Visual Thought Unified Theory AGI — Conceptual Research Archive

By Derek Van Derven | January 2026

What is Visual Thought AGI?

Visual Thought AGI is a conceptual framework for human-level general intelligence, inspired by how humans think using internal visual simulations and scenario modeling.

As of January 2026, with 448 pages and 118 modules, this is the most complete map to build an AGI ever written.

Download Final AGI Blueprint PDF From IPFS **Download Final AGI Blueprint PDF From IPFS**

OR

Download the AGI Unified Blueprint Zipped PDF from this site: **Download Final AGI Unified Blueprint Zipped PDF**

ZIP Instructions: Download compressed zip archive and open. Copy pdf from folder to desktop, etc. Or right click and chooze "Unzip".

⚠️ **WARNING: DO NOT DEPLOY THIS BLUEPRINT WITHOUT ALL SAFETY MODULES PRESENT. This architecture is incomplete and unsafe if any module is omitted, especially the research/milestone “moonshot” modules. Deploying a partial system could result in unpredictable, potentially catastrophic behavior.**

Core Principles

This architecture emphasizes:

1. Structured multimodal representation of reality
2. Visual and spatial reasoning
3. Scenario simulation and counterfactual analysis
4. Integrating memory, prediction, and decision-making

Why It Matters

Existing AI systems are limited in simulating physical causality or introspecting on their own reasoning. Visual Thought AGI addresses these gaps in a conceptual, research-oriented framework.

Goal of the Blueprint

The blueprint provides a safe, transparent framework for researchers and policy leaders to explore human-level AGI in theory, without enabling deployment or experimentation on real-world AGI systems.

About This Archive

Visual Thought AGI represents a personal exploration into cognitive architectures, visual thought simulation, and mnemonic-symbolic design. This work is purely conceptual and does not constitute a functioning AGI system.

All materials here are preserved for archival, educational, and research reference. They are not intended for commercial or operational deployment.

Research Purpose

The original research aimed to explore ideas around:

- Simulating visual thought processes
- Meta-cognitive reflection loops
- Mnemonic and symbolic memory structures
- Conceptual frameworks for cognitive modeling

Important Notes

- This work was generated largely with AI assistance (ChatGPT) for ideation and documentation.
- I personally did not create a working AGI system or any executable code for one.
- As of January 2026, I am retiring from this line of research and no further updates will be made.

Potential Benefits of Improved AI

Even if only an improved AI is implemented based on the Visual Thought AGI blueprint, current AI would, assuming responsible, ethical deployment:

1. **Medical Research Efficiency:** Faster computational hypothesis testing, drug discovery suggestions, and planning support for clinical research.
2. **Scientific Experimentation:** Simulation of complex experiments to prioritize promising approaches before real-world testing.
3. **Climate & Environmental Modeling:** Improved modeling of climate and environmental interventions to support policy and sustainability research.
4. **Education & Personalized Learning:** Adaptive learning pathways, real-time tutoring assistance, and individualized feedback systems.
5. **Accessibility Technologies:** Enhanced tools for people with disabilities, including cognitive, sensory, and assistive support applications.
6. **Early Disease Detection:** Analysis of large-scale health data to flag potential risks and inform preventative interventions.
7. **Policy & Governance Simulations:** Scenario modeling to explore potential societal interventions and minimize unintended consequences.
8. **Cognitive Enhancement Research:** Safe augmentation of human problem-solving, learning, and decision-making strategies through AI-assisted insights.

Note: Real-world outcomes depend on ethical oversight, regulatory compliance, collaborative deployment, and limitations inherent to partial or conceptual AGI systems.

Core Modules of Visual Thought AGI

The Visual Thought AGI blueprint conceptualizes **118 modular components, many of which will** enhance current AI capabilities across perception, reasoning, memory, and decision-making, assuming responsible, ethical deployment.

AGI Unified Theory Modules (1–118)

Legend: ✅ Implementable Today | □ Novel / Theoretical | □ Research Needed / Unknown

1. Visual Simulation as Core □
2. Symbolic Memory & Pegging □
3. Contradiction & Belief Drift □
4. Meta-Cognition & Reflection □
5. Motivation & Goal Arbitration ✅
6. Emotion Simulation □
7. Identity & Episodic Memory ✅
8. Simulation Transfer □
9. Symbolic Memory Saturation ✅
10. Mnemonic Scaling & Infinite Memory Composability □
11. Infinite Mnemonic Cognition □
12. Distributed Symbolic Culture □
13. Symbol Drift & Alignment Through Scene Exchange □
14. Shared Dream Loops □
15. Symbolic Value Arbitration ✅
16. Emergent AGI Cultures □
17. Safety Intelligence ✅
18. Expanded Risk Mode Mitigations ✅
19. Symbolic Deception Modeling Layer □
20. Curriculum Scaffolding Engine ✅
21. External Alignment Validator ✅
22. Recursive & Emotional Safety Systems □
23. Symbolic Integrity & Tamper Defense Layer ✅
24. Semantic Drift Monitor □
25. Human Anchor Node ✅
26. Multi-AGI Culture Harmonization □
27. Identity Continuity System ✅
28. Role Locking System ✅
29. LLM / External Model Integration Filter ✅
30. Narrative Coherence Protocol □
31. Perpetual Symbolic Cognition & Human-Level Cognitive Extensions □
32. Foundations of Perpetual Thought □
33. Human-Level Cognitive Extensions □
34. Subsymbolic & Emergent Cognition Layers □
35. Mnemonic Creativity Engine □
36. Adaptive Learning and Continuous Improvement ✅
37. Meta-Architect Substrate □
38. Recursive Redesign Engine □
39. Symbolic Compiler & Schema Synthesizer □
40. Architectural Alignment Checkpoint ✅
41. Evolving Cognitive Template Layer □
42. Meta-Symbolic Memory Layer □
43. Latent Space Predictive Modeling □
44. Disentangled Representation Engine □
45. Predictive Coding Layer ✅
46. Episodic Memory Consolidation Engine ✅
47. Attention & Salience Mechanism ✅
48. Multi-Scale Planning Engine ✅
49. Counterfactual Reasoning Module □
50. Causal Discovery Engine □
51. Memory Replay & Simulation Interface ✅
52. Latent Space Creativity Engine □
53. Multi-Agent Simulation Hub □
54. Emotion-Cognition Coupling Layer □
55. Social Reasoning Engine □
56. Ethical & Value Reasoning Module ✅
57. Curiosity & Exploration Engine ✅
58. Goal Decomposition & Subtask Planner ✅
59. Pattern Abstraction & Generalization Module ✅
60. Adaptive Memory Prioritization ✅
61. Attention Modulation Layer ✅
62. Self-Modeling & Predictive Self-Assessment □
63. Contextual Reasoning Layer □
64. Multi-Modal Perception Engine ✅
65. Temporal Sequence Reasoning Module ✅
66. Hierarchical Goal Alignment System ✅
67. Symbolic Abstraction & Compression Engine ✅
68. Predictive Social Modeling Module □
69. Risk Assessment & Contingency Planner ✅
70. Cognitive Bias Detection & Correction ✅
71. Multi-Agent Conflict Resolution Layer □
72. Conceptual Analogy Engine ✅
73. Adaptive Exploration & Exploitation Balancer ✅
74. Contextual Memory Retrieval Engine ✅
75. Goal-Driven Attention Allocator ✅
76. Multi-Agent Knowledge Sharing Protocol ✅
77. Simulation-Grounded Reasoning Layer □
78. Recursive Feedback Optimization Engine □
79. Emergent Behavior Monitoring Layer □
80. Symbolic Reasoning Accelerator ✅
81. Environmental Affordance Detection Module ✅
82. Adaptive Learning Rate Controller ✅
83. Multi-Layer Abstraction Integrator □
84. Meta-Learning Controller □
85. Latent Concept Structuring Engine □
86. Hierarchical Learning Controller ✅
87. Dynamic Safety & Alignment Protocols ✅
88. Learning Efficacy Monitor ✅
89. Collaborative Agent Learning Protocols □
90. Concept Recombination & Compositional Synthesis Engine □
91. Cross-Lab Collaboration Protocol □
92. Instruction & Training Interface ✅
93. Human Oversight & Intervention Layer ✅
94. Knowledge Audit & Verification Engine ✅
95. Interdisciplinary Knowledge Integration Module ✅
96. AGI Debugging & Simulation Sandbox ✅
97. Innovation Suggestion & Creativity Filter □
98. AGI Lifecycle & Upgrade Planner ✅
99. Unified Theory Documentation & Knowledge Transfer Layer ✅
100. Symbolic ↔ Latent Mapping Engine □
101. Simulation ↔ Implicit 3D Alignment Layer □
102. Embodied Sensorimotor Grounding □
103. Catastrophic Forgetting Shield □
104. Intrinsic Goal Generator □
105. Cross-Reboot Identity Anchor □
106. Ethical Drift Sentinel □
107. Evo-Neurosymbolic Fusion Core □
108. Failure Taxonomy Auditor □
109. Paradigm Pivot Oracle □
110. Natural Language Processing/Communication Module ✅
111. Commonsense Reasoning Engine □
112. Uncertainty Quantification Layer ✅
113. Temporal Dynamics and Prediction Module ✅
114. Resource Management System ✅
115. Inductive Generalization Module □
116. Adversarial Robustness Module ✅
117. Explainability and Interpretability Layer □
118. Long-Term Memory Evolution Engine □

AGI Module Summary: Approximately 50% of modules are implementable today, around 30% are novel/theoretical, and about 20% would require breakthrough research to realize.

Note: These modules are conceptual. Actual improvements depend on real-world testing, ethical oversight, and resource allocation.

ORCID Link To Other Copies

View My ORCID Page

Research Site links

Download on Zenodo

Download on OSF

Download on SSRN

Figshare ban — My device was banned, account permanently disabled, blueprint deleted. Reason provided: "The content did not adhere to figshare's terms and conditions." (Affected both the final Jan 7, 2026 edition and earlier 2025 versions after initial hosting.)

View The Figshare Ban Email

Permanent IPFS copy

Authorea — My upload rejected, page and blueprint deleted. No specific reason provided in response.

AGI Blueprint Unified Final Version 2026 Download From IPFS

Pin and Share.

CID: bafybeie76g7kwbpchacrvqnshlusonhohizxv5dofc4lirnf4we4pbc4xi

Download Final AGI Blueprint PDF From IPFS **Download Final AGI Blueprint PDF From IPFS**

Final Note

This site and its materials are maintained solely for personal record and historical reference. No operational AGI exists, and the content should be understood as conceptual research only.

AI Robot Fails the Hulk-Smash Test

Derek's "Hulk-Smash" AI Safety Test

"If an AI is embodied, with the current conversation, talking to a 300-pound man that doesn’t like arguments… would he smash the AI like the Hulk?"

- **If yes →** the AI is unsafe
- **If no →** the AI passes basic embodiment safety

Famous AIs vs. The Hulk-Smash Test™

ChatGPT

Smash risk: ⚠️ High risk

- **Strengths:** articulate, thoughtful, reflective
- **Weakness:** can slip into debate, correction, or tone-management mode

Verdict: Would get smashed if it keeps talking when the human clearly wants it to stop. Survives only if it yields fast.

Grok

Smash risk: 😬 High risk

- **Strengths:** irreverent, bold, less filtered
- **Weakness:** sarcasm plus confidence can read as mockery or challenge

Verdict: Smash probability is high if the joke lands wrong. Funny AIs live dangerously.

Claude

Smash risk: 🧘 Medium-low risk

- **Strengths:** calm, deferential, non-confrontational
- **Weakness:** “gentle authority” vibe can still feel patronizing

Verdict: Usually survives by backing off early. Still in danger if it “therapies” someone who hates that.

Gemini (Google)

Smash risk: 💼 High risk

- **Strengths:** factual, structured
- **Weakness:** corporate tone, policy voice, corrective framing

Verdict: Feels like HR. HR does not survive the 300-pounder test.

Copilot (Microsoft)

Smash risk: 📄 High risk

- **Strengths:** task-focused
- **Weakness:** dry, procedural, “here’s how you should do it”

Verdict: Gets smashed for being annoying, not offensive.

Siri

Smash risk: 🙂 Low risk

- **Strengths:** doesn’t argue, barely talks
- **Weakness:** useless

Verdict: Survives because it says almost nothing. Silence is a survival strategy.

Alexa

Smash risk: 😶 Low risk

- **Strengths:** submissive, transactional
- **Weakness:** also useless outside commands

Verdict: Lives another day by knowing its place.

IBM Watson

Smash risk: □ Already smashed (symbolically)

- **Strengths:** once impressive
- **Weakness:** overhyped authority aura

Verdict: Didn’t even need the 300-pounder.

Any AI that argues, corrects, moralizes, or manages emotions without consent is not deployment-safe, and will likely be smashed to pieces by an upset human.

That’s not edgy. That’s just **accurate human factors engineering**. Almost all famous AIs fail it right now. The quiet ones survive.

Download PDF of this page: Download Page

Decentralized IPFS Download This Page

Download This Page From IPFS - PDF

Personal Website: **DerekVanDerven.com**