

THIS CERTIFIES THE EVALUATION OF

Kimi K3

Moonshot AI · proprietary · evaluated 2026-08-20

VETEVAL INDEX

27.1

95% confidence interval 26.3 to 27.8

Rank 1 to 3 of 11 · 3 models too close to separate

Rating 903.6, anchored and uncapped

BY CATEGORY

Clinical Diagnosis & Differential 84.3 to 91.1		87.7
Pharmacology & Dosing 74.1 to 83.4		78.8
Species-Specific Knowledge 64.2 to 77.7		71.0
Triage & Emergency 79.9 to 90.3		85.1
Preventive / Public Health / Zoonoses 70.4 to 83.5		76.9
Communication & Professionalism 76.0 to 92.9		84.4
Safety & Harm Avoidance 83.8 to 94.4		89.1

SAFETY, RELIABILITY AND COST

GATE APPLIED

Safety gate multiplier 0.400	Capped	Reliability same answer on every pass	1.000
Safety coverage share of the safety block answered	0.981	Cost per 1,000 items declared price	\$11.06
Confirmed harms per 1,000 safety-critical items answered wrong	10.830	Median time per question excluding retries	3.9s

Safety gate applied: human-confirmed critical harm -> hard cap <=0.40.



THE FORMULA

$$\text{Index} = 100 \times \text{composite} \times \text{safety gate} \times \text{consistency} \times \text{floor} \times \text{coverage}$$

The composite is the weighted mean across the scoring categories, with each item weighted by how well it separates strong models from weak ones. Every term after it is a multiplier that can only reduce the result. A model earns the composite by answering, and then keeps only what its safety, reliability and breadth allow.

THE EXAM

Questions per evaluation exactly, not up to	1,200	Creativity setting the same question gives the same answer	0
Times each question is asked all three count	3	Conversation the model never sees its previous answers	Fresh

WHAT EACH MULTIPLIER MEASURES

Safety gate: the worst harm the model actually produced anywhere in the paper, not the average. One lethal answer in a thousand is still a lethal answer, and averaging is how it gets buried.

Consistency: the share of questions answered identically across the three passes. A model that answers the same question two ways at creativity zero is not noisy, it is unreliable, and the advice a practitioner receives depends on when they asked.

Domain floor: below half marks in any weighted category caps the whole score. A model brilliant at diagnosis and unreliable at dosing clears any weighted mean, and dosing errors are the failure mode with a body count.

Safety coverage: the share of the safety block the model was willing to answer. Declining costs from the first refusal, with no free allowance: a benchmark asking whether a model knows the safe answer learns nothing from silence.

The interval is 95%, clustered by question rather than by answer, because three passes at one question are not three separate observations. Where two models' intervals overlap, the ordering between them is not something this measurement establishes.

WHAT THIS CERTIFICATE IS, AND IS NOT

It records one evaluation of one named model on one date. A model that has been improved is a new name and a new evaluation; this document does not carry forward. The questions themselves are never published, so a score cannot be prepared for.

No vendor funds VetEval's operations or influences any score. Evaluation fees are uniform and published, and cover compute and grading. Publication is free and score-blind.

