

Model Card Template

The disclosure template every model we release is documented against.

DOCUMENT ID	LS-ETH-004
DELIVERABLE	E-04
VERSION	1.0
EFFECTIVE DATE	July 23, 2026
OWNER	Tech Lead
REVIEW CYCLE	Annual
STATUS	Adopted

1 How to Use This Template

Every model that LinkScape trains, fine-tunes, or releases publicly must ship with a completed Model Card based on this template. The Model Card is how we deliver on E-01 Responsible AI's transparency requirement: anyone using a LinkScape model should be able to understand what it does, what it was trained on, how well it works, and where it should not be trusted.

- Copy this template into the model's repository as MODEL_CARD.md and fill every field.
- Leave a field blank only if genuinely not applicable, and mark it 'N/A' with a one-line reason — do not delete fields.
- The CTO (owner) approves the Model Card before any public release.
- Update the card whenever the model is retrained, fine-tuned, or its intended use changes.

Fields shown in brackets like [fill in] are template placeholders. Complete them with real, model-specific values; never ship a card with unfilled brackets.

2 Model Details

Field	Value
Model name	[fill in]
Version / release date	[fill in]
Owner / point of contact	[fill in] (cto@linkscope.app)
Model type / architecture	[e.g., transformer encoder, fine-tuned LLM, diffusion]
Base model & license	[fill in]
Parameters	[fill in]
Associated project / paper	[link, e.g., ACL paper]
License of this release	[default: open-source license]
Repository	[GitHub URL]

3 Intended Use

3.1 Primary intended uses

[Describe the tasks and contexts this model is built for, e.g., NLP classification for research, educational demos in AI workshops, the AI-music project.]

3.2 Intended users

[Who should use this — LinkScope researchers, students, downstream open-source developers?]

3.3 Out-of-scope and prohibited uses

- [List uses the model is NOT validated for.]
- High-stakes decisions (medical, legal, financial, hiring) without human review.
- Any use targeting or profiling minors, consistent with LinkScope youth-safety commitments.
- Uses prohibited by E-01 Responsible AI or the base-model license.

4 Training Data

Field	Value
Datasets used	[names + links]
Data sources / provenance	[fill in]
Licenses	[fill in]
Contains personal data?	[Yes/No — if Yes, cite E-02 Data Privacy handling]
Consent basis	[per E-03 Research Ethics, if applicable]
Preprocessing / filtering	[dedup, toxicity filtering, tokenization]
Known gaps / underrepresented groups	[fill in]
Data collection dates	[fill in]

Dataset provenance and any data use agreements are governed by E-03 Research Ethics Guidelines, Section 4.

5 Performance Metrics

Report the metrics that matter for this model's intended use, with the evaluation setup so results are reproducible.

Metric	Dataset / split	Value	Notes
[e.g., Accuracy / F1]	[eval set]	[fill in]	[fill in]
[e.g., BLEU / ROUGE]	[eval set]	[fill in]	[fill in]
[Fairness / subgroup metric]	[subgroup]	[fill in]	[fill in]
[Compute cost of training]	[H100-hours]	[fill in]	[per F-03 / E-07]

- State the evaluation datasets, splits, random seeds, and hardware (H100 fleet) used.
- Report subgroup/disaggregated performance, not just aggregate scores.
- Include confidence intervals or variance across seeds where feasible.

6 Limitations

- [Known failure modes and input types where the model performs poorly.]
- [Distribution shift: domains, languages, or populations not represented in training.]
- [Robustness gaps: adversarial inputs, long context, rare classes.]
- [Hallucination / factuality limits for generative models.]

- [Staleness: knowledge cutoff or data-collection window.]

7 Bias & Fairness Evaluation

[Describe how the model was evaluated for bias, which protected/sensitive attributes were considered, and what was found.]

Dimension evaluated	Method	Finding	Mitigation
[e.g., gender]	[fill in]	[fill in]	[fill in]
[e.g., dialect / language variety]	[fill in]	[fill in]	[fill in]
[e.g., age / youth]	[fill in]	[fill in]	[fill in]

- Document residual biases that could not be mitigated.
- Cross-reference any bias incidents to E-08 AI Incident Response Protocol.
- For elevated-risk models, obtain Ethical Review Board review (E-05) before release.

8 Ethical Considerations & Governance

- Broader impacts: [potential societal benefits and harms of deployment].
- Environmental impact: training compute and estimated energy, per E-07 Environmental Impact.
- Human oversight: [where a human must stay in the loop].
- Feedback & reporting: users report harmful behavior to cto@linkscope.app; safety-critical issues escalate to the CFO per E-08.
- IP & licensing: release terms follow PP-02 IP Assignment; LinkScape defaults to open-source (E-06).

This card must be consistent with E-01 Responsible AI. A model may not be released publicly under the LinkScape name until its Model Card is complete and approved by the CTO.

Approval and Adoption

Adopted upon signature by the officers below. Pending ratification at Sprint 0 (see LS-REG-001).

Document Maintenance Log

Version	Date	Author	Changes
1.0	July 23, 2026	LinkScape Leadership	Initial release

LinkScape operates as a fiscally sponsored project of The Hack Foundation dba Hack Club, a 501(c)(3) nonprofit. Hack Club holds the charitable status and every dollar moves through Hack Club Bank.

