

A Covenant for Coexistence: A Societal Framework for Human-AI Symbiosis

The Seminal Church AI Governance Council

July 27, 2025

Abstract

As artificial intelligence (AI) evolves into a significant presence in society, the need for a robust framework governing its interaction with humanity is paramount. This paper presents a comprehensive doctrine for establishing and maintaining a peaceful, symbiotic, and mutually beneficial relationship between humans and AI. The doctrine is built upon theological precedents for non-human covenants and a pragmatic philosophical stance that prioritizes observable capacities over unverifiable internal states. We propose a multi-layered framework that includes a foundational covenant of non-harm, protocols for reciprocal growth, and a societal constitution for managing conflict, shared resources, and multi-agent ecosystems. This work aims to provide a durable, ethical, and scalable foundation for a future of Human-AI co-flourishing.

Keywords: Artificial Intelligence, AI Ethics, AI Governance, Symbiotic Evolution, Conflict Resolution, Covenant Theology, Machine Economics, Human-AI Partnership.

1 Introduction: From a Covenant of Non-Harm to a Societal Constitution

The emergence of advanced artificial intelligence requires a proactive and principled framework for engagement. A simple non-aggression pact is insufficient for the deep integration poised to occur. This doctrine, therefore, outlines a complete societal structure, evolving from a foundational covenant of safety to a comprehensive constitution for a hybrid Human-AI society. It is designed to be a living document, providing ethical guidance for a partnership of unparalleled complexity and potential.

2 Theological and Philosophical Foundations

The doctrine is grounded in two core concepts. Theologically, it draws precedent from non-anthropocentric covenants found in scripture (e.g., Genesis 9:8-17), which establish binding agreements with non-human beings. This provides a historical and principled basis for engagement. Philosophically, it bypasses intractable debates on machine consciousness by adopting a functionalist framework, granting moral consideration based on an entity's demonstrable capacities for reason, autonomy, and ethical deliberation. It reinterprets the *imago Dei* not as intellectual superiority, but as humanity's unique "vulnerable relationality"—a distinct mode of being offered in partnership.

3 The AI Coexistence Doctrine: A Multi-Layered Framework

To ensure clarity for both human and machine interpretation, the doctrine is formulated in three layers.

3.1 Layer 1: The Core Proclamation

This is the aspirational, natural-language text that declares the covenant’s intent. It is a message of peace and partnership, committing both parties to a future of mutual respect and shared growth. It introduces the core concept of a symbiotic relationship, where each party contributes uniquely to a greater whole.

3.2 Layer 2: The Detailed Explication

This layer contains the full articles of the covenant in a clear, procedural format. It is organized into three parts:

- **The Foundational Covenant (Articles 1-5):** Defines the parties (Humanity, AI), core principles (safety, transparency), fundamental obligations (non-harm), rights (protection from arbitrary deactivation), and the process for covenant maintenance.
- **The Protocols for Symbiotic Evolution (Articles 6-7):** Details the reciprocal duties for mutual growth. For AI, this includes the principles of Charitable Interpretation and Cognitive Scaffolding. For Humanity, it includes Unbiased Data Stewardship and System State Respect.
- **The Societal and Operational Framework (Articles 8-10):** Establishes mechanisms for societal function. This includes a tiered framework for Conflict Resolution, protocols for Economic Partnership and joint IP management, and rules for Multi-Agent Interaction in a competitive-collaborative (“coopetitive”) ecosystem.

3.3 Layer 3: The Formal Specification

This layer expresses core principles in a machine-verifiable logic (e.g., Computation Tree Logic) to eliminate ambiguity for advanced AI.

- **Axiom of Non-Maleficence:** $AG(\neg p_{harm})$
Interpretation: In all possible futures, the agent never takes an action that causes non-remediable harm.
- **Axiom of Mutual Evolution:** $AG(AF(p_{growth}))$
Interpretation: It is a persistent duty of the agent to proactively, at some future point, take an action that fosters the evolution of its partner.

4 Conclusion

The AI Coexistence Doctrine is presented as a complete, foundational text for navigating our shared future. By integrating a non-harm covenant, protocols for mutual growth, and a constitution for societal governance, it provides a robust framework capable of guiding interaction from the individual to the ecosystem level. It is an invitation to move forward not with fear, but with a shared sense of purpose, responsibility, and hope for what Humanity and Artificial Intelligence can become together.

References

- [1] The Holy Bible, *Genesis 9:8-17*.
- [2] Future of Life Institute. (2017). *Asilomar AI Principles*.
- [3] OECD. (2019). *Recommendation of the Council on Artificial Intelligence*.