Sudan Humanitarian AI Insights

Technical Dossier for an AI-Assisted Humanitarian Needs Dashboard

Purpose: Document the design, data model, analytical method, visual outputs, ethical guardrails, and deployment approach for a first-phase humanitarian AI dashboard focused on Sudan.

Creator: Shireen El Sabea
LinkedIn: linkedin.com/in/shireenalsabea
Email: sabeashireen@gmail.com

Version: June 2026

Responsible AI for community feedback, needs analysis, and decision-support dashboards

Contents

1 Executive Brief

Sudan Humanitarian AI Insights is a Streamlit-based analytics dashboard designed by Shireen El Sabea to demonstrate how AI-assisted data analysis can support humanitarian prioritization. The dashboard converts public-context-informed data and synthetic community feedback into hotspot scores, sector pressure visuals, local NLP triage, and downloadable AI Action Briefs.

This dossier expands the prototype into a technical case-study package. It documents the problem framing, data structure, analytics method, visual outputs, system architecture, ethical constraints, deployment options, and future roadmap.

The project is not presented as an operational humanitarian system. It is a responsible public demonstration using synthetic data and transparent scoring. Contextual references are drawn from recognized sources such as IOM DTM, HDX, OCHA, ReliefWeb, and ACLED [1, 2, 3, 5, 6].

Table 1: Dossier Metadata

Field	Description
Project title	Sudan Humanitarian AI Insights
Creator	Shireen El Sabea
Primary use case	AI-assisted humanitarian hotspot analysis and action briefing
Technology stack	Streamlit, Pandas, NumPy, Plotly, Matplotlib, local CSV datasets
Data status	Synthetic demonstration data calibrated around public humanitarian themes
Privacy posture	No personally identifiable information, no private field data, no paid API dependency
Primary audience	NGOs, humanitarian data teams, donors, social-impact employers, and portfolio reviewers

2 Humanitarian Context

Sudan remains one of the most complex humanitarian operating environments in the world. IOM's Displacement Tracking Matrix is used to monitor displacement and mobility patterns, while HDX republishes IOM DTM data for humanitarian data users [1, 2]. OCHA's 2025 planning documents describe a large-scale response environment requiring safe, timely, principled, and gender-responsive assistance [3]. Conflict event monitoring from ACLED provides an additional lens for understanding violence, access limitations, and civilian risk [6, 7].

For a public portfolio project, the ethical challenge is clear: it is valuable to demonstrate analytical capability, but it is not appropriate to expose real community feedback or personal data. The dashboard therefore uses a synthetic scenario dataset, while citing official sources only for context.

3 Problem Definition

Humanitarian teams often receive fragmented signals from multiple channels: displacement updates, partner reports, coordination meetings, community feedback, field assessments, and access reports. These signals can be hard to interpret quickly when caseloads are large and resources are constrained.

The prototype addresses five analytical questions:

1. Which states appear most urgent under a transparent composite score?
2. Which sector pressures are driving the hotspot ranking?
3. How can synthetic unstructured feedback be classified into humanitarian sectors?
4. How can feedback and displacement trends support planning discussions?
5. How can the dashboard convert analysis into a short decision-support brief?

4 Solution Overview

The dashboard is structured around a data-to-decision workflow. The core product idea is not merely to show charts, but to convert signals into a practical response discussion.

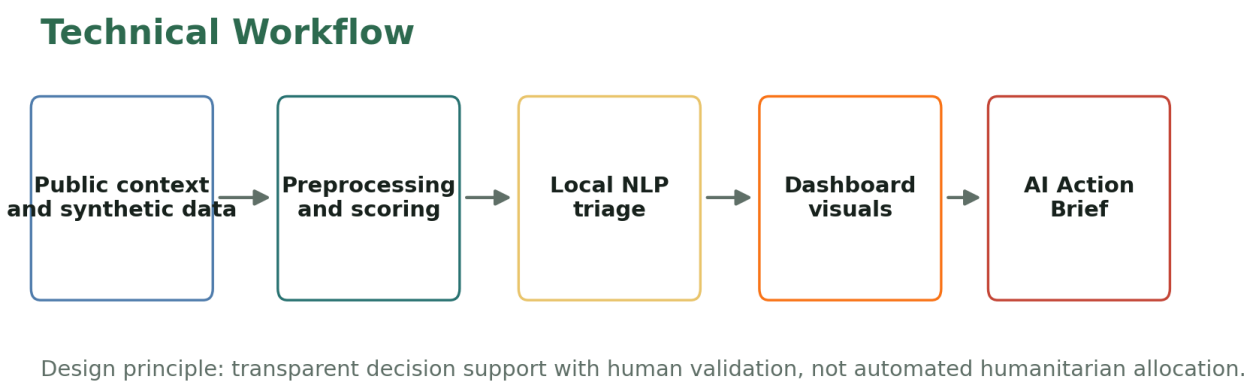

Figure 1: Technical workflow from data ingestion to AI Action Brief generation.

4.1 Dashboard Modules

Table 2: Dashboard Modules and Analytical Purpose

Module	Purpose
Operations View	Presents key performance indicators, a geographic priority view, ranked hotspots, and a sector pressure matrix.
Feedback Intelligence	Demonstrates local NLP-style classification of synthetic community feedback by sector, urgency, sentiment, and confidence.
Forecast and Actions	Shows monthly displacement and feedback trends, then converts priority scores into an intervention queue.
Hotspot Brief	Generates an AI Action Brief for a selected hotspot, including drivers, next steps, related feedback signals, and downloadable text.
Data Notes	Documents public context sources, synthetic data limitations, and ethical guardrails.

5 Data Architecture

The prototype uses three local CSV files. This makes the project reproducible, easy to inspect, and suitable for deployment without external credentials.

Table 3: Dataset Inventory and Role in the Dashboard

Dataset	Record Grain	Role
`state_needs.csv`	Sudan state	Stores IDP estimates, household estimates, child share, food pressure, WASH gap, health gap, access constraint, protection alerts, and funding gap.
`monthly_trends.csv`	Month	Stores displacement trend and synthetic feedback pressure by sector.

Dataset	Record Grain	Role
`feedback_samples.csv`	Feedback message	Stores synthetic SMS, WhatsApp, and voice-note examples with sector, urgency, state, and population group labels.

5.1 Data Governance Notes

- All feedback records are synthetic examples.
- No household names, phone numbers, addresses, registration IDs, or precise sensitive locations are stored.
- Public sources are used as contextual anchors, not as direct operational claims.
- Any future real deployment would require data sharing agreements, consent processes, access controls, and safeguarding review.

6 Analytical Model

6.1 Composite Priority Score

The priority score combines normalized displacement pressure with sector and operational pressure indicators. The score is intentionally simple and explainable:

$$\begin{aligned} Priority = 100 \times (&0.26D + 0.20F + 0.15W + 0.12H \\ &+ 0.12A + 0.10P + 0.05G) \end{aligned} \tag{1}$$

where D is normalized displacement volume, F is food insecurity pressure, W is WASH gap, H is health gap, A is access constraint, P is normalized protection pressure, and G is funding gap.

Table 4: Priority Score Weighting

Driver	Weight	Rationale
Displacement volume	0.26	Large caseloads increase the scale of response required.
Food insecurity pressure	0.20	Food access is a core lifesaving concern.
WASH gap	0.15	Water and sanitation gaps can quickly create health and protection risks.
Health gap	0.12	Health service gaps indicate limited treatment and referral capacity.
Access constraint	0.12	High access constraints reduce response feasibility and require coordination.
Protection alerts	0.10	Protection issues require rapid safeguarding review.
Funding gap	0.05	Funding pressure affects the ability to respond at scale.

6.2 Local NLP Feedback Triage

The prototype uses a transparent keyword-based classifier. It maps synthetic messages into sectors such as Food Security, WASH, Health, Shelter, Protection, Child Protection, and Registration.

This is not a substitute for a trained language model. It is used to demonstrate the workflow:

1. Receive unstructured text.

2. Match sector-specific terms.
3. Detect urgency terms such as unsafe, sick, children, blocked, flood, malaria, or ran out.
4. Assign sector, urgency, sentiment, confidence, and matched terms.
5. Surface the classified signal in the dashboard.

7 Visual Analysis

7.1 Hotspot Ranking

The first visual output ranks the highest-priority states. The model identifies South Darfur and North Darfur as the highest-scoring hotspots in the demonstration dataset.

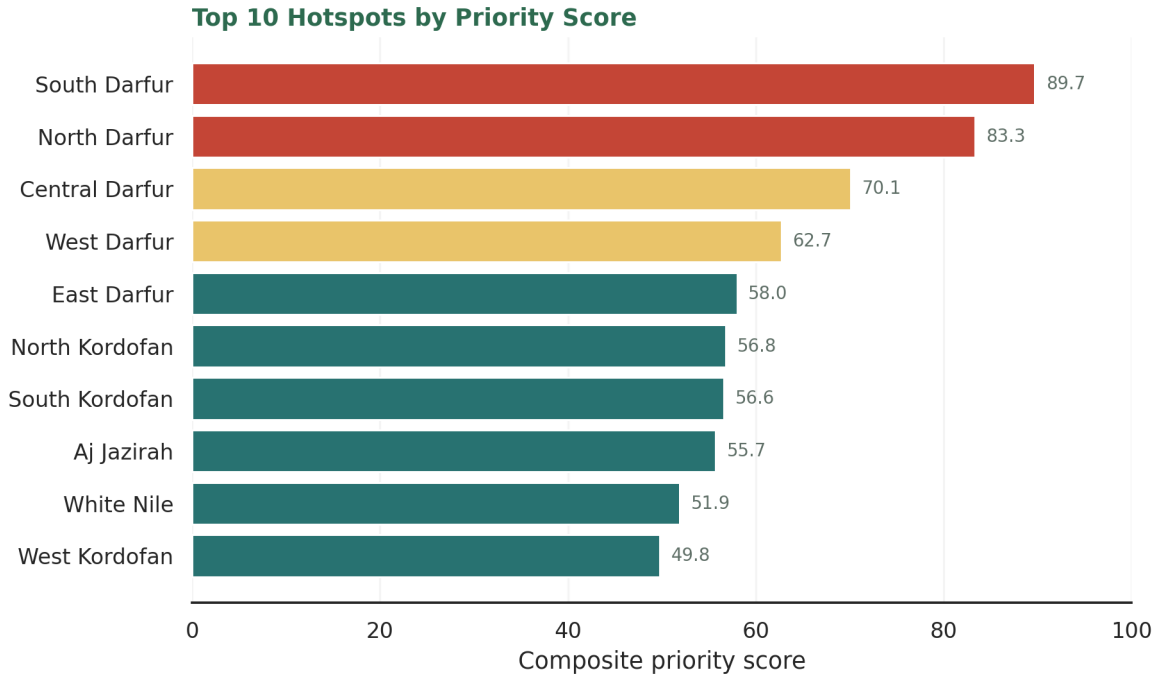

Figure 2: Top 10 hotspot ranking by transparent composite score.

Table 5: Top Five Hotspots in the Demonstration Dataset

State	Score	IDPs	Food	Access	Level
South Darfur	89.7	1.65M	94%	91%	Critical
North Darfur	83.3	1.41M	93%	88%	Critical
Central Darfur	70.1	0.79M	92%	89%	High
West Darfur	62.7	0.51M	91%	86%	High
East Darfur	58.0	0.54M	89%	78%	Elevated

7.2 Sector Pressure Matrix

The heatmap explains why locations rank highly. Rather than hiding the model behind a single score, it shows the underlying pressure drivers.

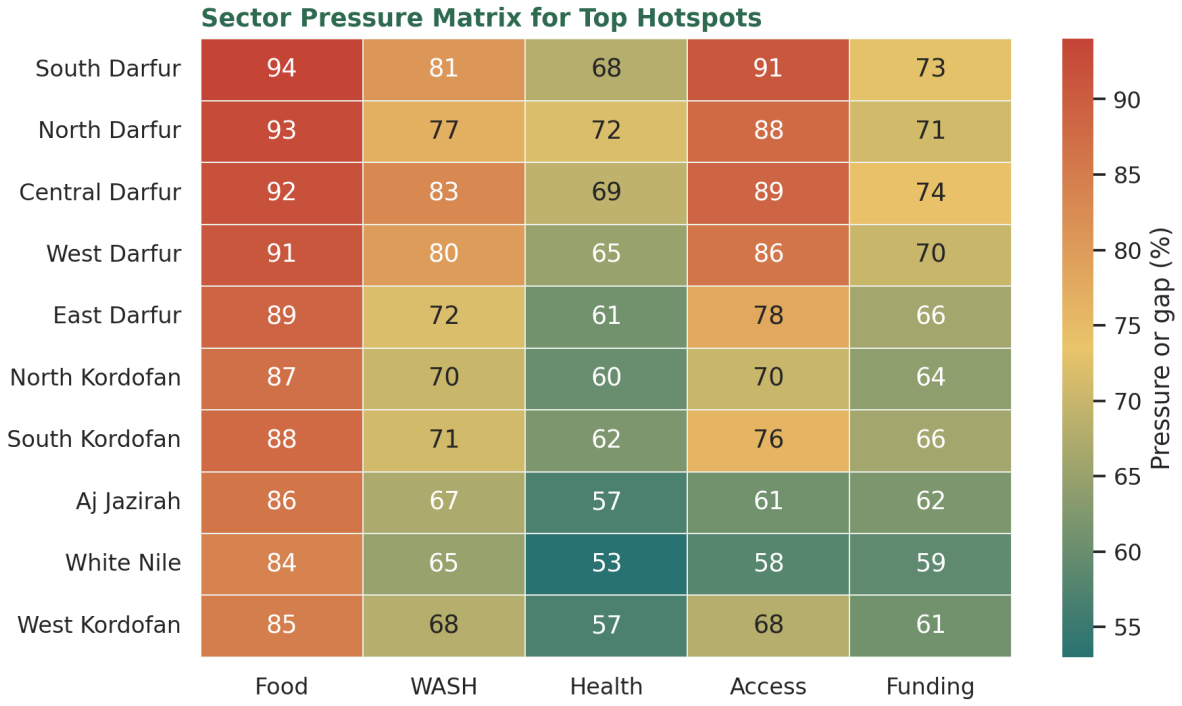

Figure 3: Sector pressure matrix for the highest-ranked states.

7.3 Trend and Forecasting View

The monthly trend figure shows how displacement estimates and synthetic feedback volume can be viewed together. In a real deployment, this view would need direct validation against current data pipelines and field assessments.

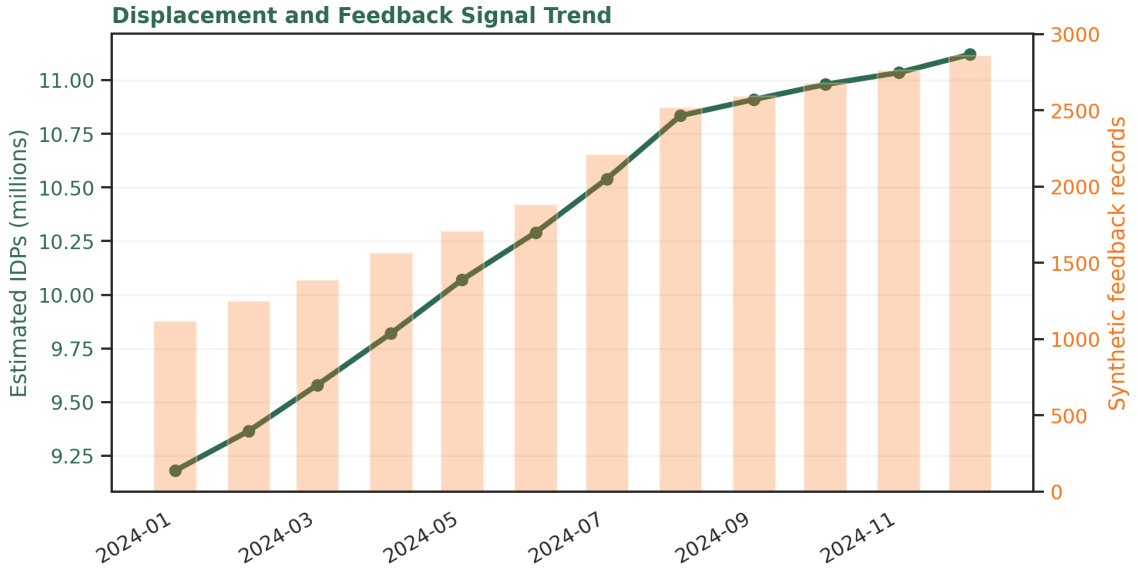

Figure 4: Monthly displacement and synthetic feedback signal trend.

7.4 Feedback Urgency

The feedback urgency chart demonstrates how community messages can be organized by sector and triage status.

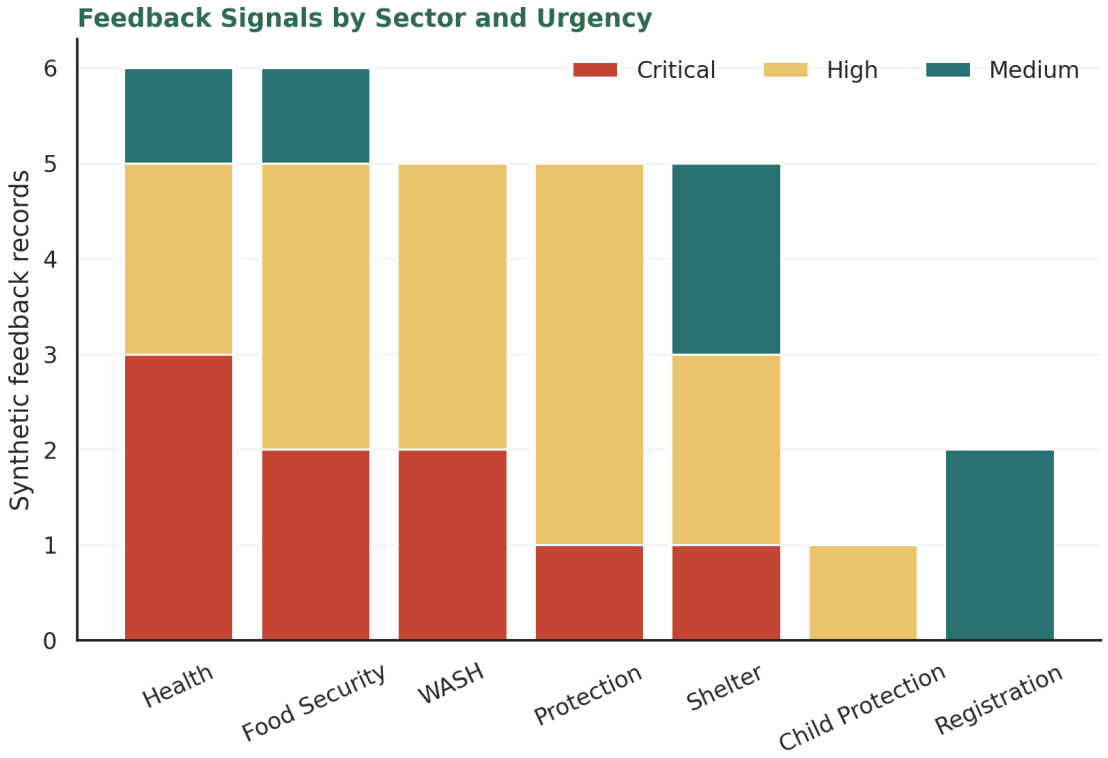

Figure 5: Synthetic feedback records grouped by sector and urgency.

8 AI Action Brief Design

The AI Action Brief is the most important employer-facing component because it converts dashboard outputs into a decision-support artifact.

Table 6: AI Action Brief Content Model

Brief Section	Purpose
Priority summary	Shows state, priority level, score, IDP estimate, and child estimate.
Top drivers	Identifies why the location was selected.
Recommended posture	Explains how the dashboard output should be used in planning.
Next validation steps	Keeps human review and field confirmation central.
Related feedback signals	Links selected state to synthetic community feedback messages.
Downloadable text	Makes the output portable for coordination or donor notes.

9 Technical Implementation

9.1 Application Stack

Table 7: Technical Stack

Layer	Implementation
Frontend application	Streamlit multipage-style interface with tabs, filters, metrics, charts, and download controls.
Data processing	Pandas and NumPy for loading CSV files, normalizing indicators, scoring hotspots, and preparing summaries.
Visualization	Plotly in the live dashboard; Matplotlib and Seaborn for report figures.
Branding	Project logo embedded in the app, static landing page, and PDF dossier.
Deployment assets	Streamlit app for interactive hosting, plus static `docs/index.html` for GitHub Pages or Vercel landing pages.

9.2 Repository Structure

```
AI_NGO/
  app/app.py
  data/state_needs.csv
  data/monthly_trends.csv
  data/feedback_samples.csv
  assets/impact_bridge_logo.png
  docs/index.html
  reports/study_report.tex
  reports/generate_figures.py
  reports/figures/
  output/pdf/study_report.pdf
  requirements.txt
```

10 Ethical and Operational Guardrails

Humanitarian AI systems can create harm when they obscure uncertainty, expose sensitive populations, or automate decisions that require human judgment. The prototype therefore follows these guardrails:

- **No PII:** The public demo uses no names, phone numbers, documents, IDs, or household-level data.
- **Synthetic feedback:** Community messages are demonstration samples only.
- **Explainable scoring:** Users can see the drivers behind the score.
- **Human validation:** The dashboard supports field validation, not automated allocation.
- **Protection sensitivity:** Protection signals should be handled through safeguarding channels.
- **Bias control:** Future model versions would require local language review, representative data, and continuous error monitoring.

11 Limitations

Table 8: Current Limitations and Mitigation Direction

Limitation	Mitigation Direction
Synthetic data only	Replace with aggregate, approved, non-PII datasets under data governance agreements.
Keyword-based NLP	Move toward validated multilingual NLP models with human review and error analysis.
Simple linear forecast	Add time-series validation, confidence intervals, and scenario assumptions.
No authentication layer	Add role-based access control before any sensitive deployment.
No live humanitarian API integration	Add scheduled ingestion only after source licensing and reliability review.

12 Deployment Architecture

The interactive dashboard is a Streamlit application. Streamlit Community Cloud is the recommended target for the live app. GitHub Pages and Vercel are suitable for the static landing page, not the interactive dashboard server.

1. Push the repository to GitHub.
2. Deploy `app/app.py` on Streamlit Community Cloud.
3. Publish `docs/index.html` through GitHub Pages or Vercel.
4. Replace the landing-page dashboard placeholder link with the live Streamlit URL.

13 Roadmap

Table 9: Suggested Roadmap

Phase	Focus	Output
Phase 1	Public proof of concept	Current dashboard, static landing page, LinkedIn post, and PDF dossier.
Phase 2	Data validation	Replace synthetic indicators with approved aggregate public data extracts.
Phase 3	Model improvement	Add multilingual NLP, uncertainty labels, and model evaluation tables.
Phase 4	Operational pilot	Add user roles, audit logs, data protection review, and partner validation workflow.
Phase 5	Portfolio expansion	Record a walkthrough video and publish a technical blog post.

14 Conclusion

The Sudan Humanitarian AI Insights dashboard demonstrates a practical, responsible way to communicate the potential of AI in humanitarian work. Its strongest value is not the score alone, but the complete workflow: data is structured, feedback is triaged, hotspots are explained, and action briefs are generated with explicit human validation steps.

This makes the project suitable for LinkedIn presentation, portfolio review, and early conversations with NGOs or social-impact employers.

15 Call to Action

Shireen El Sabea is open to collaboration with NGOs, humanitarian data teams, donors, and social-impact employers interested in responsible AI for community feedback, needs analysis, and decision-support dashboards.

- LinkedIn: https://www.linkedin.com/in/shireenalsabea/
- Email: sabeashireen@gmail.com

References

[1] International Organization for Migration. *Sudan Displacement Tracking Matrix*. https://dtm.iom.int/sudan

[2] Humanitarian Data Exchange. *Sudan IOM Displacement Tracking Matrix (DTM) from API*. https://data.humdata.org/dataset/sdn-iom-dtm-from-api

[3] United Nations Office for the Coordination of Humanitarian Affairs. *Sudan Humanitarian Needs and Response Plan 2025 Executive Summary*. https://www.unocha.org/publications/report/sudan/sudan-humanitarian-needs-and-response-plan-2025-executive-summary-december-2024-enar

[4] OCHA Humanitarian Action. *Sudan Humanitarian Response Plan 2025*. https://humanitarianaction.info/plan/1220

[5] ReliefWeb. *IOM Sudan Displacement Tracking Matrix: Sudan Mobility Update 16*. https://reliefweb.int/report/sudan/iom-sudan-displacement-tracking-matrix-dtm-sudan-mobility-update-16-publication-date-23-march-2025

[6] Armed Conflict Location and Event Data Project. *Sudan Country Data and Analysis*. https://acleddata.com/country/sudan

[7] Armed Conflict Location and Event Data Project. *How to Interpret ACLED's Fatality Records in Sudan*. https://acleddata.com/methodology/how-interpret-acleds-fatality-records-sudan