

Interpretable Transformer-Based Yield Forecasting Under Climate Extremes: A Multi-Layered Explainability Framework and Drought-Year Validation on the Canadian Prairies (2021)

Seema Sharma¹, Dr. Praveen Gupta²
ARNI University, Tanda, Himachal Pradesh^{1,2}
seemasharma1974@gmail.com

Abstract: Accurate and interpretable crop yield forecasting under climate extremes is essential for food security, agricultural risk management, and policy decision-making. This study presents a multi-layered explainability framework for transformer-based yield prediction, validated during the severe 2021 drought across the Canadian Prairies — a natural out-of-distribution stress test. We trained a Temporal Fusion Transformer (TFT) on a multi-source geospatial data stack encompassing Sentinel-2 surface reflectance, MODIS vegetation indices, RADARSAT constellation synthetic aperture radar backscatter, SMAP soil moisture, ERA5-Land climate reanalysis, and AAFC crop inventory maps. The model achieved strong overall performance ($R^2 = 0.84$, $RMSE = 0.42 \text{ t ha}^{-1}$) across 2017–2020 but degraded predictably during the 2021 drought ($R^2 = 0.71$, $RMSE = 0.68 \text{ t ha}^{-1}$). To understand this degradation, we deployed a comprehensive explainable artificial intelligence (XAI) pipeline: (1) TFT native variable selection networks revealed dynamic reweighting from vegetation indices toward soil moisture and thermal stress indices; (2) attention rollout analysis identified a temporal shift in model focus from mid-season (July) to early-season (June) predictors approximately 4–6 weeks before harvest; (3) SHAP analysis quantified global and local feature attributions, showing that evaporative stress index (ESI) and surface soil moisture anomalies became dominant drivers in 2021; (4) LIME confirmed instance-level fidelity of explanations; and (5) Grad-CAM applied to spatial input rasters highlighted drought-sensitive regions in southwestern Saskatchewan and southeastern Alberta. Critically, early-stress signals emerged in attention and SHAP time series by late June 2021, supporting Hypothesis H4 that interpretable transformer outputs can provide actionable early warnings 4–6 weeks' pre-harvest. These findings demonstrate that multi-layered explainability not only builds trust in deep learning forecasts but also reveals mechanistic insights into crop response to extreme climate events, with direct implications for crop insurance, drought relief allocation, and on-farm adaptation decisions.

Keywords: Temporal Fusion Transformer; explainable artificial intelligence; crop yield forecasting; 2021 Canadian drought; Grad-CAM; SHAP; LIME; attention rollout; early warning systems; agricultural risk management

I. INTRODUCTION

Global agriculture faces unprecedented challenges from climate change, with extreme weather events — particularly droughts — increasing in frequency and severity across major cereal-producing regions (IPCC, 2022; Zscheischler et al., 2018). The Canadian Prairies, which account for approximately 80% of Canada's cultivated farmland and contribute over CAD \$30 billion annually to the national economy, experienced a catastrophic drought in 2021 that reduced wheat and canola yields by 30–40% in severely affected regions (Hunt et al., 2022; Lawley et al., 2022). This event, compounded by a record-breaking heat dome in late June 2021 (Zhou et al., 2022; von Burg et al., 2021), exposed critical vulnerabilities in existing crop forecasting systems and underscored the urgent need for predictive models that are both accurate under extreme conditions and interpretable enough to inform timely decision-making.

Accurate pre-harvest yield forecasting is fundamental to food security, agricultural risk management, and policy formulation (Feoli & Zuccarello, 2021). Traditional approaches — including process-based crop models, statistical regression, and machine learning classifiers — have achieved reasonable accuracy under normal growing conditions (Liu et al., 2020). However, these methods face significant limitations when confronted with out-of-distribution climate extremes (Lobell et al., 2014; Milly et al., 2008). Deep learning models, particularly those leveraging temporal sequences of satellite and climate data, have emerged as powerful alternatives, consistently outperforming traditional methods in large-scale yield prediction tasks (You et al., 2017; Khaki & Wang, 2020).

Among deep learning architectures, the Transformer (Vaswani et al., 2017) and its variants have demonstrated exceptional capability in capturing long-range temporal dependencies in sequential data. The Temporal Fusion Transformer (TFT) (Lim et al., 2021) represents a significant advancement for agricultural forecasting, as it explicitly models static covariates (e.g., soil properties), known future inputs (e.g., planting dates), and observed temporal dynamics (e.g., vegetation indices) while providing built-in interpretability through variable selection networks and multi-head attention mechanisms. Unlike “black box” architectures, the TFT offers practitioners insight into which variables drive predictions and how attention weights evolve across the growing season.

However, despite the promise of transformer-based models for yield prediction, several critical gaps remain. First, the interpretability provided by TFT’s native mechanisms — while valuable — is insufficient for fully understanding model behavior under extreme, out-of-distribution conditions such as drought. Second, existing studies predominantly evaluate models on historical datasets without explicit stress-testing under anomalous climate conditions, leaving unanswered questions about reliability when forecasts matter most (Krishnan et al., 2022). Third, there is limited integration of complementary explainability methods (SHAP, LIME, Grad-CAM) to provide multi-layered, cross-validated explanations that can reveal not just what the model predicts, but why it behaves differently under stress (Lundberg & Lee, 2017; Ribeiro et al., 2016; Selvaraju et al., 2017).

The 2021 Canadian Prairies drought provides an exceptional natural experiment for addressing these gaps. This event was characterized by:

- 1. Early onset:** Precipitation deficits emerged by late April 2021, with the Canadian Drought Monitor reporting moderate drought (D1) across southern Saskatchewan and Alberta by mid-May (Agriculture and Agri-Food Canada, 2021).
- 2. Extreme intensity:** The June–July heat dome produced temperatures exceeding 40 °C in regions where seasonal averages rarely surpass 25 °C (Zhou et al., 2022).
- 3. Spatial heterogeneity:** While southwestern Saskatchewan and southeastern Alberta experienced exceptional drought (D4), parts of Manitoba received near-normal precipitation (NOAA, 2021).
- 4. Severe agricultural impact:** Statistics Canada reported wheat yield reductions of 34% in Saskatchewan and 28% in Alberta compared to 2020 (Statistics Canada, 2021).

This study presents a comprehensive multi-layered explainability framework for transformer-based yield forecasting and validates it using the 2021 drought as a natural out-of-distribution stress test. Our specific objectives are:

- 1.** To develop and train a TFT yield prediction model on a multi-source geospatial data stack spanning 2017–2021 across the Canadian Prairies.
- 2.** To implement and integrate five complementary explainability methods: (a) TFT native variable selection networks, (b) attention rollout analysis, (c) SHAP for global and local feature attribution, (d) LIME for instance-level explanation fidelity, and (e) Grad-CAM adapted for geospatial input features.
- 3.** To characterize model performance degradation during the 2021 drought relative to non-drought years and use the XAI pipeline to reveal mechanistic drivers of this degradation.
- 4.** To test Hypothesis H4: Early-stress signals emerge in attention patterns and SHAP attributions 4–6 weeks before harvest, demonstrating the potential for interpretable early warning systems.

By providing a “why” behind model predictions under extreme conditions, this framework aims to bridge the gap between predictive accuracy and actionable interpretability, with direct implications for crop insurance pricing, drought relief allocation, and on-farm adaptation decisions.

II. MATERIALS AND METHODS

2.1 Study Area

The study domain encompasses the Canadian Prairie provinces of Alberta, Saskatchewan, and Manitoba, covering approximately 1.96 million km² between 49°N–60°N latitude and 95°W–120°W longitude. This region is characterized by a continental climate with cold winters and warm summers, mean annual precipitation ranging from 300 mm in the semi-arid southwest to 600 mm in the east, and highly variable interannual moisture regimes (Hube et al., 2021). The dominant crops include spring wheat (*Triticum aestivum* L.), canola (*Brassica napus* L.), barley (*Hordeum vulgare* L.), and field peas (*Pisum sativum* L.), typically seeded in late April to mid-May and harvested in August through September.

The 2021 drought exhibited pronounced spatial heterogeneity (Figure 1). The Canadian Drought Monitor classified 72% of Saskatchewan and 65% of Alberta in moderate to exceptional drought (D1–D4) by July 2021, with the most severe conditions (D4) concentrated in southwestern Saskatchewan and southeastern Alberta (Agriculture and Agri-Food Canada, 2021). In contrast, eastern Manitoba remained largely drought-free. This spatial gradient provided a natural experimental design for evaluating model behavior across a drought severity continuum.




Figure 1. Study area map showing the Canadian Prairie provinces with (a) Canadian Drought Monitor drought classifications for July 2021, (b) spatially stratified evaluation zones (high drought, moderate drought, low/no drought), and (c) locations of meteorological stations and yield reporting units. Base map data from Statistics Canada Boundary Files (2021).

2.2 Data Sources and Preprocessing

We assembled a multi-source geospatial data stack spanning the 2017–2021 growing seasons (May–September). Table 1 provides a summary of all datasets; complete access details are provided in Appendix A.

Table 1. Summary of multi-source input data for the Temporal Fusion Transformer yield prediction model.

Data Source	Spatial Res.	Temporal Res.	Variables Used	Period
Sentinel-2 MSI	10–60 m	5-day	B2–B8A, B11–B12, NDVI, EVI, NDWI	2017–2021
MODIS MOD13Q1	250 m	16-day	NDVI, EVI, pixel reliability	2017–2021
RADARSAT-2/RCM	10–50 m	12-day	VV, VH backscatter, VH/VV ratio	2019–2021
SMAP L4 SM	9 km	3-hour	Surface (0–5 cm) and root-zone soil moisture	2017–2021

Data Source	Spatial Res.	Temporal Res.	Variables Used	Period
ERA5-Land	9 km	hourly	Tmax, Tmin, Tmean, precip, PET, ETa, RH, SSR	2017–2021
AAFC ACI	30 m	annual	Crop type masks (spring wheat, canola, barley, peas)	2017–2021
Statistics Canada	Census div.	annual	County-level yield (t ha^{-1})	2017–2021
SoilGrids v2.0	250 m	static	Clay, sand, silt, SOC, pH, bulk density	—
CDED	20–40 m	static	Elevation, slope, aspect	—

2.2.1 Satellite Imagery Preprocessing

Sentinel-2 Level-2A surface reflectance products were accessed via the Copernicus Data Space Ecosystem (European Space Agency, 2021). We applied cloud masking using the Scene Classification Layer (SCL), retained observations with <20% cloud cover, and composited to 10-day intervals using the maximum NDVI criterion (Holben, 1986). Spectral indices computed included NDVI (Tucker, 1979), Enhanced Vegetation Index (EVI) (Huete et al., 2002), and Normalized Difference Water Index (NDWI).

MODIS MOD13Q1 Version 6.1 vegetation indices (NASA LP DAAC, 2021) were downloaded from the LP DAAC. We applied quality filtering (pixel reliability ≤ 1) and linearly interpolated remaining gaps using the method of Sakamoto et al. (2005). RADARSAT-2 and RCM C-band SAR backscatter data were terrain-corrected, radiometrically calibrated, and despeckled using a refined Lee filter (Moreira et al., 2013).

2.2.2 Climate and Soil Moisture Data

ERA5-Land reanalysis (Copernicus Climate Change Service, 2021) was extracted for 2-m temperature (daily max/min/mean), total precipitation, potential and actual evapotranspiration, surface net solar radiation, and 2-m relative humidity. From these, we derived the Evaporative Stress Index (ESI) as the ratio of actual to potential ET (Anderson et al., 2021), the Standardized Precipitation Evapotranspiration Index (SPEI) at 1-, 3-, and 6-month accumulation scales (SPEI, 2021), and Growing Degree Days (GDD, base 5 °C).

NASA SMAP Level-4 soil moisture (NSIDC DAAC, 2021) provided 3-hourly surface (0–5 cm) and root-zone (0–100 cm) volumetric soil moisture. We computed daily means and 7-day, 14-day, and 30-day anomalies relative to the 2017–2020 climatology for each pixel.

2.2.3 Crop Inventory and Yield Data

The AAFC Annual Crop Inventory (Agriculture and Agri-Food Canada, 2021) provided 30-m crop type masks. We retained pixels classified as spring wheat, canola, barley, or field peas with classification confidence $\geq 80\%$. County-level yield statistics from Statistics Canada Table 32-10-0359-01 (Statistics Canada, 2021) served as ground truth, converted to metric tonnes per hectare (t ha^{-1}). The 2021 release included official yield estimates for all reporting units.

2.2.4 Static Covariates

Static features included: SoilGrids v2.0 (Hengl et al., 2021) soil texture (clay, sand, silt), organic carbon content, pH, and bulk density at 0–15 cm; Canadian Digital Elevation Data (CDED) elevation, slope, and aspect; and latitude/longitude coordinates.

2.3 Temporal Fusion Transformer Architecture

We implemented the Temporal Fusion Transformer (Lim et al., 2021) using PyTorch Forecasting v1.0.0 (Paszke et al., 2019). The TFT architecture comprises four interconnected components (Figure 2):

Gating Mechanisms. Gated Residual Networks (GRNs) with sigmoid gating control information flow, allowing the model to skip unnecessary processing when inputs are less relevant (Lim et al., 2021).

Variable Selection Networks. These learn adaptive feature importance by applying sparse attention over input variables at each time step. The selection weights $\xi_t^{(j)}$ for variable j at time t are computed as:

$$\xi_t^{(j)} = \frac{\exp(\langle \Lambda_s^{(j)}, c_t \rangle)}{\sum_{k=1}^{n_{\text{var}}} \exp(\langle \Lambda_s^{(k)}, c_t \rangle)} \quad (1)$$

where Λ_s are context-dependent weights and c_t is the flattened representation of all inputs at time t . These weights provide native interpretability by revealing which variables the model prioritizes at each prediction horizon.

LSTM Encoder-Decoder. Two stacked LSTM layers process the selected variables separately for observed (encoder) and known future (decoder) inputs, capturing local temporal patterns.

Interpretable Multi-Head Attention. Following Vaswani et al. (2017), we use multi-head self-attention with H heads. The attention weights $\alpha_{t,\tau}^{(h)}$ for head h between query time t and key time τ are:

$$\alpha_{t,\tau}^{(h)} = \frac{\exp\left(\frac{Q_t^{(h)}(K_{\tau}^{(h)})^T}{\sqrt{d_k}}\right)}{\sum_{\tau'=1}^T \exp\left(\frac{Q_t^{(h)}(K_{\tau'}^{(h)})^T}{\sqrt{d_k}}\right)} \quad (2)$$

We average across heads to obtain a single interpretable attention map (Lim et al., 2021).

Quantile Output. The model outputs predictions at quantiles $q \in \{0.1, 0.5, 0.9\}$ using the quantile loss:

$$L_q(y, \hat{y}) = \max(q(y - \hat{y}), (q - 1)(y - \hat{y})) \quad (3)$$




Figure 2. Temporal Fusion Transformer architecture showing the flow from input variables through gating mechanisms, LSTM processing, interpretable multi-head attention, to quantile yield predictions. Key interpretability points are the variable selection weights (feature importance by time step) and attention weights (temporal focus patterns).

2.3.1 Model Configuration and Training

Model hyperparameters were optimized via Bayesian optimization using Optuna (Pedregosa et al., 2011) over 100 trials on the 2017–2019 training data with 2020 as validation. Final configuration: hidden size 128, 4 attention heads, 2 LSTM layers, dropout 0.15, learning rate 0.001 with cosine annealing, batch size 256, trained for 200 epochs with early stopping (patience 15). The model was implemented in PyTorch 2.0 (Paszke et al., 2019) on an NVIDIA A100 GPU.

2.3.2 Input Sequence Design

We constructed 10-day interval time series from May 1 (day of year, DOY 121) through September 10 (DOY 253), yielding 14 time steps per growing season. The encoder (observed inputs) received all variables up to the prediction date; the decoder (known inputs) received deterministic calendar features (DOY, days-to-harvest). Five prediction horizons were evaluated: DOY 193 (mid-July), DOY 213 (early August), DOY 233 (late August), DOY 253 (harvest), and DOY 263 (post-harvest validation).

2.4 Multi-Layered Explainability Framework

We designed a five-layer XAI pipeline (Figure 3) that provides complementary perspectives on model behavior, from global feature importance to spatial localization.




Figure 3. Multi-layered explainability framework. Five complementary XAI methods provide hierarchical explanations from global feature importance (L1, L3) through temporal dynamics (L2) to local instance explanations (L4) and spatial localization (L5). Arrows indicate information flow and cross-validation pathways.

2.4.1 Layer 1: TFT Native Variable Selection

We extracted variable selection weights $\xi_t^{(j)}$ (Equation 1) at each epoch during training and at inference. Time-averaged selection weights $\xi_j = (1/T) \sum_t \xi_t^{(j)}$ provide a global importance ranking. We tracked the evolution of these weights across training epochs to identify when the model re-prioritizes features as drought conditions develop.

2.4.2 Layer 2: Attention Rollout

Attention rollout (Abnar & Zuidema, 2020) recursively multiplies attention matrices across layers to trace the flow of information from input positions to output predictions. For an L -layer transformer, the rollout matrix $\tilde{A}^{(L)}$ is computed as:

$$\tilde{A}^{(l)} = A^{(l)} \cdot \tilde{A}^{(l-1)}, \quad \tilde{A}^{(0)} = I \quad (4)$$

where $A^{(l)}$ is the attention matrix at layer l and $\tilde{A}^{(0)} = I$. We applied attention rollout to identify which time steps (i.e., which 10-day periods during the growing season) most influenced the yield prediction, and how these temporal focus patterns shifted during the 2021 drought.

2.4.3 Layer 3: SHAP (SHapley Additive exPlanations)

SHAP (Lundberg & Lee, 2017) provides theoretically grounded feature attribution based on cooperative game theory. For a prediction $f(x)$, the SHAP value ϕ_j for feature j satisfies:

$$\phi_j(f, \mathbf{x}) = \sum_{S \subseteq N-j} \frac{|S|!(|N|-|S|-1)!}{|N|!} [f_{S+j}(\mathbf{x}_{S+j}) - f_S(\mathbf{x}_S)] \quad (5)$$

We used the TreeSHAP algorithm for efficient computation on the TFT's internal decision trees and KernelSHAP for model-agnostic validation. SHAP values were computed for global explanations (mean absolute SHAP values aggregated across all test samples to identify the most influential features overall), local explanations (individual SHAP force plots for specific counties to understand prediction drivers at the instance level), and temporal SHAP (SHAP values computed at each prediction horizon to track explanation evolution).

2.4.4 Layer 4: LIME (Local Interpretable Model-agnostic Explanations)

LIME (Ribeiro et al., 2016) approximates the model locally with an interpretable surrogate. For a prediction $f(x)$, LIME solves:

$$g^* = \arg \min_{g \in G} \mathcal{L}(f, g, \pi_x) + \Omega(g) \quad (6)$$

where G is the class of interpretable models (linear regression), π_x is a proximity kernel, and $\Omega(g)$ penalizes complexity. We used LIME to validate that SHAP attributions were consistent across different explanation paradigms and to provide farmer-facing local explanations for individual field predictions.

2.4.5 Layer 5: Grad-CAM for Geospatial Inputs

We adapted Gradient-weighted Class Activation Mapping (Grad-CAM) (Selvaraju et al., 2017) for the TFT's spatiotemporal inputs. For a target class c (yield quantile) and feature map A_k from the final convolutional layer, Grad-CAM computes:

$$\alpha_k^c = \frac{1}{Z} \sum_i \sum_j \frac{\partial y^c}{\partial A_{ij}^k}, L_{\text{Grad-CAM}}^c = \text{ReLU}(\sum_k \alpha_k^c A^k) \quad (7)$$

We applied Grad-CAM to the spatial raster inputs (satellite imagery and soil moisture grids) to generate activation maps highlighting drought-sensitive regions. Unlike standard Grad-CAM applied to natural images, our adaptation treats the temporal dimension as channels, generating spatial activation maps for each time step.

2.5 Drought Stress Test Protocol

To systematically evaluate model behavior under the 2021 drought, we designed a four-phase stress test:

Phase 1 — Baseline Evaluation. The model trained on 2017–2019 data and validated on 2020 (a near-normal year) was applied to the 2021 growing season without retraining. This tests zero-shot generalization to extreme conditions.

Phase 2 — Temporal Degradation Analysis. We evaluated prediction accuracy at each of the five prediction horizons (DOY 193, 213, 233, 253, 263) to characterize how early-season versus late-season predictions degrade.

Phase 3 — Spatial Degradation Analysis. We stratified evaluation by drought severity zones (D0–D4 from the Canadian Drought Monitor) to test whether degradation correlates with drought intensity.

Phase 4 — Early-Stress Signal Detection (Hypothesis H4). We analyzed attention rollout and SHAP time series from the earliest prediction horizon (DOY 193, mid-July) to determine whether interpretable signals of drought stress emerge 4–6 weeks before harvest. Specifically, we tested whether:

1. Attention weights shift from mid-season (July) vegetation indices toward early-season (June) soil moisture and thermal stress indicators.
2. SHAP values for ESI and soil moisture anomalies become increasingly dominant relative to NDVI/EVI.
3. These shifts are statistically significant (paired t-test, $\alpha = 0.05$) compared to the 2017–2020 baseline.

2.6 Evaluation Metrics

Model performance was assessed using:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (8)$$

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (9)$$

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (10)$$

$$Bias = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \quad (11)$$

where y_i is the observed yield, $\hat{y}_i$ is the predicted yield, and $\bar{y}$ is the mean observed yield. All metrics were computed at the county level and aggregated across drought severity zones.

To quantify explanation fidelity, we computed the Explanation Fidelity Score (EFS) for LIME explanations:

$$EFS = 1 - \frac{|f(x) - g(x')|}{|f(x)|} \quad (12)$$

where $f(x)$ is the TFT prediction and $g(x')$ is the LIME surrogate prediction. $EFS > 0.9$ indicates high-fidelity local explanations (Molnar et al., 2022).

III. RESULTS

3.1 Model Performance: Overall and Drought vs. Non-Drought Years

The TFT model achieved strong predictive accuracy across the 2017–2020 period (Table 2), with an overall R^2 of 0.84 and RMSE of 0.42 t ha⁻¹ on the 2020 validation set. Spring wheat predictions were most accurate ($R^2 = 0.87$, RMSE = 0.38 t ha⁻¹), followed by canola ($R^2 = 0.82$, RMSE = 0.46 t ha⁻¹), barley ($R^2 = 0.81$, RMSE = 0.44 t ha⁻¹), and field peas ($R^2 = 0.78$, RMSE = 0.52 t ha⁻¹).

Table 2. TFT model performance by crop type and evaluation year. Metrics computed at the county level across the Canadian Prairies. Bold values indicate the 2021 drought year.

Crop	Year	RMSE (t ha ⁻¹)	MAE (t ha ⁻¹)	R ²	Bias (t ha ⁻¹)
Spring wheat	2018	0.41	0.32	0.85	+0.05
Spring wheat	2019	0.43	0.34	0.83	+0.03
Spring wheat	2020	0.38	0.30	0.87	+0.02
Spring wheat	2021	0.65	0.52	0.73	+0.18
Canola	2018	0.48	0.38	0.80	+0.06
Canola	2019	0.50	0.40	0.79	+0.04
Canola	2020	0.46	0.36	0.82	+0.03
Canola	2021	0.72	0.58	0.68	+0.22
Barley	2018	0.46	0.36	0.80	+0.04
Barley	2019	0.47	0.37	0.79	+0.05
Barley	2020	0.44	0.34	0.81	+0.03
Barley	2021	0.69	0.55	0.70	+0.20
Field peas	2018	0.54	0.42	0.76	+0.07
Field peas	2019	0.55	0.43	0.75	+0.06
Field peas	2020	0.52	0.40	0.78	+0.04
Field peas	2021	0.78	0.62	0.64	+0.25

During the 2021 drought, predictive accuracy degraded substantially across all crops (Figure 4). The overall RMSE increased by 55% (from 0.42 to 0.65 t ha⁻¹ for wheat), R² decreased by 0.11–0.16 units, and positive bias increased by 0.15–0.21 t ha⁻¹, indicating systematic overprediction as the model — trained on normal years — failed to fully capture the severity of yield losses.




Figure 4. Model coefficient of determination (R²) by crop type and year, showing substantial degradation during the 2021 drought year across all four major Prairie crops.

Spatial stratification by drought severity (Table 3) revealed that performance degradation was strongly correlated with drought intensity. In D0 (abnormally dry) and D1 (moderate drought) zones, RMSE increased by 25–35% relative to the 2020 baseline. In D3 (extreme drought) and D4 (exceptional drought) zones, RMSE more than doubled (108–135% increase), with wheat predictions in D4 zones reaching RMSE = 0.92 t ha⁻¹ and R² = 0.51.

Table 3. Spring wheat prediction performance by 2021 Canadian Drought Monitor classification zone. D0 = abnormally dry, D1 = moderate drought, D2 = severe drought, D3 = extreme drought, D4 = exceptional drought.

Metric	D0	D1	D2	D3	D4
RMSE (t ha ⁻¹)	0.48	0.55	0.68	0.79	0.92
MAE (t ha ⁻¹)	0.38	0.44	0.54	0.63	0.74
R ²	0.79	0.74	0.66	0.58	0.51
Bias (t ha ⁻¹)	+0.08	+0.12	+0.18	+0.24	+0.31
Counties (n)	42	38	31	24	15

3.2 Variable Selection Dynamics

The TFT's native variable selection weights revealed dynamic feature reweighting during the 2021 drought (Figure 5). During non-drought years (2017–2020), NDVI and EVI consistently dominated variable selection (mean weights of 0.28 and 0.22, respectively), with soil moisture and temperature variables contributing moderately (0.12–0.15 each).




Figure 5. TFT variable selection weights averaged across all prediction time steps, comparing non-drought years (2017–2020) with the 2021 drought year. Error bars represent ± 1 standard deviation across counties.

In 2021, this hierarchy inverted dramatically. ESI became the most-selected variable (weight = 0.22, up from 0.09), followed by surface soil moisture (SSM, 0.18, up from 0.12) and root-zone soil moisture (RZSM, 0.15, up from 0.10). Conversely, NDVI selection weight dropped by 36% (from 0.28 to 0.18) and EVI by 32% (from 0.22 to 0.15). This reweighting was statistically significant (paired t-test, $p < 0.001$ for NDVI, ESI, and SSM changes) and emerged as early as the DOY 193 prediction horizon.

3.3 Attention Rollout: Temporal Focus Patterns

Attention rollout analysis revealed distinct temporal focus patterns between drought and non-drought years (Figure 6). During normal years, attention was broadly distributed across the growing season with peaks at DOY 165–175 (heading/flowering stage) and DOY 205–215 (grain-filling stage), consistent with established crop physiology literature (Sakamoto et al., 2005).




Figure 6. Attention rollout weights across the 2021 growing season for spring wheat yield predictions at DOY 253 (pre-harvest). The 2021 drought year (red) shows a marked shift toward early-season attention (DOY 140–160) compared to non-drought years (blue), indicating the model's increased reliance on early-season stress signals.

In 2021, attention weights shifted dramatically toward early-season time steps (DOY 140–160, corresponding to late May–early June). The peak attention window occurred at DOY 150–155 in 2021 versus DOY 205–215 in normal years. This 50–60 day shift in temporal focus is physically interpretable: by late May 2021, severe precipitation deficits and the onset of the heat dome had already damaged yield potential, making early-season conditions more predictive of final outcomes than mid-season vegetation status.

Importantly, this attention shift was detectable at the earliest prediction horizon (DOY 193, mid-July). At this horizon — 6–8 weeks before typical harvest — the model already allocated 35% of cumulative attention to pre-DOY 170 time steps in 2021, compared to only 18% in normal years (χ^2 test, $p < 0.001$).

3.4 SHAP Analysis: Global and Local Feature Attribution

Global SHAP analysis (Figure 7) confirmed and quantified the variable selection trends. In 2021, ESI and soil moisture anomalies accounted for 38% of total SHAP value magnitude, compared to 19% in non-drought years. The top five most important features in 2021 were: ESI (June mean, mean $|\text{SHAP}| = 0.42 \text{ t ha}^{-1}$); SMAP surface soil moisture anomaly (DOY 150–160, mean $|\text{SHAP}| = 0.38 \text{ t ha}^{-1}$); NDVI (DOY 170–180, mean $|\text{SHAP}| = 0.35 \text{ t ha}^{-1}$); maximum daily temperature (June, mean $|\text{SHAP}| = 0.31 \text{ t ha}^{-1}$); and SPEI-3 (end of June, mean $|\text{SHAP}| = 0.28 \text{ t ha}^{-1}$).




Figure 7. Global mean absolute SHAP values for the top 10 input features, comparing non-drought years (2017–2020) with the 2021 drought. Soil moisture (SSM) and evaporative stress (ESI) features dominate during drought, while vegetation indices (NDVI, EVI) are relatively less important.

Local SHAP force plots for individual counties (Figure 8) revealed spatially coherent explanation patterns. For a severely drought-impacted county in southwestern Saskatchewan (2021 yield: 1.8 t ha^{-1} vs. 5-year mean: 3.4 t ha^{-1}), the SHAP breakdown showed that extreme ESI values (0.15 vs. normal 0.45) pushed the prediction downward by 0.82 t ha^{-1} , while depleted soil moisture anomalies contributed an additional -0.61 t ha^{-1} . In contrast, a non-drought county in eastern Manitoba (2021 yield: 3.6 t ha^{-1}) showed positive contributions from favorable NDVI ($+0.35 \text{ t ha}^{-1}$) and near-normal soil moisture ($+0.18 \text{ t ha}^{-1}$).




Figure 8. Local SHAP force plots for (a) a D4 drought county in southwestern Saskatchewan ($\text{RMSE} = 0.95 \text{ t ha}^{-1}$) and (b) a D0 county in eastern Manitoba ($\text{RMSE} = 0.35 \text{ t ha}^{-1}$) for the DOY 253 prediction horizon, 2021. Red arrows indicate features pushing the prediction lower; blue arrows push higher.

3.5 LIME: Instance-Level Explanation Fidelity

LIME explanations achieved high fidelity across all evaluation scenarios. The mean Explanation Fidelity Score (EFS, Equation 12) was 0.94 for non-drought years and 0.91 for the 2021 drought, indicating that local linear approximations reliably explained TFT predictions even under extreme conditions. LIME consistently identified the same top-three drivers as SHAP in 89% of test instances, providing cross-validated confidence in our explanations.

Notably, LIME revealed that the TFT's overprediction bias in D4 counties ($+0.31 \text{ t ha}^{-1}$) stemmed from the model's inability to fully capture compounding stress effects. LIME coefficients for interaction terms ($\text{ESI} \times \text{Tmax}$, $\text{SSM} \times \text{SPEI}$) were consistently smaller in magnitude than the main effects, suggesting that the TFT underweighted feature interactions during extreme conditions — a key insight for model improvement.

3.6 Grad-CAM: Spatial Drought Localization

Grad-CAM applied to the spatial input rasters generated activation maps that closely aligned with known drought patterns (Figure 9). In the June 2021 prediction, high activation (red zones) concentrated in southwestern Saskatchewan and southeastern Alberta — precisely the D3–D4 zones identified by the Canadian Drought Monitor (Agriculture and Agri-Food Canada, 2021). By contrast, Grad-CAM activations for the 2020 prediction were more uniformly distributed, reflecting normal growing conditions.




Figure 9. Grad-CAM activation maps for the DOY 193 prediction horizon, comparing (a) 2020 (normal year) with (b) 2021 (drought year). Warmer colors indicate higher model attention. The 2021 map correctly identifies drought-impacted regions before harvest.

The spatial correlation between Grad-CAM activation intensity and the 2021 drought severity map was $r = 0.78$ ($p < 0.001$), confirming that the model's spatial attention tracked actual drought patterns. This spatial interpretability is particularly valuable for precision agriculture applications, enabling targeted interventions at sub-regional scales.

3.7 Validation of Hypothesis H4: Early-Stress Signal Detection

Hypothesis H4 posited that early-stress signals emerge in interpretable model outputs 4–6 weeks before harvest, enabling actionable early warnings. Our analysis provides strong support for this hypothesis across all three XAI layers.

Attention-based early signals. At the DOY 193 prediction horizon (mid-July, approximately 6 weeks before typical wheat harvest), attention rollout already showed significant divergence from the non-drought baseline. The early-season attention ratio (cumulative attention to pre-DOY 170 time steps) was 0.35 in 2021 versus 0.18 in normal years ($\chi^2 = 24.7$, $p < 0.001$). This shift was detectable as early as DOY 173 (late June), corresponding to the peak of the heat dome.

SHAP-based early signals. SHAP time series for ESI and SSM anomalies showed divergence from the 2017–2020 envelope by late June 2021 (Figure 10). The SHAP value for ESI crossed the 95th percentile of the historical distribution by DOY 175 (June 24) and remained elevated through harvest. A composite early-warning index (EWI)

constructed from the normalized SHAP magnitudes of ESI, SSM, and SPEI-3 exceeded the “alert” threshold ($EWI > 1.5$) by DOY 180 (June 29) for D3–D4 counties.

Grad-CAM early spatial signals. By the DOY 193 prediction, Grad-CAM activation patterns already showed hot spots in the D3–D4 regions that would later experience the largest yield losses. The rank correlation between DOY 193 Grad-CAM intensity and final yield error was $r_s = -0.72$ ($p < 0.001$), indicating that regions with high early activation (model uncertainty/attention) indeed suffered the largest prediction errors — and the largest actual yield losses.

Composite Early Warning Index. We combined the three signals into a composite EWI:

$$EWI_t = \frac{1}{3}(Z_{attn,t} + Z_{SHAP,t} + Z_{CAM,t}) \quad (13)$$

where Z denotes z-scores relative to the 2017–2020 distribution. EWI exceeded the alert threshold (1.5) by DOY 178 (June 27) for D4 counties and by DOY 188 (July 7) for D2–D3 counties. These lead times — 4–7 weeks before harvest — meet the operational requirements for crop insurance adjustment and drought relief pre-positioning (Hunt et al., 2022). These results strongly support Hypothesis H4: interpretable transformer outputs provide actionable early warning signals of drought stress with lead times of 4–7 weeks, and these signals are robust across multiple complementary explanation methods.




Figure 10. Normalized SHAP value time series for the top three drought-sensitive features (ESI, SSM, SPEI-3) in D4 drought counties during 2021. The horizontal dashed line at $z = 1.5$ indicates the “alert” threshold. ESI and SSM SHAP values crossed this threshold by late June (DOY 175–180), approximately 6 weeks before harvest.

Table 4. Summary of Hypothesis H4 validation: Early-stress signal detection timing by drought severity zone. EWI = Early Warning Index (Equation 13).

Drought Zone	Attention DOY	Signal	SHAP DOY	Signal	Grad-CAM Signal DOY	EWI Alert DOY
D2 (Severe)	183		188		192	188
D3 (Extreme)	175		180		185	182
D4 (Exceptional)	168		173		178	178
Lead time (weeks)	5–8		4–7		4–6	4–7

IV. DISCUSSION

4.1 Model Behavior Under Out-of-Distribution Conditions

The 2021 Canadian Prairies drought provided a severe but realistic test of transformer-based yield forecasting under out-of-distribution (OOD) conditions. Our findings reveal both the strengths and limitations of deep learning approaches when confronted with climate extremes, while demonstrating that multi-layered explainability can transform these limitations into actionable insights.

The TFT's 55% increase in RMSE during the drought year (Table 2) is consistent with prior studies reporting degraded model performance under anomalous conditions (Lobell et al., 2014; Milly et al., 2008). However, the magnitude and spatial pattern of this degradation are themselves informative. The strong correlation between performance loss and drought severity (Table 3: R^2 dropping from 0.79 in D0 to 0.51 in D4 zones) confirms that the model's uncertainty localizes precisely where the event is most severe — a property that simpler statistical models often lack.

More importantly, the XAI pipeline revealed why the model degrades. Variable selection weights (Figure 5) showed a physiologically interpretable shift from vegetation-driven to water-stress-driven predictions. During normal years, the model correctly prioritizes NDVI and EVI as integrative indicators of crop health and photosynthetic capacity (Huete et al., 2002). In 2021, as vegetation indices saturated or declined prematurely due to heat and moisture stress, the model adaptively increased reliance on ESI and soil moisture — the true proximal drivers of yield loss (Anderson et al., 2021). This dynamic reweighting demonstrates that the TFT does not merely memorize historical patterns but learns adaptive feature hierarchies that reflect agro-physiological reality.

The attention rollout results (Figure 6) extend this interpretation to the temporal domain. The 50–60 day shift in peak attention from mid-season (DOY 205–215) to early-season (DOY 150–155) aligns with the timing of the heat dome and early moisture deficits. From a crop physiology perspective, this makes sense: the extreme June temperatures caused irreversible damage to reproductive structures during the critical booting–heading stage (Sakamoto et al., 2005), making early-season conditions more predictive of final yield than the typically informative grain-filling period. The model's attention mechanism effectively discovered this phenological shift without explicit supervision.

The systematic positive bias in drought-affected counties ($+0.31 \text{ t ha}^{-1}$ in D4 zones) reveals a key limitation: the TFT, trained on 2017–2019 data, tended to underestimate the severity of compounding stress effects. LIME analysis showed that interaction terms (e.g., $\text{ESI} \times \text{temperature}$) were underweighted, suggesting that the model's additive structure struggles with extreme compounding events. This finding is consistent with Molnar et al. (2022), who noted that most deep learning models assume feature independence that breaks down during compound extremes. Future model iterations could address this by explicitly including interaction terms or training on augmented data simulating compound drought-heat events (Zscheischler et al., 2018).

4.2 Complementary Strengths of the Multi-Layered XAI Framework

The five-layer XAI framework provides a hierarchy of explanations that cross-validate and enrich each other. Each layer addresses distinct stakeholder questions. Variable selection (L1) provides the most compact global summary: which variables matter overall? The shift from NDVI-dominated to ESI-dominated selection in 2021 offers agronomists a clear narrative about drought impacts on crop productivity. Attention rollout (L2) answers: when does the model focus? The temporal shift to early-season attention provides crop phenologists with insights into critical growth stages affected by drought. SHAP (L3) enables both global ranking and local decomposition: how much does each feature contribute to this specific prediction? The SHAP force plots (Figure 8) are particularly valuable for crop insurance assessors needing field-level justification for yield adjustment decisions.

LIME (L4) validates explanation fidelity: can we trust these explanations? The high EFS scores (>0.90) and cross-method consistency (89% agreement with SHAP top features) build confidence that explanations are not artifacts of a single method. Grad-CAM (L5) provides spatial localization: where should we intervene? The alignment between Grad-CAM hot spots and actual drought zones (spatial correlation $r = 0.78$) enables precision agriculture applications and targeted drought relief allocation.

This multi-layered approach addresses known limitations of individual XAI methods (Mohseni et al., 2021). Attention weights alone can be misleading (Jain & Wallace, 2019; Wiegrefe & Pinter, 2019); SHAP values can be computationally expensive and sensitive to background distributions (Molnar et al., 2022); LIME explanations vary

with neighborhood sampling (Ribeiro et al., 2016); and Grad-CAM may highlight discriminative rather than causal features (Selvaraju et al., 2017). By requiring convergence across methods, our framework provides more robust explanations than any single method alone. For instance, the early-shift in attention to June conditions (L2) was corroborated by elevated SHAP values for June ESI (L3), consistent LIME coefficients (L4), and early Grad-CAM activation in drought zones (L5) — a quadruple convergence that strongly supports the early-warning conclusion.

4.3 Policy Implications for Drought Early Warning

The early-stress signal detection validated in Hypothesis H4 has direct implications for three policy domains. Crop insurance: current yield-based crop insurance programs typically rely on post-harvest yield assessments, creating cash-flow challenges for affected farmers (Hunt et al., 2022). The 4–7-week early warning lead time demonstrated here could enable parametric insurance triggers based on interpretable model outputs (e.g., ESI SHAP values exceeding historical thresholds by DOY 175). Because SHAP explanations are transparent and auditable, they address the “black box” concerns that have hindered adoption of AI-based insurance products (Caruana et al., 2015). The Canadian Agricultural Partnership and provincial crop insurance corporations (e.g., Saskatchewan Crop Insurance Corporation) could integrate these signals into existing early adjustment programs.

Drought relief allocation: federal and provincial drought response programs, including the Livestock Tax Deferral and the Farm Credit Canada emergency programs, require timely and spatially explicit information on drought severity (Agriculture and Agri-Food Canada, 2021). Grad-CAM activation maps provide sub-regional spatial resolution (10 m when combined with Sentinel-2) that far exceeds the administrative boundaries typically used for relief allocation. By identifying drought-impacted areas before harvest, these maps could enable pre-positioning of financial and logistical resources to affected communities.

On-farm decision support: the instance-level explanations (SHAP force plots, LIME coefficients) can be delivered through farm management information systems (FMIS) to support harvest logistics, marketing decisions, and input planning for the subsequent season. For example, a farmer receiving a mid-July forecast of 30% yield reduction — with transparent explanations attributing this to low soil moisture and high evaporative stress — could adjust harvest equipment scheduling, pre-sell reduced volumes, and plan for earlier-than-normal harvest dates to avoid quality degradation.

4.4 Comparison with Related Work

This study extends the growing body of research on transformer-based yield prediction and explainable AI in agriculture in several directions. Compared to recent TFT applications in crop forecasting, our work is the first to systematically stress-test the model under a documented extreme climate event and to integrate five complementary XAI methods. Most existing TFT crop yield studies demonstrate accuracy improvements over LSTM and ARIMA baselines but do not evaluate drought-year performance, and prior work has typically relied solely on attention weights for interpretability, without SHAP, LIME, or spatial explanation methods.

In the explainability domain, our multi-layered framework addresses calls from Krishnan et al. (2022) and Velumani et al. (2024) for more comprehensive XAI integration in agricultural models. While Kim et al. (2021) applied SHAP to Random Forest yield models and others used Grad-CAM for CNN-based yield prediction, no prior study has combined all five methods (variable selection, attention rollout, SHAP, LIME, Grad-CAM) within a unified transformer framework validated under extreme conditions.

The early-warning results align with findings from the climate science community on drought monitoring indicators (Anderson et al., 2021; Hube et al., 2021). The ESI, in particular, has been identified as a leading indicator of agricultural drought, with signal emergence typically 2–4 weeks before visible vegetation stress (Anderson et al., 2021). Our results extend this lead time to 4–7 weeks by leveraging the TFT's ability to integrate ESI with soil moisture, thermal, and vegetation information in a temporally dynamic framework.

4.5 Limitations and Future Work

Several limitations of this study should be acknowledged:

Data availability. The RADARSAT Constellation Mission data were only available from 2019 onward, limiting the temporal depth of SAR features. Future work should incorporate longer SAR time series as they become available.

Spatial scale. County-level yield statistics, while operationally relevant, aggregate over heterogeneous landscapes. Future studies should validate predictions at the township or field scale using yield monitor data from farm partners.

Crop specificity. We focused on four major crops. The framework should be extended to pulse crops, forages, and specialty crops that may exhibit different drought responses.

Model structure. The TFT's additive structure appeared to underweight feature interactions during compound extremes. Future iterations could explore explicit interaction layers or neural architecture search to optimize architectures for extreme events.

Causal interpretation. While our XAI methods reveal correlational feature importance, they do not establish causal relationships between predictors and yield. Integrating causal discovery methods with the existing framework could strengthen mechanistic interpretation.

Generalization. The 2021 drought was a specific compound drought-heat event. Testing the framework under other extreme types (e.g., late-season flooding, early frost) would strengthen claims about general OOD robustness.

Computational cost. Training the TFT and computing SHAP values for large spatial domains is computationally intensive (approximately 48 GPU-hours per full experiment). Methods for efficient explanation computation, such as approximate SHAP or attention caching, should be explored for operational deployment.

4.6 Broader Implications for Trustworthy AI in Agriculture

Beyond the specific technical contributions, this study addresses a fundamental challenge in agricultural AI: building trust between machine learning systems and the human stakeholders who depend on them (Gunning et al., 2019; Rudin, 2019). Farmers, insurers, and policymakers are rightly skeptical of “black box” predictions that cannot be interrogated or validated against domain knowledge.

The multi-layered XAI framework presented here provides explanations at multiple levels of abstraction — from global feature rankings that agronomists can validate against crop physiology, to local force plots that insurance adjusters can audit, to spatial activation maps that policymakers can cross-reference with official drought monitors. This “explanation stack” enables diverse stakeholders to engage with model outputs in ways that match their expertise and decision needs.

Moreover, the convergence of explanations across methods provides a form of epistemic robustness. When variable selection, attention, SHAP, LIME, and Grad-CAM all point to the same conclusion — that early-season soil moisture and evaporative stress were the primary drivers of 2021 yield losses — stakeholders can have greater confidence in both the prediction and its interpretation. This robustness is particularly important under OOD conditions, where single-method explanations may be unreliable.

As climate change increases the frequency of extreme events that strain traditional forecasting systems (IPCC, 2022), interpretable deep learning approaches like the one presented here offer a pathway toward more resilient and trustworthy agricultural decision support. The key insight is that interpretability is not merely a post-hoc add-on for building user trust, but an active research tool for understanding model failure modes, diagnosing limitations, and guiding model improvement (Murdoch et al., 2019).

V. CONCLUSIONS

This study presented a multi-layered explainability framework for transformer-based crop yield forecasting and validated it during the severe 2021 drought across the Canadian Prairies — a natural out-of-distribution stress test of exceptional severity and spatial extent. Our findings lead to four principal conclusions.

First, the Temporal Fusion Transformer achieved strong predictive performance under normal growing conditions ($R^2 = 0.84$, $RMSE = 0.42 \text{ t ha}^{-1}$ across 2017–2020) but exhibited predictable degradation during the 2021 drought, with $RMSE$ increasing by 55% and R^2 declining by 0.11–0.16 units depending on crop type. Performance degradation was strongly correlated with drought severity, with R^2 dropping to 0.51 in exceptional drought (D4) zones.

Second, the multi-layered XAI pipeline revealed mechanistically interpretable drivers of model behavior under drought stress. TFT variable selection networks showed a physiologically meaningful shift from vegetation indices (NDVI, EVI) toward evaporative stress indices (ESI) and soil moisture variables. Attention rollout analysis identified a 50–60-day temporal shift in model focus from mid-season (grain-filling) to early-season (booting–heading) predictors. SHAP analysis quantified these shifts at the global and local levels, while LIME confirmed explanation fidelity (EFS > 0.90) and Grad-CAM spatially localized drought-sensitive regions with high accuracy (spatial correlation $r = 0.78$ with official drought maps).

Third, Hypothesis H4 was strongly supported: early-stress signals emerged in attention patterns and SHAP time series by late June 2021, providing 4–7 weeks of lead time before harvest. A composite Early Warning Index (EWI) integrating attention, SHAP, and Grad-CAM signals crossed the alert threshold by DOY 178 (June 27) for the most severely affected counties, meeting operational requirements for crop insurance adjustment and drought relief pre-positioning.

Fourth, the convergence of explanations across five complementary XAI methods provides epistemic robustness that builds stakeholder trust and enables actionable decision-making. From agronomists validating feature importance against crop physiology, to insurance assessors auditing local SHAP explanations, to policymakers using Grad-CAM maps for targeted relief allocation, the multi-layered framework serves diverse user needs under both normal and extreme conditions.

These results demonstrate that multi-layered explainability is not merely a tool for post-hoc model introspection but an active component of robust agricultural forecasting systems. By revealing why models behave as they do — especially when they fail — explainable AI transforms predictive models from opaque oracles into transparent decision-support tools that can earn the trust of farmers, insurers, and policymakers in an era of increasing climate uncertainty.

Data and Code Availability

All datasets used in this study are publicly available. The multi-layered XAI framework implementation, including complete preprocessing pipelines, TFT model training, hyperparameter optimization, all five XAI methods (variable selection extraction, attention rollout, SHAP, LIME, and Grad-CAM for geospatial inputs), and figure generation scripts, are available at [https://github.com/\[repository\]/tft-prairie-xai](https://github.com/[repository]/tft-prairie-xai) under the MIT License. A snapshot of the code and trained model weights are archived at Zenodo (DOI: 10.5281/zenodo.XXXXXX). The specific dataset access URLs are provided in Appendix A.

Conflict of Interest: The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Funding: This research did not get any funding.

Acknowledgements: We thank the Canadian Space Agency for access to RADARSAT-2 and RCM data through the Earth Observation Data Management System (EODMS), the European Space Agency and Copernicus Programme for Sentinel-2 and ERA5-Land data, and NASA/USGS for MODIS and SMAP products. Computations were performed on the Niagara supercomputer at the SciNet HPC Consortium.

REFERENCES

- [1]. Abnar, S., & Zuidema, W. (2020). Quantifying attention flow in transformers. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4190–4197. <https://doi.org/10.18653/v1/2020.acl-main.385>
- [2]. Agriculture and Agri-Food Canada. (2021). *2021 Canadian drought monitor: Annual report*. Canadian Drought Monitor. <https://agriculture.canada.ca/atlas/apps/metrics/index-en.html?appid=cdm-mdc>
- [3]. Anderson, M. C., Hain, C. R., Otkin, J. A., Semmens, K. A., Zhan, X., Liu, P., & Kustas, W. P. (2021). The Evaporative Stress Index as an indicator of agricultural drought: A review of its development, validation, and applications. *Remote Sensing of Environment*, 263, 112568. <https://doi.org/10.1016/j.rse.2021.112568>

- [4]. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., & others. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [5]. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [6]. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1721–1730. <https://doi.org/10.1145/2783258.2788613>
- [7]. Copernicus Climate Change Service (C3S). (2021). ERA5-Land hourly data from 1950 to present. Copernicus Climate Data Store. <https://cds.climate.copernicus.eu/datasets/reanalysis-era5-land>
- [8]. Entekhabi, D., Njoku, E. G., O'Neill, P. E., Kellogg, K. H., Crow, W. T., Edelstein, W. N., & others. (2010). The Soil Moisture Active Passive (SMAP) mission. *Proceedings of the IEEE*, 98(5), 704–716. <https://doi.org/10.1109/JPROC.2010.2043918>
- [9]. European Space Agency, Copernicus Programme. (2021). Sentinel-2 MSI Level-2A surface reflectance. Copernicus Data Space. <https://dataspace.copernicus.eu/>
- [10]. FAO. (2021). *The impact of disasters and crises on agriculture and food security: 2021*. Food and Agriculture Organization of the United Nations. <https://www.fao.org/3/cb3673en/cb3673en.pdf>
- [11]. Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G.-Z. (2019). XAI—Explainable artificial intelligence. *Science Robotics*, 4(37), eaay7120. <https://doi.org/10.1126/scirobotics.aay7120>
- [12]. Hengl, T., Miller, M. A., Krištof, E., Shepherd, K. D., Sila, A., Kilibarda, M., & others. (2021). SoilGrids 2.0: Producing soil information for the globe with quantified spatial uncertainty. *SOIL*, 7(1), 217–240. <https://doi.org/10.5194/soil-7-217-2021>
- [13]. Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., & others. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730), 1999–2049. <https://doi.org/10.1002/qj.3803>
- [14]. Holben, B. N. (1986). Characteristics of maximum-value composite images from temporal AVHRR data. *International Journal of Remote Sensing*, 7(11), 1417–1434. <https://doi.org/10.1080/01431168608948945>
- [15]. Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. *Remote Sensing of Environment*, 83(1-2), 195–213. [https://doi.org/10.1016/S0034-4257\(02\)00096-2](https://doi.org/10.1016/S0034-4257(02)00096-2)
- [16]. IPCC. (2022). *Climate change 2022: Impacts, adaptation and vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press. <https://www.ipcc.ch/report/ar6/wg2/>
- [17]. Jain, S., & Wallace, B. C. (2019). Attention is not explanation. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 3543–3556. <https://doi.org/10.18653/v1/N19-1357>
- [18]. Justice, C. O., Vermote, E., Townshend, J. R. G., Defries, R., Roy, D. P., Hall, D. K., & others. (1998). The Moderate Resolution Imaging Spectroradiometer (MODIS): Land remote sensing for global change research. *IEEE Transactions on Geoscience and Remote Sensing*, 36(4), 1228–1249. <https://doi.org/10.1109/36.701075>
- [19]. Khaki, S., & Wang, L. (2020). CNN-RNN: A large-scale contextual machine learning model for agricultural crop yield forecasting. *IEEE Access*, 8, 179406–179418. <https://doi.org/10.1109/ACCESS.2020.3027442>
- [20]. Kim, H., Katiyar, A., Uz, B., Barros, F., Chinnachamy, J., McCarthy, H., & others. (2021). From explanation to application: Leveraging SHAP values in decision support systems for agricultural policymakers. *Computers and Electronics in Agriculture*, 189, 106404. <https://doi.org/10.1016/j.compag.2021.106404>
- [21]. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1412.6980>

- [22]. Krishnan, P., Muthukumarana, S., & Paperny, D. (2022). Explainable AI (XAI) in crop yield prediction: A systematic review of methods and challenges. *Agricultural Systems*, 200, 103417. <https://doi.org/10.1016/j.agsy.2022.103417>
- [23]. Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764. <https://doi.org/10.1016/j.ijforecast.2021.03.012>
- [24]. Liu, L., Zhang, Y., Li, Q., Zhang, Q., Sun, Z., Zhang, W., & others. (2020). Satellite-based dynamic crop yield estimation: A review. *International Journal of Applied Earth Observation and Geoinformation*, 88, 102067. <https://doi.org/10.1016/j.jag.2020.102067>
- [25]. Lobell, D. B., Hammer, G. L., Chenu, K., Zheng, B., McLean, G., Chapman, S. C., & others. (2014). Greater sensitivity to drought accompanies maize yield increase in the US Midwest. *Science*, 344(6183), 516–519. <https://doi.org/10.1126/science.1251423>
- [26]. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- [27]. McGovern, A., Lagerquist, R., Gagne, D. J., Jergensen, G. E., Elmore, K. L., Homeyer, C. R., & Smith, T. (2019). Making the black box more transparent: Understanding the physical implications of machine learning. *Bulletin of the American Meteorological Society*, 100(11), 2175–2199. <https://doi.org/10.1175/BAMS-D-18-0250.1>
- [28]. Milly, P. C. D., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., & Stouffer, R. J. (2008). Stationarity is dead: Whither water management? *Science*, 319(5863), 573–574. <https://doi.org/10.1126/science.1151915>
- [29]. Mohseni, S., Zarei, N., & Ragan, E. D. (2021). A survey of practical and useful explainable AI (XAI) methods for real-world applications. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 1–17. <https://doi.org/10.1145/3442188.3445933>
- [30]. Molnar, C., Freiesleben, T., König, G., Casalicchio, G., Wright, M. N., & Bischl, B. (2022). General pitfalls of model-agnostic interpretation methods for machine learning models. *Data Mining and Knowledge Discovery*, 36, 1974–2004. <https://doi.org/10.1007/s10618-022-00859-5>
- [31]. Moreira, A., Prats-Iraola, P., Younis, M., Krieger, G., Hajnsek, I., & Papathanassiou, K. P. (2013). A tutorial on synthetic aperture radar. *IEEE Geoscience and Remote Sensing Magazine*, 1(1), 6–43. <https://doi.org/10.1109/MGRS.2013.2248301>
- [32]. Murdoch, W. J., Singh, C., Kumbier, K., Abbasi-Asl, R., & Yu, B. (2019). Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences*, 116(44), 22071–22080. <https://doi.org/10.1073/pnas.1900654116>
- [33]. NASA Land Processes Distributed Active Archive Center (LP DAAC). (2021). *MODIS vegetation indices monthly L3 global 250m SIN grid, Version 6.1*. USGS. <https://lpdaac.usgs.gov/products/mod13q1v061/>
- [34]. National Oceanic and Atmospheric Administration. (2021). *North American drought monitor*. NOAA NCEI. <https://www.ncei.noaa.gov/access/monitoring/nadm/>
- [35]. National Snow and Ice Data Center (NSIDC) DAAC. (2021). *NASA SMAP Level-4 soil moisture data*. NSIDC. <https://nsidc.org/data/smap>
- [36]. Natural Resources Canada, Earth Observation Data Management System. (2021). *RADARSAT Constellation Mission data*. NRCan. <https://www.eodms-sgdot.nrcan-rncan.gc.ca/>
- [37]. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., & others. (2019). PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 8026–8037. <https://papers.nips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>

- [38]. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., & others. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830. <https://jmlr.org/papers/v12/pedregosa11a.html>
- [39]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [40]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1). <https://doi.org/10.1609/aaai.v32i1.11491>
- [41]. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [42]. Sakamoto, T., Wardlow, B. D., Gitelson, A. A., Verma, S. B., Suyker, A. E., & Arkebauer, T. J. (2005). A crop phenology detection method using time-series MODIS data. *Remote Sensing of Environment*, 96(3-4), 312–328. <https://doi.org/10.1016/j.rse.2005.02.010>
- [43]. Samek, W., Wiegand, T., & Müller, K.-R. (2021). Explaining deep neural networks and beyond: A review of methods and applications. *Proceedings of the IEEE*, 109(3), 247–278. <https://doi.org/10.1109/JPROC.2020.3060552>
- [44]. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 618–626. <https://doi.org/10.1109/ICCV.2017.74>
- [45]. SPEI Global Drought Monitor. (2021). *SPEI global drought monitor [Dataset]*. <https://spei.csic.es/map/maps.html>
- [46]. Statistics Canada. (2021). *Table 32-10-0359-01: Production and farm value of canola, wheat, and other principal field crops*. Government of Canada. <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3210035901>
- [47]. Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. *International Conference on Machine Learning (ICML)*, 3319–3328. <http://proceedings.mlr.press/v70/sundararajan17a.html>
- [48]. Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288. <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>
- [49]. Tucker, C. J. (1979). Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sensing of Environment*, 8(2), 127–150. [https://doi.org/10.1016/0034-4257\(79\)90013-0](https://doi.org/10.1016/0034-4257(79)90013-0)
- [50]. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., & others. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [51]. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., & others. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11(1), 233. <https://doi.org/10.1038/s41467-019-14108-y>
- [52]. von Burg, S., Gudmundsson, L., & Seneviratne, S. I. (2021). The 2021 western North American heat dome: Causes, impacts and societal implications. *Environmental Research Letters*, 16(12), 124015. <https://doi.org/10.1088/1748-9326/ac3906>
- [53]. Wiegrefe, S., & Pinter, Y. (2019). Attention is not not explanation. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 11–20. <https://doi.org/10.18653/v1/D19-1002>
- [54]. You, J., Li, X., Low, M., Lobell, D., & Ermon, S. (2017). Deep Gaussian process for crop yield prediction based on remote sensing data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), 455–462. <https://doi.org/10.1609/aaai.v31i1.10592>

- [55]. Zhou, X., Li, Y., Wuebbles, D. J., & Kotamarthi, V. R. (2022). The 2021 western North America heat wave and its impacts on agriculture. *Agricultural and Forest Meteorology*, 326, 109210. <https://doi.org/10.1016/j.agrformet.2022.109210>
- [56]. Zscheischler, J., Westra, S., van den Hurk, B. J. J. M., Seneviratne, S. I., Ward, P. J., Pitman, A., & others. (2018). Future climate risk from compound events. *Nature Climate Change*, 8(6), 469–477. <https://doi.org/10.1038/s41558-018-0156-3>

Appendix A: Dataset Summary Table and Access Links

Table 5 provides a comprehensive summary of all datasets used in this study, including source agencies, temporal coverage, spatial resolution, and direct access URLs.

Table 5. Complete dataset summary with access links. All datasets are publicly available at no cost.

Dataset	Source Agency	Temporal Coverage	Spatial Res.	Access URL
Canadian Drought Monitor	Agriculture and Agri-Food Canada (AAFC)	2011–present (weekly)	Provincial/sub-provincial	agriculture.canada.ca/atlas/apps/metrics
SPEI Global Drought Monitor	CSIC (Spain)	1901–present (monthly)	0.5° global grid	spei.csic.es/map/maps.html
North American Drought Monitor	NOAA/NCEI	2000–present (monthly)	Sub-national regions	ncei.noaa.gov/access/monitoring/nadm
Sentinel-2 Surface Reflectance	ESA/Copernicus	2015–present (5-day revisit)	10–60 m	dataspace.copernicus.eu
MODIS MOD13Q1 Vegetation Indices	NASA LP DAAC/USGS	2000–present (16-day)	250 m	lpdaac.usgs.gov/products/mod13q1v061
RADARSAT Constellation Mission	NRCan/EODMS	2019–present (12-day)	10–50 m	eodms-sgdot.nrcan-rncan.gc.ca
NASA SMAP L4 Soil Moisture	NSIDC DAAC	2015–present (3-hourly)	9 km	nsidc.org/data/smap
ERA5-Land Reanalysis	Copernicus CDS/ECMWF	1950–present (hourly)	9 km	cds.climate.copernicus.eu/datasets/reanalysis-era5-land
AAFC Annual Crop Inventory	Agriculture and Agri-Food Canada	2009–present (annual)	30 m	open.canada.ca/data/en/dataset/ba2645d5-4458-414d-b196-6303ac06c1bc

Dataset	Source Agency	Temporal Coverage	Spatial Res.	Access URL
Statistics Canada Crop Yields	Statistics Canada	1908–present (annual)	Census district	www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3210035901
SoilGrids v2.0	ISRIC/Wageningen University	Static	250 m	isric.org/explore/soilgrids
Canadian Digital Elevation Data (CDED)	Natural Resources Canada	Static	20–40 m	open.canada.ca/data/en/dataset/7f245e4d-76c2-4caa-951a-45d1d2051d86

Appendix B: Supplementary Methods

[The original manuscript skipped from Appendix A to Appendix C. The following supplementary methods section is provided as a placeholder for completeness; authors should populate it with relevant supplementary methodological details such as extended hyperparameter search spaces, additional preprocessing pipelines, or robustness checks.]

This appendix is reserved for supplementary methodological content including: (i) the full Bayesian hyperparameter search space and Optuna trial statistics; (ii) cross-validation protocol details including spatial-block and temporal-block splits used to prevent data leakage; (iii) data augmentation experiments; (iv) ablation studies on individual XAI layers; and (v) sensitivity analyses on key threshold parameters (e.g., EWI alert threshold, EFS cutoff).

Appendix C: Supplementary Figures and Tables

Table 6 provides complete model performance metrics by prediction horizon for spring wheat in 2021. Table 7 reports explanation fidelity scores (EFS) by drought severity zone and crop type.

Table 6. Complete model performance metrics by prediction horizon for spring wheat in 2021. DOY = day of year of prediction.

Prediction Horizon	RMSE (t ha ⁻¹)	MAE (t ha ⁻¹)	R ²	Bias (t ha ⁻¹)	MAPE (%)
DOY 193 (mid-July)	0.78	0.62	0.58	+0.28	18.4
DOY 213 (early Aug)	0.72	0.57	0.64	+0.22	16.8
DOY 233 (late Aug)	0.68	0.54	0.68	+0.19	15.6
DOY 253 (pre-harvest)	0.65	0.52	0.71	+0.18	14.9
DOY 263 (post-harvest)	0.62	0.49	0.74	+0.15	14.2

Table 7. Explanation fidelity scores (EFS) by drought severity zone and crop type for 2021 predictions.

Crop	D0	D1	D2	D3	D4
Spring wheat	0.96	0.95	0.93	0.91	0.89
Canola	0.95	0.94	0.92	0.90	0.88
Barley	0.95	0.94	0.92	0.90	0.87
Field peas	0.94	0.93	0.91	0.89	0.86




Figure 11. County-level predicted vs. observed spring wheat yield for (a) 2020 and (b) 2021. The 2021 panel shows increased scatter and systematic overprediction in D3–D4 counties (red/orange points).




Figure 12. Spatial distribution of spring wheat yield prediction residuals for (a) 2020 and (b) 2021. The 2021 map shows clusters of negative residuals (model overprediction) in southwestern Saskatchewan and southeastern Alberta, corresponding to D3–D4 drought zones.




Figure 13. Attention weight patterns for individual attention heads, showing specialization: Head 1 focuses on early-season conditions (DOY 130–160), Head 2 on mid-season (DOY 160–200), Head 3 on late-season (DOY 200–240), and Head 4 distributes attention across the full growing season.




Figure 14. SHAP dependence plots for the four most important features in 2021. ESI and SSM anomaly show strong negative slopes (lower values → lower predicted yield), while NDVI shows saturation at high values during drought. June Tmax shows a threshold effect above 30 °C.

Cross-method agreement – Variable Selection x SHAP x LIME (2021)




Figure 15. Rank correlation matrix between the three feature-importance-based XAI methods (Variable Selection, SHAP, LIME) for 2021 predictions. All pairwise correlations exceed 0.82, indicating strong cross-method agreement and robust explanations