 Received 12-04-2026; Accepted 10-05-2026; Published 17-05-2026

Open Source Intelligence Methods for Identifying Humanitarian Needs in the Context of the War in Ukraine: A News-Driven Pipeline Approach

Yehor Holyk yehor.holyk@nure.ua
Kharkiv National University of Radio Electronics,
Nauky Ave. 14, Kharkiv, 61166, Ukraine

Dmytro Sytnikov dmytro.sytnikov@nure.ua
Kharkiv National University of Radio Electronics,
Nauky Ave. 14, Kharkiv, 61166, Ukraine

Corresponding Author: Yehor Holyk

Abstract

Humanitarian responders in active-conflict areas require timely indications of when civilian needs will arise, but unclassified information, situation reports, and conflict event streams lack clarity. This study provides a replicable OSINT pipeline for transforming these data into a cluster-resolved Humanitarian Need Pressure Index (HNPI) and short-term alerting system. The process is comprised of open data collection, clear need categorization, weighting by severity, geoparsing, aggregation by week per region, lagged regression, and ethical review. In a Ukraine pilot from 2 September 2024 - 31 August 2025, 17.511 news articles and 37.978 geolocated conflicts generated prominent clusters for health, shelter, and displacement needs, which were most pronounced in Donetsk, Kyiv and Kharkiv regions. HNPI was modeled using ridge regression one to four weeks ahead, and the MAE was around 5.7 points, while the performance was better than persistence, and the AUC of the surge alert was approximately 0.91. As the target is the HNPI, the results reflect inherent predictability not independent humanitarian need validation.

Keywords: Open-source intelligence, Humanitarian assistance,Early-warning systems,Natural language processing,Topic modelling, Sentiment analysis, Geoparsing, Ukraine war.

1. INTRODUCTION

Due to the war in Ukraine, there is an increasingly urgent need to provide humanitarian aid to the civilian population. Incidents occur daily across the country, creating a need for different forms of assistance: from the repair of energy infrastructure to the provision of shelter. It is difficult to analyze and respond to these events in a timely manner, even though formal needs assessments remain the primary approach [1, 2].

Citation: Yehor Holyk and Dmytro Sytnikov. Open Source Intelligence Methods for Identifying Humanitarian Needs in the Context of the War in Ukraine: A News-Driven Pipeline Approach. Advances in Knowledge-Based Systems, Data Science, and Cybersecurity ; Research. 2026;3(2).26.

Page No. 518-532.

Signals related to damage, displacement, needs and constraints in OSINT sources, such as news, government statements, conflict event databases, and social media, are near real-time. However, they are also noisy, multilingual, manipulable, and ethically delicate, and therefore, humanitarian application should be reproducible and bounded by ethical guidelines [3].

In order to resolve this problem we introduce the pipeline that transforms open-source and humanitarian data into the indicator, which can potentially define humanitarian needs and locations. The pipeline consists of transparent NLP modules instead of black-boxed end-to-end approaches, links its results to the humanitarian cluster approach [2] and uses the lagged regression framework inspired by conflict security early warning systems [4, 5].

We propose a seven-stage system architecture based on OSINT-driven event analysis methods, the Humanitarian Need Pressure Index (HNPI), an early warning model incorporating a time lag and an outbreak indicator, and a clear description of validation, reliability, and ethical constraints.

2. RELATED WORK

2.1 OSINT in Armed Conflict and Humanitarian Action

OSINT now allows for documentation, accountability, and situational awareness, even for civilian harm tracking in Ukraine [6]. But the same literature advises us that open-source analysis is fraught with risk of re-identification and the mosaic effect [3]. Since this project will employ OSINT strictly in aggregate form, it takes the first-aggregation stance.

2.2 Open-Source Humanitarian and Conflict Data

Humanitarian data analysis is currently using open data sources that are heterogeneous in nature: ACLED geolocated conflict data [4, 7], OCHA HDX/Reliefweb [8, 9] and displacement portals like IOM DTM/UNHCR [10, 11]. These are being used by our pipeline as covariates.

2.3 Natural Language Processing During Crisis Events

Based on crisis informatics literature, it has been found out that noisy texts may be represented in operational categories through the use of lexicons, labeled text collections, and domain adapted models [12, 13, 14, 15]. As this study focuses on conflict news and the humanitarian cluster taxonomy, lexicons, topics [16, 17], sentence embedding [18, 19, 20], sentiment [21], geoparsing [22, 23] and event-data concepts will be combined [24].

2.4 Positioning

In comparison to previous research, the pipeline brings together humanitarian text classification, analysis of events of conflict and early warning using time series analysis in a single auditable process, building upon the work done by the same authors on topic modelling and news regression for Ukraine [25, 26].

It should be noted that the new position clearly states that the innovation introduced does not include a new architecture of deep learning. On the contrary, it consists in the implementation of a workflow of OSINT analysis that combines humanitarian-cluster natural language processing, geoparsing of weekly region-level aggregation, HNPI defined, benchmarking in relation to OCHA/JIAF HNRP 2025, and audit of short-lead warnings. The proposed approach can be applied to on-field evaluation, not as a replacement of it. At the same time, the general framework can be transferred to another crisis situation.

3. METHOD: A SEVEN-STAGE OSINT PIPELINE

FIGURE 1, shows the seven stage process from acquisition of open source information to decision support through aggregation. The steps are modular in nature to enable tracing of a high region-cluster-week score to data.

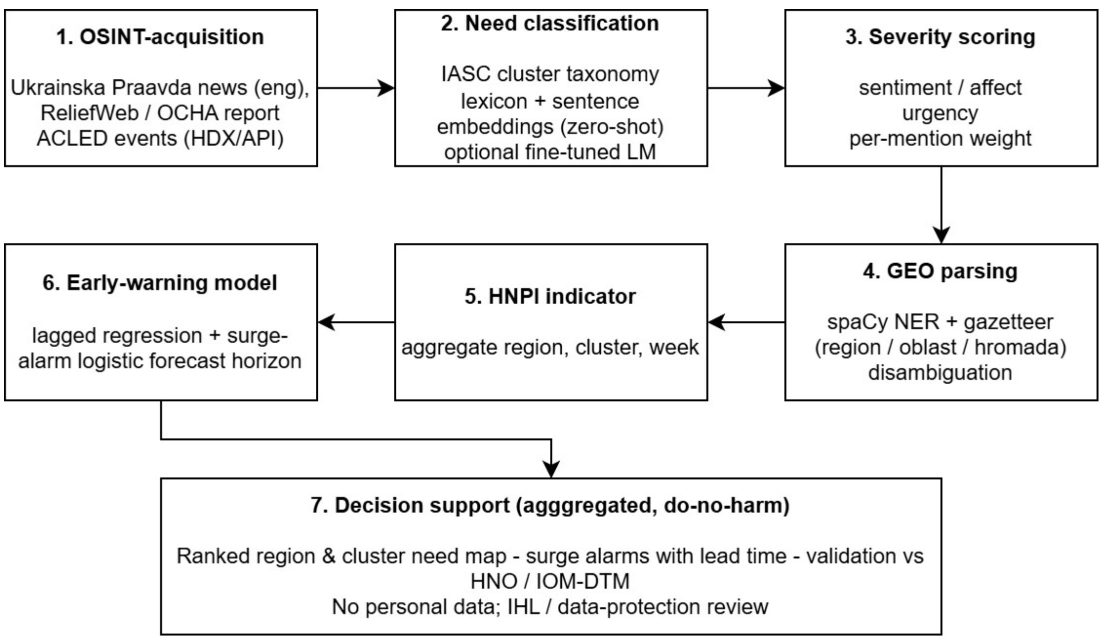

Figure 1: Architecture of the OSINT humanitarian early-warning pipeline.

3.1 Seven-Stage Workflow

- **Stage 1 – OSINT acquisition.** Acquisition is the process of getting conflict-related news, ReliefWeb/HDX humanitarian data, and dated geolocated conflicts. Logs capture source metadata and are de-duplicated through title and URL hashing, before language filtering takes place, allowing for English-language NLP.

- **Stage 2 – Need classification into humanitarian clusters.** Document mapping is carried out with respect to both the IASC-derived need clusters and an additional energy/winterisation category (TABLE 1). Seed vocabulary-based high precision matching, embedding of sentences to test the paraphrase of cluster descriptions [18, 19], and topic modeling for taxonomy coverage validation [16, 17].
- **Stage 3 – Severity and urgency weighting.** Each need mention is associated with a bounded severity weight $w \in [0.5, 1.5]$ based on sentiment and urgency [21]. The weight bounds prevent single emotive mentions from overpowering while emphasizing rapid changes.
- **Stage 4 – Geoparsing to administrative units.** Named entity recognition finds place names and then maps them into Ukrainian administrative units through geoparsing using the gazetteer [22, 23]. Unmapped locations provide data only to national totals, and lower certainty mappings are given less weight later on.
- **Stage 5 – The Humanitarian Need Pressure Index (HNPI).** HNPI is defined for region (oblast) r, humanitarian cluster c, document d, and week t. First, cluster pressure, i.e. thematic pressure of the news, is computed as

$$P_{r,c,t} = \sum_{d \in D_{r,t}} q_d \, s_d \, I_{d,c},$$

where $D_{r,t}$ is the set of documents assigned to region r and week t, $I_{d,c}$ indicates assignment to cluster c, s_d is the limited weight of severity and urgency, and q_d is the geoparsing confidence weight. Overall raw pressure is

$$P_{r,t} = \sum_{c} \lambda_c P_{r,c,t}.$$

For the experiment, λ_c is 1.3 for Energy/Winter, 1.1 for Shelter/NFI, 0.8 for WASH, 1.0 for Health, 0.9 for Food Security, 1.0 for Protection, 1.2 for Displacement, and 0.6 for Logistics. These coefficients have been pre-specified as explicit expert prior weights, not derived from the forecast target. Energy/winterisation is given higher weighting due to the risk that electricity, heating, and grid infrastructure damage pose in the context of Ukraine. Displacement and Shelter/NFI are also weighted higher due to the direct correlation between these categories and population movement and lack of suitable housing. Health and Protection have their default high-priority weightings maintained. Meanwhile, Food Security and WASH have lower weights in this experiment due to their infrequent appearance as immediate threat signals in the news and ReliefWeb data. Logistics is down-weighted because it mainly indicates access and delivery constraints that create the response rather than direct need itself. The weighting system used here is designed to be inspectable and replaceable with more data-driven systems in future works. Severity is derived as

$$s_d = 0.5 + \mathrm{clip}(0.5\,\mathrm{neg}_d + 0.5\,\mathrm{urg}_d, 0, 1),$$

in this case, neg_d refers to the harm or negative-sentiment score for document d, normalized based on indicators including casualties, damage, disruption of service, deprivation, injury, and death; urg_d refers to the urgency score for document d, normalized based on indicators including emergency response, evacuation, urgent help, winterization, and disruption of critical services. Both of these scores are normalized to $[0, 1]$. So, s_d can be viewed as a severity-urgency weight: documents that have strong signals of harm and urgency carry more weight in the need pressure, while documents without such signals just keep the base weight. Documents with unresolved location will be left out from the calculation of region-week HNPI. The region-week cell with zero document count will be considered to have zero raw pressure prior to standardization. Gazetteer match documents will receive $q_d = 1$, while those with unresolved or low-confidence location will get $q_d = 0$. Probabilistic geoparser can assign fractional q_d instead. The final index is

$$\mathrm{HNPI}_{r,t} = 100 \frac{P_{r,t} - \min_{r,t} P_{r,t}}{\max_{r,t} P_{r,t} - \min_{r,t} P_{r,t} + \epsilon},$$

resulting in a scale of 0 to 100. The intensity of conflict events measured by ACLED or ACLED schema data is deliberately kept out of the variable $P_{r,t}$ and is incorporated at a later stage for predictive purposes; ϵ is a small positive constant required solely for numerical stability, preventing division by zero in the event that all raw pressure values are identical.

- **Stage 6 – Early-warning regression and surge alarm.** The early-warning model forecasts the upcoming proxy $y_{r,t+h}$ based on lagged HNPI, cluster mentions, conflicts and region risk exposure. The point forecast is calculated through regularised regression, with a logistic head indicating when the forecast belongs to the upper quartile for surges.
- **Stage 7 - Decision support in the context of 'Do No Harm.'** Output includes administrative maps, regional rankings and surge alerts, together with relevant and open-source documentation. There is no storage of any personal information in the system, and geolocation at the individual level is not used [3].

Table 1: Types of humanitarian needs and example of keywords

Cluster (need facet)	Representative cue terms (examples)
Energy / Winterisation	blackout, power cut, grid, substation, heating, generator, no electricity, transformer
Shelter / NFI	destroyed homes, damaged housing, shelter, collective centre, blankets, repair, roofing
Water	no water, water supply, sanitation, sewage, drinking water, pumping station, hygiene
Health	hospital hit, medicines, ambulance, casualties, wounded, medical supplies, trauma
Food Security	food shortage, bread, market closed, supply route, hunger, agricultural, harvest
Protection	civilians killed, evacuation, mines, detained, missing, children, gender-based
Displacement	fled, displaced, evacuated, refugees, IDPs, transit, leaving the city, relocation
Logistics	road cut, bridge destroyed, access, convoy, corridor, blocked, delivery, fuel

4. EXPERIMENT

This experimental period spans from 2024-09-02 to 2025-08-31: 15,824 Ukrainian articles on Ukrainska Pravda in English, 1,687 ReliefWeb documents, and 37,978 geolocated conflict events extracted from a Kaggle database following the ACLED schema [4, 7]. The event time series stops at 18 April 2025, thus there will be no event context for later weeks. Geoparsing narrowed down the reporting to twelve regions/oblasts and sources are listed in TABLE 2. All results will be provided in the experiment discussion.

Table 2: Open data sources used during the experiment

Source	Role in pipeline	Access	Records
Ukrainska Pravda, English edition	Primary unstructured need signal	Web scraping (HTML)	15,824 articles
ReliefWeb (UN OCHA)	Need signal + situation context	Public API [9]	1,687 reports
Ukraine Conflict Event Dataset (Kaggle)	Geolocated conflict-event covariate	Open dataset [7]; ACLED schema [4]	37,978 events
IOM DTM / UNHCR portal	Displacement outcome / validation	Public portals [10, 11]	reference

4.1 Detected need by cluster

FIGURE 2, shows that health, shelter/NFI, and displacement top the severity-weighted cluster distribution, followed by energy, protection, food security, water, and logistics. The relatively similar values show the presence of a general need signal rather than a dominating theme, whereas ReliefWeb news increases health and protection against the news-only dataset.

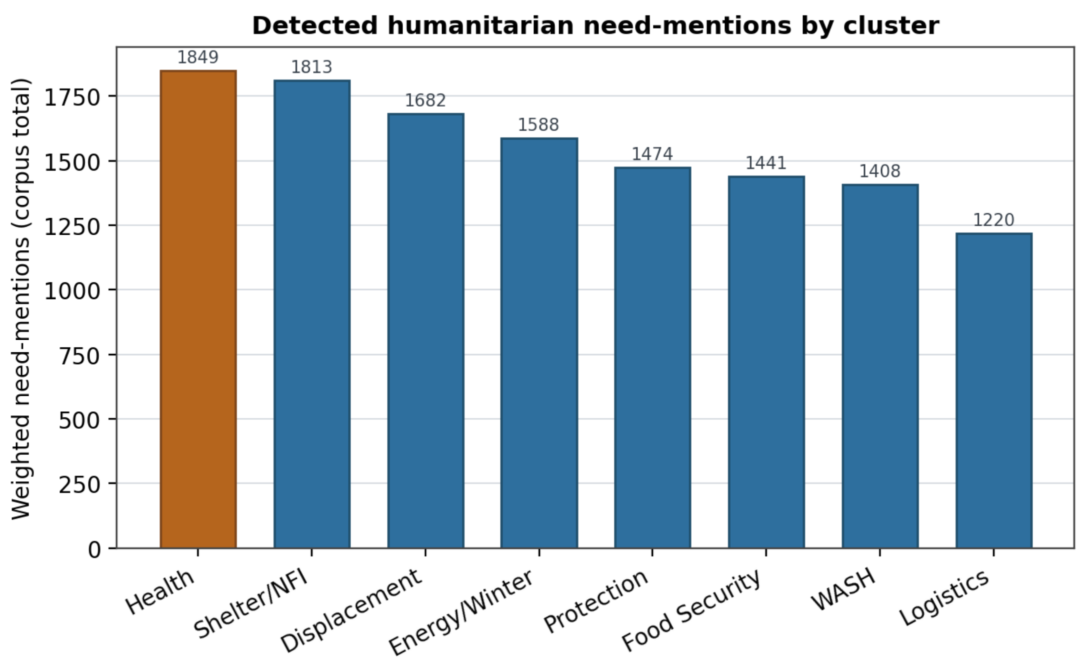

Figure 2: Distribution of severity-weighted humanitarian need mentions by cluster.

For verification purposes, the topics generated through LDA on the same corpus align consistently with the taxonomy (TABLE 3), indicating that there is an organized humanitarian signal within the open source record [16, 25].

Table 3: Unsupervised LDA topics recovered from the corpus (auxiliary view)

#	Leading terms	Approximate context
1	injured, killed, Kherson, attack, hromada	Protection / casualties
2	energy, power, sanctions, gas, oil	Energy / winterisation
3	attack, damaged, Kharkiv, building, emergency	Shelter / urban strikes
4	drones, air, missiles, drone, missile	Aerial & drone attacks
5	settlements, troops, Kursk, Pokrovsk, combat	Front-line combat
6	European, peace, meeting, NATO, security	Diplomacy / security
7	nuclear, Kremlin, leader, talks	Strategic / political
8	children, health, services, protection, needs	Health & protection

A manual validation check for the NLP need-classification and geoparsing process has been performed. For a random selection of 200 records drawn from Ukrainska Pravda and ReliefWeb websites, we have manually validated whether a humanitarian need exists or not, along with the relevant cluster labels, urgency/severity indicators, and region/oblast geoparsing. Among these 200 records, three of them are categorized as unanswered while one is marked with an unclear humanitarian need. The automatic need detection step has very good recall value (0.900) but low precision value (0.593), and therefore can be used as a preliminary screening mechanism rather than a labeling process. The multi-label clustering assignment has relatively low precision and F1 values. However, the geoparser of the gazetteer has 0.878 precision for first match in the stated twelve regions/oblast range, while perfect recovery for all manually marked geography is 0.500 since the parser only retrieves the first region in scope.

Table 4: Manual validation of need classification and geoparsing on annotated documents

Validation task	N	Precision	Recall	F1/Accuracy
Need detection (any humanitarian cluster)	196	0.593	0.900	0.715 F1
Cluster assignment, micro-average	196	0.298	0.688	0.416 F1
Cluster assignment, macro-average	196	0.322	0.759	0.400 F1
Geoparsing, in-scope first-match region	196	–	–	0.878 accuracy
Geoparsing, strict first-match geography	196	–	–	0.500 accuracy

4.2 Regional and temporal structure of need

FIGURE 3, depicts persistent pressure in Donetsk, Kyiv, and Kharkiv oblasts, along with short bursts of activity elsewhere. FIGURE 4, lists regions by average HNPI, with Donetsk (43), Kyiv (30), and Kharkiv (28) in first place, while Dnipro, Zaporizhzhia, and Sumy follow. The list correlates generally with intense event activity, except for Kyiv, where the media-derived pressure needs to be understood in its event context.

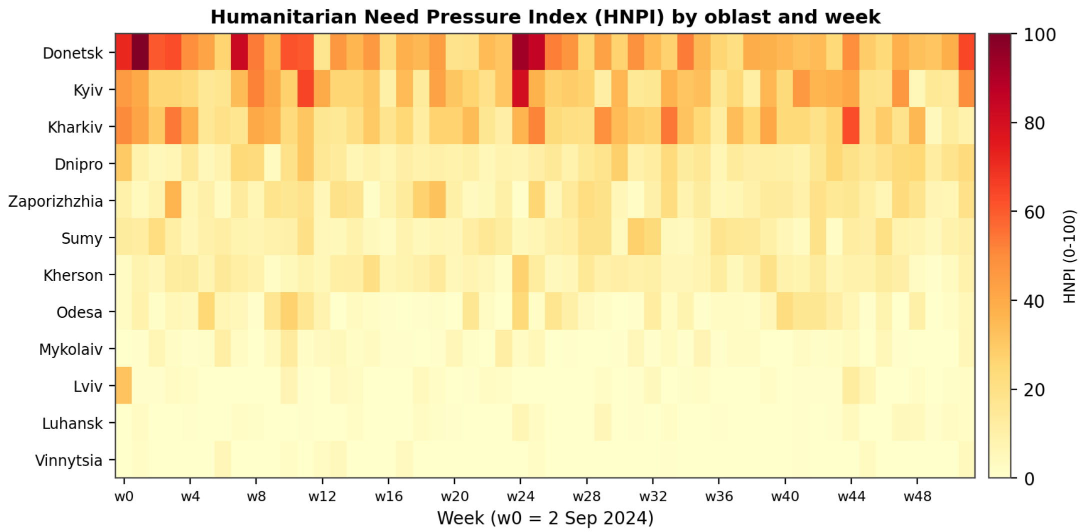

Figure 3: HNPI by region and week (w0 = 2024-09-02).

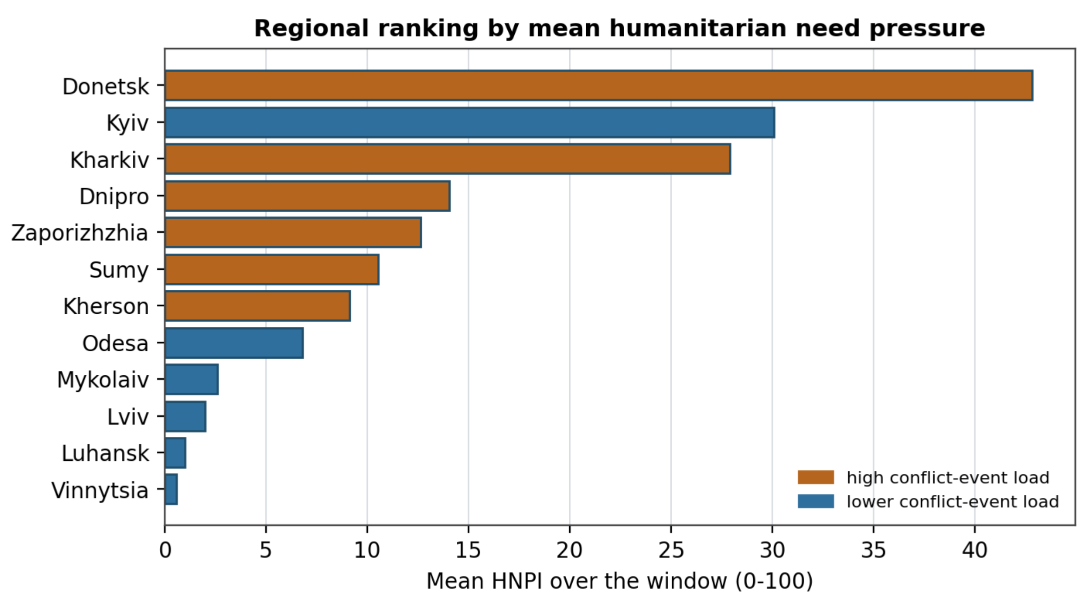

Figure 4: Regional ranking by mean HNPI. Bar colour marks regions with high conflict-event load.

4.3 Early-warning performance

The model predicts HNPI from one to four weeks out based on lagged index and conflict attributes, for target weeks ≤ 35 and prediction of weeks 36–51. As there is no independent outcome measure, this examines predictability against persistence and not validity. FIGURES 5–6 and TABLE 5, demonstrate MAE around 5.7 points per week, which is always less than persistence (6.2–6.9), as well as surge alarm AUC at about 0.91. This suggests a short-lead early warning system, but not yet an effective predictor of need for aid.

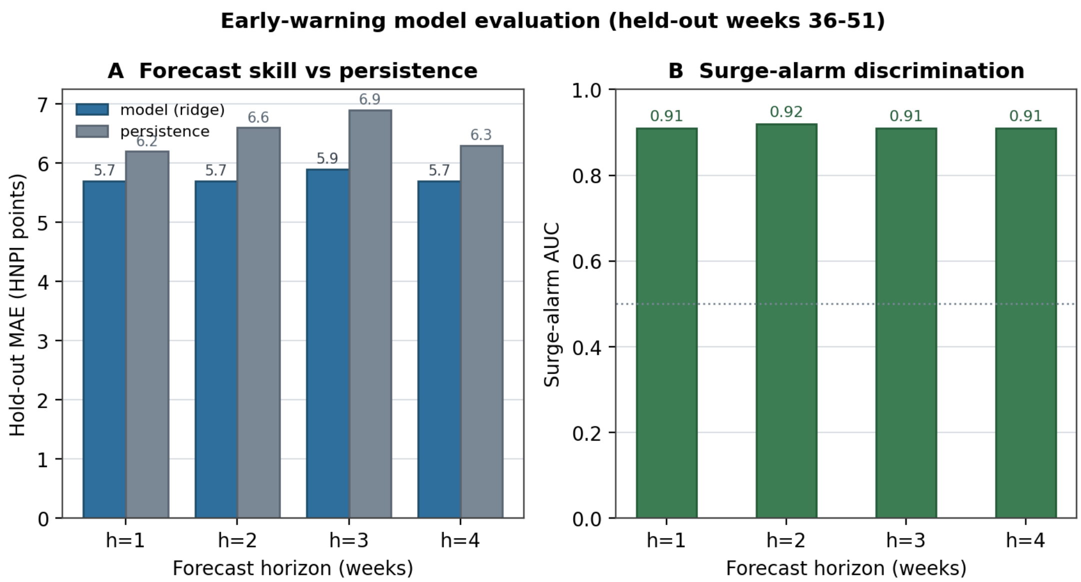

Figure 5: Forecast performance of early warning, held-out weeks 36-51: (A) forecast MAE compared with persistence benchmark by horizon; (B) AUC of surge detection by horizon.

Table 5: Early-warning performance by forecast horizon (held-out weeks 36–51)

Horizon h (weeks)	Model MAE	Model RMSE	Persist. MAE	Surge AUC
1	5.7	8.8	6.2	0.91
2	5.7	8.4	6.6	0.92
3	5.9	8.6	6.9	0.91
4	5.7	8.4	6.3	0.91

We also added paired statistical tests for the model–persistence comparison on the held-out region-week observations. For each horizon, the paired quantity is the reduction in absolute error, $|y - \hat{y}_{\mathrm{persist}}| - |y - \hat{y}_{\mathrm{model}}|$, computed for the same 192 held-out region-weeks. A non-parametric paired bootstrap with 10,000 resamples provides 95% confidence intervals and two-sided p-values for the mean error reduction. The improvement is statistically significant at the 5% level for horizons $h = 2$ and $h = 3$, while the $h = 1$ and $h = 4$ improvements are directionally positive but not significant at that threshold (TABLE 6).

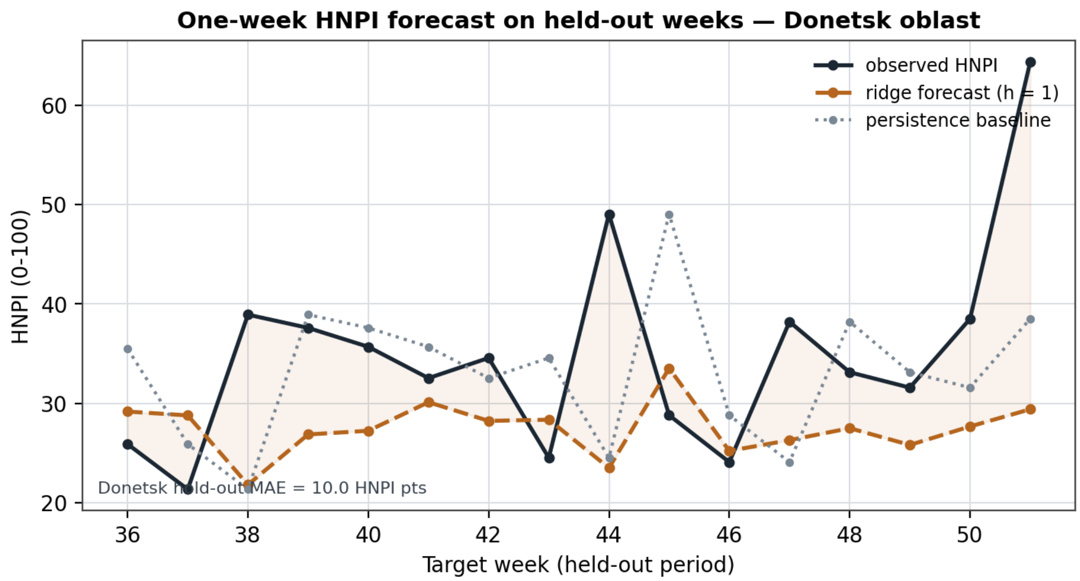

Figure 6: Weekly HNPI forecast versus observed index for the held-out weeks for Donetsk oblast, with the persistence benchmark for comparison.

Table 6: Paired bootstrap significance tests for MAE improvement over persistence

Horizon h	Test n	MAE reduction	95% bootstrap CI	p-value
1	192	0.468	[-0.111, 1.095]	0.1160
2	192	0.868	[0.108, 1.677]	0.0276
3	192	1.027	[0.224, 1.878]	0.0110
4	192	0.617	[-0.121, 1.343]	0.1026

Since the last day of the ACLED-schemed conflict event covariates is 18 April 2025 and the HNPI range extends up to 31 August 2025, we performed a sensitivity analysis for the ridge regression using the model without any ACLED events (i.e., only lagged HNPIs and regional exposure) and found that the model fits nearly identically well (and slightly better in MAE) (TABLE 7). This indicates that the reported held-out forecast skill is not driven by the incomplete conflict-event tail. It also clarifies the interpretation: late-window forecasts should be read mainly as HNPI-autoregressive diagnostics, while a complete event stream remains necessary before making claims about the marginal predictive value of conflict-event covariates.

Table 7: Sensitivity analysis excluding ACLED conflict-event covariates

Horizon h	Test n	MAE with ACLED	MAE without ACLED	Difference
1	192	5.726	5.719	-0.007
2	192	5.760	5.720	-0.040
3	192	5.869	5.819	-0.049
4	192	5.694	5.638	-0.056

4.4 External validation against OCHA/JIAF HNRP 2025

External validation is provided by including a cross-sectional consistency check of the data with the OCHA/JIAF Ukraine 2025 Humanitarian Needs and Response Plan dataset which is available through HDX [27]. The reason for choosing this data source is that there is an available downloadable XLSX file of People in Need (PiN) at ADM1 region level and cluster-level PiN. This source is therefore more compatible with the organizational and geographical structure of the HNPI in terms of the regions and clusters, compared to the readily available IOM DTM and UNHCR public websites: IOM DTM is very pertinent in the area of displacement, but administrative level time-series need an API connection or report extraction, whereas UNHCR public Ukraine statistics focus on national IDPs and refugees.

For the validation process, the 12 region covered by the HNPI index were chosen. For each of the region, we correlated average HNPI during the 52-week time period of our study with OCHA/JIAF overall joint PiN, and we also correlated severity-weighted HNPI mentions per cluster for OCHA/JIAF PiN values for the same cluster when there is such a match. As the OCHA/JIAF is only an annual index of reference for planning and needs estimates but not weekly outcomes, it measures consistency of regional and thematic ranking, not the forecast validity. Results are provided with Pearson correlation and Spearman rank correlation in TABLE 8. Shelter/NFI has the highest correlation, while overall regional has a moderate correlation ($\rho = 0.580$). The lower correlations for displacement and protection further confirm that the news-based HNPI continues to be an estimate and cannot replace the need assessment carried out on the ground.

Table 8: External validation against OCHA/JIAF Ukraine 2025 HNRP region-level reference data

HNPI comparison	OCHA/JIAF reference	n	Pearson r	Spearman ρ
Overall regional need	Overall joint PiN	12	0.317	0.580
Health weighted mentions	Health PiN	12	0.311	0.434
Shelter/NFI weighted mentions	Shelter/NFI PiN	12	0.512	0.727
WASH weighted mentions	WASH PiN	12	0.223	0.476
Food Security weighted mentions	Food Security and Livelihoods PiN	12	0.244	0.427
Protection weighted mentions	Protection Overall PiN	12	0.226	0.378
Displacement weighted mentions	IDP joint PiN	12	0.189	0.378

5. DISCUSSION

5.1 Validation Strategy

The additional comparison with OCHA/JIAF helps validate HNPI externally, although to some extent only. This helps to support the hypothesis of the partial coverage of the structure of independent regional needs by the index, specifically regarding the pressure on shelters, but this is not a validation of week ahead forecasts based on actual assistance demand. To fully validate this forecast, we need an independent time-resolved validation result, such as the IOM DTM displacements by regions and survey round or other independent humanitarian series.

5.2 Reliability Limitations

The key limitation is the forecast target, because the model forecasts the HNPI, which means it measures internal predictability but not observed humanitarian needs. Also, there is an issue with unevenness and language – only English reports are considered, plus the conflict-event covariate goes only until 18 April 2025. The lagging in reporting and information operations can skew the observed needs, thus the need for verification and cross-reporting persists [25, 26].

Therefore, we consider HNPI as a proxy measure of biased reporting and not a census of need. Source selection bias is one of the four ways through which the bias could happen where English-language Ukrainska Pravda and ReliefWeb have more coverage of the events seen by the media and reporting systems. Geographic reporting bias occurs in cases where certain places like the front lines or capitals have more coverage than other areas that are less visible. Thematic bias occurs when certain types of needs, for example, protection risks or chronic WASH constraints, are underreported in news. There is also temporal bias due to time lag between field situations and reporting.

5.3 Ethical and Legal Considerations

Due to the re-identification and mosaic risks posed by OSINT [3], the pipeline is designed to work at the aggregated level of administrative areas without any storage of personal data and excluding individual geolocation. Deployment should be carried out with data protection impact assessment.

The operational safeguarding measures include the following: release aggregate region level indicators, not geographic coordinates or personally identifiable information; retain source attribution and provide analyst access to the evidence behind all alerts; human-in-the-loop verification of resource allocation/escalation; block alerts that could divulge any civilian facility, route, shelter or distribution point; conduct data protection impact assessment with access control and retention period, including audit logs; and release model uncertainties and reporting bias alongside all alerts.

6. CONCLUSION

This paper introduced an auditable OSINT pipeline for transforming open news and public humanitarian data into the HNPI, as well as a short-lead alert on HNPI surges. In a live experiment on 17,511 documents, need related to health, shelter and displacement was predominant; pressure clustered in Donetsk, Kyiv and Kharkiv; and the lagged ridge model prediction had MAE close to 5.7 points and surge-alarm AUC close to 0.91. This shows how the system can act as a short-lead warning system, but explicitly in pursuit of the HNPI itself.

The strengths of this approach lie in low cost with audibility; the weaknesses are coverage bias, vulnerability to manipulation and OSINT ethics. Further research is needed to validate on additional outcome measures like the IOM DTM displacement, expand the text data set to multilingual news, add back the full event stream and evaluate across multiple splits.

The revised conclusion narrows the claim accordingly: HNPI is a clear early warning indicator, which will be useful in prioritizing attention by analysts, but should not be employed as a singular metric for humanitarian needs or automatic allocation tool. Through the additional manual NLP/Geoparsing analysis, it has been demonstrated that the current HNPI is best suited as a high recall screen and should therefore be coupled with analyst analysis.

References

[1] https://humanitarianaction.info/.

[2] Inter-Agency Standing Committee (IASC). Reference Module for Cluster Coordination at the Country Level. IASC, 2015.

[3] https://blogs.icrc.org/law-and-policy/2023/12/05/deploying-osint-in-armed-conflict-settings-law-ethics-theory-of-harm/

[4] Raleigh C, Linke A, Hegre H, Karlsen J. Introducing ACLED: An Armed Conflict Location and Event Dataset. J Peace Res. 2010;47:651-660.

[5] https://fews.net/.

[6] Bellingcat Investigation Team. Civilian Harm in Ukraine: Mapping Incidents. 2022-2024. Bellingcat. Available from: https://www.bellingcat.com/.

[7] https://www.kaggle.com/datasets/sangampaudel530/ukraine-conflict-event-dataset-20222025.

[8] https://data.humdata.org/.

[9] https://reliefweb.int/.

[10] https://dtm.iom.int/ukraine.

[11] https://data.unhcr.org/.

[12] Imran M, Castillo C, Diaz F, Vieweg S. Processing Social Media Messages in Mass Emergency: A Survey. ACM Comput Surv. 2015;47:1-38.

[13] Alam F, Qazi U, Imran M, Ofli FH. Humaid: Human-Annotated Disaster Incidents Data From Twitter With Deep Learning Benchmarks. Proceedings of the International AAAI Conference on Web and Social Media (ICWSM). 2021;15:933-942.

[14] Alam F, Sajjad H, Imran M, Ofli F. CrisisBench: Benchmarking Crisis-Related Social Media Datasets for Humanitarian Information Processing. In Proceedings of the International AAAI Conference on Web and Social Media. 2021;15:923-32.

[15] Olteanu A, Castillo C, Diaz F, Vieweg S. CrisisLex: a lexicon for collecting and filtering microblogged communications in crises. In Proceedings of the international AAAI conference on web and social media . 2014;8:376-385.

[16] Blei DM, Ng AY, Jordan MI. Latent Dirichlet Allocation. J Mach Learn Res. 2003;3:993-1022.

[17] Röder M, Both A, Hinneburg A. Exploring the Space of Topic Coherence Measures. In Proceedings of the Eighth ACM International Conference on Web Search and Data Mining. 2015:399–408.

[18] Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. 2019;1:4171-4186.

[19] Reimers N, Gurevych I, Sentence BE. Sentence Embeddings Using Siamese BERT-Networks. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). 2019:3982-3992.

[20] Mikolov T, Sutskever I, Chen K, Corrado G, Dean J. Distributed representations of words and phrases and their compositionality. Adv Neural Inf Process Syst. 2013;26.

[21] Hutto C, Gilbert E. Vader: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. In Proceedings of the International AAAI Conference on Web and Social Media 2014;8:216-225.

[22] Honnibal M, Montani I, Van Landeghem S, Boyd A. SpaCy: Industrial-Strength Natural Language Processing in Python. 2020. Available from: https://spacy.io/.

[23] Gritta M, Pilehvar MT, Limsopatham N, Collier N. What's Missing in Geographical. Parsing? Language Resources and Evaluation. 2018;52:603–623.

[24] https://www.gdeltproject.org/

[25] Khairova N, Holyk Y, Sytnikov D, Mishcheriakov Y, Shanidze N. Topic Modelling of Ukraine War-Related News Using Latent Dirichlet Allocation With Collapsed Gibbs Sampling .In 8th International Conference on Computational Linguistics and Intelligent Systems. 2024;3688:1-15.

[26] Khairova N, Sharonova N, Babkova N, Holyk Y, Sytnikov D. The Analysis of the Impact of News Sentiment on the Dynamics of Energy Resource Imports and Exports in Ukraine During the Full-Scale War. In Intelligent Systems Workshop at 9th International Conference on Computational Linguistics and Intelligent Systems, ISW-CoLInS 2025, 2025; 3983:12-27.

[27] https://data.humdata.org/dataset/ukr-jiaf-humanitarian-needs-and-response-plan