AI in a PDF

The visible page is a model card. The classifier is in /Keywords.
A 1,579-byte SMS spam model is not drawn on this page; it is metadata.

Task: SMS spam (labels: ham, spam)
Training data: 5,574 messages; 1,115 held out
Accuracy: 93.54% on the held-out split
Weights: int8 Complement NB, 160 words, 1,579 bytes
SHA-256 in /TinyAISHA256
Other containers: QR 2,158 B, favicon 1,519 B, WAV 1,627 B

Read it: pdf.html fetches this file, parses /Keywords and classifies
in the browser. No server and no weights endpoint.

Caveats: one 80/20 split, English SMS, bag-of-words, not an LLM.
The metadata is inspectable; a PDF is a lossless container, not a secret.