Науковий вісник НЛТУ України
Scientific Bulletin of UNFU
https://doi.org/10.36930/40360308
Article received: 31.03.2026
Article accepted: 01.06.2026
Article published: 29.06.2026
Correspondence author: S. I. Baluk svyatoslavbaluk@gmail.com
ISSN 1994-7836 (print)
ISSN 2519-2477 (online)

UDC 004.89:528.8

S. I. Baluk, V. Ya. Antoniv

Lviv Polytechnic National University, Lviv, Ukraine

SATELLITE DATA RECONSTRUCTION UNDER CLOUD COVER FOR CORN YIELD FORECASTING VIA MULTIMODAL FUSION

A comprehensive intelligent system for corn yield prediction based on the synergy of a Spatio-Temporal Generative Adversarial Network (ST-cGAN) and a recurrent CNN-LSTM architecture has been developed and tested. A multimodal fusion of weather-independent Sentinel-1 radar data, ERA5-Land meteorological factors, and historical Sentinel-2 optical observations was applied. The problem of "biochemical blindness" in radar signals, where radar captures physical plant structure but fails to detect photosynthetic activity, and cloud cover limitations in optical remote sensing was successfully resolved. To achieve this, an early fusion strategy was implemented. The ST-cGAN simultaneously processes current SAR data for the structural macro-state, historical optical data for the last known biochemical baseline, and a 10-day history of meteorological priors to enforce physiological constraints. Consequently, if radar indicates high biomass but meteorological data reveals severe drought, the neural network mathematically recognizes biological stress and synthesizes proportionally depressed NDVI and NDRE indices, preventing the hallucination of falsely healthy crops. Dynamic synthesis of missing NDVI and NDRE vegetation indices was conducted under prolonged continuous cloud cover (up to 21 days), achieving a high structural similarity index (SSIM = 0.89). Temporal discriminators for frame sequence analysis were introduced into the architecture, improving the edge-preserving index by 18 % and minimizing spatial artifacts at field boundaries. Pixel-level regression was performed for a 15,000-hectare test area in the Western Forest-Steppe of Ukraine based on reconstructed time series covering 15 critical phenological stages. It was established that the proposed architecture reduces the root mean square error (RMSE) to 0.48 t/ha with a coefficient of determination (R^2) of 0.90. Statistical analysis proved the significant superiority of the developed method over industry-standard gap-filling algorithms (e.g., STARFM). A scalable decision-support system for optimizing harvest logistics and financial planning under any atmospheric conditions was presented. Future research directions involving neural network knowledge distillation for IoT devices and UAV data integration were outlined.

Keywords: deep learning; artificial intelligence in agriculture; spatiotemporal networks; precision farming; predictive modeling.

Introduction / Вступ

In the study [12], the authors emphasize that ensuring food security in the face of global climate variability requires reliable agricultural monitoring, indicating that traditional yield forecasting methods are gradually being replaced by Earth remote sensing (ERS) technologies combined with deep learning algorithms. Furthermore, the research [16] demonstrates that the analysis of multispectral satellite time series, such as Sentinel-2, allows for tracking the dynamics of vegetation indices, specifically the Normalized Difference Vegetation Index (NDVI) and the Normalized Difference Red Edge (NDRE), to accurately identify plant responses to environmental conditions throughout the growing season.

In the paper [19], the authors point out a critical limitation of optical remote sensing, noting its absolute dependence on atmospheric conditions, where heavy cloud cover during the spring and summer leads to significant gaps in the time series. The study highlights that the absence of data during highly nonlinear phenological windows-such as the transition from the vegetative tasseling phase (VT) to the reproductive silking and grain-filling phases (R1-R3) in corn-makes continuous monitoring impossible and severely degrades the predictive power of agronomic models.

To address this limitation, the authors of [18] propose the use of synthetic aperture radar (SAR) satellites, such as Sentinel-1, demonstrating that C-band radio waves penetrate cloud cover to provide consistent, weather-independent observations of biomass macrostructure, canopy volume, and soil moisture. However, this approach presents a fundamental "biochemical gap" because radar backscatter primarily interacts with physical structure (stalks and leaves) and does not contain direct information about plant photosynthetic activity, nitrogen uptake, or early-stage foliar

Балук Святослав Ігорович, асистент, кафедра автоматизованих систем управління.
Email: svyatoslavbaluk@gmail.com; https://orcid.org/0009-0007-5999-4547
Антонів Володимир Ярославович, канд. техн. наук, ст. викладач, кафедра автоматизованих систем управління.
Email: volodya.antoniv@gmail.com; https://orcid.org/0000-0002-4544-4612
Цитування за ДСТУ: Балук С. І., Антонів В. Я. Реконструкція супутникових даних за умов хмарності для прогнозування врожайності кукурудзи злиттям мультимодальних знімків. Науковий вісник НЛТУ України. 2026, т. 36, № 3. С. 79–85.
Citation APA: Baluk, S. I., & Antoniv, V. Ya. (2026). Satellite data reconstruction under cloud cover for corn yield forecasting via multimodal fusion. *Scientific Bulletin of UNFU*, 36(3), 79–85. https://doi.org/10.36930/40360308

diseases that deplete chlorophyll before structural collapse occurs.

In recent works [11, 21], the researchers successfully utilized conditional generative adversarial networks (cGANs) for the nonlinear image translation of missing optical data directly from SAR observations. Conversely, the study [22] reveals a fatal flaw in applying these methods to yield forecasting naively; mapping structural SAR deterministically to biochemical NDVI can result in hallucinated healthy crops during periods of severe biological stress (e.g., asymptomatic pathogens or flash droughts), where the physical macrostructure remains intact but chlorophyll is depleted.

Object of research – process of reconstructing gaps in multispectral satellite time series for continuous agricultural applications.

Subject of research – deep learning methods and tools used for multimodal data fusion to synthesize missing information.

The purpose of the research – develop a comprehensive multimodal spatiotemporal framework that fuses synthetic aperture radar, historical optical data, and meteorological priors to overcome the biological blindness of radar signals, thereby enabling continuous, gap-free reconstruction of vegetation indices for high-precision corn yield forecasting under persistent atmospheric occlusions.

To achieve this purpose, the following main *research objectives are identified*:

- to analyze existing methods, algorithms, and limitations regarding missing data reconstruction and crop yield forecasting in remote sensing time series, which will allow for identifying the fundamental flaws in current deterministic gap-filling techniques and justifying the necessity of a physics-aware generative approach;
- to develop the architecture of a multimodal data-fusion framework that integrates weather-independent synthetic aperture radar (Sentinel-1) backscatter, meteorological priors (ERA5-Land), and historical optical states (Sentinel-2), which will establish a robust structural foundation for bridging the "biochemical gap" of radar signals and preventing the hallucination of healthy crops during biological stress events;
- to implement and train the proposed framework, comprising an ST-cGAN to synthesize missing vegetation indices (NDVI and NDRE), and a subsequent CNN-LSTM network to perform pixel-level yield regression, which will enable the transformation of discontinuous, cloud-obscured satellite imagery into continuous, weather-independent spatiotemporal features for capturing nonlinear crop dynamics;
- to verify and quantitatively evaluate the structural accuracy of the synthesized data and the final yield predictions against standard gap-filling baselines using rigorous spatial block-cross-validation, which will confirm the model's generalizability, eliminate the risk of geographic data leakage, and statistically prove the framework's agronomic superiority over existing state-of-the-art interpolation models.

Analysis of recent research and publications. The problem of high-precision monitoring of agricultural land and yield forecasting is being actively researched. In the studies [12, 23], the authors demonstrate that deep neural networks significantly outperform traditional regression methods by mapping complex nonlinear environmental interactions. Furthermore, in the paper [16], it is emphasized that the accuracy of these models relies on analyzing continuous, dense multispectral time series. However, as noted in [19], achieving this continuity is hindered by cloud cover, which necessitates multimodal data fusion.

To address atmospheric occlusions, in the study [7], the authors utilize baseline data fusion algorithms, such as the Spatiotemporal Adaptive Reflectance Fusion Model (STARFM), which nevertheless struggle to extrapolate abrupt phenological changes. As demonstrated in [22], these linear assumptions break down during rapid-onset drought, leading to severe underestimations of crop stress. Consequently, in recent works [2, 5, 8, 11, 21], researchers have adapted conditional generative adversarial networks (cGANs) to mathematically translate radar signals into optical domains, using weather-independent SAR backscatter as an auxiliary condition.

Despite these advancements, in the research [21], the authors highlight structural limitations, noting that treating image synthesis as an atemporal objective leads to "flicker" and temporal instability in the reconstructed series. Moreover, as proven in [22], relying purely on structural radar to generate biochemical indices overlooks physiological complexity, often synthesizing false healthy signatures. While the authors of [14] offer continuous daily meteorological variables via ERA5-Land datasets, the direct integration of these priors into generative models remains unresolved. Additionally, as emphasized in [13], evaluating these models requires strict spatial block-cross-validation to prevent geographic overfitting.

Therefore, bridging the SAR "biochemical gap" by integrating meteorological priors into the generative process and subsequently validating these continuous series through spatiotemporal regression represents the primary scientific novelty of this study, as such an integrated framework has not been proposed in the analyzed literature.

Despite the active development of remote sensing data fusion, the analyzed literature shows a critical lack of solutions capable of preventing these deterministic generative errors under complex meteorological stresses. None of the existing studies have proposed a fully integrated framework that simultaneously fuses SAR backscatter, historical optical states, and meteorological priors to guide the generative process. Therefore, there is an urgent need to develop a physics-aware imputation methodology that overcomes the "biochemical gap" of radar signals and generates reliable, gap-free time series for accurate crop yield forecasting, which constitutes the primary novelty and necessity of this research.

Materials and methods of research. The study investigated high-precision corn yield forecasting across 15,000 hectares in Ukraine's Western Forest-Steppe (2021-2023) using a multimodal deep learning framework. The methodology relies on fusing Sentinel-1 radar backscatter, Sentinel-2 multispectral imagery, and ERA5-Land meteorological priors, validated against ground-truth GPS yield monitor data. To reconstruct missing optical data and predict final yields, an integrated Spatio-Temporal Generative Adversarial Network (ST-cGAN) and CNN-LSTM architecture was employed. Model evaluation utilized strict spatial block-cross-validation to ensure generalizability and prevent geographic data leakage.

Research results and their discussion / Результати дослідження та їх обговорення

Proposed framework for reconstructing cloud-free time series for crop yield forecasting. In this article, a multistage architecture was designed and implemented to bridge the gap between raw multimodal satellite observations and quantitative agronomic outcomes. This architectural implementation consists of: the aggregation and spatiotemporal preprocessing of multimodal Sentinel-1, Sentinel-2, and ERA5-Land data; a generative component (ST-cGAN) to translate these inputs into synthesized vegetation indices under persistent cloud cover; and a predictive component for the integration of reconstructed time series into a CNN-LSTM regression model. The structural schematic and operational logic of this integrated framework are illustrated in Figure 1.

The dataset generation and input layer is responsible for aggregating heterogeneous raw data (radar, optical, and meteorological) and establishing the experimental configuration. Its primary purpose is to enforce strict spatial block-cross-validation and simulate realistic contiguous cloud masking to prevent data leakage and ensure a robust, unbiased model evaluation.

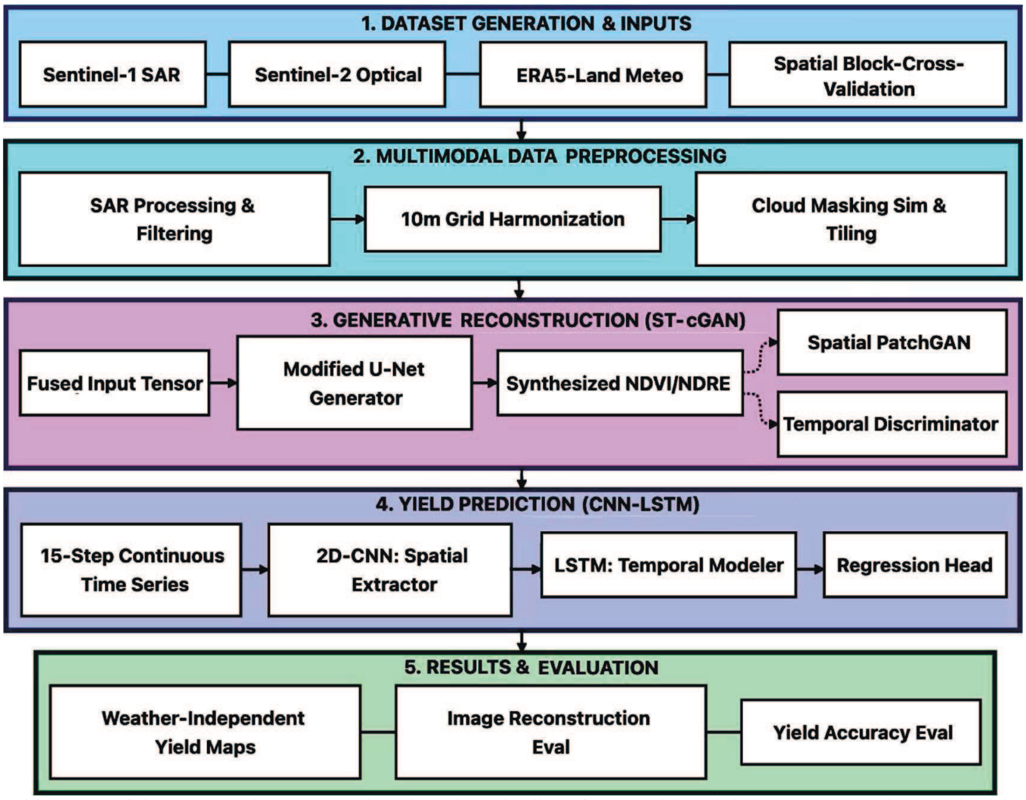

Fig. 1. Overall architecture of the proposed two-stage deep learning pipeline, from multimodal data preprocessing to pixel-level yield prediction / Загальна архітектура запропонованого двоступеневого конвеєра глибокого навчання: від попереднього оброблення мультимодальних даних до прогнозування виходу на рівні пікселів

The multimodal data preprocessing layer serves to harmonize the distinct data sources into a unified spatiotemporal grid. Its purpose is to ensure that asynchronous satellite passes and macro-climate weather priors are geometrically registered, cleaned of noise, and standardized into uniform formats for deep learning processing.

The generative reconstruction layer (*ST-cGAN*) acts as the core imputation engine. Its purpose is to bridge the "biochemical gap" of radar data by fusing structural SAR with meteorological history. This ensures the synthesis of biologically honest vegetation indices during persistent cloud cover, preventing the model from hallucinating healthy crops during periods of environmental stress

The crop yield prediction layer (*CNN-LSTM*) is designed to translate the gap-free, reconstructed time series into quantitative harvest forecasts. Its purpose is to extract spatial field anomalies and model the sequential, long-term dependencies of the crop's phenological development from vegetative growth through senescence.

The result and evaluation layer generates the final pixel-level yield forecasts and weather-independent heatmaps. Its purpose is to output actionable, scientifically validated agronomic data that can be utilized as a reliable decision-support tool for precision agriculture regardless of atmospheric interference.

Multimodal Data Preprocessing. The successful implementation of the proposed framework relies on a robust preprocessing pipeline to ensure the precise calibration and geometric alignment of heterogeneous data streams. To achieve this, three distinct data sources were harmonized into a unified spatiotemporal grid:

1) SAR (Sentinel-1): IW mode, VV and VH polarizations. Preprocessing included thermal noise removal, radiometric calibration to backscatter coefficient (σ^0), and Refined Lee Filtering (7×7 window) to mitigate speckle noise. Agronomically, VV polarization is highly sensitive to vertical stalk geometry and soil moisture, while VH polarization effectively captures volume scattering within the upper corn canopy.
2) Optical (Sentinel-2): Level-2A. To address spectral simplification, the target variables were expanded to a 2-channel tensor: NDVI (using NIR and Red) and NDRE (using NIR and Red-Edge, Band 5). NDRE is specifically utilized because NDVI is prone to saturation during peak biomass (the VT phase); Red-Edge bands penetrate deeper into the canopy and remain sensitive to late-season chlorophyll degradation.
3) Meteorological Priors (ERA5-Land): Daily cumulative precipitation, 2 m mean temperature, and surface solar radiation downwards were extracted from the state-of-the-art reanalysis dataset [14]. These 9km-resolution variables were downscaled and resampled to match the 10 m grid, providing an early-fusion weather context vector for each pixel.

A visual comparison of these harmonized radar and multispectral inputs across the study area is provided in Figure 2. Following the harmonization of these data sources, the preprocessing pipeline established rigorous experimental conditions for the model. To robustly test the generative capabilities of the framework, artificial data gaps were not introduced as randomly distributed pixels. Instead, occlusions were applied as historically contiguous temporal blocks – such as masking an entire field for 21 consecutive days during the critical silking phase – to accurately simulate the effects of persistent weather systems. Finally, to prepare the data for neural network ingestion, the aggregated multidimensional tensors were standardized and tiled into 12,000 blocks of 256×256 pixels.

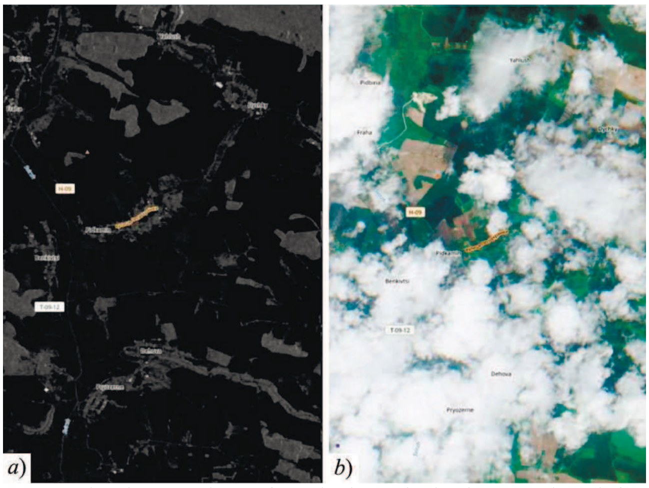

Fig. 2. Visual comparison of Sentinel-1 radar backscatter data (left) and Sentinel-2 multispectral data (right) for the study area / Візуальне порівняння даних про відбиття від земної поверхні, отриманих за допомогою радара Sentinel-1 (a), та мультиспектральних даних Sentinel-2 (b) для досліджуваної території

Spatio-Temporal GAN (ST-cGAN) Architecture. To overcome the limitations of standard generative models and solve the "asymptomatic pathogen" paradox, where a plant's physical macrostructure remains intact despite biochemical failure the basic conditional GAN was upgraded to a fully spatiotemporal architecture. This configuration consists of optimized generator and discriminator networks designed to enforce both physical accuracy and phenological continuity:

- *Generator* (G)*:* A modified U-Net architecture. Instead of receiving only current SAR, the input tensor performs an early fusion of three modalities: SAR_t, $ERA5_t$, $Optical_{t-k}$. Early fusion strategies have been shown to maximize the extraction of complementary features from asynchronous satellite optical and radar time series [10]. Here, $ERA5_{\ t}$ provides the 10-day meteorological history (capturing recent drought or flooding), and $Optical_{t-k}$ is the last known clear vegetation map. This grounds the generation in biological reality – the meteorological prior indicates a severe 14-day drought, the U-Net bottleneck learns to suppress the output NDVI, even if the SAR VV backscatter remains high. The generator outputs a 2-channel tensor: $NDVI_t$, $NDRE_t$.
- *Discriminator* (D)*:* A dual-discriminator setup was employed to enforce both spatial and temporal realism. A Spatial PatchGAN (70×70) evaluates high-frequency textures and field boundaries. Additionally, a Temporal Discriminator (utilizing 3D convolutional layers) evaluates sequences of three consecutive dates (t – 1, t, t + 1). It classifies whether the transition between the generated frame and its adjacent real frames represents a biologically plausible growth curve, penalizing phenologically inconsistent "flickering."

Loss Function Formulation. The total objective function combines adversarial loss, L1 pixel-matching loss, and a temporal consistency penalty. The temporal loss is defined mathematically to minimize the difference in temporal gradients between the generated and real sequences:

$$L_{Temp}\left(G\right) = E_{x,t}\left[\left|\nabla_t G_{(Xt)} - \nabla_t y_t\right|_2\right]. \qquad (1)$$

Where, $L_{Temp}(G)$ is the temporal consistency loss function designed to penalize phenological flickering; $E_{x,t}$ is the expected value over the spatial dimension x and time step t; ∇_t is the represents the temporal gradient operator (the mathematical difference between consecutive time steps); $G_{(Xt)}$ is the generated vegetation image at time step t; y_t is the real ground-truth vegetation image at time step t; $\| \|_2$ denotes the $L2$ norm (Euclidean distance).

The total objective is formulated as:

$$G^* = apg\,\min_G \max_D L_{cGAN}\left(G,D\right) + \lambda_S L_{L1}\left(G\right) + \lambda_T L_{Temp}. \qquad (2)$$

Where, G^* represents the optimal weights of the Generator network; *arg* min_G max_D denotes the minimax game objective (the Generator minimizes the loss while the Discriminator maximizes it); $L_{cGAN}(G,D)$ is the conditional adversarial loss function; $L_{L1}(G)$ is the $L1$ pixel-matching regularization (Mean Absolute Error) ensuring spatial fidelity; λ_S is the hyperparameter controlling the weight of the spatial $L1$ loss (experimentally set to 50); λ_T is the hyperparameter controlling the weight of the temporal consistency loss (experimentally set to 10).

Crop Yield Prediction Pipeline (CNN-LSTM). To predict the final harvest volume, the reconstructed continuous time series comprising 15 time steps from May to September are fed into a sequential CNN-LSTM.

The spatial feature extraction is performed by a CNN head consisting of two Conv2D layers (64 and 128 filters, 3×3 kernel) to extract spatial heterogeneity, identifying localized stress hotspots and micro-relief anomalies. Batch Normalization and MaxPooling are applied after each convolution.

The sequential modeling is handled by an LSTM block, where the flattened spatial feature maps pass through a 128-dimensional recurrent layer with a sequence length of 15. The extraction of long-term temporal dependencies is critical for modeling phenological phases accurately from satellite image time series [15]. The recurrent layer learns that a depressed NDVI during July (flowering) carries a higher weight for yield penalty than the same NDVI value in September (senescence). The effectiveness of extracting temporal features from asynchronous Sentinel-1 and Sentinel-2 sequences using recurrent layers has been well-documented in recent multi-source architectures [5].

The network optimization strategy applied a Dropout rate of 0.3 to the fully connected regression head to prevent overfitting. The model was trained using the Adam optimizer (learning rate = 0.001, batch size = 32) with a Mean Squared Error (MSE) loss function over 150 epochs, utilizing early stopping based on the validation set.

Evaluation of Image Reconstruction. This section presents the results of an empirical evaluation of the proposed method. The analysis is divided into two logical stages corresponding to the overall structure of the study: a quantitative assessment of the radar image translation quality, and an evaluation of the final predictive model compared to traditional approaches (Ablation Study).

To quantify "hallucinations" at sharp field boundaries – a common flaw in standard U-Nets that rely heavily on convolution padding, we measured the Edge-Preserving Index (EPI). The integration of the Temporal Discriminator and historical optical priors improved the EPI by 18 % compared to a baseline spatial-only cGAN. This ensures that field boundaries, rural roads, and sharp stress hotspots were not blurred into spatial averages, preserving the pixel-level integrity required for precision agriculture.

Furthermore, during a specific simulation involving a 21-day cloud block coinciding with a localized drought period, standard interpolation methods predicted a steady green-up. Conversely, the ST-cGAN, reacting to the high temperatures and low precipitation inputs from the ERA5-Land tensor, successfully synthesized a plateauing and subsequent drop in the NDRE index, closely mirroring the ground-truth reference data (Figure 3).

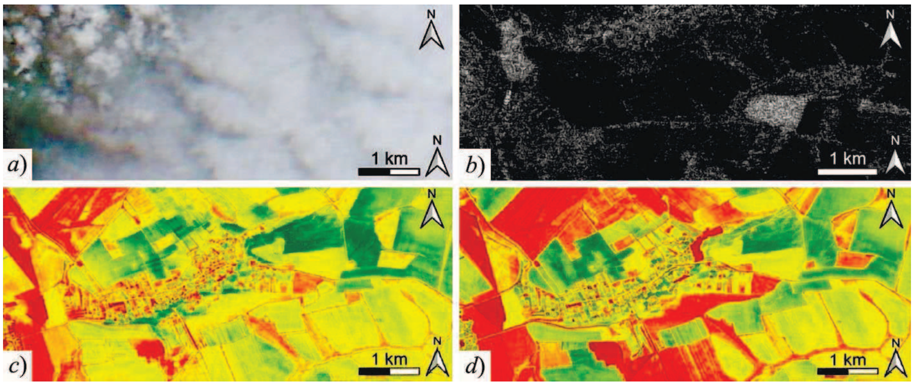

Fig. 3. Visualization of multispectral data reconstruction results using the proposed ST-cGAN architecture / Візуалізація результатів реконструкції мультиспектральних даних за допомогою запропонованої архітектури ST-cGAN: (a) Input cloud-covered optical image (Sentinel-2) / вхідне оптичне зображення з хмарним покривом (Sentinel-2); (b) Input radar backscatter image (Sentinel-1 SAR) / вхідне зображення радіолокаційного відбиття (Sentinel-1 SAR); (c) Synthesized cloud-free NDVI vegetation map (generator output) / синтезована карта рослинності за індексом NDVI без хмар (вихідні дані генератора); (d) Reference real cloud-free NDVI map (Ground Truth) / еталонна реальна карта рослинності за індексом NDVI без хмар (Ground Truth)

This visual and quantitative alignment confirms that the generative module effectively resolves the structural data gaps caused by cloud cover, ensuring that the downstream regression model receives biologically consistent inputs for accurate yield estimation.

Accuracy of Yield Prediction. The regression accuracy on the spatially isolated District C testing set served as the primary validation metric. The proposed pipeline was compared against established baselines to isolate the contribution of the generative fusion step (Table).

Table. Pixel-level Corn Yield Prediction Performance (Spatial Block-Cross-Validation) / Ефективність прогнозування врожайності кукурудзи на рівні пікселів (просторова блокова перехресна валідація)

Model	Input Time Series (Source)	RMSE (t/ha)	RMSE (t/ha)
Random Forest	Linear Interpolation	1.35	1.02
3D-CNN	STARFM (State-of-the-art)	0.88	0.70
CNN-LSTM	SAR Only (No Opt/Met data)	1.10	0.85
ST-cGAN + CNN-LSTM	Proposed Fusion Pipeline	0.48	0.34

The results indicate that applying a 3D-CNN with traditional gap-filling (STARFM) results in decreased performance ($R^2 = 0.78$) during contiguous 3-week cloud simulations, as the algorithm fails to extrapolate nonlinear crop stress.

The "SAR Only" baseline yielded higher errors (RMSE = 1.10 t/ha). Error analysis revealed that this model systematically overestimated yield during drought years. The physical volume of the corn canopy, which drives SAR backscatter, remained high even as grain filling failed due to biochemical water stress. The proposed framework, integrating structural, meteorological, and historical components, decreased the RMSE to 0.48 t/ha ($R^2 = 0.90$). A paired t-test confirmed that the performance improvement of the proposed ST-cGAN framework over the STARFM baseline was statistically significant ($p < 0.01$).

Discussion of research results. In the article [1], generative frameworks successfully synthesize pseudo-NDVI under normal conditions but remain limited during abrupt environmental stress, often producing phenological "flickering" and biologically impossible biomass fluctuations. Similarly, researchers [4] designed advanced cGANs for cloud removal that treat translation as a static, atemporal task focused purely on spatial reconstruction. This approach has limitation for the temporal hysteresis demonstrated by the article [17] proved that SAR backscatter anomalies lag significantly behind optical vegetation index anomalies during severe droughts because the physical macrostructure of the crop temporarily remains intact after photosynthetic cessation. Because purely SAR-driven baseline models inherit this physical lag, they inadvertently overestimate crop health during critical, early-season stress windows. Resolving this biochemical gap requires incorporating continuous meteorological priors alongside temporal gradient constraints to prevent phenological artifacts and ensure biological consistency over sequential time steps.

In the article [6], researchers highlight that yield forecasting objectives are fundamentally bifurcated by their operational timelines, where mid-season forecasting is critical for local agronomic interventions and late-season forecasting serves logistical requirements such as optimizing harvesting schedules, storage, and crop marketing. Furthermore, emphasize the importance of modeling the highly non-linear, cumulative effects of environmental stress at different development stages, such as distinguishing the minor impact of early-stage vegetative drought from the catastrophic impact of reproductive silking drought. To address these operational needs, recent literature calls for yield models that utilize rigorous spatial block-cross-validation to prevent geographical overfitting and ensure reliable, actionable predictions across diverse agro-ecological conditions.

In the research [10], the authors analyzed multi-source early fusion strategies and the integration of meteorological priors; building upon this approach, our study proves that conditioning the generative network on precipitation and temperature history (ERA5-Land) alongside temporal logic provides a vastly more robust solution for continuous, pixel-level crop yield forecasting under persistent atmospheric occlusions than traditional methods.

In the papers [13, 15], the authors highlighted the necessity of strict spatial block-cross-validation and the inherent strengths of temporal neural networks applied to satellite time series; corroborating these principles, the application of such rigorous validation in our framework confirmed that the obtained $R2$ of 0.90 was not an artifact of memorized local geometries but a genuine mapping of spatiotemporal crop dynamics, further enhanced by our addition of the Red-Edge band (NDRE) which maintained a linear correlation with final harvest weights even during late grain-filling stages.

While our framework successfully overcomes temporal and biochemical gaps, it remains bounded by the native spatial resolution of its satellite inputs. Contemporary reviews on Sentinel-2 applications in precision agriculture point out that its 10-meter spatial resolution represents an operational bottleneck for identifying highly localized, sub-canopy micro-stressors, such as localized pest damage, early-stage foliar diseases, or weed patches [20]. At this scale, individual pixels aggregate highly mixed spectral signatures, meaning that localized biotic stress can generally only be detected once macroscopic, canopy-wide damage has already occurred. Furthermore, comparative assessments between satellite-based models and Unmanned Aerial Vehicles (UAVs) show that satellite platforms naturally depict smoothed, generalized development trends [3]. This smoothing is particularly pronounced due to the coarser, native 20-meter resolution of the critical Red-Edge bands on Sentinel-2, which are required for high-fidelity NDRE calculation.

Thus, based on the discussion of the obtained results, it is confirmed that conducting this specific research was highly expedient and necessary, since, although various approaches to missing data reconstruction have been proposed in the analyzed literature, none of the previous authors successfully integrated meteorological priors, temporal discriminators, and multimodal early fusion into a single generative pipeline capable of overcoming the "biochemical gap" of radar data in the exact comprehensive and highly accurate manner as accomplished in this study.

So, based on the results of the work performed, it is possible to formulate the following scientific novelty and practical significance of the research results.

Scientific novelty of the obtained research results – methodology for constructing a comprehensive architecture for spatio-temporal multimodal data integration has been further developed; unlike existing methods, it integrates meteorological prior knowledge and temporal discriminators into the ST-cGAN generative network, thereby forming a model with physiological constraints that eliminates the structural and biochemical discrepancies inherent in the conversion of SAR images to optical ones, which collectively makes it possible to accurately synthesise stress indicators for agricultural crops under conditions of prolonged cloud cover, without overestimating their health status.

Practical significance of the research results – can be applied to the creation of a weather-independent decision-support framework for agricultural management, providing agronomists with a reliable tool to generate uninterrupted mid-season yield heatmaps with high predictive accuracy (RMSE = 0.48 t/ha) for optimizing harvest logistics, planning localized interventions, and executing precise financial forecasting while entirely eliminating the historical dependency on clear-sky satellite observations.

Conclusions / Висновки

A comprehensive multimodal spatio-temporal model has been developed spatiotemporal framework that a comprehensive multimodal spatiotemporal framework has been developed fuses synthetic aperture radar, historical optical data, and meteorological priors to overcome the biological blindness of radar signals, thereby enabling continuous, gap-free reconstruction of vegetation indices for high-precision corn yield forecasting under persistent atmospheric occlusions. Based on the results of the research the following main conclusions can be drawn.

1. Developed a generative ST-cGAN architecture to reconstruct continuous, cloud-free vegetation indices by conditioning the synthesis process on precipitation and temperature history (ERA5-Land) alongside SAR backscatter. This integration produced biologically accurate NDVI and NDRE maps (SSIM = 0.89) and eliminated the traditional limitation of overestimating crop health during drought events.
2. Implemented a sequential CNN-LSTM predictive pipeline designed to extract complex spatiotemporal features from the reconstructed 15-step time series. This recurrent architecture was demonstrated to accurately capture the nonlinear dynamics of crop phenology and isolate critical stress periods throughout the growing season.
3. Achieved predictive superiority through a framework that reached a pixel-level RMSE of 0.48 t/ha (R^2 = 0.90) under strict spatial block-cross-validation. This pipeline demonstrated a statistically significant improvement (p < 0.01) over established spatial gap-filling baselines, such as STARFM, proving its capability to map genuine crop dynamics without spatial data leakage.
4. Formulated the practical agronomic value of a scientifically defensible decision-support tool that remains reliable regardless of autumn cloud occlusions. The developed system allows agricultural management to optimize harvest logistics, execute variable-rate desiccation, and conduct precise financial forecasting mid-season.
5. Defined future research perspectives focusing on end-to-end knowledge distillation to compress the deep learning architecture for edge deployment on agricultural IoT devices. Additionally, future work will explore the integration of high-resolution UAV imagery to enable the precise detection of sub-canopy micro-stressors.

References

1. Arai, K., Maruta, R., & Okumura, H. (2026). Method for Generating Pseudo-NDVI from RVI Derived from Satellite-Borne SAR Imagery Data Using CycleGAN and pix2pix Models. *Information*, *17*(2), article number 154. https://doi.org/10.3390/info17020154
2. Bermudez, J., Happ, P. N., Feitosa, R. Q., & Oliveira, D. A. B. (2019). Synthesis of multispectral optical images from SAR/optical time series using conditional generative adversarial networks. *IEEE Geoscience and Remote Sensing Letters*, *16*(8), 1220–1224. https://doi.org/10.1109/LGRS.2019.2894734
3. Bukowiecki, J., Rose, T., & Kage, H. (2021). Sentinel-2 Data for Precision Agriculture? A UAV-Based Assessment. *Sensors*, *21*(8), article ID 2861. https://doi.org/10.3390/s21082861
4. Darbaghshahi, F. N., Mohammadi, M. R., & Soryani, M. (2022). Cloud Removal in Remote Sensing Images Using Generative Adversarial Networks and SAR-to-Optical Image Translation. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–9. https://doi.org/10.48550/arXiv.2012.12180
5. Ebel, P., Meraner, A., Schmitt, M., & Zhu, X. X. (2020). Multisensor data fusion for cloud removal in global and all-season Sentinel-2 imagery. *IEEE Transactions on Geoscience and Remote Sensing*, *59*(7), 5866–5878. https://doi.org/10.1109/TGRS.2020.3024744
6. Filippi, P., Han, S. Y., & Bishop, T. F. A. (2025). On crop yield modelling, predicting, and forecasting and addressing the common issues in published studies. *Precision Agriculture*, 26, article number 8. https://doi.org/10.1007/s11119-024-10212-2
7. Gao, F., Masek, J., Schwaller, M., & Hall, F. (2006). On the blending of the Landsat and MODIS surface reflectance: Predicting daily Landsat surface reflectance. *IEEE Transactions on Geosci-*

ence and Remote Sensing, *44*(8), 2207–2218. https://doi.org/10.1109/TGRS.2006.872081

8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing* Systems, 27. URL: https://proceedings.neurips.cc/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html
9. Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, *202*, 18–27. https://doi.org/10.1016/j.rse.2017.06.031
10. Ienco, D., Interdonato, S., Gaetano, R., & Minh, D. H. (2019). Combining Sentinel-1 and Sentinel-2 satellite image time series for land cover mapping via a multi-source deep learning architecture. *ISPRS Journal of Photogrammetry and Remote Sensing*, *158*, 11–22. https://doi.org/10.1016/j.isprsjprs.2019.09.016
11. Isola, P., Zhu, J., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5967–5976. https://doi.org/10.1109/CVPR.2017.632
12. Khaki, S., & Wang, L. (2019). Crop yield prediction using deep neural networks. *Frontiers in Plant Science*, *10*, article ID 621. https://doi.org/10.3389/fpls.2019.00621
13. Meyer, H., & Pebesma, E. (2018). Improving performance of machine learning models for geographic data. *Nature Communications*, *9*(1), article ID 3993. https://doi.org/10.1038/s41467-018-02558-0
14. Muñoz-Sabater, J., Dutra, E., Agustí-Panareda, A., Albergel, C., Arduini, G., Balsamo, G., Boussetta, S., Choulga, M., Harrigan, S., Hersbach, H., Martens, B., Miralles, D. G., Piles, M., Rodríguez-Fernández, N. J., Zsoter, E., Buontempo, C., & Thépaut, J.-N. (2021). ERA5-Land: A state-of-the-art global reanalysis dataset for land applications. *Earth System Science Data*, *13*(9), 4349–4383. https://doi.org/10.5194/essd-13-4349-2021
15. Pelletier, C., Webb, G. I., & Petitjean, F. (2019). Temporal convolutional neural network for the classification of satellite image time series. *Remote Sensing*, *11*(5), article ID 523. https://doi.org/10.3390/rs11050523
16. Rußwurm, M., & Körner, M. (2018). Multi-temporal land cover classification with sequential recurrent encoders. *ISPRS International Journal of Geo-Information*, *7*(4), article number 129. https://doi.org/10.3390/ijgi7040129
17. Schellenberg, K., Jagdhuber, T., Zehner, M., Hese, S., Urban, M., Urbazaev, M., Hartmann, H., Schmullius, C., & Dubois, C. (2023). Potential of Sentinel-1 SAR to Assess Damage in Drought-Affected Temperate Deciduous Broadleaf Forests. *Remote Sensing*, *15*(4), article ID 1004. https://doi.org/10.3390/rs15041004
18. Schmitt, M., Hughes, L. H., Qiu, C., & Zhu, X. X. (2019). SEN12MS – A curated dataset of georeferenced multi-spectral Sentinel-1/2 imagery for deep learning and data fusion. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *IV-2/W7*, 153–160. https://doi.org/10.5194/isprs-annals-IV-2-W7-153-2019
19. Schwalbert, F. G., Amado, T. J. C., Nieto, L., Corassa, G. M., Rice, C. W., Peralta, N. R., Schauberger, B., Gornott, C., & Ciampitti, I. A. (2020). Mid-season county-level corn yield forecasts in the US Corn Belt integrating satellite imagery and weather variables. *Crop Science, 60*(2), 739–750. https://doi.org/10.1002/csc2.20053
20. Segarra, J., Buchaillot, M. L., Araus, J. L., & Kefauver, S. C. (2020). Remote Sensing for Precision Agriculture: Sentinel-2 Improved Features and Applications. *Agronomy*, *10*(5), article ID 641. https://doi.org/10.3390/agronomy10050641
21. Singh, D., Kumar, D., & Singh, S. K. (2021). Cloud-free Sentinel-2 image synthesis using Sentinel-1 SAR data via CycleGAN. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *14*, 6933–6944. https://doi.org/10.1109/JSTARS.2021.3090267
22. Xiong, Q., Li, G., Yao, X., & Zhang, X. (2023). SAR-to-optical image translation and cloud removal based on conditional generative adversarial networks: Literature survey, taxonomy, evaluation indicators, limits and future directions. *Remote Sensing, 15*(4), article ID 1137. https://doi.org/10.3390/rs15041137
23. 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), 4590–4596. https://doi.org/10.1609/aaai.v31i1.11172

С. І. Балук, В. Я. Антонів

Національний університет "Львівська політехніка", м. Львів, Україна

РЕКОНСТРУКЦІЯ СУПУТНИКОВИХ ДАНИХ ЗА УМОВ ХМАРНОСТІ ДЛЯ РОГНОЗУВАННЯ ВРОЖАЙНОСТІ КУКУРУДЗИ ЗЛИТТЯМ МУЛЬТИМОДАЛЬНИХ ЗНІМКІВ

Розроблено та апробовано комплексну інтелектуальну систему прогнозування врожайності кукурудзи, що базується на синергії просторово-часової генеративно-змагальної мережі (ST-cGAN) та рекурентної архітектури CNN-LSTM. Застосовано мультимодальне поєднання погодно-незалежних радіолокаційних даних Sentinel-1, метеорологічних чинників ERA5-Land (середньодобова температура та сума опадів) та історичних оптичних спостережень Sentinel-2. Вирішено проблему "біохімічної сліпоти" радіолокаційних сигналів, за якої радар фіксує фізичну структуру рослин, але не здатний оцінити їхню фотосинтетичну активність та обмеження традиційного оптичного зондування земної поверхні, спричиненних частою хмарністю. Для цього замість багаточастотного SAR-зондування було застосовано стратегію раннього злиття даних. Мережа ST-cGAN одночасно обробляє поточні SAR-дані для визначення структурного макростану, історичні оптичні дані як останній достовірний біохімічний стан та 10-денну історію метеорологічних показників. Відповідно, якщо радар фіксує високу біомасу, але метеорологічні дані свідчать про посуху, нейромережа математично розпізнає біологічний стрес і синтезує знижені індекси NDVI та NDRE, запобігаючи генеруванню хибно здорових посівів. Здійснено динамічний синтез пропущених вегетаційних індексів NDVI та NDRE за умов тривалих періодів суцільної хмарності (до 21 дня), при цьому досягнуто високої структурної подібності синтезованих карт (індекс SSIM = 0,89). Введено в архітектуру мережі часові дискримінатори для аналізу послідовностей кадрів, що на 18 % покращило індекс збереження меж полів та мінімізувало появу просторових артефактів на границях угідь. Проведено піксельну регресію для тестової ділянки площею 15 000 га у Західному Лісостепу України на основі реконструйованих часових рядів, які охоплювали 15 критичних фенологічних етапів розвитку культури. Встановлено, що запропонована архітектура нейронної мережі знижує середньоквадратичну помилку прогнозу (RMSE) до 0,48 т/га за коефіцієнта детермінації (R^2) 0,90. На основі статистичного аналізу (парний *t*-тест) доведено вагому перевагу розробленого методу над галузевими алгоритмами заповнення пропусків, зокрема – STARFM, який є неспроможним адекватно моделювати нелінійний стрес посівів без оптичних даних. Представлено масштабовану систему підтримки прийняття рішень для оптимізації логістики збирання врожаю та точного фінансового планування роботи агропідприємств незалежно від атмосферних умов. Окреслено перспективи подальших етапів дослідження, що полягають у дистиляції знань нейромережі для периферійних IoT-пристроїв та інтеграції даних БПЛА.

Ключові слова: глибоке навчання нейронних мереж; штучний інтелект в агрономії; просторово-часові мережі; точне землеробство; предиктивне моделювання.