



Disarm Case Study: AI-Generated Propaganda Linking GERD Pride To "Assab/Sea Access" Narratives

Ethiopia, Aug to Oct 2025

Writer: Abel Wabella Sugebo

Title: AI and Disinformation Researcher | DISARM Analyst

Prepared by: Inform Africa/HaqCheck

About this case study

This document presents a DISARM case study on AI-generated and visual propaganda linking GERD pride to "Assab and sea access" narratives in Ethiopia (Aug to Oct 2025). It summarizes observed behaviors, content formats, and amplification pathways, grounded in an evidence log (E01–E14).

Report snapshot

- Scope: Ethiopia-focused, online information environment
- Window: August 19, 2025 to October 24, 2025
- Platforms: X, YouTube
- Languages: English, Amharic, Afaan Oromo, Mixed
- Core artifacts: AI visuals, montage videos, clipped longform interviews
- Outputs: Timeline, DISARM mapping, findings, recommendations

Contact and attribution

Inform Africa (HaqCheck / Research Unit)

Email: haqcheck@gmail.com

Phone: +1(202)-993-8087

Website: <https://informafrica.net/> and <https://haqcheck.org/>

Acknowledgement

This work was carried out in the context of the Africa Digital Rights Fund with support from CIPESA.

© 2025 Inform Africa. All rights reserved

1) Executive Summary

Between August and October 2025, a cluster of online posts circulated content that framed Ethiopia's national pride and development milestones, especially around the Grand Ethiopian Renaissance Dam (GERD), as justification for future Red Sea access via Eritrea's Assab port. The content included AI-generated or synthetic-looking visuals and patriotic montage videos, as well as clipped segments from a long-form YouTube interview that were repackaged into short, highly shareable clips on X and amplified with hashtags and explicit calls for Assab to "come back" to Ethiopia (E01, E02, E06, E11 → E10 → E09).

Using the DISARM framework, this case study documents observable behaviors across narrative development, content packaging, and amplification. The evidence set includes three primary synthetic or visual artifacts (E01, E02, E06), a repackaging and amplification chain from YouTube to X (E11 → E10 → E09), and timeline anchors around GERD inauguration coverage (E04, E05). The report also includes upstream discourse context on "geographic prison" and "sovereign access to port" framing (E12, E13), plus a counter-messaging artifact reflecting contestation from Eritrean state-aligned media (E14).

Key findings:

- Synthetic or AI-styled visuals normalized an aspirational "Assab future" narrative using tourism-like imagery and patriotic symbolism (E06, E02).
- A GERD celebration video served as an emotional bridge from development achievement to "Assab next" messaging (E01).
- Longform political commentary was clipped and repackaged into short videos for higher shareability, then amplified through quote-sharing, tags, and hashtags (E11 → E10 → E09).
- Media and influencer amplification broadened reach and reframed the issue as a matter of security and national destiny, with multilingual resonance, including Afaan Oromo audio in at least one widely viewed clip (E08).

This pattern carries information integrity and digital rights risks, including inflammatory cross-border narratives, heightened polarization, and the potential for harassment or incitement-adjacent discourse. The report concludes with practical recommendations for platforms, media, and civil society to reduce the virality of decontextualized clips, improve the handling of synthetic media, and strengthen local-language risk-review capacity.

2) Why This Matters For Digital Rights

This incident matters because it intersects with information integrity, cross-border political tensions, and the safety of online civic space. The observed content packages a sensitive geopolitical issue into emotionally resonant, easily shareable formats, including synthetic-looking visuals and decontextualized clips, which increases the risk of rapid polarization and misperception (E01, E02, E06, E11 → E10 → E09).

Key digital rights and integrity risks:

1. **Right to access reliable information**

Synthetic or AI-styled visuals can make contested narratives feel "real" through imagery rather than evidence. Tourism-like framing and patriotic montage aesthetics can normalize claims without accountability, increasing the risk of false certainty and rumor-driven belief (E06, E02, E01).

2. **Freedom of expression and safety in online civic space**

Escalatory framing around sovereignty and national destiny can harden positions, raise hostile speech, and increase pressure on journalists, researchers, and fact-checkers who attempt to correct misleading claims or encourage de-escalation (E08, E14).

3. **Cross-platform "clip migration" that strips context**

When long-form political discussion is clipped into a short video, the message can shift from analysis to a mobilizing slogan. This increases virality while reducing nuance, and can contribute to misinterpretation or weaponization of quotes (E11 → E10 → E09).

4. **Local-language resonance and uneven moderation capacity**

The presence of Afaan Oromo audio in a widely viewed clip underscores the importance of local-language risk review and contextual moderation. Without adequate capacity, harmful narratives can spread faster in linguistically specific spaces (E08).

5. **Contestation can deepen polarization.**

Counter-messaging and state-aligned responses can rapidly frame the discourse as illegitimate or irrational, reinforcing "us vs them" dynamics and reducing space for evidence-led discussion (E14).

Because these risks can escalate quickly, documenting the behaviors with a shared taxonomy (DISARM) supports responsible analysis and practical mitigation, without requiring attribution claims beyond what the evidence shows (E01–E14).

3) Research Questions

1. What narratives linked GERD pride to Assab or sea-access claims, and how were they framed emotionally? (E01, E02, E06, E08, E12, E13)
2. What content formats were used (AI-styled visuals, montage videos, clipped interviews), and how were they packaged to increase shareability? (E01, E02, E06, E11 → E10 → E09)
3. What amplification behaviors were observable on X (hashtags, tagging, quote-sharing), and how did they shape reach and framing? (E09, E10, E08, E01)
4. How did the narrative move from YouTube long-form to X short-form clips and onward to amplification? (E11 → E10 → E09)
5. What counter-messaging or contestation appeared, and how did it frame the issue? (E14)
6. How do these behaviors map to the DISARM stages based solely on the collected evidence? (E01–E14)

4) Methodology And Ethics

We collected relevant artifacts from X and YouTube and documented each item in an evidence log with a unique ID (E01–E14). For each artifact, we preserved the URL and captured screenshots or PDF copies. Where possible, we also created third-party archived links. For social posts, we recorded engagement metrics at the time of capture when those metrics were visible.

We verified artifacts by preserving the account name, post date, and content in screenshots, and by using archive links when available. For the YouTube source, we saved a timestamped link to the segment that was later clipped and shared on X (34:04), to anchor provenance and reduce the risk of misquoting (E11).

We designed the documentation process to reduce harm. The report avoids reproducing inflammatory slogans as headings and does not republish full clips. It focuses on observable behaviors and spread patterns. The DISARM framework is used to classify behaviors and techniques, not to claim attribution beyond what the evidence supports. Counter-messaging content is treated as context and labeled accordingly (E14).

5) Incident Narrative And Timeline

This case builds on a longer-running public discourse about Ethiopia's sea access. Prior reporting and analysis show that framing such as "geographic prison" and calls for "sovereign access to a port" existed before the 2025 incident window and provided language that could later be remixed into persuasive online content (E12, E13).

On September 9, 2025, Ethiopia inaugurated the Grand Ethiopian Renaissance Dam (GERD). This high-visibility event created a surge in national pride and a timely context for messages that connected development achievement to future geopolitical ambition (E04, E05).

Within the observed period, multiple formats were used to normalize and amplify Assab and sea-access narratives. On August 19, 2025, an AI-styled port image framed Assab as a desirable destination, using a casual "vacation" caption and hashtags to reduce controversy and increase shareability (E06). On March 27, 2025, a patriotic montage video circulated with slogans such as "NO MORE LANDLOCKED!" and #ASSEB, showing that the narrative and visual trope had already been circulating earlier in the year (E02).

On September 6, 2025, a widely viewed clip framed Red Sea access primarily in terms of peace and national security, with audio in Afaan Oromo, indicating multilingual resonance and the importance of local-language context in risk assessment (E08). On September 8, 2025, a GERD celebration video shared by a prominent account included on-screen messaging such as "GERD 2.0 Coming Soon!" and "See you all in Assab!", serving as a bridge from development pride to an explicit Assab message (E01).

A distinct packaging and amplification pathway was also observed from long-form video to short-form clips. A YouTube interview episode was later used as a source segment at 34:04, which was clipped and reposted on X with "No more landlocked Ethiopia" framing and tagging behavior (E11, E10). The clip was then quote-shared the following day with additional hashtags and explicit "Assab should come back" framing, illustrating how the duplicate underlying content can be sharpened into more mobilizing language during amplification (E10, E09).

Finally, contestation and counter-messaging appeared in the same period. An Eritrean state-aligned outlet published a critical response that framed Ethiopia's "Assab obsession" as beyond rationality, reflecting polarization and hardening positions within the discourse (E14).

6) Disarm Mapping

This section maps observable behaviors in the incident to DISARM stages. The mapping documents what was seen in the evidence and does not claim attribution to a single coordinating actor.

DISARM stage	Observed behavior	Evidence IDs	Confidence
Develop narratives	Sea access framed as destiny, peace, and national security; slogans normalize entitlement to Assab	E08, E09, E10, E12, E13	High
Develop content	Synthetic or AI-styled visuals and patriotic montage formats make Assab appear inevitable, desirable, or celebratory	E06, E02, E01	High
Maximise exposure	Quote-sharing, hashtags, tagging, and media amplification increase reach and reduce friction for resharing	E09, E10, E08, E01	High
Establish assets	High-visibility accounts and media pages distribute and legitimize clips	E01, E08, E10	Medium
Drive offline action	Explicit calls are present, but evidence is insufficient to claim offline mobilization outcomes	E09, E02	Low to Medium
Counter-messaging	Contestation reframes the discourse as irrational and illegitimate	E14	Medium

Table 1. DISARM mapping of observed behaviors and evidence (E01–E14)

7) Findings

First, synthetic or AI-styled visuals helped normalize Assab and sea-access claims by presenting them as aspirational and emotionally positive. The "vacation destination" framing and port imagery reduce perceived controversy and increase shareability, while montage aesthetics and slogans reinforce the sense of inevitability and national pride (E06, E02).

Second, GERD-related pride was used as a persuasive bridge to "Assab next" messaging. The GERD celebration video, which includes "See you all in Assab!", connects a development achievement to a territorial aspiration in a way that can feel natural and forward-looking to audiences (E01).

Third, long-form political discussion was repackaged into short clips that are easier to share and reframe. The observed pathway from a YouTube source segment to X sharing and quote-sharing shows how context can be reduced while slogans and hashtags sharpen the message during amplification (E11, E10, E09).

Fourth, amplification and legitimization were strengthened through media-style posting, tagging behavior, and multilingual resonance. The presence of Afaan Oromo audio and security-oriented framing suggests the narrative is being tailored for local resonance and credibility, not only for entertainment or nationalism (E08).

Finally, contestation emerged in parallel. Counter-messaging framed the discourse as irrational and illegitimate, thereby indicating polarization and reducing the space for evidence-led discussion (E14).

8) Recommendations

Platforms should strengthen labeling and friction for synthetic or AI-styled media in high-risk geopolitical discourse, especially when visuals imply territorial change or security escalation. They should also reduce the viral advantage of decontextualized political clips by promoting context links to source videos and adding friction to quote-sharing when a clip is trending without its original context (E06, E01, E11 → E10 → E09). Local-language risk-review capacity should be improved so that narratives spreading in Afaan Oromo or other languages receive timely, accurate moderation (E08).

Media and civil society actors should prebunk standard frames used in the sea-access discourse, including how aspirational imagery and patriotic montage can manufacture inevitability. Rapid evidence logging, archiving, and transparent documentation practices should be institutionalized to preserve fast-moving narratives before they are deleted or edited. Journalist and researcher safety planning should also be strengthened, as polarizing geopolitical narratives can lead to harassment of those who correct misleading claims or urge de-escalation (E06, E01, E14).

Funders and partners can support this work by investing in synthetic media literacy, verification training, and monitoring capacity for clip migration across platforms. Support for evidence-led public explainers and cross-border de-escalation information products can help reduce the dominance of inflammatory narratives while protecting civic space.

9) Limitations

Engagement metrics are point-in-time and can change quickly. This report records only what was visible at the time of capture. Some content may be edited, restricted, or deleted, which can limit access to the earliest instances; we mitigate this by capturing screenshots, PDFs, and archived links where possible.

This case study does not attribute intent or coordination to a specific actor. It documents observable behaviors mapped through DISARM, strictly based on the evidence collected (E01–E14). Context sources are used to describe background discourse and contestation, and do not serve as proof of factual claims.

10) Annex

Annex A: [Evidence log and preserved artifacts \(E01–E14\)](#)

Annex B: [Timeline table \(with evidence IDs\)](#)

Annex C: [Expanded DISARM mapping table \(optional, if requested by reviewers\)](#)