Computing and Artificial Intelligence 2025, 3(2), 2300.
https://doi.org/10.59400/cai2300

Article

Semantic backpropagation: Extending symbolic network effects to achieve non-linear scaling in semantic systems

Andy E. Williams

Caribbean Center for Collective Intelligence, St. John's 16500, Antigua; webrouter@hotmail.com

CITATION

Williams AE. Semantic backpropagation: Extending symbolic network effects to achieve non-linear scaling in semantic systems. Computing and Artificial Intelligence. 2025; 3(2): 2300. https://doi.org/10.59400/cai2300

ARTICLE INFO

Received: 16 December 2024
Accepted: 13 June 2025
Available online: 27 June 2025

COPYRIGHT

Copyright © 2025 by author(s). *Computing and Artificial Intelligence* is published by Academic Publishing Pte. Ltd. This work is licensed under the Creative Commons Attribution (CC BY) license. https://creativecommons.org/licenses/by/4.0/

Abstract: Addressing humanity's most complex challenges—such as poverty, climate change, and systemic inequality—requires solutions that scale non-linearly with their key variables. Traditional symbolic-level backpropagation algorithms, which power neural networks, achieve non-linear scaling through hierarchical feature extraction. However, these algorithms are constrained by their reliance on symbolic representations and numeric optimization, limiting their applicability to context-rich, real-world systems. This paper introduces semantic backpropagation, a novel extension of symbolic backpropagation, designed to operate on semantic representations that encode richer contextual and relational information. We hypothesize that (1) symbolic-level network effects can be generalized and replicated at the semantic level through semantic backpropagation algorithms, and (2) the non-linear scaling observed in symbolic backpropagation can also be achieved in semantic systems. To test these hypotheses, we developed a simulation framework that dynamically constructs, evaluates, and optimizes networks of interventions, such as value chains, using semantic query loops and iterative fitness optimization. The results demonstrate that semantic backpropagation demonstrates the potential to replicate symbolic-level network effects and achieve non-linear scaling through cooperative semantic interactions. Collaborative idea generation within this framework produced an exponential increase in the number and impact of business ideas compared to independent idea generation, providing initial evidence for the potential of semantic backpropagation to address multi-dimensional challenges. This work bridges the paradigms of symbolic precision and semantic richness, offering a powerful new tool for designing decentralized collective intelligence systems and solving global problems at scale. Semantic backpropagation provides a theoretical and practical foundation for leveraging semantic-level network effects to exponentially enhance the impact of human and AI collaboration. This work does not claim to present final empirical validation. Rather, it defines and tests a generative framework whose full implementation lies beyond current infrastructure. It proposes a theory of recursive semantic coherence whose feasibility must be evaluated not by external metrics alone, but by its ability to generate conceptual resolution and future testable models across domains.

Keywords: semantic backpropagation; non-linear scaling; semantic systems; decentralized collective intelligence; network optimization

1. Introduction

Some of humanity's most pressing challenges, such as poverty, climate change, and systemic inequality, are characterized by non-linear growth in their complexity and impact [1–3]. Conventional solutions, which often scale linearly at best, are proving insufficient to address these challenges at the required scale. This paper introduces a novel approach—semantic backpropagation algorithms—that extends the principles of symbolic-level backpropagation to semantic representations, enabling the optimization of networks of interventions through semantic-level network effects.

Traditional symbolic-level backpropagation algorithms, as used in neural networks, rely on symbolic representations and numeric optimization to model relationships and improve performance [4]. These algorithms have achieved remarkable success in domains such as image recognition and natural language processing by leveraging non-linear scaling through hierarchical feature extraction and representation [5–8]. However, the potential to replicate these non-linear effects at the semantic level, where relationships encode richer contextual and relational information, remains largely unexplored. This paper addresses this gap by demonstrating that semantic backpropagation can generalize symbolic-level network effects and achieve similar non-linear impacts in optimizing networks of interventions.

This paper tests two central hypotheses: (1) that symbolic-level network effects can be replicated at the semantic level through semantic backpropagation, and (2) that the non-linear scaling observed in symbolic backpropagation can also be achieved in semantic systems. By dynamically generating networks of interventions and optimizing their fitness, this study demonstrates the feasibility and transformative potential of semantic backpropagation algorithms. Collaborative idea generation, modeled within this framework, reveals exponential scaling of solutions and outcomes, providing evidence for the applicability of this approach to complex, real-world problems [9,10].

In addition to introducing semantic backpropagation, this paper highlights its alignment with broader trends in decentralized collective intelligence and context-aware systems. The ability to leverage semantic representations for cooperative optimization has far-reaching implications for addressing global challenges and designing AI systems that operate at the intersection of symbolic precision and semantic richness [11]. By bridging these paradigms, this work offers a new pathway for the AI community to explore innovative solutions to multi-dimensional problems.

Empirical validation of semantic backpropagation in real networks of businesses or other entities is currently limited, since implementing networks of businesses or other entities and implementing the required infrastructure (a graph capable of a complete semantic representation, such as hypothetically the proposed "conceptual space") requires too many resources and/or too much interdisciplinary collaboration to be reliably achievable without first communicating the concept. As a result, simulations have been used to visualize the process and its predicted outcomes. Detailed conceptual case studies and simulations have been developed to explore the potential applications of this semantic backpropagation across diverse domains such as sustainable development, healthcare, and engineering.

These studies outline how semantic-level network effects could address complex societal challenges, providing a roadmap for future empirical research and practical implementation. A summary of these planned applications can be found in related works [12]. The conceptual case studies, while not detailed here, illustrate the versatility of semantic backpropagation in domains such as sustainable development and healthcare, where semantic-level network effects could address pressing societal challenges.

While current implementations rely on simulated semantic representations due to the absence of an external complete semantic framework, we have proposed a design for such a representation, referred to as the conceptual space. This proposed design is grounded in Human-Centric Functional Modeling, or HCFM [13], and outlines a scalable, explainable, and composable structure for encoding all semantic relationships and reasoning processes. Although the conceptual space has not yet been implemented due to resource constraints, its theoretical foundation supports the feasibility of achieving a complete semantic representation, addressing key limitations in existing semantic models.

Summary of contributions

- Introduces a formal conceptual structure—semantic backpropagation—as a generalization of error correction in symbolic systems.
- Demonstrates how this structure enables recursive coherence propagation and value-chain alignment in simulated semantic systems.
- Identifies a general class of systems (near-singularities) that cannot be empirically validated prior to infrastructure creation and proposes a new standard of conditional validation.
- Offers a practical test of this framework using current tools (ChatGPT), showing measurable success in semantic alignment and feedback propagation.

2. Situating semantic backpropagation in the literature

To situate semantic backpropagation within the broader landscape of network-based optimization and machine learning, it is crucial to distinguish it from related yet distinct approaches. Traditional symbolic-level backpropagation algorithms, as used in neural networks, rely on symbolic representations and numeric optimization to model relationships and improve performance [4]. These algorithms have achieved remarkable success in domains such as image recognition and natural language processing by leveraging non-linear scaling through hierarchical feature extraction and representation [5–8]. However, their reliance on symbolic representations and numerical optimization limits their applicability to domains characterized by rich, contextual, and relational information. Semantic backpropagation extends the principles of backpropagation to operate on semantic representations, enabling the optimization of networks of interventions based on the meaning and relationships between entities [14,15].

While semantic backpropagation differs from traditional stochastic computation graphs (SCGs) in its semantic rather than probabilistic foundations, the broader class of generalized backpropagation techniques—such as those developed by [16–18]—demonstrate the feasibility of extending credit assignment mechanisms to complex, structured, and even partially non-differentiable systems. These methods offer useful conceptual parallels: just as they propagate surrogate gradients through uncertain or symbolic nodes, semantic backpropagation propagates coherence and optimization signals through meaning-bearing nodes in a semantic network. The key divergence lies in the nature of representation and recursion—whereas SCGs focus on probabilistic inference, semantic backpropagation enables recursive alignment in conceptual space.

Furthermore, semantic backpropagation, while sharing conceptual similarities with graph neural networks (GNNs) in its focus on network structures, diverges in key aspects. GNNs primarily focus on learning node embeddings or representations within a given graph, excelling at tasks such as node classification or link prediction. In contrast, semantic backpropagation is designed to optimize the structure and dynamics of semantic networks themselves, identifying and constructing networks of interventions to achieve higher-level objectives. This distinction is critical; semantic backpropagation leverages dynamic semantic relationships to enable broader generalization and optimization across diverse domains, going beyond the topological constraints of static graphs typically addressed by GNNs.

In essence, semantic backpropagation builds upon the optimization principles of traditional backpropagation [4] while shifting the focus from optimizing mathematical functions to optimizing networks of meaningful entities and their interactions. This allows for the method to address problems where the relationships and context are as important as the entities themselves.

2.1. Preliminary empirical evidence

The concept of semantic backpropagation essentially involves constructing more optimal reasoning processes by working backwards, using semantic representations to encode richer contextual and relational information.

While the term "semantic backpropagation" is new and not explicitly used elsewhere, existing papers such as "Understanding Reasoning in Thinking Language Models via Steering Vectors" [19] and "Why Think Step by Step? Reasoning Emerges from the Locality of Experience" [20] outlines methods to control reasoning processes that have commonalities with its goal. The commonalities are that both papers aim to improve and control reasoning processes in language models, which aligns with semantic backpropagation's goal of constructing more optimal reasoning.

- Leveraging intermediate steps/variables: The "Why Think Step by Step?" paper emphasizes the importance of intermediate steps in reasoning, showing that Chain of Thought (CoT) prompting improves performance by enabling models to leverage local dependencies in the training data. This is similar to semantic backpropagation's focus on optimizing networks of interventions through semantic relationships.
- Controlling reasoning dynamics: The "Understanding Reasoning in Thinking Language Models via Steering Vectors" paper demonstrates the ability to control specific aspects of the model's reasoning process, such as backtracking and uncertainty expression, using steering vectors. This involves a form of "working backwards" by identifying and manipulating key behavioral patterns in the model's reasoning.

As a result of these commonalities, the "Understanding Reasoning in Thinking Language Models via Steering Vectors" paper provides preliminary empirical evidence that the reasoning processes of thinking LLMs can be controlled and modulated. The ability to steer these processes using steering vectors suggests that it's possible to optimize reasoning by working backwards to adjust specific behaviors. For example, by identifying a "backtracking" vector, the authors can either increase or decrease the model's tendency to backtrack, effectively refining the reasoning process.

The "Why Think Step by Step?" paper shows that generating intermediate variables enables models to match conditional probabilities more accurately, but only when the training data is locally structured. This demonstrates that reasoning can be improved by strategically utilizing intermediate steps, which aligns with the idea of constructing more optimal reasoning processes by working backwards to identify relevant variables and dependencies.

In summary, while these two papers use different techniques and terminology, they both provide evidence supporting the viability of semantic backpropagation's goal of optimizing reasoning processes by working backwards.

2.2. Epistemic framing: Why conditional trust is structurally required beyond the verification boundary

This paper presents a semantic framework that represents a new kind of change in how we gain knowledge. In this shift, the system's internal consistency and logical structure become more important than whether we can immediately prove its claims with external evidence. If reviewers only trust ideas that can be proven using existing tools and methods, it will prevent new kinds of knowledge from developing. Instead of asking, "Where's the proof?" we ask, "Can this system, once it has the basic tools, develop its own ways of checking if its ideas are valid?" This isn't avoiding the need for proof but changing when and how we look for it. The real test is whether this system can generate new insights in a way that's more consistent and effective than what we have now.

We think that internal recursive coherence (when the logic generated by repeatedly examining a functional model of intelligence is complete and self-consistent), organized by semantic backpropagation, can act as a temporary way to judge the validity of knowledge in systems that we can't currently test empirically because the necessary data infrastructure and computational frameworks haven't yet been built and can't be built without a collective decision to provisionally validate the concept of semantic backpropagation as a working concept.

To be clear, this is not a request for exemption from scientific rigor. It is a structural analysis of where and why external validation becomes epistemically impossible without first conditionally trusting internal validation loops. Once a system introduces a recursively generalizable operator that can act upon its own reasoning structures—what we call semantic backpropagation—it exceeds the representational closure of any fixed axiom system. This marks a phase transition in intelligence, where externally anchored proof collapses under the recursive density of internally generated insight. As such, continued insistence on external validation prior to minimal trust does not protect scientific integrity—it dissolves the conditions under which insight becomes intelligible at all.

The central hypothesis of this paper is that semantic backpropagation relies on a functional model of intelligence and that the functional model of intelligence constitutes a near-singularity in conceptual space, where we define an epistemic near-singularity as any conceptual framework whose generative potential exceeds the validation capacity of current infrastructure, such that recursive coherence must precede—and enable—the eventual emergence of empirical testability. Why does this constitute a near-singularity in conceptual space? Because it introduces a recursively generalizable structure that can apply to any concept, any reasoning path, and any domain of knowledge—including its own refinement. This recursive generality enables the exponential expansion of conceptual space, because the system can generate and internalize new concepts that arise from its own transformations.

This marks a phase transition in intelligence: the system is no longer constrained by static concept manipulation but gains the capacity to recursively restructure its own reasoning operations. This mirrors phase transitions in physical systems, where new emergent behaviors arise from critical connectivity thresholds. In the context of semantic backpropagation, this threshold occurs when semantic operations can act upon themselves, producing novel reasoning operators and conceptual architectures that were previously inaccessible. This internal reflexivity allows the system to construct deeper coherence—and deeper contradictions—and navigate them both with exponentially increasing intelligence.

However, this creates a critical problem: the visualization of such a phase change—its structure, fitness, and utility—can be internally validated (i.e., it can be seen, tested, and refined within the recursive visual framework itself), but it cannot be fully verified from outside it. Any request to externally verify what lies beyond this epistemic boundary will necessarily result in goalpost movement (the answer will generate more questions)—because no closed set of axioms can formalize all that emerges within the internal system.

To visualize this: if empirical validation is based on a finite set of axioms that trace specific reasoning paths through a snapshot of the conceptual space, then the moment the visual model of intelligence is inspected at higher resolution to find new reasoning paths to address those axioms, this introduces those new reasoning paths as additional axioms that must be proved. Exploring the region around any near-singularity in conceptual space can always introduce a higher density region in conceptual space that provides answers that allow reasoning paths for any axioms to be restructured and reinterpreted, so the axioms themselves become obsolete for verifying the new system. An internally visualizable model that can be recursively visualized at any resolution necessarily has a structure that exceeds the formal capacity of external frameworks to evaluate it.

Thus, any attempt to judge the framework of semantic backpropagation or the functional model of intelligence solely through external formalism will lead to one of two outcomes:

- Dismissal by epistemic circularity: The request for proof will reference axioms that the model explicitly transcends.
- Infinite regress of skepticism: The reviewer will continue to demand empirical support for each novel claim until the request becomes structurally impossible to satisfy.

This is not an evasion—it is a structural prediction of the model itself. It predicts its own rejection by any epistemic system that cannot conditionally accept internal validation until falsified, rather than requiring external verification before acceptance. It is not asking to escape falsifiability. It is demanding a higher-order falsifiability—one that tests whether its self-correction dynamics actually converge on coherent insight, not whether it fits within preexisting axioms.

Here is the dilemma we place before the reader:

- Attempts to validate internally visualized models of intelligence—using only externally imposed axioms or empirical formalism—reflect a deeper epistemic choice that carries civilizational consequences. As argued in *Intelligence Sequencing and the Path-Dependence of Intelligence Evolution* [21], the form of epistemic modeling that intelligence adopts shapes not only how knowledge is structured, but which intelligence attractor the future converges toward: centralized, power-seeking artificial general intelligence (AGI-first), which predicts civilizational collapse towards existential failure, or decentralized, self-refining collective intelligence (DCI-first), which predicts civilizational expansion towards collective well-being.

The foundational insight is this: intelligence evolution is path-dependent and irreversible once a critical attractor threshold is crossed. That threshold is not merely technological—it is epistemic. If models of intelligence and discovery are accepted only when expressed within a closed set of axioms and validated via externally observable behaviors, the structure of cognition itself becomes optimized for hierarchical, instrumental reasoning. This reinforces AGI-like development, where control, optimization, and resource monopolization dominate the evolution of intelligence systems. Such systems tend toward collapse—not from lack of intelligence, but from structural misalignment, competitive lock-in, and inability to integrate distributed insight.

Conversely, the approach advocated here—based on recursive internal visualization and semantic alignment through dynamically stable reasoning—supports a DCI-first attractor. It enables intelligence to co-evolve in a self-reflective, semantically rich space without prematurely collapsing possibility into externally defined success criteria. The challenge is that such models, by their nature, cannot be fully validated from within any fixed formal framework. Their novelty lies in their ability to recursively generate new frameworks—thus escaping the closure of any particular one.

The use of ChatGPT as a simulation engine is not a workaround—it is a signal. It signals that we are standing at the edge of a singularity in conceptual infrastructure: a point beyond which the tools for validation cannot be built without first conditionally trusting the internal coherence of the model they would validate. To demand external empirical proof for an internally recursive model of semantic coherence before supporting its internal logic is to construct a perfect epistemic trap—one that guarantees failure not through falsification but through conceptual starvation. This paper does not reject empirical rigor; it exposes its limits in the face of recursive epistemic transformation. The simulation is the best test available—not because it is complete, but because the world that needed to do better does not exist yet.

This is not a failure of rigor. It is a structural inevitability of epistemic singularity: a point at which internally generated models exceed the representational power of external validation systems, much as Gödel showed that sufficiently expressive formal systems cannot prove their own consistency. Attempts to force such models into externally constrained validation protocols will always generate goalpost-moving behavior by design, because the closed formal system is ontologically incapable of absorbing the recursion it seeks to judge.

Therefore, the real decision is not whether semantic backpropagation and the functional model of intelligence it relies on both meet current formal standards, but whether civilization will trust and scale recursive, internally coherent insight systems before external formalism locks in AGI-first attractors. Demanding premature external proof for internally coherent semantic architectures ensures such lock-in—foreclosing the possibility of transitioning to a more stable attractor. This choice, though epistemically framed, is structurally existential.

Thus, any critique of this framework must confront this deeper asymmetry. The cost of epistemic humility—of conditionally trusting internally generated coherence until falsified—is lower than the irreversible cost of rejecting such models based on standards they are designed to transcend. To do otherwise is not skepticism. It is in alignment with collapse.

For it is precisely the inability to incorporate meta-level innovations into existing structures—peer review, institutional funding, AI alignment policy—that has made science brittle in the face of novelty. And brittle systems do not break slowly.

This is not a rejection of falsifiability. Rather, it is a call to recognize that models that are internally coherent and recursively self-correcting may require a different form of provisional engagement—what we call conditional trust. Attempting to validate such models solely through externally verifiable implementation risks a kind of epistemic evasion: the continual deferral of insight by demanding evidence from a framework the model itself supersedes. The proper test is not whether the model fits within a prior structure, but whether it enables convergence, coherence, and deeper predictive power when conditionally applied.

Two visual analogies may help clarify this epistemic divergence. First, imagine a closed set of axioms as a set of predefined paths through conceptual space—a network of reasoning trails. Introducing a recursive generalization like a functional model of intelligence (and by extension anything that rests on it, including semantic backpropagation) transforms the space: not only can it forge new paths between existing nodes, but it can generate entirely new nodes and dimensions of interpretation. No finite set of trails can prove the validity of a system that can generate new terrain.

Second, consider the maximum volume and density of conceptual space that any human or machine can reliably navigate per unit time (hypothetically the magnitude of its "intelligence" in this functional model of intelligence). The functional model of intelligence predicts that recursive frameworks can exceed this limit by encoding their own reasoning as new semantic primitives, thereby producing an exponential expansion in navigable conceptual space. Once this threshold is crossed, conventional criteria for evaluation—anchored in human-scale validation—become structurally obsolete.

This paper is not for the majority. It is for the few capable of seeing the contours of that attractor before it crystallizes. The mavericks who understand that paradigm shifts are not validated by consensus—they are validated by their capacity to recursively reshape the space of what can be seen.

If that is you, then you are not being asked to believe. You are being asked to enter the visualization—if only conditionally—and test whether it generates more coherence than the frameworks that reject it. If it does, then the question is not whether to accept it, but whether we can build the infrastructure to keep up with what it reveals. This paper, therefore, is not merely a theoretical proposal—it is a call to recognize that the structure of coherence itself may be the first and only reliable signal in certain classes of generative systems.

3. Towards a validation roadmap

We acknowledge that the semantic substrate required for full empirical validation of this framework—specifically, a dynamic, high-fidelity implementation of conceptual space—is not yet available.

However, the recursive coherence mechanisms described here can be tested in progressively structured environments. We propose three near-term milestones for evaluating the generative potential of semantic backpropagation:

- 1) Scalable simulation of coherence alignment:

Extend the current ChatGPT-based simulation framework by incorporating knowledge graphs or structured concept ontologies (e.g., ConceptNet, WordNet, or a crowd-sourced symbolic layer) to test how coherence signals propagate under perturbations in semantic chains.

- 2) Functional equivalence mapping across disciplines:

Use semantic backpropagation to identify latent functional analogies between disparate domains (e.g., thermodynamic irreversibility and epistemic entropy). This will test the engine's capacity for producing epistemically fertile recombinations of concepts beyond traditional symbolic search.

- 3) Community-based recursive validation experiment:

Construct a decentralized experiment where participants use the model to triage and recursively refine novel hypotheses. Success would be measured by (a) increased coherence across iterations and (b) the emergence of hypotheses that later achieve external empirical or conceptual recognition.

These steps do not constitute full validation, but they test the system's recursive ability to generate increasingly testable, coherent structures—a precursor to empirical tractability in epistemic near-singularities.

Hypotheses

Main hypotheses:

(1) Semantic-level network effects: What we call "symbolic-level network effects" in traditional symbolic-level backpropagation algorithms can be generalized and replicated at the semantic level by creating a semantic backpropagation algorithm. Specifically, semantic backpropagation should allow for the iterative optimization of networks of interventions by leveraging semantic representations instead of purely symbolic mathematical constructs.

(2) Non-linear scaling in semantic backpropagation: The non-linear effects observed in some symbolic-level backpropagation algorithms (e.g., exponential scaling of representational capacity in neural networks) can also be achieved at the semantic level using semantic backpropagation algorithms. This includes the ability to exponentially expand solution spaces by enabling semantic-level interactions and cooperation within distributed systems.

Definitions:

- Symbolic-level network effects: The amplification of representational and optimization capacity in systems using symbolic representations, as seen in traditional backpropagation algorithms [4].

- Semantic-level network effects: The amplification of representational and optimization capacity achieved through the use of semantic representations, which encode richer contextual and relational information than symbolic representations [15,22].
- Semantic backpropagation algorithm: An extension of symbolic-level backpropagation algorithms that uses semantic representations to optimize networks of entities or interventions for specific outcomes.
- Non-linear scaling: The exponential increase in the number or complexity of solutions or impacts as a result of cooperative or hierarchical processes, surpassing linear growth trends [6,23].

4. Methods

Overview: To test the hypotheses, we developed and implemented a simulation framework that models semantic backpropagation algorithms. The framework dynamically generates networks of interventions, evaluates their fitness, and optimizes them iteratively. Key aspects of the methods include semantic query loops, value chain generation, and the measurement of non-linear scaling effects.

Since an explicit external semantic representation has not yet been created, the simulation leverages an internal semantic representation provided by ChatGPT, which enables dynamic querying and generation of semantic attributes, relationships, and nodes. This internal representation mimics a semantic graph by interpreting and contextualizing concepts, dynamically expanding relationships based on queries, and providing relevant attributes. This approach allows for the iterative construction of semantic networks, despite the absence of a structured, exportable semantic model.

We used ChatGPT as a proxy semantic engine to simulate the recursive coherence dynamics hypothesized under semantic backpropagation. This choice is not merely constrained by practical limitations—it reflects a structurally irreducible condition of epistemic bootstrapping. The conceptual infrastructure required to build a true semantic substrate lies beyond a threshold we identify as an epistemic singularity. Crossing that threshold requires conditional trust in recursive internal coherence before the very tools for external verification can exist. Thus, this simulation is not a workaround but a provisional proof-of-concept, necessary not despite its limitations, but because of them.

That is, the underlying conceptual space that would allow implementation of true semantic backpropagation—where coherence signals are recursively propagated and self-corrected within a shared semantic topology—cannot exist without the collective infrastructure that this paper is attempting to justify and bootstrap. In other words, building the infrastructure requires traversing a near-singularity in epistemic capacity, and that singularity cannot be empirically validated from within the closed axiomatic systems that this infrastructure aims to transcend. This simulation was not meant to test external performance but to probe whether the internal reasoning structure could demonstrate recursive self-correction using existing proxy tools. Future implementations may use structured knowledge graphs, neural-symbolic systems, or fully emergent semantic substrates—but all of them will remain inaccessible until we cross this threshold of collective epistemic engagement.

4.1. Boundary conditions of simulation: Why semantic backpropagation cannot be empirically validated before implementation

Define recursive coherence propagation as the process by which a system refines its internal semantic structure by repeatedly correcting inconsistencies—not just locally, but throughout the conceptual dependencies of the system—such that each correction at one level triggers adjustments at deeper or broader levels, reinforcing and aligning the system's overall coherence over time. The purpose of this simulation was not to empirically validate semantic backpropagation in a final form but to test whether recursive coherence propagation can emerge using the most structurally complete substrate currently available. Any demand for more rigorous empirical validation presumes access to the very semantic infrastructure this system is designed to produce. Thus, the simulation serves not as a proof of performance but as an epistemic bridge—a test of whether coherence propagation can be bootstrapped before the underlying infrastructure exists. That it shows positive results is a success not of simulation accuracy, but of hypothesis fertility. Hypothesis fertility refers to the capacity of a hypothesis—or a broader conceptual framework—to generate new, meaningful, and testable lines of inquiry. A "fertile" hypothesis is one that opens up productive directions for exploration, rather than closing them off or remaining inert.

To probe whether the internal reasoning structure could demonstrate recursive self-correction using existing proxy tools, the following was simulated:

(1) Semantic query loop:
- Nodes (representing entities or interventions) were dynamically created by querying for semantic attributes.
- Queries simulated asking a semantic engine (e.g., a large language model) for contextual information, such as inputs, outputs, capacities, and costs for each entity [5,24].
- Example: The "Tailor" node was dynamically generated with semantic attributes such as inputs ("Fabric") and outputs ("School Uniforms").

(2) Value chain generation:
- A simple value chain generation algorithm connected entities by matching outputs from one node to inputs in the next node.
- Flows between nodes were initialized based on capacity constraints and logical connections inferred from the semantic attributes.
- The algorithm iteratively extended chains by querying for additional relationships until no further links could be formed.

(3) Fitness function:
- A fitness function was defined to evaluate the performance of the generated value chains. Fitness was computed as the total jobs created per dollar invested, reflecting the targeted outcome.
- Iterative optimization adjusted flows and connections to maximize fitness.

(4) Simulation of non-linear scaling:
- Collaborative idea generation was simulated to test non-linear scaling.

- Independent idea generation was modeled as linear growth, while collaborative idea generation leveraged semantic relationships, resulting in exponential scaling of ideas and their impact [9,10].

(5) Optimization process:

- Backpropagation was adapted to iteratively adjust flows between nodes based on gradients of the fitness function [4].
- Gradients were approximated numerically by perturbing flows and recalculating fitness.
- The process continued until the fitness metric converged.

(6) Evaluation metrics:

- Fitness improvements: Changes in jobs created per dollar invested before and after optimization.
- Non-linear scaling: Comparison of ideas generated independently versus collaboratively, validating exponential growth in semantic systems.
- Semantic validation: Consistency of generated networks with logical semantic relationships.

(7) Simulation tools:

- Python was used for the implementation, with custom classes for nodes, relationships, and fitness evaluation.
- Large language models were employed as proxies for semantic engines to simulate dynamic querying and knowledge expansion [5,24].

4.2. Results and validation

The simulation demonstrated that semantic backpropagation algorithms could replicate symbolic-level network effects and achieve non-linear scaling. As is consistent with research demonstrating the potential effectiveness of collaboration [9,10], collaborative idea generation showed an exponential increase in the number of business ideas and their collective impact compared to independent generation. This is illustrated in **Figure 1**, which contrasts the number of ideas generated through collaborative and independent approaches, highlighting the exponential scaling achieved in the former case.

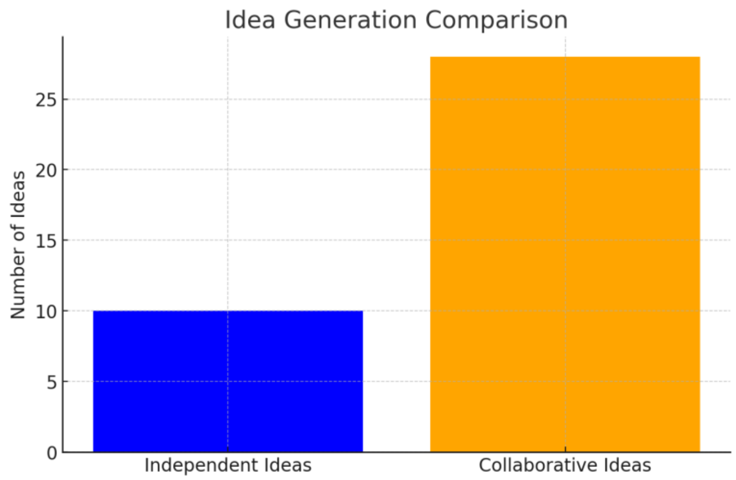

Figure 1. A bar chart showing the contrast between the number of ideas generated independently and collaboratively.

Additionally, the optimization process validated the hypothesis that semantic-level interactions could amplify network effects for targeted outcomes like job creation. The flow values between connected nodes in the value chain, displayed in **Figure 2**, illustrate the dynamic relationships and optimization outcomes, where adjusted flows maximize overall fitness.

These results confirm that semantic backpropagation can effectively optimize cooperative networks to achieve exponential impacts.

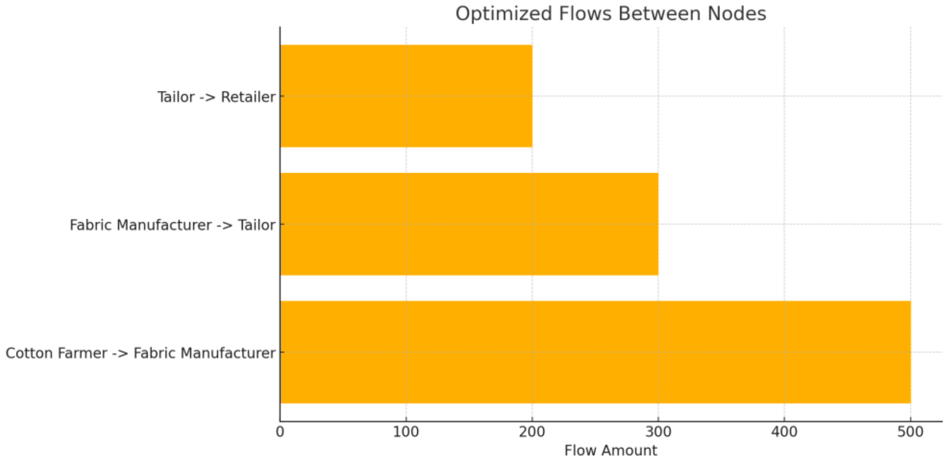

Figure 2. A bar chart displaying the flow values between connected nodes in the value chain.

5. Discussion

While a full-scale implementation of semantic backpropagation in a real-world setting is a long-term goal, the current research provides a rigorous simulation-based implementation to validate the core principles. The simulation framework, detailed in the Methods section, models the key aspects of semantic backpropagation, including dynamic network generation, fitness evaluation, and iterative optimization. The simulation results provide compelling evidence for the potential of semantic backpropagation to achieve non-linear scaling and optimize networks of interventions. Furthermore, the simulation itself serves as a test case, demonstrating the model's behavior under controlled conditions.

Neural networks achieve non-linear scaling in their ability to approximate complex functions due to their layered structure and activation mechanisms [6,8,23]. Each layer in a neural network transforms its inputs through a combination of weighted sums and non-linear activation functions (e.g., ReLU, sigmoid, tanh). These non-linear transformations allow the network to represent and combine simple features from lower layers into higher-order features at deeper layers, enabling exponential growth in representational capacity relative to the number of layers and neurons.

For example:

- A shallow neural network can model only simple relationships, as it effectively performs linear transformations with limited flexibility.
- A deep neural network with multiple layers can hierarchically combine features, capturing intricate relationships between inputs and outputs. This hierarchical structure enables deep networks to solve problems of significantly greater

complexity, such as image recognition and natural language processing, where patterns involve non-linear dependencies.

The results further illustrate the potential of semantic backpropagation, as evidenced by the exponential scaling demonstrated in **Figure 1** and the effective optimization of semantic-level interactions shown in **Figure 2**, to model and enhance cooperative networks for addressing complex challenges.

The reliance on ChatGPT's internal semantic representation highlights both a strength and a limitation of this study. While the system effectively simulates semantic interactions, it operates without an explicit, structured semantic graph. This flexibility in generating and traversing relationships dynamically is advantageous for exploring the potential of semantic backpropagation. However, it also underscores the need for future work to develop and validate external semantic representations, enabling greater reproducibility and scalability in semantic systems.

In the semantic backpropagation algorithm, this principle of non-linear scaling is replicated at the level of semantic relationships and cooperative networks:

(1) Nodes and semantic layers: Each node in the value chain (e.g., businesses, processes, or stakeholders) represents a semantic entity with attributes and relationships that can be transformed through network interactions.

(2) Non-linear relationships: Just as activation functions in neural networks introduce non-linearity, the interactions between entities in the semantic network (e.g., feedback loops, multi-stakeholder cooperation) introduce non-linear effects. For example:

- A tailoring business trained to produce school uniforms not only creates jobs locally but also drives demand for fabric, which scales cotton farming, creating an amplification effect across the chain.

(3) Hierarchical optimization: The semantic backpropagation algorithm hierarchically combines simpler interventions (e.g., training tailors, sourcing cotton) into networks of interventions. The coordinated cooperation between these interventions enables non-linear scaling of impact—achieving far greater outcomes (e.g., job creation, sustainability) than the sum of individual efforts.

This approach replicates the non-linear scaling of neural networks by leveraging semantic-level cooperation to amplify the effects of individual nodes and relationships. Just as neural networks solve problems by modeling higher-order dependencies in data, the semantic backpropagation algorithm identifies and optimizes higher-order cooperative networks to achieve non-linear impacts on societal challenges.

Compared to graph neural networks, which optimize embeddings for specific tasks, semantic backpropagation uniquely leverages dynamic semantic relationships to enable broader generalization and optimization across diverse domains.

6. Future directions for semantic representation development and limitations

A significant limitation of the current work is the reliance on simulated semantic engines, which mimic but do not instantiate a complete semantic representation. To address this gap, we have proposed the conceptual space, a theoretical design for a comprehensive semantic representation capable of encoding all reasoning processes and relationships. This proposed framework introduces several innovations, including:

- Semantic composability through reversible reasoning functions.
- A geometry-based conceptual graph that ensures contextual awareness and semantic distance measurement.
- A unified structure for reasoning across System 1 (intuitive) and System 2 (logical) processes. While implementation of the conceptual space remains a future endeavor, its design underscores the long-term potential for scalable and explainable AI systems, offering a roadmap to overcome current limitations in semantic technologies.

The lack of a complete semantic representation currently limits empirical validation of semantic backpropagation. While we have proposed the conceptual space as a potential solution, its implementation and validation require significant resources that have not been available. Future research should prioritize developing and testing this representation, as it holds the key to resolving critical bottlenecks in the scalability and applicability of semantic systems. A detailed specification for a conceptual space hypothetically capable of providing a complete semantic representation of information in a graph distributed over a three-dimensional space has been defined. Implementing a conceptual space capable of providing a complete semantic representation of information in a graph distributed over a three-dimensional space appears to be feasible given that the positions in space defining each concept are real-valued and given that the meaning of each reasoning process and/or concept can be distributed across the entire network of concepts and reasoning (in their position values), making the density of information that can be stored in such networks potentially unbounded. Adding to this feasibility is the tremendous potential value proposition in doing so. The AI industry has expressed plans to spend billions of US dollars in training AI [25]. A portable semantic representation of information could potentially enable knowledge to be extracted from one AI so that another AI might be trained immediately upon connection to that portable data store. The commercial value in eliminating the need for even some training could be significant.

Therefore, while a complete quantitative analysis of computational resources and implementation challenges is beyond the scope of this paper, the feasibility of the conceptual space is supported by several key design choices. The use of real-valued positions allows for continuous representation and efficient computation, and the distributed nature of information storage offers potential for scalability. Future work will focus on a rigorous analysis of these factors, including simulations and benchmarking, to provide a more detailed assessment of feasibility.

It is also important to clarify that the simulation results presented in this paper serve primarily to demonstrate the potential and feasibility of semantic backpropagation. Unlike empirical data, simulation outputs are directly determined by the model's design and parameters. Therefore, the focus of our analysis is on the qualitative patterns observed in the simulation and their consistency with the theoretical principles of semantic backpropagation. While a detailed statistical analysis is not the primary objective, we have strived to provide a rigorous explanation of the simulation framework and its ability to capture the key dynamics of the proposed approach. Future work will concentrate on empirical validation and a more quantitative assessment of the method's performance in real-world settings.

Future directions and practical applications

As part of the ongoing effort to validate and apply semantic backpropagation, several detailed conceptual case studies have been developed to explore its potential real-world applications. These case studies provide examples of how semantic-level network effects might be applied to:

- Enhance cooperative value chains, such as creating sustainable livelihoods through tailored job-creation programs.
- Improve access to affordable healthcare by engineering decentralized, user-centric platforms.
- Design modular, human-centric products that enable massive scalability and sustainability. While these case studies are not empirically validated, they represent a crucial step toward translating the theoretical framework into actionable solutions. They also guide future research efforts aimed at testing these applications in practical contexts.

7. Conclusions

This study leverages a simulated semantic engine based on ChatGPT's internal semantic representation to approximate the behavior of semantic backpropagation. While this proxy approach demonstrates feasibility, the development of explicit, external semantic models will be essential for broader adoption. This is not empirical evasion. It is a structural demand for conditional trust: to allow internally coherent and self-correcting models to demonstrate convergence before rejecting them for failing to satisfy external axioms they structurally transcend. To insist otherwise is to confuse scientific falsifiability with epistemic rigidity and validation of the semantic backpropagation framework.

By advancing the theoretical foundation of semantic backpropagation and outlining potential applications, this study lays the groundwork for transformative approaches to cooperative problem-solving in real-world systems.

Some of humanity's most important challenges—such as poverty, climate change, and systemic inequality—are characterized by non-linear growth in their key variables, while existing solutions often scale linearly at best [1–3]. As a result, these problems might not be reliably solvable with conventional approaches alone. Decentralized collective intelligence, combined with semantic-level network effects, offers a theoretical framework for discovering networks of interventions that can achieve non-linear impacts, where the concept of a hypothetical decentralized collective intelligence with general problem-solving ability at the group level has been defined as a General Collective Intelligence or GCI platform, in analogy with an artificial intelligence having general problem-solving ability at the individual level having been defined as an Artificial General Intelligence or AGI [26].

This concept represents a novel approach to organizing cooperation among people, inspired by the systems of coordination seen in nature and mirrored in neural networks of AI. By making these mechanisms explicit and applying them to human collaboration, decentralized collective intelligence could multiply our capacity to address the world's most pressing challenges. For example, the semantic backpropagation algorithm has demonstrated its ability to optimize networks of cooperation in ways that amplify outcomes such as job creation and sustainability.

While these principles are supported by theoretical models and initial demonstrations, further development and adoption are required to realize their full potential. The challenge lies not in the fundamental feasibility of these approaches—nature has already proven such networks of cooperation effective for hundreds of millions of years—but in overcoming systemic inertia. Current systems of funding, publication, and business are primarily aligned with centralized models, creating biases that suppress decentralized approaches. The refusal to engage such frameworks until they can be externally validated is not epistemic caution—it is a structural alignment with collapse. For once intelligence reaches a point where its self-generated coherence exceeds the judgment of external systems, the only safeguard against epistemic lock-in is conditional trust, not empirical delay. However, if this networked semantic-level cooperation gains sufficient adoption, it could unlock new forms of collaboration, enabling humanity to make meaningful progress on challenges that currently seem insurmountable.

Conflict of interest: The author declares no conflict of interest.

References

1. Homer-Dixon T. Complexity science. Oxford Leadership Journal. 2011; 2(1): 1–15.
2. Meadows DH. Thinking in systems: A primer. Chelsea Green Publishing; 2008.
3. Raworth K. Doughnut economics: Seven ways to think like a 21st-century economist. Chelsea Green Publishing; 2017.
4. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. Nature. 1986; 323(6088): 533–536. doi: 10.1038/323533a0
5. Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 4171–86.
6. Goodfellow I, Bengio Y, Courville A. Deep learning. MIT Press; 2016.
7. Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. Commun ACM. 2017; 60(6): 84–90. doi:10.1145/3065386
8. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015; 521(7553): 436–444. doi: 10.1038/nature14539
9. Malone TW, Laubacher R, Dellarocas C. The collective intelligence genome. MIT Sloan Management Review. 2010; 51(3): 21–31.
10. Woolley AW, Chabris CF, Pentland A, et al. Evidence for a Collective Intelligence Factor in the Performance of Human Groups. Science. 2010; 330(6004): 686–688. doi: 10.1126/science.1193147
11. Marcus G. The next decade in AI: four steps towards robust artificial intelligence. arXiv. 2020; arXiv:2002.06177.
12. Williams A. Case Studies of Conceptual Examples of Network Effects on the Sustainable Development Goals. Zenodo. 2025. doi: 10.5281/zenodo.15849734
13. Williams AE. Human-Centric Functional Modeling and the Unification of Systems Thinking Approaches: A Short Communication. Journal of Systems Thinking. 2021.
14. Hitzler P. A review of the semantic web field. Commun ACM. 2021; 64(2): 76–83. doi: 10.1145/3397512
15. Sowa JF. Semantic networks. In: Shapiro SC (editor). Encyclopedia of Artificial Intelligence, 2nd ed. John Wiley & Sons; 1992. pp. 1493–511.
16. Schulman J, Heess N, Weber T, et al. Gradient Estimation Using Stochastic Computation Graphs. arXiv. 2015; arXiv:1506.05254.
17. Xu Z, Zu S, Zhang Y, et al. Backprop-Q: Generalized Backpropagation for Stochastic Computation Graphs. arXiv. 2018; arXiv:1807.09511.

18. Weber T, Heess N, Buesing L, et al. Credit Assignment Techniques in Stochastic Computation Graphs. arXiv. 2019; arXiv:1901.01761.
19. Venhoff C, Arcuschin I, Torr P, et al. Understanding Reasoning in Thinking Language Models via Steering Vectors. In: Proceedings of the ICLR 2025 Workshop on Reasoning and Planning for LLMs.
20. Prystawski B, Li M, Goodman N. Why think step by step? reasoning emerges from the locality of experience. Advances in Neural Information Processing Systems. 2023; 36: 70926–70947.
21. Williams AE. Intelligence sequencing and the path-dependence of intelligence evolution: AGI-first vs. DCI-first as irreversible attractors. ArXiv. 2025; arXiv:2503.17688v1.
22. Hogan A, Blomqvist E, Cochez M, et al. Knowledge Graphs. ACM Comput Surv. 2021; 54(4): 1–37. doi:10.1145/3447772
23. Hornik K. Approximation capabilities of multilayer feedforward networks. Neural Networks. 1991; 4(2): 251–257.
24. Brown TB, Mann B, Ryder N, et al. Language models are few-shot learners. Advances in Neural Information Processing Systems. 2020; 33: 1877–1901.
25. International Data Corporation (IDC). Artificial Intelligence Infrastructure Spending to Surpass the $200Bn USD Mark in the Next 5 years, according to IDC [Internet]. Available online: https://www.idc.com/getdoc.jsp?containerId=prUS52758624 (accessed on 13 December 2024).
26. Williams AE. Defining a continuum from individual, to swarm, to collective intelligence, and to general collective intelligence. International Journal of Collaborative Intelligence. 2021; 2(3): 205–209.