

Linda - a private, local commercial evaluation

14 days from first launch · free commercial evaluation · no key, card or automatic charge

Problem. Editors, schools and teams need to review possible AI assistance without sending confidential documents to an external detector. Linda helps prioritize a human review; it cannot prove authorship.

What you try. Windows 10/11, 64-bit; app 2.0.6.0 and model pack 1.3.4. Linda Loupe runs on CPU with a small model download; Linda Assay uses a larger ensemble, with 8 GB DirectX 12 GPU memory recommended. English, Polish and Russian; document checks, local highlighting and report exports. 4 GB Assay GPU compatibility is unvalidated.

Privacy. Install the desktop app for local inference. Texts are not uploaded; optional history is stored on your machine and can be disabled/deleted. Updates and licence services use network connections. Hosted demos send text to a server and are outside this confidential pilot. Share only aggregate results unless separately authorized.

Procedure. The 14-day permission begins with your first application launch; no registration or agreed date is needed. The app has no trial timer: observe the period yourself. For a team, appoint a coordinator internally. Record the installed versions, choose retention, and freeze the app/models during the trial. Prepare a labelled, source-independent set: target 100 human + 100 AI-origin texts per target language; separate edited-human, humanized-AI and mixed-span cases. Review labels before looking at detector output. Use `precise` and `sensitive` as distinct runs; do not select thresholds on this set. The coordinator keeps originals and labels locally and shares a summary, not the texts.

Measures. False accusations of humans, AI detection and misses, uncertain outcomes, import failures, full-check median/p95 time, hardware and workflow usefulness. Count each population separately and report confidence intervals. Proposed gates: human false-positive rate $\leq 2\%$, AI detection $\geq 75\%$, uncertainty $\leq 20\%$ in each main language cell, no critical workflow failures, correct identity/counts in exports. A 14-day pilot may only establish feasibility if adequate data are unavailable.

Risks and terms. No result is the sole basis for a decision about a person. Short inserts, human editing and non-native prose can fail; global confidence is not word coverage. Published numbers are developer measurements, not independent certification. Permission covers evaluation only, with no resale, hosted third-party service or operational screening. At the end, stop commercial operational use or choose the appropriate paid licence; there is no automatic conversion. Pilot feedback does not require transferring rights to confidential data.

Next step. Download and launch Linda to start the commercial evaluation. Contact lindapro.support@proton.me for help or aggregate feedback. Personal non-commercial use stays free without a time limit. Normal licences: Personal USD 10 once, Team USD 30 once, Organization USD 500/year (paid Organization renews automatically until cancelled). Polish checkout totals include tax; final tax/amount depend on location and billing details. Trust, benchmark and licensing documents accompany this brief.