

Exploring Physics-Informed Neural Networks for Crop Yield Loss Forecasting

Miro Miranda^{1,2,*}, Marcela Charfuelan², and Andreas Dengel^{1,2}

¹Department of Computer Science, University of Kaiserslautern-Landau, Kaiserslautern, Germany

²German Research Center for Artificial Intelligence, Kaiserslautern, Germany

Abstract

In response to climate change, assessing crop productivity under extreme weather conditions is essential to enhance food security. Crop simulation models, which align with physical processes, offer explainability but often perform poorly. Conversely, machine learning (ML) models for crop modeling are powerful and scalable yet operate as black boxes and lack adherence to crop growth’s physical principles. To bridge this gap, we propose a novel method that combines the strengths of both approaches by estimating the water use and the crop sensitivity to water scarcity at the pixel level. This approach enables yield loss estimation grounded in physical principles by sequentially solving the equation for crop yield response to water scarcity, using an enhanced loss function. Leveraging Sentinel-2 satellite imagery, climate data, simulated water use data, and pixel-level yield data, our model demonstrates high accuracy, achieving an R^2 of up to 0.77—matching or surpassing state-of-the-art models like RNNs and Transformers. Additionally, it provides interpretable and physical consistent outputs, supporting industry, policymakers, and farmers in adapting to extreme weather conditions.

1 Introduction

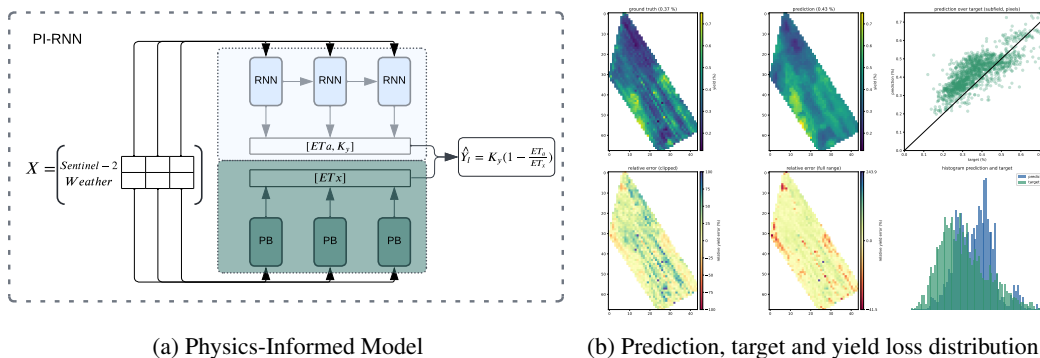


Figure 1: (a) Framework for physics-informed yield loss forecasting. Data modalities, including Sentinel-2 imagery and climate variables, are used to train an RNN that predicts biophysical properties for each time point (actual water use (ET_a), and the crop susceptibility to water scarcity (K_y)). A simulation model predicts the maximum water use (ET_x). The predicted biophysical properties are used to calculate the actual yield loss (Y_l). We leverage prior knowledge describing the relationship between water use and the relative yield loss. (b) Performance plots for visual inspection of a single field. Yield data from cereals in Switzerland is shown, harvested in 2020.

In response to climate change, closing the gap between potential and actual yields is an urgent task to sustain food security [7]. Extreme weather conditions like droughts and floodings are pressing challenges in the agricultural sector, directly affecting productivity, and causing yield losses [2, 15]. Digital agriculture has emerged as a key strategy to address these challenges, providing tools for timely, informed decision-making by policymakers, industry stakeholders, and farmers [19]. The response of crop yields to water availability has been a central focus of research for decades, serving as a critical parameter in assessing crop resilience under extreme climate conditions [17]. Traditionally, simulation models, also known as process-based models, have been employed to capture this relationship. These models build upon biological and physical principles to offer high explainability, supporting decisions in areas such as irrigation, fertilization, pest management, and disease control. However, crop simulation models often struggle with high-resolution, multidimensional data, are computationally intensive, and require calibration, rendering their applicability to high spatial resolution (pixel-level) unfeasible. Furthermore, they are often simplified representations of reality, relying on approximations to maintain computational efficiency [12], oftentimes resulting in inaccurate performances [13]. Therefore, machine learning (ML) models are increasingly utilized for crop productivity estimation [21]. Recent studies demonstrate impressive scalability and accuracy, even at pixel resolution [9, 16], by handling complex and multidimensional data [14]. However, ML models are often criticized for their black-box characteristics, limiting their interpretability [4]. Additionally, ML are seldomly designed to follow underlying physical principles of plant growth [8], thereby hardly ever providing meaningful intermediate outputs, which are essential for tracking crop responses to changing environmental conditions over time. There is a growing demand to integrate the strengths of data-driven ML approaches with the interpretability of simulation models [5, 12]. While some recent studies attempt to combine ML and prior knowledge for yield prediction, also under water-limited conditions [20], these approaches generally focus on data enrichment rather than on producing interpretable and physical consistent sequential outputs of crop physiology, similar to simulation models. Consequently, they often retain the black-box limitations of conventional ML models, providing limited insight into crop physiology by potentially violating governing physical principles. This study seeks to bridge the gap between interpretable simulation models and high-performance ML models by introducing a physics-informed approach to crop yield loss forecasting. We argue that we can learn more precisely the water use and the crop susceptibility to water scarcity, by building upon Sentinel-2 (S2) multispectral imagery, climate data, and simulations, coupled with pixel-level yield data from cereal crops harvested between 2017–2021 in Switzerland. We demonstrate yield loss estimation at the pixel level, grounded on physical principles by sequentially solving the equation for crop yield response to water scarcity [6], using an enhanced loss function. Initial results indicate high potential, achieving an R^2 of up to 0.77, matching or surpassing the performance of state-of-the-art (SOTA) ML models for yield prediction, including RNNs and Transformers, while providing interpretable and physical consistent outputs, supporting industry, policymakers, and farmers in adapting to extreme weather conditions.

2 Methodology

Problem Definition The water use, also called *evapotranspiration* (ET), is the sum of all biophysical processes in which liquid water is converted to water vapor from various surfaces, including topsoil and vegetation. For a specific crop type, the Food and Agriculture Organization (FAO) differentiates between the maximum (ET_x) and the actual ET (ET_a). The former represents the water use under standard and non-limiting environmental conditions and is solely impacted by climate conditions and crop specifics, achieving full productivity, i.e., disease-free, well-fertilized, and under optimum soil water conditions. In contrast, the ET_a represents the water use under limiting environmental conditions, resulting in a reduction in water use and ultimately causing a reduction of crop productivity. This reduction in productivity is defined as yield loss and typically expressed in relative terms. Various factors cause productivity limiting conditions, including soil infertility, soil salinity, limited soil water content, diseases, and poor management. Especially, in the light of extreme weather conditions, the frequency of severe droughts and floodings is expected to increase, causing either water scarcity or water abundance, negatively impacting ET and crop productivity [12, 2]. Assessing the reduction in ET, is a major challenge in agriculture. A detailed description of biophysical processes of the ET is described by the FAO-56 method [1], using the Penman–Monteith equation, recommended for daily ET (mm/day) estimation. The ET under limiting condition is defined by:

$$ET_a = K_s \cdot ET_x, \quad (1)$$

with K_s being a stress coefficient, resulting in a reduction of the ET. In an earlier work, the FAO described the relationship between water use and the relative yield loss [6], stating that the relative reduction in ET is related to the relative reduction in yield:

$$Y_l = \left(1 - \frac{Y_a}{Y_x}\right) = K_y \left(1 - \frac{ET_a}{ET_x}\right), \quad (2)$$

with Y_a as the actual yield, Y_x the maximum yield, or potential yield, and K_y as the crop yield response factor. This factor represents the effect of a reduction in ET on the crop yield by capturing the complex relationship between crop water use and productivity. More specifically, $K_y > 1$ indicates high sensitivity to water deficits with a proportionally larger yield reduction, and $K_y < 1$ indicates higher resilience to water deficits. Different studies exist, that empirically estimated K_y coefficients for various crops. However, often reporting differences, making the equation difficult to solve analytically. Furthermore, K_y values change over the growing period, since many crops exhibit variable susceptibility to water scarcity over the growing period. More importantly, while ET_x can be more or less accurately estimated, by following [1], estimating the ET_a is very difficult with high precision, making Eq. 2 difficult to solve accurately in reality.

Architecture & Evaluation To address this limitation, we propose to learn both ET_a and K_y using a Recurrent Neural Network (RNN). We constrain the network by physical information such that the reduction in ET corresponds to the observed reduction in yield. We argue that simulated ET_x values are sufficiently accurate. This is supported by previous studies [3], allowing us to focus at ET_a and K_y at pixel resolution. We leverage an RNN, more specifically a LSTM [11] backbone architecture, where each hidden state is passed to a sequential layer with 128 hidden units, incorporating a linear layer, batch normalization, and dropout. Finally, two linear layers are incorporated with a single output channel each, reflecting K_y and ET_a , respectively. Physical constraints are imposed by integrating information based on the ET and the relative reduction in crop yield, resulting in a physics-informed RNN (PI-RNN). To estimate the yield reduction at time step t_i as a function of a reduction in ET, a two component loss term is employed:

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_l + \lambda_2 \mathcal{L}_{phys}, \quad (3)$$

$$\text{with: } \mathcal{L}_l = \mathbb{E} \left[(\hat{y}_a - y_a)^2 \right]$$

$$\text{and: } \mathcal{L}_{phys} = \mathbb{E} \left[\underbrace{1_{\{ET_a < 0\}} \cdot (ET_a)^2}_{\text{lower bound penalty}} + \underbrace{1_{\{ET_a > ET_x\}} \cdot (ET_a - ET_x)^2}_{\text{upper bound penalty}} + \underbrace{1_{\{0 \leq ET_a \leq ET_x\}} \cdot (ET_a - ET_x)^2}_{\text{within bounds MSE}} \right].$$

With $1\{\cdot\}$ as the indicator function, which equals 1 if the condition inside the braces is true and 0 otherwise. The first component pushes the network to predict K_y and ET_a that, by using Eq. 2 the predicted yield loss (\hat{y}_a) is close to the actual yield loss (y_a) of the ground truth data. The second component forces the network to produce ET_a values bounded between $[0, ET_x]$ and sufficiently close to ET_x , to solve Eq. 2 with biophysical consistency. Moreover, λ is a hyperparameter that controls the weighting of both terms. Figure 1 (a) depicts the overall architecture of the PI-RNN.

Data & Evaluation As ground truth data, combine harvester yield data is used. More specifically, 54098 yield samples from 54 yield maps of cereal crops are used. This data was harvested in Switzerland between 2017 and 2021, containing georeferenced data points with information about the yield in $10m$ pixel resolution. An example yield map is depicted in Figure 1 (b). For detailed information about the dataset and preprocessing, we refer the reader to [18]. We define the maximum yield sample as the yield potential (Y_x). The reduction for each sample is expressed in relative terms as $(1 - \frac{y_a}{Y_x})$. For the simulations, we employ the FAO paper-56 [1] that simulates ET_x over time.

For training, S2 L2A multispectral time series data is employed using 10 spectral bands, available from seeding to harvesting with $10m$ resolution as the model input. Additionally, weather data is incorporated, including the total precipitation and temperature, derived from [10]. Data modalities are early fused using the raw time series of S2 images, by aggregating weather features between S2 time steps.

For each experiment, a K-fold cross-validation (K=10) is performed, where results correspond to the average across folds. For quantitative evaluation, standard regression metrics are used. This includes the coefficient of determination (R^2), root-mean-square error (RMSE), mean absolute error (MAE). To evaluate the potential of the proposed method, we compare results against current SOTA models

for crop yield prediction, including RNN (LSTM) [16] and Transformers [9]. Note that these models cannot predict physical consistent outputs, only the final relative yield loss.

3 Results

In Table 1, we evaluate the yield loss prediction performance by comparing the predicted crop yield loss to the actual yield loss and compare the results against a RNN and Transformer model. All models are trained on S2 and weather data. We demonstrate that our method outperforms the RNN model on various regression metrics, such as R^2 and RMSE. In detail, we demonstrate a 2 percentage point (p.p.) improvement. Compared to the Transformer model, our method achieves similar results on all metrics. Only a minor decrease of 1 p.p. in R^2 is observed. Additionally, we showcase the performance of the simulation model, and highlight its weak performance to predict the actual yield loss effectively. This is because the water use cannot be predicted at the pixel level, as it depends on weather data with a $30km$ spatial resolution. Additionally, the simulation models struggle to estimate actual water consumption with high accuracy (see Figure 2 (a)). Qualitatively, Figure 1 (b) depicts an example yield map, showcasing accurate predictions at pixel level with high spatial variability, close to the target data. To further assess whether the inclusion of S2 data enhances the accuracy of crop yield loss estimation at the pixel level, we trained several models with different data modalities. Results are depicted in Table 2, underlying the improved performance, when S2 data is used. In detail, a significant improvement of 21 p.p. in R^2 is shown over the model trained solely on weather data.


Figure 2 (a) illustrates the sequential estimations of ET_a and K_y values up to 120 days before harvest, alongside simulated ET_x and ET_a values. The Figure indicates that K_y values remain below 1, suggesting a higher resilience to water scarcity in this dataset. Furthermore, it is depicted that estimated ET_a is consistently lower than the simulated ET_x , indicating yield loss due to water limitation. In Figure 2(b), the relationship between ET and predicted yield loss is shown. The Figure demonstrates that yield loss decreases during the growing period in correlation with increased

Model	Modalities	MAE	RMSE	R2
RNN		0.05	0.08	0.75
Transformer	Sentinel-2 + Weather	0.05	0.07	0.78
PI-RNN (ours)		0.05	0.07	0.77
Simulation	Weather	0.45	0.47	-6.92

Table 1: Overview table of yield loss prediction performance of different models. All models are trained using S2 and weather data.

Model	Modalities	MAE	RMSE	R2
	Weather	0.08	0.11	0.46
PI-RNN (ours)	Sentinel-2	0.05	0.07	0.75
	Sentinel-2 + Weather	0.05	0.07	0.77

Table 2: Overview table of yield loss prediction performance of the PI-RNN (our) model, using different input modalities.



(a) Predictions and simulations of crop properties.

(b) R^2 scores and yield loss prediction.

Figure 2: Visualization of sequential estimations and predictions of biophysical crop properties and model performance up to 120 days before harvest.

ET, as formally defined. Additionally, yield loss predictions become more accurate as the harvest approaches.

4 Conclusion

Informed Neural Networks hold significant potential for crop yield modeling, offering enhanced adaptability to challenging environmental conditions. We presented a novel approach to modeling crop productivity under environmental constraints by incorporating biophysical properties and demonstrated promising experimental outcomes. To further assess the validity of the proposed method, additional experiments on larger and more diverse datasets are necessary, encompassing a wide range of yield limiting conditions.

References

- [1] Richard G Allen, Luis S Pereira, Dirk Raes, Martin Smith, et al. Crop evapotranspiration-guidelines for computing crop water requirements-fao irrigation and drainage paper 56. *Fao, Rome*, 300(9):D05109, 1998.
- [2] Naveen Kumar Arora. Impact of climate change on agriculture production and its sustainable solutions. *Environmental sustainability*, 2(2):95–96, 2019.
- [3] Jiabing Cai, Yu Liu, Tingwu Lei, and Luis Santos Pereira. Estimating reference evapotranspiration with the fao penman–monteith equation using daily weather forecast messages. *Agricultural and Forest Meteorology*, 145(1-2):22–35, 2007.
- [4] Davide Castelvecchi. Can we open the black box of ai? *Nature News*, 538(7623):20, 2016.
- [5] Tirtharaj Dash, Sharad Chitlangia, Aditya Ahuja, and Ashwin Srinivasan. A review of some techniques for inclusion of domain-knowledge into deep neural networks. *Scientific Reports*, 12(1):1040, 2022.
- [6] J Doorenbos and AH Kassam. Yield response to water. *Irrigation and drainage paper*, 33:257, 1979.
- [7] Tony Fischer, Derek Byerlee, and Greg Edmeades. Crop yields and food security: will yield increases continue to feed the world? In *Capturing Opportunities and Overcoming Obstacles in Australian Agronomy*, 2012.
- [8] Erhu He, Yiqun Xie, Licheng Liu, Weiye Chen, Zhenong Jin, and Xiaowei Jia. Physics guided neural networks for time-aware fairness: an application in crop yield prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14223–14231, 2023.
- [9] Patrick Helber, Benjamin Bischke, Peter Habelitz, Cristhian Sanchez, Deepak Pathak, Miro Miranda, Hiba Najjar, Francisco Mena, Jayanth Siddamsetty, Diego Arenas, et al. Crop yield prediction: An operational approach to crop yield modeling on field and subfield level with machine learning models. In *IGARSS - IEEE International Geoscience and Remote Sensing Symposium*, pages 2763–2766. IEEE, 2023.
- [10] Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [12] Yinhong Kang, Shahbaz Khan, and Xiaoyi Ma. Climate change impacts on crop yield, crop water productivity and food security—a review. *Progress in natural Science*, 19(12):1665–1674, 2009.
- [13] Guoyong Leng and Jim W Hall. Predicting spatial and temporal variability in crop yields: an inter-comparison of machine learning, regression and process-based models. *Environmental research letters: ERL [Web site]*, 15(4):044027, 2020.

- [14] Miro Miranda, Deepak Pathak, Marlon Nuske, and Andreas Dengel. Multi-modal fusion methods with local neighborhood information for crop yield prediction at field and subfield levels. In *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*, pages 4307–4311. IEEE, 2024.
- [15] Amy Molotoks, Pete Smith, and Terence P Dawson. Impacts of land use, population, and climate change on global food security. *Food and Energy Security*, 10(1):e261, 2021.
- [16] Deepak Pathak, Miro Miranda, Francisco Mena, Cristhian Sanchez, Patrick Helber, Benjamin Bischke, Peter Habelitz, Hiba Najjar, Jayanth Siddamsetty, Diego Arenas, Michaela Vollmer, Marcela Charfuelan, Marlon Nuske, and Andreas Dengel. Predicting Crop Yield with Machine Learning: An Extensive Analysis of Input Modalities and Models on a Field and Sub-Field Level. In *IGARSS- IEEE International Geoscience and Remote Sensing Symposium*, pages 2767–2770, 2023.
- [17] Luis S Pereira, Richard G Allen, Martin Smith, and Dirk Raes. Crop evapotranspiration estimation with fao56: Past and future. *Agricultural water management*, 147:4–20, 2015.
- [18] Gregor Perich, Mehmet Ozgur Turkoglu, Lukas Valentin Graf, Jan Dirk Wegner, Helge Aasen, Achim Walter, and Frank Liebisch. Pixel-based yield mapping and prediction from Sentinel-2 using spectral indices and neural networks. *Field Crops Research*, 292:108824, 2023.
- [19] Ribana Roscher, Lukas Roth, Cyrill Stachniss, and Achim Walter. Data-Centric Digital Agriculture: A Perspective. *arXiv preprint arXiv:2312.03437*, 2023.
- [20] Guanyuan Shuai and Bruno Basso. Subfield maize yield prediction improves when in-season crop water deficit is included in remote sensing imagery-based models. *Remote Sensing of Environment*, 272:112938, 2022.
- [21] Thomas Van Klompenburg, Ayalew Kassahun, and Cagatay Catal. Crop yield prediction using machine learning: A systematic literature review. *Computers and Electronics in Agriculture*, 177:105709, 2020.