INTERNATIONAL MULTI DISCIPLINARY JOURNAL FOR RESEARCH & DEVELOPMENT
eISSN 2394-6334 Impact factor: 7,854 Volume 12, issue 10 (2025)

SYNAPTIC PLASTICITY AND CRITICAL LEARNING PERIODS IN DEEP CONVOLUTIONAL NETWORKS: BRIDGING BIOLOGICAL MECHANISMS WITH MACHINE REPRESENTATION

Dr. Elias V. Thorne
Department of Computational Neuroscience, Institute of Advanced Systems

Dr. Sarah J. Mendez
Department of Computational Neuroscience, Institute of Advanced Systems

Abstract: The Background: Deep Convolutional Neural Networks (CNNs) have achieved remarkable success in static classification tasks. However, unlike biological neural systems, they often struggle with continuous learning and adaptability, prone to catastrophic forgetting when introduced to new data distributions.
Methods: This study proposes a hybrid architecture, the "Plastic-CNN" (P-CNN), which integrates principles of synaptic plasticity and adult neurogenesis into standard deep learning frameworks. We compare the performance of standard architectures (VGG-16, AlexNet) against the P-CNN using the iNaturalist 2017 dataset for fine-grained visual classification and a financial dataset for dynamic fraud detection. We further analyze the role of "critical learning periods" by modulating information bottlenecks during the initial training epochs.
Results: The P-CNN demonstrated a statistically significant improvement in long-term feature retention compared to static baselines. Specifically, the inclusion of neurogenic node-addition layers reduced validation loss during domain shifts. The analysis confirms that applying information constraints during the early "critical period" of training creates more robust generalized representations, mirroring biological sensory development.
Conclusion: Incorporating biological constraints such as synaptic plasticity and neurogenesis does not merely mimic the brain but offers a tangible computational advantage for addressing the stability-plasticity dilemma in artificial intelligence. These findings suggest that future architectures should prioritize dynamic topology over static depth.

Keywords: Deep Learning, Synaptic Plasticity, Convolutional Neural Networks, Neurogenesis, Critical Learning Periods, Artificial Intelligence, Adaptive Algorithms.

1. Introduction

The trajectory of artificial intelligence, particularly within the domain of computer vision, has been defined by the pursuit of deeper, more abstract representations of data. Since the seminal work on backpropagation applied to handwritten zip code recognition [3], the field has moved toward architectures of increasing depth and complexity. The development of AlexNet [4] marked a paradigm shift, demonstrating that deep convolutional neural networks (CNNs) could outperform traditional computer vision techniques by a significant margin. This was further refined by architectures such as VGG [5], which emphasized the importance of depth and small convolutional filters in capturing hierarchical features. However, despite these advances, a fundamental dichotomy remains between artificial neural networks (ANNs) and the biological systems they vaguely emulate: the balance between stability and plasticity.

In biological systems, learning is not a static process that terminates once a threshold of accuracy is reached. Instead, the mammalian brain exhibits continuous synaptic plasticity, allowing it to acquire new information without immediate destruction of previously stored memories—a phenomenon described as "taming the beast" of synaptic input [7]. Conversely, standard deep learning models, once trained, are typically frozen. If they are retrained on new data, they suffer from catastrophic forgetting, where the adjustment of weights to accommodate new patterns erases the optimized configuration for previous ones. This rigidity limits the deployment of CNNs in dynamic environments where the data distribution shifts over time, such as in real-time fraud detection [1] or ecological monitoring via datasets like iNaturalist [6].

This paper investigates the integration of biological mechanisms—specifically synaptic plasticity, critical learning periods, and neurogenesis—into deep convolutional architectures. We posit that the "static" nature of current state-of-the-art CNNs is a limitation derived from computational convenience rather than theoretical necessity. Recent advances in understanding convolutional networks [2] allow us to visualize feature maps, yet these maps remain fixed after training. By contrast, research into the critical learning periods of deep networks [9] suggests that the initial phases of training establish an "information bottleneck" that defines the network's future plasticity. If we view the training of a CNN not as a simple optimization problem but as a developmental process analogous to biological maturation, new avenues for architecture optimization emerge.

Furthermore, the role of neurogenesis—the birth of new neurons in the adult brain—has been identified as a crucial component in memory retention and the encoding of temporal sequences [12, 13]. While traditional algorithms often utilize fixed topologies, the introduction of dynamic node generation could theoretically allow artificial systems to segregate new information from old, thereby mitigating interference. This study aims to bridge the gap between these biological insights and machine representation. We propose a comparative analysis, examining how biologically inspired modifications to standard algorithms affect performance, stability, and convergence. By revisiting the foundational principles of backpropagation [3] through the lens of modern neuroscience [7, 8], we seek to develop a more robust framework for learning rich features without the immediate necessity for massive, repetitive labeling.

2. Literature Review

2.1 The Evolution of Convolutional Architectures

The resurgence of neural networks in the current century is largely attributable to the successful implementation of Convolutional Neural Networks (CNNs). The fundamental architecture, established by LeCun et al. [3], utilized shared weights and pooling layers to reduce computational complexity while retaining spatial hierarchies. This approach remained relatively dormant until Krizhevsky et al. [4] leveraged GPU acceleration to train a massive deep network, achieving state-of-the-art results on ImageNet. Their work highlighted the necessity of non-linear activation functions (ReLU) and dropout regularization to prevent overfitting. Subsequent research by Zeiler and Fergus [2] provided a visual methodology for understanding these networks, revealing that lower layers act as edge and texture detectors while upper layers represent high-level concepts. Simonyan and Zisserman [5] pushed this further with VGG, demonstrating that very deep networks with small filters could achieve superior localization and classification. However, Gu et al. [2] note in their review of recent advances that while architectures have become more complex, the fundamental learning rule—backpropagation via stochastic gradient descent—has remained largely unchanged.

2.2 Biological Plasticity and Stability

In contrast to the fixed weights of a deployed CNN, biological synapses are in a constant state of flux. Abbott and Nelson [7] describe synaptic plasticity as a complex interplay of potentiation and depression that allows neural circuits to self-regulate. This regulation is crucial; without it, Hebbian learning (cells that fire together, wire together) would lead to runaway excitation. Abraham and Robins [8] discuss this in the context of ANN simulations, noting that retention of memory in artificial systems is often achieved by lowering the learning rate to near zero, which effectively halts the acquisition of new knowledge. This "stability-plasticity dilemma" is a core challenge. Biological systems solve this partly through active inference and predictive coding [10], where the brain actively constructs hypotheses about sensory input, updating internal models only when prediction errors occur.

2.3 Critical Periods and Information Theory

A pivotal concept in recent literature is the existence of "critical periods" in deep learning, a term borrowed from developmental biology. Achille et al. [9] demonstrated that deep networks exhibit a temporary phase during early training where the topology is highly sensitive to input statistics. Once this period passes, the network's ability to learn new, distinct tasks diminishes significantly. This mirrors the biological critical periods for vision or language acquisition. This suggests that the timing of data presentation is as critical as the data itself. If the initial weights are set during a period of high plasticity, they form a "skeleton" of representation that guides all future learning.

2.4 Neurogenesis and Temporal Encoding

Perhaps the most radical departure from standard connectionism is the concept of neurogenesis. Altman [14] provided early autoradiographic evidence of cell proliferation in adult brains, challenging the dogma that the central nervous system is fixed post-development. Aimone et al. [12, 13] expanded on this by proposing a computational function for these new neurons: temporal encoding. They argue that young, highly excitable neurons provide a timestamp to memories, allowing the hippocampus to distinguish between similar events that occurred at different times. In the context of artificial intelligence, this implies that a fixed network topology is insufficient for tasks involving temporal drift. Adams et al. [10] further suggest that cortical architecture is modular and hierarchical, supporting the idea that new modules (or nodes) could be recruited for new tasks.

2.5 Feedback and Control Systems

Finally, the stability of any learning system, biological or artificial, can be viewed through the lens of control theory. Astrom and Murray [15] define feedback systems as essential for managing uncertainty and disturbance. In neural networks, the error gradient serves as the feedback signal. However, as networks become deeper, the feedback signal can vanish or explode, leading to instability. Biological systems utilize complex inhibitory loops to maintain homeostasis [7]. Patel [1] highlights the importance of such stability in applied fields like fraud detection, where algorithms must adapt to new fraud patterns without destabilizing the detection of established methods.

3. Methodology

3.1 Architecture Design: The Plastic-CNN

To test the hypothesis that biological constraints improve representation, we designed a hybrid architecture termed the "Plastic-CNN" (P-CNN). The baseline architecture was a modified VGG-16 [5]. The P-CNN variant introduced two key modifications:

1. Synaptic Plasticity Layer: Replacing standard dense layers with "plastic" layers where weight updates are governed not only by gradients but by a local decay term λ, simulating the metabolic cost of maintaining a synapse.

2. Neurogenic Module: A dynamic mechanism where the network monitors the error distribution in the penultimate layer. If the loss remains above a threshold τ for n epochs (indicating an inability to learn a new pattern with current capacity), a new node is initialized. This node is connected with high plasticity (high learning rate) while existing nodes are "frozen" or set to a lower learning rate, mimicking the hyper-excitability of young neurons described by Aimone et al. [12].

3.2 Datasets and Preprocessing

We utilized two distinct datasets to evaluate the generalizability of the approach:

- iNaturalist 2017 [6]: This dataset contains large-scale species classification data. It was selected for its fine-grained categories and visual complexity. The long-tail distribution of species poses a significant challenge for standard CNNs, making it ideal for testing plasticity. We utilized a subset of 100,000 images across 500 classes.

- Financial Fraud Dataset: Following the context of Patel [1], we utilized a synthetic transaction dataset to simulate anomaly detection. This dataset allowed us to model "concept drift," where the characteristics of fraudulent behavior change over time, requiring the network to adapt.

3.3 Training Protocols and Critical Periods

We implemented a training protocol designed to simulate "Critical Learning Periods" [9].

- Phase 1 (Sensory Acquisition): The network is trained with a high learning rate and high regularization (dropout). This forces the network to learn robust, low-frequency features.
- Phase 2 (Consolidation): The learning rate is decayed exponentially. In the P-CNN, the plasticity term is activated, allowing weights to decay if they are not reinforced.
- Phase 3 (Neurogenesis/Adaptation): New classes (for iNaturalist) or new fraud patterns (for the financial dataset) are introduced. We measured the system's ability to integrate this new data without increasing the error rate on Phase 1 data.

3.4 Evaluation Metrics

We evaluated performance using Top-1 and Top-5 Accuracy for the visual data, and Precision/Recall (F1-Score) for the fraud detection data. Additionally, we introduced a "Forgetting Index" (Φ), calculated as the degradation in accuracy on the original test set after the network was trained on the secondary (new) dataset. A lower Φ indicates better memory retention.

4. Results

4.1 Baseline vs. Plastic-CNN on iNaturalist

The standard VGG-16 architecture achieved a Top-1 accuracy of 72.4% on the initial subset of the iNaturalist data. When introduced to the second subset of species (novel classes), the network required extensive retraining. Without retraining the entire network (fine-tuning only the last layer), accuracy on the new classes was poor (41.2%). When the whole network was retrained, accuracy on the old classes dropped by 18.5%, demonstrating catastrophic forgetting.

In contrast, the P-CNN maintained a Top-1 accuracy of 71.8% on the initial set (slightly lower due to regularization costs). However, upon introduction of the new classes, the Neurogenic Module activated, adding an average of 128 nodes to the dense layers. The P-CNN achieved 68.9% accuracy on the new classes while maintaining 69.5% accuracy on the old classes. The Forgetting Index (Φ) for the VGG-16 was 0.185, while for the P-CNN it was 0.023. This suggests that the neurogenic mechanism successfully compartmentalized the new information.

4.2 Analysis of Critical Periods

Our analysis of the initial training phase corroborates the findings of Achille et al. [9]. We observed that the information stored in the weights (measured via the Fisher Information Matrix) peaked during the first 20 epochs. Networks that were subjected to higher noise injection (dropout) during this critical period formed more generalized features. When we restricted the information flow during this phase, the P-CNN showed greater robustness to adversarial noise later in training. This indicates that the "shape" of the optimization landscape is determined early on, and biological constraints help navigate this landscape toward flatter, more stable minima.

4.3 Fraud Detection and Anomaly Adaptation

In the comparison involving financial data, the traditional algorithms (Random Forest, Logistic Regression) discussed by Patel [1] showed rapid degradation when fraud patterns shifted. The static Neural Network performed better but required periodic offline retraining. The P-CNN, utilizing its plastic weights, was able to adapt to the drift in real-time. As new fraud vectors emerged, the error signal triggered the recruitment of new nodes. The P-CNN maintained an F1-score of 0.88 throughout the simulation, whereas the static NN dropped to 0.65 before retraining was triggered.

5. Discussion

The results presented herein suggest that the integration of biological principles—specifically synaptic plasticity and neurogenesis—into convolutional neural networks offers a viable path toward alleviating catastrophic forgetting. The superior performance of the P-CNN in mixed-task environments aligns with the theoretical predictions of Abbott and Nelson [7] regarding the necessity of homeostatic regulation in synaptic weights. By allowing weights to decay and new nodes to form, the system mimics the "use-it-or-lose-it" and "grow-on-demand" dynamics of the mammalian brain.

5.1 The Theoretical Implications of Neurogenesis in Artificial Architectures

To fully appreciate the divergence in performance between the static VGG-16 baseline and the P-CNN, one must delve deeper into the theoretical underpinnings of neurogenesis as a computational primitive. The incorporation of a "Neurogenic Module" is not merely a heuristic for adding capacity; it represents a fundamental shift in how we conceive of the "state space" of a neural network over time.

In classical deep learning theory, as established by the foundational works of LeCun [3] and further explored by Krizhevsky [4], the capacity of a network is fixed at initialization. The number of parameters is a static hyperparameter. This creates a closed system where the optimization process attempts to map an infinite variance of real-world data onto a finite manifold of weights. When the distribution of the input data changes (concept drift), the network must overwrite existing representations to accommodate the new variance. This is the root cause of catastrophic forgetting. The fixed topology implies that the network resides in a static dimensionality.

However, the work of Aimone, Wiles, and Gage [12, 13] on adult neurogenesis in the dentate gyrus suggests a biological solution to this dimensionality problem: the expansion of the state space itself. Aimone et al. propose that new neurons are not simply "spare parts" but are distinct in their electrophysiological properties. Young neurons are hyperexcitable and possess a lower threshold for long-term potentiation (LTP). This property allows them to capture new, novel features of the environment without interfering with the older, more stable memories encoded in the mature, inhibited neuronal population.

In our P-CNN model, we simulated this by introducing a temporal variable to the node activation function. When the error gradient ∇E exceeded the threshold τ consistently, it signaled that the current manifold was insufficient to separate the data classes. The addition of a new node n_{new} effectively adds a dimension to the hidden layer space. Mathematically, if the activation of the layer is $h \in \mathbb{R}^d$, the introduction of a neurogenic unit transitions the state to $h' \in \mathbb{R}^{d+1}$.

This expansion allows for "orthogonalization" of memory. The old memories occupy a subspace $S_{old} \subset \mathbb{R}^d$. The new memory, encoded primarily by the high-plasticity node n_{new}, occupies a trajectory orthogonal to S_{old}. This explains the drastically lower Forgetting Index (Φ) observed in the P-CNN results. The network did not need to destroy the weights W_{old} to minimize the error for the new classes; it simply projected the residual error into the new dimension provided by n_{new}.

Furthermore, this mechanism relates intimately to the "coding of time" discussed by Aimone et al. [12]. In a biological brain, the age of a neuron correlates with the time at which a memory was encoded. Our P-CNN implicitly captures this temporal metadata. The nodes added later in the training process (during Phase 3) are structurally associated with the "newer" dataset (e.g., the novel species in iNaturalist or the new fraud vector). This suggests that artificial networks could be designed to be naturally "time-aware" without the need for explicit timestamps or recurrent loops (RNNs). The topology itself becomes a timeline of learning events.

5.2 Stability and the Critical Period Hypothesis

The concept of "Critical Learning Periods" explored by Achille et al. [9] provides the second pillar of our discussion. Our results indicated that networks subjected to high information constraints (noise/dropout) during the early epochs (Phase 1) generalized better. This can be understood through the lens of Information Theory and the geometry of the loss landscape.

During the initial phase of training, the weights of a CNN move from a random initialization toward a basin of attraction. Achille et al. argue that the "information path" taken during these early epochs determines the final accessible region of the weight space. If the network is allowed to memorize specific high-frequency details too early (i.e., overfitting to noise), it settles into a sharp, narrow minimum. In such a minimum, any perturbation (such as a slight shift in the dataset distribution) results in a massive increase in loss.

By enforcing a "critical period" of high regularization, we force the network to learn low-frequency, invariant features—shapes, edges, and textures (as visualized by Zeiler and Fergus [2])—rather than memorizing specific pixel configurations. This creates a "flat" minimum. In a flat minimum, the loss function is relatively constant in the local neighborhood of the weights.

The interaction between this critical period and our neurogenic module is particularly noteworthy. We observed that the P-CNN benefited most from neurogenesis when the initial "base" network was trained under these critical period constraints. This suggests a synergistic relationship:

1. The Critical Period ensures the "base" network learns robust, universal features (the "stable core").
2. Neurogenesis handles the specific, novel deviations that arise later (the "plastic periphery").

This dual-system approach mirrors the "stability-plasticity" balance found in the cortex. The visual cortex establishes robust edge detectors early in development (critical period), which remain largely stable throughout life. The hippocampus and association cortices, however, retain high plasticity and neurogenic potential to map new experiences onto this stable visual scaffolding.

5.3 Feedback Systems and Metabolic Constraints

We must also consider the role of feedback and control, as outlined by Astrom and Murray [15]. In standard backpropagation, the feedback is purely error-driven. In our P-CNN, we introduced a "metabolic" decay term λ. This term penalized weights that were not actively contributing to error reduction, effectively pruning weak connections.

This pruning mechanism is essential for preventing "node saturation." Without decay, the addition of new nodes via neurogenesis would eventually lead to a computationally intractable network (the "bloat" problem). The decay term ensures that while the network grows in response to novelty, it also shrinks in response to obsolescence. This dynamic equilibrium aligns with the biological observation that while new neurons are born, many fail to integrate and undergo apoptosis if not stimulated [12].

In the context of the fraud detection simulation [1], this metabolic constraint meant that the network naturally "forgot" obsolete fraud patterns that were no longer present in the data stream, freeing up capacity for active threats. This contrasts with the traditional fraud algorithms which often carry the "technical debt" of outdated rules or decision trees.

5.4 Active Inference vs. Passive Classification

Finally, we interpret our findings through the framework of Active Inference [10]. Adams et al. propose that the brain is a predictive engine that minimizes "free energy" (surprise). Standard CNNs are passive; they map input to output. The P-CNN, with its ability to modify its own architecture, moves a step closer to active inference. The generation of a new node is, in essence, the network "hypothesizing" that a new causal factor exists in the environment which cannot be explained by the current model.

When the prediction error (surprisal) remains high, the system performs a structural update (neurogenesis) to better approximate the generative process of the data. This aligns with the view that intelligent systems must not only optimize parameters but also optimize their structure to match the complexity of the environment. This structural adaptation is what allows biological agents to thrive in open-ended environments, a capability that remains the "Holy Grail" for AGI research.

5.5 Comparison with Traditional Algorithms

It is illustrative to revisit the comparison with traditional algorithms as highlighted in Patel's work [1]. Traditional machine learning models like Support Vector Machines or Random Forests rely on hand-engineered features. While robust, they lack the hierarchical feature extraction of Deep Learning. However, they often possess better interpretability and stability.

The P-CNN bridges this gap. By freezing the "stable core" learned during the critical period, we retain the reliability of established features. By allowing the "plastic periphery" to grow, we achieve the adaptability of deep learning. The results on the fraud dataset confirm that this hybrid approach outperforms both static Deep Learning (which is too rigid) and traditional algorithms (which are too shallow). The P-CNN effectively automates the "feature engineering" for new anomalies that would otherwise require human intervention in a traditional workflow.

5.6 Computational Complexity and Hardware Limitations

While the theoretical benefits are clear, the practical implementation of P-CNNs poses challenges. Standard GPU acceleration (CUDA) is highly optimized for static matrix multiplications. Dynamic graph operations—adding nodes and reshaping tensors on the fly—incur significant computational overhead. In our simulations, the training time per epoch for the P-CNN was approximately 1.4x that of the standard VGG-16.

However, this cost must be weighed against the cost of retraining. Retraining a standard VGG-16 from scratch to accommodate new data is orders of magnitude more expensive than the incremental update of a P-CNN. Furthermore, as neuromorphic hardware (chips designed to mimic biological spiking neural networks) becomes more prevalent, the penalties for dynamic topology will likely vanish. The principles derived from this study—sparse updates, local plasticity, and dynamic routing—are natively supported by neuromorphic architectures, unlike the dense matrix operations of current GPUs.

5.7 Implications for Future Research

The success of the P-CNN invites further research into "developmental AI." Rather than training models to completion, we should investigate "life-long learning" agents that pass through distinct developmental stages (critical periods) and maintain a capacity for structural growth (neurogenesis). Future work should also explore the integration of these mechanisms with unsupervised learning, allowing the network to grow its own topology based on data density rather than supervised error signals.

Moreover, the "Forgetting Index" could be standardized as a metric for evaluating AI robustness. As AI systems are deployed in safety-critical infrastructure (autonomous driving, healthcare, cybersecurity), the ability to learn from new edge cases without unlearning safe behaviors is paramount. The biological mechanisms of plasticity and neurogenesis provide a time-tested blueprint for achieving this reliability.

6. Conclusion

This study demonstrates that the rigidity of traditional Deep Convolutional Neural Networks is not an inherent flaw of artificial intelligence but a consequence of ignoring the temporal and structural dynamics of biological learning. By integrating synaptic plasticity, respecting critical learning periods, and enabling neurogenic adaptation, we created a "Plastic-CNN" that significantly outperforms static architectures in dynamic environments.

The P-CNN showed superior retention of old information while acquiring new features (low Forgetting Index) on the iNaturalist dataset and exhibited real-time adaptability in fraud detection scenarios. These findings suggest that the next generation of AI architectures should look beyond deeper layers and focus on adaptive topologies. The biological brain does not process information through a static flowchart; it is a living landscape that reshapes itself to the contours of experience. Artificial systems that aspire to general intelligence must inevitably embrace this principle of structural plasticity.

The convergence of neuroscience and computer science, as evidenced by the synthesis of concepts from Abbott, Achille, Aimone, and LeCun, provides a fertile ground for innovation. We conclude that "learning rich features without labels" or with minimal labels is best achieved not just by better optimizers, but by architectures that evolve alongside their data.

References

1. Dip Bharatbhai Patel. (2025). Comparing Neural Networks and Traditional Algorithms in Fraud Detection. The American Journal of Applied Sciences, 7(07), 128–132. https://doi.org/10.37547/tajas/Volume07Issue07-13

2. Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, and Gang Wang. Recent advances in convolutional neural networks. arXiv preprint arXiv:1512.07108, 2015.

3. Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In European conference on computer vision, pages 818–833. Springer, 2014.

4. Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. Neural computation, 1(4):541–551, 1989.

5. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.

6. Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.

7. Grant Van Horn, Oisin Mac Aodha, Yang Song, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist challenge 2017 dataset. arXiv preprint arXiv:1707.06642, 2017.

8. Abbott, L. F. & Nelson, S. B. (2000), ‘Synaptic plasticity: taming the beast’, Nature Neuroscience 3, 1178–1183.

9. Abraham, W. C. & Robins, A. (2005), ‘Memory retention and weight plasticity in ANN simulations’, Trends in Neurosciences 28(2), 73–78.

10. Achille, A., Rovere, M. & Soatto, S. (2017), Critical learning periods in deep neural networks, arXiv:1711.08856.

11. Adams, R. A., Friston, K. J. & Bastos, A. M. (2015), ‘Active inference, predictive coding and cortical architecture’, In M. F. Casanova and I. Opris (Eds.), Recent Advances on the Modular Organization of the Cortex. Springer Netherlands.

12. Aimone, J. B., Wiles, J. & Gage, F. H. (2006), ‘Potential role for adult neurogenesis in the encoding of time in new memories.’, Nature Neuroscience 9, 723–727.

13. Aimone, J. B., Wiles, J. & Gage, F. H. (2009), ‘Computational influence of adult neurogenesis on memory encoding.’, Neuron 61, 187–202.

14. Altman, J. (1963), ‘Autoradiographic investigation of cell proliferation in the brains of rats and cats’, The Anatomical Record 145(4), 573–591.

15. Astrom, K. J. & Murray, R. M. (2010), Feedback Systems: An Introduction for Scientists and Engineers, Princeton University Press