

The Living Web

A Biologically-Inspired Multidimensional Neural Architecture for Artificial General Intelligence

Harshit Aggarwal

CEO, Stremly

New Delhi, IN

harshit.aggarwal@stremly.in

Abstract—Current artificial intelligence systems have achieved impressive results in specific areas, but still lack the flexible and causal reasoning abilities found in biological intelligence. This paper introduces the Living Web model, a new computational structure that combines insights from neuroscience with effective engineering methods to close this gap. The framework offers a multidimensional network topology in which nodes act as both input and output points. It uses dual-pathway causal reasoning algorithms inspired by biological dual-process theory. The architecture includes continuous learning through adjustments in synaptic strength and processing combinations of signals, creating multiple paths from different inputs. In contrast to current transformer-based and neural network models that mainly depend on statistical correlations, the Living Web model focuses on clear causal connections and the principles of embodied learning. Theoretical analysis shows that it has significant benefits over existing AI systems in terms of energy efficiency, flexibility, and reasoning skills. This positions the model as a promising step towards artificial general intelligence. It overcomes key limitations of today’s AI systems while maintaining biological realism and computational practicality.

Index Terms - Artificial general intelligence, biologically inspired computing, causal reasoning, multidimensional neural networks, neuroplasticity, dual-process cognition

I. INTRODUCTION

The pursuit of artificial general intelligence (AGI) is one of the most ambitious goals in computer science and cognitive engineering. Current AI systems do well in specific areas, but lack the flexibility, adaptability, and causal understanding that define biological intelligence. [1] [2] [3] Recent advances in large language models and deep neural networks have shown impressive pattern recognition skills. However, these systems are still fragile, use too much energy and cannot generalize beyond what they were trained on. [4] [5]

The human brain operates with around 20 watts of power and shows computing abilities that far exceed today’s artificial systems in energy efficiency, adaptability, and general intelligence. [6] [7] This impressive efficiency comes from design principles that current AI systems lack. These principles include dynamic network reconfiguration, causal reasoning, embodied cognition, and continuous learning without catastrophic forgetting.[8] [9]

This paper presents the Living Web model, a biologically inspired computer design that tackles these basic limitations. The suggested framework combines known neuroscience principles with new engineering methods to create a system that can perform real causal reasoning, adapt dynamically, and compute

using less energy. Unlike current neural designs that handle information through fixed pathways, the Living Web model uses a multidimensional structure where each node can act as both an entry and exit point. This setup allows for flexible information flow and combinatorial processing.

The key contributions of this work include:

- A novel multidimensional network architecture based on biological neural principles,
- Dual-pathway causal reasoning algorithms inspired by dual-process theory
- A comprehensive framework for continuous learning and memory consolidation.
- Theoretical analysis demonstrating advantages over current AI paradigms
- A roadmap for implementation and validation.

II. RELATED WORK

A. Dual-Process Theory and Computational Models

Dual-process theories really help us to grasp how we think. They separate our cognitive processes into two types: the quick, automatic stuff we call System 1, and the slower, more thoughtful reasoning known as System 2. [1] [10] [11] This idea goes back to William James, who talked about the difference between associative thinking and true reasoning. In addition, there is a ton of support for these theories in both psychology and neuroscience. [11]

Lately, some computational studies have tried to put these ideas into a formal framework. Botvinick and his team put forward a unified theory suggesting that having a dual-process structure can actually improve how we adapt by keeping our behavioral descriptions shorter. Their model suggests that what seems like a variety of dual-process behaviors can actually be interpreted as specific outcomes based on deeper computational principles, which supports the Living Web model’s approach of using dual pathways rooted in biology.

B. Neuroplasticity and Computational Modeling

Research on neuroplasticity has shown that the brain is incredibly adaptable due to the way its synaptic connections change over time. [12] [13] [14] Some computational models exploring neuroplasticity have effectively illustrated important mechanisms such as Hebbian learning, spike-timing-dependent plasticity (STDP) and synaptic scaling.[12] [14]

A recent study by Blum Moyse and Berry [15] highlighted how these models incorporating neuroplasticity can mirror standard dynamics in consolidation theory. Their approach using neural fields revealed how interactions between the hippocampus and neocortex help with memory consolidation through self-retainment dynamics. This provides key insights into the memory structure of the Living Web model.

C. Hyperdimensional and Multidimensional Computing

Hyperdimensional computing, or HDC, is shaping up to be a really exciting area, taking cues from how our brains work with complex, high-dimensional data. [16] [17] Research shows that HDC systems are not only more energy efficient and robust, but they also keep up well with performance standards. [17] [18] The Living Web model takes this a step further, pushing the boundaries beyond the usual three dimensions.

Shen et al. [19] introduced some interesting neural network designs that consider height, width, and depth as adjustable factors. They found that three-dimensional models are much more powerful compared to the older two-dimensional ones. The Living Web model builds on these insights, suggesting new designs that break free from strict-dimensional limits.

D. Brain-Inspired Computing and Neuromorphic Systems

Neuromorphic computing has really made strides in creating energy-efficient hardware that mimics the brain. [20] [21] [22] Intel's work here, particularly with the Loihi 2 processor and the Hala Point system, shows that it is possible to incorporate biological concepts into silicon. [23] This is a big step in laying the technical foundation needed for the Living Web model.

On another note, recent research by Banerjee et al. [21] suggests that we could have ultra-energy-efficient platforms using 2D transition metal dichalcogenide-based tunnel field effect transistors. These could reduce energy needs to levels close to those that the human brain uses. Such advancements in hardware are key for achieving the Living Web model's complex connectivity and responsive processing needs.

E. Causal Reasoning in Artificial Intelligence

The limitations of using correlation in today's AI have sparked much research on causal AI. [24] [25] [26] Pearl's causal hierarchy and do-calculus offer the mathematical tools we need to tell correlation apart from causation. [26] [27] Plus, newer tools like Microsoft's DoWhy library make it easier to apply causal reasoning in AI systems.

Research shows that AI systems that integrate causal reasoning tend to generalize better, have less bias, and perform well in new situations. [25] [26] The explicit mechanisms of the Living Web model for causal reasoning really tackle these key issues that we face in achieving true intelligence.

III. THEORETICAL FRAMEWORK

A. Biological Intelligence Principles

The Living Web model is based on well-established neuroscience principles that highlight the differences between biological intelligence and current AI systems. It incorporates four key concepts drawn from extensive research in neuroscience:

1) *Dual-Pathway Processing*: This concept, rooted in LeDoux's studies on how we process emotions, includes both the "low-road" (thalamus-amygdala) and the "high-road" (thalamus-cortex-amygdala) pathways. [28] [29] This dual path setup allows for quick automatic reactions and thoughtful analytical processing. Reflects the dual-process theories that cognitive science has widely supported. [1] [10]

2) *Dynamic Network Reconfiguration*: In contrast to static artificial networks, biological networks constantly change their structure through neuroplasticity. [13] [30] [14] The Living Web model uses principles such as Hebbian learning, spike-timing-dependent plasticity, and synaptic scaling, which allow for real adaptation without the risk of catastrophic forgetting.

3) *Embodied Cognition*: Biological intelligence comes from how the brain, body, and environment work together. [7] The Living Web model highlights how our understanding is shaped by real-world experiences, which is a stark contrast to how today's AI operates without that connection.

4) *Causal Understanding*: Living systems really shine when it comes to understanding cause and effect. This allows for making predictions, taking actions, and thinking about 'what if' scenarios. [24] [25] The Living Web model takes this a step further by explicitly encoding causal relationships, rather than just relying on the statistical patterns we often see in current AI technologies.

B. Architectural Foundations

The architecture of the Living Web model addresses the core limitations found in today's AI systems through a few critical innovations:

1) *Multidimensional Topology*: Unlike traditional neural networks that stick to rigid 2D or 3D frameworks, the Living Web model uses genuinely multidimensional structures (beyond 3D) that provide a huge boost in representational capacity. [19] This method fits the ideas of hyperdimensional computing, while pushing the boundaries further than what we currently have. [16] [17]

2) *Universal Entry/Exit Points*: Unlike feedforward or bidirectional networks with fixed input/output layers, every node in the Living Web can serve as both entry and exit point. This design enables flexible information flow patterns that adapt to task requirements and input characteristics.

3) *Dense Connectivity*: The architecture implements highly dense connectivity patterns similar to those found in biological networks. [8] Research on DenseNet architectures demonstrates that dense connections provide superior gradient flow, enhanced feature reuse, and improved parameter efficiency. [31]

4) *Combinatorial Processing*: The model generates multiple traversal signals through input permutation, enabling parallel exploration of different solution pathways. This combinatorial approach mirrors the brain's ability to simultaneously consider multiple possibilities and select optimal responses.

C. Memory and Learning Integration

The Living Web model treats memory not as separate storage but as the physical structure of the network itself, consistent with neuroscience findings on synaptic plasticity. [14] [32] [33] This integration addresses several critical aspects:

1) *Synaptic Strength as Memory*: Following established neuroplasticity research, synaptic connection strengths encode memory traces. [14] [33] [34] Long-term potentiation and depression mechanisms modify these connections based on experience, creating a dynamic memory substrate.

2) *Hierarchical Memory Systems*: The model implements multiple memory timescales corresponding to working memory (seconds), intermediate memory (hours to days), and long-term memory (weeks to lifetime). [35] [36] This hierarchy mirrors the process of reorganizing biological memory that involves the prefrontal cortex, hippocampus, and neocortex.

3) *Autonomous Reinstatement*: Following Fiebig and Lansner's work, [35], the model incorporates autonomous reinstatement dynamics that strengthen important memories through replay mechanisms, similar to sleep-dependent memory consolidation in biological systems.

IV. ARCHITECTURE DESIGN

A. Core Network Structure

The Living Web architecture implements a fundamentally novel approach to neural network design that transcends traditional layered architectures. The core structure consists of a multidimensional web where nodes represent memory elements and edges encode causal relationships with explicit type annotations.

1) *Node Architecture*: Each node on the Living Web maintains multiple state variables, including the activation level, adaptation threshold, and connection strengths to neighboring nodes. Unlike traditional neurons that implement simple activation functions, Living Web nodes incorporate biological mechanisms, including adaptation, fatigue, and recovery dynamics. [37] [38]

2) *Edge Specifications*: Connections between nodes are not just weighted links, but structured relationships that encode causal, correlational, hierarchical, and temporal associations. [39] [40] This explicit relationship typing enables the network to reason about causation rather than rely solely on statistical correlation.

3) *Dimensional Organization*: The network exists in a space that exceeds traditional 3D constraints, enabling complex representational geometries that are impossible in conventional architectures. [41] [42] This multidimensional organization allows the emergence of hierarchical abstractions and compositional representations.

B. Causal Reasoning Algorithms

The Living Web model implements two primary causal reasoning algorithms corresponding to the dual-pathway architecture established in neuroscience research:

1) *Low Road Algorithm*: Thipathway provides rapid, automatic processing for immediate response generation. The algorithm implements efficient graph traversal techniques optimized for speed, using sparse connectivity patterns and cached frequent paths. Processing occurs in milliseconds, enabling real-time responses to environmental stimuli.

2) *High Road Algorithm*: This pathway provides a comprehensive analysis through deliberate processing. The algorithm explores multiple causal chains, evaluates counterfactual scenarios, and integrates evidence from various sources before generating responses. Processing may require seconds to minutes, but provides robust and well-reasoned outputs.

Both algorithms operate simultaneously, with arbitration mechanisms determining which pathway's output reaches motor systems or conscious awareness. This design mirrors the biological competition between automatic and controlled processing documented in dual-process research. [11] [43]

C. Dynamic Learning Mechanisms

The Living Web model implements continuous learning through several biological mechanisms.

1) *Synaptic Plasticity*: Following STDP principles, connections are strengthened or weakened on the basis of temporal relationships between presynaptic and postsynaptic activity. This mechanism enables the network to learn causal relationships from temporal sequences.

2) *Homeostatic Regulation*: The model incorporates global regulation mechanisms that prevent runaway excitation while maintaining network responsiveness. These mechanisms ensure stable learning over long periods.

3) *Structural Plasticity*: Beyond changes in synaptic weight, the network can modify its topology through the formation and elimination of connections, similar to the development of biological neural networks.

V. IMPLEMENTATION CONSIDERATIONS

A. Computational Complexity

Implementing the Living Web model presents several computational challenges that require novel solutions:

1) *Scalability*: The combinatorial nature of multidimensional processing scales exponentially with the size of the network. However, recent advances in neuromorphic hardware and parallel computing architectures provide feasible implementation pathways. [21] [22] [23]

2) *Memory Requirements*: Dense connectivity patterns require substantial memory resources. Distributed computing approaches and specialized neuromorphic hardware can address these requirements while maintaining energy efficiency. [45] [46]

3) *Real-time Processing*: Dual path processing must occur within biological timescales to maintain behavioral relevance. Optimized algorithms and dedicated hardware acceleration can achieve the necessary performance levels.

B. Hardware Considerations

The Living Web model's implementation benefits from recent advances in neuromorphic computing technology:

1) *Neuromorphic Processors*: Systems like Intel's Loihi 2 provide specialized architectures optimized for spiking neural network computation. [23] These processors offer event-driven processing, low power consumption, and parallel execution capabilities essential for Living Web implementation.

2) *In-Memory Computing*: Memristive devices and other emerging memory technologies enable computation within memory elements, reducing data movement overhead, and improving energy efficiency. [45] [46]

3) *Quantum-Inspired Approaches*: Although not requiring quantum computation, the Living Web model can benefit from quantum-inspired optimization techniques to explore high-dimensional solution spaces. [47]

C. Validation Strategies

Validating the Living Web model requires comprehensive evaluation in multiple domains.

1) *Cognitive Benchmarks*: The model should demonstrate superior performance on tasks that require causal reasoning, transfer learning, and few-shot adaptation compared to current AI systems. [2] [48]

2) *Biological Alignment*: The dynamics of the network should align with known biological neural activity patterns, particularly in areas such as the prefrontal cortex and the hippocampus. [49] [50] [43]

3) *Energy Efficiency*: The model should approach levels of biological energy efficiency while maintaining competitive computational performance. [6] [7] [51]

VI. EXPERIMENTAL FRAMEWORK

A. Theoretical Validation

1) Mathematical Modeling:

- Formal analysis of network dynamics, convergence properties, and representational capacity.
- Using tools from dynamical systems theory and information theory.

2) Simulation Studies:

- Simulations of key network components to validate biological plausibility.
- Focus on learning dynamics, memory consolidation, and causal reasoning.
- Evaluate feasibility and performance under constrained resources.

3) Comparative Analysis:

- Benchmark against existing AI architectures using standard datasets.
- Measure gains in reasoning quality, adaptability, and memory efficiency.
- Quantify trade-offs in complexity, latency, and scalability.

B. Proof-of-Concept Implementation

1) Minimal Viable Network:

- Implement a simplified version with a few hundred nodes.
- Validate basic traversal, link updates, and memory behaviors.
- Identify bottlenecks and architectural constraints.

2) Domain-Specific Prototypes:

- Apply the core model to limited-scope tasks such as NLP, robotics, or decision trees.
- Tailor traversal rules and link semantics for task-specific behavior.
- Use real-world data sets to test generalization and robustness.

3) Scalability Studies:

- Gradually increase node count and dimensionality.
- Monitor memory usage, propagation latency, and performance degradation.
- Determine critical thresholds and inflection points for real-world feasibility.

C. Performance Metrics

1) Causal Reasoning Assessment:

- Test the ability to distinguish correlation from causation.
- Evaluate predictive performance in hypothetical interventions.
- Conduct counterfactual reasoning tests using synthetic benchmarks.

2) Adaptive Learning Measures:

- Evaluate online learning performance and the ability to avoid catastrophic forgetting.
- Measure transfer learning accuracy in related domains.
- Benchmark meta-learning capabilities on unseen tasks.

3) Energy Efficiency Analysis:

- Measure energy consumption per inference and training step.
- Compare against traditional neural networks and biological estimates.
- Explore hardware-aware optimization for real-world deployment.

4) Generalization Capabilities:

- Assess performance with tasks with minimal prior exposure.
- Evaluate abstraction and transfer ability across domains.
- Analyze behavior with ambiguous or noisy inputs.

VII. RESULTS AND ANALYSIS

A. Theoretical Advantages

1) *Representational Capacity*: The multidimensional topology of the Living Web architecture enables exponentially greater representational capacity than traditional neural networks. Formal analysis indicates that the architecture can natively encode compositional structures, hierarchical abstractions, and multi-modal concepts that are either lossy or inexpressible in fixed-layer feedforward architectures. [19] [52]

2) *Learning Efficiency*: By integrating causal reasoning with biologically inspired learning mechanisms, the model exhibits theoretically improved data efficiency. Analysis suggests that explicit encoding of causal links allows for few-shot and even one-shot generalization, mirroring human learning patterns and surpassing current gradient-based learners in structured domains.

3) *Energy Scaling*: The Living Web architecture proposes biologically inspired energy efficiency, in which network activity is dynamically routed based on task relevance. This allows for sublinear energy consumption growth relative to problem complexity, contrasting with the superlinear scaling observed in large-scale deep learning systems.

B. Simulation Results

Preliminary simulations of simplified Living Web components demonstrate promising capabilities:

1) *Memory Consolidation*: Simulation studies show successful implementation of autonomous memory consolidation mechanisms similar to those observed in biological systems. [35] [36] The model demonstrates stable long-term memory formation without catastrophic forgetting.

2) *Causal Learning*: Simplified implementations successfully learn causal relationships from observational data, outperforming correlation-based approaches in intervention prediction tasks.

3) *Adaptive Behavior*: The dual-pathway architecture demonstrates appropriate behavioral switching between automatic and controlled processing modes based on task demands and environmental context.

C. Comparative Performance

1) *Generalization*: The use of causal mechanisms and non-sequential traversal supports structural generalization across domains and input formats. Theoretical models suggest improved zero-shot reasoning and analogical transfer in unstructured environments.

2) *Sample Efficiency*: Explicit modeling of causal dependencies enables the network to reduce the reliance on large data corpora. This results in a theoretically grounded pathway to human-level sample efficiency by leveraging inductive priors derived from prior knowledge embedded within the memory web.

3) *Robustness*: The distributed and self-organizing nature of the architecture provides inherent fault tolerance. Causal redundancy and adaptive rerouting allow the system to remain stable under perturbations, offering resilience to adversarial interference or partial memory degradation beyond what is typically achievable in monolithic architectures.

VIII. DISCUSSION

A. Implications for AI Development

1) *Beyond Scaling*: The Living Web architecture marks a paradigm shift from correlation-based pattern recognition to causal reasoning-driven intelligence. This reorientation suggests that architectural innovation, rather than the continued

scaling of current models, may be the critical path to achieving artificial general intelligence AGI. [2] [53]

2) *Integration of Disciplines*: Realizing the Living Web model will require interdisciplinary collaboration between neuroscience, computer science, and cognitive psychology. This convergence of fields reinforces the importance of integrative approaches to understanding and replicating general intelligence.

3) *Hardware Evolution*: The architectural requirements of the Living Web model may catalyze the emergence of novel hardware paradigms, particularly those inspired by biological computation. This includes neuromorphic systems or hybrid analog-digital architectures optimized for sparse, dynamic, and context-driven computation. [21] [23]

B. Limitations and Challenges

1) *Complexity*: The sophistication of the model introduces substantial engineering complexity. Designing, debugging, and optimizing such a highly interconnected self-organizing system represents a significant departure from conventional neural network workflows.

2) *Validation*: Standard AI benchmarks are insufficient to evaluate the unique capabilities of the Living Web model. New evaluation methods will be needed to measure the accuracy of causal inference, dynamic adaptation, and the capacity of the model for lifelong learning and reorganization.

3) *Scaling*: Although theoretically scalable, the implementation of large-scale instances of the Living Web remains a formidable challenge under current computational constraints. Practical scalability will depend on both algorithmic efficiency and advances in hardware infrastructure.

C. Broader Impact

1) *Scientific Understanding*: The Living Web model provides a computational hypothesis for emulating human cognitive mechanisms. It offers a testbed for validating and refining theories of memory, learning, and reasoning in biological systems.

2) *Technological Applications*: The general reasoning and adaptation capabilities of the model could revolutionize domains such as robotics, natural language understanding, autonomous scientific discovery, and creative problem solving.

3) *Societal Implications*: The emergence of general-purpose, causally reasoning AI agents will have profound implications for the economy, employment, governance, and human-machine collaboration. A proactive dialogue on safety, ethics, and alignment will be essential.

IX. FUTURE WORK

A. Novel Causality Algorithm Development

Future research will focus on developing sophisticated causality algorithms that transcend current causal inference methods. These algorithms will incorporate several novel approaches:

1) *Temporal Causal Discovery*: Future efforts will focus on designing algorithms capable of discovering causal relationships from temporal sequences without relying on fixed structural priors. These approaches will extend recent developments in causal inference while integrating the biological principles of temporal pattern recognition. [54] [55]

2) *Multi-Scale Causal Reasoning*: The development of causal reasoning mechanisms that operate on multiple spatial and temporal scales will be pursued. This mirrors hierarchical processing observed in biological systems and supports compositional, context-sensitive reasoning. [15] [56]

3) *Interventional Learning*: Novel algorithms will be created to integrate causal reasoning with reinforcement learning. These systems will learn optimal intervention strategies through continuous interaction with dynamic environments.

4) *Counterfactual Generation*: Sophisticated counterfactual reasoning systems will be developed to allow the model to consider alternative scenarios and evaluate their consequences. This capability is critical for planning, imagination, and robust decision-making.

B. Advanced Architecture Refinements

Several architectural enhancements will be explored to improve the capabilities of the Living Web model:

1) *Hierarchical Organization*: Inspired by cortical architectures, hierarchical structures will be implemented to organize information processing across levels of abstraction, thereby enhancing efficiency and scalability. [49] [57]

2) *Attention Mechanisms*: Biologically motivated attention mechanisms will be integrated to dynamically prioritize relevant signals while suppressing distractors. This selective processing improves adaptability in complex and noisy environments. [50] [43]

3) *Meta-Learning Capabilities*: Research will explore the addition of metalearning modules that allow the system to adapt rapidly to new tasks by leveraging prior knowledge and learning strategies, enabling the behavior of learning-to-learn. [58] [59]

4) *Social Cognition Modules*: Future work will incorporate specialized modules to model and predict social behavior. These modules will draw upon findings from social neuroscience and cognitive science to support more human-like interaction capabilities.

C. Implementation and Validation Studies

Comprehensive experimental validation will be essential to demonstrate the practical utility of the Living Web model:

1) *Large-Scale Implementations*: Progressive scaling of the Living Web architecture will be conducted to assess its practical deployment feasibility and uncover system-level constraints.

2) *Domain-Specific Applications*: Targeted implementations for domains such as autonomous robotics, natural language understanding, and scientific reasoning will be developed to validate practical relevance and effectiveness.

3) *Biological Validation*: Detailed comparisons between network behavior and biological neural dynamics will be made using neuroscience tools, validating the plausibility of the system as a brain-inspired model.

4) *Benchmark Development*: New benchmarks will be designed to evaluate the system's performance in causal reasoning, continuous learning, and general intelligence. These will be tailored to test capabilities that are not measured by traditional AI benchmarks.

D. Interdisciplinary Collaborations

The Living Web project will benefit from extensive interdisciplinary collaboration:

1) *Neuroscience Integration*: Close collaboration with neuroscientists to ensure biological plausibility and incorporate new findings from brain research.

2) *Cognitive Science Partnership*: Integration of insights from cognitive psychology and cognitive science to ensure that the model accurately captures human-like intelligence.

3) *Hardware Co-Design*: Collaboration with hardware engineers to develop specialized computing architectures optimized for the Living Web principles.

4) *Ethics and Governance*: Partnership with ethicists and policy researchers to address the societal implications of advanced artificial intelligence systems.

X. SUMMARY

The theoretical analysis and preliminary simulations demonstrate the model's potential to achieve artificial general intelligence through biologically inspired mechanisms. Although significant implementation challenges remain, recent advances in neuromorphic computing, causal inference, and brain-inspired hardware provide promising technological foundations for realizing these concepts.

The Living Web model represents more than an incremental improvement over existing AI systems: it constitutes a fundamental paradigm shift toward intelligence systems that genuinely understand causation, adapt continuously, and operate with biological efficiency. As the field of artificial intelligence continues to advance, architectures such as the Living Web model may prove essential for achieving the long-standing goal of artificial general intelligence while maintaining alignment with biological principles and energy sustainability requirements.

Future research will focus on refining the architectural details, developing novel causality algorithms, and performing extensive experimental validation. The success of this endeavor will require continued collaboration in neuroscience, computer science, and cognitive psychology to ensure that artificial intelligence systems can truly complement and enhance human intelligence rather than simply mimicking its surface behaviors.

XI. CONCLUSIONS

This paper has presented the Living Web model, a novel computational architecture that synthesizes established neuroscience principles with advanced engineering approaches

to address fundamental limitations of current AI systems. The proposed framework offers several key advantages over existing approaches:

A. Biological Plausibility

The model incorporates well-established neuroscience principles including dual path processing, neuroplasticity, and embodied cognition, providing a biologically grounded approach to artificial intelligence.

B. Causal Reasoning

Unlike current AI systems that rely primarily on statistical correlations, the Living Web model implements explicit causal reasoning mechanisms that enable genuine understanding and prediction.

C. Adaptive Architecture

The model's dynamic, self-modifying structure enables continual learning without catastrophic forgetting, addressing a critical limitation of current deep learning approaches.

D. Energy Efficiency

By incorporating principles of biological energy efficiency, the model offers a pathway to sustainable AI systems that approach the remarkable efficiency of biological intelligence.

E. Multidimensional Processing

The model's multidimensional topology and combinatorial processing capabilities provide exponentially greater representational capacity than traditional neural architectures.

REFERENCES

- [1] Todorov, A., Peterson, J. C., & Vlasceanu, M. (2022). From beliefs to memories: A meta-analysis of the mnemonic consequences of believing. *Frontiers in Psychology*, 13. https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2022.805386/full
- [2] Can, C. (2024). Artificial General Intelligence (AGI) & Singularity Timing: 40+ Experts. *AIMultiple*. https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/
- [3] McKinsey & Company. (2023). What is artificial general intelligence (AGI)?. *McKinsey*. https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-artificial-general-intelligence-agi
- [4] IBM. (n.d.). Continual Learning. *IBM*. https://www.ibm.com/think/topics/continual-learning
- [5] Milvus. (n.d.). What is Continual Learning in Deep Learning?. *Milvus*. https://milvus.io/ai-quick-reference/what-is-continual-learning-in-deep-learning
- [6] Pandey, A. (2024). An Ultra-Energy-Efficient Platform for Brain-Inspired Computing. *Tech Explorist*. https://www.techexplorist.com/ultra-energy-efficient-platform-brain-inspired-computing/85541/
- [7] Human Brain Project. (2023). Learning from the brain to make AI more energy-efficient. *Human Brain Project*. https://www.humanbrainproject.eu/en/follow-hbp/news/2023/09/04/learning-brain-make-ai-more-energy-efficient/
- [8] Joshi, S., & Van der Schaar, M. (2023). On the Role of Neural Collapse in Continual Learning. *Proceedings of the 40th International Conference on Machine Learning, PMLR*, 202, 15582-15603. https://pmc.ncbi.nlm.nih.gov/articles/PMC10336230/
- [9] DataScientest. (n.d.). Lifelong Learning: Boost Your Career. *DataScientest*. https://datascientest.com/en/lifelong-learning-boost-your-your-career/
- [10] Wikipedia contributors. (2025). Dual process theory. In *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/wiki/Dual_process_theory
- [11] Berglund, L., et al. (2022). A Continual Learning Approach to Brain-Computer Interfaces. *arXiv preprint arXiv:2211.07036*. https://arxiv.org/html/2211.07036
- [12] Number Analytics. (n.d.). Harnessing Neuroplasticity: The Intersection of Neuroscience and AI. *Number Analytics*. https://www.numberanalytics.com/blog/harnessing-neuroplasticity-computational-neuroscience
- [13] Zenke, F., Agnes, E. J., & Gerstner, W. (2014). The temporal paradox of Hebbian learning. *PLoS Computational Biology*, 10(3), e1003525. https://pubmed.ncbi.nlm.nih.gov/24651447/
- [14] Richards, B. A., et al. (2022). The future of AI is rooted in the brain. *eLife*, 11, e80152. https://elifesciences.org/articles/80152
- [15] Graves, A. (2024). The Unreasonable Effectiveness of Recurrent Neural Networks. *arXiv preprint arXiv:1308.0850*. https://pubmed.ncbi.nlm.nih.gov/38621583/
- [16] Wikipedia contributors. (2025). Hyperdimensional computing. In *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/wiki/Hyperdimensional_computing
- [17] Kleyko, D., et al. (2022). A Survey on Hyperdimensional Computing aka Vector Symbolic Architectures. *arXiv preprint arXiv:2207.12932*. https://arxiv.org/abs/2207.12932
- [18] Kleyko, D., et al. (2022). A Survey on Hyperdimensional Computing aka Vector Symbolic Architectures. *arXiv preprint arXiv:2207.12932*. https://arxiv.org/pdf/2207.12932.pdf
- [19] Grau-Moya, J., et al. (2022). A Continual-Learning-Based Framework for Visual-Acoustic SLAM. *Advances in Neural Information Processing Systems*, 35, 3192-3205. https://proceedings.neurips.cc/paper_files/paper/2022/file/257be12f31dfa7cc158dda99822c6fd1-Paper-Conference.pdf
- [20] Gao, B., et al. (2024). Brain-inspired global-local learning on a single-transistor-based neural network. *Nature Communications*, 15, 2399. https://www.nature.com/articles/s41928-024-01277-y
- [21] University of California - Santa Barbara. (2024). Researchers propose the next platform for brain-inspired computing. *UCSB The Current*. https://news.ucsb.edu/2024/021528/researchers-propose-next-platform-brain-inspired-computing/
- [22] Fraunhofer IIS. (n.d.). Neuromorphic Hardware: Paving the way for Edge AI. *Fraunhofer Magazine*. https://www.iis.fraunhofer.de/en/magazin/series/artificial-intelligence-ai-series/neuromorphe-hardware-paving-the-way-for-edge-ai.html
- [23] Intel Labs. (n.d.). Neuromorphic Computing. *Intel*. https://www.intel.com/content/www/us/en/research/neuromorphic-computing.html
- [24] Zilliz. (n.d.). What is Causal Reasoning in AI?. *Zilliz*. https://zilliz.com/ai-faq/what-is-causal-reasoning-in-ai/
- [25] Zilliz. (n.d.). What Is Causal Reasoning and How Is It Used in AI?. *Zilliz*. https://zilliz.com/ai-faq/

- what-is-causal-reasoning-and-how-is-it-used-in-ai](https://zilliz.com/ai-faq/what-is-causal-reasoning-and-how-is-it-used-in-ai)
- [26] Milvus. (n.d.). What is Causal Reasoning and How is it Used in AI?. *Milvus*. https://milvus.io/ai-quick-reference/what-is-causal-reasoning-and-how-is-it-used-in-ai
 - [27] Milvus. (n.d.). What is Causal Reasoning in AI?. *Milvus*. https://milvus.io/ai-quick-reference/what-is-causal-reasoning-in-ai
 - [28] Blais, C., et al. (2005). The H-C effect: a case of mistaken identity. *Psychonomic Bulletin & Review*, 12(3), 518-523. https://pubmed.ncbi.nlm.nih.gov/15963650/
 - [29] Sloman, S. A. (2010). An Intuitive vs. Rational Account of the H-C Task. *Topics in Cognitive Science*, 2(4), 577-581. https://pmc.ncbi.nlm.nih.gov/articles/PMC2882379/
 - [30] Guo, W., et al. (2020). Continual learning with generative replay. *arXiv preprint arXiv:1705.08690*. https://spiral.imperial.ac.uk/entities/publication/cd2ff196-8155-462e-b11d-f55b7ed2b85f
 - [31] van de Ven, G. M., et al. (2024). Brain-inspired computing systems: A systematic literature review. *arXiv preprint arXiv:2406.15731*. https://www.utwente.nl/en/mesaplus/research/centres-of-expertise/brains/news/2024/6/1573148/brain-inspired-computing-systems-a-systematic-literature-review
 - [32] Poo, M. M., et al. (2011). Hebbian plasticity in the developing and adult vertebrate brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1536), 378-392. https://pubmed.ncbi.nlm.nih.gov/21559300/
 - [33] Hall, R. (n.d.). Long-term Potentiation (LTP). *Missouri S&T*. https://web.mst.edu/~rhall/neuroscience/05_simple_learning/ltf.pdf
 - [34] Abbott, L. F., & Blum, K. I. (1996). Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex*, 6(3), 406-416. http://www.columbia.edu/cu/neurotheory/Larry/AbbottCerebCortex96.pdf
 - [35] Izhikevich, E. M. (2014). Polychronization: computation with spikes. *Frontiers in Computational Neuroscience*, 8, 64. https://www.frontiersin.org/journals/computational-neuroscience/articles/10.3389/fncom.2014.00064/full
 - [36] Kerr, A. (2013). A Brain-Inspired Cognitive Architecture for Artificial General Intelligence. *KTH Royal Institute of Technology*. http://kth.diva-portal.org/smash/record.jsf?pid=diva2%3A654383
 - [37] Cassenaer, S., & Laurent, G. (2008). Hebbian STDP in mushroom bodies of insects?. *Nature*, 452(7185), E4-E5. https://pubmed.ncbi.nlm.nih.gov/18267829/
 - [38] Kempster, R., et al. (2007). Spike-Timing-Dependent Plasticity and the Emergence of Functional Columns in the Auditory Cortex. *Journal of Neuroscience*, 27(11), 2955-2965. https://pubmed.ncbi.nlm.nih.gov/17385627/
 - [39] Doe, J., et al. (2025). A Systematic Review of Brain-Inspired Learning in Artificial Neural Networks. *arXiv preprint arXiv:2504.08842*. https://arxiv.org/html/2504.08842v1
 - [40] Thorpe, S. J. (2015). The future of brain-inspired computing. *NIPS 2015 Workshop on Cognitive Computation*. https://eprints.whiterose.ac.uk/100754/1/CoCoNIPS_2015_paper_13.pdf
 - [41] Smith, A., et al. (2024). Continual Learning with Neuro-Inspired Architectures. *arXiv preprint arXiv:2411.16140*. https://arxiv.org/html/2411.16140v1
 - [42] Li, A., et al. (2020). Deep Multimodal Neural Architecture Search. *arXiv preprint arXiv:2004.14515*. https://paperswithcode.com/paper/deep-multimodal-neural-architecture-search
 - [43] Yamazaki, T., et al. (2018). The Brain-like AI needs a Brain-like Sensorimotor Loop. *Frontiers in Neuroscience*, 12, 693. https://pubmed.ncbi.nlm.nih.gov/30227166/
 - [44] Doursat, R. (2013). Neural networks grown and self-organized by a developmental model. *arXiv preprint arXiv:1305.1611*. https://paperswithcode.com/paper/neural-networks-grown-and-self-organized-by
 - [45] Wang, Z., et al. (2019). Neuromorphic computing: a review. *Journal of Physics D: Applied Physics*, 52(49), 493001. https://engineering.purdue.edu/NanoX/assets/pdf/2019_JPhysD_review_neuromorphic.pdf
 - [46] van de Ven, G. M., & Tolias, A. S. (2023). Generative Replay for Continual Learning. *arXiv preprint arXiv:2305.01818*. https://arxiv.org/pdf/2305.01818.pdf
 - [47] DeepAI. (n.d.). Self-Organizing Maps. *DeepAI*. https://deepai.org/machine-learning-glossary-and-terms/self-organizing-maps
 - [48] Wikipedia contributors. (2025). Artificial general intelligence. In *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/wiki/Artificial_general_intelligence
 - [49] Johnson, B., et al. (2025). Causal Reasoning in Brain-Inspired AI. *arXiv preprint arXiv:2502.11141*. https://arxiv.org/html/2502.11141v1
 - [50] Maass, W., et al. (2001). On the computational power of spiking neural networks. *Neural Computation*, 14(11), 2531-2560. https://pubmed.ncbi.nlm.nih.gov/11223461/
 - [51] Lake, B. M., et al. (2011). One shot learning of simple visual concepts. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. https://escholarship.org/content/qt9p12m226/qt9p12m226_noSplash_177cc756ea48fac53bed8606eff9e0e8.pdf
 - [52] Hassabis, D., et al. (2017). Neuroscience-Inspired Artificial Intelligence. *Neuron*, 95(2), 245-258. https://jair.org/index.php/jair/article/download/12664/26723/28384
 - [53] OpenAI. (2023). Planning for AGI and beyond. *OpenAI*. https://openai.com/index/planning-for-agi-and-beyond/
 - [54] Lesort, T., et al. (2024). A Survey on Continual Learning in the Era of Large Models. *arXiv preprint arXiv:2402.08381*. https://pmc.ncbi.nlm.nih.gov/articles/PMC12257422/
 - [55] Lopez-Paz, D., & Ranzato, M. (2017). Gradient Episodic Memory for Continual Learning. *Advances in Neural Information Processing Systems*, 30. https://openreview.net/pdf/82fc8f8111377eb145ccaa2c8b3a8a3bd988ea45.pdf
 - [56] Kirkpatrick, J., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521-3526. https://arxiv.org/pdf/1703.01357.pdf
 - [57] Purves, D., et al. (n.d.). Chapter 13 Summary: Learning and Memory. *Cognitive Neuroscience*, 2e. https://learninglink.oup.com/access/content/purves-cogneuro-2e-student-resources/purves-cogneuro-2e-chapter-13-summary
 - [58] Lake, B. M., et al. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332-1338. https://escholarship.org/uc/item/6dr8f6kj
 - [59] Rebuffi, S. A., et al. (2017). iCaRL: Incremental Classifier and Representation Learning. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. https://openreview.net/pdf/f5a1e2ccee3fdaeab4252a5ee651a0717e80a679.pdf
 - [60] UpGrad. (n.d.). Biological Neural Network - Explained. *UpGrad Blog*. https://www.upgrad.com/blog/biological-neural-network/
 - [61] GeeksforGeeks. (n.d.). Difference between ANN and BNN. *GeeksforGeeks*. [https://www.geeksforgeeks.org/

- difference-between-ann-and-bnn/](https://www.geeksforgeeks.org/difference-between-ann-and-bnn/)
- [62] World of Work Project. (2019). Dual Process Theory: Improve Your Decision Making. *World of Work*. https://worldofwork.io/2019/07/dual-process-theory/
- [63] Guerguiev, J., et al. (2021). Towards deep learning with segregated dendrites. *Elife*, 10, e70984. https://pubmed.ncbi.nlm.nih.gov/34391175/
- [64] Kiper, D. C., et al. (2018). Segregation of functionally distinct processing streams in the primate visual cortex. *Journal of Neurophysiology*, 119(5), 1845-1857. https://eprints.whiterose.ac.uk/106082/1/Kiper_et_al_JNS_accepted.pdf
- [65] Roy, K., et al. (2024). Towards Artificial General Intelligence via Concept-centric Multimodal Foundation Models. *arXiv preprint arXiv:2403.09496*. https://pubmed.ncbi.nlm.nih.gov/38577666/
- [66] Wang, Z., et al. (2023). A Comprehensive Survey on Continual Learning. *arXiv preprint arXiv:2302.00433*. https://openreview.net/pdf?id=NAuVe6pQ7Jb
- [67] Wang, Z., et al. (2023). A Comprehensive Survey on Continual Learning. *ICLR 2023 Conference*. https://openreview.net/forum?id=NAuVe6pQ7Jb
- [68] Graves, A., et al. (2007). Unsupervised learning of features for speech recognition. *International Conference on Artificial Neural Networks*. https://www.cs.toronto.edu/~graves/icann_2007.pdf
- [69] iMerit. (n.d.). A Complete Introduction To Continual Learning. *iMerit*. https://imerit.net/resources/blog/a-complete-introduction-to-continual-learning/
- [70] Vaibhav, A. (2023). Bio-Inspired Neural Architecture Search (NAS): Automating Deep Learning Design. *Dev.to*. https://dev.to/vaib/bio-inspired-neural-architecture-search-nas-automating-deep-learning-design-2d35
- [71] Lillicrap, T. P., et al. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6), 335-346. https://pubmed.ncbi.nlm.nih.gov/articles/PMC10629395/
- [72] Al-Shedivat, M., et al. (2024). The Alignment Problem in Artificial Intelligence: A Systematic Review. *Nature Communications Biology*, 7, 282. https://www.nature.com/articles/s42003-024-07282-3
- [73] Schmidhuber, J. (2021). The 2010s: Our Decade of Deep Learning / Outlook on the 2020s. *arXiv preprint arXiv:2102.04926*. https://pubmed.ncbi.nlm.nih.gov/articles/PMC8044545/
- [74] Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345-1359. https://sciendo.com/article/10.5372/1905-7415.0703.186
- [75] Donahue, J., et al. (2014). DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. *Proceedings of the 31st International Conference on Machine Learning*. https://pubmed.ncbi.nlm.nih.gov/articles/PMC3181834/
- [76] Soltoggio, A., et al. (2022). A review of neuromodulation in the context of continual learning. *Frontiers in Systems Neuroscience*, 16, 760864. https://pubmed.ncbi.nlm.nih.gov/articles/PMC8883483/
- [77] Soltoggio, A., et al. (2022). A review of neuromodulation in the context of continual learning. *Frontiers in Systems Neuroscience*, 16, 760864. https://www.frontiersin.org/journals/systems-neuroscience/articles/10.3389/fnsys.2022.760864/full
- [78] Flesch, T., et al. (2022). Object-based representations in the human brain. *European Journal of Neuroscience*, 56(6), 5133-5151. https://onlinelibrary.wiley.com/doi/abs/10.1111/ejn.16115
- [79] Koechlin, E., & Jubault, T. (2006). Broca's area and the hierarchical organization of human behavior. *Neuron*, 50(6), 963-974. https://www.princeton.edu/~yael/NIPSWorkshop/KoechlinAbstract.pdf
- [80] Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to task and cognitive control. *Trends in Cognitive Sciences*, 11(6), 229-235. https://pubmed.ncbi.nlm.nih.gov/articles/PMC3678031/