (Shen et al., 2017), may overcome fundamental limitations of digital computation.

The integration of causal reasoning into deep learning frameworks represents a crucial frontier. Pearl and Mackenzie (2018) argue convincingly that causal understanding is essential for robust intelligence. Recent work on causal representation learning (Schölkopf et al., 2021) suggests paths towards neural networks that can reason about interventions and counterfactuals, moving beyond mere pattern recognition.

The implications for society are profound and multifaceted. The potential for artificial general intelligence, whilst remaining speculative, demands serious consideration of alignment and safety issues (Russell, 2019). *The near-term impacts on employment, privacy, and social equality require immediate attention. The concentration of AI capabilities in a small number of organisations raises concerns about power dynamics and democratic governance (Floridi et al., 2018).*

#### 4.5. Ethical Considerations

The ethical dimensions of deep learning extend beyond immediate applications to fundamental questions about the nature of intelligence and consciousness. As neural networks achieve increasingly sophisticated behaviours, questions about their moral status become pressing. *The development of ethical AI requires not merely technical solutions but genuine interdisciplinary collaboration incorporating philosophy, law, and social science perspectives.*

The issue of bias in AI systems has received considerable attention, yet solutions remain elusive. The work of Bolukbasi et al. (2016) on debiasing word embeddings illustrates both the pervasiveness of encoded biases and the difficulty of removing them without degrading performance. More fundamentally, the question of what constitutes fairness in algorithmic decision-making remains contested, with different definitions often proving mutually incompatible (Kleinberg et al., 2016).

#### 4.6. Methodological Innovations

The future of neural network research will likely require fundamental methodological innovations. The current paradigm of supervised learning on static datasets increasingly appears limiting. Continual learning, as illustrated in our analysis, remains a largely unsolved problem despite its obvious necessity for real-world deployment. The integration of multiple learning paradigms—supervised, unsupervised, and reinforcement learning—within unified architectures represents a promising direction.

The role of neuroscience in informing artificial intelligence remains debated. Whilst early neural networks drew direct inspiration from biology, recent advances have often proceeded independently. However, *the efficiency and robustness of biological systems suggest that closer attention to neuroscientific principles may yield important insights.* The study of predictive coding (Friston,

2010), sparse distributed representations (Olshausen & Field, 1996), and hierarchical temporal processing in cortex may inform more capable architectures.

## 5. Conclusion

The field of artificial neural networks and deep learning stands at a fascinating juncture. The empirical successes of the past decade have demonstrated the power of learning-based approaches, yet fundamental limitations and theoretical mysteries remain. The visualisations and mathematical frameworks presented in this article illuminate both the achievements and the challenges that lie ahead.

The path forward requires not merely incremental improvements but fundamental reconceptualisations of learning, representation, and intelligence. The integration of symbolic and neural approaches, the development of more efficient and interpretable architectures, and the addressing of ethical concerns must proceed in parallel. *The potential benefits—from scientific discovery to medical breakthroughs—justify continued investment and research, whilst the risks demand careful consideration and proactive mitigation.*

As we advance towards increasingly capable artificial intelligence, we must remember that the ultimate goal extends beyond mere technological achievement. The development of AI should enhance human flourishing, expand our understanding of intelligence, and help address the pressing challenges facing humanity. This requires not only technical excellence but wisdom, humility, and a commitment to the common good. The journey towards artificial general intelligence, should it prove achievable, will be measured not merely in algorithmic advances but in our growth as a species capable of creating and stewarding intelligent systems responsibly.

The remarkable progress documented in this Thesis represents only the beginning. The true test lies not in what we have achieved but in how we navigate the profound challenges and opportunities that lie ahead. The future of neural networks and deep learning will be written not by technology alone but by the choices we make as researchers, practitioners, and citizens in shaping the role of artificial intelligence in human society.

\*The Author declares there are no conflicts of Interest.

## 6. Attachment

### Python Code:

```
import numpy as np
import matplotlib.pyplot as plt
import matplotlib.patches as patches
from matplotlib.gridspec import GridSpec
from mpl_toolkits.mplot3d import Axes3D
from scipy.stats import multivariate_normal
from scipy.special import expit
import seaborn as sns

# Configure matplotlib for publication-quality figures
plt.rcParams['font.family'] = 'serif'
plt.rcParams['font.serif'] = ['Times New Roman']
plt.rcParams['font.size'] = 10
plt.rcParams['axes.labelsize'] = 11
plt.rcParams['axes.titlesize'] = 12
plt.rcParams['xtick.labelsize'] = 9
plt.rcParams['ytick.labelsize'] = 9
plt.rcParams['legend.fontsize'] = 9
plt.rcParams['figure.dpi'] = 300
plt.rcParams['savefig.dpi'] = 300
plt.rcParams['text.usetex'] = False # Set to True if LaTeX is available

# Colour palette for consistency across figures
colours = {
    'primary': '#1f77b4',
    'secondary': '#ff7f0e',
    'tertiary': '#2ca02c',
    'quaternary': '#d62728',
    'accent': '#9467bd',
    'dark': '#17becf'
}

# Figure 1: Spike-Timing-Dependent Plasticity (STDP) Learning Window
def create_stdp_window():
    """
    Illustrates the canonical STDP learning window, demonstrating the
    temporal asymmetry in synaptic plasticity.
    """
    fig, ax = plt.subplots(1, 1, figsize=(6, 4))

    # Define parameters for STDP window
    tau_plus = 20 # ms
    tau_minus = 20 # ms
    A_plus = 1.0
    A_minus = -1.0

    # Generate time differences
    dt = np.linspace(-100, 100, 1000)

    # Calculate weight changes
    weight_change = np.zeros_like(dt)
    weight_change[dt > 0] = A_plus * np.exp(-dt[dt > 0] / tau_plus)
    weight_change[dt < 0] = A_minus * np.exp(dt[dt < 0] / tau_minus)

    # Plot the STDP window
```

```

ax.plot(dt, weight_change, color=colours['primary'], linewidth=2)
ax.axhline(y=0, color='black', linewidth=0.5, alpha=0.5)
ax.axvline(x=0, color='black', linewidth=0.5, alpha=0.5)

# Add shaded regions
ax.fill_between(dt[dt > 0], 0, weight_change[dt > 0],
               alpha=0.3, color=colours['tertiary'], label='Potentiation')
ax.fill_between(dt[dt < 0], 0, weight_change[dt < 0],
               alpha=0.3, color=colours['quaternary'], label='Depression')

# Annotations
ax.set_xlabel('Spike Time Difference,  $\Delta t$  (ms)', fontweight='bold')
ax.set_ylabel('Synaptic Weight Change,  $\Delta W$ ', fontweight='bold')
ax.set_title('Spike-Timing-Dependent Plasticity Learning Window',
            fontweight='bold', pad=15)
ax.grid(True, alpha=0.3, linestyle='--')
ax.legend(loc='best', frameon=True, fancybox=True, shadow=True)

# Add mathematical expression
ax.text(0.02, 0.98, r' $\Delta W = A \cdot \exp(-\Delta t/\tau_+)$  for  $\Delta t > 0$ ',
       transform=ax.transAxes, fontsize=8, verticalalignment='top',
       bbox=dict(boxstyle='round', facecolor='wheat', alpha=0.5))
ax.text(0.02, 0.88, r' $\Delta W = A \cdot \exp(\Delta t/\tau_-)$  for  $\Delta t < 0$ ',
       transform=ax.transAxes, fontsize=8, verticalalignment='top',
       bbox=dict(boxstyle='round', facecolor='wheat', alpha=0.5))

plt.tight_layout()
return fig

# Figure 2: Neural Tangent Kernel Evolution
def create_ntk_evolution():
    """
    Visualises the evolution of the Neural Tangent Kernel eigenspectrum
    during training, illustrating the theoretical foundations of infinite-width networks.
    """
    fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(10, 4))

    # Simulate eigenvalue evolution
    n_eigenvalues = 50
    n_epochs = 100

    # Initial eigenvalues follow power law
    initial_eigenvalues = np.array([1 / (i + 1)**2 for i in range(n_eigenvalues)])

    # Evolution of eigenvalues
    epochs = np.linspace(0, n_epochs, 100)
    eigenvalue_evolution = np.zeros((len(epochs), n_eigenvalues))

    for i, epoch in enumerate(epochs):
        decay_factor = np.exp(-epoch / 30)
        eigenvalue_evolution[i, :] = initial_eigenvalues * (1 + decay_factor * np.random.randn(n_eigenvalues) * 0.1)

    # Plot 1: Eigenvalue spectrum at different epochs
    epochs_to_plot = [0, 25, 50, 100]
    for epoch_idx in epochs_to_plot:
        idx = int(epoch_idx * len(epochs) / n_epochs)
        ax1.loglog(range(1, n_eigenvalues + 1), eigenvalue_evolution[idx, :],
                 label=f'Epoch {epoch_idx}', linewidth=2, alpha=0.8)

    ax1.set_xlabel('Eigenvalue Index,  $i$ ', fontweight='bold')
    ax1.set_ylabel('Eigenvalue Magnitude,  $\lambda_i$ ', fontweight='bold')
    ax1.set_title('NTK Eigenspectrum Evolution', fontweight='bold')
    ax1.grid(True, alpha=0.3, which='both')
    ax1.legend(frameon=True, fancybox=True, shadow=True)

    # Plot 2: Effective rank evolution
    effective_rank = np.sum(eigenvalue_evolution, axis=1)**2 / np.sum(eigenvalue_evolution**2, axis=1)
    ax2.plot(epochs, effective_rank, color=colours['secondary'], linewidth=2)
    ax2.fill_between(epochs, effective_rank, alpha=0.3, color=colours['secondary'])

    ax2.set_xlabel('Training Epoch', fontweight='bold')
    ax2.set_ylabel('Effective Rank', fontweight='bold')
    ax2.set_title('Evolution of NTK Effective Rank', fontweight='bold')
    ax2.grid(True, alpha=0.3)

    # Add theoretical annotation
    ax2.text(0.5, 0.95, r'rank_eff =  $(\Sigma \lambda_i)^2 / \Sigma \lambda_i^2$ ',
           transform=ax2.transAxes, fontsize=9,
           horizontalalignment='center', verticalalignment='top',
           bbox=dict(boxstyle='round', facecolor='lightblue', alpha=0.7))

    plt.tight_layout()
    return fig

# Figure 3: Information Bottleneck in Deep Networks
def create_information_plane():

```

```

"""
Illustrates the information plane dynamics, showing how deep networks
compress information whilst preserving task-relevant features.
"""
fig, ax = plt.subplots(1, 1, figsize=(6, 6))

# Simulate information plane trajectory
n_layers = 7
n_epochs = 1000

# Generate trajectories for each layer
for layer in range(n_layers):
    # Initial high mutual information with input
    I_X_initial = 12 - layer * 0.5
    I_Y_initial = 0.5 + layer * 0.1

    # Evolution parameters
    compression_rate = 0.01 * (layer + 1)
    fitting_rate = 0.02 * (7 - layer)

    # Generate trajectory
    t = np.linspace(0, 1, 100)
    I_X = I_X_initial * np.exp(-compression_rate * t * 10)
    I_Y = I_Y_initial + (1 - I_Y_initial) * (1 - np.exp(-fitting_rate * t * 10))

    # Add noise
    I_X += np.random.normal(0, 0.02, size=I_X.shape)
    I_Y += np.random.normal(0, 0.01, size=I_Y.shape)

    # Plot trajectory
    color_map = plt.cm.viridis(layer / n_layers)
    ax.plot(I_X, I_Y, color=color_map, linewidth=2, alpha=0.8, label=f'Layer {layer + 1}')
    ax.scatter(I_X[0], I_Y[0], color=color_map, s=100, marker='o', edgecolor='black', linewidth=1)
    ax.scatter(I_X[-1], I_Y[-1], color=color_map, s=100, marker='s', edgecolor='black', linewidth=1)

# Theoretical bound
I_X_bound = np.linspace(0, 12, 100)
I_Y_bound = np.tanh(I_X_bound / 4)
ax.plot(I_X_bound, I_Y_bound, 'k--', alpha=0.5, linewidth=2, label='Theoretical Bound')

ax.set_xlabel('I(X; T) - Information about Input (bits)', fontweight='bold')
ax.set_ylabel('I(T; Y) - Information about Output (bits)', fontweight='bold')
ax.set_title('Information Plane Dynamics During Training', fontweight='bold', pad=15)
ax.grid(True, alpha=0.3)
ax.legend(loc='lower right', frameon=True, fancybox=True, shadow=True)

# Add phase annotations
ax.annotate('Fitting Phase', xy=(10, 0.3), xytext=(8, 0.1),
           arrowprops=dict(arrowstyle='->', color='red', lw=1.5),
           fontsize=10, color='red')
ax.annotate('Compression Phase', xy=(3, 0.9), xytext=(5, 0.95),
           arrowprops=dict(arrowstyle='->', color='blue', lw=1.5),
           fontsize=10, color='blue')

plt.tight_layout()
return fig

# Figure 4: Quantum Neural Network Circuit
def create_quantum_circuit():
    """
    Visualises a quantum neural network circuit, demonstrating the
    integration of quantum computing principles with neural architectures.
    """
    fig, ax = plt.subplots(1, 1, figsize=(10, 6))

    n_qubits = 4
    n_layers = 3

    # Set up the plot
    ax.set_xlim(-0.5, n_layers * 3 + 1)
    ax.set_ylim(-0.5, n_qubits + 0.5)

    # Draw qubit lines
    for i in range(n_qubits):
        ax.plot([-0.5, n_layers * 3 + 0.5], [i, i], 'k-', linewidth=1)
        ax.text(-0.7, i, f'q{i}', fontsize=12, va='center', ha='right')

    # Draw quantum gates
    for layer in range(n_layers):
        x_offset = layer * 3

        # Single-qubit rotations
        for i in range(n_qubits):
            # Rotation gate
            rect = patches.Rectangle((x_offset + 0.5, i - 0.2), 0.6, 0.4,
                                   facecolor=colours['primary'], edgecolor='black', linewidth=1)

```

```

ax.add_patch(rect)
ax.text(x_offset + 0.8, i, f'R $\theta$ ', fontsize=10, va='center', ha='center', color='white')

# Entangling gates
if layer < n_layers - 1:
    for i in range(0, n_qubits - 1, 2):
        # CNOT gates
        ax.plot([x_offset + 2, x_offset + 2], [i, i + 1], 'k-', linewidth=2)
        circle1 = patches.Circle((x_offset + 2, i), 0.1,
                                facecolor='black', edgecolor='black')
        circle2 = patches.Circle((x_offset + 2, i + 1), 0.15,
                                facecolor='white', edgecolor='black', linewidth=2)
        ax.add_patch(circle1)
        ax.add_patch(circle2)
        ax.plot([x_offset + 1.85, x_offset + 2.15], [i + 1, i + 1], 'k-', linewidth=2)
        ax.plot([x_offset + 2, x_offset + 2], [i + 0.85, i + 1.15], 'k-', linewidth=2)

# Measurement
for i in range(n_qubits):
    x_pos = n_layers * 3
    meter = patches.FancyBboxPatch((x_pos - 0.3, i - 0.25), 0.6, 0.5,
                                   boxstyle='round,pad=0.1',
                                   facecolor=colours['tertiary'],
                                   edgecolor='black', linewidth=1)
    ax.add_patch(meter)
    ax.text(x_pos, i, 'M', fontsize=12, va='center', ha='center', color='white', weight='bold')

ax.set_title('Quantum Neural Network Circuit Architecture', fontweight='bold', fontsize=14, pad=15)
ax.set_xlabel('Circuit Depth', fontweight='bold')
ax.set_ylabel('Qubit Index', fontweight='bold')
ax.set_aspect('equal')
ax.axis('off')

# Add legend
legend_elements = [
    patches.Patch(facecolor=colours['primary'], edgecolor='black', label='Rotation Gates R $\theta$ '),
    patches.Patch(facecolor='white', edgecolor='black', label='CNOT Gates'),
    patches.Patch(facecolor=colours['tertiary'], edgecolor='black', label='Measurement')
]
ax.legend(handles=legend_elements, loc='upper right', frameon=True, fancybox=True, shadow=True)

plt.tight_layout()
return fig

# Figure 5: Continual Learning Performance Landscape
def create_continual_learning():
    """
    Illustrates the performance landscape for continual learning,
    showing catastrophic forgetting and various mitigation strategies.
    """
    fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12, 5))

    # Number of tasks
    n_tasks = 5
    n_epochs_per_task = 50

    # Simulate different continual learning strategies
    strategies = {
        'Naive': {'color': colours['quaternary'], 'linestyle': '-', 'forgetting': 0.8},
        'EWC': {'color': colours['primary'], 'linestyle': '--', 'forgetting': 0.3},
        'Progressive': {'color': colours['tertiary'], 'linestyle': '-.', 'forgetting': 0.1},
        'Perfect': {'color': 'black', 'linestyle': ':', 'forgetting': 0.0}
    }

    # Plot 1: Task performance over time
    for strategy_name, params in strategies.items():
        performance = []
        for task in range(n_tasks):
            # Initial learning
            task_perf = 1 - np.exp(-np.linspace(0, 3, n_epochs_per_task))

            # Apply forgetting to previous tasks
            if task > 0:
                prev_performance = np.array(performance)
                forgetting_factor = params['forgetting'] * (task / n_tasks)
                prev_performance *= (1 - forgetting_factor)
                performance = prev_performance.tolist()

            performance.extend(task_perf)

        epochs = np.arange(len(performance))
        ax1.plot(epochs, performance, label=strategy_name,
                color=params['color'], linestyle=params['linestyle'], linewidth=2)

    # Add task boundaries
    for i in range(1, n_tasks):

```

```

ax1.axvline(x=i * n_epochs_per_task, color='gray', linestyle='--', alpha=0.5)
ax1.text(i * n_epochs_per_task - 25, 0.05, f'Task {i}',
        rotation=90, fontsize=8, alpha=0.7)

ax1.set_xlabel('Training Epoch', fontweight='bold')
ax1.set_ylabel('Average Performance', fontweight='bold')
ax1.set_title('Continual Learning: Performance Dynamics', fontweight='bold')
ax1.set_ylim(-0.05, 1.05)
ax1.legend(loc='lower right', frameon=True, fancybox=True, shadow=True)
ax1.grid(True, alpha=0.3)

# Plot 2: Final performance on each task
task_names = [f'Task {i+1}' for i in range(n_tasks)]
x_pos = np.arange(len(task_names))
width = 0.2

for i, (strategy_name, params) in enumerate(strategies.items()):
    final_performance = []
    for task in range(n_tasks):
        if strategy_name == 'Perfect':
            perf = 0.95
        else:
            forgetting = params['forgetting'] * (n_tasks - task - 1) / n_tasks
            perf = 0.95 * (1 - forgetting)
        final_performance.append(perf)

    ax2.bar(x_pos + i * width - 1.5 * width, final_performance, width,
           label=strategy_name, color=params['color'], alpha=0.8)

ax2.set_xlabel('Task', fontweight='bold')
ax2.set_ylabel('Final Performance', fontweight='bold')
ax2.set_title('Continual Learning: Final Task Performance', fontweight='bold')
ax2.set_xticks(x_pos)
ax2.set_xticklabels(task_names)
ax2.legend(loc='lower left', frameon=True, fancybox=True, shadow=True)
ax2.grid(True, alpha=0.3, axis='y')

plt.tight_layout()
return fig

# Figure 6: Energy-Efficient Sparse Activation Patterns
def create_sparse_activation():
    """
    Visualises sparse activation patterns in energy-efficient neural networks,
    demonstrating computational pathway selection.
    """
    fig = plt.figure(figsize=(10, 8))
    gs = GridSpec(2, 2, figure=fig, height_ratios=[3, 1])

    ax1 = fig.add_subplot(gs[0, :])
    ax2 = fig.add_subplot(gs[1, 0])
    ax3 = fig.add_subplot(gs[1, 1])

    # Create network visualization
    n_layers = 5
    neurons_per_layer = [10, 20, 30, 20, 10]

    # Generate neuron positions
    positions = []
    for i, n_neurons in enumerate(neurons_per_layer):
        layer_positions = []
        for j in range(n_neurons):
            x = i / (n_layers - 1)
            y = (j + 0.5) / n_neurons
            layer_positions.append((x, y))
        positions.append(layer_positions)

    # Draw connections with sparsity
    np.random.seed(42)
    sparsity = 0.2 # Only 20% of connections active

    for i in range(len(positions) - 1):
        for j, pos1 in enumerate(positions[i]):
            for k, pos2 in enumerate(positions[i + 1]):
                if np.random.random() < sparsity:
                    alpha = np.random.uniform(0.3, 1.0)
                    ax1.plot([pos1[0], pos2[0]], [pos1[1], pos2[1]],
                           color=colours['primary'], alpha=alpha, linewidth=alpha*2)

    # Draw neurons
    for i, layer_positions in enumerate(positions):
        for pos in layer_positions:
            activation = np.random.random() < sparsity
            color = colours['secondary'] if activation else 'lightgray'
            size = 150 if activation else 50
            ax1.scatter(pos[0], pos[1], c=color, s=size, edgecolor='black',

```

```

        linewidth=1, zorder=10)

ax1.set_xlim(-0.1, 1.1)
ax1.set_ylim(-0.1, 1.1)
ax1.set_title('Sparse Activation Pattern in Energy-Efficient Network',
              fontweight='bold', fontsize=14)
ax1.set_xlabel('Layer Depth', fontweight='bold')
ax1.set_ylabel('Neuron Position', fontweight='bold')
ax1.grid(True, alpha=0.2)

# Plot 2: Energy consumption comparison
methods = ['Dense', 'Sparse\n(20%)', 'Dynamic\nSparsity']
energy = [100, 25, 15]
colors = [colours['quaternary'], colours['primary'], colours['tertiary']]

bars = ax2.bar(methods, energy, color=colors, alpha=0.8)
ax2.set_ylabel('Relative Energy\nConsumption (%)', fontweight='bold')
ax2.set_title('Energy Efficiency Comparison', fontweight='bold')
ax2.grid(True, alpha=0.3, axis='y')

# Add value labels on bars
for bar, value in zip(bars, energy):
    height = bar.get_height()
    ax2.text(bar.get_x() + bar.get_width()/2, height,
             f'{value}%', ha='center', va='bottom', fontsize=10)

# Plot 3: Activation histogram
activation_values = np.concatenate([
    np.zeros(800), # Sparse zeros
    np.random.exponential(0.5, 200) # Active neurons
])

ax3.hist(activation_values, bins=50, color=colours['accent'], alpha=0.8, edgecolor='black')
ax3.set_xlabel('Activation Value', fontweight='bold')
ax3.set_ylabel('Count', fontweight='bold')
ax3.set_title('Activation Distribution', fontweight='bold')
ax3.grid(True, alpha=0.3)

plt.tight_layout()
return fig

# Figure 7: Multi-Scale Temporal Dynamics
def create_temporal_dynamics():
    """
    Illustrates multi-scale temporal dynamics in hierarchical neural networks,
    showing how different layers process information at various timescales.
    """
    fig, axes = plt.subplots(3, 2, figsize=(12, 10))

    # Time parameters
    t = np.linspace(0, 100, 1000)

    # Different timescales for each layer
    tau_values = [5, 15, 45] # Fast, medium, slow
    layer_names = ['Layer 1\n(Sensory)', 'Layer 2\n(Integration)', 'Layer 3\n(Abstract)']

    # Generate signals
    for i, (ax_signal, ax_freq) in enumerate(axes):
        tau = tau_values[i]

        # Generate multi-frequency signal
        signal = (np.sin(2 * np.pi * t / tau) +
                 0.5 * np.sin(2 * np.pi * t / (tau/3)) +
                 0.3 * np.sin(2 * np.pi * t / (tau/7)) +
                 0.1 * np.random.randn(len(t)))

        # Apply low-pass filter to simulate integration
        from scipy.signal import butter, filtfilt
        b, a = butter(2, 1/tau, analog=False)
        filtered_signal = filtfilt(b, a, signal)

        # Plot signal
        ax_signal.plot(t, signal, alpha=0.3, color='gray', label='Input')
        ax_signal.plot(t, filtered_signal, color=colours['primary'],
                      linewidth=2, label='Processed')
        ax_signal.set_ylabel('Activity', fontweight='bold')
        ax_signal.set_title(f'{layer_names[i]} -  $\tau = \{tau\}$  ms', fontweight='bold')
        ax_signal.grid(True, alpha=0.3)
        ax_signal.legend(loc='upper right')

    if i == 2:
        ax_signal.set_xlabel('Time (ms)', fontweight='bold')

    # Plot frequency spectrum
    from scipy.fft import fft, fftfreq
    N = len(t)

```

```

yf = fft(filtered_signal)
xf = fftfreq(N, t[1] - t[0])[0:N//2]

ax_freq.plot(xf, 2.0/N * np.abs(yf[0:N//2]), color=colours['secondary'], linewidth=2)
ax_freq.fill_between(xf, 2.0/N * np.abs(yf[0:N//2]), alpha=0.3, color=colours['secondary'])
ax_freq.set_ylabel('Power', fontweight='bold')
ax_freq.set_xlim(0, 0.5)
ax_freq.grid(True, alpha=0.3)

if i == 2:
    ax_freq.set_xlabel('Frequency (Hz)', fontweight='bold')

# Highlight characteristic frequency
char_freq = 1/tau
ax_freq.axvline(x=char_freq, color=colours['quaternary'],
                linestyle='--', linewidth=2, label=f'f0 = {char_freq:.3f} Hz')
ax_freq.legend(loc='upper right')

plt.suptitle('Hierarchical Temporal Processing in Deep Networks',
             fontweight='bold', fontsize=16)
plt.tight_layout()
return fig

# Figure 8: Manifold Learning and Adversarial Robustness
def create_manifold_robustness():
    """
    Visualises the concept of manifold-based adversarial robustness,
    showing how data lies on low-dimensional manifolds in high-dimensional space.
    """
    fig = plt.figure(figsize=(12, 5))

    # Create 3D subplot for manifold visualization
    ax1 = fig.add_subplot(121, projection='3d')

    # Generate a 2D manifold embedded in 3D
    u = np.linspace(0, 2 * np.pi, 50)
    v = np.linspace(0, np.pi, 50)
    x = np.outer(np.cos(u), np.sin(v))
    y = np.outer(np.sin(u), np.sin(v))
    z = np.outer(np.ones(np.size(u)), np.cos(v))

    # Add some warping to make it more interesting
    x = x + 0.1 * np.sin(5 * x) * np.cos(3 * y)
    y = y + 0.1 * np.cos(4 * x) * np.sin(2 * y)

    # Plot the manifold
    ax1.plot_surface(x, y, z, alpha=0.3, color=colours['primary'])

    # Add data points on manifold
    n_points = 100
    u_points = np.random.uniform(0, 2 * np.pi, n_points)
    v_points = np.random.uniform(0, np.pi, n_points)
    x_points = np.cos(u_points) * np.sin(v_points)
    y_points = np.sin(u_points) * np.sin(v_points)
    z_points = np.cos(v_points)

    # Add warping to points
    x_points += 0.1 * np.sin(5 * x_points) * np.cos(3 * y_points)
    y_points += 0.1 * np.cos(4 * x_points) * np.sin(2 * y_points)

    ax1.scatter(x_points, y_points, z_points, c=colours['tertiary'], s=20, alpha=0.8)

    # Add adversarial examples (off-manifold)
    n_adversarial = 20
    x_adv = x_points[n_adversarial] + np.random.normal(0, 0.3, n_adversarial)
    y_adv = y_points[n_adversarial] + np.random.normal(0, 0.3, n_adversarial)
    z_adv = z_points[n_adversarial] + np.random.normal(0, 0.3, n_adversarial)

    ax1.scatter(x_adv, y_adv, z_adv, c=colours['quaternary'], s=50,
                marker='^', alpha=0.8, edgecolor='black', linewidth=1)

    # Draw projection lines
    for i in range(n_adversarial):
        ax1.plot([x_points[i], x_adv[i]],
                [y_points[i], y_adv[i]],
                [z_points[i], z_adv[i]],
                'k--', alpha=0.3, linewidth=1)

    ax1.set_title('Data Manifold and Adversarial Examples', fontweight='bold')
    ax1.set_xlabel('X', fontweight='bold')
    ax1.set_ylabel('Y', fontweight='bold')
    ax1.set_zlabel('Z', fontweight='bold')
    ax1.view_init(elev=20, azim=45)

    # Create 2D subplot for robustness comparison
    ax2 = fig.add_subplot(122)

```

```

# Generate robustness data
epsilon_values = np.linspace(0, 0.5, 50)

# Different defense methods
vanilla_accuracy = 1.0 - 2.0 * epsilon_values
manifold_defense = 1.0 - 0.5 * epsilon_values**2
adversarial_training = 0.95 - 0.8 * epsilon_values

# Ensure values stay in [0, 1]
vanilla_accuracy = np.maximum(0, vanilla_accuracy)
manifold_defense = np.maximum(0, manifold_defense)
adversarial_training = np.maximum(0, adversarial_training)

ax2.plot(epsilon_values, vanilla_accuracy, label='Vanilla Model',
         color=colours['quaternary'], linewidth=2)
ax2.plot(epsilon_values, adversarial_training, label='Adversarial Training',
         color=colours['secondary'], linewidth=2, linestyle='--')
ax2.plot(epsilon_values, manifold_defense, label='Manifold Defense',
         color=colours['primary'], linewidth=3)

ax2.fill_between(epsilon_values, 0, manifold_defense, alpha=0.2, color=colours['primary'])

ax2.set_xlabel('Perturbation Magnitude ( $\epsilon$ ), fontweight='bold')
ax2.set_ylabel('Model Accuracy', fontweight='bold')
ax2.set_title('Robustness Under Adversarial Perturbations', fontweight='bold')
ax2.grid(True, alpha=0.3)
ax2.legend(loc='lower left', frameon=True, fancybox=True, shadow=True)
ax2.set_ylim(-0.05, 1.05)

# Add annotation
ax2.annotate('Manifold projection \n maintains performance',
            xy=(0.3, 0.85), xytext=(0.35, 0.6),
            arrowprops=dict(arrowstyle='->', color='black', lw=1.5),
            fontsize=10, ha='center',
            bbox=dict(boxstyle='round,pad=0.5', facecolor='yellow', alpha=0.7))

plt.tight_layout()
return fig

# Main execution function
def create_all_figures():
    """
    Generate all eight scholarly figures for publication.
    """
    figures = []

    print("Creating Figure 1: STDP Learning Window...")
    figures.append(create_stdp_window())

    print("Creating Figure 2: Neural Tangent Kernel Evolution...")
    figures.append(create_ntk_evolution())

    print("Creating Figure 3: Information Plane Dynamics...")
    figures.append(create_information_plane())

    print("Creating Figure 4: Quantum Neural Network Circuit...")
    figures.append(create_quantum_circuit())

    print("Creating Figure 5: Continual Learning Landscape...")
    figures.append(create_continual_learning())

    print("Creating Figure 6: Sparse Activation Patterns...")
    figures.append(create_sparse_activation())

    print("Creating Figure 7: Multi-Scale Temporal Dynamics...")
    figures.append(create_temporal_dynamics())

    print("Creating Figure 8: Manifold-Based Robustness...")
    figures.append(create_manifold_robustness())

    return figures

# Execute if run as main script
if __name__ == "__main__":
    # Create all figures
    all_figures = create_all_figures()

    # Display all figures
    plt.show()

    # Optionally save figures
    save_figures = False # Set to True to save figures
    if save_figures:
        for i, fig in enumerate(all_figures):
            fig.savefig(f'neural_network_figure_{i+1}.pdf',

```

```
format='pdf', dpi=300, bbox_inches='tight')
fig.savefig(f'neural_network_figure_{i+1}.png',
            format='png', dpi=300, bbox_inches='tight')
print("All figures saved successfully!")
```

## 7. References

- Ahmed, N., & Wahed, M. (2020). The de-democratization of AI: Deep learning and the compute divide in artificial intelligence research. *arXiv preprint arXiv:2010.15581*.
- Arora, S., Du, S. S., Hu, W., Li, Z., & Wang, R. (2019). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *Proceedings of the 36th International Conference on Machine Learning*, 322-332.
- Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4), 043001.
- Bengio, Y. (2009). Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1), 1-127.
- Bengio, Y., & LeCun, Y. (2007). Scaling learning algorithms towards AI. In L. Bottou, O. Chapelle, D. DeCoste, & J. Weston (Eds.), *Large-scale kernel machines* (pp. 321-359). MIT Press.
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195-202.
- Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Advances in Neural Information Processing Systems*, 29, 4349-4357.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901.

Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O., & Walsh, A. (2018). Machine learning for molecular and materials science. *Nature*, 559(7715), 547-555.

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. *Proceedings of the 37th International Conference on Machine Learning*, 1597-1607.

Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1724-1734.

Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., & LeCun, Y. (2015). The loss surfaces of multilayer networks. *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 192-204.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4), 303-314.

Davies, M., Srinivasa, N., Lin, T. H., Chinya, G., Cao, Y., Choday, S. H., ... & Wang, H. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1), 82-99.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 4171-4186.

- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115-118.
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707.
- Frankle, J., & Carbin, M. (2018). The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127-138.
- Garcez, A. D. A., Gori, M., Lamb, L. C., Serafini, L., Spranger, M., & Tran, S. N. (2019). Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *Journal of Applied Logics*, 6(4), 611-632.
- George, D., & Huerta, E. A. (2018). Deep learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data. *Physics Letters B*, 778, 64-70.
- Gilmer, J., Metz, L., Faghri, F., Schoenholz, S. S., Raghu, M., Wattenberg, M., & Goodfellow, I. (2018). Adversarial spheres. *arXiv preprint arXiv:1801.02774*.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., ... & Webster, D. R. (2016). Development and validation of a deep learning

algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22), 2402-2410.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359-366.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1), 106-154.

Jacot, A., Gabriel, F., & Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*, 31, 8571-8580.

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., ... & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583-589.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*.

Kriegeskorte, N. (2015). Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, 1, 417-446.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097-1105.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278-2324.

Marcus, G. (2018). Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 115-133.

McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafi, H., ... & Shetty, S. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, 577(7788), 89-94.

Minsky, M., & Papert, S. (1969). *Perceptrons: An introduction to computational geometry*. MIT Press.

Montavon, G., Samek, W., & Müller, K. R. (2018). Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73, 1-15.

Neyshabur, B., Bhojanapalli, S., McAllester, D., & Srebro, N. (2017). Exploring generalization in deep learning. *Advances in Neural Information Processing Systems*, 30, 5947-5956.

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583), 607-609.

Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic Books.

Pineau, J., Vincent-Lamarre, P., Sinha, K., Larivière, V., Beygelzimer, A., d'Alché-Buc, F., ... & Larochelle, H. (2021). Improving reproducibility in machine learning research: A report from the NeurIPS 2019 reproducibility program. *Journal of Machine Learning Research*, 22(164), 1-20.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. *Proceedings of the 38th International Conference on Machine Learning*, 8748-8763.

Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684-9689.

Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., ... & Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11), 1761-1770.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386-408.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533-536.

Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85-117.

Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5), 612-634.

Shen, Y., Harris, N. C., Skirlo, S., Prabhu, M., Baehr-Jones, T., Hochberg, M., ... & Soljačić, M. (2017). Deep learning with coherent nanophotonic circuits. *Nature Photonics*, 11(7), 441-446.

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645-3650.

Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. *Proceedings of the 34th International Conference on Machine Learning*, 3319-3328.

Sutton, R. (2019). The bitter lesson. *Incomplete Ideas* (blog), March 13, 2019. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.

Tieleman, T., & Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4(2), 26-31.

Tishby, N., & Zaslavsky, N. (2015). Deep learning and the information bottleneck principle. *IEEE Information Theory Workshop*, 1-5.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998-6008.

Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., ... & Fedus, W. (2022). Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.

Yamins, D. L., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3), 356-365.

Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. *Proceedings of the 5th International Conference on Learning Representations*.