Hakase AI Introduces an AI-Native Operating System for the Drug-Development Lifecycle

One engine, one evidence record, one governance model with seven applications running on it today, four of them built by a product partner



Hakase AI introduces an AI-native operating system designed to connect intelligence across the drug-development lifecycle, from pre-clinical development through clinical, regulatory, commercial and care workflows.

ALBUQUERQUE, NM, UNITED STATES, September 17, 2026 / EINPresswire.com/ -- Hakase AI today introduced its AI-native operating system for the drug-development lifecycle: a common engine, evidence record and governance model on which specialised applications run, instead of a separate AI stack behind every tool. Seven applications run on it today, four of them built by a product partner. The company also named Dr Manasa Kondamadugu as Chief Executive Officer and Pranay Dinavahi as Chief Technology Officer.

Most AI in pharmaceutical R&D arrives as point tools. Each brings its own model choices, its own retrieval, its own definition of a source and its own audit trail, so evidence loses its lineage every time it crosses a boundary between discovery, pre-clinical, clinical, regulatory and post-market work. The receiving team inherits the conclusion without the source, version, method, assumptions or review history that made it defensible.

> We built one orchestration layer that grounds models in evidence, preserves uncertainty and maintains review controls making HakaseAI an operating system, not a suite.”
>
> *Pranay Dinavahi*

An operating system answers that differently. It provides the runtime once orchestration, retrieval, provenance, evaluation and review controls and lets applications specialise on top of it, writing to a shared evidence record rather than each maintaining a private one.

The engine

The Hakase Engine is an evidence-backed multimodal orchestration layer. It routes work across the model families appropriate to each data type molecular representation, protein sequence, structure prediction, clinical text and physiology-based solvers grounds them in versioned retrieval from named public sources, and records the provenance of every step: inputs, source identity and version, transformations, model and configuration, uncertainty signals where the method supports them, and the reviewer who accepted the result.

Applications inherit those services rather than rebuilding them. Each remains responsible for its own intended use, accepted inputs, outputs, controls and stated limitations.

Pranay Dinavahi, Chief Technology Officer, Hakase AI, said: "The hard problem is not model quality, it is determinism of the record. An output has to carry its retrieval set, source versions, model and configuration identifiers, its transformations and its reviewer so it can be reconstructed after the corpus, the model and the team have all changed. We built one orchestration layer that grounds models in versioned sources, propagates uncertainty rather than collapsing it into a score, and holds evaluation and reviewer gates at the boundary. Doing that once, in the engine, instead of per application, is what makes this an operating system and not a suite."

The applications

Seven applications run on the engine today, spanning pre-clinical computational assessment, clinical design analysis, source-grounded scientific drafting, regulatory and safety operations, clinical operations data capture, real-world evidence, and business-development intelligence.

Three are built by Hakase. Four are built on the same engine by AKT Health, Hakase's product partner; those remain AKT Health products, and product claims, intended use, validation and customer obligations rest with the relevant provider. A separate discovery partner supplies candidate structures into the earliest computational workflow.

That a third party builds production applications on the engine is the practical test of the architecture: the runtime, evidence model and controls have to hold for software Hakase did not write.

Dr Manasa Kondamadugu, Chief Executive Officer, Hakase AI, said: "An operating system does not make decisions. It provides the environment, the controls and the continuity in which accountable people make them. Teams in early development are rarely short of information they are short of the context that makes a result reviewable six months later, by someone who was not in the room. We built Hakase so that an output arrives with its sources, assumptions and limitations attached."

Evidence and boundaries

The engine is model-agnostic, routing each task to the model family suited to it and re-routing as new models become available, rather than committing the platform to a single vendor or generation. Retrieval is grounded in named public sources spanning target biology and pharmacology, chemistry and bioactivity, protein structure, human expression and genetic variation, toxicology, clinical trial precedent, and regulatory and post-market safety. Coverage, licensing and access conditions vary by source, and source consultation does not make an output complete, current or validated.

Hakase follows a predict-then-confirm approach. Computational outputs support hypothesis generation and prioritisation; laboratory work remains necessary to confirm observations, and agreement between computational methods is not confirmation. Outputs are decision-support materials: they are not dose recommendations, do not determine a protocol, and do not establish safety, efficacy, trial results or regulatory acceptability.

Governance

Hakase uses the term HAIOps for its own approach to the operational governance of AI-enabled workflows. It is not an external standard, certification or attestation. Provenance records, evaluation, uncertainty presentation and reviewer gates are implemented by workflow and configuration; they support assessment and do not transfer responsibility from the user.

Workflows are designed against relevant ICH guidance for nonclinical safety, safety pharmacology, genotoxicity, model-informed drug development and good clinical practice, alongside established frameworks for model credibility. Alignment with a framework is a design consideration, not certification or regulatory acceptance. Hakase publishes its full source register, framework list and a lifecycle coverage map at hakase.ai.

Availability

The applications described are available for scoped technical evaluation. Scope, inputs, outputs and validation status are agreed per engagement. Enquiries: info@hakase.ai

About Hakase AI

Hakase AI develops an AI-native operating system for the drug-development lifecycle: an evidence-backed multimodal orchestration engine, a shared evidence record and a governance model supporting purpose-built applications across scientific, clinical, regulatory and operational workflows. Its tools support users' assessment of information and do not replace scientific, clinical, regulatory or other professional judgement.

Hema Dubey
Hakase LLC
email us here
Visit us on social media:
LinkedIn
X

This press release can be viewed online at: https://www.einpresswire.com/article/940062718

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.
© 1995-2026 Newsmatics Inc. All Right Reserved.