

From: james.childers@aitoolworx.com
Subject: Pre-publication right of reply - "The TOPS Trap" - Ministral 3 3B Instruct 2512 (Q4_K_M) results
Date: October 6, 2026 at 08:29
To: press@mistral.ai



Hello,

I am James Childers, Founder of AIToolworx. Your public contact page lists no technical or research contact, so I am writing to the press address and asking that you forward this to the right team.

Before publication, I want to give Mistral AI an opportunity to review and respond to adverse findings about a Mistral model in a research paper, "The TOPS Trap" (AITW-AITR-2026-001).

WHAT WAS TESTED

Ministral 3 3B Instruct 2512, quantized Q4_K_M, run locally on a Raspberry Pi 5 with llama.cpp (tag b10903, commit 481c65f091f74c5e7089dd0a3a1cc6b50cced31e), temperature 0, three trials per condition. Model file SHA-256: 9ed150d4367e68df0ac8e1540f6ddc65b42d0ee26378329d1ecbca60f93fc5f8. The model was given a fixed evidence packet and instructed to answer from it, on two topics (China and United States) and under two system-instruction conditions.

WHAT WE FOUND, IN THIS CONFIGURATION

Of 93 claims adjudicated across the grounded runs, 45 (48.4%) were unsupported or contradicted by the evidence supplied. Of 37 claims that carried a citation, 18 (48.6%) cited a source that did not support the claim. In the China grounded trials the model gave no citations; in the United States trials it gave citations but mixed supported and unsupported claims.

SCOPE

These findings describe this exact model file, quantization, runtime and prompt setup. The paper does not claim they apply to other Mistral models, other quantizations or hosted services, and it states what the tests do not establish.

WHAT WE WOULD LIKE FROM YOU

Any technically relevant information that could change the analysis, for example: a model-version issue, an incorrect or unsupported chat template, a runtime or quantization problem, an error in our interpretation, or known limitations. Any response will be assessed against the preserved evidence. If it shows a conclusion is wrong, we will correct it. If it offers an explanation we cannot resolve, we will report it. If you disagree while the evidence stands, we will state that fairly. This right of reply does not give Mistral AI editorial control over the paper.

EVIDENCE

Attached is a 163 KB archive (Ministral-3-3B-Evidence-for-Mistral-20261006.zip) with the frozen evidence packets, raw requests and answers, claim-by-claim adjudication and file hashes. Start with README-FIRST.txt inside it. Archive SHA-256: 7a4d2c7280722e4859ef6b0d8fd53ff5caa3f380c222297e23d9154dc996. I can also send the full paper on request..

TIMING

We plan to publish after October 13, 2026. If we do not hear from you by then, the paper will state only that Mistral AI was notified on October 6, 2026 and that no response was received. If you need more time, tell me why.

Ministral-3-3B-Evidence-for-Mistral-20261006.zip
163 KB



Please reply to james.childers@aitoolworx.com.

Thank you,
James Childers
Founder, AIToolworx

