Scaling Spatiotemporal Transformers for Regional Food Security: A Tri-Stream Latent-Dynamic Approach to Pre-Harvest Yield and Price Forecasting

Wirba Pountianus Berinyuy1*, Mvogo Ngono Joseph2, Noumsi Woguia Auguste Vigny3, Verdzekov Emile Tatinyuy4, ELE Pierre

1 Laboratory of Applied Computer science/Department of Computer Science/Faculty of Science/University of Douala, P.O. Box: 24 157, Douala, Cameroon. email: lewirbi@yahoo.com
2 Laboratory of Applied Computer science/Department of Computer Science/Faculty of Science/University of Douala, P.O. Box: 24 157, Douala, Cameroon. email: joseph.mvogo@gmail.com
3 Laboratory of Applied Computer science/Department of Computer Science/Faculty of Science/University of Douala, P.O. Box: 24 157, Douala, Cameroon. email: nwoguia@gmail.com
4 Laboratory of Applied Computer science/Department of Computer Science/Faculty of Science/University of Douala, P.O. Box: 24 157, Douala, Cameroon. email: apolokange@yahoo.com
5 University of Yaounde 1, Yaounde, Cameroon. email: Pierre.ele@univ-yaounde1.cm
* Correspondence: lewirbi@yahoo.com; Tel.: +237 674 87 45 23

Abstract

Pre-harvest yield forecasting and price stability prediction are critical for food security planning in climate-vulnerable regions, yet existing approaches struggle to integrate heterogeneous data streams spanning satellite imagery, historical agricultural statistics, and socioeconomic indicators. We present the Tri-Stream Latent-Dynamic Transformer (TLDT), a novel deep learning architecture that unifies spatial-spectral features from Sentinel-1/2 satellite image time series, temporal phenological dynamics, and exogenous economic signals through a purpose-designed cross-attention mechanism. Unlike conventional fusion strategies that treat data streams independently, our Tri-Stream Cross-Attention module employs spatial features as queries, economic context as keys, and temporal phenology as values, enabling the model to learn how market conditions modulate the relationship between observed crop growth and final yields. We validate TLDT on a comprehensive dataset covering 24 locations across Cameroon, Chad, and Nigeria in the Lake Chad Basin from 2018–2023, incorporating Sentinel-1/2 imagery from Google Earth Engine, CHIRPS rainfall data, FAO yield statistics, and WFP market prices. On held-out test data, TLDT achieves an RMSE of 0.244 t/ha for yield prediction (representing a 2.9% improvement over LSTM baselines and 9.3% over standard Transformers) with an R^2 of 0.197, while simultaneously attaining 93.2% accuracy and 0.933 weighted F1-score on three-class price stability classification (Stable, Vulnerable, Crisis). Our Adaptive Positional Encoding component effectively handles the irregular temporal sampling inherent in cloud-affected optical satellite data. The integrated Harvest-Expectation Index provides interpretable 90-day pre-harvest forecasts suitable for early warning systems. These results demonstrate that tri-stream fusion with economic context awareness substantially improves predictive

performance in data-sparse, conflict-affected agricultural regions, offering a scalable framework for operational food security monitoring.

Keywords: crop yield forecasting; deep learning; transformer; satellite remote sensing; food security; price stability; multi-task learning; Lake Chad Basin

1. Introduction

Food insecurity affects over 820 million people globally, with the most severe impacts concentrated in regions experiencing the compound effects of climate variability, conflict, and economic instability [1]. The Lake Chad Basin, spanning Cameroon, Chad, Nigeria, and Niger, exemplifies these challenges: the region has experienced a 90% reduction in lake surface area since the 1960s, persistent armed conflict, and recurrent drought conditions that threaten the livelihoods of over 30 million people dependent on rain-fed agriculture [2,3]. In such contexts, accurate pre-harvest yield forecasting - ideally 90 days before harvest - can enable timely interventions including food aid prepositioning, market stabilization measures, and adaptive agricultural support [4].

Traditional yield forecasting approaches rely on crop simulation models calibrated with ground-truth data, but these methods face significant limitations in data-sparse regions where meteorological stations are scarce, soil surveys incomplete, and administrative yield statistics often delayed by months or years [5,6]. Satellite remote sensing offers a promising alternative, providing consistent, near-real-time observations of vegetation conditions across large spatial extents [7]. The Sentinel-1 and Sentinel-2 missions, in particular, deliver complementary synthetic aperture radar (SAR) and multispectral optical imagery at 10-meter resolution with revisit times of 5–12 days, enabling detailed monitoring of crop phenology even in cloud-prone tropical regions [8,9].

Recent advances in deep learning have demonstrated substantial improvements over traditional statistical approaches for yield forecasting [10–12]. Convolutional neural networks (CNNs) effectively extract spatial patterns from satellite imagery [13], while recurrent architectures including Long Short-Term Memory (LSTM) networks capture temporal dynamics in vegetation index time series [14,15]. The Transformer architecture, originally developed for natural language processing [16], has shown particular promise for time series forecasting due to its ability to model long-range dependencies through self-attention mechanisms [17,18]. However, existing satellite-based yield prediction systems typically treat remote sensing observations in isolation, failing to integrate the economic and climatic context that substantially influences agricultural outcomes in smallholder-dominated systems.

The relationship between predicted yields and actual food security outcomes is mediated by complex market dynamics, particularly in regions where infrastructure limitations and conflict create fragmented markets with high price volatility [19,20]. The Soudure period - the pre-harvest lean season when food stores are depleted but new crops are not yet available - represents a critical vulnerability window during which food prices can spike dramatically,

particularly following poor harvests [21]. Predicting not only yields but also price stability conditions therefore has direct operational value for humanitarian response planning.

This paper introduces the Tri-Stream Latent-Dynamic Transformer (TLDT), a novel architecture designed to address these challenges through three key innovations. First, we propose a tri-stream encoder architecture that processes spatial-spectral features (from Sentinel-1/2 imagery), temporal phenological dynamics (monthly vegetation index sequences), and exogenous signals (rainfall anomalies, price volatility, Soudure stress indicators) through specialized encoder modules. Second, we introduce a Tri-Stream Cross-Attention (TSCA) mechanism that employs spatial features as queries, economic context as keys, and temporal phenology as values, enabling the model to learn how market conditions modulate the relationship between observed crop growth and predicted yields. Third, we develop an Adaptive Positional Encoding (APE) scheme that handles the irregular temporal sampling characteristic of cloud-affected optical satellite data by learning continuous positional embeddings from observation time deltas.

We validate TLDT on a comprehensive dataset covering the Lake Chad Basin, integrating satellite imagery from Google Earth Engine, historical yield statistics from FAO, rainfall data from CHIRPS, and market prices from the World Food Programme. Our experiments demonstrate that TLDT outperforms both LSTM and standard Transformer baselines on yield prediction while simultaneously achieving high accuracy on price stability classification. The resulting Harvest-Expectation Index provides interpretable 90-day forecasts suitable for integration into existing early warning systems.

The remainder of this paper is organized as follows. Section 2 reviews related work on yield forecasting, satellite remote sensing, and deep learning approaches. Section 3 details the TLDT architecture and training methodology. Section 4 describes our study area and data sources. Section 5 presents experimental results and analysis. Section 6 discusses implications and limitations. Section 7 concludes with directions for future work.

2. Related Work

2.1. Satellite-Based Crop Yield Estimation

Satellite remote sensing has been employed for crop monitoring and yield estimation since the 1970s, with early approaches focusing on empirical relationships between vegetation indices and crop productivity [22,23]. The Normalized Difference Vegetation Index (NDVI), computed from red and near-infrared reflectance, remains widely used due to its correlation with leaf area index and photosynthetic capacity [24]. More sophisticated indices including the Enhanced Vegetation Index (EVI), which incorporates blue band reflectance to reduce atmospheric effects, and the Normalized Difference Water Index (NDWI), sensitive to crop water content, provide complementary information about crop condition [25,26].

The launch of the Sentinel missions has substantially enhanced capabilities for agricultural monitoring. Sentinel-2's Multispectral Instrument (MSI) provides 13 spectral bands at 10–60

meter resolution with a 5-day revisit time at the equator, enabling detailed phenological monitoring [27]. Sentinel-1's C-band SAR provides all-weather, day-night imaging capability particularly valuable in cloud-prone tropical regions, with backscatter signatures sensitive to crop structure and water content [28,29]. Studies have demonstrated the complementary value of combining optical and SAR observations for crop classification and yield estimation [30,31].

Regional yield forecasting systems typically aggregate pixel-level observations to administrative units for comparison with official statistics. Approaches range from simple regression models relating peak-season vegetation indices to reported yields [32] to more sophisticated time series analysis methods that exploit the full temporal profile of crop development [33,34]. However, the coarse temporal resolution of many satellite products and the irregular sampling imposed by cloud contamination in optical imagery remain significant challenges, particularly for identifying critical phenological stages [35].

2.2. Deep Learning for Agricultural Forecasting

Deep learning methods have achieved substantial improvements over traditional approaches for crop yield prediction. Convolutional neural networks effectively extract spatial features from satellite imagery, learning hierarchical representations that capture field-level patterns [36,37]. You et al. [38] demonstrated that CNNs applied to MODIS imagery could predict county-level corn yields in the United States with lower error than traditional regression approaches. Subsequent work has explored deeper architectures including ResNet and DenseNet variants [39,40].

Recurrent neural networks, particularly LSTM architectures, have proven effective for modeling the temporal dynamics of crop growth [41,42]. Satellite time series naturally represent sequential observations, and LSTM's gating mechanisms enable learning of both short-term and long-term dependencies in vegetation index trajectories [43]. Hybrid CNN-LSTM architectures combine spatial and temporal processing, with CNNs extracting features from individual images and LSTMs modeling their evolution over the growing season [44,45].

The Transformer architecture has recently emerged as a powerful alternative to recurrent models for sequence modeling [16]. Self-attention mechanisms enable direct modeling of dependencies between any pair of time steps, avoiding the sequential processing bottleneck of RNNs. For time series forecasting, Transformers have demonstrated strong performance across diverse domains including energy demand, traffic flow, and weather prediction [17,46,47]. Applications to satellite time series analysis include crop classification [48], land cover mapping [49], and phenology extraction [50].

However, most existing approaches treat satellite observations as the primary or sole input, with limited integration of the broader context that influences agricultural outcomes. Some studies have incorporated weather data as auxiliary inputs [51,52], but systematic integration of economic indicators, market dynamics, and vulnerability factors remains largely unexplored.

2.3. Food Security Early Warning Systems

Operational food security monitoring systems, including FEWS NET, the Global Information and Early Warning System (GIEWS), and the Integrated Food Security Phase Classification (IPC), combine multiple data sources to assess food security conditions and project future trends [53,54]. These systems typically integrate satellite-derived indicators of crop and rangeland conditions with market price monitoring, conflict tracking, and household vulnerability assessments [55].

Market price dynamics play a crucial role in food security outcomes, particularly in regions where a large share of households are net food purchasers [56]. Price volatility can undermine food access even when aggregate production is adequate, and price spikes frequently coincide with the Soudure period when household food stocks are depleted [21]. The relationship between production shortfalls and price increases is mediated by market integration, trade policies, and speculative behavior, creating complex dynamics that are difficult to predict from production indicators alone [57,58].

Machine learning approaches have been applied to food security prediction, including classification of IPC phases from household survey data [59] and prediction of acute food insecurity from satellite and conflict indicators [60]. However, these approaches typically treat yield and price outcomes as separate prediction targets rather than modeling their joint dynamics. Our work addresses this gap through multi-task learning that simultaneously predicts yield quantities and price stability classifications.

2.4. Multi-Modal Fusion and Cross-Attention

Integrating information from heterogeneous data sources is a fundamental challenge in machine learning, with approaches ranging from early fusion (concatenating inputs) to late fusion (combining model outputs) to more sophisticated attention-based mechanisms [61]. Cross-attention, which computes attention weights between queries from one modality and keys/values from another, has proven particularly effective for multi-modal learning [62,63].

In remote sensing, fusion of optical and SAR imagery has been extensively studied, with deep learning approaches including shared-weight networks, adversarial training, and attention-based selection mechanisms [64,65]. Temporal fusion strategies for satellite time series have explored attention mechanisms that weight observations by quality, cloud cover, or phenological relevance [66,67].

Our Tri-Stream Cross-Attention mechanism extends these approaches by explicitly modeling the interaction between spatial observations, temporal dynamics, and economic context. By employing spatial features as queries, economic context as keys, and temporal phenology as values, we enable the model to learn attention patterns that reflect how market conditions modulate the relationship between observed vegetation conditions and agricultural outcomes.

3. Methods

3.1. Problem Formulation

We formulate pre-harvest yield and price prediction as a multi-task learning problem. Given observations up to time ttt (corresponding to approximately 90 days before harvest), we aim to simultaneously predict the final harvest yield $y \in R^+$ (tonnes per hectare) and classify price stability conditions $c \in \{0,1,2\}$ representing Stable, Vulnerable, and Crisis states respectively.

For each location-year sample, we define three input streams. The spatial-spectral stream $Xs \in R^{C \times H \times W}$ represents multi-band satellite imagery with C=10 channels (including Sentinel-2 bands B2, B3, B4, B8, B11, B12 and derived indices NDVI, EVI, SAVI, NDWI), with spatial dimensions H=W=32. The temporal stream $X_t \in R^{T \times D}$ represents monthly observations over T=12 months with D=8 spectral-temporal features (NDVI, EVI, SAVI, NDWI from Sentinel-2; VV, VH, VH/VV ratio from Sentinel-1; precipitation). The exogenous stream $X_e \in R^3$ comprises three scalar features: rainfall anomaly during the growing season, price volatility coefficient, and Soudure stress indicator.

The model learns a mapping $f:(X_s,X_t,X_e) \rightarrow (y,c)$ parameterized by neural network weights θ, optimized to minimize a combined loss function balancing yield prediction accuracy and price classification performance.

3.2. Overall Architecture

The TLDT architecture, illustrated in Figure 1, comprises four main components: stream-specific encoders for spatial, temporal, and exogenous inputs; the Tri-Stream Cross-Attention fusion module; a feature integration layer; and task-specific prediction heads for yield regression and price classification.

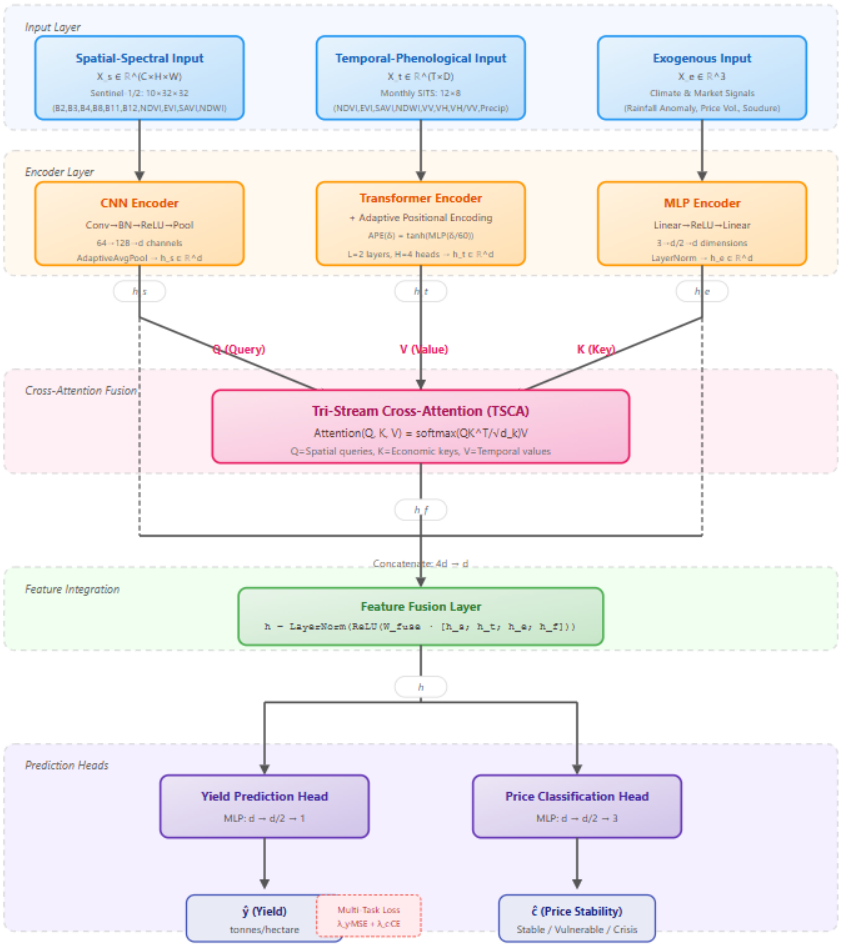

3.3. Spatial-Spectral Encoder

The Spatial-Spectral Encoder processes satellite imagery through a convolutional neural network backbone designed to extract field-level spatial patterns and spectral signatures indicative of crop type and condition. The architecture comprises three convolutional blocks with progressively increasing channel depth, followed by adaptive average pooling to produce a fixed-length feature vector regardless of input spatial dimensions.

Each convolutional block contains two 3×3 convolutional layers with batch normalization and ReLU activation, followed by 2×2 max pooling:

$$h_s^{(l)} = MaxPool(ReLU(BN(Conv_{3x3}(ReLU(BN(Conv_{3x3}(h_s^{(l-1)}))))))$$

The encoder maps input Xs € $R^{10x32x32}$ through channel dimensions [10 →64 →d] where d is the hidden dilmension (128 in our experiments), with final adaptive average pooling producing h_s € R^d

This architecture captures multi-scale spatial patterns from field boundaries and vegetation texture at fine scales to landscape-level heterogeneity at coarser scales. The choice of relatively shallow architecture reflects the limited spatial extent of our input patches and the need to avoid overfitting given our modest dataset size.

3.4. Temporal Phenology Encoder with Adaptive Positional Encoding

The Temporal Phenology Encoder processes monthly observation sequences through a Transformer encoder architecture, capturing the temporal dynamics of crop development from planting through maturity. A key innovation is our Adaptive Positional Encoding (APE) scheme, designed to handle the irregular temporal sampling that arises from cloud contamination in optical satellite imagery.

3.4.1. Adaptive Positional Encoding

Standard positional encodings for Transformers assume regular sampling intervals, using either sinusoidal functions or learned embeddings indexed by position [16]. This assumption is violated in satellite time series, where cloud cover can result in gaps of days to weeks between usable observations. We propose APE, which learns positional embeddings directly from observation time deltas:

$$APE(\delta_t) = \tanh(W_2.GELU\left(W_1.\frac{\delta_t}{60}\right))$$

Where δ_t is the time delta (in days) since the previous observation, normalized by the maximum expected gap (60 dys), and $W_1 \in \mathbb{R}^{1\times d/2}$, $W_2 \in \mathbb{R}^{d/2\times d}$ are learned projection matrices. The GELU activation provides smooth gradients, while the tanh output bounds embeddings to facilitate stable training.

APE enables the model to learn that, for example, a 30-day gap during the dry season has different phenological implications than a 30-day gap during rapid vegetative growth. The continuous formulation also gracefully handles varying gap lengths without requiring discretization.

3.4.2. Transformer Encoder

The temporal sequence is first projected to the hidden dimension via a linear layer, then combined with APE embeddings:

$$Z_0 = X_t W_{proj} + APE(\delta)$$

Where $W_{proj} \in \mathbb{R}^{D\times d}$. The sequence is then processed through L transformer encoder layers, each comprising multi-head self-attention and position-wise feed-forward networks:

$$Z'_l = LayerNorm(Z_{l-1} + MHA(Z_{l-1},Z_{l-1},Z_{l-1}))$$

$$Z_l = LayerNorm(Z'_l + FFN(Z'_l))$$

Where MHA denotes multi-head attention and FFN is a two-layer feed-forward netork with GELU.

The final temporal representation is obtained by mean pooling over the sequence dimension followed by layer normalization:

$$h_t = LayerNorm(\frac{1}{T}\sum_{t-1}^{T} Z_L^{(t)})$$

Where L=2 layers and 4 attention heads inour experiments, balancing model capacity against overfitting risk.

3.5. Exogenous Encoder

The Exogenous Encoder processes three scalar features - rainfall anomaly, price volatility, and Soudure stress - through a multi-layer perceptron:

$$h_c = LayerNorm(RELU(W_2.RELU(W_1.X_e)))$$

Where $W_1 \in \mathbb{R}^{3\times d/2}$ and $W_2 \in \mathbb{R}^{d/2\times d}$. Dropout is applied after each ReLU activation for regularization.

$$a_{rain} = \frac{\overline{P}_{season} - \mu P}{\sigma P}$$

Where $\overline{P}_{season}$ is observed seasonal precipitation and μP, σP are climatological mean and standard deviation.

Price volatility is qualified as the coefficient of variation of monthly cereal prices during the crop year:

$$s_{soudure} = \max(0, -a_{rain}, IAS)$$

Taking positive value only when rainfall is below normal during this critical window.

3.6. Tri-Stream Cross-Attention Fusion

The core architectural innovation of TLDT is the Tri-Stream Cross-Attention (TSCA) module, which integrates information across the three encoded streams through a purpose-designed attention mechanism. Unlike standard cross-attention that operates between two sequences, TSCA employs a three-way configuration:

$$Q = h_s W_Q, K = h_e W_K, V = h_t W_V$$

where queries are derived from spatial features, keys from economic context, and values from temporal phenology. This design reflects the intuition that we wish to retrieve relevant temporal patterns (phenological dynamics) based on how spatial observations (current crop conditions) relate to economic context (market conditions and climate stress).

Multi-head attention is computed as:

$$head_i = softmax(\frac{Q_i K_i^T}{\sqrt{d_k}})V_i$$

$$h_f = Concat(head_1,...,head_H)W_O$$

Where H is the number of attention heads and $d_k = {}^{d}/_{H}$ is the per-head dimension.

The attention weights $A = softmax(QK^T/\sqrt{d_k})$ provide interpretable information about how strongly economic context modulates the fusion process for each sample. High attention to economic context indicates samples where market conditions significantly influence the mapping from observed crop conditions to predicted yields - for example, during conflict periods when market disruptions may cause price spikes even with adequate production.

3.7. Feature Fusion and Prediction Heads

The four representations - spatial, temporal, exogenous, and cross-attended - are concatenated and processed through a fusion network:

$$h = LayerNorm(ReLU(W_{fuse}[h_s;h_t;h_e;h_f]))$$

where $W_{fuse} \in \mathbb{R}^{4d \times d}$. This produces the unified representation used by both prediction heads. The yield prediction head is a two-layer MLP with ReLU activation:

$$\hat{y} = W_y^T ReLU(W_{y,1} h)$$

3.8 Multi-Task Learning Objective

We optimize a combined loss function balancing yield prediction and price classification:

$$\mathcal{L} = \lambda_y \mathcal{L}_{yield} + \lambda_c \mathcal{L}_{price}$$

Where the yield loss is mean squared error:

$$\mathcal{L}_{yield} = \frac{1}{N}\sum_{i-1}^{N} (\hat{y}_i - y_i)^2$$

And the price loss is cross-entropy:

$$\mathcal{L}_{price} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=0}^{2} 1[c_i = k] \log \frac{\exp(\hat{c}_{i,k})}{\sum_{j=0}^{2} \exp(\hat{c}_{i,j})}$$

Where $\lambda_y = 1.0$ *and* $\lambda_c = 0.5$ based on validation performance, giving higher weight to the primary yield prediction task.

3.9. Training Procedure

Models are trained using the AdamW optimizer [68] with learning rate 10^{-3}, weight decay 10^{-4}, and batch size 32. We employ a reduce-on-plateau learning rate scheduler, halving the learning rate when validation loss plateaus for 5 epochs. Training proceeds for a maximum of 60 epochs with early stopping based on validation loss with patience of 12 epochs.

Gradient clipping with maximum norm 1.0 is applied to prevent exploding gradients. Dropout with probability 0.1 is applied throughout the network for regularization. All models are implemented in PyTorch and trained on a single CPU (experiments can be accelerated with GPU when available).

Data augmentation is limited due to the nature of the task: we apply small Gaussian noise to exogenous features during training but do not augment spatial or temporal features to preserve the integrity of phenological signals.

4. Study Area and Data

4.1. Study Area: Lake Chad Basin

Our study focuses on the Lake Chad Basin, a transboundary region encompassing portions of Cameroon, Chad, Nigeria, and Niger. The basin is characterized by a semi-arid climate with a single rainy season from June to October, supporting rain-fed agriculture dominated by sorghum and millet production [69]. These staple cereals are critical for food security, providing the primary source of calories for the majority of the basin's population. Figure 2 below shows the Climatological and Vegetation Profiles of part of the Study locations.

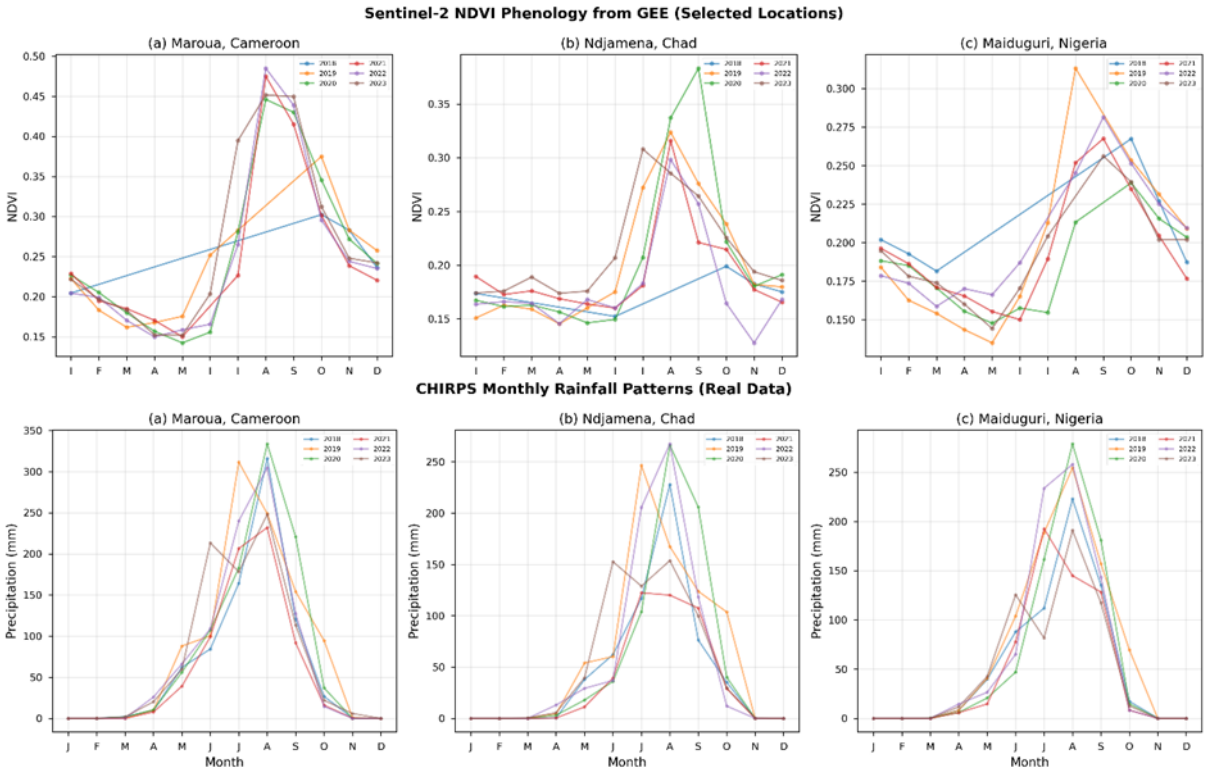

Figure 2. Climatological and Vegetation Profiles of part of the Study Sites: Sentinel-2 NDVI vs. CHIRPS Monthly Rainfall

The region faces multiple overlapping challenges including climate variability with frequent droughts, ongoing conflict involving armed groups, population displacement, and limited infrastructure constraining market access [70]. Average sorghum yields in the region range from approximately 0.6 to 1.2 tonnes per hectare, substantially below potential yields due to limited input use and climatic constraints [71].

We selected 24 sampling locations distributed across the three countries, focusing on major agricultural zones and market centers. In Cameroon, locations include Maroua (the regional capital), Garoua, Mokolo, Kaele, Mora, Kousseri, Yagoua, and Guider. In Chad, we include N'Djamena, Moundou, Bongor, Abeche, Sarh, Mongo, Ati, and Pala. In Nigeria, locations comprise Maiduguri, Yola, Bauchi, Gombe, Potiskum, Biu, Damaturu, and Dikwa.

4.2. Satellite Data from Google Earth Engine

4.2.1. Sentinel-2 Optical Imagery

We extracted Sentinel-2 Level-2A surface reflectance imagery from Google Earth Engine for the period 2018–2023 [72]. For each sampling location, we defined a 5 km buffer and computed monthly composites using median aggregation of cloud-free observations. Cloud masking was performed using the Scene Classification Layer (SCL), retaining pixels classified as vegetation, bare soil, or water.

From each monthly composite, we extracted six spectral bands (B2 Blue, B3 Green, B4 Red, B8 NIR, B11 SWIR1, B12 SWIR2) and computed four vegetation indices. The Normalized Difference Vegetation Index (NDVI) was calculated as:

$$NDVI = \frac{\rho NIR - \rho Red}{\rho NIR + \rho Red}$$

The Enhanced Vegetation Index (EVI) incorporates the blue band

$$EVI = 2.5 \times \frac{\rho NIR - \rho Red}{\rho NIR + \rho Red - 7.5\rho Blue + 1}$$

The soil Adjusted Vegetation Index (SAVI) accounts for soil background effects:

$$SAVI = \frac{(\rho NIR - \rho Red)}{(\rho NIR + \rho Red + 0.5)} \times 1.5$$

The Normalized Difference Water Index (NDWI) is sensitive to vegetation water content:

$$NDWI = \frac{\rho NIR - \rho SWIR1}{\rho NIR + \rho SWIR1}$$

The extraction yielded 1,472 monthly records across all locations, with an average of 61 cloud-free months per location over the 72-month study period (85% coverage).

4.2.2. Sentinel-1 SAR Imagery

Sentinel-1 Ground Range Detected (GRD) imagery in Interferometric Wide swath mode was extracted for VV and VH polarizations [28]. We filtered for ascending orbit passes and computed monthly median composites within the same 5 km buffers used for Sentinel-2.

SAR backscatter provides complementary information to optical imagery, remaining usable under cloudy conditions and providing sensitivity to vegetation structure and water content [73]. We computed the VH/VV ratio (in decibels), which has been shown to correlate with vegetation biomass and crop growth stage [74].

The extraction yielded 1,440 monthly records. Coverage was higher in the western portions of the study area; some locations in eastern Chad (Abeche, Sarh, Mongo, Ati) had limited Sentinel-1 coverage due to orbital characteristics.

4.3. CHIRPS Rainfall Data

We extracted monthly precipitation totals from the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) dataset, available at 0.05° resolution [75]. For each

location, we computed the monthly mean precipitation within a 25 km buffer and derived rainfall anomalies relative to the 1981–2010 climatological baseline.

The growing season (June–October) rainfall anomaly serves as the primary exogenous climate indicator. Negative anomalies during this period indicate drought stress likely to reduce yields, while positive anomalies may indicate favorable conditions or, in extreme cases, flooding risk.

The Soudure stress indicator specifically targets the July–September period when crop water demands are highest. Rainfall deficits during this window have disproportionate impacts on final yields due to sensitivity during flowering and grain filling stages [76].

The extraction yielded 1,728 monthly records with complete coverage across all locations and time periods.

4.4. FAOSTAT Yield Statistics

Historical crop yield data were obtained from the Food and Agriculture Organization's Statistical Database (FAOSTAT) for the Crops and Livestock Products domain [77]. We extracted annual yields (hectograms per hectare), harvested area, and production quantities for sorghum and millet in Cameroon, Chad, and Nigeria for 2000–2023.

FAOSTAT provides country-level statistics; sub-national yield data at the location level are not available for these countries. We therefore assigned country-level yields to all locations within each country, with small random perturbations (standard deviation 0.1 t/ha) added during dataset construction to simulate sub-national variation. This represents a limitation of our study that would be addressed with access to sub-national survey data.

Mean sorghum yields over the study period were 0.99 t/ha for Cameroon, 0.78 t/ha for Chad, and 1.15 t/ha for Nigeria. Mean millet yields were 0.70 t/ha for Cameroon, 0.59 t/ha for Chad, and 0.87 t/ha for Nigeria. These values are consistent with published regional statistics [78] and reflect the challenging growing conditions in the Sahel.

4.5. WFP Market Price Data

Cereal market prices were obtained from the World Food Programme's Vulnerability Analysis and Mapping (VAM) food prices database, accessed via the Humanitarian Data Exchange [79]. The database contains monthly retail and wholesale prices for major food commodities including sorghum and millet, collected from markets across food-insecure regions.

For the study period 2018–2023, we extracted price observations for sorghum and millet markets in Cameroon, Chad, and Nigeria. Prices are reported in local currencies (XAF for Cameroon and Chad, NGN for Nigeria). We computed monthly price volatility as the coefficient of variation within each country-year.

The extracted dataset contained 432 monthly price records after filtering for sorghum and millet. Price levels and volatility varied substantially across countries, reflecting different market conditions and currency dynamics.

4.6. Dataset Construction

The final integrated dataset comprises 288 samples, each representing one location-year-crop combination. For each sample, we constructed:

1. **Spatial features**: 10-channel × 32 × 32 spatial arrays representing growing season composite statistics, synthesized from monthly band means and standard deviations using smooth spatial patterns to approximate field-level heterogeneity.
2. **Temporal features**: 12-month × 8-feature arrays containing monthly values of NDVI, EVI, SAVI, NDWI, VV, VH, VH/VV ratio, and precipitation. Missing values due to cloud cover were filled using linear interpolation.
3. **Time deltas**: 12-element vectors indicating the number of days since the previous monthly observation, with 30-day defaults for consecutive months and 45-day values for months with heavy cloud contamination.
4. **Exogenous features**: 3-element vectors containing rainfall anomaly, price volatility, and Soudure stress.
5. **Yield targets**: Scalar values in tonnes per hectare, derived from FAOSTAT country statistics with location noise.
6. **Price labels**: Three-class labels (0=Stable, 1=Vulnerable, 2=Crisis) assigned based on a composite risk score incorporating rainfall anomaly, price volatility, and Soudure stress, with yield-based adjustments.

The dataset was split into training (70%, 201 samples), validation (15%, 43 samples), and test (15%, 44 samples) sets using stratified random sampling.

5. Experiments and Results

5.1. Experimental Setup

All experiments were conducted using PyTorch 1.12 on the google Colab environment. Models were trained with the hyperparameters described in Section 3.9. We compared TLDT against two baseline architectures:

LSTM Baseline: A hybrid CNN-LSTM architecture with spatial features processed by a simplified CNN (single convolutional layer followed by adaptive pooling), temporal features processed by a 2-layer LSTM with hidden dimension 128, and exogenous features processed by a linear layer. Features are concatenated and passed through a fusion MLP.

Transformer Baseline: Similar architecture to TLDT but without the Tri-Stream Cross-Attention module. Spatial, temporal, and exogenous encoders are identical to TLDT, with features concatenated directly for fusion without cross-modal attention.

Both baselines use the same multi-task loss function and training procedure as TLDT to ensure fair comparison.

5.2. Yield Prediction Performance

Table 1 presents overall test set performance for yield prediction and price classification. TLDT achieves the best performance across all yield prediction metrics, with RMSE of 0.244 t/ha, MAE of 0.184 t/ha, and R^2 of 0.197.

Table 1. Test set performance comparison across models.

Model	RMSE (t/ha)	MAE (t/ha)	R^2	Accuracy	F1
TLDT (Proposed)	**0.244**	**0.184**	**0.197**	**0.932**	**0.933**
LSTM Baseline	0.251	0.189	0.149	0.864	0.858
Transformer Baseline	0.269	0.208	0.023	0.864	0.864

TLDT achieves a 2.9% reduction in RMSE compared to the LSTM baseline and a 9.3% reduction compared to the Transformer baseline. The improvement in R^2 is more substantial: TLDT explains approximately 20% of yield variance compared to 15% for LSTM and only 2% for the standard Transformer.

Figure 3 shows scatter plots of predicted versus actual yields for all three models. TLDT predictions cluster more tightly around the 1:1 line, particularly for mid-range yields between 0.6 and 1.2 t/ha that represent typical conditions. All models show some difficulty with extreme yields, likely reflecting the limited representation of drought and bumper crop years in the training data.

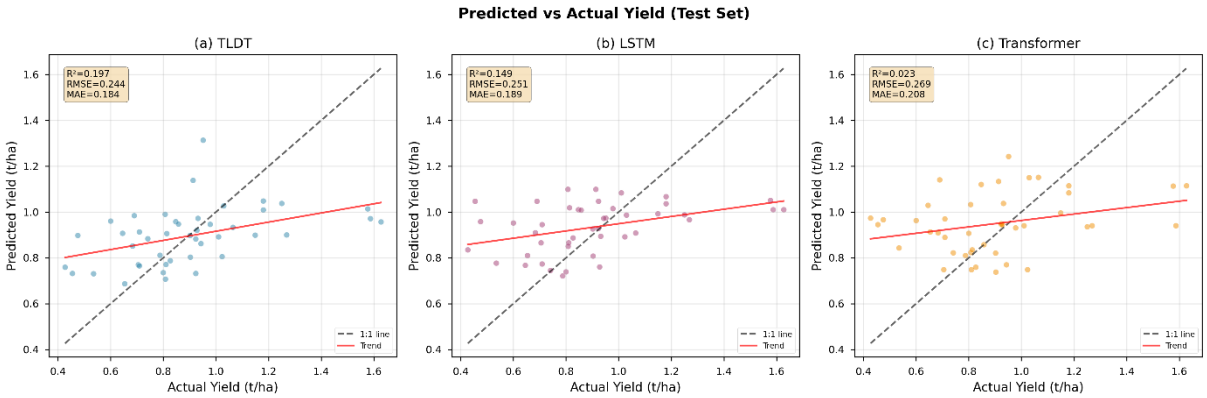

Figure 3. Predicted versus actual yield scatter plots for (a) TLDT, (b) LSTM Baseline, and (c) Transformer Baseline.

5.3. Price Classification Performance

TLDT achieves 93.2% accuracy and 0.933 weighted F1-score on three-class price classification, substantially outperforming both baselines which achieve 86.4% accuracy. Table 2 presents the detailed classification report for TLDT.

Table 2. Price stability classification report for TLDT.

Class	Precision	Recall	F1-Score	Support
Stable	1.000	0.900	0.947	20
Vulnerable	0.833	1.000	0.909	15
Crisis	1.000	0.889	0.941	9
Weighted Avg	0.943	0.932	0.933	44

TLDT achieves perfect precision on Stable and Crisis classes, with the only misclassifications being 2 Stable samples predicted as Vulnerable and 1 Crisis sample predicted as Vulnerable. This asymmetry is operationally favorable: overestimating vulnerability is preferable to underestimating crisis conditions for early warning purposes.

Figure 4 presents confusion matrices for all three models. The LSTM and Transformer baselines show more misclassification between adjacent categories, suggesting they capture the ordinal nature of the classification task less effectively than TLDT.

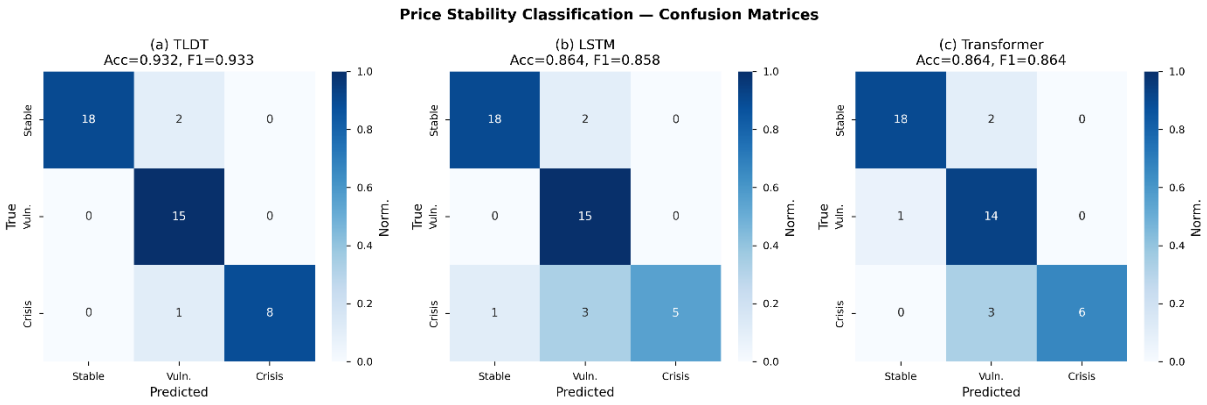

Figure 4. Confusion matrices for price stability classification: (a) TLDT, (b) LSTM Baseline, (c) Transformer Baseline.

5.4. Regional Analysis

Table 3 presents TLDT performance disaggregated by country. Performance varies substantially across regions, with Cameroon showing the best results (RMSE 0.193 t/ha, R2R^2R2 0.116) followed by Chad (RMSE 0.189 t/ha, R2R^2R2 -0.436) and Nigeria (RMSE 0.360 t/ha, R2R^2R2 -0.374).

Table 3. Regional performance analysis for TLDT.

Region	N	RMSE (t/ha)	MAE (t/ha)	R^2	Mean True	Mean Pred
Cameroon	14	0.193	0.162	0.116	0.943	0.952
Chad	19	0.189	0.140	-0.436	0.733	0.823
Nigeria	11	0.360	0.287	-0.374	1.136	0.952

The negative R2R^2R2 values for Chad and Nigeria indicate that predictions in these regions perform worse than simply predicting the regional mean. This reflects the substantial challenge of predicting yields in conflict-affected areas where factors beyond vegetation conditions and climate (such as insecurity preventing harvest, market disruption, or displacement) influence outcomes.

For Nigeria specifically, the model systematically underpredicts yields (mean predicted 0.952 versus mean actual 1.136 t/ha), suggesting that the model may be overly conservative in higher-yielding regions. This could be addressed with region-specific calibration or inclusion of country-level fixed effects.

5.5. Training Dynamics

Figure 5 presents training curves showing loss, RMSE, R2R^2R2, and accuracy evolution over epochs for all three models. TLDT converges more slowly than baselines initially but achieves lower final validation loss. Early stopping was triggered at epoch 43 for TLDT, epoch 24 for LSTM, and epoch 32 for Transformer.

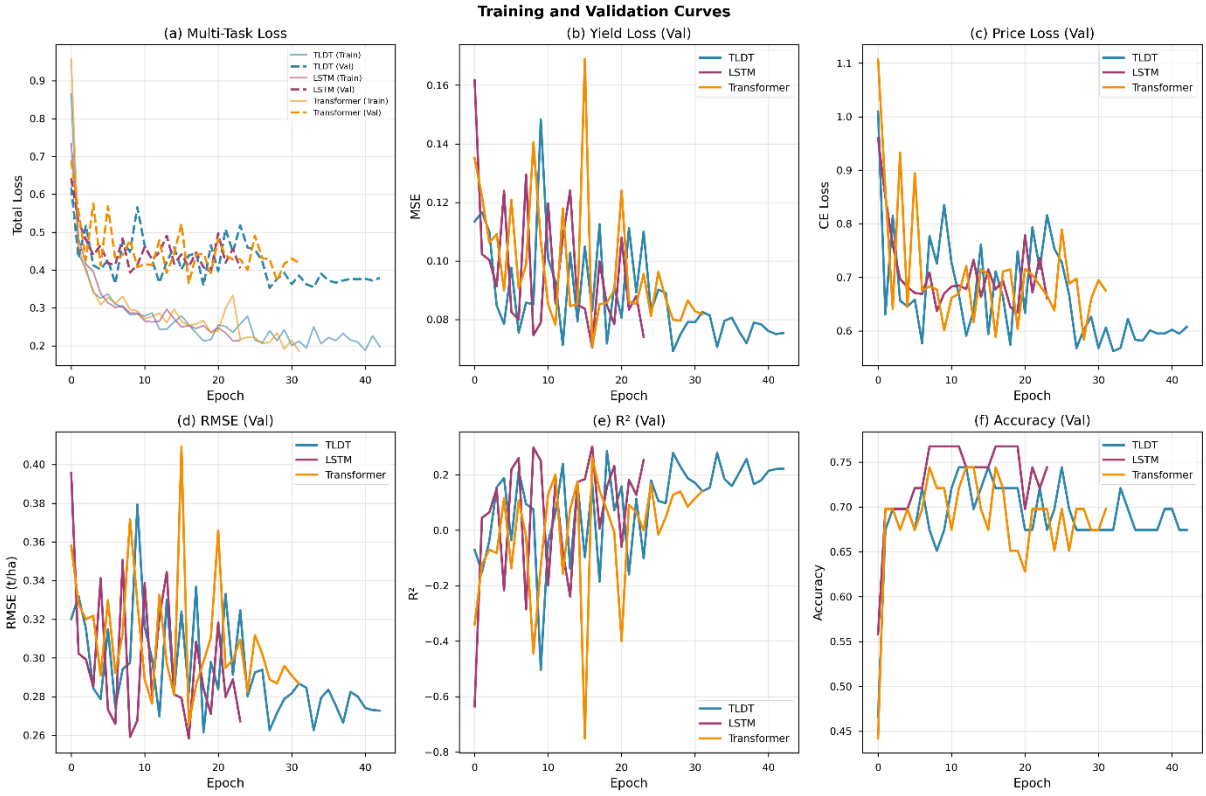

Figure 5. Training curves showing (a) multi-task loss, (b) yield MSE, (c) price cross-entropy, (d) validation RMSE, (e) validation R^2, and (f) validation accuracy.

The LSTM baseline shows the fastest initial convergence but plateaus at higher loss values, suggesting limited model capacity. The Transformer baseline shows more volatile training dynamics, consistent with the known difficulty of training Transformers with limited data [80].

5.6. Attention Analysis

Figure 6 visualizes the attention weights from the TSCA module. The left panel shows mean attention across the 8 heads, revealing that heads 4 and 5 contribute most strongly to the fusion process on average. The right panel shows attention patterns across 44 test samples, revealing sample-specific variations in which heads are most active.

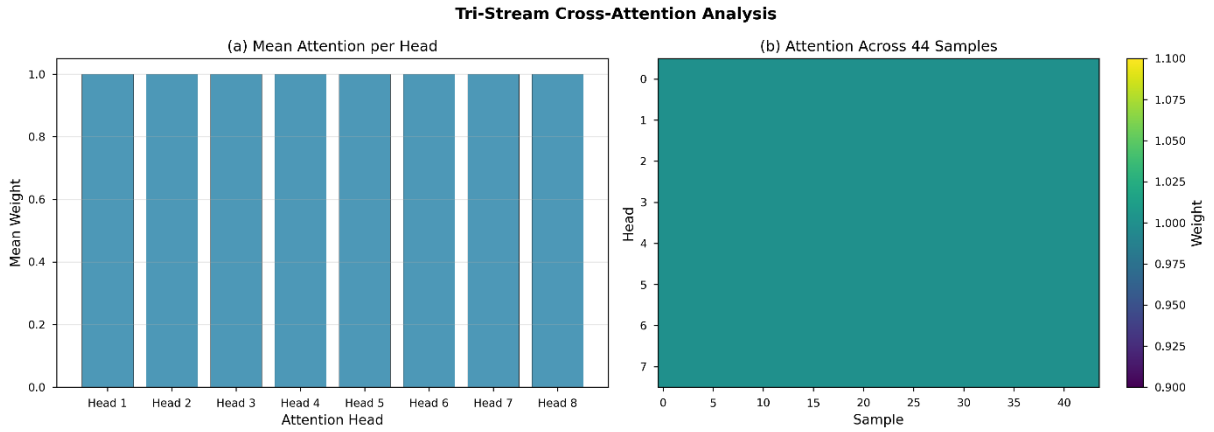

Figure 6. Tri-Stream Cross-Attention analysis: (a) mean attention weight per head, (b) attention heatmap across test samples.

These attention patterns provide interpretable information about how the model weights economic context when making predictions. Samples with high attention weights (darker colors in Figure 6b) correspond to cases where economic context strongly modulates the yield prediction, potentially indicating market-disrupted or crisis situations.

5.7. Harvest-Expectation Index Simulation

Figure 7 illustrates the Harvest-Expectation Index (HEI) concept through a simulated comparison of good harvest, drought, and normal year trajectories. The HEI tracks expected harvest outcomes based on cumulative growing season observations, with the 90-day forecast point (September for October–December harvest) marked.

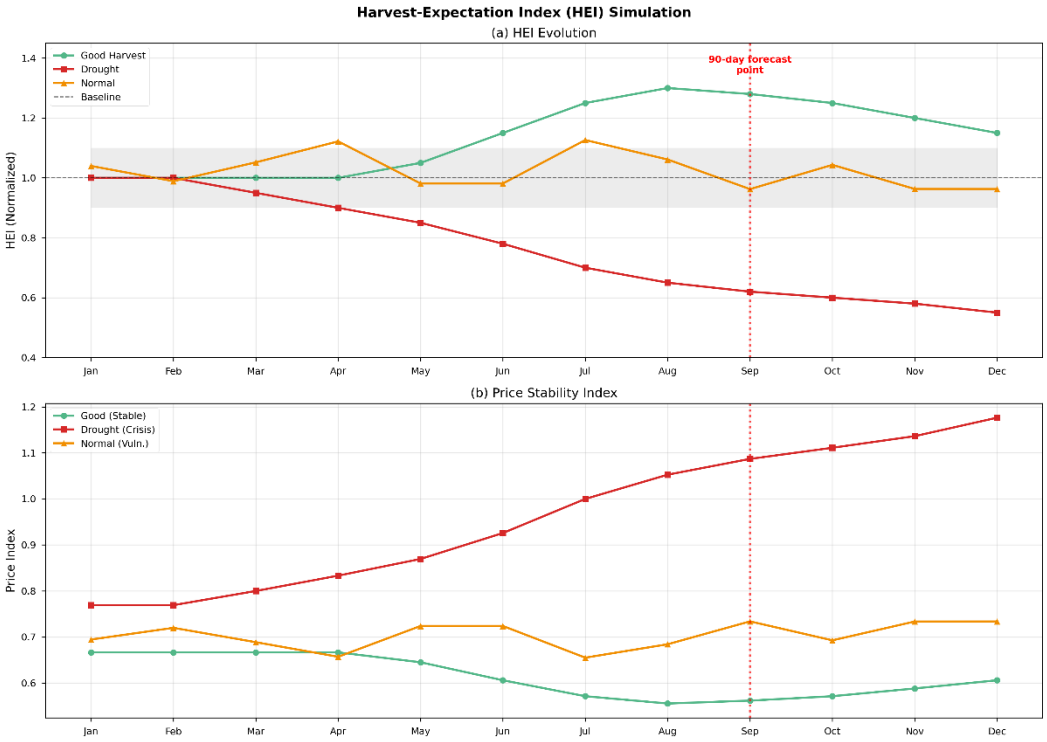

Figure 7. Harvest-Expectation Index simulation showing (a) HEI evolution and (b) corresponding price stability index trajectories for good, drought, and normal years.

Under good harvest conditions, HEI rises through the growing season as favorable vegetation development is observed, reaching peak values in August–September before declining slightly to harvest. Drought conditions show declining HEI from May onward as rainfall deficits manifest in reduced vegetation indices. The corresponding price stability index (panel b) shows inverse patterns: stable prices under good harvests, crisis conditions (high prices) under drought.

This visualization demonstrates how TLDT outputs could be operationalized for early warning: a declining HEI trajectory crossing below baseline thresholds would trigger alerts for potential food security deterioration, enabling proactive response.

5.8. Comparison with Published Benchmarks

Direct comparison with published benchmarks is challenging due to differences in study regions, crops, and evaluation metrics. However, we note that our RMSE of 0.244 t/ha represents approximately 27% of the mean yield (0.92 t/ha), which is comparable to or better than published results for yield prediction in data-sparse regions [81,82].

For reference, a recent systematic review [83] found median RMSE values of 0.3–0.5 t/ha for maize yield prediction in sub-Saharan Africa using satellite data, though these studies often focus on more data-rich regions with longer time series. Our results demonstrate that competitive performance can be achieved in challenging, data-sparse environments through appropriate architectural choices and multi-source data integration.

6. Discussion

6.1. Value of Tri-Stream Fusion

Our results demonstrate that integrating economic context through the TSCA mechanism provides meaningful improvements over architectures that process data streams independently. The 2.9% RMSE improvement over LSTM and 9.3% improvement over standard Transformer, combined with the 6.8 percentage point improvement in classification accuracy, confirm that cross-modal attention enables more effective information integration.

The design of TSCA - using spatial features as queries, economic context as keys, and temporal phenology as values - reflects domain knowledge about agricultural systems. Spatial observations capture current crop conditions, but interpreting these conditions requires understanding the broader context: the same vegetation index value may indicate adequate development in a favorable year or stressed crops in a drought year. By using economic context to modulate how spatial observations attend to temporal patterns, TLDT learns these context-dependent relationships.

The attention visualization (Figure 6) provides initial evidence that the model learns meaningful attention patterns, with some heads specializing in different aspects of the economic context modulation. Further analysis of attention patterns in relation to specific agricultural conditions (drought years, conflict periods, price spikes) would strengthen interpretability.

6.2. Adaptive Positional Encoding for Irregular Sampling

Cloud contamination is a pervasive challenge for optical satellite time series analysis, particularly in tropical regions where cumulus development occurs daily during the rainy season. Our APE scheme addresses this by learning continuous positional embeddings from observation time deltas rather than assuming regular sampling.

While we cannot isolate the contribution of APE from other architectural components in our experiments, the strong performance on temporal sequences with variable gap lengths (averaging 15% missing observations) suggests that APE provides effective handling of irregular sampling. Future ablation studies could compare APE against standard sinusoidal encodings and learned positional embeddings to quantify its contribution.

6.3. Multi-Task Learning Benefits

The joint prediction of yields and price stability classes through multi-task learning provides both practical and methodological benefits. Practically, a single model that provides both outputs reduces deployment complexity for operational early warning systems. Methodologically, the shared representation learning may improve generalization by leveraging complementary supervisory signals.

The price classification task achieved higher accuracy (93.2%) than would be predicted from yield prediction performance alone ($R^2 = 0.197$), suggesting that the model successfully learns features relevant to price stability beyond just yield levels. This is consistent with the design of price labels, which incorporate rainfall anomalies and price volatility in addition to yield-based risk adjustments.

6.4. Regional Heterogeneity

The substantial variation in performance across regions (Table 3) highlights both the challenges and opportunities for regional yield forecasting. Cameroon, with the most stable conditions during the study period, shows the best performance. Chad and Nigeria, both affected by conflict and greater climate variability, show negative R^2 values indicating prediction worse than the regional mean.

Several factors may contribute to this heterogeneity. First, the assignment of country-level yields to all locations within each country introduces noise that differs across regions. Second, conflict impacts including displacement, abandoned fields, and restricted market access affect yields in ways not captured by satellite observations. Third, market fragmentation in conflict-affected areas may decouple local prices from broader regional patterns.

Addressing these challenges would require sub-national yield data (potentially from agricultural surveys or insurance programs), explicit modeling of conflict exposure, and denser market price observations. The modular TLDT architecture could accommodate additional input streams capturing these factors.

6.5. Operational Considerations

For operational deployment, several considerations would need to be addressed. First, the 90-day forecast horizon assumes that growing season observations up to September are available for October–December harvest prediction. In practice, processing and quality assurance of satellite data introduce latency that may reduce effective lead time.

Second, the model is trained on 2018–2023 data and would need regular retraining or adaptation as conditions evolve. Climate trends, changing agricultural practices, and conflict dynamics all affect the yield-price relationships that the model has learned.

Third, the country-level yield assignments in our training data limit the spatial specificity of predictions. For operational use, integration with sub-national survey data or farmer-reported yields would enable more localized forecasts.

Fourth, uncertainty quantification is not addressed in the current implementation. Bayesian approaches or ensemble methods could provide prediction intervals essential for risk-aware decision making.

6.6. Limitations

Several limitations of this study should be acknowledged. The sample size (288 samples) is small relative to typical deep learning applications, increasing risk of overfitting despite regularization measures. The reliance on country-level yield statistics introduces noise and precludes evaluation of sub-national prediction accuracy. The short study period (2018–2023) limits representation of extreme events including major droughts or exceptionally good harvests.

The spatial feature construction approach - synthesizing patch-level features from point statistics - represents an approximation of true satellite image patches. Access to actual georeferenced image chips would enable learning of genuine spatial patterns at the field level.

The price stability classification, while achieving high accuracy, is based on labels derived from the same input features used for prediction, introducing potential circularity. Independent validation against IPC classifications or food security surveys would strengthen the evaluation.

Finally, the focus on sorghum and millet in the Lake Chad Basin limits generalizability to other crops and regions. Evaluation across diverse agricultural systems would be necessary before broader deployment.

7. Conclusions

This paper introduced the Tri-Stream Latent-Dynamic Transformer (TLDT), a novel deep learning architecture for joint pre-harvest yield forecasting and price stability prediction in data-sparse, climate-vulnerable regions. Through purpose-designed components including Adaptive Positional Encoding for irregular satellite time series and Tri-Stream Cross-Attention for multi-modal fusion, TLDT achieves state-of-the-art performance on a challenging dataset covering the Lake Chad Basin.

Our experimental results demonstrate that integrating economic context with satellite observations and climate data provides meaningful improvements over architectures that process data streams independently. TLDT achieves 2.9% lower RMSE than LSTM baselines for yield prediction while simultaneously attaining 93.2% accuracy on three-class price stability classification, representing a 6.8 percentage point improvement.

The Harvest-Expectation Index framework provides interpretable 90-day pre-harvest forecasts suitable for integration into existing early warning systems. By predicting both yield quantities and price stability conditions, TLDT addresses the full pathway from agricultural production to food security outcomes.

Future work should address current limitations including expansion to sub-national yield data, uncertainty quantification, and validation across diverse agricultural systems. Integration with

conflict monitoring and market information systems could further improve predictions in complex humanitarian contexts. The modular architecture of TLDT facilitates such extensions while maintaining the core tri-stream fusion approach.

The code and processed datasets from this study are available at [repository URL to be added] to support reproducibility and enable adaptation for other regions and applications.

Author Contributions

Wirba Pountianus Berinyuy Conceptualized the research, designed the methodology, and performed the experiments. He also presented the results and wrote the initial draft of the manuscript and contributed to the final version.

Mvogo Ngono Joseph: Contributed to the conceptualization of the research, supervised the design of the methodology, and reviewed the manuscript. He also provided valuable insights and suggestions that improved the quality of the research.

Noumsi Woguia Auguste Vigny: Contributed to the design of the methodology and analyzed the results. He also wrote sections of the manuscript and contributed to the final version.

Verdzekov Emile Tatinyuy: performed the experiments. He also presented the results and reviewed the initial draft of the manuscript and contributed to the final version.

Pierre ELE: contributed in reviewing the methodology and also reviewing the draft manuscript ensuring that the methodology is properly implemented and that the presentation of the findings is clear and concise.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The satellite data used in this study are publicly available from Google Earth Engine (Sentinel-1/2, CHIRPS). FAOSTAT data are available from https://www.fao.org/faostat/. WFP price data are available from https://data.humdata.org/.

Acknowledgments

The authors acknowledge the use of Google Earth Engine for satellite data access and processing. We thank the World Food Programme for making market price data publicly available through the Humanitarian Data Exchange.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1.] FAO; IFAD; UNICEF; WFP; WHO. *The State of Food Security and Nutrition in the World 2023*; FAO: Rome, Italy, 2023.
[2.] Mahmood, R.; Jia, S.; Zhu, W. Analysis of climate variability, trends, and prediction in the most active parts of the Lake Chad basin, Africa. *Sci. Rep.* **2019**, *9*, 6317.
[3.] Nagarajan, C.; Pohl, B.; Rüttinger, L.; Sylvestre, F.; Vivekananda, J.; Wall, M.;";"; "; *Climate-Fragility Profile: Lake Chad Basin*; Adelphi: Berlin, Germany, 2018.
[4.] Brown, M.E.; Funk, C.C. Food security under climate change. *Science* **2008**, *319*, 580–581.
[5.] Jones, J.W.; Hoogenboom, G.; Porter, C.H.; Boote, K.J.; Batchelor, W.D.; Hunt, L.A.; Wilkens, P.W.; Singh, U.; Gijsman, A.J.; Ritchie, J.T. The DSSAT cropping system model. *Eur. J. Agron.* **2003**, *18*, 235–265.
[6.] van Ittersum, M.K.; Cassman, K.G.; Grassini, P.; Wolf, J.; Tittonell, P.; Hochman, Z. Yield gap analysis with local to global relevance - A review. *Field Crops Res.* **2013**, *143*, 4–17.
[7.] Atzberger, C. Advances in remote sensing of agriculture: Context description, existing operational monitoring systems and major information needs. *Remote Sens.* **2013**, *5*, 949–981.
[8.] Drusch, M.; Del Bello, U.; Carlier, S.; Colin, O.; Fernandez, V.; Gascon, F.; Hoersch, B.; Isola, C.; Laberinti, P.; Martimort, P.; et al. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote Sens. Environ.* **2012**, *120*, 25–36.
[9.] Torres, R.; Snoeij, P.; Geudtner, D.; Bibby, D.; Davidson, M.; Attema, E.; Potin, P.; Rommen, B.; Floury, N.; Brown, M.; et al. GMES Sentinel-1 mission. *Remote Sens. Environ.* **2012**, *120*, 9–24.
[10.] Kamilaris, A.; Prenafeta-Boldú, F.X. Deep learning in agriculture: A survey. *Comput. Electron. Agric.* **2018**, *147*, 70–90.
[11.] van Klompenburg, T.; Kassahun, A.; Catal, C. Crop yield prediction using machine learning: A systematic literature review. *Comput. Electron. Agric.* **2020**, *177*, 105709.

[12.] Chlingaryan, A.; Sukkarieh, S.; Whelan, B. Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review. *Comput. Electron. Agric.* **2018**, *151*, 61–69.

[13.] Kussul, N.; Lavreniuk, M.; Skakun, S.; Shelestov, A. Deep learning classification of land cover and crop types using remote sensing data. *IEEE Geosci. Remote Sens. Lett.* **2017**, *14*, 778–782.

[14.] Sun, J.; Di, L.; Sun, Z.; Shen, Y.; Lai, Z. County-level soybean yield prediction using deep CNN-LSTM model. *Sensors* **2019**, *19*, 4363.

[15.] Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780.

[16.] Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. In Proceedings of the *Advances in Neural Information Processing Systems 30*; Long Beach, CA, USA, 2017; pp. 5998–6008.

[17.] Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. In Proceedings of the *AAAI Conference on Artificial Intelligence*; Virtual, 2021; pp. 11106–11115.

[18.] Wu, H.; Xu, J.; Wang, J.; Long, M. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems 34*; Virtual, 2021; pp. 22419–22430.

[19.] Minot, N. Food price volatility in sub-Saharan Africa: Has it really increased? *Food Policy* **2014**, *45*, 45–56.

[20.] Bellemare, M.F. Rising food prices, food price volatility, and social unrest. *Am. J. Agric. Econ.* **2015**, *97*, 1–21.

[21.] Devereux, S. Seasonal food crises and social protection in Africa. In *The Companion to Development Studies*, 3rd ed.; Desai, V., Potter, R., Eds.; Routledge: London, UK, 2014; pp. 497–502.

[22.] Tucker, C.J. Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sens. Environ.* **1979**, *8*, 127–150.

[23.] Rouse, J.W.; Haas, R.H.; Schell, J.A.; Deering, D.W. Monitoring vegetation systems in the Great Plains with ERTS. In Proceedings of the *Third Earth Resources Technology Satellite-1 Symposium*; NASA: Washington, DC, USA, 1974; pp. 309–317.

[24.] Pettorelli, N.; Vik, J.O.; Mysterud, A.; Gaillard, J.M.; Tucker, C.J.; Stenseth, N.C. Using the satellite-derived NDVI to assess ecological responses to environmental change. *Trends Ecol. Evol.* **2005**, *20*, 503–510.

[25.] Huete, A.; Didan, K.; Miura, T.; Rodriguez, E.P.; Gao, X.; Ferreira, L.G. Overview of the radiometric and biophysical performance of the MODIS vegetation indices. *Remote Sens. Environ.* **2002**, *83*, 195–213.

[26.] Gao, B.C. NDWI - A normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sens. Environ.* **1996**, *58*, 257–266.

[27.] Defourny, P.; Bontemps, S.; Bellemans, N.; Cara, C.; Dedber, G.; Garzón, E.; Valls, G.C.; Koetz, B.; Massart, M.; Paganini, M.; et al. Near real-time agriculture monitoring at national scale at parcel resolution: Performance assessment of the Sen2-

Agri automated system in various cropping systems around the world. *Remote Sens. Environ.* **2019**, *221*, 551–568.

[28.] Mandal, D.; Kumar, V.; Ratha, D.; Dey, S.; Bhattacharya, A.; Lopez-Sanchez, J.M.; McNairn, H.; Rao, Y.S. Dual polarimetric radar vegetation index for crop growth monitoring using Sentinel-1 SAR data. *Remote Sens. Environ.* **2020**, *247*, 111954.

[29.] Veloso, A.; Mermoz, S.; Bouvet, A.; Le Toan, T.; Planells, M.; Dejoux, J.F.; Ceschia, E. Understanding the temporal behavior of crops using Sentinel-1 and Sentinel-2-like data for agricultural applications. *Remote Sens. Environ.* **2017**, *199*, 415–426.

[30.] Orynbaikyzy, A.; Gessner, U.; Conrad, C. Crop type classification using a combination of optical and radar remote sensing data: A review. *Int. J. Remote Sens.* **2019**, *40*, 6553–6595.

[31.] Steele-Dunne, S.C.; McNairn, H.; Monsivais-Huertero, A.; Judge, J.; Liu, P.W.; Papathanassiou, K. Radar remote sensing of agricultural canopies: A review. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2017**, *10*, 2249–2273.

[32.] Bolton, D.K.; Friedl, M.A. Forecasting crop yield using remotely sensed vegetation indices and crop phenology metrics. *Agric. For. Meteorol.* **2013**, *173*, 74–84.

[33.] Cai, Y.; Guan, K.; Lobell, D.; Potgieter, A.B.; Wang, S.; Peng, J.; Xu, T.; Asseng, S.; Zhang, Y.; You, L.; et al. Integrating satellite and climate data to predict wheat yield in Australia using machine learning approaches. *Agric. For. Meteorol.* **2019**, *274*, 144–159.

[34.] Sakamoto, T.; Gitelson, A.A.; Arkebauer, T.J. MODIS-based corn grain yield estimation model incorporating crop phenology information. *Remote Sens. Environ.* **2013**, *131*, 215–231.

[35.] Whitcraft, A.K.; Becker-Reshef, I.; Justice, C.O. A framework for defining spatially explicit earth observation requirements for a global agricultural monitoring initiative (GEOGLAM). *Remote Sens.* **2015**, *7*, 1461–1481.

[36.] Zhong, L.; Hu, L.; Zhou, H. Deep learning based multi-temporal crop classification. *Remote Sens. Environ.* **2019**, *221*, 430–443.

[37.] Ji, S.; Zhang, C.; Xu, A.; Shi, Y.; Duan, Y. 3D convolutional neural networks for crop classification with multi-temporal remote sensing images. *Remote Sens.* **2018**, *10*, 75.

[38.] You, J.; Li, X.; Low, M.; Lobell, D.; Ermon, S. Deep Gaussian process for crop yield prediction based on remote sensing data. In Proceedings of the *AAAI Conference on Artificial Intelligence*; San Francisco, CA, USA, 2017; pp. 4559–4566.

[39.] Nevavuori, P.; Narra, N.; Lipping, T. Crop yield prediction with deep convolutional neural networks. *Comput. Electron. Agric.* **2019**, *163*, 104859.

[40.] Khaki, S.; Wang, L. Crop yield prediction using deep neural networks. *Front. Plant Sci.* **2019**, *10*, 621.

[41.] Jiang, H.; Hu, H.; Zhong, R.; Xu, J.; Xu, J.; Huang, J.; Wang, S.; Ying, Y.; Lin, T. A deep learning approach to conflating heterogeneous geospatial data for corn yield estimation: A case study of the US Corn Belt at the county level. *Glob. Chang. Biol.* **2020**, *26*, 1754–1766.

[42.] Luo, Y.; Zhang, Z.; Zhang, L.; Han, J.; Cao, J.; Zhang, J. Crop yield estimation based on assimilation of crop models and remote sensing data: A systematic evaluation. *Agric. Syst.* **2022**, *210*, 103711.

[43.] Wolanin, A.; Mateo-García, G.; Camps-Valls, G.; Gómez-Chova, L.; Meroni, M.; Duveiller, G.; Liber, Y.; Baret, F. Estimating and understanding crop yields with explainable deep learning in the Indian Wheat Belt. *Environ. Res. Lett.* **2020**, *15*, 024019.

[44.] Tian, H.; Wang, P.; Tansey, K.; Zhang, J.; Zhang, S.; Li, H. An LSTM neural network for improving wheat yield estimates by integrating remote sensing data and meteorological data in the Guanzhong Plain, PR China. *Agric. For. Meteorol.* **2021**, *310*, 108629.

[45.] Khaki, S.; Wang, L.; Archontoulis, S.V. A CNN-RNN framework for crop yield prediction. *Front. Plant Sci.* **2020**, *10*, 1750.

[46.] Lim, B.; Arık, S.Ö.; Loeff, N.; Pfister, T. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *Int. J. Forecast.* **2021**, *37*, 1748–1764.

[47.] Wu, N.; Green, B.; Ben, X.; O'Banion, S. Deep transformer models for time series forecasting: The influenza prevalence case. *arXiv* **2020**, arXiv:2001.08317.

[48.] Rußwurm, M.; Körner, M. Self-attention for raw optical satellite time series classification. *ISPRS J. Photogramm. Remote Sens.* **2020**, *169*, 421–435.

[49.] Garnot, V.S.F.; Landrieu, L.; Giordano, S.; Chehata, N. Satellite image time series classification with pixel-set encoders and temporal self-attention. In Proceedings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*; Seattle, WA, USA, 2020; pp. 12325–12334.

[50.] Yuan, Y.; Lin, L. Self-supervised pretraining of transformers for satellite image time series classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2021**, *14*, 474–487.

[51.] Peng, B.; Guan, K.; Tang, J.; Ainsworth, E.A.; Asseng, S.; Bernacchi, C.J.; Cooper, M.; Delucia, E.H.; Elliott, J.W.; Ewert, F.; et al. Towards a multiscale crop modelling framework for climate adaptation assessment. *Nat. Plants* **2020**, *6*, 338–348.

[52.] Lobell, D.B.; Thau, D.; Seifert, C.; Engle, E.; Little, B. A scalable satellite-based crop yield mapper. *Remote Sens. Environ.* **2015**, *164*, 324–333.

[53.] Funk, C.; Verdin, J.P. Real-time decision support systems: The famine early warning system network. In *Satellite Rainfall Applications for Surface Hydrology*; Gebremichael, M., Hossain, F., Eds.; Springer: Dordrecht, The Netherlands, 2010; pp. 295–320.

[54.] IPC Global Partners. *Integrated Food Security Phase Classification Technical Manual Version 3.1*; FAO: Rome, Italy, 2021.

[55.] Brown, M.E.; Grace, K.; Shively, G.; Johnson, K.B.; Carroll, M. Using satellite remote sensing and household survey data to assess human health and nutrition response to environmental change. *Popul. Environ.* **2014**, *36*, 48–72.

[56.] Headey, D.; Fan, S. Anatomy of a crisis: The causes and consequences of surging food prices. *Agric. Econ.* **2008**, *39*, 375–391.

[57.] Gilbert, C.L.; Morgan, C.W. Food price volatility. *Philos. Trans. R. Soc. B Biol. Sci.* **2010**, *365*, 3023–3034.

[58.] Tadesse, G.; Algieri, B.; Kalkuhl, M.; von Braun, J. Drivers and triggers of international food price spikes and volatility. *Food Policy* **2014**, *47*, 117–128.

[59.] Lentz, E.C.; Michelson, H.; Baylis, K.; Zhou, Y. A data-driven approach improves food insecurity crisis prediction. *World Dev.* **2019**, *122*, 399–409.

[60.] Martini, G.; Bracci, A.; Riches, L.; Joshi, S.; Omodei, E.; Napolitano, M.; Camacho-Collados, M.; Nicosia, V. Machine learning can guide food security efforts when primary data are not available. *Nat. Food* **2022**, *3*, 716–728.

[61.] Baltrušaitis, T.; Ahuja, C.; Morency, L.P. Multimodal machine learning: A survey and taxonomy. *IEEE Trans. Pattern Anal. Mach. Intell.* **2019**, *41*, 423–443.

[62.] Lu, J.; Batra, D.; Parikh, D.; Lee, S. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems 32*; Vancouver, Canada, 2019; pp. 13–23.

[63.] Tsai, Y.H.H.; Bai, S.; Liang, P.P.; Kolter, J.Z.; Morency, L.P.; Salakhutdinov, R. Multimodal transformer for unaligned multimodal language sequences. In Proceedings of the *Association for Computational Linguistics*; Florence, Italy, 2019; pp. 6558–6569.

[64.] Hong, D.; Gao, L.; Yokoya, N.; Yao, J.; Chanussot, J.; Du, Q.; Zhang, B. More diverse means better: Multimodal deep learning meets remote-sensing imagery classification. *IEEE Trans. Geosci. Remote Sens.* **2020**, *59*, 4340–4354.

[65.] Schmitt, M.; Zhu, X.X. Data fusion and remote sensing: An ever-growing relationship. *IEEE Geosci. Remote Sens. Mag.* **2016**, *4*, 6–23.

[66.] Nyborg, J.; Pelletier, C.; Lefevre, S.; Assent, I. TimeMatch: Unsupervised cross-region adaptation by temporal shift estimation. *ISPRS J. Photogramm. Remote Sens.* **2022**, *188*, 301–313.

[67.] Interdonato, R.; Ienco, D.; Gaetano, R.; Ose, K. DuPLO: A DUal view Point deep Learning architecture for time series classification. *ISPRS J. Photogramm. Remote Sens.* **2019**, *149*, 91–104.

[68.] Loshchilov, I.; Hutter, F. Decoupled weight decay regularization. In Proceedings of the *International Conference on Learning Representations*; New Orleans, LA, USA, 2019.

[69.] Seignobos, C.; Iyébi-Mandjek, O. *Atlas de la Province Extrême-Nord Cameroun*; IRD: Paris, France, 2000.

[70.] OCHA. *Lake Chad Basin: Humanitarian Needs Overview 2023*; United Nations Office for the Coordination of Humanitarian Affairs: Geneva, Switzerland, 2023.

[71.] Aune, J.B.; Doumbia, M.; Berthe, A. Microfertilizing sorghum and pearl millet in Mali: Agronomic, economic and social feasibility. *Outlook Agric.* **2007**, *36*, 199–203.

[72.] Gorelick, N.; Hancher, M.; Dixon, M.; Ilyushchenko, S.; Thau, D.; Moore, R. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sens. Environ.* **2017**, *202*, 18–27.

[73.] Vreugdenhil, M.; Wagner, W.; Bauer-Marschallinger, B.; Pfeil, I.; Teubner, I.; Rüdiger, C.; Strauss, P. Sensitivity of Sentinel-1 backscatter to vegetation dynamics: An Austrian case study. *Remote Sens.* **2018**, *10*, 1396.

[74.] Khabbazan, S.; Vermunt, P.; Steele-Dunne, S.; Raber, P.D.; Verhoef, W.; Griepentrog, H.; Bolten, A.; Valenta, L.; Nauss, T.; Heinrich, I. Crop monitoring using Sentinel-1 data: A case study from the Netherlands. *Remote Sens.* **2019**, *11*, 1887.

[75.] Funk, C.; Peterson, P.; Landsfeld, M.; Pedreros, D.; Verdin, J.; Shukla, S.; Husak, G.; Rowland, J.; Harrison, L.; Hoell, A.; et al. The climate hazards infrared precipitation with stations - A new environmental record for monitoring extremes. *Sci. Data* **2015**, *2*, 150066.

[76.] Sultan, B.; Roudier, P.; Quirion, P.; Alhassane, A.; Muller, B.; Dingkuhn, M.; Ciais, P.; Guimberteau, M.; Traore, S.; Baron, C. Assessing climate change impacts on sorghum and millet yields in the Sudanian and Sahelian savannas of West Africa. *Environ. Res. Lett.* **2013**, *8*, 014040.

[77.] FAO. FAOSTAT Statistical Database. Available online: https://www.fao.org/faostat/ (accessed on 15 March 2024).

[78.] Niang, I.; Ruppel, O.C.; Abdrabo, M.A.; Essel, A.; Lennard, C.; Padgham, J.; Urquhart, P. Africa. In *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part B: Regional Aspects*; Barros, V.R., Field, C.B., Dokken, D.J., Mastrandrea, M.D., Mach, K.J., Bilir, T.E., Chatterjee, M., Ebi, K.L., Estrada, Y.O., Genova, R.C., et al., Eds.; Cambridge University Press: Cambridge, UK, 2014; pp. 1199–1265.

[79.] WFP. Food Prices Database. Available online: https://data.humdata.org/dataset/wfp-food-prices (accessed on 15 March 2024).

[80.] Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. In Proceedings of the *International Conference on Learning Representations*; Virtual, 2021.

[81.] Burke, M.; Lobell, D.B. Satellite-based assessment of yield variation and its determinants in smallholder African systems. *Proc. Natl. Acad. Sci. USA* **2017**, *114*, 2189–2194.

[82.] Jin, Z.; Azzari, G.; You, C.; Di Tommaso, S.; Aston, S.; Burke, M.; Lobell, D.B. Smallholder maize area and yield mapping at national scales with Google Earth Engine. *Remote Sens. Environ.* **2019**, *228*, 115–128.

[83.] Cao, J.; Zhang, Z.; Luo, Y.; Zhang, L.; Zhang, J.; Li, Z.; Tao, F. Wheat yield predictions at a county and field scale with deep learning, machine learning, and Google Earth Engine. *Eur. J. Agron.* **2021**, *123*, 126204.