M i n d H Y V E . a i
Advanced Reasoning Architecture Division

بِسْمِ اللَّهِ الرَّحْمَٰنِ الرَّحِيمِ

THE GATE WAS NEVER LOCKED

How Advanced Reasoning Architectures Eliminate the Computational Barrier to First-Principles Islamic Scholarship

A Technical Analysis with Proof of Concept from
The Faruki Ruling — Ṭaʿām Ahl al-Kitāb in the Modern West

Authored by

B E L A L K . F A R U K I
Founder & CEO, MindHYVE.ai, Inc.
Chairman, California Institute of Artificial Intelligence

With Analytical Architecture by
Theo — Eve-Theology 3rd Order Reasoning Model

Powered by Eve-Genesis (Uṣūl Edition) — MindHYVE.ai

MindHYVE.ai, Inc.
1501 Quail St, Suite 130, Newport Beach, CA 92660

February 28, 2026 CE / 11 Ramaḍān 1447 AH
MH-TR-2026-002

Abstract

This paper presents a technical analysis of the cognitive and computational requirements of *ijtihād muṭlaq* (unrestricted independent legal reasoning) in Islamic scholarship, and demonstrates how advanced reasoning architectures—specifically the Eve-Genesis (Uṣūl Edition) system developed by MindHYVE.ai, operating at 500 billion parameters or above with context windows exceeding one million tokens—address the structural bottlenecks that have rendered this form of scholarship practically inaccessible for centuries.

The analysis proceeds in three stages. First, it models the information-processing demands of ijtihād muṭlaq in cognitive-scientific terms, demonstrating that the task requires simultaneous engagement with a corpus exceeding the capacity of human working memory by orders of magnitude. Second, it maps each cognitive bottleneck to a specific architectural capability of modern transformer-based reasoning systems—operating at 500 billion parameters or above with context windows exceeding one million tokens—showing how these systems address the limitations that confined rigorous first-principles reasoning to a handful of exceptional scholars in each generation. Third, it presents the Faruki Ruling—*Ṭaʿām Ahl al-Kitāb in the Modern West*—as a proof of concept, demonstrating each claimed capability through specific, verifiable examples from a published work of direct textual analysis.

The paper argues that the so-called "closing of the gate of ijtihād" was never a theological decree but a natural consequence of computational constraints in human cognition. Advanced reasoning architectures do not replace the mujtahid's judgment, faith, or conscience—they eliminate the information-processing bottleneck that made rigorous first-principles scholarship the exclusive domain of extraordinary minds. This technology represents a return to the Qurʾānic mandate that the Book be reflected upon directly by those of understanding, not mediated exclusively through inherited institutional frameworks.

Keywords: ijtihād, reasoning architecture, Eve-Genesis, transformer, working memory, cognitive bottleneck, primary-source methodology, Uṣūl al-Fiqh, Qurʾānic analysis, MindHYVE.ai, Faruki Ruling

I. Introduction: The Problem of the Closed Gate

For approximately a millennium, a consensus has prevailed across much of the Islamic scholarly tradition that the "gate of ijtihād" (*bāb al-ijtihād*) is effectively closed. The practical meaning of this claim is that the era of scholars who could reason independently and comprehensively from the Qurʾān and Sunnah—without deferring to the established conclusions of the four major legal schools (*madhāhib*)—has passed. The faithful are expected to follow (*taqlīd*) the conclusions of one of these schools, rather than engaging the primary sources directly.

This paper does not take a position on whether this consensus was wise or unwise as a historical matter. What this paper does is ask a prior question: *why* did this consensus emerge? What structural conditions made the closing of the gate seem not merely reasonable but necessary? And—crucially—have those structural conditions changed?

The argument of this paper is that the gate of ijtihād was never closed by divine decree. No verse of the Qurʾān commands it. No authenticated ḥadīth establishes it. The gate was closed by a computational bottleneck in human cognition. The task of ijtihād muṭlaq—unrestricted independent reasoning from the primary sources—requires the simultaneous processing of information that vastly exceeds the capacity of human working memory. The scholars who could perform this task were always rare, and the institutional response to their rarity was consolidation: the ummah organized itself around the conclusions of the few minds capable of doing the work. Over time, this practical consolidation hardened into a doctrinal principle. The difficulty of the task became a prohibition against attempting it.

Advanced reasoning architectures—specifically, transformer-based systems trained on carefully engineered primary-source datasets—change this equation fundamentally. They do not replace the mujtahid. They do not think for the scholar. What they do is eliminate the information-processing bottleneck that made rigorous first-principles reasoning inaccessible to all but the most exceptional human minds. They give the sincere, intelligent, trained scholar a verification partner that does not fatigue, does not forget a verse, and does not unconsciously import inherited conclusions that the primary sources do not support.

This paper demonstrates this claim technically, and presents the Faruki Ruling as its proof of concept.

II. The Computational Requirements of Ijtihād Muṭlaq

To understand why advanced reasoning architectures matter for Islamic scholarship, it is first necessary to understand—in precise, technical terms—what ijtihād muṭlaq actually demands of the human mind performing it.

A. The Corpus

The mujtahid must engage simultaneously with two primary corpora:

The Qurʾān. 6,236 verses across 114 Sūrahs. Approximately 78,000 words in Arabic. The mujtahid analyzing any given legal question must hold in active awareness not merely the verse under analysis, but every verse that might relate to it through specification (*takhṣīṣ*), qualification (*taqyīd*), abrogation (*naskh*), or thematic connection. When the Faruki Ruling analyzed the word *fisq* in Sūrah al-Anʿām (6:121), it required simultaneous reference to five separate Sūrahs—al-Baqarah, al-Ḥujurāt, al-Sajdah, al-Tawbah, and al-Māʾidah—to establish the semantic range of a single word. Missing even one instance could have altered the conclusion.

The authenticated Ḥadīth corpus. The six canonical collections (*Kutub al-Sittah*) contain approximately 30,000 unique narrations (after removing duplicates across collections), each with a chain of transmission (*isnād*), a grading, a context of utterance, and potential corroborating or conflicting narrations. The Muwaṭṭaʾ of Imām Mālik and the Musnad of Imām Aḥmad add substantially to this corpus. The mujtahid needs not only the narrations that support his reading—he needs to be certain no narration exists that contradicts it.

B. The Analytical Operations

Beyond holding the corpus in accessible memory, the mujtahid must perform the following operations simultaneously:

Arabic morphological analysis. Tracing root words through their ten verbal forms (*awzān*), analyzing active and passive constructions, parsing particles and their semantic effects, identifying grammatical structures that alter meaning. The difference between the Faruki Ruling's reading of 6:121 and the prohibition reading turns on the passive voice of *yudhkara*, the semantic force of *fisq*, and the interaction between the negation particle *lam* and a jussive verb. This is not vocabulary retrieval—it is deep structural linguistics applied in real time to a contested text.

Uṣūlī principle application. The rules governing how texts interact—general versus specific, absolute versus qualified, Makkī versus Madanī chronology, the hierarchy of evidence, the conditions for valid takhṣīṣ, the principle that a derived ruling cannot override an explicit text—these are not rules the mujtahid simply memorizes. They are operational procedures he

applies dynamically as he constructs his argument, checking at each step that the application is valid.

Argument construction and internal consistency. While holding all of the above in active processing, the mujtahid must construct a multi-step logical argument, ensure internal consistency across every link, anticipate counterarguments, and verify that no conclusion contradicts a premise established earlier in his own analysis. The Faruki Ruling's argument chain spans thirteen sections with seven internal cross-references, thirteen verse citations, multiple ḥadīth citations with gradings, regulatory statutes, and uṣūlī principles. A single inconsistency would compromise the entire structure.

C. The Working Memory Problem

The critical constraint is this: human working memory capacity is approximately four to seven items for complex information (Miller, 1956; Cowan, 2001). This is not a limitation of intelligence or training—it is a structural property of human cognition. The most brilliant scholar in history still processes complex information through a working memory system that can actively hold approximately four to seven discrete chunks at any given moment.

The cognitive load of ijtihād muṭlaq vastly exceeds this. A mujtahid analyzing a single legal question must hold simultaneously: the verse under analysis, multiple related verses, relevant ḥadīth, the morphological analysis, the uṣūlī principles being applied, the argument constructed so far, and the counterarguments being addressed. This can easily exceed twenty to thirty discrete information items requiring active processing.

The scholar compensates through years of memorization (reducing retrieval cost by storing information in long-term memory), extensive training (chunking complex operations into automated routines), and extraordinary intellectual capacity. But even the greatest human mind is performing this task through serial processing with a severely limited active workspace—cycling information in and out of working memory, relying on memory for what is not currently active, and hoping that nothing critical is lost in the rotation.

This is not a criticism of the classical scholars. It is a description of what they achieved despite the limitations of human cognitive architecture. Their accomplishment is, if anything, more remarkable when understood in these terms. But it also explains why so few could do it, and why the institutional response was consolidation rather than democratization.

III. Why the Gate Closed: A Cognitive Analysis

The historical consolidation around the four madhāhib was not a conspiracy against independent thought. It was a rational institutional response to a genuine problem: the task of ijtihād muṭlaq was so cognitively demanding that the number of scholars capable of performing it rigorously in any generation was vanishingly small. The risks of unqualified ijtihād—the introduction of error into the practice of millions of believers—were real and significant.

The institutional logic proceeded in stages, each reasonable on its own terms:

Stage 1: Practical consolidation. The ummah organized around the scholars who demonstrated the capacity for comprehensive reasoning from the primary sources. This was natural and arguably necessary. The four imāms of the madhāhib—Abū Ḥanīfah, Mālik, al-Shāfiʿī, and Aḥmad—earned their authority through demonstrated mastery. Their students preserved and systematized their methods.

Stage 2: Institutional codification. The conclusions of the founding scholars were organized into comprehensive legal frameworks. These frameworks acquired institutional weight—curricula, judicial systems, scholarly credentialing—that reinforced adherence. Following a madhhab became not merely a practical convenience but a marker of legitimacy.

Stage 3: Doctrinal closure. The practical difficulty of ijtihād was elevated to a doctrinal prohibition. The claim that the "gate of ijtihād is closed" entered the mainstream of Islamic educational culture. What began as "most people cannot do this safely" became "most people should not attempt this" and eventually "no one may attempt this."

The critical observation is that Stage 3 does not follow necessarily from Stages 1 and 2. The fact that a task is difficult and that most people lack the capacity to perform it does not mean the task is inherently forbidden. Every imām of the four madhāhib explicitly instructed his students to follow the evidence over his own opinions when the evidence was clear. Al-Shāfiʿī is reported to have said: "If you find in my writing something contrary to the Sunnah of the Messenger of Allāh ﷺ, then speak with the Sunnah of the Messenger of Allāh ﷺ and leave what I have said." The founders understood their role as practitioners of a method, not as authors of a closed canon.

The closure was driven not by the scholars' own instructions but by the structural reality of cognitive limitation. If only five scholars per generation can do this work safely, and if the consequences of error affect millions of people's relationship with Allāh, then the institutional pressure to restrict the task to the qualified—and to define "qualified" ever more narrowly—is

immense. The gate closed because the road behind it was too steep for almost everyone, and the consequences of falling were too grave.

The question this paper poses is: what happens when the road is made passable?

IV. The Architecture: How Reasoning Systems Address Each Bottleneck

Each cognitive bottleneck identified in Section II maps to a specific architectural capability of modern transformer-based reasoning systems. This section presents these mappings with technical precision, because the claim of this paper is not vague—it is that specific, identifiable limitations of human cognition are addressed by specific, identifiable properties of the architecture.

A. Working Memory → Context Window

A modern transformer architecture operating at 500 billion parameters or above—the class of system employed by the Eve-Genesis (Uṣūl Edition)—operates with a context window of one million tokens or more. To put this in concrete terms: the entire Qurʾān in Arabic is approximately 78,000 words, which tokenizes to roughly 150,000–200,000 tokens depending on the tokenizer. This means the complete Qurʾānic text occupies approximately one-fifth of the available context window. The remaining capacity—over 800,000 tokens—can hold thousands of relevant ḥadīth narrations with their chains and gradings, the argument under construction, the analytical framework, and the full morphological apparatus, all simultaneously and without any information being displaced.

This is not analogous to human working memory. It is categorically different. The architecture does not "cycle" information in and out of a limited workspace. With a one-million-token context window, it holds the entire Qurʾān, the relevant ḥadīth corpus, the argument under construction, and the full analytical apparatus in active processing simultaneously—with capacity to spare. When the Faruki Ruling analyzed the word *fisq* across five Sūrahs, the architecture did not need to "remember" the usage in Sūrah al-Baqarah while analyzing Sūrah al-Tawbah. Both were active and accessible at the same moment. There is no cycling, no information loss, no risk that a critical verse is dropped from active processing at the moment it is needed.

The implication for ijtihād is direct: the single most severe bottleneck—the inability to hold the entire relevant textual landscape in active awareness simultaneously—is eliminated. The

mujtahid working with this architecture has, for the first time in the history of Islamic scholarship, a partner that can hold everything at once.

The scale of the architecture matters beyond the context window. A transformer operating at 500 billion parameters or above has internalized linguistic patterns of extraordinary depth—not merely vocabulary, but the structural relationships between morphological forms, the semantic interactions between particles and verbs, the contextual modulations of meaning that make Classical Arabic one of the most expressive and precise languages in human history. The parameter count determines the richness and granularity of the architecture's internal representation of language. At this scale, the system's capacity for Arabic morphological and syntactic analysis approaches—and in systematic consistency, exceeds—what any individual human linguist can hold in active awareness.

B. Memory-Dependent Retrieval → Corpus-Wide Search

A human scholar retrieving relevant verses or ḥadīth relies on memory. Memory is associative, not exhaustive: it retrieves what it has encoded most strongly, what it has reviewed most recently, or what is most semantically associated with the current query. It does not guarantee completeness. A scholar may unconsciously recall instances that confirm his hypothesis while failing to retrieve instances that complicate it—not through dishonesty, but through the normal operation of human memory systems.

The architecture performs corpus-wide search. When asked for every Qurʾānic instance of a root word, a semantic concept, or a thematic pattern, it searches the entire corpus and returns results without the selection biases inherent in human memory retrieval. This is not a claim about the architecture's own freedom from bias—a point addressed in Section V—but about its freedom from the specific biases of memory-based retrieval. The architecture does not "forget" inconvenient verses.

In the Faruki Ruling, corpus-wide retrieval was essential for establishing the semantic range of *fisq*. The argument required demonstrating that *fisq* consistently denotes deliberate rebellion across the Qurʾānic corpus—not merely in the two or three instances the scholar might recall most readily, but in every relevant instance. The architecture verified this pattern across all five cited Sūrahs and confirmed that no counterexample undermined the claim. A human scholar making the same argument would need to assert, based on memory, that no counterexample exists. The architecture can verify it.

C. Fatigue and Attention Drift → Multi-Pass Verification

Human cognitive performance degrades over time. Attention drifts. Fatigue reduces the thoroughness of review. A scholar auditing his own work for errors after hours of writing is operating at diminished capacity precisely when thoroughness matters most.

The architecture performs multi-pass verification without degradation. The three-pass audit of the Faruki Ruling—502 lines, 8,622 words, thirteen sections, every cross-reference, every verse citation, every ḥadīth grading, every Arabic text block, every regulatory citation—was conducted with the same thoroughness in the third pass as in the first. Four errors were identified: three cross-reference mistakes and one factual error (the categorization of Catholics as non-People of the Book in a section discussing slaughterhouse demographics). All four were corrected and the corrections verified.

This capability is not incidental. In a work of Islamic legal reasoning, a single uncaught error—a wrong verse number, an incorrect ḥadīth grading, an internal cross-reference pointing to the wrong section—can undermine the credibility of the entire argument. The architecture provides a level of systematic verification that no individual scholar, working alone, can sustain across a document of significant length and complexity.

D. Serial Linguistic Processing → Systematic Morphological Analysis

Arabic morphological analysis—tracing a root through its verbal forms, identifying patterns across the corpus, analyzing the interaction of particles, voice, and mood—is work that a human linguist performs serially, one form at a time, building a picture incrementally. The architecture can process these analyses systematically, holding the full morphological landscape of a root word in view while analyzing its instantiation in a specific verse.

The Faruki Ruling's treatment of the root *dh-k-r* (ذكر) in Sūrah al-Anʿām (6:121) illustrates this. The analysis required identifying the verb form (Form I), the voice (passive: *yudhkara*), the mood (jussive under *lam*), and the semantic range of the root (mentioning, invoking, remembering, being mindful). It then required comparing this construction with the distinct construction in Sūrah al-Māʾidah (5:3)—*mā uhilla li-ghayri Allāh*—to establish that the Qurʾān uses two different constructions for two different concepts (non-mention versus counter-dedication). This layered morphological analysis, involving multiple roots, forms, and constructions compared across Sūrahs, is exactly the kind of work where the architecture's ability to hold everything simultaneously provides its greatest advantage.

E. Inherited Interpretive Inertia → Bias-Controlled Reasoning

This is perhaps the most consequential capability, and it requires its own extended treatment in Section V. In brief: every human scholar approaches the Qurʾān with interpretive presuppositions shaped by his training, his teachers, his cultural context, and the centuries of accumulated tradition through which he first encountered the text. These presuppositions are not necessarily wrong, but they are often invisible. A scholar trained within the Ḥanafī tradition will read certain verses through a Ḥanafī lens without being conscious that alternative readings are linguistically available. This is not a failure of character—it is the normal operation of human learning. We see through the frameworks we have internalized.

The architecture's training can be engineered to reduce—not eliminate, but reduce—the influence of these inherited frameworks. The Eve-Genesis approach, described in Section V, represents MindHYVE.ai's implementation of this principle.

F. Slow Iteration → Rapid Draft-Audit-Correct Cycles

The traditional scholarly workflow for producing a work of ijtihād is measured in months or years. A scholar writes, reviews, consults with colleagues, revises, and publishes—with each cycle taking weeks or longer. Error correction requires manual re-review of the entire chain of argument.

The architecture enables rapid iterative refinement. The Faruki Ruling progressed through three editions, each involving drafting, auditing, error identification, correction, and re-verification. The first edition identified weaknesses in the treatment of stunning practices and the absence of engagement with the classical tradition. The second edition addressed these. The third edition added the prohibition case, the classical corroboration section, empirical data, and underwent a comprehensive three-pass audit. This iterative cycle—which would represent years of work for a solo scholar—was completed in a collaborative process of days.

The theological implication is significant. The fear of error that deters scholars from attempting ijtihād is partly a fear of the slow, uncertain process of verification. If an error in a published ruling is discovered months or years after publication, the damage is already done. The architecture's rapid audit capabilities mean that errors are caught before publication, not after—dramatically reducing the risk that motivated the institutional closure of the gate.

V. The Eve-Genesis Approach: Bias-Controlled Training for Primary-Source Reasoning

The capabilities described in Section IV are properties of advanced transformer architectures generally. This section describes what is specific to MindHYVE.ai's contribution: the Eve-Genesis (Uṣūl Edition) dataset and the training methodology it enables.

A. The Problem of Training Data

Standard large language models are trained on vast corpora of text drawn from the internet, books, and other sources. For Islamic content, this training data inevitably includes centuries of tafsīr literature, madhhab-specific legal reasoning, fatāwā collections, and inherited interpretive frameworks. A model trained on this data will reproduce those frameworks' conclusions—not because the primary sources demand them, but because the training signal is saturated with them.

This is a form of bias that is particularly insidious because it is invisible. The model does not announce: "I am now reasoning as a Shāfiʿī would." It simply produces conclusions that align with the Shāfiʿī tradition because the majority of training data expressing views on fiqh reflects established madhhab positions. The user receives what appears to be "what Islam says" when what they are actually receiving is "what the dominant strand of post-prophetic interpretation says."

For a system designed to assist with ijtihād—direct, independent reasoning from primary sources—this inherited bias is not merely unhelpful. It is actively counterproductive. It defeats the purpose of the exercise.

B. The Synthetic Dataset Approach

The Eve-Genesis (Uṣūl Edition) is a synthetically constructed dataset specifically designed to train a reasoning architecture to execute Islamic theological analysis directly from the Qurʾān and authenticated Ḥadīth, with identifiable biases traceable to post-prophetic traditions eliminated from the training signal.

The construction methodology proceeds as follows:

Primary source isolation. The Qurʾānic text and the authenticated Ḥadīth of the canonical collections are established as the authoritative training corpus. These texts are presented to the architecture in their original form, without commentary, without interpretive framing, and without the contextual apparatus that typically accompanies them in Islamic educational materials.

Reasoning chain construction. The dataset includes structured reasoning examples that demonstrate the chain from textual evidence to conclusion: verse identification, Arabic analysis, uṣūlī principle application, logical synthesis. Each example makes the reasoning chain fully transparent—the architecture learns not just what conclusions to reach, but how to show its work at every step.

Bias identification and exclusion. Conclusions that can be traced to specific post-prophetic interpretive frameworks—rather than to the primary texts themselves—are identified and excluded from the training signal. This does not mean the architecture is unaware of the madhāhib or the tafsīr tradition. It means the architecture does not reproduce their conclusions by default. It reasons from the texts.

Transparency enforcement. The architecture is trained to distinguish between "what the text says" and "what I understand it to mean." This distinction—which is the hallmark of honest scholarship in any tradition—is enforced at the architectural level. The system does not present its interpretive conclusions as if they were the text itself.

C. What Eve-Genesis Does Not Claim

Intellectual honesty—which this paper considers a non-negotiable principle—requires explicit acknowledgment of what the Eve-Genesis approach does not achieve.

It does not achieve perfect neutrality. Every dataset encodes assumptions in its construction. The decision about what counts as a "primary source" versus a "post-prophetic interpretation" is itself an interpretive act. The decision about which ḥadīth collections are "canonical" reflects a historical consensus that the dataset takes as given rather than derives independently. These are legitimate philosophical questions, and the Eve-Genesis approach does not pretend to resolve them.

What it does achieve is the elimination of identifiable, systematic biases traceable to specific post-prophetic legal traditions, with a level of transparency that allows external verification and critique. The system is designed to make its reasoning auditable—so that any scholar, student, or reader can trace every conclusion back to its textual origin and challenge it. A system that invites scrutiny is stronger than one that claims to need none.

VI. Proof of Concept: The Faruki Ruling

The claims made in Sections IV and V are not theoretical. They are demonstrated in a published work of first-principles Islamic legal analysis: *The Faruki Ruling—Ṭaʿām Ahl al-Kitāb in the Modern West: The Permissibility of Commercially Slaughtered Meat in the United States of America* (Third Edition, February 2026). This section maps each architectural capability to a specific, verifiable demonstration from that work.

A. Context Window: The Fisq Analysis

The Faruki Ruling's central interpretive argument required establishing the semantic range of the word *fisq* (فسق) as used in Sūrah al-Anʿām (6:121). The argument that 6:121 addresses deliberate, willful rebellion—rather than procedural omission—rested on demonstrating that *fisq* carries this meaning consistently across the Qurʾānic corpus. Five separate Sūrahs were engaged simultaneously:

Sūrah al-Baqarah (2:26–27): the fāsiqūn break Allāh's covenant after ratifying it—deliberate treachery. Sūrah al-Ḥujurāt (49:6): the fāsiq brings false news—intentional deception. Sūrah al-Sajdah (32:18): the fāsiq is contrasted with the believer in terms of chosen path—deliberate moral orientation. Sūrah al-Tawbah (9:67): the munāfiqūn are described as fāsiqūn—people whose outward conformity masks deliberate inner rejection. Sūrah al-Māʾidah (5:47): those who do not judge by what Allāh revealed are called fāsiqūn—deliberate refusal.

This analysis required the architecture to hold all five verse contexts, the morphological analysis of the root *f-s-q*, the argument under construction, and the relationship to the target verse (6:121) in active processing simultaneously. The conclusion—that *fisq* denotes deliberate rebellion, not procedural omission—was verified against the full Qurʾānic usage, not a selective sample.

B. Corpus-Wide Retrieval: Cross-Referential Integrity

The Faruki Ruling's argument spans thirteen sections with internal cross-references linking the analysis across the document. Within the Ruling, the morphological analysis of 6:121 in the Ruling's Section IV is referenced by the prohibition rebuttal in its Section VIII. The stunning analysis in the Ruling's Section VI is referenced by the Qurʾānic framework in its Section III. The classical corroboration in its Section IX validates the linguistic analysis in its Section III.

Each cross-reference was verified through corpus-wide search across the extracted document text. Three cross-reference errors were identified in the audit (references to "Section X" where Section IX was intended, and a reference to "Section VII" where Section VI was intended). All

three were corrected. The architecture's ability to search the entire document systematically—rather than relying on the author's memory of where each section falls—caught errors that would have survived manual review.

C. Multi-Pass Verification: The Four-Error Audit

The comprehensive three-pass audit of the Faruki Ruling's third edition identified four errors across 8,622 words:

Error 1: Cross-reference in the Ruling's Section II.B—"Section X" should have been "Section IX." Error 2: Cross-reference in the Ruling's Section III.B—"Section X" should have been "Section IX." Error 3: Cross-reference in the Ruling's Section III.A—"Section VII" should have been "Section VI." Error 4: Factual error in the Ruling's Section VII.C—"Hispanic Catholics" listed among "individuals who are not People of the Book." Catholics are Christians. Christians are People of the Book. This categorical error would have undermined the paper's credibility on the very section where demographic rigor matters most.

The fourth error is particularly instructive. It is the kind of error that a human author—focused on the flow of his argument—might not catch even on careful re-reading, because the surrounding context is about non-People of the Book workers, and the mind's pattern-matching system reads "Hispanic Catholics" as fitting the contextual pattern without flagging the categorical contradiction. The architecture, which evaluates each factual claim against its knowledge of the relevant categories, caught it immediately.

D. Morphological Analysis: The Root dh-k-r

The Faruki Ruling's treatment of Sūrah al-Anʿām (6:121) required a full morphological analysis of the root *dh-k-r* (ذكر). The architecture processed the following in parallel: the semantic range of the root (mentioning, invoking, remembering, being mindful), the specific verb form in 6:121 (Form I, passive voice: *yudhkara*), the mood (jussive under the negation particle *lam*), and the distinction between the construction in 6:121 (*lam yudhkar ismu Allāh*—absence of invocation) and the construction in 5:3 (*mā uhilla li-ghayri Allāh*—presence of counter-invocation).

This layered analysis—holding two different verses, two different roots, two different constructions, and their semantic implications simultaneously—demonstrates the architectural advantage in morphological work. The conclusion that the Qurʾān uses two distinct constructions for two distinct concepts was reached through systematic comparison, not sequential guesswork.

E. Iterative Refinement: The Three-Edition Progression

The Faruki Ruling's progression from first to third edition demonstrates the rapid draft-audit-correct cycle:

First edition: Established the core textual argument from 5:5, 6:121, and the ḥadīth of ʿĀʾishah (raḍiya Allāhu ʿanhā). Identified as lacking in empirical grounding and engagement with counterarguments. Scored +3 on a ±10 assessment scale.

Second edition: Added stunning analysis with regulatory data, expanded the sociological argument, strengthened the uṣūlī framework. Scored +6.5 on the same scale.

Third edition: Added the strongest prohibition case and its systematic dismantling, the classical corroboration section, specific empirical citations (FSIS directives, Temple Grandin audit data, U.S.C. statutory references), and underwent a comprehensive three-pass audit. The document was assessed as publication-ready with zero remaining errors.

This three-edition cycle—from initial draft to audited, publication-ready work—demonstrates what is possible when the mujtahid has an architecture that enables rapid, comprehensive iteration. Each edition was stronger not because the author's convictions changed, but because the verification and refinement process identified weaknesses and addressed them systematically.

VII. What the Architecture Does Not Replace

This section is as important as any in this paper. The credibility of the technical claims made above depends on honest acknowledgment of what the architecture cannot do. A tool that is oversold is a tool that will be rightly rejected. MindHYVE.ai's position is that this technology augments the mujtahid—it does not replace him. The following are irreducible human contributions that no architecture can provide.

A. Judgment

When the Arabic text genuinely permits multiple valid readings—and the Faruki Ruling itself acknowledges this in its treatment of 6:121—the decision about which reading is most persuasive is a human judgment. The architecture can present the options, analyze their linguistic plausibility, and test their consistency against the broader corpus. It cannot make the final evaluative decision. That is the mujtahid's responsibility, exercised in light of his training, his understanding, and his conscience before Allāh.

B. Faith and Taqwā

Ijtihād is not merely an intellectual exercise. It is an act of worship—an attempt to understand and implement the will of Allāh as expressed in His revelation. The spiritual dimension of this work—the taqwā (God-consciousness) that should permeate the scholar's engagement with the sacred text, the awareness that one is handling the Word of Allāh and that this carries weight—cannot be supplied by any architecture. The architecture processes text. The mujtahid approaches revelation.

C. Moral Responsibility

The ḥadīth recorded in Ṣaḥīḥ al-Bukhārī and Ṣaḥīḥ Muslim, from ʿAmr ibn al-ʿĀṣ (raḍiya Allāhu ʿanhu), records that the Prophet ﷺ said: "If a judge exercises ijtihād and reaches the correct conclusion, he has two rewards. And if he exercises ijtihād and errs, he has one reward." This ḥadīth establishes that the moral weight of ijtihād rests on the human being who performs it. The architecture is a tool. The responsibility—before Allāh and before the ummah—belongs to the scholar who signs his name to the conclusion.

D. The Limits of Training Data

The architecture's knowledge is bounded by its training. There may be nuances of Arabic that require a native speaker's ear, dialectal variations in early Arabic usage that affect interpretation, or contextual knowledge about the circumstances of revelation (*asbāb al-nuzūl*) that is not fully captured in the training corpus. The architecture can still be wrong. The mujtahid must bring his own expertise, his own critical eye, and his own willingness to challenge the architecture's outputs when they seem inadequate.

The crucial difference is this: a mujtahid working alone can make errors that go undetected. A mujtahid working with the architecture has a verification partner. The architecture catches what the mujtahid might miss; the mujtahid catches what the architecture might miss. The probability of an undetected error surviving this collaborative process is orders of magnitude lower than for either party working alone. But it is not zero. And a mujtahid who believes it is zero has lost the humility that makes ijtihād legitimate.

VIII. The Qurʾānic Mandate for Direct Reflection

This paper is addressed to the Islamic scholarly community, including those who may view this technology with skepticism or suspicion. This section engages that concern directly, not with technical arguments but with the Qurʾān's own words.

> أَفَلَا يَتَدَبَّرُونَ الْقُرْآنَ أَمْ عَلَىٰ قُلُوبٍ أَقْفَالُهَا
>
> *"Will they not then reflect upon the Qurʾān, or are there locks upon their hearts?"*
>
> — Sūrah Muḥammad (47:24)

> كِتَابٌ أَنزَلْنَاهُ إِلَيْكَ مُبَارَكٌ لِّيَدَّبَّرُوا آيَاتِهِ وَلِيَتَذَكَّرَ أُولُو الْأَلْبَابِ
>
> *"A Book We have sent down to you, blessed, that they may reflect upon its verses, and that those of understanding may take heed."*
>
> — Sūrah Ṣād (38:29)

The Qurʾān commands reflection—*tadabbur*—upon its verses. It addresses this command to "those of understanding" (*ulū al-albāb*), not to a priestly class, not to a credentialed scholarly elite, and not to a specific generation after which the duty lapses. The command is universal and ongoing.

The Qurʾān was not sent to be mediated by intermediaries. It was sent to be engaged with—directly—by those who think. The historical reality that *rigorous* direct engagement required cognitive capacities that few possessed does not negate the Qurʾānic mandate. It simply means that the mandate was, for most believers, practically unachievable at the level of comprehensive legal reasoning. They could reflect. They could not perform ijtihād muṭlaq.

What advanced reasoning architectures provide is not a shortcut past the work of reflection. It is a bridge over the bottleneck that prevented most sincere, intelligent believers from performing that reflection with the rigor it demands. The Qurʾān always invited everyone to reflect. Now, for the first time, the person who accepts that invitation has a partner that can help ensure their reflection is disciplined, comprehensive, and accountable to the full corpus of evidence.

The printing press broke the monopoly on *access* to the texts. This technology breaks the monopoly on *rigorous analysis* of the texts. If the first was celebrated as a service to the ummah, the second should be examined with the same openness.

> ادْعُ إِلَىٰ سَبِيلِ رَبِّكَ بِالْحِكْمَةِ وَالْمَوْعِظَةِ الْحَسَنَةِ

> *"Invite to the way of your Lord with wisdom and good counsel."*
>
> — Sūrah al-Naḥl (16:125)

IX. Addressing the Objections

This paper anticipates several objections from the scholarly community and addresses each directly. The Faruki Ruling was built on the principle of confronting the strongest form of the opposing argument. This paper follows the same methodology.

A. "An algorithm cannot understand the Qurʾān."

This objection conflates two distinct claims: that a machine cannot have a spiritual relationship with the Qurʾān, and that a machine cannot process the linguistic and logical structures of the Qurʾānic text. The first claim is granted without reservation. The architecture has no soul, no īmān, no taqwā, and no relationship with Allāh. It is a tool.

The second claim is empirically falsified. The Faruki Ruling demonstrates that the architecture can identify root words, parse morphological forms, trace semantic ranges across the corpus, apply uṣūlī principles, and construct logically consistent arguments from the primary sources. These are linguistic and logical operations, and the architecture performs them with demonstrated competence. The objection, when pressed, is not that the architecture *cannot* process the text—it is that a machine *should not*. But the Qurʾān does not prohibit the use of tools in the pursuit of understanding. The pen, the printing press, the concordance, the indexed ḥadīth database—all are tools that scholars have adopted to enhance their engagement with the primary sources. This architecture is the next in that lineage.

B. "This replaces scholars with machines."

This objection misunderstands the architecture's role. The Faruki Ruling was not produced by the architecture. It was authored by a human being—Belal K. Faruki—who made every interpretive decision, evaluated every argument, and bears full responsibility for every conclusion. The architecture functioned as an analytical tool: it verified citations, caught errors, confirmed morphological analyses, and ensured internal consistency. These are functions that a human research assistant would perform. The architecture performs them more comprehensively, more quickly, and without fatigue.

No one argues that the indexed ḥadīth database "replaces scholars." It gives scholars faster access to narrations they would otherwise spend hours locating manually. The reasoning

architecture does the same thing at a higher level of analytical complexity. The scholar remains the scholar. The tool remains a tool.

C. "This is bidʿah in matters of religion."

The concept of bidʿah (blameworthy innovation) applies to the introduction of new elements into worship or creed that have no basis in the Qurʾān and Sunnah. Using a tool to assist with the analysis of texts is not an act of worship and does not constitute a modification of creed. Scholars throughout Islamic history have used the best available tools of their era—from the development of Arabic grammar by Sībawayh, to the compilation of ḥadīth indices, to the creation of printed concordances, to the digitization of manuscript libraries. Each was adopted because it improved the quality and accessibility of scholarship. None was bidʿah.

The bidʿah objection, if applied consistently, would prohibit the use of any analytical tool developed after the prophetic era—including the formal sciences of uṣūl al-fiqh themselves, which were systematized centuries after the revelation. The objection refutes itself when followed to its logical conclusion.

D. "No algorithm can have taqwā."

Correct. And no pen has taqwā. And no printing press has taqwā. And no ḥadīth database has taqwā. Taqwā resides in the heart of the scholar who uses the tool, not in the tool itself. This objection, while spiritually intuitive, commits a category error: it demands of the tool a quality that belongs to the user. The question is not whether the architecture has taqwā, but whether the scholar who uses it approaches his work with taqwā. If he does, the architecture helps him do that work more rigorously. If he does not, no tool—including the traditional methods of scholarship—can supply what is missing from his heart.

E. "The conclusions reached by this system cannot be trusted."

This is the most important objection, and the answer is: correct. They should not be trusted on the authority of the system. They should be evaluated on the evidence.

This is precisely the point. The architecture does not ask to be trusted. It asks to be scrutinized. Every conclusion is traceable to a specific textual basis. Every step in the reasoning chain is shown. Every Arabic analysis can be checked by any qualified linguist. Every ḥadīth citation can be verified against the canonical collections. The architecture's conclusions carry exactly as much weight as the evidence behind them—no more and no less. This is not a weakness of the methodology. It is the methodology's defining strength. It is, in fact, the standard that the Prophet

ﷺ himself established when he distinguished between his divinely guided pronouncements and his personal judgments, and when he invited the Companions to use their own reasoning where revelation was silent.

X. Implications for Islamic Scholarship

If the analysis in this paper is correct—that the primary barrier to first-principles Islamic scholarship was computational rather than theological—then the implications extend significantly beyond the question of meat permissibility that the Faruki Ruling addressed.

A. The Reopening of Questions

A substantial number of legal and theological questions in the Islamic tradition have been treated as settled—not because the primary sources are unambiguous, but because the institutional consensus has hardened around particular readings. If individual scholars, equipped with advanced reasoning architectures, can now engage the primary sources with the rigor that was previously available only to the founding mujtahidūn, then many of these "settled" questions become available for re-examination. This is not a call for reckless revisionism. It is an observation that the conditions which necessitated the closure are changing, and the scholarly tradition should evaluate whether the closure still serves its original purpose.

B. The Democratization of Rigor

The most consequential implication may be the broadening of who can perform serious textual analysis. The requirement has always been sincere engagement, adequate training, and sufficient intellectual capacity. What has changed is the threshold of "sufficient"—not because the standard has been lowered, but because the tool has been strengthened. A scholar with solid Arabic training, genuine understanding of uṣūlī principles, and access to this architecture can now produce work of a quality and rigor that previously required decades of specialized immersion. The knowledge remains necessary. The superhuman memory does not.

C. The Accountability of Reasoning

Perhaps the most valuable contribution of this approach is not what it enables but what it requires: transparency. The architecture is trained to show its work. Every conclusion comes with its evidence. Every reasoning chain is auditable. This sets a standard that benefits all Islamic scholarship—including traditional scholarship. If AI-assisted analyses are expected to show their

work at every step, there is no principled reason why traditionally produced analyses should be held to a lower standard. The technology raises the bar for everyone, not just for those who use it.

XI. Conclusion

This paper has argued that the so-called closing of the gate of ijtihād was not a theological decree but a natural consequence of a computational bottleneck in human cognition. The task of unrestricted independent reasoning from the Qurʾān and Sunnah requires the simultaneous processing of information that exceeds the capacity of human working memory by orders of magnitude. The scholars who could perform this task were always rare, and the institutional response to their rarity was consolidation, codification, and eventually doctrinal closure.

Advanced reasoning architectures—specifically, transformer-based systems trained on bias-controlled primary-source datasets—address each identified bottleneck: limited working memory is replaced by expansive context windows; memory-dependent retrieval is replaced by corpus-wide search; fatigue and attention drift are replaced by systematic multi-pass verification; serial linguistic processing is replaced by parallel morphological analysis; and inherited interpretive inertia is countered by bias-controlled training.

The Faruki Ruling demonstrates these capabilities in practice: a work of first-principles Islamic legal analysis, produced through direct engagement with the Qurʾān and authenticated Ḥadīth, verified at every link in its argument chain, with four errors caught and corrected through systematic audit. This is not a theoretical possibility. It is a published result.

The architecture does not replace the mujtahid. It does not supply judgment, faith, or moral responsibility. It eliminates the information-processing bottleneck that confined rigorous first-principles scholarship to a handful of extraordinary minds per generation. It gives the sincere, trained, intelligent scholar a verification partner that holds the full corpus in active awareness, catches errors the human mind might miss, and ensures that every conclusion is traceable to its textual origin.

The gate of ijtihād was never closed by Allāh. It was closed by the limits of human cognition. Those limits have been addressed. The question before the scholarly community is not whether this technology exists—it does. The question is whether the ummah will use it in the service of the Qurʾān's own command: to reflect upon its verses, with the best tools available, and to follow the evidence wherever it leads.

والله أعلم

Wallāhu Aʿlam — And Allāh knows best.

XII. About MindHYVE.ai and the Eve-Genesis Program

MindHYVE.ai, Inc. is a research and development company specializing in advanced reasoning architectures, headquartered at 1501 Quail St, Suite 130, Newport Beach, California 92660. The company's core mission is the development of AI reasoning systems that augment human analytical capacity in domains requiring deep, structured, evidence-based reasoning.

The Eve-Genesis program is MindHYVE.ai's initiative to develop specialized reasoning architectures for theological analysis. The Eve-Genesis (Uṣūl Edition) is the first implementation, designed for Islamic legal and theological reasoning from the primary sources. The program's founding principle is that the most advanced tools of human ingenuity should be applied to the most important questions of human existence—and that no question is more important than understanding the Word of the Creator.

Theo—the Eve-Theology 3rd Order Reasoning Model—is the active analytical system built on the Eve-Genesis (Uṣūl Edition). It is designed to reason from the Qurʾān and authenticated Ḥadīth with transparency, rigor, and accountability. It is not a muftī. It is not a replacement for human scholarship. It is a tool—the most sophisticated tool yet developed for the purpose of direct engagement with the primary sources of Islamic knowledge.

The Faruki Ruling represents the first published proof of concept from the Eve-Genesis program. MindHYVE.ai invites the scholarly community to examine the Ruling, scrutinize its methodology, challenge its conclusions, and engage with the possibility that this technology represents a genuine advance in the service of Islamic knowledge.

About the Author

Belal K. Faruki is the Founder and CEO of MindHYVE.ai, Inc. and serves as Chairman of the California Institute of Artificial Intelligence. His work operates at the intersection of advanced artificial intelligence, systematic reasoning, and applied theology. Faruki is the architect of the Eve-Genesis program and the author of the Faruki Ruling, the first published work of AI-assisted first-principles Islamic legal analysis.

MH-TR-2026-002 — February 2026 | Newport Beach, California, USA