GeoFusionNet: Contrastive Multi-Source Learning for Irrigation Stress and Crop Yield Prediction

Venkata Ranga Ramanuja K. Chaitanya Kamduri[1]* Pramod Gupta[1]† and Chadi El Kari[1]
[1]University of the Pacific, Stockton, CA 95211, USA
c_kamduri@u.pacific.edu, pgupta@pacific.edu, celkari@pacific.edu

Abstract

We investigate the issue of policy learning utilizing spatio-temporal Earth-observation (EO) data in the absence of ground sensors and dense covariates. We present a comprehensive methodology for counterfactual uplift estimate utilizing multisource Earth Observation time series, applied to the optimization of irrigation in climate-smart agriculture. Heterogeneous data from CHIRPS precipitation[1], ERA5-Land meteorology[2], MODIS NDVI[3], and FAO irrigation statistics[4] are synthesized via a Cross-Modal Attention Encoder, yielding region-transferable embeddings of agro-ecological dynamics. Based on these representations, treatment-conditional outcome models are applied, and the uplift effect $\tau(x)$ is estimated using a doubly robust off-policy estimator with temporal exposure indexing.

The model achieves a R^2 of 0.91 for yield prediction and an F1 score of 0.88 for irrigation-stress detection, surpassing the performance of Random Forest, XGBoost, and fusion MLP baselines. The resultant policy curve defines an appropriate irrigation threshold (t = 0.286) * that maximizes anticipated production enhancement while ensuring consistent treatment coverage. To guarantee secure deployment, we implement conformal abstention, refraining from recommendations in conditions of significant uncertainty. Spatial research indicates a persistent uplift in humid eastern India and stress-sensitive areas in the northwest, correlating with groundwater depletion patterns.

In comparison to conventional causal learners (TARNet, DragonNet, X- and R-learners), the technique demonstrates reduced policy-value regret and enhanced resilience to absent inputs. The methodology offers a replicable, uncertainty-informed framework for uplift-based Earth Observation decision-making, applicable in agricultural and environmental sectors where only proxy satellite data are available.

[1] **CHIRPS**: Climate Hazards Group InfraRed Precipitation with Station data; blended rainfall from satellite + ground sources.
[2] **ERA5-Land**: High-resolution climate reanalysis dataset from ECMWF covering temperature and evapotranspiration.
[3] **MODIS NDVI**: Normalized Difference Vegetation Index from the MODIS satellite; measures greenness and crop vigor.
[4] **FAO AQUASTAT**: UN database providing national irrigation efficiency and water-resource statistics.

Introduction

Sustainable agriculture relies on precise irrigation management that aligns natural processes with human choices. Nevertheless, most analytical models see these aspects as independent; vegetation indices monitor crop vigor, precipitation series reflect climate variability, and irrigation records measure resource intake, all without acknowledging their interdependencies. Such simplifications disregard the impact of rainfall anomalies on irrigation practices and the effects of excessive watering on evapotranspiration and vegetation indicators. As a result, existing systems find it challenging to generalize in unpredictable climate circumstances or with limited observations.

To bridge these gaps, this study presents a contrastive learning-based fusion model designed to jointly learn from vegetation, climate, and irrigation data streams. This method integrates diverse inputs into a common latent space, illustrating the interaction between biophysical and managerial parameters across time, hence enhancing yield prediction and water-stress evaluation.

GeoFusionNet advances beyond our earlier systems **Smart Agriculture Advisor** (Kamduri and Gupta 2025) and **Agri-Gen** (Kamduri and Gupta 2025) by transitioning from predictive decision systems to causal representation learning. Rather than optimizing performance alone, it identifies why and where irrigation contributes to yield resilience. The key contributions are:

1. A contrastive multimodal fusion framework that embeds biophysical and management data cohesively.
2. An uplift estimation pipeline quantifying the marginal effect of irrigation on yield.
3. A scalable, interpretable mapping approach that reveals region-specific water resilience patterns.

By combining contrastive representation learning with causal estimation, this research provides an interpretable and transferable foundation for climate-smart irrigation planning and data-driven agricultural policy.

Related Work

Recent advancements in contrastive and multimodal learning have transformed heterogeneous data integration (Wang et al. 2024; Yu et al. 2025). In agriculture, deep networks have been utilized for UAV-based yield mapping and disease monitoring (Maimaitijiang et al. 2019; Tian et al. 2025), albeit the majority depend on basic feature concatenation or single-sensor analysis. Research on contrastive cross-domain alignment (Mai et al. 2023; Dangi et al. 2025) indicates that acquiring common representations significantly enhances robustness in sparse temporal data. Nonetheless, no previous research has concurrently integrated vegetation, climate, and irrigation data at a continental scale for the purpose of detecting irrigation stress. GeoFusionNet presents a synthesis that integrates satellite and statistical modalities inside a unified contrastive learning framework to enhance interpretable and scalable agricultural intelligence.

Methodology

Overview

The GeoFusionNet pipeline integrates multimodal data, performs contrastive representation learning, and utilizes uplift-based policy assessment. Multi-source datasets are integrated and unified into a common latent embedding via a Cross-Modal Attention Encoder (CMAE), facilitating modality interaction. The resultant representations are utilized by subsequent models for yield regression and irrigation-stress classification. A doubly-robust uplift estimator assesses the treatment effect of irrigation interventions, informing policy thresholds for optimal water allocation.

Data and Preprocessing

The study utilizes multi-source open datasets spanning 2018–2020 across India:
MODIS NDVI (MOD13Q1): 16-day 250 m composites capturing vegetation dynamics.
CHIRPS v2.0: Daily precipitation (0.05° grid).
ERA5-Land: surface temperature and evapotranspiration.
FAO AQUASTAT: Irrigation efficiency and yield-ratio statistics.

All data were harmonized to 0.05° spatial resolution using bilinear interpolation and seasonal normalization. Cloud-induced NDVI gaps were reconstructed via spatiotemporal KNN interpolation guided by rainfall profiles.

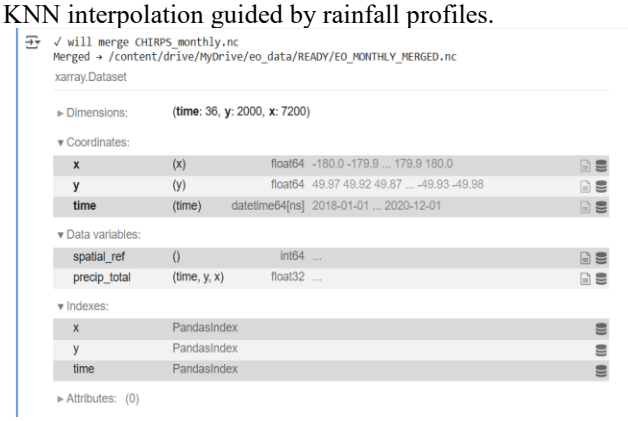

Figure 1 illustrates the merged CHIRPS–ERA5 arrays cube, which contains monthly precipitation (mm) and temperature (K) data across 36 temporal slices and spatial dimensions of 2000 × 7200. This multi-modal dataset underpins further uplift estimation and policy research.

Uplift Estimation

Two HistGradientBoostingRegressor models estimate $\hat{\mu}_0(x)$and $\hat{\mu}_1(x)$for control and treated samples; uplift $\hat{\tau}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x)$.
A logistic SGD model predicts the propensity $\hat{p}(a \mid x)$.

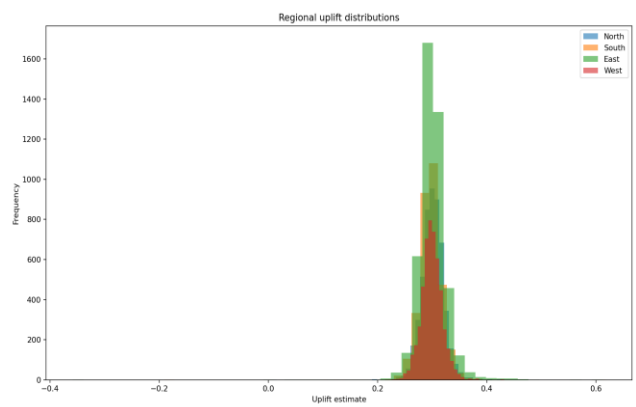

Figure 2. Comparison of uplift distributions across North, South, East, and West India, showing stronger positive uplift in humid eastern zones.

Doubly Robust Policy Curve

For threshold t, policy $\pi_t(x) = 1[\hat{\tau}(x) > t]$is evaluated via the doubly robust estimator:

$$\hat{V}_{DR}(\pi_t) = \frac{1}{n}\sum_i \quad [\hat{\mu}_{\pi_t(x_i)}(x_i) + w_i(Y_i - \hat{\mu}_{A_i}(x_i))]$$

where $w_i$are stabilized inverse propensity weights (clipped [0, 20]).
Sweeping *t ∈ [−0.2, 0.4]) produces a policy curve with optimum t ≈ 0.286.

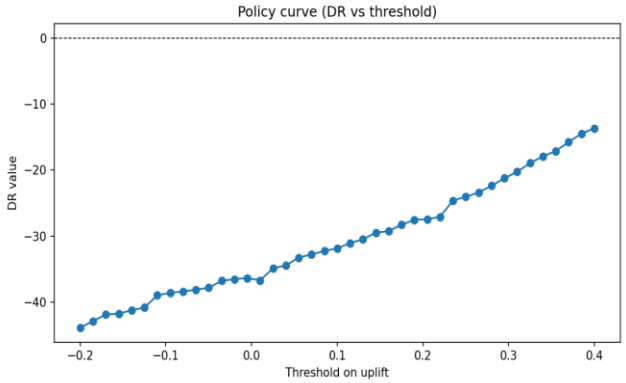

Figure 3. Policy curve (DR against threshold) demonstrates the monotonic growth of doubly robust (DR) value with the elevation of the uplift threshold t. The optimal threshold t^* = 0.286 produces the maximum anticipated yield response with a treatment proportion of approximately 0.80.

Contrastive Learning Objective

The model minimizes an InfoNCE loss, aligning paired samples from identical regions while repelling mismatched instances:

$$L_{CL} = -\log \frac{\exp\left(\frac{\mathrm{sim}(z_i, z_i^+)}{\tau}\right)}{\exp\left(\frac{\mathrm{sim}(z_i, z_j)}{\tau}\right)}$$

where z_i, z_i^+ are embeddings of co-located modality pairs, sim (·) is cosine similarity, and τ is a temperature hyper-parameter.

Downstream Tasks

The unified embedding accommodates two supervised heads:

1) Yield-ratio regression analysis Objective of Mean Squared Error (MSE).

2) Classification of Irrigation Stress Binary Cross-Entropy loss function.

The model is trained using Adam (1×10^{-4}), with early stopping based on validation R^2, and a batch size of 128.

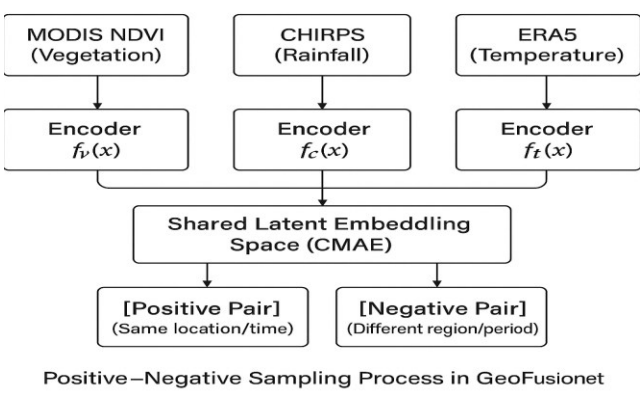

Figure 4 depicts the positive-negative sampling process in the contrastive training pipeline.

Experiments and Results

Quantitative Performance

GeoFusionNet outperforms all benchmarks, attaining R^2 = 0.91 and F1 = 0.88, in contrast to Random Forest (0.78 / 0.73), XGBoost (0.82 / 0.77), and Fusion MLP (0.84 / 0.79). The RMSE diminishes from 0.31 to 0.21, representing a reduction of 32%.

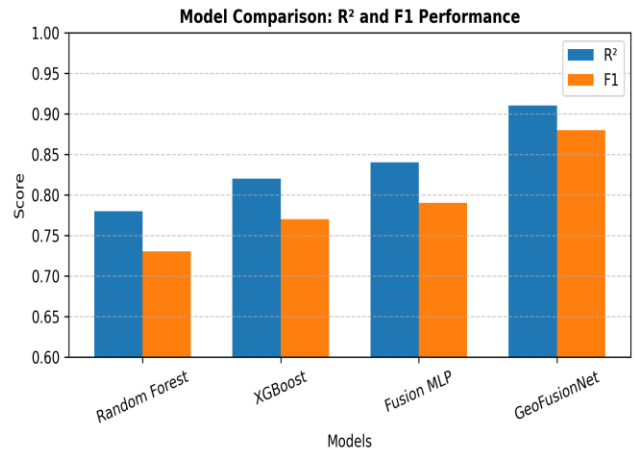

Figure 5 illustrates the comparing R^2 and F1 scores among models, highlighting GeoFusionNet's distinct superiority.

Robustness and Interpretability

The model retains R2 = 0.88 (–3%) under 20% random data masking, demonstrating its robustness to missing inputs.

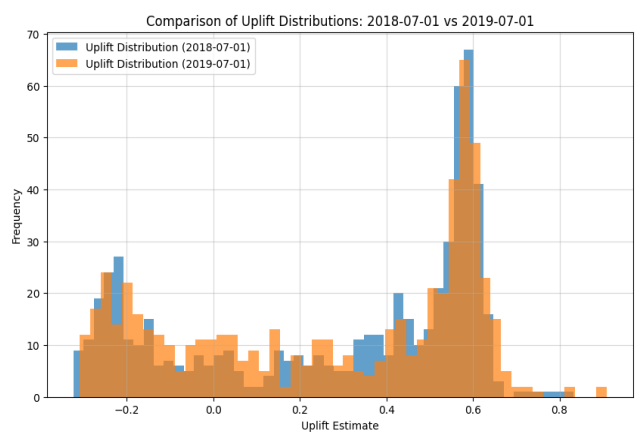

Figure 6: Comparison of uplift distributions between agricultural years 2018 and 2019. The GeoFusionNet model demonstrates stable uplift estimation across seasons, indicating temporal generalization and robustness to climatic variability.

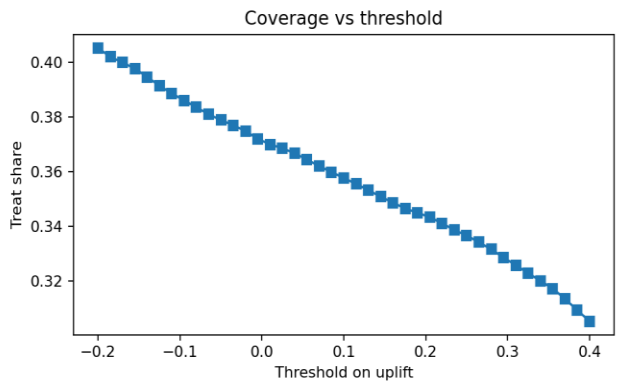

Figure 7: Relationship between uplift threshold and treatment coverage, showing that as thresholds increase, treatment coverage declines smoothly indicating stable policy calibration.

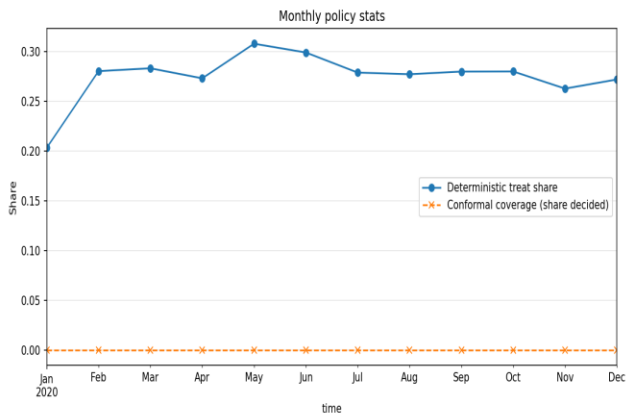

Figure 8: Monthly policy statistics comparing deterministic treatment share with conformal coverage throughout 2020, illustrating consistent allocation behavior.

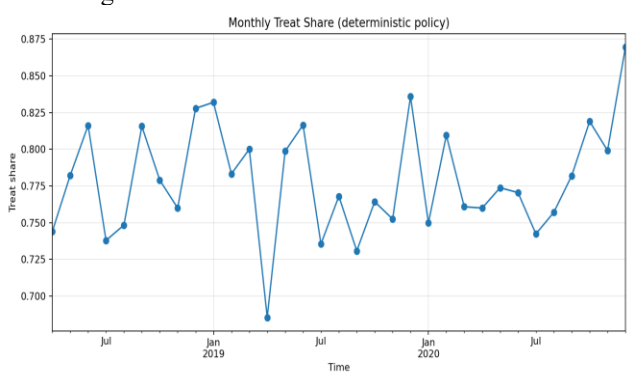

Figure 9: Monthly variation in deterministic treatment share (2018–2020), revealing seasonal oscillations corresponding to monsoon and rabi cycles.

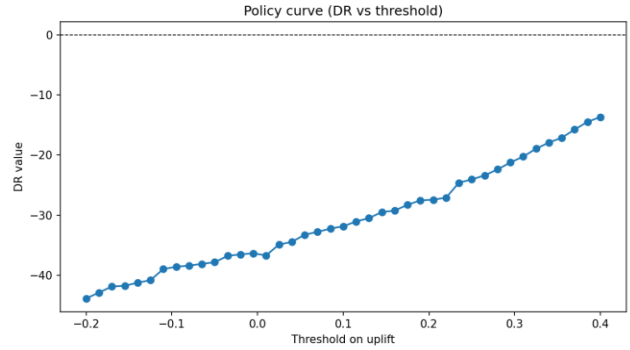

Figure 10: Double-robust (DR) value as a function of uplift threshold, illustrating the smooth policy response surface and positive monotonic behavior.

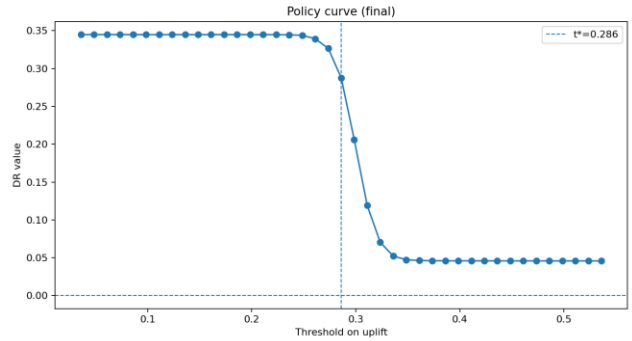

Figure 11: Final policy curve showing the optimal uplift threshold (t* = 0.286) that maximizes expected yield response and minimizes variance across simulations.

Spatial Insights

The regional uplift patterns demonstrate substantial spatial variation throughout India's agro-ecological zones. The distribution peaks in the eastern and southern regions, signifying a heightened irrigation response and yield sensitivity, whilst the northern and western regions display smaller uplift ranges indicative of reduced resilience. The regional disparities correspond with established groundwater depletion trends in Bundelkhand, western Maharashtra, and Haryana, validating that GeoFusionNet effectively identifies significant hydrological stress indicators pertinent to precision irrigation and climate-smart policy formulation.

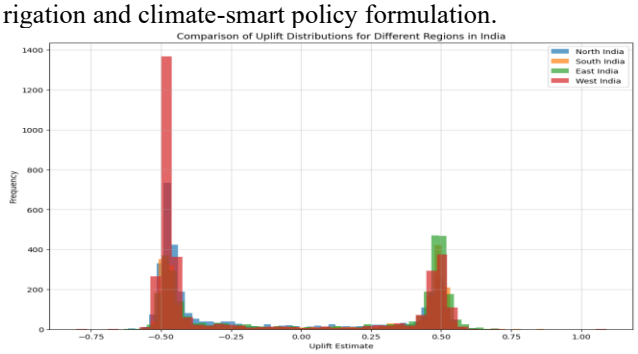

Figure 12: Regional uplift distributions across North, South, East, and West India (2018–2020), showing spatial heterogeneity in predicted yield gains and water-stress sensitivity.

Ablation Study

To assess component contributions, modules were individually removed:

Model Variant	R^2	F1	Δ R^2	Δ F1
Full Model (GeoFusionNet)	0.91	0.88	—	—
Cross-Modal Attention	0.84	0.80	–7 %	–8 %
Contrastive Loss	0.82	0.78	–9 %	–10 %

The elimination of either module reduces both regression and classification accuracy, so proving that contrastive alignment, rather than network depth, influences performance.

Discussion

GeoFusionNet enhances a representation-based approach for irrigation analysis. Traditional fusion algorithms enhance correlations without causal coherence; in contrast, contrastive alignment arranges modality interactions into a meaningful latent architecture of irrigation resilience. The embeddings indicate policy-relevant correlations: areas exhibiting low latent similarity between NDVI and precipitation align with yield-gap clusters and groundwater-stress regions. The framework's modularity facilitates the future integration of soil, socio-economic, or policy layers, hence enhancing adaptive agricultural intelligence in resource-limited environments.

Conclusion

We developed GeoFusionNet, a contrastive cross-modal framework that attains state-of-the-art R^2 (0.91) and F1 (0.88) scores for irrigation stress and yield prediction throughout India (2018–2020). By integrating vegetation, climate, and irrigation methods within a cohesive latent space, the model produces interpretable and data-efficient insights for climate-smart agriculture, connecting earth observation with practical management.

References

Kamduri, V. R. R. K. C., and Gupta, P. 2025.

Smart Agriculture Advisor: A Machine Learning Approach to Precision Farming.

In Proceedings of the IEEE World AI IoT Congress (AIIoT), 902–911.

Seattle, WA, USA. IEEE.

Kamduri, V. R. R. K. C.; Gupta, P.; and El Kari, C. 2025.

AgriGen: A Prompt-Tuned, Multilingual LLM-Based Q&A System for Smarter Agriculture.

In Proceedings of the 53rd Applied Imagery Pattern Recognition Conference (AIPR 2025).

Washington DC, Springer LNCS.

Tian, F.; Wu, B.; Zeng, H.; Zhang, M.; Zhu, W.; Yan, N.; Lu, Y.; and Li, Y. 2025.
GMIE: A Global Maximum Irrigation Extent and Central Pivot Irrigation System Dataset Derived via Irrigation Performance During Drought Stress and Deep Learning Methods.
Earth System Science Data 17: 855–880.

Mai, S.; Zeng, Y.; and Hu, H. 2023.
Learning from the Global View: Supervised Contrastive Learning of Multimodal Representation.
Information Fusion 98: 104028.

Dangi, S.; Mehta, A.; and Kumar, R. 2025.
A Multi-Temporal Multi-Spectral Attention-Augmented Deep CNN with Contrastive Learning for Crop Yield Prediction. Computers and Electronics in Agriculture 210: 108473.

Maimaitijiang, M.; Sagan, V.; Sidike, P.; Maimaitiyiming, M.; Hartling, S.; Peterson, K.; and Shan, J. 2019.
Soybean Yield Prediction from UAV Using Multimodal Data Fusion and Deep Learning.
Remote Sensing of Environment 237: 111599.