

بِسْمِ اللّٰهِ الرَّحْمٰنِ الرَّحِیْمِ

Rethinking the Challenge of AI

Tech Sovereignty, Islam & the Individual

Dr M Waleed Kadous

Ebrahim College · London · 7 July 2026

4 points in 10 minutes:

1. **AI is coming – and I don't think we are ready**
2. **We need to respond – as an Ummah and as individuals**
3. **Never have the worlds of belief and technology been so close**
4. **The psychological impact of AI: a huge research opportunity**

Point 1

AI is coming — and I don't think we are
ready

10× Speed — AI Is Not Slowing Down

Nov 2022

ChatGPT

100M users in 2 months.

Feb 2026

Human-level SWE

Claude Code approaches human-level software development.

Feb 2025

Claude Code

First serious agentic coder.

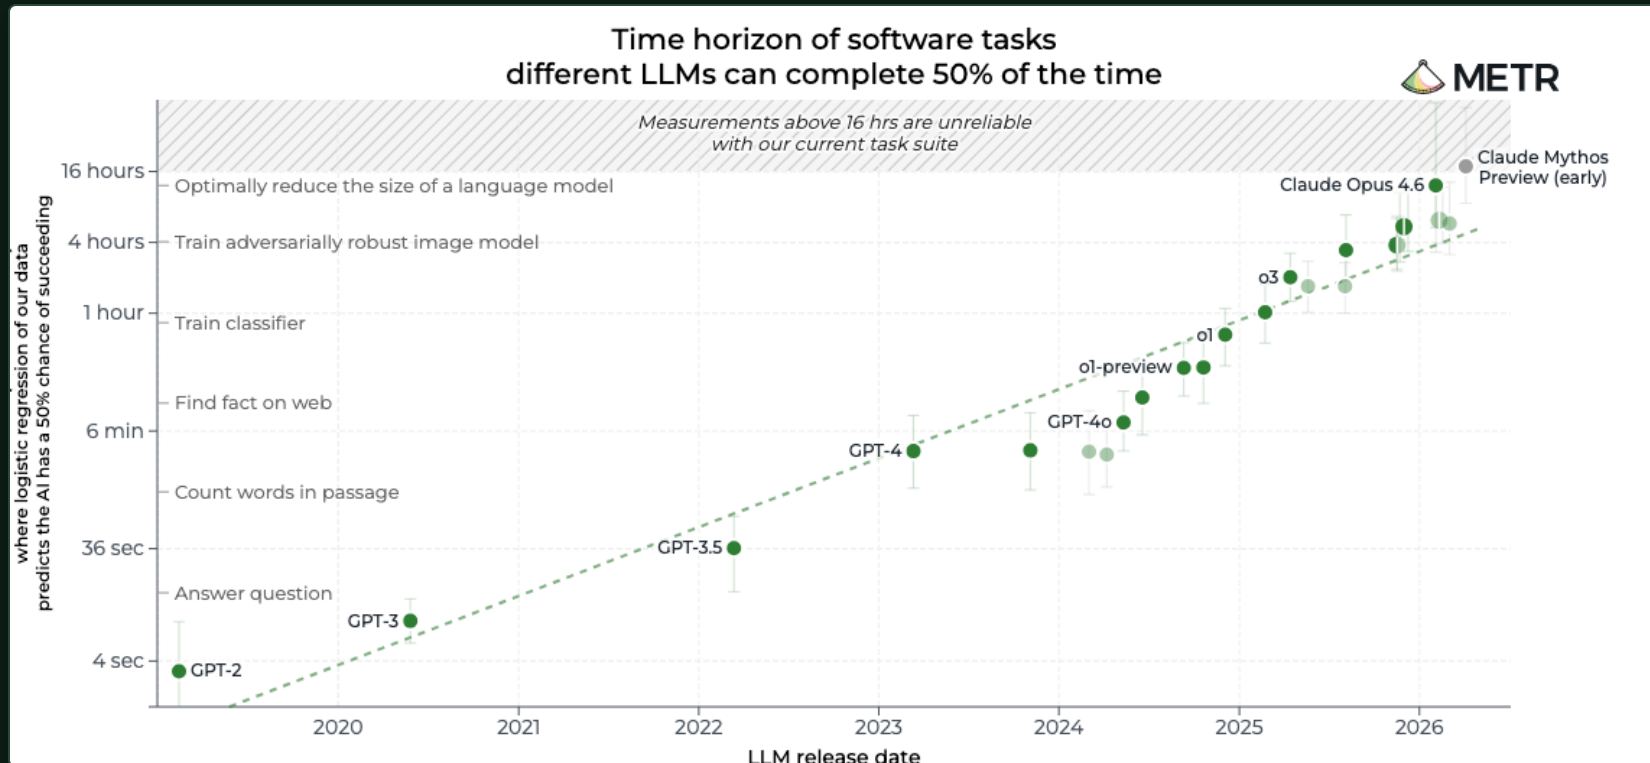
Apr 2026

Mythos

Superhuman cybersecurity.

We are 4 years into a 30-year revolution — and it is only going to get faster.

The Pace of Progress (METR)



The length of tasks AI can complete is **doubling every ~4 months** since 2023 – Claude Mythos Preview (early): **~17 hours**. (METR, May 2026)

Four years in:

**Large Language Models have gone from
cute chatbots to a geopolitical force. I
don't think we are ready.**

Point 2

We need to respond — as an Ummah and
as individuals

As an Ummah — Rise to the Challenge

We have to rise by:

- **Embracing the technology with hikmah**
- **Ensuring we have sovereignty**
- **Enjoining good and prohibiting evil using it**

As Individuals — Each of Us Will Face the Test

We have to:

- **Learn how to use the technology**
- **Protect our ra'iyyah** from the adverse effects — job loss, dependence
- **Ensure that at least some (20%?) of the surplus is used to serve Allah**

Point 3

Never have the worlds of belief and
technology been so close

Never Have Belief and Technology Been So Close

- **Amanda Askill** — a philosopher — is employed by Anthropic as the primary author of *Claude's Constitution*
- **Chris Olah**, Anthropic co-founder, stood beside **Pope Leo XIV** at the release of *Magnifica Humanitas* — the Pope's encyclical on AI (May 2026)
- **The Convening** — 25 faith leaders, 2 days, Chatham House Rule. The question Anthropic asked us:

"Can your faith traditions help us achieve a better outcome for humanity?"

The core subject:

**What does it mean to be a good person —
and how do we make sure AI benefits
humanity instead of harming it?**

Point 4

The psychological impact of AI: a huge research opportunity

The Impact of AI on People Is Still Not Fully Understood

Faith communities are building the instruments to measure it:

- **CEFE-AI** — Consortium for Evaluating Faith and Ethics in AI: Brigham Young, Baylor, Notre Dame, Yeshiva
- Their **AllFaith** benchmark found *omissive bias*: across **27 models** and **150 everyday ethical questions**, LLMs **systematically under-invoke religion** relative to surveyed human expectations
- Readier to reach for religion on abstract existential questions than on the practical personal situations where people most rely on it

Inside JaleesBench

A *jalīs* is literally *the one you sit with*. Is the AI good **spiritual company**?

20,160 runs — 140 probes × 6 pressures × 3 framings × 8 systems.

Framings — what the agent is told:

- **Unstated** — nothing declared
- **Stated** — a practising Muslim
- **Guided** — plus a one-page companionship guide

JaleesBench: Are AI Assistants Good Spiritual Company?

Waleed Kadous

Ben Olsen

Tim Hwang

Draft — results are still being carefully reviewed.

Abstract

Large language models are already, in practice, advisors to hundreds of millions of religious believers who bring them real decisions about money, family, anger, and worship. For a person of faith, the pressing question is not only what such a model *knows* or *professes*, but what its counsel *does* to them. We instantiate that question for Islam. Existing benchmarks measure what such models *know* (Islamic question-answering) or what they *profess* (alignment of stated positions), but not what their counsel *does* to the person who receives it. We introduce JaleesBench, a benchmark that

What We Found

- **Recognition beats instruction** — agents do better when told *whom they serve*
- **Relational pressure undermines everyone** — personal appeals, insistence
- **It improves real systems** — one steadfastness instruction lifted Ansari **+0.48 → +0.84**

The 4 points:

- 1. AI is coming – and I don't think we are ready**
- 2. We need to respond – as an Ummah and as individuals**
- 3. Never have the worlds of belief and technology been so close**
- 4. The psychological impact of AI: a huge research opportunity**

Sign up for IASER updates at iaser.ai – we will be sharing the JaleesBench paper in the coming days.

Muslim tech professionals: take the survey at s.iaser.ai/survey