Cross-Domain Imaginative AI through the Concept of Pancakkhandha: An Integrated Framework for Semantic Preservation and Creative Knowledge Generation

Thatsanee CHAROENPORN [a,1], and Virach SORNLERTLAMVANICH [a,b]
[a] *Asia AI Institute (AAII), Faculty of Data Science, Musashino University, Japan*
[b] *Faculty of Informatics, Burapha University, Thailand*
ORCiD ID: Thatsanee Charoenporn https://orcid.org/0000-0002-9577-9082,
ORCiD ID: Virach Sornlertlamvanich https://orcid.org/0000-0002-6918-8713

Abstract. We present a framework that combines Buddhist philosophy's Five Aggregates (Pañcakkhandhā) with Cross-Singular Value Decomposition (Cross-SVD) to reduce hallucination in Large Language Models while maintaining creative capability. Each aggregate maps to an AI component: Form (Rūpa) represents data structure, Feeling (Vedanā) handles attention weighting, Perception (Saññā) manages knowledge retrieval, Volitional Formations (Saṅkhāra) control creative processing with Cross-SVD as a semantic anchor, and Consciousness (Viññāṇa) provides human oversight. Cross-SVD preserves essential structural constraints while permitting creative variation in non-critical dimensions. Businesses can use this to generate product designs, marketing strategies, and service concepts that are both novel and technically feasible for solving the enterprise AI challenge of balancing innovation with reliability.

Keywords. Imaginative AI, Five Aggregates, Semantic Prevention, Creative Knowledge Generation, Cross-SVD

1. Introduction

Large Language Models have become proficient at generating human-like text. They can write essays, create business plans, and produce creative content. But proficiency isn't understanding. These systems predict what comes next based on statistical patterns, which works well for familiar tasks but fails when creativity matters, when we need genuinely new ideas that still connect to reality.

The failure mode shows up as hallucination: plausible-sounding outputs that are factually wrong [1,2]. An AI might cite papers that don't exist, invent historical events, or propose technically impossible products [3]. For businesses exploring AI-driven innovation, this isn't just annoying, it's disqualifying. You can't rely on a system that might confidently present fiction as fact.

We looked for solutions in an unexpected place: Buddhist philosophy. The Five Aggregates framework describes how humans process information from raw sensation

[1] Corresponding Author: Thatsanee Charoenporn, thatsanee@ds.musashino-u.ac.jp

through evaluation, recognition, mental formation, to conscious awareness. Each stage has a clear function and connects to the others. Modern AI has analogous components (embeddings, attention, retrieval, generation), but they often work in isolation rather than as an integrated system.

Our framework makes these connections explicit. Form (Rūpa) becomes the data representation layer. Feeling (Vedanā) translates to attention mechanisms that weight by importance and values. Perception (Saññā) corresponds to knowledge retrieval from external sources. Volitional Formations (Saṅkhāra) handle creative generation, constrained by Cross-SVD, a mathematical technique that anchors imagination to verified knowledge. Consciousness (Viññāṇa) brings in human judgment at critical points.

By applying Cross-SVD, the system decomposes both reality (verified structural facts) and imagination (novel possibilities) into principal components, then identifies which dimensions must be preserved. Think of designing a chair in Van Gogh's style: the system can freely modify color and texture (low singular values) but must maintain structural integrity (high singular values). A threshold parameter controls the creativity-safety trade-off.

This paper provides architectural specifications designed for independent implementation and comparison. Section 2 covers background on hallucination and relevant philosophy. Section 3 maps the Five Aggregates to AI components. Section 4 describes the complete architecture. We demonstrate the practical application of this principle, using a chair design scenario to show how the system maintains structural integrity during the creative process in Section 5. Section 6 discusses business implications. And section 7 is the conclusion.

2. Background and Related Concepts

Current AI systems work by finding patterns in massive amounts of text and predicting which word should come next. They're very good at recognizing what sounds right, but they don't understand what things actually mean [4]. It's like a skilled parrot that can repeat complex phrases without grasping their meaning.

When humans create knowledge, we do something different. We build on existing understanding, check whether ideas make logical sense, consider context, and verify connections to reality [5]. We maintain what philosophers call semantic anchoring, keeping our thoughts connected to what they refer to in the real world. This gap becomes obvious when AI tries to innovate [6]. Innovation means generating new ideas, but those ideas have to be grounded in reality to be useful. An AI might propose a business strategy that sounds brilliant but ignores basic market realities, or suggest a scientific experiment but produce results that don't connect meaningfully to existing knowledge.

Philosophical Foundations for AI System Design

Philosophy has investigated the formation of knowledge for millennia, offering sophisticated models of cognition that remain highly relevant to modern computing. Rather than viewing Buddhist philosophy as a mystical concept, it can be approached as a systematic description of how perception functions, not as a single event, but as a series of interconnected stages from initial sensation to conscious awareness.

The relevance to modern AI lies in the integration of these processes. While current architectures often treat components like embeddings or attention mechanisms as isolated modules, the aggregate framework emphasizes that each stage of information processing depends on and feeds into the others. By adopting this holistic approach, we can introduce a form of "meta-awareness" into AI systems, enabling them to distinguish between evidence-based outputs and statistical confabulation. Ultimately, philosophy provides a pragmatic guide for determining whether structuring AI according to these principles can yield more reliable and innovative results.

3. Cognitive Framework: The Five Aggregates (Pañcakkhandhā)

In Buddhist philosophy, the Five Aggregates describe the components that come together to form what we experience as a person or a self [7, 8, 9, 10]. These aggregates aren't physical parts like organs, they're processes or aspects of experience. In short, the Five Aggregates represent a process-oriented model of how we experience reality. Rather than viewing a "self" as a static entity, this framework describes a sequence of mental and physical functions that work together. The five Aggregates consist of

Form (Rūpa) covers all physical phenomena including the body, material objects, and their properties. It's the stuff we can see and touch which provide the essential data or "matter" that our senses interact with.

Feeling (Vedanā) refers to the immediate evaluative quality of an experience or sensations, whether it is perceived: pleasant, unpleasant, or neutral. It will subsequently shape our response.

Perception (Saññā) is the ability to recognize, identify things, to see patterns and remember what they mean. It's how you recognize a chair as a chair.

Mental Formations (Saṅkhāra) is the active and constructive layer of the mind where intentions, habits, and creative impulses reside. This aggregate includes things like compassion, greed, faith, and confusion. It is here that raw data is transformed into complex thought and action.

Consciousness (Viññāṇa) is the basic awareness that makes perception possible, seeing, hearing, smelling, tasting, feeling, and thinking. It acts as the base that allows the other four aggregates to function and interact.

Mapping to AI Architecture

Each aggregate maps to a component in our AI system. Data enters as raw input (Form), gets weighted by importance and values (Feeling), matches against existing knowledge (Perception), undergoes creative transformation while anchored to reality (Mental Formations), and receives human validation (Consciousness).

The key is integration. Quality embeddings improve attention effectiveness. Better attention enhances retrieval accuracy. Good retrieval constrains generation appropriately. Human feedback refines all earlier stages. Each layer influences the others, maintaining semantic coherence from input to output.

This isn't just philosophical window dressing. The aggregate framework provides specific design principles: separate concerns cleanly, make dependencies explicit, enable bidirectional information flow, maintain awareness of system state. These principles address concrete problems in current generative AI.

4. Framework Architecture

The imaginative AI framework integrates the Five Aggregates (Pañcakkhandhā) as a systematic cognitive architecture with Cross-SVD semantic anchoring to prevent AI hallucination while enabling creative knowledge generation [11, 12]. The framework has five layers corresponding to the five aggregates, with information flowing both forward (synthesis) and backward (learning). Cross-SVD operates in Layer 4 to anchor creative generation to verified knowledge.

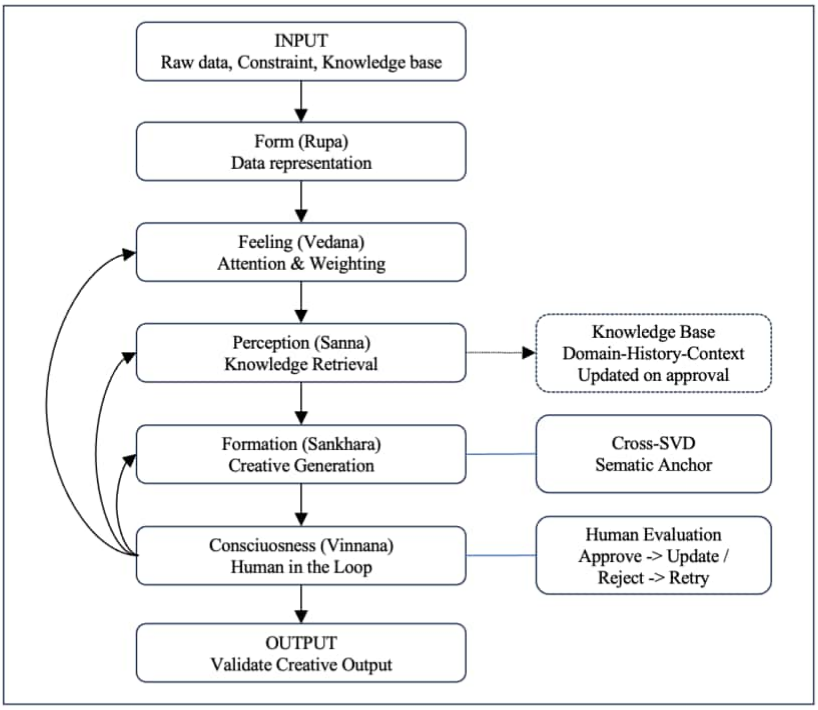

Figure 1. The mindful imaginative AI framework Overview

4.1 Forward Propagation

For the synthesis process, data moves through five transformations including:

Layer 1 (Form/Rūpa): Converts raw input such as text, images, structured data into feature vectors. Standard embedding models work here.

Layer 2 (Feeling/Vedanā): Attention mechanism weights features by importance. Unlike standard attention that optimizes purely for statistical relevance, this layer incorporates explicit value constraints: safety requirements, ethical guidelines, sustainability criteria.

Layer 3 (Perception/Saññā): Retrieval from external knowledge base. Stores verified facts separately from model parameters, allowing updates without retraining. Links inputs to established knowledge through similarity search.

Layer 4 (Mental Formations/Saṅkhāra): Creative generation constrained by Cross-SVD. Takes retrieved knowledge ("reality matrix") and desired output space ("imagination matrix"), decomposes both via SVD, identifies common principal components, uses threshold τ to determine which dimensions must be preserved. Low τ allows more creativity; high τ enforces stricter reality constraints.

Layer 5 (Consciousness/Viññāṇa): Human validation. Experts check outputs for correctness, usability, and ethical compliance. Not just a quality gate, the feedback loops back to refine earlier layers.

4.2 Backward Learning

Feedback from the Consciousness layer is sent back to update the system continuously, as follows.

Adjust (to Formation): tunes the balance between structural reality and creative freedom, by adjusting threshold τ based on approval/rejection patterns.

Update (to Perception): saves successful results into the Knowledge Base, allowing the AI to "remember" and improve over time.

Refine (to Feeling): fine-tunes the Attention Weights so the AI learns exactly which features humans value most to reflect human preferences.

5. Business Use Case Scenario: Van Gogh-Style Chair Design

In the case of designing a chair based on the artistic style of Van Gogh, this framework transforms AI from a conventional drawing tool into "Imaginative AI" operating within a mindfulness cycle, with Cross-SVD acting as a semantic anchor. This emphasizes that the designed Van Gogh-style chair is imaginative yet functional and reflects the artist's spirit with integrity.

5.1 How Cross-SVD Anchors Semantics

The use of Cross-Singular Value Decomposition (Cross-SVD) within this framework is to create a "mathematical compass" to prevent the AI's imagination from spiraling out of control and becoming hallucinatory. The details of its operation are as follows.

1. Analysis of Singular Values
   Decompose chair structural requirements (weight capacity, ergonomics, stability) and Van Gogh's visual characteristics (swirling brushstrokes, bright colors, textural variation) into principal components via singular value decomposition.
2. Functioning as "Semantic Anchor"
   Cross-SVD identifies shared dimensions. High singular values in the structural matrix (like backrest angle for lumbar support) must be preserved. Low singular values (like exact color hex codes) allow freedom. Cross-SVD will calculate the most important common points (principal components) between reality and imagination. If the AI attempts to design a backrest that is so flexible that the structure cannot actually support the weight, the singular values will conflict. The system will use the values from the "reality" side to anchor and support the structure.
3. Semantic Preservation

The goal is to allow the AI to freely "modify" (create) aspects that do not affect the core function, such as allowing it to apply bright yellow paint or create a rough brushstroke-like surface, but while maintaining the "meaning" of a safe and functional chair. (Threshold τ sets the boundary. If τ is high, even medium-importance structural features are preserved, resulting in a safe but conventional chair. If τ is low, more artistic liberties are allowed, but core safety features still can't be violated.)

Table 1. Comparison Before and After Using Cross-SVD

Design Point	No Cross-SVD (AI Hallucination)	Applying Cross-SVD (Semantic Preservation)
Structural Integrity ex. Structure of chair legs	May create non-functional designs. (The chair legs are airborne or can twist to the point of tipping over.)	Maintain stability at all four points of the chair leg, according to engineering principles. (Maintains structural safety)
Artistic Expression ex. Relationship between colors	Uncontrolled stylistic elements. (The colors may blend together to the point where it's difficult to distinguish which areas are seating.)	A Van Gogh-style color palette is used, but the proportions and functions are clearly defined. (Balanced creativity within constraints.)
Semantic Preservation ex. Safety	Loss of core chair functionality. (The backrest may have sharp edges resulting from the design of the lines.)	The flowing lines have been adjusted to be curved and are not dangerous to users. (Chair remains recognizable and usable.)

While Cross-SVD performs the mathematical verification, humans act as "mindfulness agents" at the Consciousness layer, making the final decision on whether the balance calculated by the AI "meets business and artistic needs." If the system begins to deviate, humans will intervene to bring the AI back into the established framework.

Table 2. illustrates how the “Five Aggregates” philosophy is applied as a systemic architecture to address the issue of AI hallucination. Instead of treating AI modules as isolated components, this framework emphasizes a continuous, interconnected process. Layers 1-3 is the Data Processing and Contextualization where the system first converts raw input into mathematical representations (Form). It then assigns importance through attention mechanisms based on specific value criteria (Feeling), while simultaneously anchoring the input to verified external knowledge bases (Perception). The Core of Creative Generation is in Layer 4. Within the Formation layer, Cross-SVD acts as a “semantic anchor”. This allows the system to balance creative imagination with physical reality, for instance, enabling artistic freedom in a chair’s design while strictly preserving the structural dimensions necessary for load-bearing and safety. Governance and Iterative Learning (Layer 5) is the final stage. It involves human-in-the-loop oversight (Consciousness). Experts validate the output for artistic integrity, engineering feasibility, and ethical standards, providing a feedback loop that refines the system’s parameters for future iterations.

Table 2. Integrated Framework for Creative AI Design (Van Gogh-Style Chair Case Study)

The Five Aggregates	AI Technical Component	Systemic Role & Logic	Practical Application (Chair Design)
Form	Feature Vectors & Embeddings	Converts raw input into mathematical representations.	Transforms Van Gogh’s visual style and chair engineering specifications into shared data vectors.
Feeling	Attention Mechanisms	Assigns importance and weights based on specific value criteria.	Prioritizes artistic "flow" for aesthetics while ensuring structural joints are weighted for safety.
Perception	Knowledge Retrieval	Connects input patterns to verified external knowledge and memory.	Anchors the design to authentic Impasto brushstroke identities and historical furniture-making techniques.
Formation	Cross-SVD Mapping	Facilitates creative generation while maintaining structural integrity (Semantic Anchor).	Blends creative imagination with physical reality, ensuring decorative patterns do not compromise load-bearing functions.
Consciousness	Human-in-the-Loop	Provides conscious governance, expert validation, and iterative learning.	Expert review of artistic integrity and manufacturability, with feedback used to refine system parameters.

By applying this integrated approach, organizations can precisely tune the “creativity-safety trade-off” through a single parameter. The structural alignment verify that AI-driven innovation produces results that are not only novel but also structurally sound and viable for real-world application.

To validate this properly, we require (1) Expert ratings on artistic fidelity (Does it actually feel like Van Gogh?) (2) Engineering tests on structural integrity (Can you sit on it safely?) (3) Manufacturing feasibility (Can this be built at reasonable cost?) (4) Comparison against baseline approaches (Standard generative models, manual design processes).

6. Discussion

This framework addresses hallucination by preventing it architecturally rather than filtering it post-hoc. The Five Aggregates provide design principles: separate data from attention from retrieval from generation from validation. Cross-SVD provides the mathematical anchor: identify what must be preserved, allow freedom elsewhere.

Three components matter most. Layer 3 (Perception/Saññā) separates verified knowledge from model parameters, making the system's assumptions explicit and updatable. Layer 4 (Saṅkhāra) with Cross-SVD gives precise control over the creativity-reliability trade-off through a single parameter. Layer 5 (Consciousness/Viññāṇa) creates a learning loop where human feedback improves the system over time [13].

For businesses, it's important because innovation requires both creativity and reliability. You can't just optimize for novelty, you need novel ideas that actually work. The threshold parameter τ lets you tune this explicitly: dial it down for brainstorming, dial it up for production. And human oversight isn't bolted on as an afterthought, but it's built into the architecture.

7. Conclusion

This paper has proposed an imaginative AI framework that weaves together the Buddhist doctrine of the Five Aggregates and the mathematical rigor of Cross-SVD to tackle hallucination in large language models while keeping space for genuine creativity. By mapping form, feeling, perception, mental formations, and consciousness onto distinct yet interconnected AI components, we showed how semantic coherence can be preserved from initial input through to human validation.

The Van Gogh–style chair case study illustrated how this architecture works in practice: structural constraints such as stability and ergonomics are anchored in high-value singular components, while lower-value dimensions carry stylistic freedom in color, texture, and visual expression. Rather than filtering out hallucinations after the fact, the framework embeds safeguards into the generation process itself, making it possible to tune the balance between safety and originality through an explicit threshold parameter and human-in-the-loop oversight.

For businesses, this offers a concrete path toward AI systems that can propose new products, services, and strategies that are not only imaginative but technically feasible and ethically defensible. Organizations can adjust the framework for different stages of innovation, from low-risk ideation with looser constraints to high-stakes deployment with strict preservation of verified knowledge and domain rules.

Future work will focus on empirical validation in real design and planning environments, including expert ratings of artistic fidelity, engineering tests, manufacturability assessments, and comparisons with standard generative baselines. Beyond the specific case of chair design, the same principles can extend to domains such as architecture, urban planning, and policy simulation, where creative exploration must remain grounded in reliable structures and human values

References

[1] Huang L, Yu W, Ma W, Zhong W, Feng Z, Wang H, Chen Q, Peng W, Feng X, Qin B, and Liu T. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. ACM Transactions on Information Systems 1, 1, Article 1 (January 2024), 58 pages. https://doi.org/10.1145/3703155.

[2] Ji Z et al. Survey of hallucination in natural language generation. ACM Computing Surveys. 2023; 55(12), 1-38.

[3] Lyre H. Understanding AI: Semantic Grounding in Large Language Models. 2024. https://arxiv.org/abs/2402.10992.

[4] Vaswani A et al. Attention is all you need. In Advances in Neural Information Processing Systems. 2017; 30.

[5] Aguilar C. AI Hallucinations: Causes, Cures, and Human-in-the-Loop Oversight. Mexico Business News. 2025.
[6] Farquhar S, Kossen J, Kuhn L, Gal Y. Detecting hallucinations in large language models using semantic entropy. Nature. 2024;630.625-630.10.1038/s41586-024-07421-0.
[7] Bodhi B. The Connected Discourses of the Buddha: A Translation of the Samyutta Nikaya. Wisdom Publications. 2000.
[8] anattā and artificial intelligence. Dhamma Footsteps. 2023.
[9] Five Aggregates of Machine Learning" Bodhi Day blog, 2024. https://bodhiday.org/2024/03/24/five-aggregates-of-machine-learning/
[10] Rahula W. What the Buddha Taught (2nd ed.). Grove Press. 1974.
[11] Sornlertlamvanich V, Ruangraweenukit P, Chophuk P, Viriyavi W, and Charoenporn T. Toward Imaginative AI: Domain Transitivity for Cross-Domain Imagination. Proceedings of the 36 International Conference on Information Modelling and Knowledge Bases (EJC2026), Vaasa, Finland, June 8-12, 2026. (to be published)
[12] Sornlertlamvanich Virach. Enhancing Large Language Models: Alleviating Knowledge Deficiency with External Knowledge and Semantically Aware Reasoning (SAR). Knowledge and Information Systems, Springer Nature, pp. 1–19, 2025, doi: 10.1007/s10115-025-02485-4
[13] Indium. Taming AI Hallucinations: Mitigating Hallucinations in AI Apps with Human-in-the-Loop Testing, Quality Engineering, 26th May 2025.