

MARLP: Time-series Forecasting Control for Agricultural Managed Aquifer Recharge

Yuning Chen

ychen372@ucmerced.edu
University of California, Merced
Merced, CA, USA

Kang Yang

kyang73@ucmerced.edu
University of California, Merced
Merced, CA, USA

Zhiyu An

zan7@ucmerced.edu
University of California, Merced
Merced, CA, USA

Brady Holder

Luke Paloutzian*
University of California, Agriculture
and Natural Resources
Parlier, CA, USA

Khaled M. Bali

kmbali@ucanr.edu
University of California, Agriculture
and Natural Resources
Parlier, CA, USA

Wan Du ✉

wdu3@ucmerced.edu
University of California, Merced
Merced, CA, USA

ABSTRACT

The rapid decline in groundwater around the world poses a significant challenge to sustainable agriculture. To address this issue, agricultural managed aquifer recharge (Ag-MAR) is proposed to recharge the aquifer by artificially flooding agricultural lands using surface water. Ag-MAR requires a carefully selected flooding schedule to avoid affecting the oxygen absorption of crop roots. However, current Ag-MAR scheduling does not take into account complex environmental factors such as weather and soil oxygen, resulting in crop damage and insufficient recharging amounts. This paper proposes MARLP, the first end-to-end data-driven control system for Ag-MAR. We first formulate Ag-MAR as an optimization problem. To that end, we analyze four-year in-field datasets, which reveal the multi-periodicity feature of the soil oxygen level trends and the opportunity to use external weather forecasts and flooding proposals as exogenous clues for soil oxygen prediction. Then, we design a two-stage forecasting framework. In the first stage, it extracts both the cross-variate dependency and the periodic patterns from historical data to conduct preliminary forecasting. In the second stage, it uses weather-soil and flooding-soil causality to facilitate an accurate prediction of soil oxygen levels. Finally, we conduct model predictive control (MPC) for Ag-MAR flooding. To address the challenge of large action spaces, we devise a heuristic planning module to reduce the number of flooding proposals to enable the search for optimal solutions. Real-world experiments show that MARLP reduces the oxygen deficit ratio by 86.8% while improving the recharging amount in unit time by 35.8%, compared with the previous four years.

*Email: beholder@ucanr.edu, ltpaloutzian@ucanr.edu.

✉ Wan Du is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0490-1/24/08
<https://doi.org/10.1145/3637528.3671533>

CCS CONCEPTS

• Applied computing → Agriculture; Forecasting; • Information systems → Data stream mining; Sensor networks.

KEYWORDS

Model Predictive Control, Time Series, Forecasting, Causal Learning, Agriculture

ACM Reference Format:

Yuning Chen, Kang Yang, Zhiyu An, Brady Holder, Luke Paloutzian, Khaled M. Bali, and Wan Du. 2024. MARLP: Time-series Forecasting Control for Agricultural Managed Aquifer Recharge. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3637528.3671533>




Figure 1: The benefits of applying Ag-MAR.

1 INTRODUCTION

Groundwater is an important resource for agricultural stability, for example, providing up to 60% of the water supply in dry years for California [20]. Due to recurring global droughts in recent years, groundwater pumping has increased significantly, exceeding the natural recharge rate and leading to insufficient water supply in underground aquifers [29]. This poses a threat to the food security of many regions [14, 24]. Therefore, careful management and conservation of groundwater resources are highlighted in many regions worldwide. In California, the Sustainable Groundwater Management Act (SGMA) has been introduced, aiming to achieve a balance between extraction and recharge within the next 20 years [10].

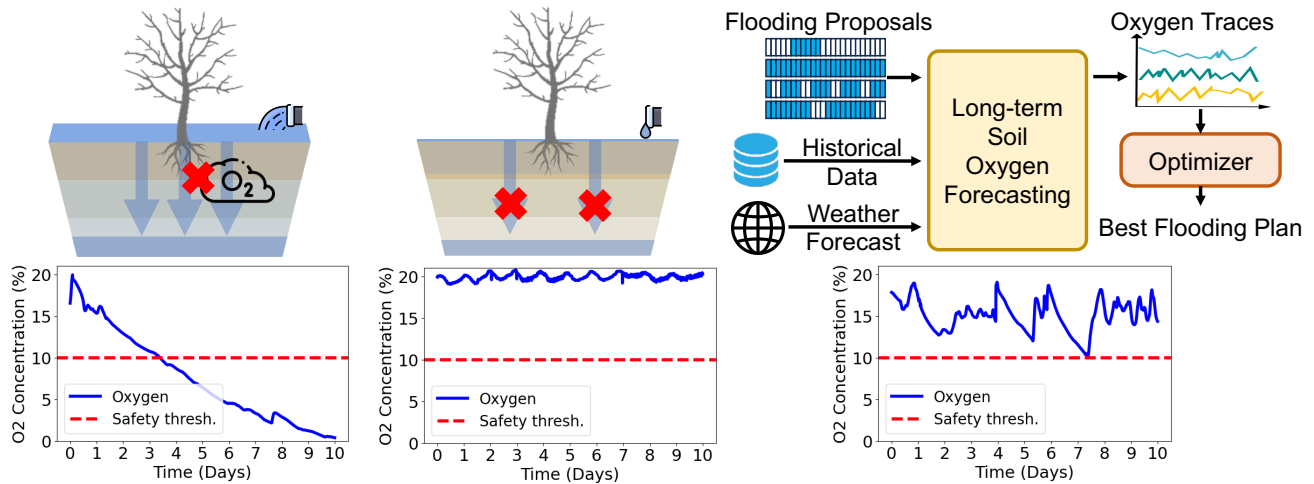


Figure 2: The illustration of oxygen fluctuation in continuous flooding, regular irrigation, and intermittent flooding.

Managed Aquifer Recharge (MAR) is a technique used to redirect excess surface water into underground aquifers during raining seasons, helping to replenish groundwater sources and reduce the impacts of excessive water withdrawal. This method is particularly adopted in areas close to rivers where the land is primarily used for farming, known as Agricultural Managed Aquifer Recharge (Ag-MAR). Ag-MAR has been recognized as an effective method for restoring water levels in depleted aquifers, enhancing the sustainability of crop yield, as well as for other advantages such as conditioning the soil before planting seasons and enhancing habitats for bird populations [14, 35], as illustrated in Figure 1.

The effective implementation of Ag-MAR remains a challenging problem. Flooding farmland reduces the soil’s oxygen content, as water hinders the dissolution of oxygen, limiting its availability to crop roots. Crops have specific tolerance thresholds for the soil oxygen level; if the level drops below the threshold, the crop root starts to decay, significantly damaging the crops. Consequently, Ag-MAR needs to optimize two objectives at the same time, i.e., maximizing the amount of water recharged to the underground aquifer while keeping the soil oxygen level above a predefined threshold. We show examples of the soil oxygen level trends under three different scenarios in Figure 2. From left to right of the figure: first, if the flooding continues for an continuous long period, the soil oxygen level continuously drops over the tolerance threshold for the plant root, resulting in root rot and future yield reduction. Second, if the water amount is slight, like sprinkler irrigation or light rain, it may only achieve a balance with evapotranspiration (ET) (water lost to air from soil surface and plants), resulting in insufficient aquifer recharge. Finally, the optimized solution is to flood on intermittent cycles, alternating between substantial flooding and drying periods so that the oxygen can diffuse into the soil.

Ideally, such a schedule should take into account multiple pieces of information, e.g., current soil oxygen level, future weather, and soil type. Given that the water permeation through soil is a continuous and long-lasting process, the effect of flooding actions can not be evaluated immediately, but will result in a delayed, long-term effect. Therefore, it is crucial to model and predict each flooding event’s

long-term impact, along with potential environmental dynamics, to determine the optimal flooding schedule. However, building an accurate prediction and control method that considers all of the above information remains an open challenge.

Analyzing our four-year real-world dataset in alfalfa fields¹ revealed two key observations that motivate our design, which we will describe in detail in Section 2. In short, we found out that:

- Soil oxygen exhibits a multi-periodicity pattern, modulated by environmental factors and flooding actions.
- Due to the strong causality between weather-soil and flooding-soil, exogenous clues such as weather forecasts and flooding proposals can boost oxygen prediction.

In light of these observations, we devise MARLP, a model predictive control (MPC) system for Ag-MAR. The core is a causality-aware long-term forecasting model that features a two-stage learning scheme. First, it integrates cross-variate and periodicity learning, generating a preliminary self-consistent multi-variate forecasting. In this step, environment-related periodicity is handled by segmenting the 1D data and reshaping them into a 2D format to facilitate learning interperiod-variation, while the action-triggered periodicity is filtered out. The exogenous clues are then used to calibrate the oxygen prediction via a causality-aware projection module. By combining them, the final oxygen prediction adapts well to the fluctuating rhythms of environmental factors.

This predictive capability sets the stage for the subsequent MPC workflow, which provides predictive outputs to any flooding proposals for days in the future. The optimizer can accordingly choose the best flooding strategy that both mitigates oxygen-deficit risks and maximizes the recharging to the underground aquifer. However, due to the long forecast window, the total number of flooding proposals is exponentially huge, which cannot be searched by brute-force or approximation algorithms. To this end, we propose a domain-specific heuristic planning module that filters the invalid flooding proposals in advance. The number of proposals has been

¹Alfalfa is a classical crop for Ag-MAR as it does not require any nitrogen fertilizer after establishment, gaining all necessary nitrogen from biological N_2 fixation and root uptake. This characteristic significantly minimizes the risk of nitrate leakage [43].




Figure 3: The rationale of soil oxygen and water content variation during flooding.

reduced to thousands, making it practical for the system to iterate all proposals and find the best in real-time.

To demonstrate the effectiveness of MARLP in predicting oxygen variations and scheduling flood actions, we perform statistical comparisons on past datasets, real-world control trials, and large-scale simulations. The experimental results show that the proposed algorithm and scheme outperform the state-of-the-art models and can provide practical and trustworthy decisions.

Our contributions are summarized as follows:

- We formulate the Ag-MAR optimization as a long-term model predictive control problem and design a customized time-series forecasting model that utilizes the external predictive input and extracts the periodicity features of data traces.
- We propose MARLP, an MPC workflow based on the model. To handle the large action space, we devise a heuristic planning scheme, making the forecasting-based search practical.
- We conduct extensive experiments to demonstrate the effectiveness of MARLP on both recharging amounts and safety warranty. We release the five-year dataset and code repository at https://github.com/ycucm/kdd24_marlp.

2 BACKGROUND

2.1 Ag-MAR Optimization

Ag-MAR is the most applied MAR scheme for two primary reasons: 1) The nearby river areas are usually fully plotted with fields, and applying water to these fields enables minimal water transportation; 2) the riverside fields typically represent the lowest points in the area, ensuring minimal lift work. For the area of resource-demand mismatch, Ag-MAR is carried out during the off-season when the surface water flow is adequate [46], and crops are not in their active growth phase.

Ideally, we should flood as much as possible into the field, but excessive ponding will cause oxygen deficiency and root rot[50]. To coordinate these contradictory objectives, water use must be intermittent, so that the soil can dry out and the oxygen level can take time to recover, as shown in Figure 2. The control objective is to maximize the amount of water recharged while maintaining a healthy oxygen level for the root zone. So this problem could




Figure 4: Multi-periodicity and causal relationship pattern.

be formulated as an optimization problem while the output is a flooding decision series, indicating whether and how much water to apply at each timestamp:

$$\begin{aligned}
 & \text{maximize} && \sum_{i \in T} (f_i F + p_i - ET_i - \Delta S_i) \\
 & \text{subject to} && O_i > O_{safe}, && \forall i \in T, \\
 & && f_i \in \{0, 1\}, && \forall i \in T,
 \end{aligned} \tag{1}$$

where f_i is a Boolean variable that indicates if the flooding is conducted in the i -th time step, F is the flooding gain (mm) per time step, p_i is the precipitation gain (mm), and ET_i is the evapotranspiration loss (ET) in unit time. ΔS_i is the change in soil storage (mm) (dependent on the available water capacity (AWC) of the soil). Surface runoff is not considered due to the flat field. Note that ET_i is the combination of surface evaporation and plant transpiration. The oxygen level at any time is the accumulated results of past environmental factors:

$$O_j = \mathcal{F}(f_i, p_i, ET_i, \Delta S_i, \text{ for } i \leq j) \tag{2}$$

Figure 3 shows the principle of how soil water content ($\theta(m^3/m^3)$) and oxygen percentage change within a flooding event. In the beginning, the water saturates the soil and exceeds the surface, the oxygen level drops because the water has squeezed the gas in the soil (t_1 to t_2). This period continues after the valve is turned off. t_3 represents the turning point of the oxygen level, which

In our experimental field, after 10 minutes of turning on the valve, the entire field surface would be flooded, i.e., the soil water content reaches the peak value and the oxygen starts to drop since no gas exchange may occur. The flooding lasts for t_s in total, which can be controlled by our strategy.

Why predicting for a long-term? To choose the best flooding strategy that maximizes the water amount while reducing the risks of oxygen-deficit situations, the consequences of applying water should be accurately predicted, until the next time that the oxygen level recovers to the dry level. Given that this recovery process may be slowed by high atmosphere humidity or interrupted by precipitation, the forecasting window should be long enough to cover the entire recovery process.

2.2 Key Observations

Throughout our five-year in-field Ag-MAR dataset, exemplified in Figure 4, we've revealed two key observations for helping long-term prediction.

OBSERVATION 1. *Soil oxygen exhibits a multi-periodicity pattern, caused by different sources.*

Daily periodicity: This dynamic is subject to the impact of micro-biofunctions, whose behavior is modulated by environmental conditions and our strategic flooding interventions, e.g., elevated temperatures invigorate microbial activities. This leads to a more rapid consumption of oxygen, especially during the day when the temperature is at its zenith and microbial metabolic activities peak. Overall, the multi-periodicity pattern makes the prediction task far from interpretable and straightforward.

Action-triggered periodicity: Post saturation, the oxygen level forms periodic V-curves after saturation, first drops, and then gradually recovers, as shown in Figure 3.

OBSERVATION 2. *There is a strong causality between flooding, weather, and soil oxygen.*

Precipitation affects soil oxygen levels by saturating the soil, displacing air from pore spaces, and reducing aeration, leading to anaerobic conditions that affect plant and microbial respiration. Thus, predicting soil oxygen levels benefits from considering the causality between weather and soil oxygen levels. Fortunately, modern weather forecasting reports that synthesize global atmospheric modeling are becoming more and more reliable, especially in scenarios with sudden and severe rainfall events [32, 73]. They can be used as external clues to facilitate oxygen level inference.

3 LONG-TERM SOIL OXYGEN PREDICTION

In this section, we introduce a long-term time-series forecasting model customized for soil oxygen prediction.

3.1 Model Overview

The model architecture is shown in Figure 5, consisting of three major components: dependency extraction backbone, multi-periodicity block, and causal projection. The iTransformer dependency learning backbone extracts the cross-variate and temporal dependencies. The latent space representation is then passed towards the multi-periodicity block to learn the inter- and intra-periodicity features. To this end, a forecasting result would be produced, which is self-consistent among all variates. Then the partially observable future clues, i.e., future flooding and external weather forecasts, are utilized to calibrate the oxygen prediction.

The historical records can be represented as $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\} \in \mathbb{R}^{T \times N}$, with T time steps and N variates. They are soil oxygen concentration, soil water content, flooding history, and weather records, i.e., air temperature, precipitation, humidity, and wind speed. Besides, some partially observable future clues are included, specifically the flooding plan and weather forecasts for future S time steps $\mathbf{Y} = \{\mathbf{x}'_{T+1}, \dots, \mathbf{x}'_{T+S}\} \in \mathbb{R}^{S \times N'}$. With both inputs, we predict the oxygen in future S time steps $\mathbf{z} = \{\mathbf{x}_{T+1}, \dots, \mathbf{x}_{T+S}\} \in \mathbb{R}^{S \times N}$.

3.2 Multi-Periodicity Block

Periodicity lies inherently in real-world time series [63]. Due to the lack of application contexts, existing works assume that the periodicity would be kept. However, this is not always the case. The action-triggered periodic patterns would change when the action patterns are changed. Therefore, we identify the action-triggered

periodicity in Fast Fourier Transform (FFT) analysis and filter it out in advance.

To harness the full potential of periodicity within other periodic patterns, we adopt a structured approach, TimesBlock, to first perform data segmentation and reorganization [63]. The data is segmented according to the daily frequency bin in FFT results to isolate periodic components, which are then reorganized into a 2D format that aligns their intra-period indexes. In this reorganized structure, each column of the tensor represents a discrete time point within a single period, and each row correlates to the same phase across different periods. This configuration allows the model to differentiate and learn from both intra-period and inter-period variations. This transformation overcomes the inherent limitations of 1D time-series data representation, enhancing the learning of temporal patterns in microbial activities. Inception blocks are then implemented to extract and learn the periodicity from a specific time segment period/frequency.

3.3 Causal Projection with Exogenous Clues

Trustworthy weather forecasts and flooding plans are partially observable factors of the future. To combine their insights with history-inferred results, we compare them with preliminary weather and flooding forecasting outputs from history-based prediction modules. Considering the temporal self-consistency of $\alpha_{T+n} \in \mathbf{x}'_{T+n}$ with all other $\beta_{T+n} \in \mathbf{x}'_{T+n}$, if an external clue $\hat{\alpha}_{T+n}$ is trustworthy, $\Delta\alpha_{T+n} = \hat{\alpha}_{T+n} - \alpha_{T+n}$ can be leveraged to fill the confidence gap between β_{T+n} and the groundtruth if causality holds between them.

The causality between temperature, precipitation, soil water, and oxygen holds intuitively and empirically. Discovering causal relationships within dozens of sequential variates is relatively easy, especially when their physical interpretations are specified. However, the challenge remains in leveraging this causality to enhance time-series forecasting. We categorize the variables into three tiers based on their causal relationships, where upper-tier variables cause the subsequent lower-tier ones, which can't be reversed. Flooding and weather factors are in the top layer, followed by soil moisture in the second layer, and soil oxygen in the bottom layer. After setting the causal layers between variates, we apply Granger causality [54] learning for each pair of variates to learn the parameters:

$$\Delta\beta(t) = \sum_{j=1}^p A_j \Delta\alpha(t-j) + \sum_{j=1}^p A'_j \frac{d\Delta\alpha(t-j)}{dt} + E(t) \quad (3)$$

where $E(t)$ is the residual, A_j and A'_j are linear parameters, we only model the first-derivative causality since the context is clear, e.g., increased soil water to reduce the oxygen diffusion speed, higher temperature leads to faster soil water emission, etc. During inference, we iterate through layers, up to bottom, to calculate $\Delta\beta$ for all β without external clues, i.e., soil water and soil oxygen. These values are utilized to calibrate raw outputs to achieve a new consistency between soil oxygen forecasting and external clues.

3.4 Dependency Extraction Backbone

Understanding dependencies between soil oxygen levels and other variables is crucial for predicting soil oxygen level. While transformer-based methods excel at uncovering dependencies in time series, they are less effective at identifying relationships across different

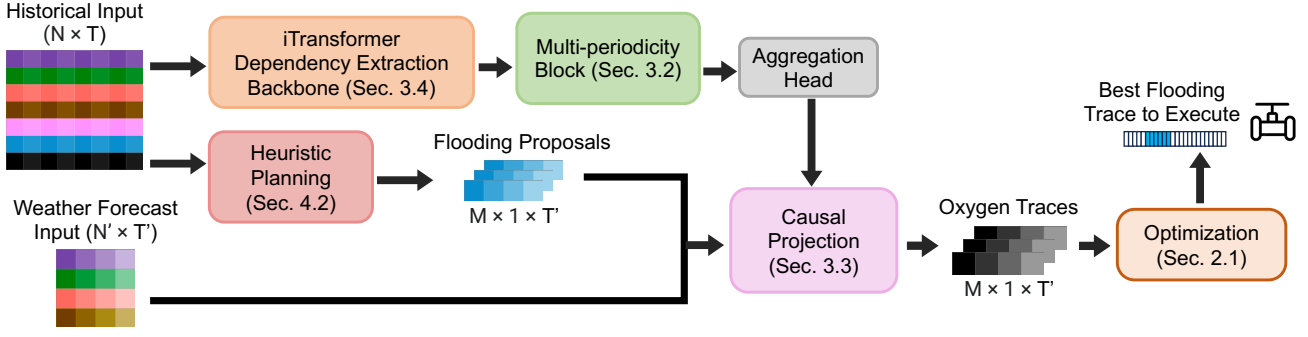


Figure 5: The illustration of long-term soil oxygen prediction architecture and how it facilitates control.

variables. To this end, inversed transformer (iTransformer) [38] is adopted as the backbone to learn the dependencies among variables via inversed layer normalization and attention mechanism.

The layer normalization is applied across time steps rather than features, which preserves the distinct temporal dynamics of each variable, ensuring the learning of the patterns inherent to the data. The feed-forward networks serve to distill complex temporal features from each variate, allowing the attention module to work effectively. The attention mechanism of iTransformer is carefully calibrated to work with the tokenized series of variates. It avoids the traditional composite token format to enhance the model’s ability to map out the dependencies among multiple variables. We modify the decoder to be combined with the multi-periodicity block and to be trained end-to-end.

In-situ Model Update. Considering distribution shifts between regions, fields, and other environmental dynamics, we use the newly collected data to adapt the forecasting model, which enhances the scalability for wide adoption.

4 MODEL PREDICTIVE CONTROL

In this section, we integrate the prediction of soil oxygen into our MPC workflow to guide the recharging actions.

4.1 Workflow

The workflow of MARLP is shown in Figure 2. The heuristic planning generates flooding traces with predefined rules and constraints to propose potential flooding schedules. The long-term oxygen forecasting module then incorporates historical data and weather forecasts to simulate the consequential oxygen trace of all flooding traces. These flooding proposals and the oxygen level predictions are then analyzed by the optimizer according to the optimization objective specified by Eq.1. Once the best flooding plan is identified, the optimizer sends the actions to the actuator. The workflow can be conducted again anytime, not necessarily after the plan is fully executed.

To handle the errors in external clues like weather forecasts, we have the decisions updated every 10 minutes to recalibrate them timely, significantly reducing the impact of unexpected dynamics. This scheduling interval can be adjusted to balance the timeliness and the computation overhead. Note that agile re-scheduling doesn’t conflict with the necessity of long-term forecasting because flooding actions can never be revoked.

4.2 Heuristic Planning

The essential part of MPC is to estimate the optimal flooding trace among all flooding proposals, while not consuming an intolerable amount of computation. Traditionally, this is done by adopting stochastic searching methods such as shooting-based methods[4] or cross-entropy methods[1]. Considering a planning period that may extend beyond 120 hours, with a decision required every 10 minutes, the number of potential action sequences can reach 2^{720} , which is too large for a stochastic searching method to be effective. To address this, we propose schemes based on an understanding of saturating actions in the system. These schemes are derived from two key principles:

(1) **Flooding Duration Constraint:** To ensure the effectiveness of groundwater recharging, once flooding begins, it must continue for at least a minimum duration before ceasing. This constraint can be represented as:

$$\forall t \in T, \quad (F(t-1) = 0 \wedge F(t) = 1) \Rightarrow (\forall \tau \in [t, t + \Delta t_{\min_flood}), F(\tau) = 1) \quad (4)$$

$$\forall t \in T, \quad (F(t) = 1) \Rightarrow (\exists \tau \in [t, t + \Delta t_{\max_flood}), F(\tau) = 0) \quad (5)$$

Here, $F(t)$ is a binary indicator function where $F(t) = 1$, if flooding is occurring at time t , T , is the total observation time period, and Δt_{\min_flood} is the minimum required duration for continuous flooding.

(2) **Idle Period Duration Constraint:** Between two floodings, the idle interval must be long enough that the oxygen can effectively diffuse, otherwise it should not take the interval, but should grasp the chance to flood more. This constraint ensures an adequate drying period for the soil oxygen recovery:

$$\forall t \in T, \quad (F(t-1) = 1 \wedge F(t) = 0) \Rightarrow (\forall \tau \in [t, t + \Delta t_{\min_idle}), F(\tau) = 0) \quad (6)$$

In this case, $F(t) = 0$ indicates an idle (non-flooding) period at time t , and Δt_{\min_idle} represents the minimum required duration for the idle period to allow for soil aeration.

By implementing these constraints, we can mathematically define and enforce the necessary spacing between flooding and idle periods within the MPC framework. By filtering out suboptimal action traces, we significantly reduced the size of the search space



Figure 6: The alfalfa field of 2023 experiments.

to several thousand traces, which can be effectively brute-forced to find the optimal trace.

5 IN-FIELD EXPERIMENT SETUP

As shown in Figures 6 and 13, we build a real experimental testbed to verify the effectiveness of MARLP in practical scenarios. It includes sensory data collection, transmission and decision making modules. These components function as a comprehensive in-field experiment system that enables online evaluation of MARLP, offering more precise and realistic insights than those obtained from simulations or offline evaluations.

Sensor Data Collection: The oxygen and moisture sensors are placed at critical points throughout the facility. These sensors continuously measure real-time oxygen and moisture data, enabling prompt adjustments to maintain optimal oxygen levels.

Sensor Data Transmission: The collected oxygen and moisture data are transmitted to a central server through LoRa networks, ensuring long-range, low-power wireless communication [66, 67, 69]. The implementation details are in Appendix A. We configured the transmission parameters to ensure that all measured sensory data are accurately received by our central server [68, 70].

Inference Server. This server is configured on the cloud [65] with an Intel(R) Core(TM) i9-11900KF CPU and an NVIDIA GEFORCE RTX 3080 Ti GPU. It executes the MPC algorithm and forecasting neural networks with continuously aggregated sensor data and queried weather forecast data from open-source web API [48] in real time.

System Overhead. The power consumption of each sensor node for sensing and communicating is 64 mW, which can be easily covered by a solar panel, mitigating the need to change batteries. The gateway and server could be purchased as a service during flooding seasons. Overall, the system requires less than \$400 to implement and \$50/year to maintain, ensuring easy adoption, even for small farms.

6 EVALUATION

We ask the following questions to evaluate MARLP through in-field experiments and large-scale simulations:

- RQ1 How effective is MARLP in predicting oxygen curves?
- RQ2 How effective is MARLP in control performance?
- RQ3 Can MARLP be effectively generalized across different soil types, plant species, and weather patterns?
- RQ4 How effective is each design component of MARLP?

To answer these questions, we evaluate the predictive capacity on datasets of past years in Section 6.1 and assess the control performance of MARLP during the in-field deployment in Section 6.2.

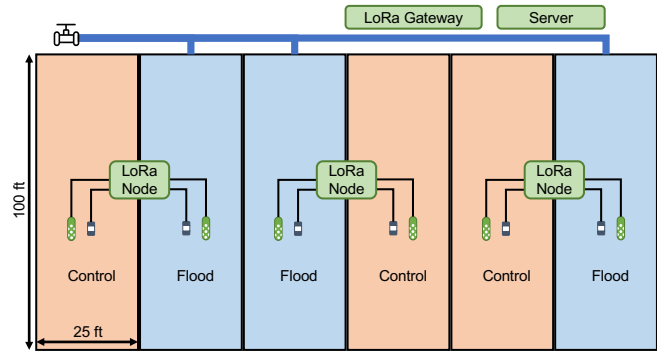


Figure 7: The illustration of in-field deployment.

Then, we investigate the performance under different factors in Section 6.3, followed by the ablation study in Section 6.4.

6.1 Predictive Capability

We choose four recent, representative and highly performed time-series forecasting models as baselines: TimesNet [63], PatchTST [45], DLinear [72], and iTransformer [38], covering convolutional, linear, and transformer-based methodologies. We evaluate all baseline forecasting models on five datasets from five years within three fields. The prediction sequence length is set as 720, which represents five days. Given the sensor reading interval of 10 minutes, the forecasting window is 120 hours, which can fully observe oxygen recovery in most cases. Table 2 shows dataset statistics, including the collection area, period, flooding strategy, and sequence length. Ag-MAR actions lie between growing seasons, which is roughly from January to April for alfalfa in California, US. From 2020 to 2023, the field is flooded with constant intervals, e.g., once a week. Trials in 2024 are controlled by MARLP.

Table 1 reports the MSE, MAE, as well as the mean absolute error of the peak time (PTE) and peak value (PVE) among all forecasting models. The unit of peak time error is an hour. In Table 1, it is evident that MARLP consistently achieves the highest performance (highlighted with red values) when incorporating weather forecast data. This underscores the effectiveness of our long-term soil oxygen prediction algorithm, which is vital for the MPC.

6.2 In-field Control Experiments

We perform the first in-field deployment of the Ag-MAR control scheme, as illustrated in Section 5. The quality of control is evaluated under two factors: **oxygen deficit ratio (ODR)** and **recharging amount**. ODR calculates the ratio of time that the soil oxygen level is below the safety threshold, which should be zero in ideal cases. The best recharging amount may vary according to weather conditions, so the optimal recharging amount can not be asserted, instead, we can only compare the schemes with each other. All in-field experiments are conducted in alfalfa fields, with the safety threshold of oxygen concentration as 10%.


Table 3 compares the control performance of MARLP with the weekly flooded scheme in 2020-2023 as the baseline. The weekly flooded scheme resulted in an average oxygen deficit ratio of 2.72%, while controlling with MARLP yields an ODR of 0.36%. At the same

Table 1: Comparison of predictive performance between MARLP and baselines, input and output sequence length: 720 (120 h). The best and second best performances in each dataset are in red and blue colors.

Model	Metric	agmar2020		agmar2021		agmar2022		agmar2023		agmar2024	
		w/ WF	w/o WF	w/ WF	w/o WF	w/ WF	w/o WF	w/ WF	w/o WF	w/ WF	w/o WF
MARLP (ours)	MSE (↓)	0.331	0.579	0.295	0.351	0.706	0.930	0.309	0.421	1.130	1.735
	MAE (↓)	0.347	0.604	0.428	0.469	0.657	0.792	0.271	0.376	0.848	1.059
	peak_time (↓)	13.409	29.750	28.114	37.063	22.390	32.277	10.967	22.381	44.092	50.385
	peak_value (↓)	0.723	1.116	0.515	0.962	1.034	1.610	0.317	0.521	1.526	1.692
TimesNet[63]	MSE (↓)	0.576	0.603	0.527	0.483	0.883	0.854	1.130	1.025	1.295	1.149
	MAE (↓)	0.603	0.621	0.581	0.563	0.746	0.725	0.924	0.868	0.935	0.841
	peak_time (↓)	33.060	32.522	35.670	45.182	30.537	42.625	40.228	44.371	70.404	42.336
	peak_value (↓)	1.448	1.327	0.789	0.938	1.789	1.763	0.984	1.126	1.528	1.271
PatchTST/64[45]	MSE (↓)	0.465	0.490	0.337	0.374	1.545	0.864	0.342	0.399	1.311	1.480
	MAE (↓)	0.536	0.547	0.456	0.483	0.988	0.736	0.431	0.466	0.913	0.989
	peak_time (↓)	39.127	44.855	45.299	43.963	29.936	26.776	28.925	25.388	55.352	61.752
	peak_value (↓)	1.251	1.233	0.790	0.704	1.728	1.589	0.523	0.581	1.611	1.502
DLinear[72]	MSE (↓)	2.370	2.401	1.814	1.670	4.079	4.944	0.322	0.327	1.598	1.730
	MAE (↓)	1.391	1.409	1.180	1.105	1.828	2.042	0.446	0.450	1.074	1.115
	peak_time (↓)	74.776	74.501	42.604	38.959	56.322	55.657	43.076	41.139	50.871	52.143
	peak_value (↓)	2.237	2.359	1.912	1.781	2.716	2.802	0.756	0.756	1.833	1.903
iTransformer[38]	MSE (↓)	0.537	0.621	0.387	0.565	0.864	0.855	0.477	1.749	1.151	1.732
	MAE (↓)	0.581	0.622	0.490	0.602	0.753	0.754	0.527	0.989	0.853	1.056
	peak_time (↓)	53.088	41.893	47.908	44.291	27.952	30.450	29.978	39.140	46.569	51.133
	peak_value (↓)	1.472	0.938	1.066	0.773	1.885	1.639	0.776	1.214	1.535	1.703

Table 2: The collection settings of Ag-MAR datasets.

Year	2020	2021	2022	2023	2024
Area(ft^2)	590*280	590*280	590*280	284*132	150*100
Flooding	Const.	Const.	Const.	Const.	MARLP
Duration	2/20-4/2	2/12-3/31	1/19-4/8	2/28-4/6	1/19-4/4
Sequence	6086	6902	11455	5389	11001



(a) Weekly flooding scheme fails to adapt to the weather conditions. (b) MARLP can avoid potential oxygen deficits while grasping flooding chances.

Figure 8: An example of how MARLP handles precipitation.**Table 3: Control performance comparison.**

	Const.	MARLP
Oxygen Deficit Ratio	2.72%	0.36%
Recharging Amount (inch per week)	7.640	10.371

time, MARLP increased the recharging amount (inch per week) from 7.64 to 10.371, with a 35.8% improvement.

Figure 8 provides an example of comparison between the weekly flooding scheme and MARLP to show the reason behind the high effectiveness of MARLP. The red and purple bars represent the amount of precipitation and flooding, respectively. Each bar is 10 minutes wide, so the visual areas indicate the total amount of water input. In Figure 8 (a), the weekly-based approach ignores the heavy rain forecast after flooding, resulting in the unexpected oxygen deficit below 10% on day 1 and day 8. At the same time, it misses the opportunity to flood during days 3-6, when there is no rainfall. Figure 8 (b) shows the performance of MARLP. On day 1 and day 2, when the rainfall is negligible, it conducts a 14-hour recharge for 3.1 mm. The oxygen low peak is 10.16%, without exceeding the safety threshold, as depicted by the red dashed line. Knowing the heavy rain on day 4 of the forecast, it discarded aggressive proposals and stopped flooding for 5 days to avoid danger. This shows that MARLP can foresee the consequences of each proposal and avoid risky flooding while making full use of flooding opportunities.

6.3 Large-scale Simulation

To evaluate the generalization capability of all prediction models, we utilize a simulator to replicate a diverse array of factors, including soil types, crop species, and regional climate.² Specifically, we simulate the water content transition in the soil using HYDRUS [56], a classical simulator for soil flux [6]. Then we model the oxygen diffusion process, root respiration, and microbial respiration based on empirical equations [13]. The evaluation targets the oxygen deficit ratio and recharging amount in the simulation.

²We didn't conduct in-field A/B tests due to the limits of field resources.




Figure 9: Control performance for soil types.




Figure 10: Control performance for crop species.

6.3.1 Impact of Soil Types. To check the generalizability of MARLP on different soil, we mimic three representative soil types: sand, loam, and silt according to the soil texture triangle defined by USDA [58]. We simulate each soil texture as a testbed, where all models are evaluated with MPC architecture the same as MARLP. The crops and weather remain the same as in the 2023 in-field experiment. Figure 9(a) shows that MARLP achieves the lowest oxygen deficit ratio than the baselines for all three soil textures, and Figure 9(b) shows that MARLP achieves a high recharging amount at the same time. Although DLinear achieves a higher recharging amount on loam and slit, it performs significantly worse than MARLP in terms of oxygen deficit ratio. The superior robustness based on the causality-aware forecasting of MARLP holds on other soil types. Note that the soil texture has a significant impact on the general trend of the potential recharging amount because soil with relatively high percolation rates can accelerate the atmosphere exchange process.

6.3.2 Impact of Crop Species. We choose walnut trees, grapes and almond trees, with distinct root densities, depths and oxygen-deficit tolerances [47], to assess how the crop diversity influences the control performance. Walnuts and almonds are less tolerant to flooding than alfalfa, hence they are more susceptible to high oxygen-deficit ratio. Figure 10(a) shows that MARLP achieves the lowest oxygen-deficit ratio than the baselines for all three crop species, and Figure 10(b) shows that MARLP achieves a high recharging amount at the same time. Therefore, MARLP has the best generalization capability among all methods.

6.3.3 Impact of Regional Climatic. Although California is at the forefront of Ag-MAR adoption, this practice is becoming increasingly popular in other parts of the world that face similar hydrological challenges, particularly those with a Mediterranean climate. We

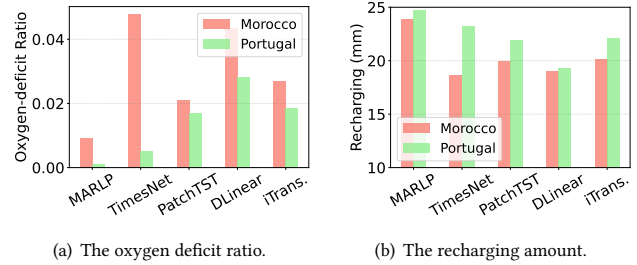


Figure 11: Control performance for climatic features in different regions.

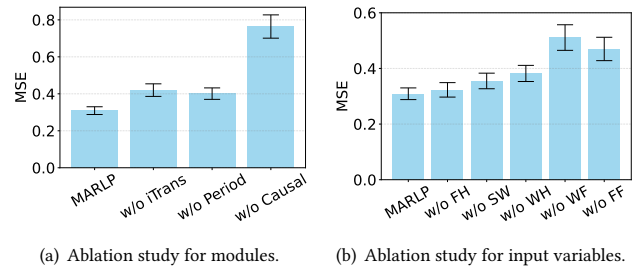


Figure 12: Ablation study for modules and input variables.

broaden its scope to include the High-Atlas region in Morocco [8] and the Algarve region in Portugal [57], incorporating soil and weather data from these diverse global regions that are actively exploring Ag-MAR. The experimental periods for the simulations are both Feb. 28th - Apr. 6th, 2023, aligning with our real-world experiments in California. We use HYDRUS [56] to simulate the soil dynamics of these two sites and use five models to conduct flooding control separately, with the same goal settings. The results are shown in Figure 11. The control performance of all models exhibits drops since the High-Atlas encounters a few sharp precipitations due to high altitude, while Algarve is located by the sea, with stronger daily periodicity. MARLP keeps achieving the best control performance. It underscores the generalizability of MARLP and its potential to be adapted worldwide.

6.4 Ablation Study

In this section, we systematically tested the performance of MARLP by excluding each design module to evaluate their individual effectiveness. Additionally, we examined how each input variable contributes to the system. These tests revealed the necessity of involving these clues and casual strength between the soil oxygen and each of them. We plot the mean and standard deviation of MSE across all datasets.

Design module ablations. We construct ablated versions of the model by substituting the iTransformer with a vanilla Transformer, removing the periodicity block and the causal module. Figure 12(a) illustrates that the performance of each ablated version diminishes, with the version lacking the causal module experiencing the most significant reduction. This decline is because of the pronounced causality inherent in Ag-MAR data.

Impact of Input Features. Figure 12(b) shows the influence of various inputs on the MSE of oxygen forecasting. The abbreviations FH, SW, WH, WF, and FF represent flooding history, soil water content, weather history, weather forecast, and future flooding, respectively. All inputs play a pivotal role in enhancing the final performance, with weather forecasts and future flooding making the most significant contributions. Consequently, optimal forecasting necessitates a synergistic integration of historical data.

Overall, this ablation study affirms the importance of each factor in improving oxygen forecasting and underscores the significance of causal-aware forecasting that incorporates external indicators.

7 DISCUSSIONS AND FUTURE WORKS

Ag-MAR studies. Ag-MAR has been studied in terms of environmental benefits and potential risks [6, 24, 30], the fitness of different soil types, crops, and geological regions. The preliminary results are positive and bring it towards global adoption [42]. However, its current implementation is empirical [23], lacking a standardized and automatic workflow to achieve the best flooding. Based on the oxygen pattern analysis, MARLP provides a systematical solution that can work seamlessly with sensor systems, without the need for expert knowledge. Recent studies also illustrate that the oxygen-deficit tolerance level is temperature dependent [7]. Future works can integrate this feature to enhance the practicality of MARLP.

Multi-variate time-series prediction algorithms. The evolution of multivariate time series prediction algorithms has been marked by significant milestones, starting with the development of the autoregressive integrated moving average (ARIMA) model [9], progressing through the use of recurring neural networks (RNN) [12, 27], and further evolving with the introduction of Transformer [59]. NS-Transformer [39] proposes series stationarization and de-stationary attention to handle the distribution shift and boost the performance of the Transformer on time-series prediction. Dish-TS [21] proposes a general paradigm to alleviate distribution shifts in time series. ESG [71] integrates learning pairwise correlations and temporal dependency in one framework. iTransformer [38] optimizes the extraction of temporal and cross-variate dependency by swapping two modules. Instead of conducting universal forecasting, other works focus on different objectives. TimesNet [63] and DEPTS [22] aims to learn periodic patterns in time series. TSMixer [19] utilizes the MLP architecture to reduce computing overhead. IPOC [11] innovates with ensemble learning, offering real-time adaptable confidence interval predictions. MARLP distinguishes itself with a two-stage method that first predicts based on historical patterns, then utilizes external clues for causality-aware calibration.

Causal discovery and inference for time-series data. Causal relationship widely exists in sequence data, e.g., the classical causality between milk price and butter price [5]. Granger causality is proposed to handle the delay of the causal impact [54]. [28] proposes Bayesian forecasting with time-varying causal models, but it only works for short terms. CASTOR [52] introduces time-lagged links into GNN to enhance Granger causality modeling. REASON [62] utilizes graph neural networks to extract layered causal relationships. CORAL [61] is a framework that automatically updates the root cause analysis model. Instead of discovering and quantifying causality, our solution focuses on applying causality to long-term time-series forecasting with external clues.

Time-series forecasting for real applications. Time-series forecasting is not only a critical question in agricultural practices, but also vital in other industries such as electricity, weather, finance, traffic, and human-computer interaction [37, 40, 41, 53, 60]. For example, RAPT [53] performs a prediction on medical data for the diagnosis of pregnancy complications. ClimODE [60] uses physics-informed neural ODE to simulate global atmosphere dynamics for climate forecasting. We will upgrade the sensor systems to enable neural ODE on Ag-MAR oxygen modeling in future works.

MPC based on long-term forecasting. MPC has been used for application scenarios from the microsecond-level horizon (e.g., embedded system voltage control [36]) to the hour-level horizon (e.g., building control [2, 4], grid control [44]). Existing works utilize MPC for irrigation, but the planning horizon and forecast window are less than one day [17, 18]. Reinforcement learning [15, 16, 55] based control can achieve effective control, but lacks reliability and requires a large amount of data to converge, making it unsuitable for Ag-MAR. To the best of our knowledge, this work is the first MPC work to consider a planning horizon of several days.

Robustness and Efficiency. Robustness is a major concern for applied machine learning [31, 49, 51]. The predicted state trajectory generated as part of the MPC planning process allows the safety check of the trajectory, which offers a higher level of reliability [3]. In future works, we may apply safe guarantee mechanisms like a Gaussian-based uncertainty check [5], to achieve better tradeoffs between reliability and efficiency. Furthermore, given the typically extensive search spaces involved in real-world, long-term forecasting control, efficiency becomes a pivotal aspect. Unlike RL, which can separate agent models from environmental models to improve efficiency [25, 33, 34], time series modeling offers limited scope to balance performance with efficiency. However, recent advances in state space models (SSM) have significantly mitigated computational burdens [26]. Future research could explore the application of SSMs in scenarios where real-time requirements are critical.

8 CONCLUSION

This paper introduces MARLP, a model predictive control system for Ag-MAR, employing heuristic planning and a long-term oxygen forecasting module to optimize groundwater recharge and soil oxygen levels. Benefiting from a causality-aware forecasting model, MARLP effectively manages environmental variables in real-time, enhancing water use efficiency in agriculture. The successful deployment of MARLP, which reduces the oxygen deficit ratio and improves the total amount of applied water, showcases the system's potential in precision agriculture and sustainable resource management.

ACKNOWLEDGMENTS

This work was supported in part by a UC Merced Fall 2023 Climate Action Seed Competition grant, and a UC Merced Spring 2023 Climate Action Seed Competition grant. Kang Yang was supported by a financial assistance award approved by the Economic Development Administration's Farms Food Future program. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- [1] Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, and J Zico Kolter. 2018. Differentiable mpc for end-to-end planning and control. *Advances in neural information processing systems* 31 (2018).
- [2] Zhiyu An, Xianzhong Ding, and Wan Du. 2024. Go Beyond Black-box Policies: Rethinking the Design of Learning Agent for Interpretable and Verifiable HVAC Control. *arXiv preprint arXiv:2403.00172* (2024).
- [3] Zhiyu An, Xianzhong Ding, and Wan Du. 2024. Reward Bound for Behavioral Guarantee of Model-based Planning Agents. *arXiv preprint arXiv:2402.13419* (2024).
- [4] Zhiyu An, Xianzhong Ding, Arya Rathee, and Wan Du. 2023. CLUE: Safe Model-Based RL HVAC Control Using Epistemic Uncertainty Estimation. In *ACM BuildSys*.
- [5] Titus O Awokuse and Xiaohong Wang. 2009. Threshold effects and asymmetric price adjustments in US dairy markets. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* 57, 2 (2009), 269–286.
- [6] Khaled M Bali, Abdelmoneim Zakaria Mohamed, Sultan Begna, Dong Wang, Daniel Putnam, Helen E Dahlke, and Mohamed Galal Eltarabily. 2023. The use of HYDRUS-2D to simulate intermittent Agricultural Managed Aquifer Recharge (Ag-MAR) in Alfalfa in the San Joaquin Valley. *Agricultural Water Management* 282 (2023), 108296.
- [7] AL Barta and RM Sulc. 2002. Interaction between waterlogging injury and irradiance level in alfalfa. *Crop science* 42, 5 (2002), 1529–1534.
- [8] Houssne Bouimouass, Sarah Tweed, Vincent Marc, Younes Fakir, Hamza Sahraoui, and Marc Leblanc. 2024. The importance of mountain-block recharge in semiarid basins: An insight from the High-Atlas, Morocco. *Journal of Hydrology* (2024), 130818.
- [9] George EP Box and Gwilym M Jenkins. 1968. Some recent advances in forecasting and control. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 17, 2 (1968), 91–109.
- [10] California Department of Water Resources. 2014. Sustainable Groundwater Management Act (SGMA). <https://water.ca.gov/programs/groundwater-management/sgma-groundwater-management>.
- [11] Jiaodong Chen, Yang Luo, Xiuqi Huang, Fuxin Jiang, Yangguang Shi, Tieying Zhang, and Xiaofeng Gao. 2023. IPOC: An Adaptive Interval Prediction Model based on Online Chasing and Conformal Inference for Large-Scale Systems. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 202–212.
- [12] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014).
- [13] FJ Cook and JH Knight. 2003. Oxygen transport to plant roots: modeling for physical understanding of soil aeration. *Soil Science Society of America Journal* 67, 1 (2003), 20–31.
- [14] Helen E Dahlke, Andrew G Brown, Steve Orloff, Daniel H Putnam, and Toby O'Geen. 2018. Managed winter flooding of alfalfa recharges groundwater with minimal crop damage. *California Agriculture* 72, 1 (2018).
- [15] Xianzhong Ding, Alberto Cerpa, and Wan Du. 2023. Exploring deep reinforcement learning for holistic smart building control. *ACM Transactions on Sensor Networks* (2023).
- [16] Xianzhong Ding, Alberto Cerpa, and Wan Du. 2023. Multi-zone HVAC Control with Model-Based Deep Reinforcement Learning. *arXiv preprint arXiv:2302.00725* (2023).
- [17] Xianzhong Ding and Wan Du. 2022. DRLIC: Deep Reinforcement Learning for Irrigation Control. In *ACM/IEEE IPSN*.
- [18] Xianzhong Ding and Wan Du. 2023. Optimizing irrigation efficiency using deep reinforcement learning in the field. *ACM Transactions on Sensor Networks* (2023).
- [19] Vijay Ekambaram, Arindam Jati, Nam Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. TSMixer: Lightweight MLP-Mixer Model for Multivariate Time Series Forecasting. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*. 459–469.
- [20] Alvar Escrivá-Bou, Brian Gray, Sarge Green, Thomas Harter, Richard Howitt, Duncan MacEwan, and N Seavy. 2017. Water Stress and a Changing San Joaquin Valley. *Public Policy Institute of California*. https://www.ppic.org/content/pubs/report/R_0317EHR.pdf (2017).
- [21] Wei Fan, Pengyang Wang, Dongkun Wang, Dongjie Wang, Yuanchun Zhou, and Yanjie Fu. 2023. Dish-ts: a general paradigm for alleviating distribution shift in time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [22] Wei Fan, Shun Zheng, Xiaohan Yi, Wei Cao, Yanjie Fu, Jiang Bian, and Tie-Yan Liu. 2021. DEPTS: Deep Expansion Learning for Periodic Time Series Forecasting. In *International Conference on Learning Representations*.
- [23] Yonatan Ganot and Helen E Dahlke. 2021. A model for estimating Ag-MAR flooding duration based on crop tolerance, root depth, and soil texture data. *Agricultural Water Management* 255 (2021), 107031.
- [24] Yonatan Ganot and Helen E Dahlke. 2021. Natural and forced soil aeration during agricultural managed aquifer recharge. *Vadose Zone Journal* 20, 3 (2021), e20128.
- [25] Kevin Gmelin, Shikhar Bahl, Russell Mendonca, and Deepak Pathak. 2023. Efficient RL via Disentangled Environment and Agent Representations. In *Proceedings of the 40th International Conference on Machine Learning*.
- [26] Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752* (2023).
- [27] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [28] Biwei Huang, Kun Zhang, Mingming Gong, and Clark Glymour. 2019. Causal discovery and forecasting in nonstationary environments with state-space models. In *International conference on machine learning*. PMLR, 2901–2910.
- [29] Scott Jasechko, Hansjörg Seybold, Debra Perrone, Ying Fan, Mohammad Shamsudduha, Richard G Taylor, Othman Fallatah, and James W Kirchner. 2024. Rapid groundwater decline and some cases of recovery in aquifers globally. *Nature* 625, 7996 (2024), 715–721.
- [30] George Kourakos, Helen E Dahlke, and Thomas Harter. 2019. Increasing groundwater availability and seasonal base flow through agricultural managed aquifer recharge in an irrigated basin. *Water Resources Research* 55, 9 (2019), 7464–7492.
- [31] Nathaniel Lahn, Sharath Raghvendra, and Kaiyi Zhang. 2024. A Combinatorial Algorithm for Approximating the Optimal Transport in the Parallel and MPC Settings. *Advances in Neural Information Processing Systems* 36 (2024).
- [32] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirmberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. 2023. Learning skillful medium-range global weather forecasting. *Science* (2023), eadi2336.
- [33] Guangchen Lan, Dong-Jun Han, Abolfazl Hashemi, Vaneet Aggarwal, and Christopher G Brinton. 2024. Asynchronous federated reinforcement learning with policy gradient updates: Algorithm design and convergence analysis. *arXiv preprint arXiv:2404.08003* (2024).
- [34] Guangchen Lan, Han Wang, James Anderson, Christopher Brinton, and Vaneet Aggarwal. 2024. Improved Communication Efficiency in Federated Natural Policy Gradient via ADMM-based Gradient Updates. *Advances in Neural Information Processing Systems* 36 (2024).
- [35] Elad Levintal, Maribeth L Kniffin, Yonatan Ganot, Nisha Marwaha, Nicholas P Murphy, and Helen E Dahlke. 2023. Agricultural managed aquifer recharge (Ag-MAR)—a method for sustainable groundwater management: A review. *Critical Reviews in Environmental Science and Technology* 53, 3 (2023), 291–314.
- [36] Eyke Liegmann, Petros Karamanakos, and Ralph Kennel. 2021. Real-time implementation of long-horizon direct model predictive control on an embedded system. *IEEE Open Journal of Industry Applications* 3 (2021), 1–12.
- [37] Shun Liu, Kexin Wu, Chufeng Jiang, Bin Huang, and Danqing Ma. 2023. Financial time-series forecasting: Towards synergizing performance and interpretability within a hybrid machine learning approach. *arXiv preprint arXiv:2401.00534* (2023).
- [38] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. 2023. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. *arXiv preprint arXiv:2310.06625* (2023).
- [39] Yong Liu, Haixu Wu, Jianmin Wang, and Mingsheng Long. 2022. Non-stationary Transformers: Exploring the Stationarity in Time Series Forecasting. (2022).
- [40] Yilin Liu, Shijia Zhang, and Mahanth Gowda. 2021. When video meets inertial sensors: Zero-shot domain adaptation for finger motion analytics with inertial sensors. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*. 182–194.
- [41] Yilin Liu, Shijia Zhang, Mahanth Gowda, and Srihari Nelakuditi. 2022. Leveraging the properties of mmwave signals for 3d finger motion tracking for interactive iot applications. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 6, 3 (2022), 1–28.
- [42] Nisha Marwaha, George Kourakos, Elad Levintal, and Helen E Dahlke. 2021. Identifying agricultural managed aquifer recharge locations to benefit drinking water supply in rural communities. *Water Resources Research* 57, 3 (2021), e2020WR028811.
- [43] Nicholas Paul Murphy. 2022. *Examining Nitrate Leaching Potential and Nitrogen Cycle Dynamics under Agricultural Managed Aquifer Recharge in the Central Valley of California*. University of California, Davis.
- [44] James R Nelson and Nathan G Johnson. 2020. Model predictive control of microgrids for real-time ancillary service market participation. *Applied Energy* 269 (2020), 114963.
- [45] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2022. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*.
- [46] Richard G Niswonger, Eric D Morway, Enrique Triana, and Justin L Huntington. 2017. Managed aquifer recharge through off-season irrigation in agricultural regions. *Water Resources Research* 53, 8 (2017), 6970–6992.
- [47] AT O'Geen, Matthew BB Saal, Helen E Dahlke, David A Doll, Rachel B Elkins, Allan Fulton, Graham E Fogg, Thomas Harter, Jan W Hopmans, Chuck Ingels, et al. 2015. Soil suitability index identifies potential areas for groundwater banking on agricultural lands. *California Agriculture* 69, 2 (2015).
- [48] open-meteo. 2014. Free Weather API. <https://open-meteo.com/>.

- [49] Abhijeet Phatak, Sharath Raghvendra, Chittaranjan Tripathy, and Kaiyi Zhang. 2023. Computing all optimal partial transports. In *International Conference on Learning Representations*.
- [50] Jiangxiao Qiu, Samuel C Zipper, Melissa Motew, Eric G Booth, Christopher J Kucharik, and Steven P Loheide. 2019. Nonlinear groundwater influence on biophysical indicators of ecosystem services. *Nature Sustainability* 2, 6 (2019), 475–483.
- [51] Sharath Raghvendra, Pouyan Shirzadian, and Kaiyi Zhang. 2024. A New Robust Partial p -Wasserstein-Based Metric for Comparing Distributions. *arXiv preprint arXiv:2405.03664* (2024).
- [52] Abdellah Rahmani and Pascal Frossard. 2023. Castor: Causal Temporal Regime Structure Learning. *arXiv preprint arXiv:2311.01412* (2023).
- [53] Houxing Ren, Jingyuan Wang, Wayne Xin Zhao, and Ning Wu. 2021. Rapt: Pre-training of time-aware transformer for learning robust healthcare representation. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*. 3503–3511.
- [54] Anil Seth. 2007. Granger causality. *Scholarpedia* 2, 7 (2007), 1667.
- [55] Zhihao Shen, Kang Yang, Wan Du, Xi Zhao, and Jianhua Zou. 2019. DeepAPP: A Deep Reinforcement Learning Framework for Mobile Application Usage Prediction. In *ACM SenSys*.
- [56] J Simnek, M Sejna, and M Th Van Genuchten. 1999. *The HYDRUS-2D software package for simulating the two-dimensional movement of water, heat, and multiple solutes in variably-saturated media: Version 2.0*. US Salinity Laboratory, Agricultural Research Service.
- [57] Kath Standen, Luís Costa, Rui Hugman, and José Paulo Monteiro. 2023. Integration of Managed Aquifer Recharge into the Water Supply System in the Algarve Region, Portugal. *Water* 15, 12 (2023), 2286.
- [58] NRCS USDA. 1999. United States department of agriculture. *Natural Resources Conservation Service. Plants Database*. <http://plants.usda.gov> (accessed in 2000) (1999).
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [60] Yogesh Verma, Markus Heinonen, and Vikas Garg. 2023. ClimODE: Climate Forecasting With Physics-informed Neural ODEs. In *The Twelfth International Conference on Learning Representations*.
- [61] Dongjie Wang, Zhengzhang Chen, Yanjie Fu, Yanchi Liu, and Haifeng Chen. 2023. Incremental causal graph learning for online root cause analysis. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2269–2278.
- [62] Dongjie Wang, Zhengzhang Chen, Jingchao Ni, Liang Tong, Zheng Wang, Yanjie Fu, and Haifeng Chen. 2023. Interdependent Causal Networks for Root Cause Localization. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5051–5060.
- [63] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *International Conference on Learning Representations (ICLR)*.
- [64] Xianjin Xia, Qianwu Chen, Ningning Hou, Yuanqing Zheng, and Mo Li. 2023. XCopy: Boosting Weak Links for Reliable LoRa Communication. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*. 1–15.
- [65] Yifei Xu, Yuning Chen, Xumiao Zhang, Xianshang Lin, Pan Hu, Yunfei Ma, Songwu Lu, Wan Du, Zhuoqing Mao, Ennan Zhai, et al. 2024. CloudEval-YAML: A Practical Benchmark for Cloud Configuration Generation. *Proceedings of Machine Learning and Systems* 6 (2024), 173–195.
- [66] Kang Yang, Yuning Chen, Xuanren Chen, and Wan Du. 2023. Link quality modeling for lora networks in orchards. In *Proceedings of the 22nd International Conference on Information Processing in Sensor Networks*. 27–39.
- [67] Kang Yang, Yuning Chen, and Wan Du. 2024. OrchLoc: In-Orchard Localization via a Single LoRa Gateway and Generative Diffusion Model-based Fingerprinting. In *ACM MobiSys*.
- [68] Kang Yang and Wan Du. 2022. LLDPC: A Low-Density Parity-Check Coding Scheme for LoRa Networks. In *ACM SenSys*.
- [69] Kang Yang and Wan Du. 2024. A Low-Density Parity-Check Coding Scheme for LoRa Networking. *ACM Transactions on Sensor Networks* (2024).
- [70] Kang Yang, Miaomiao Liu, and Wan Du. 2024. RALoRa: Rateless-Enabled Link Adaptation for LoRa Networking. *IEEE/ACM Transactions on Networking* (2024), 1–16.
- [71] Junchen Ye, Zihan Liu, Bowen Du, Leilei Sun, Weimiao Li, Yanjie Fu, and Hui Xiong. 2022. Learning the evolutionary and multi-scale graph structure for multivariate time series forecasting. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*.
- [72] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting?. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 11121–11128.
- [73] Yuchen Zhang, Mingsheng Long, Kaiyuan Chen, Lanxiang Xing, Ronghua Jin, Michael I Jordan, and Jianmin Wang. 2023. Skilful nowcasting of extreme precipitation with NowcastNet. *Nature* 619, 7970 (2023), 526–532.

A SENSOR NODE DEPLOYMENT



Figure 13: The illustration of the solar-powered sensor node.

The sensor node is powered by solar energy. As shown in Figure 13, it consists of several key components. The Arduino Uno serves as the central controller, managing the operations of sensor readings and LoRa signal modulation. It is connected to KE-25 oxygen sensors and IRROMETER Watermark 200SS soil moisture sensors, which are deployed in the soil. The InAir9B LoRa radio is connected to the Arduino Uno via a relay board, which enables low-power long-range communication. The rechargeable battery along with the solar charger ensures minimum maintenance efforts. Key components are hosted in the waterproof box to protect them from damage and fast aging caused by environmental factors. The spreading factor (SF) [64] of the LoRa transmission is 8.