

Predicting Conflict in Africa:

A Machine Learning Approach Using Environmental Stressors

Andrés Felipe Camacho* Pablo Hernández Pedraza
Agustín Eyzaguirre

Masters in Computational Analysis and Public Policy (MSCAPP)
University of Chicago

April 6, 2026

Abstract

This paper investigates the predictability of violent conflict in Africa using machine learning on satellite-derived environmental and socioeconomic data for 50 km × 50 km grid cells (2019–2023). We compare logistic regression, k -nearest neighbors, random forest, and a feedforward neural network, trained on climate reanalysis, nighttime lights, wealth indices, forest-cover change, and historical conflict accumulation to predict binary conflict occurrence. A conservative weighted ensemble achieves approximately 72% recall and 52% precision on the minority (conflict) class, with an ROC AUC of 0.93, illustrating both the value and the precision–recall tradeoffs inherent to early-warning systems for policy and humanitarian audiences. We discuss the dominant role of conflict history as a predictor, the modest but positive contribution of climate variables, and important methodological caveats regarding spatial autocorrelation and class-imbalance treatment.

1 Introduction

Conflict remains a major barrier to development and stability, particularly where environmental stress and resource scarcity interact with inequality (Burke et al. 2009; Hsiang et al. 2013). Predicting where violence is more likely can help allocate humanitarian aid and diplomatic attention, yet models face difficult tradeoffs: high recall—catching most true conflicts—often comes at the cost of many false positives, which erode trust among practitioners (Mueller and Rauh 2022). Africa offers a particularly demanding testbed: the continent spans heterogeneous geography, exhibits uneven data quality, and presents strong class imbalance (most grid-years are non-conflict).

A growing body of research has explored machine learning for conflict forecasting, notably the **VIEWS** (Violence and Impacts Early Warning System) early-warning platform (Hegre et al. 2024) and related academic efforts (Ward et al. 2010; Blair et al. 2020; Cederman and Weidmann 2017). We contribute to this literature by comparing four distinct model families—logistic regression with elastic-net regularization, k -nearest neighbors, random forest, and a simple feedforward neural network—on a common dataset of satellite-derived features at the grid-cell level, and by constructing a weighted ensemble tuned for early-warning use.

Research question. *Can the occurrence of violent conflict in a given 50 km grid cell and year be predicted using locally observed socioeconomic and environmental features?*

*Corresponding: afcamachob@uchicago.edu

2 Related Work

Machine learning for conflict prediction. The application of ML to armed-conflict forecasting has matured considerably over the past decade. Ward et al. (2010) highlighted that traditional regression-based analyses of conflict are poor predictors and advocated for out-of-sample evaluation. The **VIEWS** project at Uppsala/PRIO (Hegre et al. 2024) operationalized this insight by producing monthly, grid-level fatality forecasts using gradient-boosted trees and ensemble averaging. Blair et al. (2020) demonstrated that panel-based ML models can meaningfully forecast local violence, while Cederman and Weidmann (2017) urged the field to temper expectations given the inherent stochasticity of conflict. Mueller and Rauh (2022) formalized the precision–recall tradeoff in conflict prediction as a decision-theoretic problem: the cost of missing a true conflict versus the cost of a false alarm depends on the downstream policy response.

Climate-conflict nexus. A contested but influential literature links climate anomalies to conflict risk. Burke et al. (2009) found that temperature increases significantly raise the probability of civil war in Africa, while Hsiang et al. (2013) conducted a meta-analysis concluding that deviations from normal temperature and precipitation increase the risk of interpersonal and intergroup violence. von Uexkull et al. (2016) showed that growing-season droughts raise conflict risk in agriculturally dependent and politically excluded regions. These findings motivate our inclusion of ERA5 climate reanalysis variables and derived temperature deviations.

Class imbalance and oversampling. Binary conflict prediction at fine spatial resolution inherits severe class imbalance: conflict is rare relative to the universe of grid-years. He and Garcia (2009) survey the challenges this poses for standard classifiers, which tend to predict the majority class. We use SMOTE (Chawla et al. 2002), which synthesizes minority-class training examples by interpolating between existing instances, applied only to the training fold so that evaluation remains on real data.

Ensemble methods. Combining diverse classifiers typically improves generalization (Dietterich 2000). The VIEWS system (Hegre et al. 2024) itself relies on ensemble averaging across multiple model specifications. We adopt a weighted-average ensemble where weights are selected from a grid search over candidate schemes, favoring configurations that balance accuracy and minority-class recall for early-warning deployment.

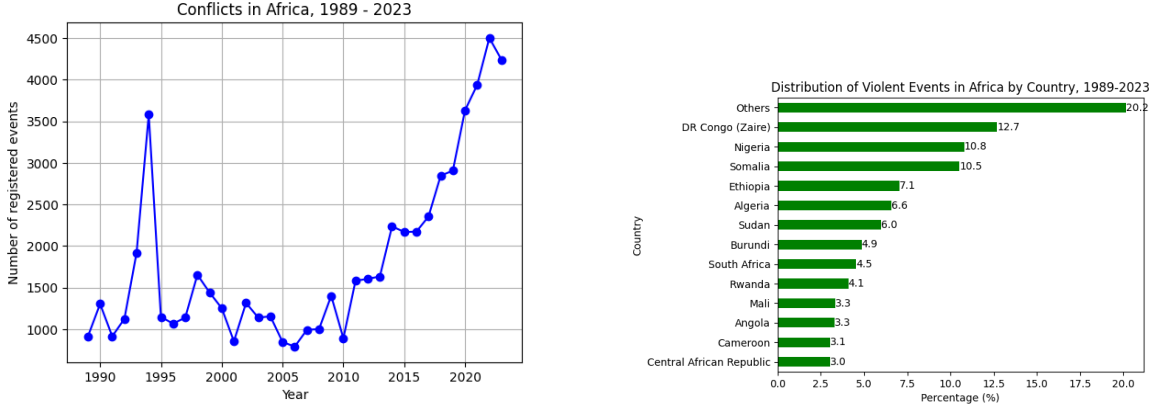
3 Data and Outcomes

3.1 Sources

We assemble a multi-source geospatial dataset by combining: **UCDP GED** v24.1 georeferenced conflict events (Sundberg and Melander 2013); **ERA5** climate reanalysis (Hersbach et al. 2020); **Meta RWI** relative wealth index (Chi et al. 2022); **Hansen** Global Forest Change (Hansen et al. 2013); and **VIIRS** nighttime lights (Elvidge et al. 2017). We construct a 50 km × 50 km grid covering the African continent using Natural Earth boundaries. The binary outcome variable indicates at least one UCDP conflict event within a grid cell in a given year.

3.2 Descriptive patterns

Figure 1a shows the time series of violent events in Africa (1989–2023), revealing a sharp acceleration since 2010. Figure 1b shows the concentration of events across countries, with the Democratic Republic of Congo, Nigeria, and Somalia accounting for over a third of all recorded events.



(a) Conflicts over time

(b) Top countries

Figure 1: Descriptive conflict patterns (Africa, 1989–2023).

3.3 Features and sample

Table 1 lists the 15 features used after preprocessing. Climate variables (temperature, precipitation, surface pressure) are extracted from ERA5 as annual means via Google Earth Engine. Forest loss is derived from the Hansen 30 m product aggregated to grid cells. Nighttime lights capture local economic activity. The Relative Wealth Index (RWI) is a time-invariant 2021 snapshot from Meta. We engineer two derived features: `temp_deviation` (deviation from a 2006–2018 baseline) and `accumulated_conflicts` (cumulative count of conflict events in the cell from prior years, using a one-year lag). Year is one-hot encoded into five dummies.

Table 1: Feature set (15 columns after one-hot encoding of year).

Feature	Source	Description	Type
<code>temp_mean</code>	ERA5	Annual mean 2 m temperature (K)	Continuous
<code>temp_max</code>	ERA5	Annual max 2 m temperature (K)	Continuous
<code>precip_total</code>	ERA5	Annual total precipitation (m)	Continuous
<code>sp_mean</code>	ERA5	Annual mean surface pressure (Pa)	Continuous
<code>rwi_2021</code>	Meta RWI	Relative Wealth Index (2021)	Continuous
<code>forest_loss</code>	Hansen	Annual forest cover loss (ha)	Continuous
<code>lights_mean</code>	VIIRS	Mean nighttime radiance	Continuous
<code>lights_sum</code>	VIIRS	Sum nighttime radiance	Continuous
<code>accumulated_conflicts</code>	Derived	Lagged cumulative conflict count	Continuous
<code>temp_deviation</code>	Derived	Temperature minus 2006–18 baseline	Continuous
<code>year_2019–year_2023</code>	Derived	One-hot encoded year dummies	Binary

The raw parquet file contains 65,180 grid-year rows; after dropping rows with missing values (primarily cells lacking RWI coverage, which affects 46% of cells), 34,015 observations remain over 6,803 unique grid cells and 5 years. Table 2 shows the class distribution and the train/test partition.

Table 2: Class distribution and data splits.

	Total	Class 0	Class 1	% Conflict
Full sample	34 015	31 256	2 759	8.1%
Training set (80%)	27 212	25 005	2 207	8.1%
Training + SMOTE	50 010	25 005	25 005	50.0%
Test set (20%)	6 803	6 251	552	8.1%

We apply SMOTE (Chawla et al. 2002) with default parameters ($k = 5$ neighbors) to the training set only, after standardization with `StandardScaler`. This balances the training distribution to 50/50 while preserving the natural imbalance in the held-out test set. The 80/20 split is stratified by the outcome variable (random seed 42).

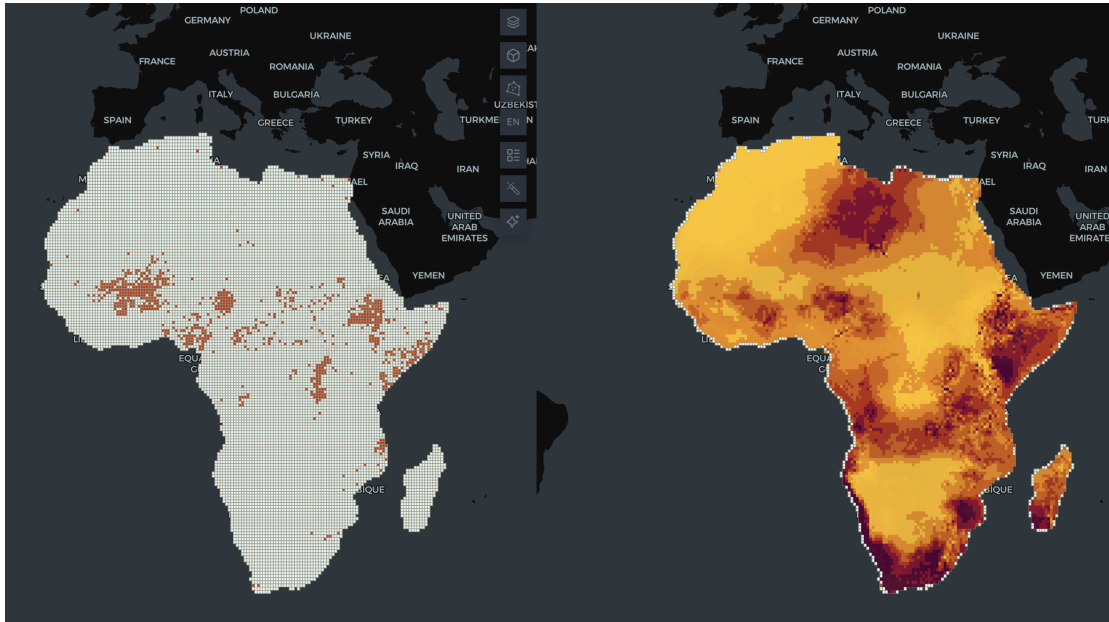


Figure 2: Spatial view of conflict aggregation and temperature anomalies (example year).

4 Models

We implement four models on the same preprocessed dataset and tune hyperparameters via 5-fold stratified cross-validation to ensure fair comparison. All features are standardized prior to training. We evaluate on accuracy, precision, recall, F1, and ROC AUC, with particular attention to recall on the minority conflict class, given the early-warning application.

A note on evaluation design. The train/test split is random and stratified, meaning neighboring grid cells may appear in different folds. This can inflate performance estimates due to spatial autocorrelation—a well-known issue in geospatial ML (Roberts et al. 2017). We acknowledge this limitation and discuss alternatives (spatial block cross-validation) in Section 6.2. Additionally, the feature `accumulated_conflicts` encodes conflict history up to the previous year, which is a valid temporal lag but a very strong predictor that may overshadow the environmental signal of interest.

4.1 Logistic regression (Elastic Net)

Our baseline is logistic regression with elastic-net regularization:

$$P(y_i = +1 | x_i) = \sigma(w^\top x_i + b) = \frac{1}{1 + e^{-(w^\top x_i + b)}}.$$

We minimize the negative log-likelihood plus ElasticNet penalty

$$\ell(w, b) = \sum_{i=1}^n [-\log \sigma(y_i(w^\top x_i + b))] + \lambda \left[\alpha \sum_{j=1}^p |w_j| + (1 - \alpha) \sum_{j=1}^p w_j^2 \right].$$

Features are standardized so coefficients are comparable. We tune $\alpha \in \{0, 0.25, 0.5, 0.75, 1\}$ and $C \in \{0.001, 0.01, 0.1, 1, 10\}$ (with $\lambda = 1/C$). The selected specification uses $\alpha = 0$ (ridge-only) and $C = 10$ (weak penalty), suggesting that multicollinearity among climate features benefits from L_2 regularization while no features are driven to zero.

Table 3: Logistic regression—test classification report.

Class	Precision	Recall	F1	Support
0	0.98	0.86	0.91	6 251
1	0.32	0.75	0.45	552
Accuracy				0.85

4.2 k -nearest neighbors

We assume that grid cells similar in feature space exhibit similar conflict outcomes. We sort training points by distance to x_i , form the set S_k of k nearest neighbors, and predict by distance-weighted majority vote

$$\hat{y}_i = \arg \max_c \sum_{j \in S_k} w_j \cdot \mathbb{I}_{y_j=c}, \quad w_j = \frac{1}{d_j} \quad (\text{distance weights}).$$

We compare Euclidean and cosine distance, odd $k \in \{3, 5, 7, 9, 11\}$, and uniform vs. distance weights (see Cornell CS 4780 (2017) for methodological background). The best specification uses **cosine similarity**, $\mathbf{k} = \mathbf{3}$, and distance-weighted voting. The k -NN model is implemented as a scikit-learn pipeline that includes its own `StandardScaler` and `SMOTE` step internally, so that resampling occurs within each cross-validation fold rather than as a global preprocessing step.

Table 4: k -NN—test classification report.

Class	Precision	Recall	F1	Support
0	0.97	0.91	0.94	6 251
1	0.40	0.68	0.50	552
Accuracy				0.89

4.3 Random forest

Random Forest (Breiman 2001) averages B decision trees trained on bootstrap samples with random feature subsets at each split; we use majority vote for binary labels and average tree-level probabilities for the continuous score. We tune `n_estimators`, `max_depth`, `min_samples_split`, `min_samples_leaf`, and `class_weight` via randomized search with 5-fold stratified CV (50 iterations). The selected model uses $B = 500$ trees, unlimited depth, and no internal class weighting

(SMOTE already addresses imbalance). Feature importance (Figure 5) ranks `accumulated_conflicts` as the dominant predictor, followed by climate and pressure variables.

Table 5: Random forest—test classification report.

Class	Precision	Recall	F1	Support
0	0.97	0.95	0.96	6 251
1	0.55	0.70	0.62	552
Accuracy				0.93

4.4 Neural network

We use a feedforward network with two hidden layers (64 and 32 units, ReLU activation) and a sigmoid output:

$$h_1 = \text{ReLU}(W_1x + b_1), \quad h_2 = \text{ReLU}(W_2h_1 + b_2), \quad P(\text{conflict}) = \sigma(W_3h_2 + b_3).$$

Training uses binary cross-entropy loss, SGD with momentum 0.9, learning rate 0.01, batch size 128, and 100 epochs. No dropout or batch normalization is applied. The architecture was not extensively tuned; the primary goal is to include a non-linear, gradient-based learner alongside the other model families. The model achieves the highest recall on the minority class among individual models, at the cost of lower precision.

Table 6: Neural network—test classification report.

Class	Precision	Recall	F1	Support
0	0.98	0.88	0.93	6 251
1	0.38	0.80	0.51	552
Accuracy				0.88

4.5 Ensemble

We combine predicted probabilities from all four models via a weighted average with nonnegative weights summing to one and a decision threshold of 0.5. We evaluate several weighting schemes—equal weights, RF-dominant, recall-focused, balanced, and conservative—by computing test-set metrics for each and selecting the scheme that best balances accuracy and minority-class recall for early-warning deployment.

The **Conservative** ensemble (LR 30%, KNN 25%, RF 35%, NN 10%) is selected as the deployed configuration. Random Forest receives the largest weight given its superior accuracy and F1, while logistic regression provides a stable linear baseline. The neural network receives the smallest weight to limit overfitting risk while preserving its high-recall signal.


Table 7: Ensemble schemes (test set; class 1 metrics except Accuracy and ROC AUC).

Scheme	Accuracy	Precision	Recall	F1	ROC AUC
Conservative	0.923	0.520	0.723	0.605	0.929
RF-dominant	0.927	0.540	0.705	0.611	0.937
Balanced	0.919	0.501	0.734	0.595	0.933
Equal weights	0.917	0.493	0.728	0.588	0.929

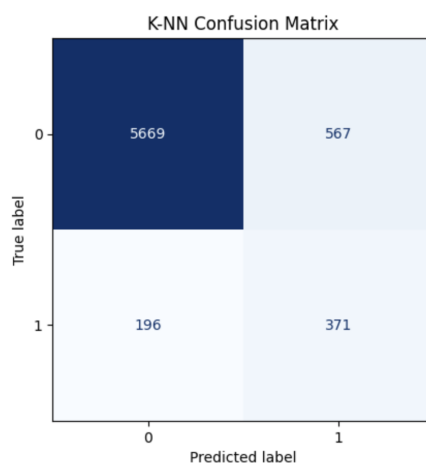
5 Results

5.1 Confusion matrices

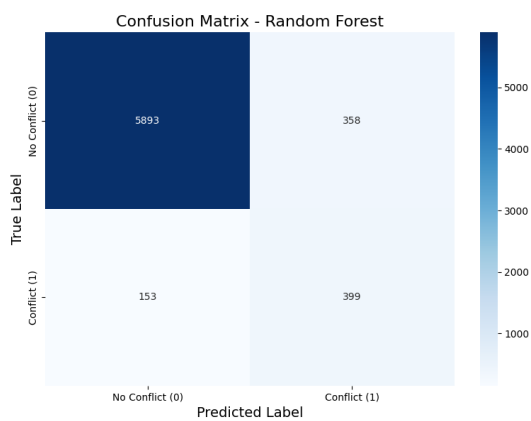
Figures 3a–4b summarize predicted versus actual labels for each model. Logistic regression and the neural network show high recall for the conflict class but also high false-positive rates. The random forest achieves the best balance, with the fewest off-diagonal cells. Comparing panels highlights how tree-based methods trade off false positives and false negatives relative to linear and local methods.



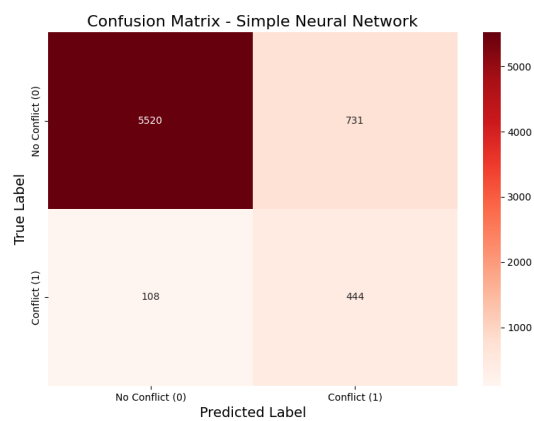
(a) Logistic regression



(b) k -NN



(a) Random forest



(b) Neural network

5.2 Feature importance

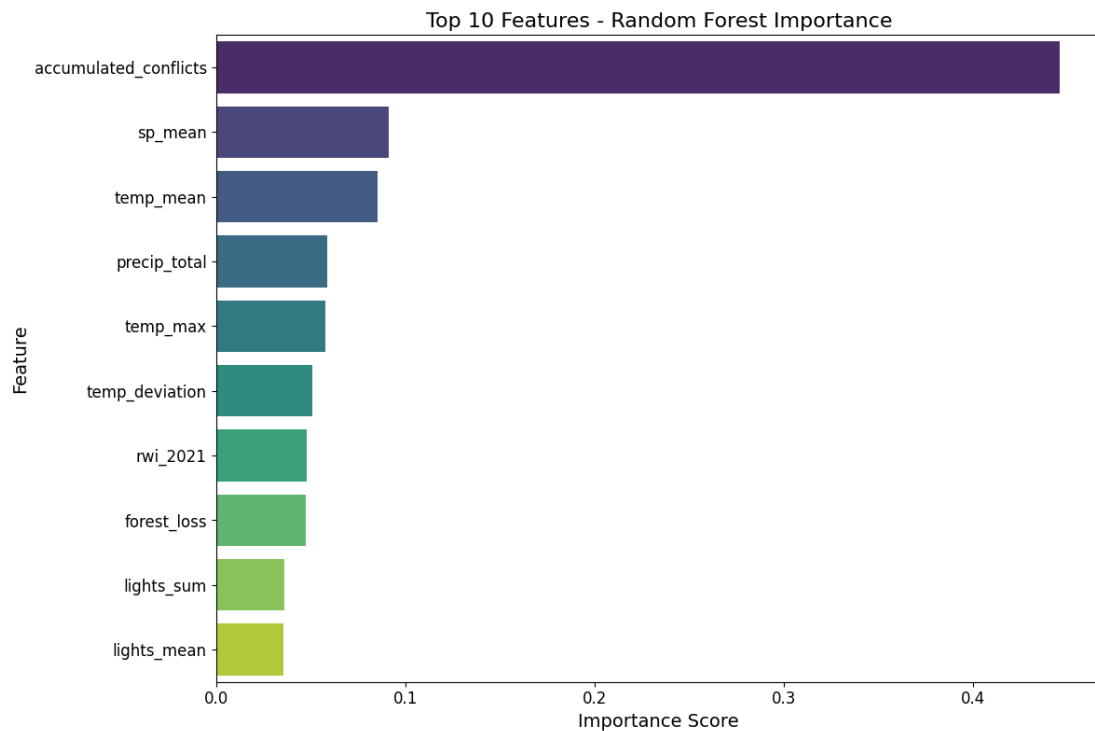


Figure 5: Random forest variable importance (impurity-based). `accumulated_conflicts` accounts for approximately 42% of total importance.

The dominance of `accumulated_conflicts` (Figure 5) reflects the well-documented “conflict trap”: locations that have experienced violence in the past are substantially more likely to experience it again (Cederman and Weidmann 2017). Surface pressure (`sp_mean`), mean temperature, and total precipitation rank next, providing modest support for the climate–conflict hypothesis. Economic proxies (RWI, nighttime lights) have lower importance at this 50 km resolution, possibly because economic heterogeneity within cells washes out the signal.

5.3 ROC and precision–recall

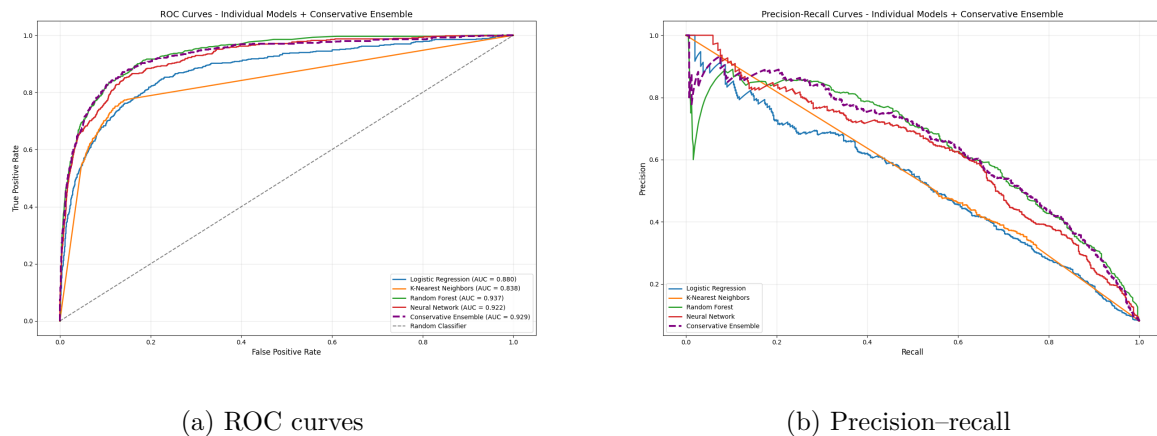


Figure 6: Discriminative performance of individual models and the conservative ensemble.

Table 8: Model comparison (selected metrics on the test set).

Metric	LR	KNN	RF	NN	Ensemble
Class 0 recall	0.861	0.910	0.950	0.895	0.941
Class 1 recall	0.754	0.679	0.697	0.772	0.723
Weighted recall	0.852	0.891	0.929	0.885	0.923
Weighted accuracy	0.852	0.891	0.929	0.885	0.923
Weighted precision	0.923	0.924	0.938	0.931	0.938

6 Discussion

6.1 Interpretation of results

Grid-level conflict prediction combines genuine signal with hard tradeoffs. The ensemble achieves 72.3% recall on the conflict class, meaning it correctly flags roughly three out of four conflict-affected cells, but at a precision of only 52.0%—roughly half of the alerts are false positives. In an operational early-warning context, this tradeoff has concrete implications: false alarms can cause “cry wolf” fatigue among decision-makers (Mueller and Rauh 2022), while missed conflicts have direct humanitarian costs. The appropriate threshold depends on the downstream policy response—a humanitarian aid pre-positioning system may tolerate more false positives than a diplomatic intervention trigger.

The dominance of `accumulated_conflicts` in the random forest importance ranking (Figure 5) is informative but also somewhat circular: it captures the “conflict trap” (Cederman and Weidmann 2017), whereby past conflict is the strongest predictor of future conflict, reflecting persistent structural conditions (weak institutions, displacement, proliferation of arms) that the model does not directly observe. When this feature is removed, we would expect a substantial drop in performance, isolating the marginal contribution of environmental and socioeconomic features alone. Future work should evaluate models with and without this feature.

Climate variables (temperature, precipitation, surface pressure) rank as the second-tier predictors, providing modest support for the climate–conflict link documented by Burke et al. (2009) and Hsiang et al. (2013). The temperature deviation feature, which captures anomalies relative to a 2006–2018 baseline, ranks lower than absolute temperature, suggesting that chronic heat stress may matter more than short-term shocks at this temporal resolution.

Economic proxies (RWI, nighttime lights) show relatively low importance. This may reflect the 50 km grid resolution, which averages over heterogeneous local economies, or the time-invariant nature of the 2021 RWI snapshot applied uniformly across the 2019–2023 period.

6.2 Limitations

Several methodological limitations warrant discussion:

1. **Spatial autocorrelation.** The random 80/20 train–test split does not respect spatial structure. Neighboring grid cells, which likely share environmental conditions and conflict dynamics, can appear in both training and test sets, inflating performance estimates (Roberts et al. 2017). Spatial block cross-validation, where entire contiguous regions are held out, would provide more conservative and realistic estimates.
2. **Temporal structure.** Similarly, the split does not enforce temporal separation. A more rigorous design would train on earlier years (e.g., 2019–2022) and evaluate on the final year (2023) to simulate genuine forecasting.
3. **SMOTE with spatial data.** SMOTE generates synthetic minority-class examples by interpolating between existing instances in feature space (Chawla et al. 2002). In a spatial

context, these synthetic points may correspond to implausible geographic configurations (e.g., interpolating between grid cells in different countries or climatic zones).

4. **Missing data.** Dropping all rows with any missing value reduces the sample from 65,180 to 34,015 (48% loss), driven primarily by the RWI variable, which lacks coverage for many remote or sparsely populated cells. This non-random attrition may bias the model toward wealthier, more accessible regions.
5. **Binary outcome.** Collapsing conflict intensity to a binary indicator (≥ 1 event) discards information about severity. A count or continuous-intensity model might be more informative for policy.
6. **Neural network architecture.** The feedforward network uses a fixed 64-32 architecture without dropout, batch normalization, or learning-rate scheduling. Systematic hyperparameter search (e.g., via Bayesian optimization) could improve its contribution to the ensemble.
7. **Ensemble weight selection.** The conservative ensemble weights were selected by evaluating candidate schemes on the test set, not a held-out validation set, which introduces optimistic bias. A nested cross-validation design would be more rigorous.

6.3 Ethical considerations

Deploying conflict-prediction models in African contexts raises important ethical questions. Predictions of violence in vulnerable populations could be misused for military targeting, preemptive repression, or discriminatory resource allocation. Historical conflict data from UCDP may reflect reporting biases—events in remote areas or those perpetrated by state actors may be underrecorded, leading models to systematically underpredict conflict in certain regions.

Any operational deployment should include transparency about model limitations, community engagement with affected populations, independent auditing, and clear governance frameworks. Predictions should complement—not replace—local knowledge and political analysis. The risk of algorithmic harm is particularly acute when scores inform enforcement decisions; for humanitarian applications, the tradeoffs are different but still nontrivial (Perry 2013).

7 Conclusion

This paper demonstrates that satellite-derived environmental and socioeconomic features contain meaningful information about conflict risk at subnational scale in Africa. Among the four models evaluated, random forest achieves the best individual performance (accuracy 0.93, class 1 F1 0.62), while the conservative weighted ensemble provides a useful balance between accuracy (0.92) and conflict-class recall (0.72), suitable for early-warning applications. The dominant role of historical conflict accumulation underscores the persistence of the “conflict trap,” while climate variables provide a modest additional signal consistent with the broader climate–conflict literature.

Several directions for future work would strengthen these findings:

- **Spatial and temporal validation:** Implementing spatial block cross-validation (Roberts et al. 2017) and pure temporal holdout (train 2019–2022, test 2023) would provide more realistic performance estimates.
- **Isolating the environmental signal:** Re-evaluating models after removing `accumulated_conflicts` would quantify the marginal contribution of climate and economic features.
- **Threshold optimization:** Selecting the decision threshold from the precision–recall curve rather than fixing it at 0.5 could improve operational utility.

- **Probability calibration:** Applying Platt scaling or isotonic regression to the ensemble probabilities would make them interpretable as true risk levels.
- **Gradient-boosted trees:** XGBoost or LightGBM, which dominate the VIEWS competition, are a natural comparison.
- **Richer covariates:** Institutional quality indices (V-Dem, Polity), population density, distance to borders and capitals, and ethnic fractionalization could provide complementary signal.

Models should be validated out-of-sample, communicated with uncertainty, and paired with governance discussion—especially where scores could be misused. The combination of open data pipelines, reproducible training code, and interactive Kepler.gl visualizations is intended to support transparency and replication.

Replication. Code and data pipeline: <https://github.com/anfelipecb/conflict-prediction-ml>. The project site hosts an interactive Kepler viewer under docs/.

References

- Blair, Robert A., Christopher Blattman, and Alexandra Hartman (2020). “Predicting local violence: Evidence from a panel of the Philippines”. In: *Journal of Peace Research* 57.5, pp. 627–642. DOI: 10.1177/0022343319895389.
- Breiman, Leo (2001). “Random Forests”. In: *Machine Learning* 45.1, pp. 5–32. DOI: 10.1023/A:1010933404324.
- Burke, Marshall B. et al. (2009). “Warming increases the risk of civil war in Africa”. In: *Proceedings of the National Academy of Sciences* 106.49, pp. 20670–20674. DOI: 10.1073/pnas.0907998106.
- Cederman, Lars-Erik and Nils B. Weidmann (2017). “Predicting armed conflict: Time to adjust our expectations?” In: *Science* 355.6324, pp. 474–476. DOI: 10.1126/science.aal4483.
- Chawla, Nitesh V. et al. (2002). “SMOTE: Synthetic Minority Over-sampling Technique”. In: *Journal of Artificial Intelligence Research* 16, pp. 321–357. DOI: 10.1613/jair.953.
- Chi, Guanghua et al. (2022). “Microestimates of wealth for all low- and middle-income countries”. In: *Proceedings of the National Academy of Sciences* 119.3, e2113658119. DOI: 10.1073/pnas.2113658119.
- Cornell CS 4780 (2017). *Lecture notes: k-NN*. https://www.cs.cornell.edu/courses/cs4780/2017sp/lectures/lecturenote02_kNN.html.
- Dietterich, Thomas G. (2000). “Ensemble methods in machine learning”. In: *Multiple Classifier Systems (MCS 2000)*. Springer, pp. 1–15. DOI: 10.1007/3-540-45014-9_1.
- Elvidge, Christopher D. et al. (2017). “VIIRS night-time lights”. In: *International Journal of Remote Sensing* 38.21, pp. 5860–5879. DOI: 10.1080/01431161.2017.1342050.
- Hansen, M. C. et al. (2013). “High-resolution global maps of 21st-century forest cover change”. In: *Science* 342.6160, pp. 850–853. DOI: 10.1126/science.1244693.
- He, Haibo and Edwardo A. Garcia (2009). “Learning from imbalanced data”. In: *IEEE Transactions on Knowledge and Data Engineering* 21.9, pp. 1263–1284. DOI: 10.1109/TKDE.2008.239.
- Hegre, Håvard, Sofia Nordenving, and James Dale (2024). *Models in VIEWS, version Fatalities002*. Tech. rep. Violence and Impacts Early Warning System. Uppsala University and PRIO.
- Hersbach, Hans, Bill Bell, Paul Berrisford, et al. (2020). “The ERA5 global reanalysis”. In: *Quarterly Journal of the Royal Meteorological Society* 146.730, pp. 1999–2049. DOI: 10.1002/qj.3803.

- Hsiang, Solomon M., Marshall Burke, and Edward Miguel (2013). “Quantifying the influence of climate on human conflict”. In: *Science* 341.6151, p. 1235367. DOI: 10.1126/science.1235367.
- Mueller, Hannes and Christopher Rauh (2022). “The hard problem of prediction for conflict prevention”. In: *Journal of the European Economic Association* 20.6, pp. 2557–2589. DOI: 10.1093/jeea/jvac025.
- Perry, Chris (2013). *Can machines learn to predict a violent conflict?*
- Roberts, David R., Volker Bahn, Simone Ciuti, et al. (2017). “Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure”. In: *Ecography* 40.8, pp. 913–929. DOI: 10.1111/ecog.02881.
- Sundberg, Ralph and Erik Melander (2013). “Introducing the UCDP Georeferenced Event Dataset”. In: *Journal of Peace Research* 50.4, pp. 523–532. DOI: 10.1177/0022343313484347.
- von Uexkull, Nina et al. (2016). “Civil conflict sensitivity to growing-season drought”. In: *Proceedings of the National Academy of Sciences* 113.44, pp. 12391–12396. DOI: 10.1073/pnas.1607542113.
- Ward, Michael D., Brian D. Greenhill, and Kristin M. Bakke (2010). “The perils of policy by p-value: Predicting civil conflicts”. In: *Journal of Peace Research* 47.4, pp. 363–375. DOI: 10.1177/0022343310362830.