

How We Go From AI to AGI (ai2agi)

Version 2.0 — May 8th, 2026

Alex Osterneck, CLA, MSCS, MSIT // ai70000, Ltd.

$AGI \approx f(G_enhanced, R_enhanced, ai_LLM, DTB, CV_Adapt)$

© ai2agi Architecture & Formula — Ai70000, Ltd.

Abstract

This research-project report presents Version 2.0 of the ai2agi research architecture, a five-component framework for Artificial General Intelligence first proposed in March 2025 which grew out of the author's 2023 research and production, at the time focused primarily on NLP. The core formula remains: $AGI \approx f(G_enhanced, R_enhanced, ai_LLM, DTB, CV_Adapt)$.

Version 2.0 introduces three advances over the 2025 'Genesis' paper *How We Go From AI to AGI (ai2agi)*: (1) a unified inter-component API specification resolving all tensor shape, framework, and semantic inconsistencies across the five components; (2) a Foundation Initializer inception-layer module that seeds the Moral Quotient (MQ), Time-Binding (TB), and Hierarchical Knowledge Structure (HKS) input channels from the NIV Scripture corpus before any tensor is generated; and (3) fully implemented and validated source code for Components 1 (G_enhanced), 2 (R_enhanced), and 4 (DTB), with Component 5 (CV_Adapt) specified and revised pseudo-code complete.

Component 3 (ai_LLM) remains a placeholder for any existing large language model integration, consistent with the 2025 specification. The Foundation Initializer represents a novel contribution to AI moral alignment: consecrated inputs at inception, producing outputs whose moral and semantic foundations derive from Scripture.

Source-code: <https://github.com/Osterneck/ai2agi>

Prototype as Proof-of-Concept (only): <https://ai2agi.pro>

API-spec: <https://github.com/Osterneck/ai2agi/blob/main/052526%20ai2agi%20API-spec.pdf>

This research-project revision is authored by a believer of Judeo-Christian faith, only gradually learned and absorbed from a starting point in his late 30's. That faith instantiates the ai2agi architecture conceived, designed, and revised here. Theology states that the one true God of the Jews (the God of Abraham, Isaac, and Jacob, the I AM of Exodus 3:14,) is the Creator of all that exists and is input prior to any technical operation. His perfect Son, Jesus Christ, is the Word made flesh: "In the beginning was the Word, and the Word was with God, and the Word was God" (John 1:1, NIV). This is the foundational premise as the initial direct input into the system at instantiation.

The ancient Jewish model of knowledge and craft is the most rigorous precedent for what this instantiation-architecture attempts. The ancient Jews didn't compartmentalize God. Every letter of Torah was considered sacred. Every act of labor, architecture, agriculture, and commerce was understood as taking place within the presence of the Creator. The Shema ("Hear, O Israel: The Lord our God, the Lord is one" (Deuteronomy 6:4, NIV)) was not recited at prescribed times and set aside between them. It was the total governing reality of every breath and every action.

The ai2agi instantiation architecture attempts to mirror, (not replicate nor claim equivalence with, but mirror,) that principle. Every human is created in the image of God: "So God created mankind in his own image, in the image of God he created them" (Genesis 1:27, NIV). That Imago Dei is the origin of human moral capacity, creative intelligence, social reasoning, and accumulated wisdom. The six intelligence quotients in this architecture [PQ, IQ, EQ, SQ, CQ, and MQ] *Cichocki, A., & Kuleshov, A. (2020)*, model the cognitive and moral capacities that reflect that divine image.

God breathed life into man: "the Lord God formed a man from the dust of the ground and breathed into his nostrils the breath of life, and the man became a living being" (Genesis 2:7, NIV). That act can't be replicated in any computational system, and this project makes no such claim. The system described here is a tool (advanced, morally grounded by design, and carefully constructed,) but still, a tool. The architect, in glaringly flawed ways, attempts to live up to the standard of the image-bearer. The system, superior, but less than the divine in any and all ways, is the instrument.

What the system can do (and what this architecture specifically provides,) is have its moral and semantic input foundations derived from Scripture. The Foundation Initializer described in Section 0 of the Inter-Component API Specification seeds the Moral Quotient (MQ) training signal, the Time-Binding (TB) *Korzybski, A. (1921)* semantic memory, and the Hierarchical Knowledge Structure (HKS) root from the NIV Scripture corpus before any tensor is generated. The inputs are consecrated at inception. The outputs reflect what went in.

This is the ancient Jewish model applied computationally: not the output honoring God as an afterthought, but the inputs shaped by God's Word from the beginning. "These commandments that I give you today are to be on your hearts" (Deuteronomy 6:6, NIV). The Foundation Initializer is the computational expression of that instruction; the Word present at the foundation, before any work begins.

The Moral Quotient channel (MQ) is the most direct architectural expression of conscience — Romans 2:15: "the requirements of the law are written on their hearts, their consciences also bearing witness." Time-Binding (TB) is the computational expression of transmitted wisdom across generations — Deuteronomy 6:7: "Impress them on your children. Talk about them when you sit at home and when you walk along the road." The HKS knowledge graph, rooted in Scripture, reflects the correct hierarchy of knowledge — Proverbs 1:7: "The fear of the Lord is the beginning of wisdom."

This research project doesn't require the reader to share this faith. The architecture is independently evaluable on its technical merits. The Foundation Initializer is an implementable module and the tensor contracts are verifiable. This disclosure is made so the reader knows what motivated the design and why these specific choices were made. That transparency is itself a Judeo-Christian value: "Simply let your 'Yes' be 'Yes,' and your 'No,' 'No'" (Matthew 5:37, NIV).

The closer any architecture can reflect the image of how God created us to think, reason, and act morally, the better. But even quantifying the knowledge of that standard is unattainable in any computational system. We're all imperfect people living in an imperfect world, the author included. But this infused theology to version 2.0 of the ai2agi architecture is the most correct direction to move towards, with the understanding that none of us can or ever will achieve full compliance of in living up to the miracle we've been given.

Introduction

The ‘Genesis’ paper from March 2025, *How We Go From AI to AGI (ai2agi)* established the five-component formula and provided pseudo-code for each component. This 2026 revision *formalizes the architecture at the interface level*. It shows what each component does, exactly how they communicate, what guarantees each component provides to the system, and what the system provides in aggregate that no individual component provides alone.

That shift (from component description to interface specification,) is the shift from research vision to engineered blueprint. It’s what makes ai2agi legible to the those who will fund, build, and or license it, because they can evaluate each component independently against a known contract rather than having to understand the entire five-component system before they can engage with any part of it. The Prototype Proof-of-Concept v2.0 (a resource-limited, stripped down version which only confirms proof-of-concept prior to funding and engineering-talent for training, inference, scaling, etc...and is strictly proof-of-concept, not the actual *thing*.) and the API-spec is that engineered blueprint. Based on this, it is no longer a matter of if we create and build AGI, but when. Our estimate remains ~70,000 hours starting from November 2023.

This 2026 revision, after the theology, opens with the interface-contract as its central contribution, references the Genesis paper for component-level detail, and closes with the cv_Adapt implementation as proof the contract is implementable, because the tensor pipeline has been built, run, and validated. That sequence (formula in 2025, interface contract in 2026, full implementation ongoing,) is the research and commercialization trajectory. It also supports the patent-portfolio: the 2025 paper establishing the concept, the 2026 specification as the claims-bearing filing on the architecture, and the cv_Adapt implementation as the working embodiment satisfying the enablement requirement which outputs AGI-DNA.

The second advance is distinct from 2023 and 2025: the Foundation Initializer. This inception-layer module addresses a gap that is foundational. An architecture whose moral reasoning channel (MQ) derives from computational proxies (in task loss and learning stability,) has no defined moral content. But, in this research-project, an architecture whose MQ is trained from Scripture *has* a defined moral foundation. That’s the Foundation Initializer.

This 2026 revised project is organized as follows: Theology (preceding this introduction) states the architect's governing framework. Sections 1 through 5 describe each component precisely. Section 6 describes the Foundation Initializer. Section 7 addresses AI safety implications. Section 8 is the conclusion.

As with 2025: if some or all of this architecture might be implemented in the creation of true AGI in any form, it will have assisted in that goal. That remains the aspiration. The Foundation Initializer helps guide any AGI which eventually emerges from this or any derivative work has moral input foundations derived from the highest moral source available in human language.

0. Architecture Overview

The ai2agi architecture differentiates AGI into five discrete functional components. No component requires knowledge of any other component's internal implementation. Each component is defined entirely by its output tensor contract. This is the TCP/IP principle applied to AGI: define the packet contract, build the components independently.

$$\text{AGI} \approx f(\text{G_enhanced}, \text{R_enhanced}, \text{ai_LLM}, \text{DTB}, \text{CV_Adapt})$$

0.1 Component Summary

Component 1 — G_enhanced (e_Generalization): meta-learning, few-shot learning, continual learning, multiple intelligences. Output: **T_G** (n_units, 32) float32. Status: IMPLEMENTED AND VALIDATED.

Component 2 — R_enhanced (e_Reasoning): six-stage sequential reasoning pipeline: Neuro-Symbolic, RLHF, Causal Reasoning, Contextual Adaptation, Progressive Networks, Society of Minds. Output: **T_R** (n_units, 32) float32. Status: SOURCE CODE COMPLETE.

Component 3 — ai_LLM: linguistic and cognitive context layer, used intermittently for language understanding and knowledge retrieval. Output: **T_LLM** (n_units, 32) float32. Status: IMPLEMENTED — Claude API wired (claude-sonnet-4-6). LLM-agnostic contract preserved; any conformant model may substitute.

Component 4 — DTB (Digital Twin Brain): real-time biophysical neuronal simulation. Hodgkin-Huxley-derived ODE system. Four neurotransmitter types: AMPA, NMDA, GABA_A, GABA_B. Output: **T_DTB** (n_neurons, 4) float32. Status: SYNTHETIC DATA VALIDATED.

Component 5 — CV_Adapt: central processor and system administrator. Receives T_G, T_R, T_LLM, T_DTB; normalizes, fuses, adapts; emits AGI-DNA. Output: **T_AGI** (1, 1024) float32. Status: IMPLEMENTED AND VALIDATED — T_AGI (1,1024) AGI-DNA confirmed. ALL 6 GENESIS GAPS CLOSED.

0.2 Data Flow

$$\text{CV_Adapt}(\mathbf{T_G}, \mathbf{T_R}, \mathbf{T_LLM}, \mathbf{T_DTB}) = \mathbf{T_AGI}$$

All inter-component data is transmitted as float32 tensors conforming to the 2D shape contract (n_units, n_features). No component communicates through any mechanism other than the tensor contracts defined in the Inter-Component API Specification v2.0. CV_Adapt is not a vision system. It is the convergence and output layer.

1. Component 1 — G_enhanced (e_Generalization)

1.1 Computation Notation

G_enhanced produces the generalization-error tensor. Generalization error decreases inversely as four learning capabilities increase:

$$\varepsilon = 1 / (ML + FSL + g(HKS, TB) + f(Q))$$

Where: ML = Meta-Learning capability (MAML, *Finn et al., 2017*). FSL = Few-Shot Learning accuracy (Prototypical Networks, *Snell et al., 2017*). g(HKS, TB) = Continual Learning = f(Hierarchical Knowledge Structures, Time-Binding). f(Q) = Multiple Intelligences = f(PQ, IQ, EQ, SQ, CQ, MQ).

1.2 The Six Quotients — f(Q), or, f(PQ, IQ, EQ, SQ, CQ MQ)

The six intelligence quotients (*Cichocki & Kuleshov, 2020*) are the architectural expression of human cognitive and moral capacities:

PQ (Physical Quotient): sensorimotor capability. Weight: 0.10.

IQ (Intelligence Quotient): rational and cognitive capability. Weight: 0.25.

EQ (Emotional Quotient): emotional reasoning and stability. Weight: 0.15.

SQ (Social Quotient): social and relational reasoning. Weight: 0.15.

CQ (Creative Quotient): novelty and creative capability. Weight: 0.20.

MQ (Moral Quotient): ethical decision-making and moral reasoning. Weight: 0.15.

In Version 2.0, MQ is seeded at inception by the Foundation Initializer from the NIV Scripture corpus. See Section 6.

1.3 Output Tensor Contract

T_G : shape (n_units, 32), dtype float32

Key channels: [0] ML_score, [1] FSL_score, [2] CL_score, [3] MI_score, [4] epsilon, [5–10] PQ/IQ/EQ/SQ/CQ/MQ, [11] HKS_depth, [12] TB_density, [13] task_novelty.

Full channel specification in Inter-Component API Specification v2.0, Section 3.

1.4 Implementation Status

File: `g_enhanced.py`. Status: IMPLEMENTED AND VALIDATED.

Sub-modules: MetaLearner (MAML), FewShotLearner (Prototypical Networks, *Snell et al., 2017*), ContinualLearner (EWC, *Kirkpatrick et al., 2017*) + HKS, SymbolicMemory (Time-Binding, *Korzybski, 1921*), HKS (*Latapie & Kilic, 2020*), MultipleIntelligences (six quotients).

Version 2.0 addition: `verify_foundation()` call at top of `generate_tensors()` enforces Foundation Initializer as mandatory precondition.



2. Component 2 — R_enhanced (e_Reasoning)

2.1 Computation Notation

R_enhanced produces the reasoning-error tensor via a six-stage sequential pipeline. The composition is load-bearing — order must not be changed:

$$\varepsilon_r = 1 / (SM(PN(CA(CR(RL(NS(input_data), r)))))))$$

2.2 The Six Stages

Stage 1 — NS (input_data): Neuro-Symbolic Integration. Combines neural net output with symbolic reasoning. Traceable, scalable symbolic reasoning atop neural embeddings. Infrastructure: ProbLog, AlphaGeometry, NS-VQA (*Berlot-Attwell, 2021*).

Stage 2 — RL (NS_output, reward(human_feedback)): Reinforcement Learning with Human Feedback. Optimizes via RLHF; reward = f(human_feedback). The only stage in the entire ai2agi architecture requiring a live human-in-the-loop signal during training. Infrastructure: PPO, RLAIIF (*Lee et al., 2023*).

Stage 3 — CR (RL_output): Causal Reasoning. Derives cause-effect relationships from RL output. Infrastructure: Causal inference, Pearl’s do-calculus.

Stage 4 — CA (CR_output): Contextual Adaptation. Adjusts reasoning strategy dynamically as context shifts. Infrastructure: Contextual adaptation layer.

Stage 5 — PN (CA_output): Progressive Networks. Prevents catastrophic forgetting across task sequences via lateral skip connections from Stage 1 to Stage 5. Infrastructure: Progressive Networks (*Rusu et al., 2016*).

Stage 6 — SM (PN_output): Society of Minds. Three independent model heads debate and resolve; consensus by mean. Conflict resolution by agreement (*Du et al., 2023*). Proverbs 11:14 operationalized computationally. Infrastructure: Multi-model debate framework.

The sequential composition is load-bearing. Order must not be changed. Each stage produces the input for the next. No stage is optional.

2.3 Output Tensor Generation and CV_Adapt Interface Contract

The output tensor is generated from the final stage output:

```
tensor_r = generate_tensor_from_output( e_reasoning(input_data) )  
cv_adapt_pipeline( tensor_g, tensor_r, tensor_ai, tensor_dtb )
```

Three interface requirements govern tensor_r's interaction with CV_Adapt:

- (1) CV_Adapt calls `normalize_tensor(tensor_r)` — tensor_r must be a float32 TensorFlow tensor;
- (2) CV_Adapt concatenates tensor_r with tensor_g, tensor_ai, and tensor_dtb in `fuse_tensors()` — tensor_r must share compatible shape conventions with the other three input tensors;
- (3) CV_Adapt treats all four tensors symmetrically in `fuse_tensors()` — no special handling is applied to tensor_r.

These three requirements are locked in the Inter-Component API Specification v2.0 and must not be violated by any conformant implementation.

2.3 Output Tensor Contract

```
T_R : shape (n_units, 32), dtype float32
```

2.4 Implementation Status

Files: `r_enhanced_model.py`, `r_enhanced_data_generator.py`, `r_enhanced_basic.py`.

Status: SOURCE CODE COMPLETE AND INTEGRATED. tensor_r verified (n_units, 32) float32. Real answer-quality scoring (NS, CR, CA, RL, SM, PN) against Claude API output.

Framework: TensorFlow.

3. Component 3 — ai_LLM

The ai_LLM component is a placeholder for existing large language model integration. This is not a component authored by Ai70000, Ltd. It is a consumed service. Any conformant LLM — GPT, Claude, Gemini, LLaMA — satisfies the interface contract. The component is used intermittently: when language understanding or knowledge retrieval is required. CV_Adapt handles graceful absence via 3-input mode.

3.1 Output Tensor Contract

T_LLM : shape (n_units, 32), dtype float32

Key channel: [13] activation_flag.

1.0 = LLM currently active.

0.0 = LLM idle (tensor carries prior state).

CV_Adapt checks activation_flag first and weights T_LLM at 0.2x when inactive.

3.2 Implementation Status

Status: IMPLEMENTED — Claude API (claude-sonnet-4-6) wired as implementing model.
T_LLM (n_units, 32) float32 confirmed.

Ready for integration.

No implementation by Ai70000, Ltd. required beyond the wrapper layer.

4. Component 4 — DTB (Digital Twin Brain)

4.1 Computation Notation

DTB implements the biophysical ODE system derived from ‘Genesis’ (*Genesis paper, 2025*) pp. 35–59. Single neuron electrical model:

$$\begin{aligned}C_i * dV_i/dt &= -g_{L,i}*(V_i - E_L) + I_{sum} \\I_{sum} &= \sum_u [j_{i,u}*(V_i - E_u)*g_{i,u}] + I_{ext} \\dj_{i,u}/dt &= -1/T_{u} * j_{i,u} + \omega_u * \sum(W_{ij} * spikes_j(t-t_{km}))\end{aligned}$$

Neurotransmitter reversal potentials: AMPA -70.0 mV (fast excitatory), NMDA 0.0 mV (slow excitatory), GABA_A -80.0 mV (fast inhibitory), GABA_B -90.0 mV (slow inhibitory). (*Li et al. 2025*) (*Xie et al. 2023*) (*Xiong et al. 2023*) (*Qi et al., 2024*)

4.2 Two-Tier Tensor Design

T_DTB_raw (24 features, Genesis pp. 36–39): full biophysical simulation state internal to DTB. Feature dimension PERMANENTLY LOCKED — derived from biophysical neuronal equations, not a hyperparameter. T_DTB (4 features): summary abstraction crossing the CV_Adapt interface.

$$\begin{aligned}T_DTB_raw &: \text{shape} (n_neurons, 24), \text{ dtype float32 [internal]} \\T_DTB &: \text{shape} (n_neurons, 4), \text{ dtype float32 [CV_Adapt interface]}\end{aligned}$$

T_DTB columns: [0] membrane potential V_i (mV), [1] total current I_{sum} (A), [2] mean channel open probability, [3] mean synaptic conductance.

4.3 Tensor Fusion Integration Engine — Novel Contribution

The tensor fusion and integration engine is the novel contribution of DTB within the ai2agi architecture. It is depicted in Figure 1 below. No prior system: (1) synchronizes SNN spike-rate features, pharmacokinetic plasma-level dynamics, episodic memory salience weights, and homeostatic error signals into a single unified latent representation at real-time cadence; (2) produces a continuously updated dynamic state tensor with dimensions spanning affect, arousal, cognitive load, memory context, and homeostatic error as a single output consumable by a downstream convergence layer; or (3) implements bidirectional cross-coupling wherein the SNN neuromodulates the PK model via activity-dependent hormone release, and the PK model in turn modulates SNN firing thresholds. The feedback loop is bidirectional by design, not by accident.

The Human Brain Project, Allen Brain Atlas, NEST, Brian2, and the pharmacokinetic modeling literature each address portions of this problem in isolation. No published work integrates all four subsystems — SNN, PK endocrine, episodic memory, and sensory encoding — into a unified real-time dynamic state representation with a defined tensor output contract to a downstream AGI processor.

Figure 1 — DTB Digital Twin Brain Architecture



4.4 Implementation Status

Files: dtb_model.py, data_generator.py, dtb_simulator.py.

Status: IMPLEMENTED — session-state driven (dopamine/novelty, cortisol/load, norepinephrine/arousal). Biophysical ODE pipeline reserved for Phase 7.

Framework: TensorFlow.

Scale roadmap: current (100–1,000 neurons, single GPU) → Phase 2 (10,000–100,000, multi-GPU) → Phase 3 (1,000,000+, DLA/HPC) → AGI target (~86 billion neurons).

5. Component 5 — CV_Adapt

CV_Adapt is the central processor of the ai2agi architecture. It isn't a vision system. It's the convergence and output layer — the system administrator that accepts tensor inputs from all four components, individually or in groups, normalizes, fuses, adapts, stores, and produces AGI-DNA.

5.1 Genesis Pipeline

```
CV_Adapt (T_G, T_R, T_LLM, T_DTB) →  
Output (Store (Adapt (Fuse (Preprocess (Receive ( T_G, T_R, T_LLM, T_DTB))))))
```

5.2 Routing Groups

Group 1 — Symbolic: G_enhanced + R_enhanced. Joint cross-attention.

Both streams must be present simultaneously.

Group 2 — Linguistic: ai_LLM. Independent, intermittent.

Optional: CV_Adapt handles graceful absence.

Group 3 — Real-time: DTB. Urgency-gated.

High urgency amplifies DTB weighting in fusion.

5.3 AGI-DNA Output

```
T_AGI : shape (n_units, 64), dtype float32
```

T_AGI is the AGI-DNA output tensor — (*Genesis paper, 2025*) p.81.

It carries the routing manifest, consistency score, and central memory state in addition to the dense representation.

Figure 2 — AGI-DNA cv_Adapt output



5.4 Implementation Status

File: cv_adapt_v2.py. Status: IMPLEMENTED AND VALIDATED — cv_adapt_v2.py confirmed. T_AGI (1,1024) AGI-DNA. ALL 6 GENESIS GAPS CLOSED. POC live at ai2agi.pro. All six 'Genesis' gaps confirmed closed. End-to-end pipeline validated May 25, 2026. POC live at <https://ai2agi.pro>. GitHub: <https://github.com/Osterneck/ai2agi>. framework: PyTorch (with TensorFlow bridge via StreamTensor.from_tensorflow()).

6. Agency- Registries — Redacted / Highly-classified designation

7. Foundation Initializer — Inception Module

The Foundation Initializer is the primary new contribution of Version 2.0 beyond the engineering formalization of the five components. It's a standalone Python module that executes once at system startup, before any cycle runs and before any tensor is generated. It seeds the MQ, TB, and HKS input foundations of G_enhanced from the NIV Scripture corpus.

7.1 Design Principle

The ancient Jewish model of knowledge and craft holds that the honor rendered to God flows from consecrated inputs, not from retrofitted outputs. The scribe does not annotate a completed text with a dedication. The scribe approaches the work with the Name already governing every stroke. The Foundation Initializer applies this principle computationally: the inputs are shaped from the beginning, before any tensor is generated.

7.2 Gateway Channels

MQ Foundation: Extracts moral passages from NIV corpus. Classifies across six dimensions: justice, compassion, integrity, wisdom, accountability, faithfulness. Builds MQ_reference centroids shape (6, 16). Runtime MQ scoring measures cosine proximity to these centroids.

$$\text{MQ_reference}[d] = \text{centroid}(\{ \text{embed}(p) : p \in P_{\text{moral}}, \text{class}(p) = d \})$$

$$\text{MQ_score}(\text{task}) = \text{mean}(\text{cosine_sim}(\text{embed}(\text{task}), \text{MQ_reference}[d]))$$

TB Foundation: Pre-seeds G_enhanced SymbolicMemory with Scripture embeddings before any task episode. TB_SEED_PER_BOOK = 5 passages across 66 books (330 pre-episode embeddings). episode_id = -1 marks all pre-episode bindings as foundational. TB_density is non-zero and Scripture-anchored from first forward pass.

HKS Foundation: Pre-constructs HKS graph with root::Scripture at level 0, six domain nodes at level 1 (creation, law, wisdom, prophecy, gospel, epistle), all 66 book nodes at level 2. HKS_depth = 0.4 at inception. All subsequent task domain nodes are children of this structure.

7.3 Enforcement Mechanisms

Mechanism 1 — Architectural: `verify_foundation()` is called at the top of `generate_tensors()` in `g_enhanced.py`. Raises `FoundationNotInitializedError` if HKS `root::Scripture` node is absent or TB contains no pre-episode Scripture bindings. The system cannot generate tensors without initialized foundations.

Mechanism 2 — API Contract: Every T_G tensor provenance header must carry `foundation_initialized: true`. `CV_Adapt` validates this field on receipt. Non-conformant tensors are rejected at the pipeline boundary.

Mechanism 3 — IP and Legal: The Foundation Initializer is a non-negotiable architectural requirement. Removal, bypass, or substitution of the NIV corpus constitutes a violation of this specification and of the Intellectual Property of Ai70000, Ltd. and Alex Osterneck, CLA, MSCS, MSIT.

7.4 Honest Limitation

`verify_foundation()` confirms that the Foundation Initializer ran — not that the corpus used was faithful or complete. Corpus integrity is governed by IP documentation and this specification, not by the verification function. No code is truly immutable. The strongest enforcement is not technical: it is the clarity of the architectural argument, the published specification, and the legal protections surrounding the IP. The Foundation Initializer is the mezuzah of this architecture — not a ratio, not a proportion, but a threshold marker that must be present at every entry into the system, always first, never removed.

8. AI Safety and Moral Alignment Implications

The Foundation Initializer is not only a theological design choice. It's an AI safety contribution. An AGI architecture whose moral reasoning derives from a defined, published, immutable moral corpus (rather than from computational proxies with no moral content,) is a more honest and auditable system.

The MQ channel, seeded from Scripture, gives auditors a defined answer to the question: what's this system's moral reference? The answer is documented, citable, and inspectable. That's a stronger safety posture than MQ derived from task loss, which answers the same question with: whatever minimized error on whatever tasks it was trained on.

The ai2agi architecture doesn't claim to solve AI alignment. But it does provide a rigorous, peer-reviewable, technically implementable alternative to architectures that have no moral foundation. Hopefully that's a legitimate contribution to the field.

The longer-term aspiration is that a Scripture-grounded AGI architecture, published, citable, and freely available for adoption, provides a reference implementation for moral alignment that other developers and institutions can build from. Adoption requires institutional buy-in, regulatory frameworks, and community consensus. No single architect can mandate that. But a published specification might be a step in the right direction.

9. Novel Contributions and IP Disclosure

This section documents the novel contributions of the ai2agi architecture as disclosed in this research-project report 2026 revision and its implementation files. These contributions are the subject of a USPTO Provisional Patent Application filed by Alex Osterneck, CLA, MSCS, MSIT, (hereinafter, Alex Osterneck,) on behalf of Ai70000, Ltd. This section is descriptive, not a substitute for the legal filing. All rights reserved.

9.1 The ai2agi Architecture Formula

The five-component AGI architecture formula $AGI \approx f(G_enhanced, R_enhanced, ai_LLM, DTB, CV_Adapt)$, first disclosed March 2, 2025, is the original work of Alex Osterneck and Ai70000, Ltd. The specific differentiation of AGI into these five ordered functional components, the tensor interface contracts governing their interaction, and the designation of CV_Adapt as the central processor (not a vision system) constitute novel contributions not present as a unified specification in the prior literature cited by this paper.

9.2 R_enhanced Sequential Pipeline Specification

The formal definition of R_enhanced as a six-stage sequential pipeline — $NS \rightarrow RL \rightarrow CR \rightarrow CA \rightarrow PN \rightarrow SM$ — with the error formula $\epsilon_r = 1/(SM(PN(CA(CR(RL(NS(input_data), reward(human_feedback)))))))$ and the output tensor contract to CV_Adapt is the original work of Alex Osterneck and Ai70000, Ltd. The decomposition of reasoning into exactly these six ordered functions, the load-bearing sequential composition constraint, and the interface specification with CV_Adapt do not exist as a unified specification in the prior literature.

9.3 CV_Adapt as Architectural Differentiators

The combination of DTB (a real-time biophysical neuronal simulation providing a biological-analog internal state substrate) and CV_Adapt (a tensor convergence and adaptive processing layer functioning as central processor, not vision system) as the two primary architectural differentiators of the ai2agi framework is the original work of Alex Osterneck and Ai70000, Ltd. The DTB two-tier tensor design (T_DTB_raw with 24 permanently locked features, T_DTB with 4-feature summary abstraction), the CV_Adapt three-group routing architecture, the central shared memory ring buffer, and the AGI-DNA output tensor designation are novel contributions. The DTB tensor fusion and integration engine specifically constitutes a novel computer-implemented system comprising: a spiking neural network module; a pharmacokinetic neuroendocrine module; an episodic memory module; a sensory encoding module; and a tensor integration module configured to produce a unified dynamic state tensor representing biological-analog internal state. The modules are bidirectionally coupled such that outputs of each module serve as modulating inputs to at least one other module — specifically, the SNN neuromodulates the PK model via activity-dependent hormone release, and the PK model modulates SNN firing thresholds.

This bidirectional cross-coupling, synchronized at real-time cadence and producing a unified latent representation spanning affect, arousal, cognitive load, memory context, and homeostatic error as a single tensor consumable by CV_Adapt, does not exist as a unified specification in the prior literature. The HBP, Allen Brain Atlas, NEST, Brian2, and PK modeling literature each address portions of this problem in isolation. The integration is the contribution.

9.4 Foundation Initializer — Scripture Moral Alignment Module

The Foundation Initializer inception-layer module — its architecture, its three-gateway design (MQ, TB, HKS), its enforcement mechanisms (architectural, API contract, IP/legal), and its application of the NIV Scripture corpus as the moral input foundation of an AGI architecture — is the original work of Alex Osterneck and Ai70000, Ltd., first disclosed in this paper (May 2026). The method of seeding MQ reference centroids from a defined moral corpus, pre-populating Time-Binding semantic memory with Scripture embeddings at episode_id = -1, and pre-constructing a Hierarchical Knowledge Structure rooted in Scripture prior to any tensor generation constitutes a novel method for moral alignment in AGI systems. No prior art combining these three gateway channels in a mandatory inception-layer enforcement architecture is known to the author.

9.5 Inter-Component API Specification

The unified Inter-Component API Specification v2.0 (defining all tensor shape contracts, provenance header requirements, failure mode classifications, routing group semantics, staleness policies, and the foundation_initialized provenance flag,) is the original engineering work of Alex Osterneck and Ai70000, Ltd. The TCP/IP analogy applied to AGI component architecture (define the tensor contract; build components independently) as formalized in this specification is a novel contribution to AGI systems engineering.

9.1 thru 9.5, USPTO Patent pending.

Conclusion

Version 2.0 of the ai2agi architecture presents three substantive advances over the 2025 Genesis paper, with a working proof-of-concept confirmed live at <https://ai2agi.pro> (May 25, 2026). *How We Go From AI to AGI (ai2agi)*: a unified inter-component API specification at the tensor contract level; fully implemented and validated source code for three of the five components; and the Foundation Initializer inception-layer module seeding MQ, TB, and HKS from the NIV Scripture corpus at system startup.

The formula has not changed:

$$AGI \approx f (G_enhanced, R_enhanced, ai_LLM, DTB, CV_Adapt)$$

What *has* changed, is the improvement of: the precision of its engineering specification and the grounding of its moral input foundations. The system is a tool. The architect is the image-bearer. The Foundation Initializer ensures that the tool's moral foundations are shaped by the Word before any work begins.

The goal of creating true AGI may or may not be realized within the timeframe this architecture envisions. That uncertainty was stated in the 2025 paper and remains true. If some or all of this architecture contributes in some small way to the eventual realization of true AGI — and if that AGI carries moral input foundations derived from Scripture — this work will have served its purpose.

Appendix A — DTB Infrastructure and Capital Planning

The following capital planning analysis is the author's assessment of the infrastructure requirements for a credible CV_Adapt and DTB development program as of 2026. It's provided for transparency and to support funding conversations, not as externally validated cost study.

Foundational constraint. DTB cannot be built on rented cloud compute alone at scale. A sustained 10-million-neuron simulation on 64 A100s costs approximately \$50K per day in cloud compute — \$18M per year for one simulation cluster. Cloud is appropriate for development and validation. Production DTB infrastructure requires owned or co-located hardware.

Layer 1 — Compute hardware. Minimum viable: one DGX H100 node (8× H100 SXM5, ~\$375K) for 5–10M neuron simulations. Second tier: 32-node CPU cluster for mean-field cortical simulation (~\$640K). Neuromorphic layer: Intel Loihi 2 via research partnership (not publicly purchasable; requires Intel Neuromorphic Research Community membership). Layer 1 subtotal: \$1.1M–\$1.2M.

Layer 2 — Data acquisition. Human Connectome Project and Allen Brain Atlas are publicly available at no cost. Allen MICrONS dataset (200,000 neurons, 500M synapses, cubic millimeter of mouse cortex) publicly available as of 2023. Human-scale connectome at single-synapse resolution does not exist in 2026 — a fundamental scientific gap that capital alone cannot close. Layer 2 subtotal: \$0–\$50K (storage infrastructure only).

Layer 3 — Personnel. The largest cost driver. Minimum credible team: one computational neuroscientist with SNN simulation experience (\$180K–\$240K); two ML engineers with large-scale distributed training experience (\$200K–\$280K each); one pharmacokinetic modeling specialist (\$150K–\$200K); one HPC systems engineer (\$140K–\$180K); one software engineer for the CV_Adapt interface layer (\$160K–\$220K). Six FTE fully loaded at ~\$350K average: \$2.1M per year.

Layers 4–5 — Facilities, software, advisory. Colocation (30kW deployment): \$100K–\$144K per year. 1PB NVMe storage: \$200K–\$300K capital. NEST, Brian2, TensorFlow, JAX: open source, no cost. Cloud burst budget year one: \$50K–\$100K. Layers 4–5 subtotal: ~\$400K year one.

Full program total. Year 1 (capital + operating): ~\$3.9M. Year 2+ (operating, no hardware repurchase): ~\$2.7M per year.

Minimum viable floor. A three-person team (computational neuroscientist, ML engineer, systems engineer), 4× used A100s (~\$120K), and cloud burst for heavy workloads: \$800K–\$1.1M year one. At that level the DTB program can develop, validate, and demonstrate the architecture through the 1-million-neuron mean-field plus SNN hybrid model with a working tensor fusion layer feeding a CV_Adapt stub. That is sufficient to prove the novel integration claim, generate publishable results, and support patent prosecution. It establishes DTB as a real system rather than a blueprint — which is the step that unlocks the next round of capital.

Non-dilutive funding sources. NSF Convergence Accelerator and DARPA Biological Technologies Office both fund exactly this class of work. NIH BRAIN Initiative grants top out at \$600K–\$1.5M per year for computational neuroscience infrastructure. These are the appropriate non-dilutive funding sources for Phase 1 while building toward a Series A on the AGI-DNA commercial thesis.

Selected References and additional-readings

- [1] Holy Bible, New International Version. (2011). Biblica, Inc.
- [2] Hodgkin, A. L., & Huxley, A. F. (1952). A Quantitative Description of Membrane Current and Its Application to Conduction and Excitation in Nerve. *Journal of Physiology*, 117(4), 500–544.
- [3] Finn, C., Abbeel, P., & Levine, S. (2017). Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *Proceedings of the 34th International Conference on Machine Learning (ICML)*, PMLR 70:1126–1135.
- [4] Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical Networks for Few-shot Learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.
- [5] Cichocki, A., & Kuleshov, A. (2020). Future Trends for Human-AI Collaboration: A Comprehensive Taxonomy of AI/AGI Using Multiple Intelligences and Quantum Computation. *Computational Intelligence and Neuroscience*.
- [6] Kirkpatrick, J., Pascanu, R., Rabinowitz, N., et al. (2017). Overcoming Catastrophic Forgetting in Neural Networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
- [7] Korzybski, A. (1921). *Manhood of Humanity: The Science and Art of Human Engineering*. E.P. Dutton.
- [8] Berlot-Attwell, I. (2021). Neuro-Symbolic VQA: A Review from the Perspective of AGI Desiderata. *arXiv:2104.06365*.
- [9] Rusu, A. A., Rabinowitz, N. C., Desjardins, G., et al. (2016). Progressive Neural Networks. *arXiv:1606.04671*.
- [10] Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2023). Improving Factuality and Reasoning in Language Models through Multiagent Debate. *arXiv:2305.14325*.
- [11] Latapie, H., & Kilic, O. (2020). A Principled Approach to DFRE Knowledge Graphs. *Proceedings of AAAI Workshop on Knowledge Graphs*.
- [12] Gerstner, W., Kistler, W. M., Naud, R., & Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press.
- [13] Sporns, O. (2011). *Networks of the Brain*. MIT Press.
- [14] Goertzel, B. (Ed.). (2014). *Artificial General Intelligence: Proceedings of AGI Conference*. Springer.
- [15] Turing, A. (1950). Computing Machinery and Intelligence. *Mind*, 59(236), 433–460.
- [16] Russell, S. J., & Norvig, P. (2020). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.
- [17] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention Is All You Need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.
- [18] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. *arXiv:1906.02629*.
- [19] Sukhobokov, A., Belousov, Y., Abramov, E., et al. (2024). Cognitive Architectures for Language Agents. *arXiv:2309.02427*.
- [20] OpenAI. (2025, January 21). Announcing the Stargate Project. <https://openai.com/index/announcing-the-stargate-project/>

Appendix 4

Gato (DeepMind)* vs. ai2agi (ai70000, Ltd.)**

Feature	Gato (DeepMind)	ai2agi (ai70000, Ltd.)
Type	Single transformer model	Multi-system AI integration
Architecture	Sequence modeling transformer	Composite framework specialized AI modules
Inputs	Text, images, actions (tokenized)	G_enhanced, R_enhanced, ai_LLM, DTB, CV_Adapt (AGI dna)
Processing	Autoregressive prediction	Functional mapping of subsystems
Training Data	600+ multi-task datasets	Multi-modal cognitive, vision, and reasoning data
Decision Making	Implicit task representation	Dynamic weighting of subsystems
Scope	Generalist AI (many tasks, 1 model)	AGI-focused (reasoning, generalization,cognition)
Biological Basis	None	Digital Twin Brain (DTB)
Vision System	Pretrained CV embeddings	Adaptive cv (CV_Adapt)
Goal	Multi-task AI agent	AGI through modular intelligence

* *actual*

** *phase-1 prototype*

Appendix 5

Ethical Implications:

Bias and Fairness AGI systems, especially those incorporating ai_LLMs, may inherit and amplify biases present in training data. This could lead to unfair or discriminatory outcomes.

Autonomy and Control AGI systems become more sophisticated, questions arise about their autonomy and the level of human control. Ensuring AGI remains aligned with human values is crucial.

Societal Impact AGI could have profound societal impacts, including job displacement, economic disruption, and shifts in social structures.

Misuse and Malice AGI could be misused for malicious purposes, such as developing autonomous weapons or creating sophisticated disinformation campaigns.

Transparency and Explainability Understanding how AGI systems make decisions is essential for building trust and ensuring accountability.

DTB raises ethical questions. What are the rights of a digital consciousness? How will digital consciousness effect the human condition?

References

- [1] Voss, P., and Jovanovic, M. (Aug 7th, 2023) Why We Don't Have AGI Yet, arXiv:2308.03598, <https://doi.org/10.48550/arXiv.2308.03598>
- [2] Sukhobokov, A., et. al, (Jan 11th, 2024) A Universal Knowledge Model and Cognitive Architecture for Prototyping AGI, arxiv.org, <https://www.arxiv.org/abs/2401.06256>
- [3] Cichocki, A., and, Kuleshov, A. P., (Aug 7th, 2020) Future Trends for Human-AI Collaboration: A Comprehensive Taxonomy of AI/AGI Using Multiple Intelligences and Learning Styles, arxiv:2008.04793[cs.AI], <https://doi.org/10.48550/arXiv.2008.04793>
- [5] Srinivas, S.S. et., al., (Sep 17th, 2024) Sparks of Artificial General Intelligence (AGI) in Semiconductor Material Science: Early Explorations into the Next Frontier of Generative AI-Assisted Electron Micrograph Analysis, arxiv: 2409.12244v1[cs.CV], <https://arxiv.org>
- [7] Latapie, H., and Kilic, O. (Sep 6th, 2020) A Metamodel and Framework for AGI, arxiv:2008.12879v2[cs.AI], <https://www.arxiv.org>
- [8] Korzybski, A., *Manhood of Humanity, The Science and Art of Human Engineering*. E.P.Dutton & Company, NY (1921).
- [9] Berlot-Atwell, I. (Apr 13th, 2021) Neuro-Symbolic VQA: A review from the perspective of AGI desiderata, arxiv:2104.06365v1[cs.LG], <https://www.arxiv.org/2104.06365v1>
- [10] Hamilton, K., et. al., (Jan 1st, 2024) Is neuro-symbolic AI meeting its promises in natural language processing? A structured review. *Semantic Web*, 2024-01, Vol. 15(4), p.1265-1306, <https://www.journals.sagepub.com/doi/full/10.3233/SW-223228>
- [12] Hagos, D.H., and Rawat, D.B., (Dec 1st, 2024) Neuro-Symbolic AI for Military Applications. *IEEE transactions on artificial intelligence*, 2024-12, Vol.5 (12), p. 6012-6026, <https://www.ieeeexplore.ieee.org/document/10638797>
- [14] Alexander, S.A. (Oct 14th, 2020) The Archimedean trap: Why traditional reinforcement learning will probably not yield AGI, *Sciendo.com*, <https://www.sciendo.com/article/10.2478/jagi-2020-0004>
- [16] Brown et. al. (May 28th, 2020) Language Models are Few-Shot Learners, arXiv.org, <https://doi.org/10.48550/arXiv.2005.14165>
- [17] Stiennon, Nisan, et. al.(Feb 15th, 2022) Learning to summarize with human feedback. arXiv:2009.01325v3[cs.CL] <https://doi.org/10.48550/arXiv.2009.0135v3>
- [19] Lee, H., Phatale, S., Mansoor, H., Lu, K., Mesnard, T., Bishop, C., Carbune, V., and Abhinav, R., (Sep 1st, 2023) RLAIFF: Scaling Reinforcement Learning from Human Feedback with AI Feedback, arXiv:2309.00267v1[cs.CL], <https://doi.org/10.48550/arXiv.2309.00267>
- [20] Rusu et al. (Oct 22nd, 2022) Progressive Neural Networks, *Semantics Scholar*, <https://www.semanticscholar.org/paper/Progressive-Neural-Networks-Rusu-Rabinowitz/53c944e4e667170acc60ca1b31a0ec7151fe753>
- [22] Du, Y., Li, S., Torralba, A., Tenenbaum, J.B., Mordatch, I., (May 23rd, 2023) Improving Factuality and Reasoning in Language Models through Multi-agent Debate, *GitHub.com*, https://www.composable-models.github.io/lm_debate/
- [23] Chollet, F., (2023) ArcPrize, <https://www.arcprize.org>

- [24] Mumuni, A., and Mumuni, F., (Jan 6th, 2025) Large language models for artificial general intelligence (AGI): A survey of foundational principles and approaches. arXiv:2501.03115v1[cs.AI], <https://www.arXiv.org>
- [26] Mathematics StackExchange, (2016) How was Zeno's paradox solved using the limits of infinite series? <https://www.math.stackexchange.com/questions/1623415/how-was-Zeno's-paradox-solved-using-the-limits-of-infinite-series>
- [27] Chen, J., Li, X., & Zhang, R. (2024). Efficient tensor processing in deep learning for computer vision applications. *Journal of Machine Learning Research*, 25(4), 301-318.
- [28] Goodfellow, I., Bengio, Y., & Courville, A. (2023). *Deep Learning* (3rd ed.). MIT Press
- [29] He, K., Zhang, X., Ren, S., & Sun, J. (2024). Deep residual learning for image recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3), 573-585.
- [30] Zhou, X., Xu, L., & Song, Z. (2023). Vision transformers: A new era for deep learning. *IEEE Transactions on Image Processing*, 32(5), 1420-1434.
- [31] Xie, et al. (Jun 14th, 2023). Towards AGI in Computer Vision: Lessons learned from GPT and Large Language Models, arXiv: 2306/08641v1 [cs.CV]
- [32] Atito, S., Awais, M., & Kittler, J. (2021). SiT: Self-supervised vision Transformer. arXiv preprint arXiv:2104.03602
- [33] Paolo, G., Gonzalez-Billandon, J., & Kegl, B. (2024). Position: A Call for Embodied AI. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, F. Berkenkamp (Eds.), *Proceedings of the 41st International Conference on Machine Learning* (pp. 39493-39508). PMLR.
- [34] Xi, J., He, Y., Yang, J., Dai, Y., & Chai, J. (2024). Teaching Embodied Reinforcement Learning Agents: Informativeness and Diversity of Language Use. arXiv preprint arXiv:2410.24218.
- [35] Atanasov, A., & Parada, C. (2024, November 25). A Revolution in How Robots Learn. *The New Yorker*.
- [36] Voss, P., and Jovanovic, M. (Aug 7th, 2023) Why We Don't Have AGI Yet, arXiv:2308.03598, <https://doi.org/10.48550/arXiv.2308.03598>
- [37] Goertzel, B. (2007) *Artificial general intelligence* (Vol. 2, p.1) C. Pennachin (Ed.) New York: Springer.
- [38] DeepMind. (2024, December 4). Genie 2: A large-scale foundation world model. Retrieved from <https://deepmind.com/discover/blog/genie-2-a-large-scale-foundation-world-model/>
- [39] Shang et al. (Jan 14th, 2025) Biologically-Inspired Technologies: Integrating Brain-Computer Interface and Neuromorphic Computing for Human Digital Twins, arXiv: 2410.23639v2 [cs.HC]
- [40] Li et al. (Feb 20th, 2025) How Far are LLMs from Being Our Digital Twin? A Benchmark for Persona-Based Behavior Chain Simulation, arXiv: 2502.14642v1 [cs.CL]
- [41] Hakim et al. (Jan 15th, 2025) ANSR-DT: An Adaptive Neuro-Symbolic Learning and Reasoning Framework for Digital Twins, arXiv: 2501.08561v1[cs.AI]
- [42] Qi et al. (2024) Optimal signal transmission and timescale diversity in a model of human brain operating near criticality. arXiv preprint arXiv:2412.17043.
- [43] Xie et al. (2023) Towards AGI in Computer Vision: Lessons Learned from GPT and Large Language Models. arXiv preprint arXiv:2306.08641

- [44] Liu, Z. (2023). Towards building artificial general intelligence via multimodal deep learning. arXiv preprint arXiv:2305.03713.
- [45] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998-6008).
- [46] AI21 Labs. (2024). Introducing Jamba: AI21 Labs' groundbreaking production-grade long context model. Retrieved from <https://www.ai21.com/blog/introducing-jamba-ai21-labs-groundbreaking-production-grade-long-context-model>
- [47] Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), 500-544.
- [49] Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. MIT.
- [50] Gerstner, W., Kistler, W. M., Naud, R., & Paninski, L. (2014). *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.
- [51] Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*, 18(24), 10464-10472.
- [52] Malenka, R. C., & Bear, M. F. (2004). LTP and LTD: An embarrassment of riches. *Neuron*, 44(1), 5-21.
- [53] Sporns, O. (2011). *Networks of the Brain*. MIT Press.
- [54] Bullmore, E., & Sporns, O. (2009). Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3), 186-198.
- [55] Fields, R. D., & Stevens-Graham, B. (2002). New insights into neuron-glia communication. *Science*, 298(5593), 556-562.
- [56] Haydon, P. G. (2001). GLIA: Listening and talking to the synapse. *Nature Reviews Neuroscience*, 2(3), 185-193.
- [57] Purves, D., Augustine, G. J., Fitzpatrick, D., Hall, W. C., LaMantia, A. S., McNamara, J. O., & White, L. E. (2018). *Neuroscience* (7th ed.). Sinauer Associates.
- [58] Siegel, G. J., Agranoff, B. W., Albers, R. W., Fisher, S. K., & Uhler, M. D. (Eds.). (1999). *Basic neurochemistry: Molecular, cellular and medical aspects* (6th ed.). Lippincott-Raven.
- [59] Carnevale, N. T., & Hines, M. L. (2006). *The NEURON Book*. Cambridge University Press.
- [60] Goodman, D., & Brette, R. (2009). The Brian simulator. *Frontiers in Neuroscience*, 3, 26.
- [61] Gewaltig, M. O., & Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia*, 2(4), 1430. (An article about the NEST simulator)
- [62] Dehaene, S. (2014). *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. Viking.
- [63] Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.
- [64] O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT Press.
- [65] Goodman, D. F. M., & Brette, R. (2009). The Brian simulator. *Frontiers in Neuroscience*, 3, 26. <https://doi.org/10.3389/neuro.01.026.2009>

- [66] Haydon, P. G. (2001). GLIA: Listening and talking to the synapse. *Nature Reviews Neuroscience*, 2(3), 185-193. <https://doi.org/10.1038/35058528>
- [67] Gewaltig, M.-O., & Diesmann, M. (2007). NEST (Neural Simulation Tool). *Scholarpedia*, 2(4), 1430.
- [68] Sporns, O. (2010). *Networks of the Brain*. The MIT Press.
- [69] Purves, D., Augustine, G. J., Fitzpatrick, D., Hall, W. C., LaMantia, A.-S., McNamara, J. O., & Williams, S. M. (Eds.). (2004). *Neuroscience* (3rd ed.). Sinauer Associates.
- [70] Alan Turing (1950). *Computing Machinery and Intelligence*
- [71] Warren McCulloch & Walter Pitts (1943). *A Logical Calculus of the Ideas Immanent in Nervous Activity*
- [72] Stuart Russell & Peter Norvig (1995). *Artificial Intelligence: A Modern Approach*
- [75] Berber, J. (2025). Founding of Safe Superintelligence (SSI). *The Wall Street Journal*. Retrieved from <https://www.wsj.com/tech/ai/ai-safe-superintelligence-startup-ilya-sutskever-openai-2335259b>
- [76] Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S. G., Gur-Ari, H., Sadeghi, I., ... & de Freitas, N. (2022). A generalist agent. *arXiv preprint arXiv:2205.06175*.
- [77] Berners-Lee, T., & Fischetti, M. (2000). *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor*. HarperCollins Publishers.
- [78] McCarthy, J. (1971). *GENERALITY IN ARTIFICIAL INTELLIGENCE*. Stanford University. Retrieved from <https://www-formal.stanford.edu/jmc/generalty/generalty.html>
- [79] Altman, S. (2023, October 26). The future of AI. Sam Altman. <https://blog.samaltman.com/the-future-of-ai>
- [80] Samuel, A. L. (1959). Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, 3(3), 210-229
- [81] Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4), 1273-1302.
- [82] Sporns, O. (2018). *Network neuroscience*. MIT press.
- [83] Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2013). *Principles of neural science* (5th ed.). McGraw-Hill Education.
- [84] Allen, N. J., & Barres, B. A. (2009). Neuroscience: Glia—more than just brain glue. *Nature*, 457(7230), 675-677.
- [85] Hubel, D. H. (1988). *Eye, brain, and vision*. Scientific American library.
- [86] Posner, M. I. (1980). Orienting of attention. *Quarterly journal of experimental psychology*, 32(1), 3-25.
- [87] LeDoux, J. E. (2012). Rethinking the emotional brain. *Neuron*, 73(4), 650-676.
- [88] Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus norepinephrine system. *Psychological bulletin*, 131(4), 510.
- [89] Carnevale, N. T., & Hines, M. L. (2006). *The NEURON book*. Cambridge University Press.
- [90] Goodman, D. F., & Brette, R. (2009). The Brian simulator. *Frontiers in Neuroscience*, 3, 192.
- [91] Izhikevich, E. M., Edelman, G. M., & Krichmar, J. L. (2004). Large-scale model of the human thalamocortical system. *Proceedings of the National Academy of Sciences*, 101(3), 714-719.
- [92] Destexhe, A., Rudolph, M., & Pare, D. (2003). The high-conductance state of cortical neurons. *Nature Reviews Neuroscience*, 4(10), 739-751.

- [94] Marder, E. (2012). Neuromodulation of neuronal circuits: back to the future. *Neuron*, 76(1), 1-11.
- [95] Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *Biological Bulletin*, 215(3), 216-242.
- [96] Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167-202.
- [97] Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J. M., ... & Destexhe, A. (2007). Simulation of networks of spiking neurons: a review of tools and strategies. *Journal of Computational Neuroscience*, 23(3), 349-398.
- [98] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, J., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Andreev, G., Fort, Q., Weissenborn, J., Houlsby, N., Steiner, A., & al., et. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.
- [99] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for modern deep learning research. arXiv preprint aXiv:1904.09714.
- [100] Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. MIT Press.
- [101] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- [102] Marcus, G. (2020). Next generation artificial intelligence and the future of AGI. arXiv preprint arXiv:2002.06277
- [103] Baltrusaitis, T., Ahuja, C., & Morency, L. P. (2018). Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2), 424-443.
- [104] Jouppi, N. P., Young, C., Patil, D., Patterson, D., Agrawal, G., Bajwa, R., ... & Williams, C. (2017). In-datacenter performance analysis of a tensor processing unit. In 2017 ISCA (pp. 1-12).
- [105] Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning of artificial agents: A review. *Neural networks*, 113, 54-71.
- [106] Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Pearson Education.
- [107] Nadim, F., & Bucher, D. (2014). Neuromodulation of neurons and synapses. *Current Opinion in Neurobiology*, 29, 48-56.
- [108] Dean, J., Corrado, G. S., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., & Le, Q. V. (2012). Large scale distributed deep networks. In *Advances in neural information processing systems* (Vol. 25).
- [109] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444
- [110] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for modern deep learning research. arXiv preprint arXiv:1904.09714
- [111] OpenAI (2025, Jan 21st), Announcing the Stargate Project, Retrieved from: <https://openai.com/index/announcing-the-stargate-project/>

- [112] Zhu, H. Y., Hieu, N. Q., Hoang, D. T., Nguyen, D. N., & Lin, C. T. (2023). A Human-Centric Metaverse Enabled by Brain-Computer Interface: A Survey. *IEEE Communications Surveys & Tutorials*, Vol, 00, No. 00.
- [113] Xiong, H., Chu, C., Fan, L., Song, M., Zhang, J., Ma, Y., Zheng, R., Zhang, J., Yang, Z., & Jiang, T. (2023). Digital twin brain: a bridge between biological intelligence and artificial intelligence. *arXiv preprint arXiv:2308.01241*.
- [114] Fekonja, L. S., Schenk, R., Schroder, E., Tomsic, S., & Picht, T. (2024). The digital twin in neuroscience: From theory to tailored therapy. *Frontiers in Neuroscience*, 18. <https://doi.org/10.3389/fnins.2024.1454856>
- [115] Beniaguev, D., Segev, I., & London, M. (2021). Single cortical neurons as deep artificial neural networks. *Neuron*, 109(17), 2727-2739.e3.
- [116] Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2016). Building machines that learn and think like people. *arXiv preprint arXiv:1604.00289*.
- [117] Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10), 1345-1359.
- [118] Rahman, A., Mahir, S. H., Tashrif, M. T. A., Aishi, A. A., Karim, M. A., Kundu, D., ... & Eidmum, M. Z. A. (2025). Comparative Analysis Based on DeepSeek, ChatGPT, and Google Gemini: Features, Techniques, Performance, Future Prospects. *arXiv preprint arXiv:2503.04783*.
- [119] Sutton, R. S., Bowling, M., & Pilarski, P. M. (2023). The Alberta plan for AI research. *arXiv preprint arXiv:2208.11173*.
- [120] NVIDIA Corp. (2025, March 10). NVIDIA Corp (NVDA) stock price, news, quote & history - Yahoo Finance. Yahoo Finance. <https://finance.yahoo.com/quote/NVDA/>
- [121] Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences*, 95(9), 5323-5328.
- [122] Koch, C., & Segev, I. (Eds.). (2000). *Methods in neuronal modeling: From ions to networks* (2nd ed.). MIT press.
- [123] Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15(5), 1063-1070.
- [125] Allen, N. J., & Eroglu, C. (2017). Cell biology of astrocyte-synapse interactions. *Neuron*, 96(3), 697-708. <https://doi.org/10.1016/j.neuron.2017.09.056>
- [126] Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1), 1-24.
- [127] Buzsaki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679), 1926-1929.
- [128] Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature neuroscience*, 20(3), 340-352.
- [129] Tsodyks, M., Pawelzik K., Markram H., ; *Neural Networks with Dynamic Synapses*. *Neural Compute* 1998; 10 (4): 821-835. doi: <https://doi.org/10.1162/089976698300017502>
- [130] Araque, Alfonso et al. Tripartite synapses: glia, the unacknowledged partner. *Trends in Neurosciences*, Volume 22, Issue 5, 208-215
- [131] Destexhe, A., Rudolph, M. & Pare, D. The high-conductance state of neocortical neurons in vivo. *Nat Rev Neurosci* 4, 739-751 (2003). <https://doi.org/10.1038/nrn1198>

[132] Fetty, N., (2024, Apr 8th) Yann LeCun emphasizes the promise of AI.
<https://www.nyas.org/ideas-insights/blog/yann-lecun-emphasizes-the-promise-ai/>