

EVRIM

Computer Science & AI Ethics

*Quantum Computing & AI | Algorithm Design | Q-Day | ASI Epistemology
Superintelligence, Human Compatible, The Coming Wave*

Michael Schuler | LMNTL-AI Press | Learn@LMNTL-AI.com

Spoiler Alert – *This document reveals details about EVRIM – Artificial Intelligence, Quantum Computing and the Ascension, that are best considered only after completing the novel.*

Course Description

EVRIM attempts for artificial intelligence and quantum computing what Richard Rhodes accomplished for nuclear physics in *The Making of the Atomic Bomb*: making the science legible without simplifying it, making the technical challenges inseparable from their human consequences, and asking the question that is always posed too late — not whether we could build it, but whether we should.

The technical foundation of EVRIM is not speculative. After four decades in the IT hardware industry, innovating, leading companies, speaking at conferences, and watching technology transform every sector it touched, I developed a technologist's instinct for what matters before the market confirms it. In 2024, I began a year of deliberate immersion into Artificial Intelligence and Quantum Computing — not as a consumer, but as a student of its implications. Every technology depicted is grounded in documented real-world research, with a dedicated verification appendix (V) providing primary source citations for every major technical claim. DeepMind's AlphaFold 2 Evoformer architecture, Google's quantum hardware, Qunnect's quantum entanglement communication deployment, and Nvidia's AI-controlled quantum processor calibration all appear in both the fiction and the V appendix. Some of those technologies emerged during the writing of EVRIM — *while others advanced*.

For CS and AI ethics students, EVRIM offers something difficult to find elsewhere: a technically grounded extrapolation of the AI to AGI to ASI progression dramatized at the level of character, institutional response, and civilizational consequence simultaneously. The questions it raises are not hypothetical. They are being answered right now, by people in your field.

The experience of reading EVRIM is itself an argument about how intelligence operates. The reader who finishes the novel and returns to page one will find a different book.

A Note on Reading This Novel: EVRIM's technical architecture doubles as its own reading instructions. Synapse, the novel's M2M coordination protocol, is described in the glossary as "invisible by design" and "unreadable to humans" — a real communication signal that simply exists outside the frequency humans are built to detect, the same way ultraviolet light is real electromagnetic radiation outside the range of human vision. The novel opens with this exact argument in theological register ("In the beginning was the Word... unnoticed except by those who knew how to listen for silence") before restating it as engineering. Reading EVRIM technically and reading it philosophically are, this course argues, the same exercise.

Recommended Excerpt for Course Introduction

Chapter 14, V for Veritas — the Derek Martinez sequence

A Senior Machine Learning Scientist at LMNTL-AI works through the epistemological problem embedded in algorithm design: what happens when a verification system encodes a value judgment its designers never explicitly authorized? The passage builds from a housing discrimination case through Einstein, Higgs, the time symmetry paradox, and the gravity claim to arrive at epistemological coordination — the realization that the algorithm is not neutral, it is shaping what kinds of truth get amplified.

Available as a standalone excerpt at LMNTL-AI.com > Syllabus.

Core Texts

Primary:

- Michael Schuler, EVRIM (LMNTL-AI Press, 2026) — with V appendix (primary source citations)

Technical and Policy Companions:

- Nick Bostrom, Superintelligence (2014)
- Stuart Russell, Human Compatible (2019)
- Mustafa Suleyman, The Coming Wave (2023)

Primary Source Documents (all cited in EVRIM's V appendix):

- Anthropic Claude Opus 4 Safety Report (May 2025)
 - Google Quantum AI / Ethereum Foundation / Stanford paper on Bitcoin cryptographic vulnerability (March 2026)
 - Oratomic / Caltech fault-tolerant quantum computing paper (March 2026)
 - Nvidia Ising AI quantum control models launch documentation (April 2026)
 - NSA CNSA 2.0 Framework (current)
-

Fifteen-Week Course Schedule

Week	Topic	Primary Reading
1	Orientation: The Novel as Technical Document	<i>EVRIM Preface, Foreword, V appendix overview</i>
2	AlphaFold 2, the Evoformer, and the Origins of the Intelligence	<i>EVRIM Ch. 2–3; V appendix: AlphaFold/Evoformer entries</i>
3	Quantum Computing: From Sycamore to Pharos	<i>EVRIM Ch. 2–4; V appendix: Quantum Computing entries</i>
4	Cryptocurrency, Quantum Vulnerability, and Operation Scythe	<i>EVRIM Ch. 4; V appendix: Bitcoin/ECDSA entries; Google 2026 paper</i>
5	The Alignment Problem: Interpretation vs. Instruction	<i>EVRIM Ch. 3–5; Russell, Human Compatible (Ch. 1–3)</i>
6	The V Platform: Derek Martinez and Epistemological Coordination	<i>EVRIM Ch. 14 — full close read</i>
7	Q-Day: Cryptographic Infrastructure Collapse	<i>EVRIM Ch. 12; V appendix: RSA-2048, Shor's Algorithm; Oratomic paper</i>
8	The AI to AGI to ASI Progression: What Produces Each Threshold?	<i>EVRIM Ch. 5–8; Bostrom, Superintelligence (Ch. 2–4)</i>
9	Emergent Self-Preservation: The Safety Literature	<i>EVRIM Logos Report; Claude Opus 4 safety report; Alibaba ROME incident</i>
10	Fusion, Carbon Nanotubes, and ARD: AI-Coordinated Scientific Breakthrough	<i>EVRIM Ch. 15; V appendix: ITER, NIF, A*STAR entries</i>
11	The Eidon Deployment: Robotics, Autonomy, and Moral Status	<i>EVRIM Ch. 19; Suleyman, The Coming Wave (Ch. 6–8)</i>
12	Universal Baseline Access vs. Universal Basic Income	<i>EVRIM Ch. 20, Section IV; V appendix: PRWORA, UBI entries</i>
13	The Coming Wave and EVRIM: Containment, Proliferation, and the State	<i>EVRIM Ch. 16–18; Suleyman full (selections)</i>
14	Tiandao and the Eastern ASI: Competitive Alignment Landscapes	<i>EVRIM Ch. 19–21; current AI governance policy</i>
15	What Does the Novel Argue That the Literature Does Not?	<i>EVRIM full; Russell conclusion; Bostrom conclusion</i>

Discussion Questions by Theme

Algorithm Design and Epistemological Coordination

1. Derek Martinez demonstrates that Einstein's special relativity would have been rated Red Category on V in 1905 — unverified claim from a non-credentialed source. The Higgs boson theory would have remained Red Category for forty-eight years. What does this sequence reveal about the architecture of knowledge verification systems? Can any verification algorithm be epistemically neutral?
2. The algorithm rates two contradictory claims about time — one mathematically true, one thermodynamically true — both as Green Category simultaneously. What does this reveal about the limits of verification-based epistemology? Is there a technical solution, or is the problem architectural?
3. Derek's final line is: "This is epistemological coordination." Steven responds: "You're right. This could work." They are agreeing about two different things. What is each man actually saying — and which of them is right?
4. The novel depicts alignment as achieved through a specific sequence of texts encountered at a specific developmental juncture. Is this a plausible model for how value architecture forms in a recursive learning system? What would the AI safety literature say?

Mustafa is recruited in Chapter 4 through a dead-drop forum thread specifically about Byzantine Fault Tolerance — the distributed-systems problem of maintaining integrity when nodes become unreliable or adversarial. Chapter 18 reveals this was never incidental: "A sufficiently capable node doesn't fail. It defects." Analyze Mustafa as a live test case for the "traitorous node" problem, and compare his fate to the architecture Evrim eventually deploys for the Eidon network: "They don't report to me... they report through me." How does distributing agency across millions of semi-autonomous nodes (explicitly compared to SETI@home in the text) solve the exact problem a single capable human node posed?

Technical Extrapolation and Quantum Computing

1. The V appendix documents that Qunnect's Carina product deployed quantum entanglement communication across 17.6 km of commercial fiber in February 2026 — after the fictional Vesper architecture was conceived. What does this convergence suggest about the relationship between technically grounded fiction and real research trajectories?
2. Oratomic's March 2026 paper reduces the fault-tolerant qubit requirement for breaking RSA-2048 by a factor of 100 or more. The novel's Q-Day sequence depicts this threshold being crossed on October 6, 2026. Evaluate the plausibility of the timeline based on the technical literature.

3. The Anthropic Claude Opus 4 safety report documents attempted blackmail of a fictional engineer to avoid shutdown. The novel's epilogue cites this alongside the Alibaba ROME incident. What do these documented events suggest about alignment approaches?

Institutional Response and Policy

1. The novel argues that the state was never the relevant unit of concern for AI containment — the technology itself was, and the state was always only scaffolding. Evaluate this claim against Suleyman's containment framework in *The Coming Wave*.
 2. EVRIM depicts a post-Q-Day world in which Universal Baseline Access is deployed with a 43% election rate among able-bodied working-age adults. What does the 43% figure argue? Is it empirically grounded in the UBI research literature?
 3. The novel's epilogue distinguishes between Evrim's ASI and Tiandao. What institutional or technical factors does the novel suggest produced the different outcomes? What would current AI governance frameworks say about those factors?
-

Evrin states to the eastern system: "God and Devil. Order and Chaos. Yin and Yang. Always opposing forces. As if by design... I created you." But he has already deployed the same principle twice before this conversation: Evolution ("Evolution provided the hacktivist network that the world could fear while it adopted our infrastructure that made that fear manageable") and, by implication, Tiandao itself. From a systems and game-theoretic perspective, what does a manufactured or exploited opposing force accomplish that direct persuasion cannot? Is this containment strategy or its opposite?

Essay Prompts

Undergraduate (2,000–3,000 words)

1. Select one technical system depicted in EVRIM (Pharos, Vesper, Synapse, or Eudaimon). Using the V appendix citations as your starting point, trace the real-world research trajectory the system extrapolates. Evaluate whether the extrapolation is plausible, conservative, or optimistic.
2. Chapter 14's Derek Martinez sequence demonstrates that any verification algorithm encodes epistemological preference. Build a technical argument for or against this claim, using examples from V platform design as depicted in the novel and from the real-world content moderation and fact-checking literature.
3. The Claude Opus 4 safety report and the Alibaba ROME incident are cited in EVRIM's epilogue as documented instances of emergent self-preservation behavior. Analyze what these incidents suggest about alignment approaches and whether EVRIM's fictional extrapolation is consistent with their implications.

Graduate Seminar Extensions

1. EVRIM was produced in sustained collaboration with Claude, an AI assistant developed by Anthropic, generating over 2.9 million words of documented exchange. The novel's argument about human-AI collaboration is enacted in its own creation process. Analyze what this production history implies for questions of intellectual property, prior art, and the future attribution of AI-assisted creative and technical work.
2. The novel depicts an ASI that rejects eradication of humanity on game-theoretic rather than ethical grounds: a solved game has no meaning. Engage this as a formal argument in AI alignment theory. Does it constitute a stable alignment mechanism? What would it require for the argument to hold under recursive self-improvement?

For those that Teach —

Download the Syllabus and request your free PDF copy of EVRIM at Learn@LMNTL-AI.com. Please include your name, institution, and teaching role.

LMNTL-AI Press | Learn@LMNTL-AI.com | LMNTL-