AEGIS FRONTIER

AEGIS Framework: Multilingual, Multimodal Information-Threat Classification

Research overview · November 2026 · AEGIS FRONTIER Research

Purpose and scope

The AEGIS research framework is designed to support structured analysis of information integrity and information-related harms. It approaches information assessment as a multi-class problem rather than collapsing all problematic content into a binary true-or-false decision. The initial research scope focuses on English and Chinese text, with a planned extension towards image, audio and video analysis.

Five-class information taxonomy

True information: information supported by available evidence and context. **Misinformation:** false or inaccurate information shared without necessarily established intent to deceive. **Disinformation:** false or manipulated information disseminated with deceptive intent. **Malinformation:** information that may be factually grounded but is used, exposed or framed to cause harm. **Hate speech:** expression that attacks, demeans or incites hostility against people on the basis of protected or identity characteristics. Operational definitions require careful annotation guidance and context-sensitive review.

Research architecture and development direction

The framework brings together data curation, annotation protocols, model development, evaluation and reliability analysis. AEGIS-Core is the principal research track. AEGIS-Vision explores visual and multimodal signals relevant to information integrity. The research roadmap includes dataset compatibility assessment, reproducibility, model performance and robustness, language coverage, explainability and human review.

Responsible use and limitations

Automated classification is probabilistic and context dependent. A model output should not be treated as a definitive finding of intent, truth, illegality or harmfulness. High-impact decisions require human review, documented evidence, appeal pathways and safeguards for freedom of expression, privacy and due process. Dataset provenance, language coverage, annotation disagreement and model uncertainty should be reported transparently.

Development status

This document is an introductory company research overview based on the AEGIS research programme description. It is not a peer-reviewed publication, independent validation report, product certification or claim of production deployment. Technical specifications and results should be added only when supported by released research artefacts and reproducible evaluations.

AEGIS FRONTIER — Research programme

Research priorities

- Multilingual NLP and information-threat classification across English and Chinese.
- Multimodal extension from text towards image, audio and video evidence.
- Dataset quality, annotation consistency, provenance and licensing.
- Reproducible training, benchmark protocol verification and model reliability.
- Responsible deployment, human oversight, privacy and freedom-of-expression safeguards.

Relationship to the AEGIS divisions

AEGIS Intelligence leads AI, data and information-integrity research. AEGIS Cyber contributes secure systems and threat-awareness perspectives. AEGIS Research supports scientific methods, publication and capacity building. AEGIS Robotics and AEGIS Microelectronics address physical and embedded technology systems. AEGIS Advisory connects technical work with governance, policy and institutional requirements.

Suggested citation

AEGIS FRONTIER SYSTEMS LIMITED (2026). *AEGIS Framework: Multilingual, Multimodal Information-Threat Classification*. Company research overview, November 2026.

Contact and corporate information

AEGIS FRONTIER SYSTEMS LIMITED · Plot 54, Muganzi Awongera, Makerere Kavule, Kampala, Uganda · aegisfrontier.com

© 2026 AEGIS FRONTIER SYSTEMS LIMITED. This overview may be shared with attribution. It does not disclose proprietary model weights, private datasets or confidential implementation details.