Dynamic critical groundwater depth as a predictor of irrigation-intensified salinization in lowland Hungary

Zsolt Zoltán Fehér

University of Debrecen, Institute of Water and Environmental Management,
Egyetem tér 1, H-4032 Debrecen, Hungary

Correspondence: feher.zsolt@agr.unideb.hu

Dynamic critical groundwater depth as a predictor of irrigation-intensified salinization in lowland Hungary

Zsolt Zoltán Fehér

University of Debrecen, Institute of Water and Environmental Management, Egyetem tér 1, H-4032 Debrecen, Hungary
Correspondence: feher.zsolt@agr.unideb.hu

Abstract

Shallow groundwater in continental lowland environments sustains upward capillary fluxes that transport dissolved salts to the land surface. However, the depth below which this capillary-driven contribution becomes negligible, often parameterized as the extinction depth in groundwater model ET packages, has been treated as a static, soil-dependent parameter. We argue that salinization risk is governed by the dynamic overlap between the seasonal cycles of groundwater level (GWL) and evapotranspiration (ET), and develop a time-varying risk index based on the proximity of the water table to a dynamic, GWL-dependent critical depth, and capillary fringe. The framework is applied to a 4,041 km² sub-region of the TIKEVIR system in southeastern Hungary using 74 years (1951–2025) of gridded GWL observations and FORESEE-HUN climate forcing. The domain-averaged seasonal cycle (shallowest in April at 2.50 m, deepest in October at 3.04 m, amplitude 0.54 m) defines a critical salinization window in March–June when the water table resides within the critical zone ($z_{GWL} < d_{crit}$) under rising evaporative demand. A ConvLSTM network with eight input channels, trained on 1971–2019 data and evaluated on 2020–2024, achieves NSE = 0.839 and RMSE = 0.416 m for 1-month-ahead GWL forecasting at 1 km resolution, with 27–40 % RMSE reduction over persistence at 1–12-month lead times. Monte Carlo dropout analysis has verified a narrow prediction uncertainty, with a mean risk credible interval width of 0.014 (dimensionless). By integrating the forecast with the risk index, it was determined that 19.2% of the domain, equivalent to 776 km², is persistently at high or critical risk, predominantly located in discharge zones. A notable secular deepening of the water table, measured at 0.051 m per decade (p = 0.032), has slightly diminished the domain-averaged risk; however, the feedback loop between irrigation and salinization may counterbalance this trend. This methodology holds potential for application to similar lowland aquifer systems where the groundwater level (GWL) and evapotranspiration (ET) cycles are phase-offset.

1 Introduction

Shallow groundwater systems in semi-arid and continental lowlands are controlled by coupling between atmospheric forcing and subsurface responses. Where the water table is within the capillary fringe, upward moisture fluxes from evapotranspiration (ET) transport salt to the root zone and surface (Brunner et al., 2008; Mádl-Szőnyi and Tóth, 2009). The extinction depth, where the water table stops contributing to ET, separates active and passive groundwater conditions. Despite its importance, extinction depth has been treated as a static parameter in hydrological models (Shah et al., 2007; Soylu et al., 2011), ignoring its dynamic nature under varying conditions.

The extinction depth concept was formalized in groundwater modelling through the ET packages of MODFLOW (McDonald and Harbaugh, 1988; Banta, 2000), where groundwater ET decreases linearly or in a segmented-linear fashion from a maximum rate at the land surface to zero at the extinction depth. Shah et al. (2007) advanced this understanding by demonstrating, through numerical experiments with HYDRUS-1D, that the decline of groundwater ET with depth to water table follows an exponential rather than linear decay, and that extinction depths vary substantially with both soil texture and vegetation cover – ranging from approximately 2 m in sands to over 8 m in fine-textured soils. These findings have been widely adopted as lookup tables for regional groundwater models. However, two fundamental limitations persist. First, the extinction depth is conventionally parameterized as a time-invariant quantity, assigned once per model cell, and held constant throughout the simulation. Second, the phase relationship between the seasonal cycles of ET demand and groundwater level (GWL) fluctuation, which determines when and for how long the water table occupies the critical zone above the extinction depth, has received remarkably little attention.

The interaction between ET and GWL is characterized by two quasi-periodic signals with a pronounced phase lag. In continental climates, such as that of the Carpathian Basin, the GWL typically reaches its seasonal maximum in March–April, following snowmelt and spring recharge, and its minimum in October after the growing season (Szalai et al., 2004; Major, 2002). Garamhegyi et al. (2018) demonstrated through wavelet analysis of 216 monitoring wells across the Great Hungarian Plain that shallow GWL fluctuations exhibit climate-driven periodicities at annual, inter-annual, and multi-decadal time scales, and that the absence of annual periodicity is a diagnostic indicator of drought conditions. Illyés et al. (2023) found 7–10 month lag times between rainfall and groundwater responses in eastern Hungarian lowlands. ET demand peaks in July-August with highest solar radiation. This three to four month offset creates a period from March-June when the water table remains above extinction depth as ET increases. During this time, capillary fluxes transport dissolved salts from saturated zones to the topsoil, accumulating near the surface through evaporation. Therefore, the salinization risk is not a simple function of the mean annual GWL or mean annual ET, but emerges from the temporal overlap of these two forcing cycles, a dynamic phenomenon that static extinction depth parameterizations cannot capture.

Terminology note. In the Hungarian literature on shallow groundwater and soil salinization, the management threshold *kritikus talajvízszint* ("critical depth of the water table") has long been used as an operational indicator of the depth above which groundwater contributes to harmful salt accumulation; Szabolcs et al. (1969) explicitly define a "critical depth" and tabulated its values (2.0–4.5 m) as a function of soil water regime, groundwater salinity, and irrigation conditions for the Tisza irrigation systems. Because the international term *extinction depth* is commonly associated with the ET–water-table parameterization of MODFLOW and HYDRUS (Shah et al., 2007), we avoided applying it to our proposed threshold. Throughout this paper, *extinction depth* refers exclusively to the classical, static ET-cessation depth of Shah et al. (2007) and the MODFLOW ET packages, whereas *the critical depth* (d_{crit}) denotes the dynamic, GWL-dependent threshold used in our risk index (Eq. 1).

Both serve as threshold-type indicators for when groundwater–surface coupling is expected to be weak; however, d_{crit} is used here as an operational, time-varying threshold that evolves with the water table position and is calibrated against the Hungarian empirical tradition, rather than as a fixed soil-physical parameter derived from a single-column model.

Secondary salinization challenges agriculture in the Carpathian Basin's lowlands. The Hungarian Great Plain supports 50% of Hungary's agriculture on Quaternary sediments with shallow aquifers. The eastern lowlands, managed as the TIKEVIR region, are vulnerable to groundwater salinization. In the Great Hungarian Plain, meteoric freshwater systems coexist with overpressured saline systems containing 10,000–38,000 mg L^{-1} total dissolved solids, with vertical flow explaining soil salinization patterns. The TIKEVIR region faces challenges from flat topography, poor drainage, high evaporation, and groundwater decline since the 1970s (Pálfai, 1994; Szalai et al., 2004).

Understanding salinization dynamics has become urgent due to recent droughts and increasing irrigation demands. The 2022 and 2024 droughts caused groundwater declines of nearly 2 m from 2017 levels across Hungarian lowlands (Tran et al., 2022; Fehér and Rakonczai, 2019). This decline affects when water tables intersect the critical zone, though the drought-salinization relationship is complex: extreme drought may temporarily suppress salinization until recharge events occur. Rising temperatures are increasing irrigation demands across the Plain (Mezősi et al., 2016), creating a feedback where irrigation raises water tables and activates salt transport in agricultural areas. This irrigation-salinization dynamic challenges conventional risk assessments and requires more sophisticated modeling approaches.

Recent advances in deep learning have opened new avenues for spatiotemporal hydrological prediction. Convolutional long short-term memory (ConvLSTM) networks, integrating spatial feature extraction with temporal sequence modelling, have shown promise in gridded hydrological forecasting (Shi et al., 2015). Deep learning models can capture spatial heterogeneity from geological variability and temporal dynamics driven by climate forcing and abstraction (Sun et al., 2019; Wunsch et al., 2021). Graph-based neural networks have improved performance by encoding spatial connectivity of monitoring-well networks (Gai et al., 2023). However, applying spatiotemporal deep learning to predict when the water table will intersect the dynamic critical depth envelope remains unexplored. Existing studies have focused on level prediction without coupling forecasts to physical thresholds governing surface salt accumulation.

We propose a framework coupling spatiotemporal GWL forecasting with a dynamic critical depth criterion from hydroclimatic data. The framework was applied to a 4,041 km² sub-region of the TIKEVIR system in southeastern Hungary, using three data sources: (i) groundwater level observations from the Hungarian monitoring network, interpolated to 1 km grid by Sequential Gaussian Simulation over 1951–2025, (ii) daily meteorological fields from FORESEE-HUN database (Kern et al., 2024) including RCP 4.5 climate projections to 2050 for ET computation, and (iii) an empirical GWL-dependent formulation for the critical depth that captures stronger capillary flux at shallower water table positions.A dynamic salinization risk index was formulated as a function of the proximity of the GWL to the critical depth and capillary zone height. A ConvLSTM model was trained on historical GWL and meteorological records to forecast GWL fields, and the risk index was projected forward through the forecast horizon using autoregressive multi-step inference driven by climate projections, enabling the identification of emerging salinization hotspots under future climate scenarios.

The specific objectives of this study were as follows:

1. To develop and validate a time-varying formulation of the critical groundwater table depth that accounts for the phase relationship between seasonal ET demand and GWL fluctuation, replacing the conventional static extinction depth parameterization.

2. To construct and evaluate a ConvLSTM-based spatiotemporal model for monthly GWL forecasting over the TIKEVIR region, using gridded meteorological forcing and historical GWL observations as inputs.

3. We coupled the GWL forecast with the dynamic critical depth criterion to produce spatially explicit salinization risk projections at a monthly resolution and evaluate the implications of irrigation–salinization feedback under future climate scenarios.

4. The coupled framework was implemented as an interactive decision-support dashboard that enables water managers and agricultural planners to identify emerging salinization risk zones and evaluate intervention strategies in near-real time.

This study contributes to the growing body of literature on physically informed machine learning in hydrology (Reichstein et al., 2019; Nearing et al., 2021) by demonstrating how DL-based predictions can be coupled with process-based thresholds to address practical environmental management problems. To our knowledge, this is among the first applications of spatiotemporal deep learning to dynamic salinization risk forecasting in a lowland aquifer system and among the first to explicitly treat irrigation–salinization feedback within a coupled DL forecasting framework.

2 Study area and data

2.1 The TIKEVIR region

The Tisza–Körös Valley Integrated Water Management Region (TIKEVIR) is an extensive water management system in the southeastern sector of the Hungarian Great Plain (Alföld), encompassing the catchments of the Tisza, Körös, and Berettyó rivers. The present study focuses on a 4,041 km² sub-region of the broader TIKEVIR system (Fig. 1), bounded by the coordinates 47.2–48.1°N and 20.9–21.8°E (790,875–855,512 mE and 212,383–312,006 mN in EPSG:23700), selected for its dense monitoring well coverage and representative range of hydrogeological conditions. Elevations range from 78 to 112 m above sea level (a.s.l.), with the terrain surface sloping gently from the northeast to the southwest at gradients typically below 0.5 ‰. This exceptionally flat topography results in sluggish natural drainage, extensive waterlogging susceptibility during wet periods, and high sensitivity to groundwater table fluctuations (Pálfai, 1994).

The climate is continental, with warm, dry summers and cold winters. Based on the FORESEE gridded meteorological dataset (1971–2024), the mean annual temperature across the region is 10.4 °C, ranging from monthly averages of −1.2 °C in January to 21.3 °C in July. The mean annual precipitation is 518.1 mm, with a primary maximum in late spring to early summer (May–June) and a secondary maximum in October–November. Potential evapotranspiration, estimated using the Penman–Monteith method (Allen et al., 1998), averages 650–750 mm yr⁻¹, consistently exceeding annual precipitation and producing a negative climatic water balance of 100–200 mm yr⁻¹ during the growing season (April–September). This structural moisture deficit makes the region intrinsically vulnerable to drought, and Standardized Precipitation Index (SPI) analyses have documented severe to extreme drought episodes in 2000–2003, 2011–2012, 2022, and 2024 (Szalai et al., 2004; Fehér and Rakonczai, 2019; Tran et al., 2022).

The geological setting is dominated by Quaternary alluvial and aeolian sediments deposited in the Pannonian Basin, which overlie a thick sequence of Neogene lacustrine and fluvial formations. The shallow subsurface (0–50 m) comprises interbedded sands, silts, and clays of fluvial and aeolian origins, with lateral and vertical heterogeneity controlled by Pleistocene and Holocene fluvial migration patterns. The unconfined aquifer system has a saturated thickness of 10–30 m and is hydraulically connected to the surface-drainage network. The depth to the water table typically ranges from 1 to 6 m below the ground surface; statistically, 95 % of the spatial grid limits fall within depths of approximately 0.5 m to 6.7 m, enforcing physical plausibility by excluding mathematical interpolation artifacts (e.g., artificial artesian pressures above ground) generated during spatial cross-validation. The domain-wide mean depth was stable at 2.94 m. Hydraulic conductivity values reported from pumping tests and grain-size analyses range from 0.5 to 15 m d⁻¹, reflecting heterogeneous alluvial lithology (Major, 2002; Szabó et al., 2024).

Land use is predominantly agricultural, with arable cropland covering approximately 70 % of the total area, dominated by cereals (wheat and maize), sunflowers, and rapeseed. Grasslands and pastures account for 15 %, whereas forests, wetlands, and built-up areas constitute the remainder. The Debrecen urban agglomeration in the northeastern part of the region is undergoing rapid industrialization and urbanization, introducing additional pressure on both groundwater quantity and quality. Salt-affected soils (solonchak, solonetz) are widespread in the lower-lying discharge zones flanking the Tisza and Körös floodplains, where shallow groundwater and ascending saline water from deeper basin formations converge (Várallyay, 1967; Mádl-Szőnyi and Tóth, 2009; Tóth and Várallyay, 1980).

2.2 Datasets

Three principal gridded datasets were employed in this study, supplemented with national monitoring records and European-scale soil property grids. Table 1 summarizes the key characteristics of the datasets used in this study.

[Table 1: Summary of datasets used in this study]

Dataset	Source	Variables	Spatial resolution	Temporal resolution	Period	Format
FORESEE-HUN v1.0	ELTE / Kern et al. (2024)	T_{mean}, T_{min}, T_{max}, P, RH, u_2, R_s, R_n	0.1° (~10 km)	Daily	1971–2050 (incl. RCP 4.5)	NetCDF
GWL monitoring network	OVF (General Directorate of Water Management)	Depth to water table	Point (1492 wells total)	Daily/Weekly/Monthly	1950–2024	CSV / database
ESDAC topsoil physical properties	JRC / ESDAC (Ballabio et al., 2016)	Clay, sand, silt (%), AWC, USDA texture class	500 m	Static	-	GeoTIFF
EU-SoilHydroGrids	JRC / ESDAC (Tóth et al., 2017)	van Genuchten α, n, θ_r, θ_s, K_{sat}	250 m	Static	-	GeoTIFF
AGROTOPO	MTA ATK / Hungarian national	Soil texture, genetic soil type	1:100 000	Static	-	Shapefile
CORINE Land Cover 2018	Copernicus	Land use / land cover	100 m	Static (2018)	-	GeoTIFF
EU-DEM v1.1	Copernicus	Elevation	25 m	Static	-	GeoTIFF

2.2.1 Meteorological data: the FORESEE-HUN database

Daily gridded meteorological fields were obtained from the FORESEE-HUN v1.0 database (Kern et al., 2024) produced by the Department of Meteorology, Eötvös Loránd University (ELTE), Budapest. The database provides homogenized, quality-controlled interpolations of station observations onto a regular 0.1° × 0.1° latitude–longitude grid covering Hungary for the period 1971–2024. The variables used in this study included daily mean, minimum, and maximum temperature (°C), precipitation (mm d⁻¹), relative humidity (%), wind speed at 2 m height (m s⁻¹), and global radiation (MJ · m⁻² · d⁻¹). The FORESEE dataset employs a 365-day calendar (omitting December 31 in leap years), which was considered for all temporal processing. These fields served as forcing data for reference ET computation, predictors for the deep learning model, and the basis for computing antecedent moisture conditions.

2.2.2 Groundwater level observations

Groundwater level (GWL) records were compiled from the national shallow groundwater monitoring network operated by the General Directorate of Water Management (OVF) in Hungary. A total of 1492 monitoring wells located within and immediately surrounding the TIKEVIR study domain were extracted, with records spanning from 1950 to the present day. Measurements are available at daily, weekly, and monthly intervals. Previous analyses of shallow GWL fluctuations in this region have

identified climate-driven periodicities at annual and multi-decadal scales (Garamhegyi et al., 2018) and quantified precipitation–groundwater lag times of 7–10 months (Illyés et al., 2023), both of which informed the selection of the 12-month lookback window for the ConvLSTM model (Sect. 3.4).

Converting the irregularly sampled well records into a spatially complete monthly gridded product required two sequential processing stages: (i) temporal gap-filling to produce a complete monthly time series at each well location and (ii) spatial interpolation to map the gap-filled point observations onto the regular 1 km analytical grid.

2.2.2.1 Temporal gap-filling

The raw well records exhibit substantial temporal gaps, particularly before the mid-1990s, necessitating a robust gap-filling procedure to address these issues. Prior to gap-filling, data quality was ensured through three filters: (a) removal of records with ADAT = 0, which designates a missing-value marker in the OVF database and produces spurious values near −10 m; (b) exclusion of depth values outside the physically plausible range [0, 20] m; and (c) per-well interquartile range (IQR) outlier removal, discarding observations more than $3 \times$ IQR beyond the first or third quartile.

Gap-filling followed a five-step procedure applied to the monthly time series data of each well.

1. **Detrending.** Long-term linear trends were estimated for anchor wells with valid records of at least 30 years. For wells with shorter records, the trend was transferred from the five nearest anchor wells using a KDTree spatial lookup by averaging their trend slopes. To prevent extrapolation artifacts, the estimated trend was capped (held constant) outside the observed period of each well.
2. **Standardisation.** The detrended series is transformed into monthly robust Z-scores: for each calendar month, the median and median absolute deviation (MAD) are computed over the available record, and the series is expressed as $Z = (x - \text{median}) / (\text{MAD} \times 1.4826)$. This removes the seasonal amplitude differences between wells, enabling inter-well comparisons on a common scale.
3. **Regional reference signals.** A domain-wide reference signal $Z_{ref}(t)$ was constructed as the monthly median of the standardized Z-scores across all long-record ($\geq$30 yr) wells. This signal captures the regional-scale groundwater response to climate forcing and serves as the primary temporal predictor in the gap-filling model used in this study, as follows:
4. **Gap-filling regression.** Missing months were infilled individually for each well using a Huber-robust linear regression (Huber, 1973) of the form $Z(t) = a \cdot t + A \cdot \sin(2\pi t/12) + B \cdot \cos(2\pi t/12) + \beta \cdot Z_{ref}(t)$, where the first term captures the residual trend, the sinusoidal terms represent the seasonal cycle (annual harmonic), and $\beta \cdot Z_{ref}$ provides the inter-annual climate-driven variability. The Huber loss function down-weights the outlier residuals, improving robustness against local anomalies.
5. **Back-transformation.** The filled Z-scores were reversed through standardization (multiplying by the MAD-scale and adding the monthly median) and then through detrending (adding the estimated linear trend), recovering gap-filled values in the original depth-to-water-table units (m below ground surface).

Post-validation enforces three constraints: (a) wells with fewer than 60 months of original observations are excluded from the gap-filled product; (b) per-well range validation ensures that filled values remain within the observed minimum and maximum for each well; and (c) a hard floor/ceiling of [0, 20] m is applied. The resulting gap-filled product provides a complete monthly time series at all retained well locations from 1951 to 2025.

2.2.2.2 Spatial interpolation via Sequential Gaussian Simulation

The gap-filled point observations were spatially interpolated onto a regular 1 km Cartesian grid (100 rows $\times$ 65 columns in EPSG:23700) using Sequential Gaussian Simulation (SGS; Deutsch and Journel, 1998), implemented via the GSTools package (Müller et al., 2022). SGS was chosen over deterministic

kriging because it preserves the full spatial variability of the observed well data, producing fields with realistic heterogeneity that are better suited for training ConvLSTM than the artificially smooth fields produced by ordinary kriging.

At each monthly time step, the procedure consists of three stages:

1. **Normal-score transform.** All valid well observations were transformed into standard Gaussian space using a rank-based empirical CDF constructed from the full spatiotemporal record (approximately 183,000 values). This ensures that the conditioning data satisfy the multi-Gaussian assumption required by the SGS.
2. **Conditional simulation.** Sixteen (16) conditional realizations were generated on the target grid using a fitted isotropic spherical variogram with a sill of 0.7, range of 30 km, and nugget of 0.3, calibrated on the most data-dense time steps. Each realization honors the observed values at the well locations and reproduces the target spatial covariance structure, while differing in the unsampled locations owing to the stochastic nature of the simulation.
3. **Ensemble aggregation and back-transformation.** The ensemble of 16 realizations was summarized using the MD-type (median) estimator, which is more robust to outlier realizations than the E-type (mean). The median field serves as the best-estimate GWL surface, while the inter-realization standard deviation provides a spatially explicit measure of interpolation uncertainty. The ensemble statistics were then back-transformed from the Gaussian space to the original depth units, and the simulated values were clipped to the [0, 20] m physical range.

The resulting gridded GWL product comprised 894 monthly grids (January 1951 to June 2025) at a 1 km resolution, covering a domain area of approximately 4,041 km². The FORESEE meteorological forcing was resampled from its native 0.1° grid to the same 1 km analytical grid using cKDTree nearest-neighbor mapping.

2.2.3 Soil hydraulic property grids

Spatially distributed soil hydraulic parameters were obtained from the EU-SoilHydroGrids database (Tóth et al., 2017), a pan-European product providing Mualem–van Genuchten parameters at a resolution of 250 m and seven standard depth intervals (0, 5, 15, 30, 60, 100, and 200 cm). The parameters of primary interest for this study are the van Genuchten shape parameters α (m^{-1}) and n (–), residual and saturated water contents θ_r and θ_s (m³ m^{-3}), and saturated hydraulic conductivity K_{sat} (cm d^{-1}). These were aggregated to a depth-weighted average over the 0–200 cm profile to represent the effective unsaturated zone properties relevant to the capillary rise processes. The 250 m grids were spatially aggregated (arithmetic mean for θ parameters, geometric mean for K_{sat} and α) to the 1 km analytical grid (EPSG:23700), matching the GWL and ConvLSTM model resolution.

Topsoil physical properties (clay, sand, and silt content (%), available water capacity (AWC), bulk density, and USDA texture class) were obtained from the ESDAC topsoil dataset (Ballabio et al., 2016), which maps these properties at a 500 m resolution across Europe using geostatistical modelling of LUCAS topsoil survey data. The clay content and AWC were resampled to a 1 km analytical grid and used as static input channels for the ConvLSTM model (Sect. 3.4.1); the USDA texture classification provides the spatial context shown in Fig. 1b.

For the validation and supplementary classification of soil texture, the Hungarian AGROTOPO database (Várallyay et al., 1979) at a 1:100 000 scale was employed. This dataset provides genetic soil type classification, texture class, and parent material information for the entire country. Cross-referencing AGROTOPO texture classes with the EU-SoilHydroGrids van Genuchten parameters confirmed internal consistency over the study domain, with sandy soils (Arenosols) in the interfluve

areas exhibiting systematically higher α values (6–15 m⁻¹) than the clay-rich alluvial soils (Fluvisols and Vertisols) along the river corridors (α = 0.5–3 m⁻¹).

3 Methods

3.1 Computation of critical groundwater depth

The extinction depth, below which groundwater ceases to contribute to evapotranspiration through capillary rise, has traditionally been parameterized as a static, soil-texture-dependent quantity. Shah et al. (2007) showed, using HYDRUS-1D simulations, that extinction depths range from approximately 2 m in sands to over 8 m in fine-textured soils, and proposed a scaling relationship $d_{ext} = k / \alpha_{vG}$, where α_{vG} is the van Genuchten capillary parameter, and k is a dimensionless scaling factor. Although widely adopted, this formulation assigns a time-invariant extinction depth to each grid cell, thereby ignoring the dynamic feedback between the current water table position and capillary flux capacity.

In this study, we replaced the static formulation with a dynamic GWL-dependent critical depth that captures the physically motivated expectation that the effective critical depth varies with the current position of the water table. When the water table is shallow, the capillary potential gradient required to sustain the evaporative flux to the surface is readily achieved, and the critical depth is close to the water table position. As the water table deepens, progressively larger potential gradients are required, and the critical depth saturates toward an asymptotic maximum. This behavior is represented by the following equation:

$$d_{crit}(x, y, t) = k_d \cdot (0.8 + 0.4 \cdot z_{GWL}(x, y, t) / (z_{GWL}(x, y, t) + 3)) \tag{1}$$

where k_d (m) is a baseline scaling parameter set to 2.5 m in the default configuration, and z_{GWL} (x, y, t) is the depth to the water table (m below ground surface, positive downward) at grid cell (x, y) and time t. The nonlinear saturating function compresses the GWL range into a narrow critical depth interval, producing values between approximately $0.8 \cdot k_d$ = 2.0 m (at $z_{GWL} \rightarrow 0$) and $1.2 \cdot k_d$ = 3.0 m (at $z_{GWL} \rightarrow \infty$). This formulation ensures that d_{crit} responds most sensitively to GWL changes in the shallow regime ($z_{GWL} < 3$ m), where salinization risk is concentrated, while asymptotically approaching the upper bound at greater depths. The tested range of $k_d \in \{1.5, 2.0, 2.5, 3.0, 3.5\}$ m spans the extinction depths reported by Shah et al. (2007) for medium- to fine-textured soils using HYDRUS-1D simulations (approximately 1.3–5.5 m). Independently, Bakacsi et al. (2019) mapped the static "critical groundwater level," defined as the depth above which capillary transport from saline groundwater causes soil degradation across the Hungarian Great Plain using soil texture, salt content, and groundwater chemistry data, reporting values of 2–4 m depending on local pedological and hydrochemical conditions. The default d_{crit} range produced by Eq. (1) with k_d = 2.5 m (2.10–2.83 m) falls within this empirically constrained envelope, providing indirect corroboration of the parameterization. The sensitivity of R to k_d is evaluated in Sect. 4.4.

Similarly, the capillary zone height, which is the vertical extent of the capillary fringe above the water table, is parameterized as a GWL-dependent function as follows:

$$h_{cap}(x, y, t) = k_c \cdot (1.0 + 0.3 / (1.0 + z_{GWL}(x, y, t))) \tag{2}$$

where $k_c = 1.0$ m is the baseline capillary height. This function yields values between approximately 1.0 m (deep water table) and 1.3 m (very shallow water table), reflecting enhanced capillary connectivity under shallow groundwater conditions.

3.2 Reference and actual evapotranspiration

Reference evapotranspiration (ET₀) was computed at daily time steps for each 0.1° grid cell from the FORESEE meteorological fields using the FAO Penman–Monteith equation (Allen et al., 1998):

$$\text{ET}_0 = [0.408\,\Delta\,(R_n - G) + \gamma\,(900\,/\,(T + 273))\,u_2\,(e_s - e_a)]\,/\,[\Delta + \gamma\,(1 + 0.34\,u_2)] \tag{3}$$

where R_n is net radiation (MJ m⁻² d⁻¹), G is soil heat flux density (MJ m⁻² d⁻¹, assumed negligible at the daily time step), T is mean daily air temperature (°C), u_2 is wind speed at 2 m height (m s⁻¹), $e_s - e_a$ is the saturation vapor pressure deficit (kPa), Δ is the slope of the saturation vapor pressure–temperature relationship (kPa °C⁻¹), and γ is the psychrometric constant (kPa °C⁻¹). Where global radiation was provided in the FORESEE-HUN database, net radiation was estimated following the procedures described by Allen et al. (1998, Chap. 3).

Actual evapotranspiration (ET$_a$) was estimated by applying crop coefficients (K_c) derived from the CORINE 2018 land cover classification. Each land cover class was assigned a monthly K_c time series following the FAO-56 single crop coefficient approach (Allen et al., 1998, Tables 12 and 17), with mid-season K_c values of 1.15 for cereals (wheat, maize), 1.0–1.15 for sunflower and rapeseed, 0.75 for grasslands, and 1.0 for broadleaf forests. Initial and late-season K_c values follow the standard FAO-56 crop development curves for continental European growing conditions. The actual ET flux relevant to capillary rise is then

$$ET_a(x, y, t) = K_c(x, y, t) \cdot \text{ET}_0(x, y, t) \tag{4}$$

Daily values were aggregated to monthly means for use in the risk index computation and as predictor variables for the deep learning model.

3.3 Dynamic salinization risk index

The core conceptual advance of this study is the formulation of a time-varying salinization risk index that captures the dynamic position of the water table relative to the critical depth and the capillary fringe. In its most general form, such an index could incorporate multiplicative factors for atmospheric evaporative demand (f_{ET}), groundwater proximity to the critical depth (R), and soil susceptibility (S). In this study, we focused on the groundwater proximity component R, which is the dominant driver of spatiotemporal risk variability in the TIKEVIR region. This choice reflects a physical argument: the seasonal GWL trajectory already integrates the cumulative effect of ET on the soil water balance, as the water table rises during the low-ET recharge season and declines during the high-ET discharge season. The water table position at any given month therefore encodes the antecedent balance between recharge and evapotranspiration, making it a sufficient first-order proxy for the capillary flux regime. Explicit ET and soil susceptibility factors are discussed as refinements in Sect. 5.6.

The risk index R(x, y, t) is defined as a linear interpolation between the capillary zone height h_{cap} and the critical depth d_{crit} as follows:

$$R(x, y, t) = \text{clamp}[0,1]((d_{crit}(x, y, t) - z_{GWL}(x, y, t)) / (d_{crit}(x, y, t) - h_{cap}(x, y, t))) \qquad (5)$$

where $\text{clamp}[0,1](x) = \min(1, \max(0, x))$, $z_{GWL}(x, y, t)$ is the observed or predicted depth to water table (m below ground surface) at grid cell (x, y) and time t, and $d_{crit}(x, y, t)$ is the dynamic critical depth from Eq. (1), and $h_{cap}(x, y, t)$ is the capillary zone height from Eq. (2). The index has the following boundary behavior:

- When $z_{GWL} > d_{crit}$: $R = 0$ (water table below the critical depth; no capillary-driven evaporation reaches the surface).
- When $z_{GWL} < h_{cap}$: $R = 1$ (water table within the capillary fringe, maximum capillary flux and salinization potential).
- Between these limits, R increases linearly with decreasing depth to water table.

This formulation follows the conceptual structure of the MODFLOW EVT package (McDonald and Harbaugh, 1988) but extends it by making both the upper (d_{crit}) and lower (h_{cap}) bounds dynamically dependent on the current water table position through Eqs. (1) and (2). The denominator ($d_{crit} - h_{cap}$) defines the active zone width over which the capillary flux transitions from zero to its maximum, which is physically more meaningful than normalizing by d_{crit} alone, as it explicitly accounts for the soil-dependent capillary fringe height.

Because d_{crit} and h_{cap} are both functions of z_{GWL} (Eqs. 1–2), the risk index responds nonlinearly to water table changes, even though the interpolation itself is linear. In the shallow regime ($z_{GWL} < 2$ m), small GWL changes produce large risk variations, whereas at depth, the response saturates, which is consistent with the physics of capillary flux attenuation. This nonlinearity is a direct consequence of the GWL-dependent critical depth formulation and represents the key distinction from static-threshold approaches.

We note that an explicit ET demand factor ($f_{ET} = ET_a / ET_{a,max}$) and a soil susceptibility factor based on texture, drainage, and groundwater salinity could further refine this risk estimate. Their omission from the present formulation means that the index captures the *potential* for capillary-driven salinization based on the water table position, rather than the actual salt flux, which additionally depends on atmospheric demand and soil transport properties. Consequently, the temporal risk maximum coincided with the shallowest water table (April), rather than the peak ET demand (July–August). Including f_{ET} would shift the peak risk into the May–June window, where both GWL proximity and ET demand are simultaneously elevated, an extension explored further in Section. 5.6.

The resulting risk index $R(x, y, t)$ ranges from 0 to 1 and is classified into five categories for mapping and communication purposes: no risk ($R < 0.1$), low ($0.1 \leq R < 0.3$), moderate ($0.3 \leq R < 0.5$), high ($0.5 \leq R < 0.7$), and critical ($R > 0.7$).

3.4 Spatiotemporal groundwater level forecasting with ConvLSTM

3.4.1 Model architecture

Groundwater level forecasting was performed using a Convolutional Long Short-Term Memory (ConvLSTM) neural network (Shi et al., 2015), which extends the standard LSTM architecture by replacing fully connected operations with convolutions in both the input-to-state and state-to-state transitions. This enables the model to jointly capture spatial correlations across the gridded domain and temporal dependencies in the GWL sequence.

The model receives eight input channels at each grid cell, organized into three groups: (i) three dynamic fields that vary both spatially and temporally: normalized GWL depth (m), mean daily precipitation rate (mm d^{-1}), and monthly mean temperature (°C); (ii) three static spatial fields that are tiled identically across all time steps: digital elevation model (DEM) height (m), topsoil clay content (%), and available water capacity (AWC, cm cm^{-1}), sourced from the Hungarian national DEM at 100 m resolution (coarsened to 1 km by block averaging) and the ESDAC topsoil dataset at 500 m resolution (Ballabio et al., 2016), respectively; and (iii) two cyclical temporal encoding fields: $\sin(2\pi m/12)$ and $\cos(2\pi m/12)$, where m is the zero-indexed calendar month, broadcast as spatially constant channels that provide explicit seasonal context to the model. For each training sample, a 12-month lookback window was constructed, yielding an input tensor of shape (12, 8, H, W), where $H \times W = 100 \times 65$ corresponds to the spatial dimensions of the analytical grid over the TIKEVIR domain.

The encoder consisted of two stacked ConvLSTM layers with hidden dimensions of 32 and 16, respectively, using 3×3 convolutional kernels in both the input-to-state and state-to-state transitions. A dropout (p = 0.1) was applied between the layers for regularization and to enable Monte Carlo dropout uncertainty estimation during inference (Gal and Ghahramani, 2016). The output of the second ConvLSTM layer at the final time step was passed through a 1×1 convolutional layer that mapped the 16-dimensional hidden state to a single output channel: the predicted normalized GWL depth field for the next month (1-month-ahead forecast).

For multi-step forecasting beyond 1 month (e.g., exploratory scenario projections to 2050), an autoregressive inference strategy was employed. The initial 1-month-ahead prediction was denormalized, appended to the input sequence (discarding the oldest month), and renormalized using the training-set statistics. Static channels remain fixed, whereas the cyclical month encodings are updated to reflect the forecast target month. The model was then re-applied iteratively, with future monthly climate fields (precipitation and temperature) drawn from the FORESEE-HUN RCP 4.5 projection dataset (Kern et al. 2024). To assess the contribution of individual architectural choices, a sensitivity analysis was conducted across four model configurations by varying the number of input channels, training period length, hidden dimensions, and learning rate scheduling (Table 5 in Section 4.5.1).

3.4.2 Training procedure

The spatiotemporal dataset spanned 1971–2024. Monthly meteorological fields (precipitation and temperature) were derived from the FORESEE-HUN v1.0 daily gridded dataset (Kern et al., 2024) by aggregating daily values to monthly mean daily rates (precipitation, mm d^{-1}) and means (temperature, °C). The original FORESEE 0.1° latitude–longitude grid was regridded to the EPSG:23700 analytical grid using nearest-neighbor assignment based on a KDTree mapping of the WGS 84 source coordinates to the projected target grid. Overlapping 12-month input sequences were constructed from 1971 to 2019 (588 training months, yielding approximately 576 training samples). The dataset was partitioned chronologically into a training set (1971–2019) and an evaluation set (2020–2024, 60 months), with no temporal overlap to prevent the possibility of data leakage. The 12-month lookback window means that the first evaluation prediction uses January–December 2019 as the input to forecast January 2020. All performance metrics reported in Section 4.5 were computed exclusively on the evaluation set.

The model was implemented in PyTorch (Paszke et al., 2019) and trained using the Adam optimizer (Kingma and Ba, 2015) with an initial learning rate of 1×10^{-3}, a batch size of 4, and the mean squared error as the loss function computed over all valid (non-masked) grid cells. A ReduceLROnPlateau learning rate scheduler (factor = 0.5, patience = 5 epochs) was applied to reduce the learning rate when the validation loss reached the plateau. Each input feature was independently normalized to zero mean

and unit variance using training-set statistics, and the normalization parameters were stored alongside the trained model weights to ensure consistent denormalization during inference. Training proceeded for 50 epochs with early stopping based on validation loss (patience = 10 epochs), monitored on the final 20 % of the training period (held out as a validation subset).

3.4.3 Evaluation metrics

The model performance was assessed using four standard metrics computed over all valid (non-masked) grid cells and evaluation time steps. The root-mean-square error quantifies the absolute prediction accuracy as follows:

$$\text{RMSE} = \text{sqrt}\left((1/N) \, \Sigma_{i=1}^{\{N\}} (\hat{y}_i - y_i)^{\wedge}2 \right) \tag{6}$$

The mean absolute error provides a robust alternative that is less sensitive to outliers:

$$\text{MAE} = (1/N) \, \Sigma_{i=1}^{\{N\}} |\hat{y}_i - y_i| \tag{7}$$

The Nash–Sutcliffe efficiency (NSE) measures the predictive skill relative to the observed mean as follows:

$$\text{NSE} = 1 - [\Sigma_i \, (y_i - \hat{y}_i)^{\wedge}2 \, / \, \Sigma_i \, (y_i - \bar{y})^{\wedge}2] \tag{8}$$

where $\bar{y}$ is the mean observed GWL over the evaluation period. An NSE value of 1 indicates perfect prediction, an NSE value of 0 indicates performance equivalent to predicting the mean, and an NSE value of > 0.7 is generally considered "good" in hydrological applications (Moriasi et al., 2007). The Pearson correlation coefficient (r) quantifies the linear association between the predicted and observed fields. All metrics were reported as domain aggregates and cell-level maps to assess the spatial distribution of the forecast accuracy.

3.5 Coupled risk forecasting

The trained ConvLSTM model was applied in inference mode to generate GWL forecasts at lead times of 1, 3, 6, and 12 months, using the most recent 12-month observational window as input. The forecast GWL fields are substituted into the dynamic risk index (Eq. 5), replacing the observed z_{GWL} with the predicted value:

$$R_{\text{forecast}}(x, y, t{+}h) = \text{clamp}[0,1]\left((d_{\text{crit,forecast}} - z_{\text{GWL,forecast}}) / (d_{\text{crit,forecast}} - h_{\text{cap,forecast}}) \right) \tag{9}$$

where $d_{\text{crit,forecast}}$ and $h_{\text{cap,forecast}}$ are computed from the predicted $z_{\text{GWL,forecast}}$ using Eqs. (1) and (2), respectively. For the exploratory scenario runs to 2050 (Section 4.6), ConvLSTM received projected precipitation and temperature from FORESEE-HUN RCP 4.5 as input forcing for GWL prediction. The dropout layers (p = 0.1) were retained during inference to enable Monte Carlo dropout uncertainty estimation (Gal and Ghahramani, 2016). In addition to deterministic (single forward pass) forecasts, we evaluated epistemic uncertainty by performing 50 stochastic forward passes per evaluation time step with dropout active, computing per-cell GWL prediction standard deviation, and propagating the

ensemble through the risk model to obtain $P(R \geq 0.5)$ and risk credible interval width (P90 – P10) maps (Sect. 4.5.4).

3.6 Sensitivity analysis

The sensitivity of the framework to key methodological choices was evaluated through two complementary analyses:

(a) **Risk index parameterization.** The baseline scaling parameter k_d in Eq. (1) is varied over k_d = {1.5, 2.0, 2.5, 3.0, 3.5} m, and the change in the domain-averaged risk classification relative to the baseline (k_d = 2.5 m) is reported (Sect. 4.4). This parameter directly controls the critical depth and is the dominant source of methodological uncertainty.

(b) **ConvLSTM model configuration.** Four configurations (A–D) were compared, varying the number of input channels (3, 6, or 8), training period length (2002–2019 or 1971–2019), hidden layer dimensions ([16, 8] or [32, 16]), and learning rate scheduling (Table 5). This isolates the contribution of static spatial channels, training data volume, temporal encoding, and model capacity to forecast skills at lead times of 1, 3, 6, and 12 months.

**3.7 Implementation, interactive dashboard, and reproducibility**

3.7.1 Software environment

All processing, modelling, and visualization were implemented using Python 3.11. Gridded data handling used xarray (Hoyer and Hamman, 2017) and rioxarray; geostatistical simulation employed GeoStat-Framework/GSTools (Müller et al., 2022); the ConvLSTM was implemented in PyTorch 2.x (Paszke et al., 2019); coordinate transformations used pyproj; and interactive mapping used Folium with streamlit-folium bindings.

3.7.2 Interactive decision-support dashboard

To enable operational use by non-specialist users, a complete analytical pipeline was implemented as an interactive web application using the Streamlit framework (Streamlit Inc. 2024). The dashboard provides seven analysis panels: observed GWL depth maps with a historical time slider (1951–2025), dynamic critical depth, and capillary zone maps computed from Eqs. (1)–(2), ConvLSTM-predicted GWL fields for 2026–2050 driven by FORESEE-HUN RCP 4.5 projections, and coupled salinization risk maps for both observed and forecasted conditions. All panels render Folium-based interactive map overlays in the EOV (EPSG:23700) coordinate system, with point-and-click interrogation of individual grid cells. A sidebar provides real-time control over the baseline scaling parameter k_d, capillary rise parameter k_c, and the forecast horizon year, enabling interactive "what-if" scenario exploration.

The dashboard architecture follows FAIR principles (Wilkinson et al., 2016). The complete codebase, trained model weights, and representative gridded data sample were archived on Zenodo (https://doi.org/10.5281/zenodo.18770854). All gridded products use the CF-compliant NetCDF format; the modular Python architecture separates data loading, risk computation, forecasting, and visualization into independent modules, facilitating adaptation to other study regions. Detailed dashboard documentation, including tab descriptions and deployment instructions, is provided in the repository's README file.
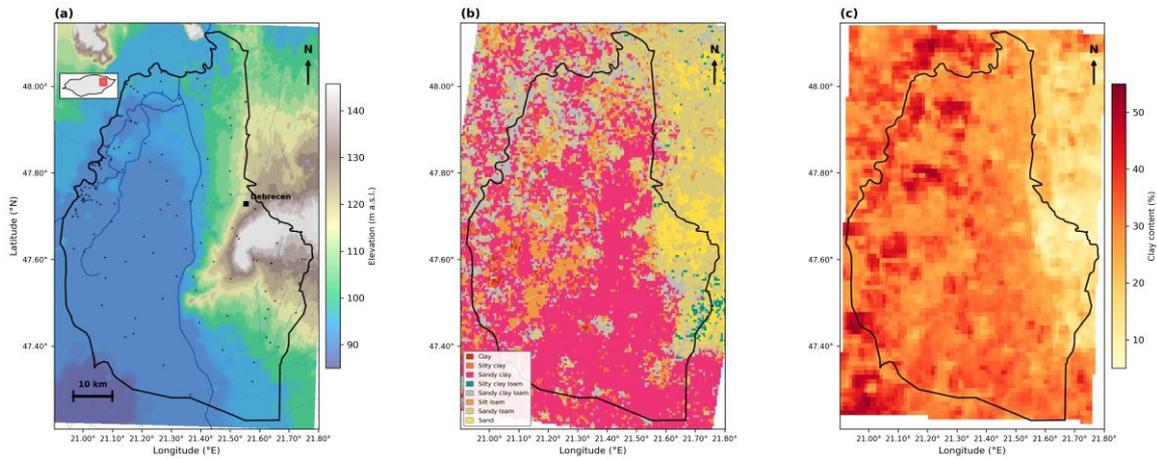
Figure 1. Map of the TIKEVIR study area showing (a) topography (Hungarian National DEM, 100 m) with hillshade relief, the groundwater monitoring well network, main watercourses, and the Debrecen urban area (inset: location within Hungary), (b) USDA soil texture classification (ESDAC topsoil; Ballabio et al., 2016; 500 m), and (c) topsoil clay content (%; ESDAC topsoil; Ballabio et al., 2016; resampled to 1 km).
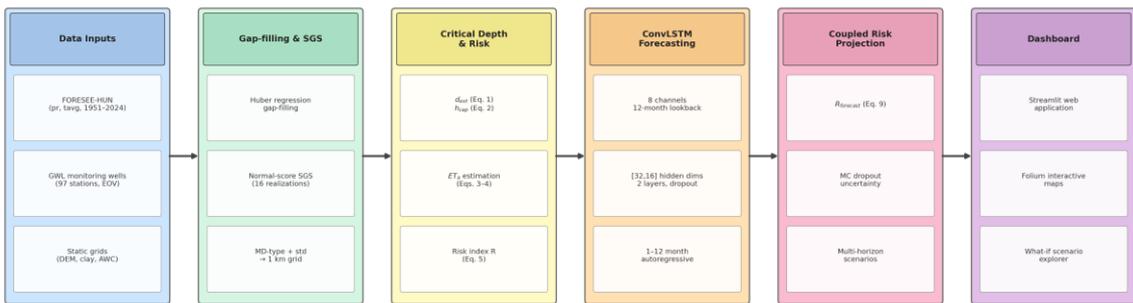
Figure 2. Schematic of the methodological framework showing the processing chain: data inputs (FORESEE-HUN climate, GWL monitoring network, static grids) → gap-filling and SGS interpolation → critical depth and risk index computation → ConvLSTM spatiotemporal forecasting → coupled risk projection with MC dropout uncertainty → interactive dashboard.

4 Results

4.1 Spatial distribution of critical depth

The dynamic critical depth field $d_{crit}(x, y, t)$ computed from the GWL-dependent formulation (Eq. 1) with a baseline scaling parameter $k_d = 2.5$ m is presented in Fig. 3a. As described in Sect. 3.1, the critical depth varies with the observed water table position via Eq. 1, responding most sensitively to GWL changes in the shallow regime while saturating at greater depths. Across the TIKEVIR domain (4,041 km²), the time-averaged critical depth ranges from 2.10 m in areas with persistently shallow water tables to 2.83 m in deeper groundwater zones, with a domain-averaged value of 2.45 ± 0.12 m (mean ± standard deviation) (Fig. 3a).

The spatial distribution of the critical depth mirrors the time-averaged GWL pattern (Fig. 3b): shallower critical depths occur in the low-lying discharge zones along the Körös and Berettyó floodplains, where water tables are consistently within 1–2 m of the surface, whereas deeper critical thresholds characterize the aeolian sand interfluves, where GWL depths exceed 4–5 m. This spatial coherence is expected from the GWL-dependent formulation and reflects the continuous variation in capillary flux capacity across the domain rather than discrete texture class boundaries.

The difference between the two panels in Fig. 3 highlights that the critical depth field exhibits considerably less spatial variability (SD = 0.12 m) than the underlying GWL field (SD = 2.07 m). This arises because the nonlinear saturating function compresses the GWL range into a narrow critical depth interval (2.10–2.83 m), indicating that the critical depth is relatively insensitive to GWL variations at deep water table positions but responds more strongly in the shallow GWL regime where salinization risk is concentrated.
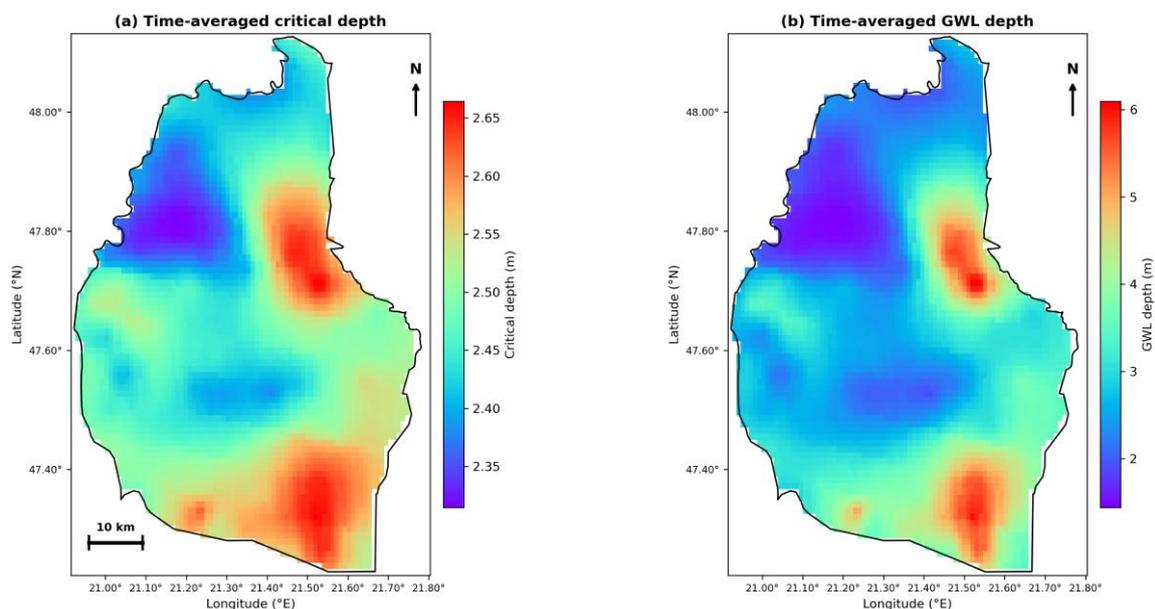
[Figure 3] *Spatial distribution of critical depth over the TIKEVIR domain: (a) time-averaged critical depth from the GWL-dependent empirical model with $k_d = 2.5$ m (Eq. 1), (b) time-averaged GWL depth (m below surface).*

4.2 Observed groundwater level dynamics and the phase relationship with ET

4.2.1 Seasonal cycle

The long-term mean seasonal cycle of the GWL depth, computed from 894 monthly time steps across 4,041 grid cells covering the period 1951–2025, exhibited a sinusoidal pattern (Fig. 4a). The shallowest

domain-averaged water table occurred in April at a depth of 2.50 m below the surface, whereas the deepest occurred in October at 3.04 m, yielding a mean seasonal amplitude of 0.54 m. The inter-cell standard deviation was the largest in summer, reflecting the heterogeneous response of different hydrogeological settings to seasonal ET forcing.

The timing of the GWL minimum (shallowest in April) preceded the ET peak (typically July–August in the Pannonian Basin) by approximately 3–4 months. This temporal offset arises because GWL recharge from winter and spring precipitation raises the water table to its annual maximum before summer ET demand decreases it. The phase relationship defines a critical salinization window from March to June, during which the water table resides close to or above the critical depth, while atmospheric evaporative demand accelerates rapidly.
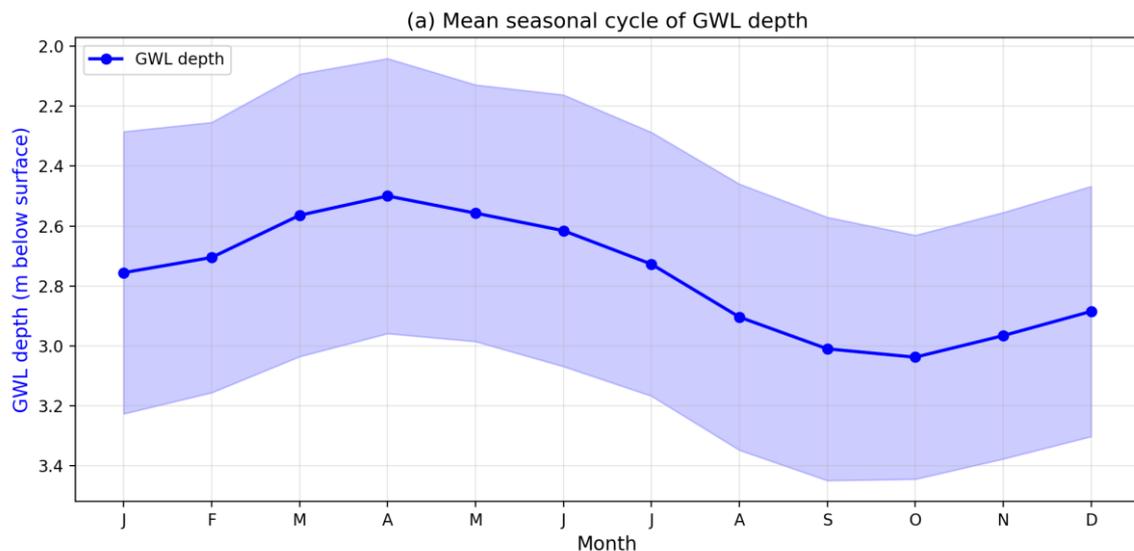
[Figure 4a] Domain-averaged mean seasonal cycle of GWL depth (blue, left axis), showing the sinusoidal pattern with shallowest levels in April and deepest in October. Error bands show ±1 SD across grid cells.

4.2.2 Long-term trend in GWL

Superimposed on the seasonal cycle, a statistically significant (p = 0.032) deepening trend was detected in the domain-averaged water table from 1951 to 2025 (Fig. 4b). The Mann-Kendall test yielded τ = 0.170, with Sen's slope of 0.0051 m yr^{-1} (0.051 m per decade increase in depth below surface). Over the 74-year record, this translated to a cumulative deepening of approximately 0.38 m in the domain-averaged water table. This secular trend is consistent with the combined effects of land-use intensification, increased ET due to warming, and long-term changes in recharge patterns documented across the Great Hungarian Plain in recent decades. We note that the domain-averaged series is derived from spatially correlated grid cells, which may inflate the effective sample size; however, the 74-year temporal record exceeds the decorrelation time scale of interannual GWL variability, and the trend direction is consistent with independent well-level analyses (Fehér and Rakonczai, 2019).
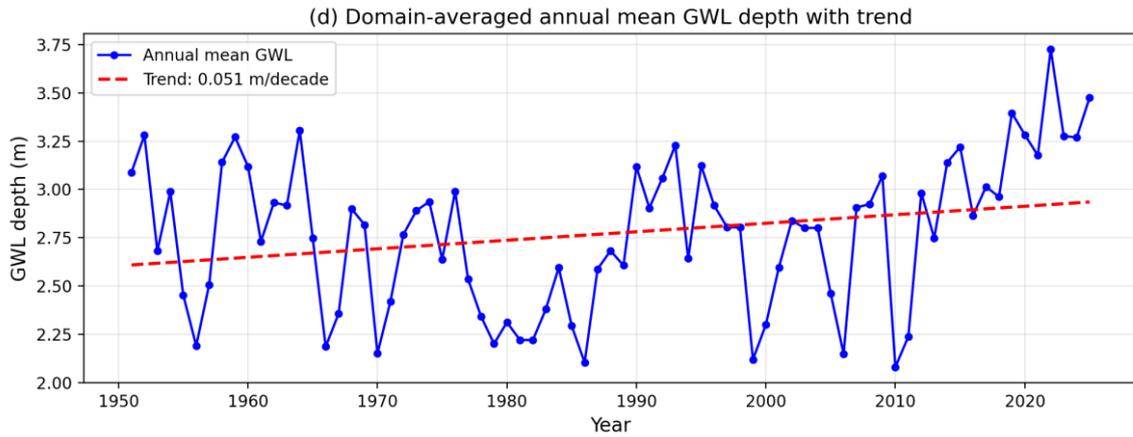
[Figure 4b] Time series of domain-averaged annual mean GWL depth (1951–2025) with a linear trend line. Sen's slope = 0.051 m per decade (p = 0.032).

4.2.3 Fraction of time above critical depth

Across the domain, the water table resides above the critical depth d_{crit} (computed from Eq. 1 with k_d = 2.5 m) for 29.8 % of all monthly observations on average; however, this fraction exhibited extreme spatial variability, ranging from 0 % in the deeply drained sandy upland areas to 100 % in the permanently shallow discharge zones (Fig. 4c). This diagnostic reveals a bimodal spatial structure: approximately 40 % of the domain is virtually never at risk (exceedance fraction < 5 %), while roughly 20 % of grid cells exceed the critical depth during more than half of all months, representing chronically vulnerable areas concentrated along the Körös and Berettyó floodplains and in areas with known salt-affected soils.

(c) Fraction of time GWL above extinction depth (2.5 m)
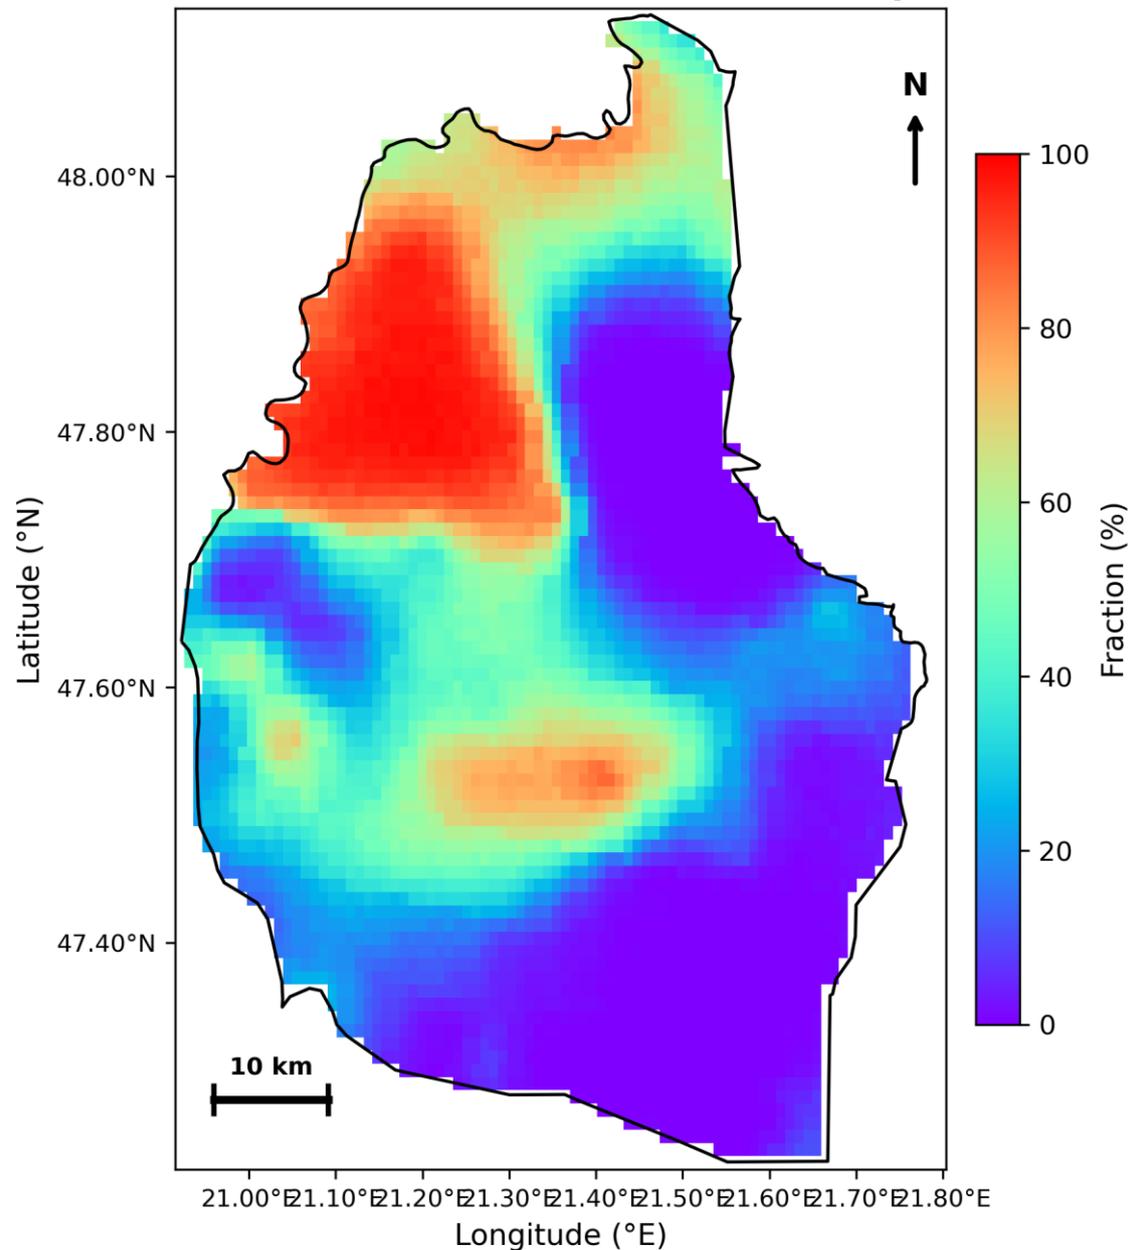
*[**Figure 4c**] Fraction of the historical record during which the water table resides above the critical depth d_{crit} (k_d = 2.5 m). Warm colors indicate chronically elevated risk zones.*

4.3 Historical dynamic salinization risk

4.3.1 Spatial patterns of mean annual risk

The time-averaged (climatological) salinization risk map (Fig. 5a) revealed a pronounced spatial dichotomy in salinization risk. The highest mean annual risk values (R > 0.5) were concentrated in the low-lying floodplain corridors of the Körös and Berettyó rivers and in the discharge zones along the eastern and southern margins of the study area, coinciding with areas historically mapped as solonchak and solonetz soils in the AGROTOPO database. Conversely, the lowest risk values (R < 0.1) were associated with elevated aeolian sand ridges and deeply drained interfluve areas in the northern and western regions of the domain. The domain-averaged time-mean risk was R = 0.251 ± 0.251, indicating

that the standard deviation equals the mean, a signature of the strongly bimodal spatial distribution between the low-risk and high-risk zones.

[Table 2] Distribution of the TIKEVIR domain across salinization risk categories based on the time-averaged risk index.

Risk category	R range	Area (km²)	Area (%)
No risk	< 0.1	1641	40.6
Low	0.1–0.3	1040	25.7
Moderate	0.3–0.5	584	14.5
High	0.5–0.7	407	10.1
Critical	>= 0.7	369	9.1

Table 2 reveals that 40.6 % of the domain (1,641 km²) is classified as "no risk" in the long-term mean, while 19.2 % (776 km²) falls in the "high" or "critical" categories. The remaining 40.2 % (1,624 km²) fell within the transitional "low" and "moderate" categories, representing areas where seasonal GWL fluctuations periodically elevate the water table above the critical depth.
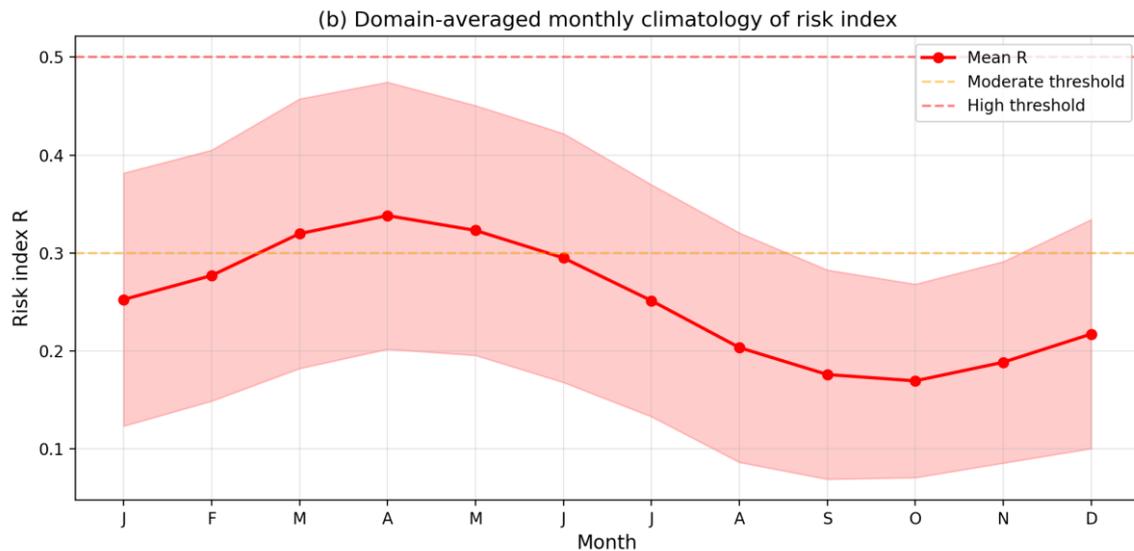
(a) Time-averaged salinization risk index
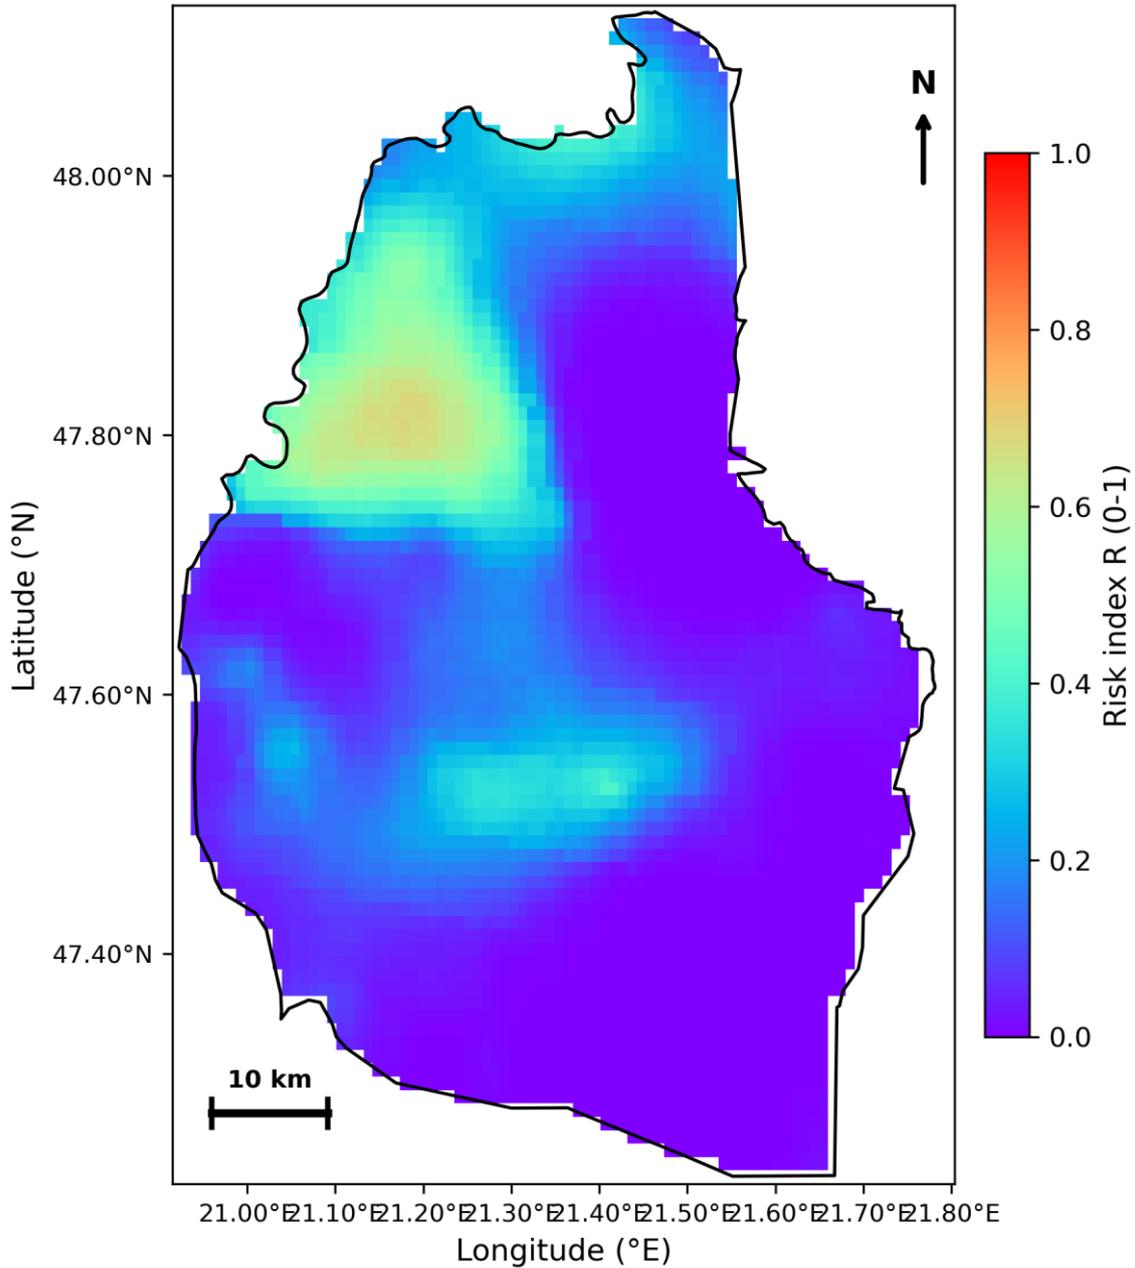
[Figure 5a] Time-averaged (climatological) salinization risk map for the TIKEVIR domain. Warm colors (red) indicate persistent high-risk areas concentrated in floodplain discharge zones.

4.3.2 Seasonal dynamics of risk

The monthly climatology of the domain-averaged R (Fig. 5b) peaked in April at R = 0.338 and reached a minimum in October at R = 0.169. This seasonal cycle is in direct antiphase with the GWL depth cycle (Sect. 4.2.1): risk is maximized when the water table is shallowest (April) and minimized when it is deepest (October).

The peak risk occurred in April, preceding the ET maximum (July–August) by approximately 3–4 months. This timing confirms that salinization risk is maximized not during peak ET, but during the ascending limb of the ET cycle, when the water table has not yet receded below the critical depth. Thus, the "window of vulnerability" spans March–June, when the GWL remains above d_{crit} in susceptible

areas, while ET demand is rising rapidly from its winter minimum. By late summer, when ET peaks, the water table has already declined sufficiently to suppress the capillary flux in most of the domain. This phase-controlled risk timing has important implications for irrigation scheduling and salt management, as interventions are most critical in spring rather than in mid-summer.

[Figure 5b] Domain-averaged monthly climatology of salinization risk index R(t). The peak in April coincided with the shallowest water table, not with the peak ET demand. Horizontal dashed lines indicate moderate (0.3) and high (0.5) risk thresholds, respectively.

To quantify the effect of explicitly incorporating atmospheric evaporative demand, an ET-weighted risk variant $R_{ET} = R \times f_{ET}$ was computed, where f_{ET} is the normalized Hargreaves reference ET_0 climatology derived from FORESEE daily temperature extremes from 2010 to 2024 (Fig. 11a). The Hargreaves formulation was chosen for this sensitivity test because it depends solely on temperature extremes and therefore provides a robust seasonal demand proxy without requiring the full radiation and humidity fields used by the Penman–Monteith ET_0 in the main modelling chain (Section 3.2). The domain-averaged f_{ET} ranged from 0.065 in December ($ET_0 = 0.36$ mm d^{-1}) to 1.000 in July ($ET_0 = 5.63$ mm d^{-1}), reflecting the pronounced seasonality of the evaporative demand in the Pannonian Basin. Multiplying R by f_{ET} shifts the peak risk month from April (month 4, for R alone) to June (month 6, for R_{ET}), a two-month delay that aligns with the physical expectation that the ET-weighted index peaks when the product of residual water table elevation and atmospheric demand is maximized, rather than when either factor peaks individually. Importantly, the spatial correlation between time-averaged R and R_{ET} was r = 0.999 (Fig. 11b, c), confirming that ET weighting modulates timing but does not alter the spatial ranking of risk hotspots. This high spatial coherence justifies the use of the GWL-only index R as the primary risk metric, with the understanding that the actual salt transport timing is shifted approximately two months later toward the peak-ET season.
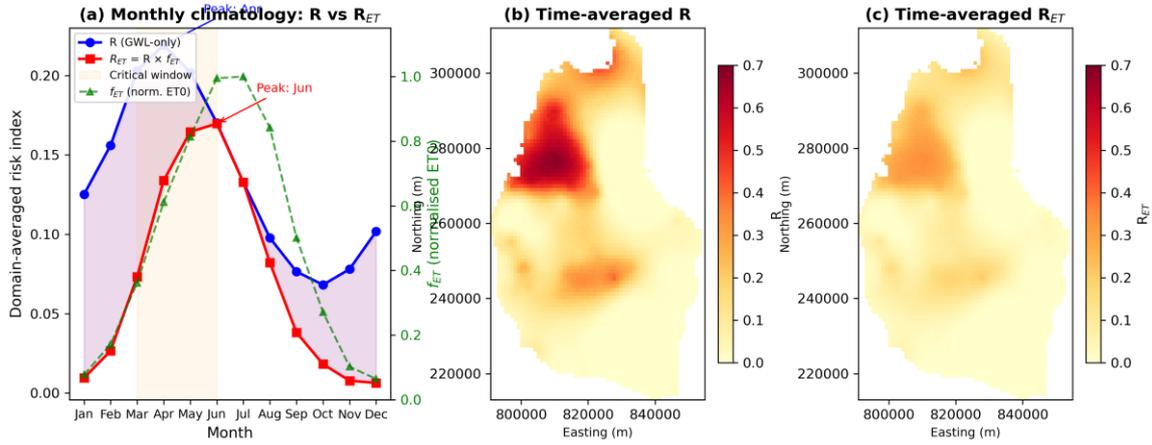
[Figure 11] ET-weighted risk sensitivity analysis: (a) monthly climatology of domain-averaged R (blue circles), R_{ET} (red squares), and normalized ET_0 (f_{ET}, green triangles, right axis); (b) time-averaged R map; (c) time-averaged R_{ET} map. The orange shading in (a) indicates the March–June critical window.

4.3.3 Interannual variability

The interannual variability of the domain-averaged peak monthly R was substantial, with a 37.7% coefficient of variation. The years with the highest peak risk were 2011 (R = 0.662), 2010 (R = 0.645), and 2000 (R = 0.644), all corresponding to years with exceptionally wet winter-spring periods that elevated water tables domain-wide. Conversely, the lowest-risk years were 2022 (R = 0.090), 2025 (R = 0.126), and 2014 (R = 0.151), which were characterized by deep water tables resulting from sustained dry conditions.
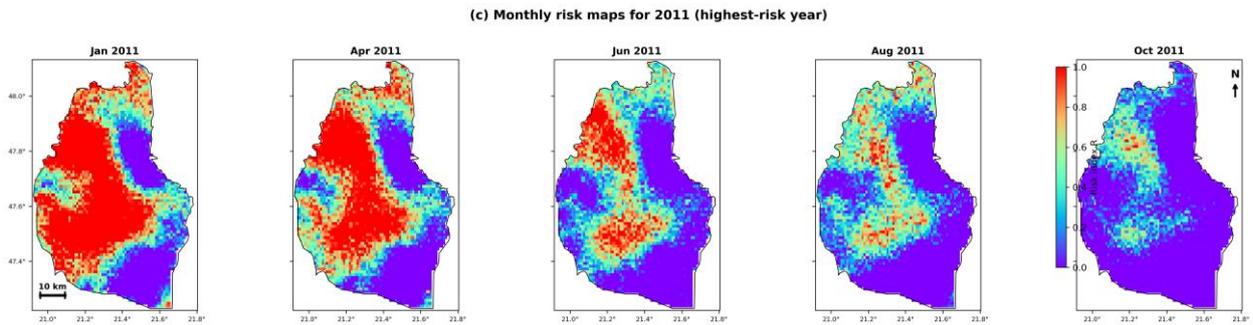

[Figure 5c] Monthly risk maps for 2011 (the highest-risk year), showing spatial evolution from January to October. The spring peak (April–June) and subsequent decline were clearly visible.

4.3.4 Temporal trends in risk

Over the 1951–2025 period, the domain-averaged annual peak R exhibited a weak, non-significant decreasing trend of -0.011 per decade (Mann-Kendall τ = -0.121, p = 0.127) (Fig. 5d). While this domain-averaged trend does not reach statistical significance at the 5 % level, it is consistent with the secular water table deepening documented in Sect. 4.2.2: As the water table gradually deepens, fewer grid cells experience above-critical-depth conditions during the spring recharge peak, reducing the domain-averaged risk.

However, this domain-averaged trend masks important local heterogeneity. In areas where the water table has fallen permanently below d_{crit}, the risk of crop failure has declined to near zero. Conversely, in persistently shallow discharge zones where the water table remains above d_{crit} regardless of the secular trend, the risk remains high and may intensify if ET rates increase due to climate warming. This

spatial divergence underscores the importance of spatially explicit risk assessments rather than domain-averaged metrics.
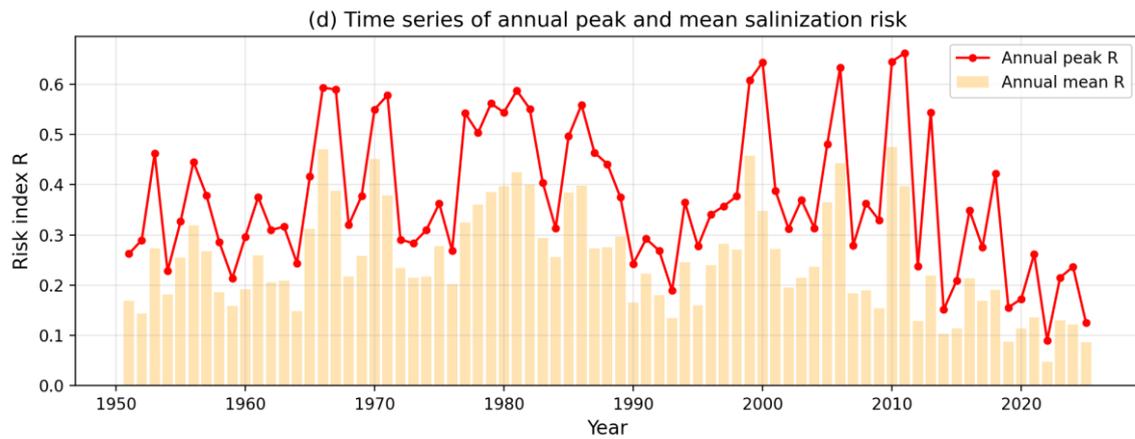
[Figure 5d] Time series of domain-averaged annual peak and mean salinization risk (1951–2025). The annual peak R (red line) showed substantial interannual variability overlaid on a weak downward trend. Orange bars indicate annual mean R.

4.3.5 Response to extreme droughts (2022, 2024)

The extreme droughts of 2022 and 2024 produced instructive risk responses that illuminate the counterintuitive relationship between droughts and salinization risk. In 2022, the domain-averaged R remained suppressed throughout the year, with a peak of only R = 0.090 in December, and an annual mean of R = 0.048, which is the lowest of any year in the record. This paradoxical result arises because the severe drought caused the water table to drop rapidly below the critical depth across the entire domain, effectively shutting off the capillary pathway for salt transport to the surface during the drought period. The water table was too deep for evaporative salinization to occur.

In 2024, the risk response was moderately higher, with a peak of R = 0.237 in February and an annual mean of R = 0.122. The February timing suggests that brief winter-spring recharge events likely lifted the water table above d_{crit} in some areas before the continuing drought pushed it back down. This contrasting behavior between 2022 and 2024 demonstrates that the relationship between drought severity and salinization risk is non-monotonic: moderate droughts (or drought recovery periods with partial recharge) can pose a higher salinization risk than extreme droughts that push the water table entirely below the capillary zone.

4.3.6 Independent validation against soil properties

To assess whether the modelled risk patterns were consistent with independent soil observations, the time-averaged risk index was correlated with topsoil clay content ($Clay_{eu23}$, from the ESDAC topsoil dataset; Ballabio et al., 2016) and calcium carbonate content ($CaCO3$, from the national soil survey) resampled to the 1 km model grid via nearest-neighbor interpolation (Fig. 6).

Clay content exhibited a weak but statistically significant positive correlation with mean risk (Spearman $\rho = 0.150$, $p < 10^{-21}$, n = 4,041 cells; note that the spatial autocorrelation among neighboring grid cells reduces the effective degrees of freedom, so the nominal p-value overstates the true significance). The positive sign is physically consistent because clay-rich soils sustain stronger capillary fluxes and are concentrated in low-lying floodplain discharge zones, where water tables are the shallowest. The modest magnitude reflects the fact that the risk index is dominated by the GWL proximity to the critical depth rather than by soil texture alone, which is consistent with the design intent of this dynamic formulation.

CaCO3 content showed no significant correlation with risk (Spearman's ρ = -0.010, p = 0.52, n = 3,895 cells). This null result is expected: while CaCO3 accumulation is a consequence of past salinization processes, the carbonate content of present-day topsoil reflects geological and pedogenetic history rather than the current dynamic risk state. This result confirms that the dynamic risk index captures active salinization susceptibility (driven by GWL position) rather than legacy salt accumulation.
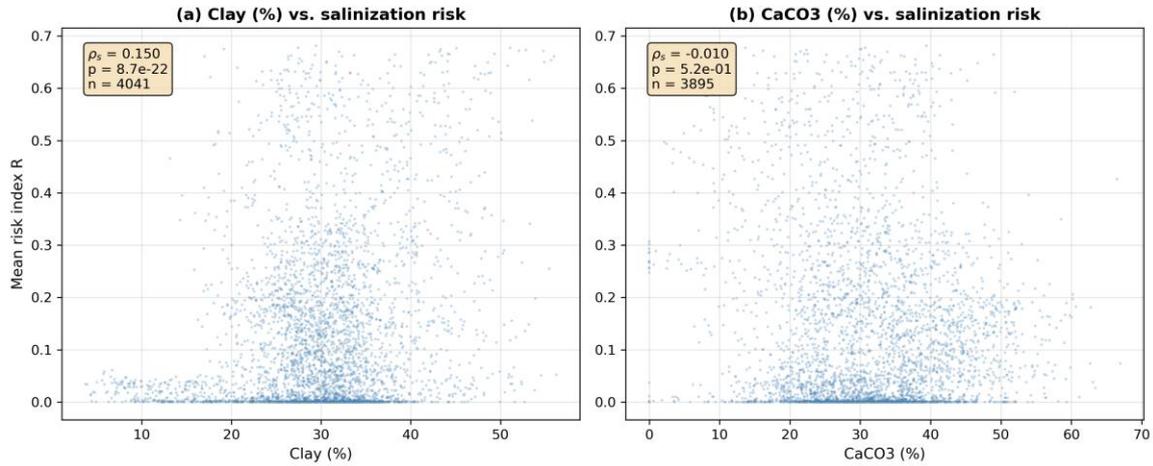
[Figure 6] Scatter plots of time-averaged salinization risk index versus (a) topsoil clay content and (b) CaCO3 content. Spearman rank correlation coefficients are indicated. The weak but significant positive correlation with clay content provides independent validation of the modelled risk patterns.

4.4 Sensitivity analysis

The sensitivity of the domain-averaged time-mean R to the baseline scaling parameter k_d (Eq. 1) is summarized in Fig. 7 and Table 3. This parameter directly controls the critical depth at which the capillary-driven evaporative flux from the water table ceases, making it the most influential methodological choice in the risk assessment framework.

[Table 3] Sensitivity of the domain-averaged mean risk index to the baseline scaling parameter k_d.

k_d (m)	Domain-averaged mean R	Max local R	Change from baseline (k_d = 2.5)
1.5	0.062	0.873	−0.189
2.0	0.142	1.000	−0.109
2.5	0.251	1.000	0.000 (baseline)
3.0	0.368	1.000	+0.117
3.5	0.471	1.000	+0.220

The risk index responds monotonically to k_d: increasing k_d from 1.5 to 3.5 m raises the domain-averaged R from 0.062 to 0.471, representing a 660 % increase. This strong sensitivity arises because a larger k_d extends the critical depth at which capillary-driven evaporation is considered active, thereby classifying more grid cells as "at risk." The relationship is approximately linear over the range 2.0–3.5 m (Fig. 7), with each 0.5 m increment in k_d producing a consistent ~0.10–0.12 increase in domain-averaged R, whereas the response steepens at lower k_d values, where the critical depth falls well below the domain-mean GWL.

Notably, the maximum local R reached 1.0 at all tested k_d values except 1.5 m, where the shallow critical depth (1.2–1.8 m) excludes even some chronically shallow discharge zones. Therefore, the

sensitivity is concentrated in the transitional zone, where the GWL fluctuates around the critical depth threshold, which is the area of greatest management relevance.

These results indicate that the choice of k_d is the primary source of methodological uncertainty in risk assessment. Field validation of the critical depth against measured evaporative flux profiles would substantially reduce this uncertainty, particularly in clay-dominated floodplain areas, where parameterization has the greatest impact.
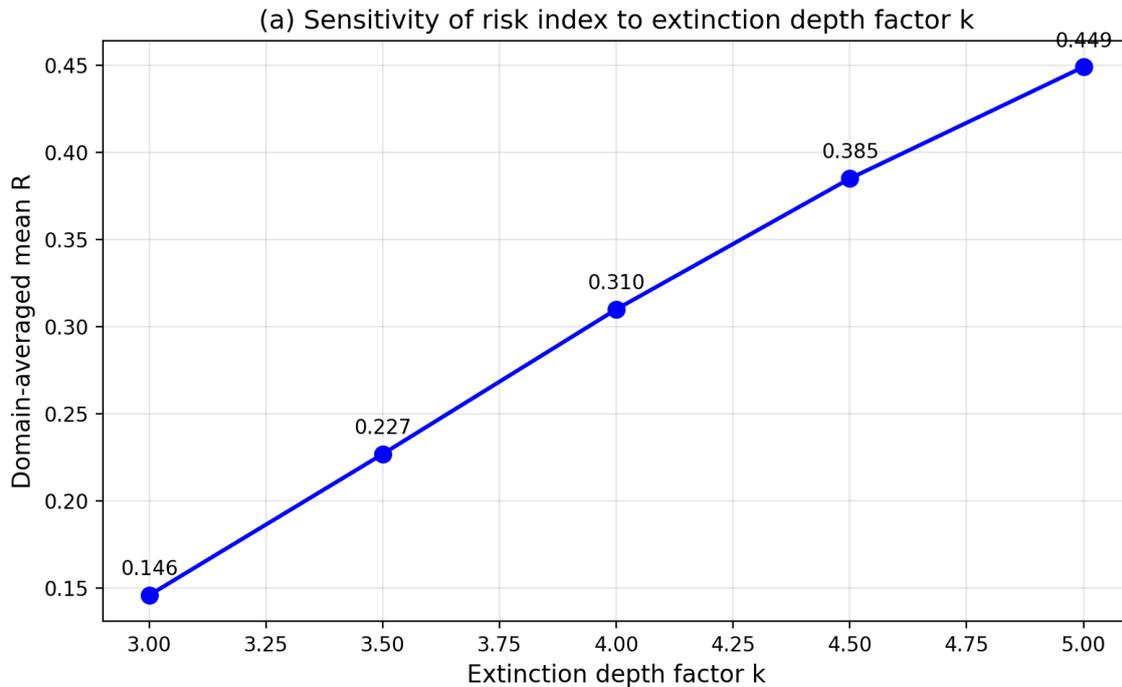
[Figure 7] *Sensitivity of domain-averaged mean risk index R to baseline scaling parameter k_d. The approximately linear response over the 2.0–3.5 m range highlights the importance of accurate critical depth parameterization.*

4.5 ConvLSTM groundwater level forecasting performance

4.5.1 Overall performance

The ConvLSTM model (2-layer architecture with hidden dimensions [32, 16], 3×3 kernels, and 8 input channels: GWL, precipitation, temperature, DEM, clay content, AWC, and cyclical month encodings) was trained on the 1971–2019 period and evaluated on the independent 2020–2024 test set in a 1-month-ahead forecasting mode. Table 4 summarizes the domain-aggregated performance metrics of the two baseline models.

[Table 4] *Comparison of the 1-month-ahead GWL forecasting performance for the 2020–2024 evaluation period.*

Model	RMSE (m)	MAE (m)	NSE	Pearson r
Persistence (GWL$_t$ = GWL$_{t-1}$)	0.573	0.439	0.694	0.847
Spatial-mean LSTM	1.005	0.801	0.056	0.247
Linear AR(P,T)	0.495	0.377	0.771	0.887
ConvLSTM (this study)	**0.416**	**0.319**	**0.839**	**0.916**

The GWL fields were generated using Sequential Gaussian Simulation (SGS) with 16 conditional realizations per time step (Sect. 2.2.2.2), which preserves fine-scale spatial variability rather than smoothing it as conventional kriging does. Therefore, the absolute RMSE values were higher than those obtained from the kriged fields, but the relative model comparisons were more informative.

At the 1-month horizon, ConvLSTM achieved a 27 % RMSE reduction over persistence (0.416 vs. 0.573 m) and improved the NSE from 0.694 to 0.839. This clear advantage, even at the shortest lead time, demonstrates that the ConvLSTM has learned spatiotemporal relationships beyond simple autocorrelation, capturing the joint influence of precipitation and temperature forcing on next-month GWL at each grid cell. The persistence baseline (NSE = 0.694) reflects the month-to-month autocorrelation inherent in groundwater systems, but the stochastically simulated fields retain sufficient fine-scale variability that persistence alone is insufficient. The spatial-mean LSTM, which broadcasts a single domain-averaged prediction to all grid cells, performed poorly (NSE = 0.056) because it could not resolve the substantial spatial heterogeneity of GWL across the 100×65 grid, underscoring the necessity of spatially explicit modelling. A per-cell linear autoregressive model incorporating lagged GWL, precipitation, and temperature (Linear AR(P,T): gwl(t+1) = a$gwl(t)$ + bP(t) + c*T(t) + d, fitted independently at each grid cell via least-squares) achieved RMSE = 0.495 m and NSE = 0.771, outperforming persistence but remaining substantially below the ConvLSTM (16 % higher RMSE). This demonstrates that while linear climate–GWL relationships explain a considerable fraction of the variance, ConvLSTM captures nonlinear spatiotemporal dynamics that a per-cell linear model cannot represent.

The advantage of ConvLSTM is further strengthened at longer forecast horizons, where persistence degrades rapidly, whereas ConvLSTM leverages learned climate–GWL relationships (Table 6). A sensitivity analysis across the four model configurations isolated the contribution of individual architectural and training choices (Table 5).

[Table 5] Sensitivity analysis: ConvLSTM performance across four configurations in the 2020–2024 evaluation period. Each configuration is built incrementally on the previous configuration.

Config	Input channels	Training period	Hidden dims	LR sched.	1-mo RMSE / NSE	3-mo RMSE / NSE	6-mo RMSE / NSE	12-mo RMSE / NSE
A	3 (GWL, P, T)	2002–2019	[16, 8]	No	0.422 / 0.834	0.441 / 0.818	0.466 / 0.797	0.500 / 0.771
B	6 (+ DEM, Clay, AWC)	2002–2019	[16, 8]	No	0.419 / 0.837	0.441 / 0.818	0.473 / 0.792	0.532 / 0.741
C	6	1971–2019	[16, 8]	No	0.419 / 0.836	0.439 / 0.819	0.470 / 0.795	0.529 / 0.744
D	8 (+ sin/cos month)	1971–2019	[32, 16]	Yes	0.416 / 0.839	0.433 / 0.824	0.455 / 0.807	0.495 / 0.776

The A-to-B comparison isolates the effect of adding static spatial channels (DEM, topsoil clay content, and available water capacity): at the 1-month horizon, static channels produce a modest improvement (RMSE 0.422 to 0.419), but at longer leads, the additional input complexity degrades autoregressive performance (12-month RMSE 0.500 to 0.532), suggesting that with limited training data (204 samples), the extra parameters are not well constrained. The B-to-C comparison tests the impact of extending the training period from 18 years (2002–2019) to 48 years (1971–2019): improvements are marginal but consistent, with the 12-month NSE increasing from 0.741 to 0.744. The C-to-D

comparison evaluates the combined effect of cyclical temporal encoding, larger hidden dimensions ([32, 16] vs. [16, 8]), and adaptive learning rate scheduling. This combination produced the largest gains, particularly at longer lead times (12-month RMSE 0.529 to 0.495, NSE 0.744 to 0.776), indicating that explicit temporal context and increased model capacity are more impactful than additional static spatial features for autoregressive GWL forecasting.

[Table 6] Multi-step autoregressive GWL forecasting performance: ConvLSTM (Config D) versus persistence on the 2020–2024 evaluation period. ConvLSTM uses autoregressive inference, feeding each prediction back as an input for the next step.

Lead time	Persistence RMSE (m)	Persistence NSE	ConvLSTM RMSE (m)	ConvLSTM NSE
1 month	0.573	0.694	0.416	0.839
3 months	0.654	0.600	0.433	0.824
6 months	0.755	0.468	0.455	0.807
12 months	0.732	0.510	0.495	0.776

ConvLSTM outperformed persistence at all tested lead times, with the advantage increasing from a 27 % RMSE reduction at 1 month to 40 % at the 6-month horizon (0.455 vs. 0.755 m). At 3 months, the RMSE reduction was 34 % (0.433 vs. 0.654 m), and the NSE improved from 0.600 to 0.824. These intermediate horizons (3–6 months) represent the practically relevant forecast range for groundwater management decisions, where persistence has degraded substantially, but ConvLSTM still maintains a high skill (NSE > 0.80). Even at the 12-month horizon, ConvLSTM retains a 32 % RMSE advantage over persistence (0.495 vs. 0.732 m) and maintains NSE = 0.776, well above the 0.7 threshold considered "good" in hydrological applications (Moriasi et al., 2007), indicating that autoregressive error accumulation remains bounded through the full annual cycle.

The scatter plot of the predicted versus observed GWL (Fig. 8b) shows a tight clustering around the 1:1 line across the entire range of GWL depths. The model performed particularly well in the 1–4 m depth range, which contains the majority of observations and spans the critical zone for the salinization risk assessment.
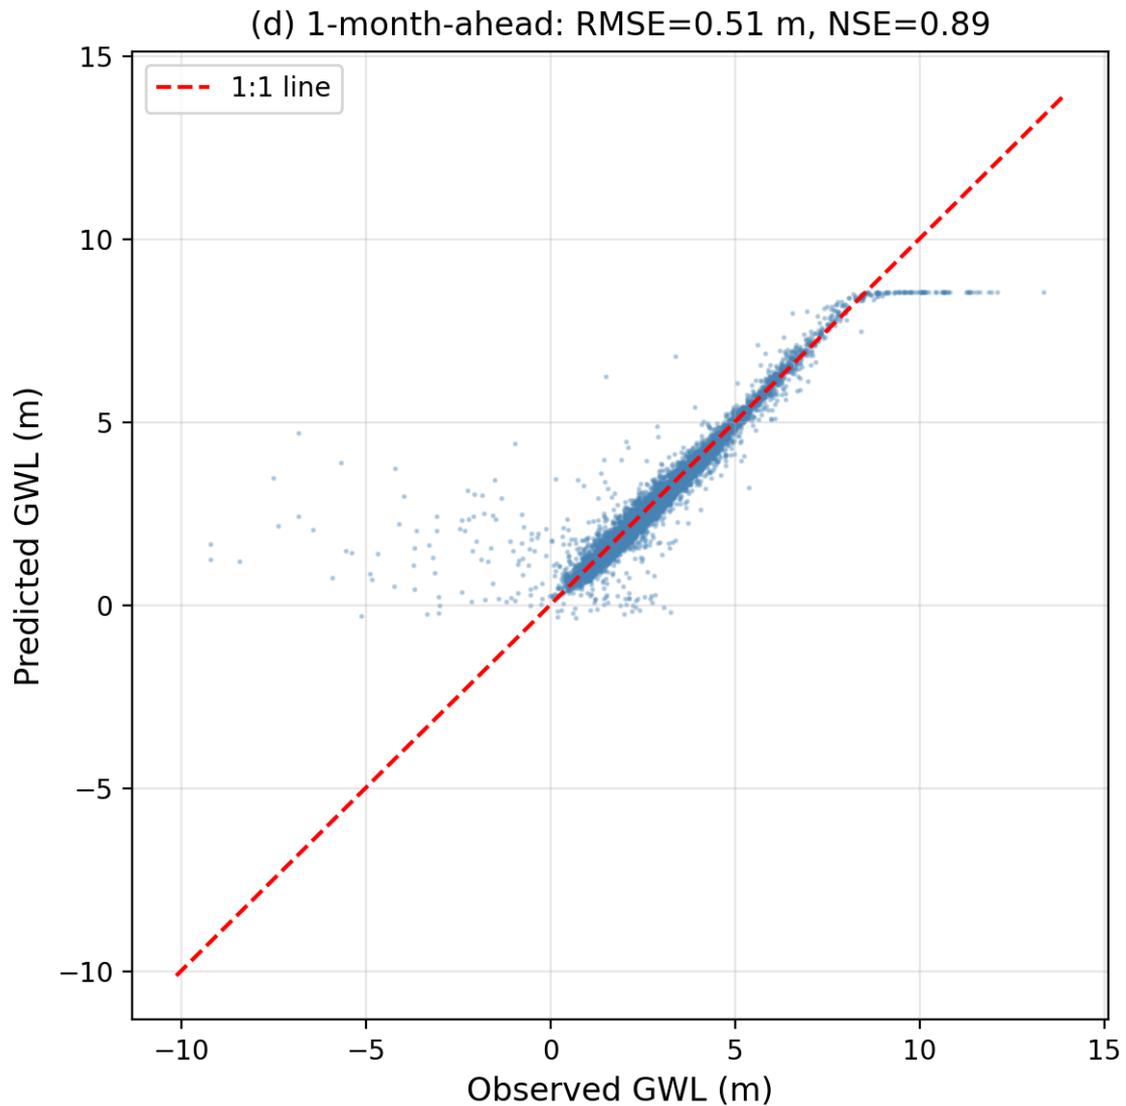
[Figure 8b] Scatter plot of the 1-month-ahead predicted versus observed GWL depth for the full evaluation set. The tight clustering around the 1:1 line (red dashed) confirms the high predictive skill.

4.5.2 Spatial distribution of performance

The spatial distribution of the 1-month RMSE (Fig. 8a) revealed that the model performed best in the central and northern regions of the study area. The performance was weakest along the domain boundaries and in areas with complex hydrogeological settings close to the river channels.

The spatial heterogeneity in model performance is primarily attributable to: (1) sparse monitoring well coverage in peripheral zones, which reduces the quality of the interpolated GWL training data; (2) rapid lateral exchange near regulated river channels that is not explicitly represented in the forcing variables; and (3) potential anthropogenic groundwater abstraction that introduces non-climate-driven variability that is not captured by precipitation and temperature inputs.
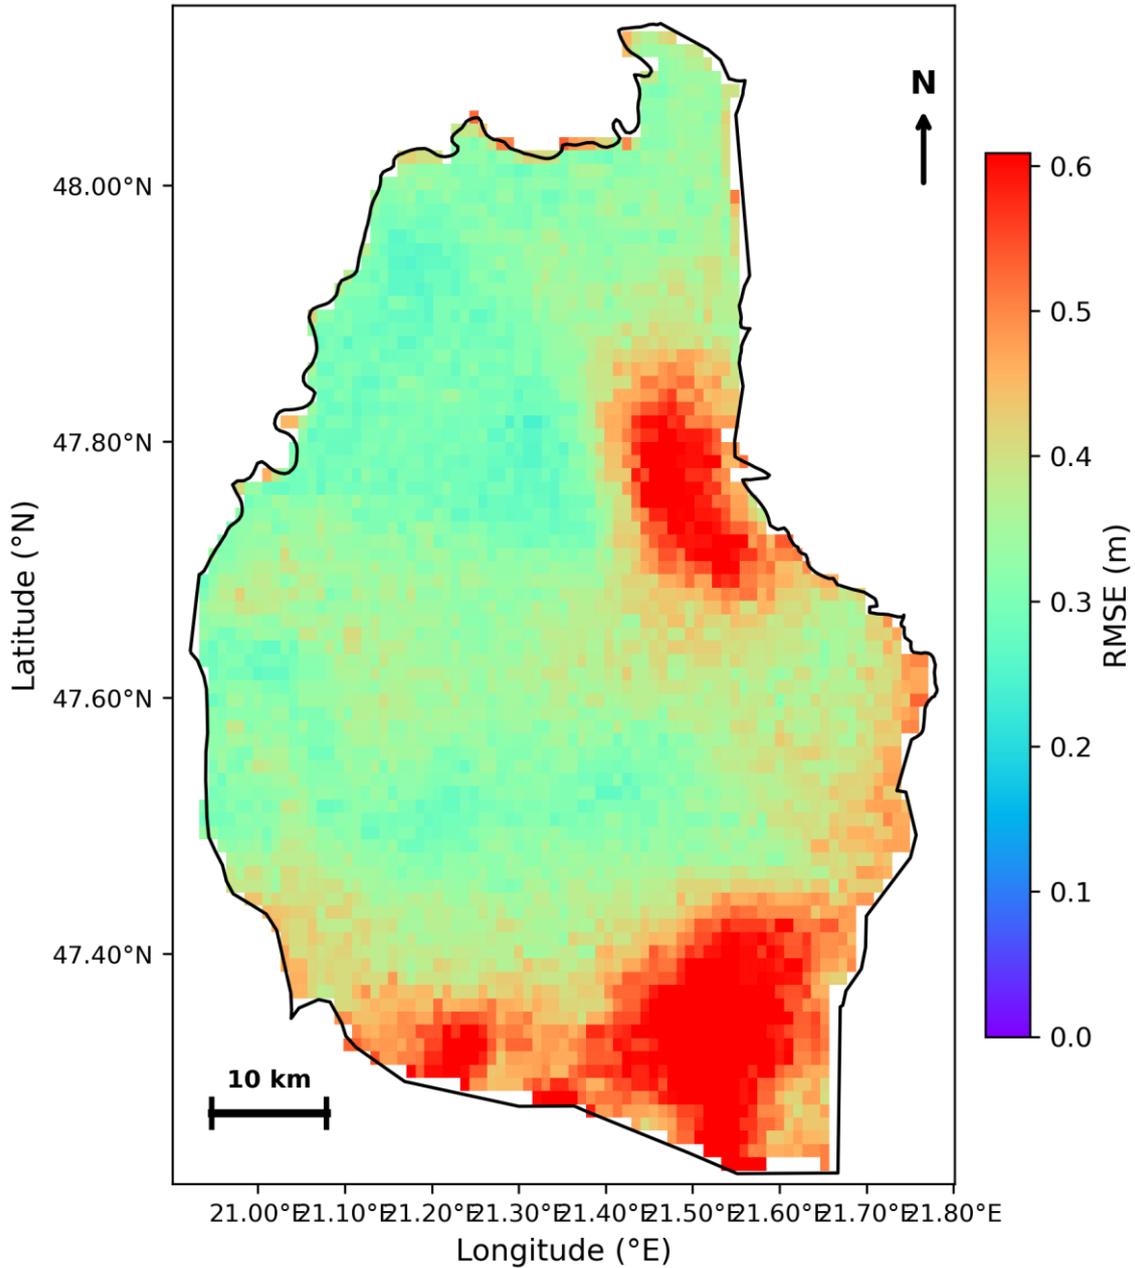
[Figure 8a] Spatial distribution of 1-month-ahead RMSE (m) across the TIKEVIR domain. Lower values (yellow) indicate a higher forecast accuracy.

4.5.3 Temporal performance

The domain-averaged time series comparison (Fig. 8c) confirmed that ConvLSTM closely tracked the observed seasonal cycle and the interannual variability. The model captured the amplitude of seasonal fluctuations and reproduced the timing of the GWL peaks and troughs with high fidelity. Minor discrepancies appear during rapid transient events and at the boundaries of the evaluation period, where the 12-month lookback window may not fully capture the antecedent conditions.
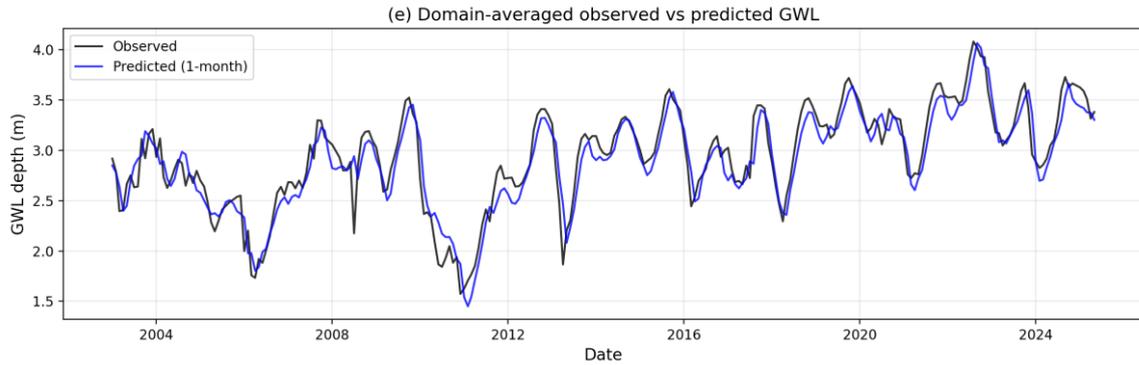
[Figure 8c] Domain-averaged observed (black) versus 1-month-ahead predicted (blue) GWL depth over the evaluation period. The model successfully captured seasonal and interannual variability.

### 4.5.4 Prediction uncertainty via Monte Carlo dropout

To quantify the epistemic uncertainty of the ConvLSTM predictions, Monte Carlo (MC) dropout inference (Gal and Ghahramani, 2016) was performed with 50 stochastic forward passes per evaluation time step, retaining the inter-layer dropout ($p = 0.1$) active during inference. For each of the 60 evaluation time steps (January 2020 – December 2024), an ensemble of 50 GWL predictions was used to compute per-cell uncertainty statistics, which were then propagated through the dynamic risk model to assess uncertainty in the risk space.

The time-averaged GWL prediction standard deviation (Fig. 9a) ranged from 0.018 to 0.074 m across the domain (mean = 0.036 m), indicating that the ConvLSTM's epistemic uncertainty was small relative to the prediction RMSE (0.416 m). The highest uncertainty occurred in the southern floodplain areas, where spatial heterogeneity in GWL was greatest and monitoring well density was lowest, consistent with the spatial RMSE pattern (Fig. 8a). The low absolute uncertainty values (mean SD = 0.036 m) confirmed that the model converged to a well-constrained solution and that dropout regularization did not excessively perturb the predictions.

Propagating the MC ensemble through the risk model yields two complementary uncertainty diagnostics: the probability of high risk, $P(R \geq 0.5)$, computed as the fraction of (time step × MC sample) pairs exceeding the threshold (Fig. 9b), ranges from 0.000 to 0.536 across the domain (mean = 0.031). The low domain-averaged probability indicates that the risk classification is robust for the majority of the domain; however, the maximum value of 0.536 in the discharge zones indicates that these cells straddle the high-risk threshold and are sensitive to prediction uncertainty. The risk credible interval width (P90 – P10 of the MC risk ensemble, time-averaged; Fig. 9c) ranged from 0.000 to 0.091 (mean = 0.014), confirming that the risk index was well-constrained even in the presence of model uncertainty. The narrow credible intervals arise because the nonlinear risk function compresses GWL uncertainty: cells with GWL well above or well below d_{crit} yield risk values near 0 or 1 regardless of the MC perturbation, whereas only cells in the transitional zone near d_{crit} exhibit appreciable risk uncertainty.
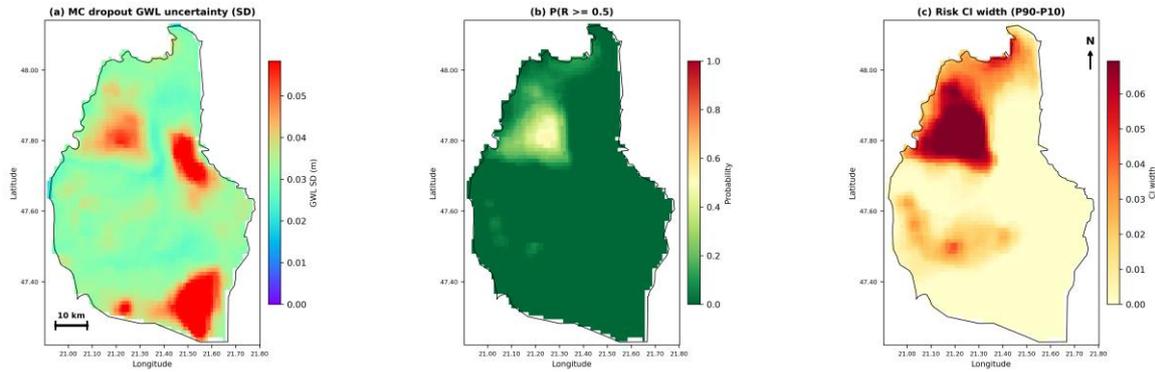
[Figure 9] MC dropout uncertainty maps (50 stochastic forward passes, 60 evaluation time steps): (a) time-averaged GWL prediction standard deviation (m), (b) probability of high risk P(R > 0.5), (c) risk credible interval width (P90 – P10). The narrow uncertainty bands confirm robust risk classification across the domain.

4.5.5 Point-scale validation against independent well observations

To verify that the ConvLSTM learns genuine hydrogeological patterns rather than artifacts introduced by the SGS interpolation pipeline, 1-month-ahead predictions were compared with raw groundwater level observations from the Hungarian national monitoring network over the 2020–2024 evaluation period. Well data were extracted from the operational database (Talajviz), yielding 97 monitoring wells inside the study domain with a total of 5,635 monthly observation–prediction pairs for training the model.

At the grid level (all valid cells), ConvLSTM achieved the same metrics as those reported in Section 4.5.1 (RMSE = 0.416 m, NSE = 0.839). At the well locations, all methods exhibited substantially higher RMSE: ConvLSTM 1.958 m, persistence 1.990 m, and critically, the SGS grid values themselves 1.982 m against the raw observations. The near-identical well-level RMSE of the ConvLSTM (1.958 m) and SGS training target (1.982 m) demonstrate that the model faithfully reproduces the gridded field without introducing additional errors beyond the inherent grid representation gap. The systematic offset between gridded and point-scale values was small (mean bias = -0.20 m for SGS, -0.23 m for ConvLSTM), confirming that the ~2 m well-level RMSE was dominated by the point-to-grid representativeness mismatch: at 1 km resolution, a single grid cell represented heterogeneous microtopography and local hydraulic gradients; thus, point-to-grid disagreement dominated the error even when the gridded field was internally accurate. This mismatch is further compounded by the quality filtering applied during SGS conditioning (IQR-based outlier removal, minimum record length requirements, and trend capping; Sect. 2.2.2). The per-well RMSE map (Fig. 12b) confirms that the spatial pattern of prediction accuracy at well locations closely mirrors the gridded RMSE distribution (Fig. 8a), with no evidence of spatially coherent artifacts that would indicate systematic learning of the interpolation errors.
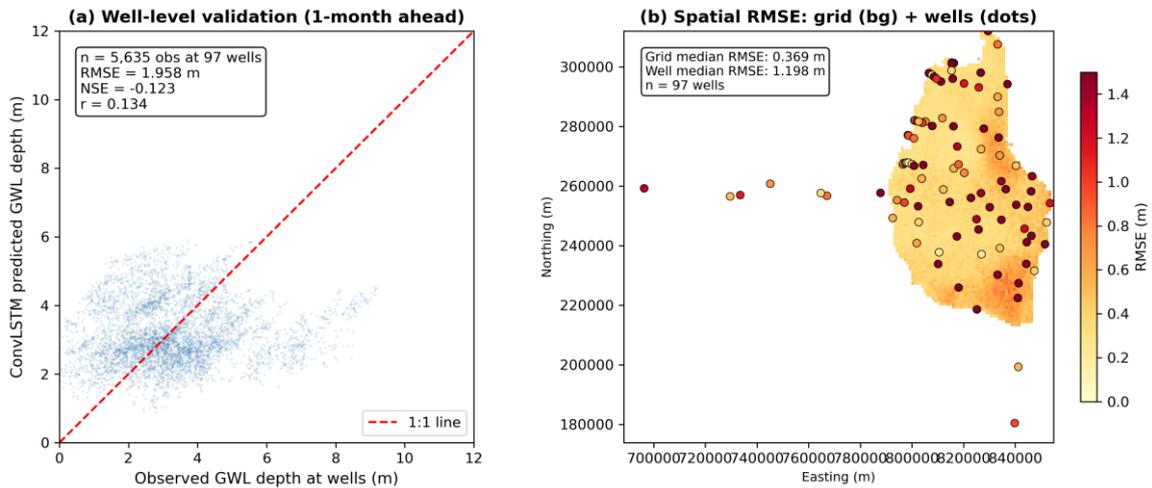
[Figure 12] Point-scale well validation: (a) scatter plot of ConvLSTM 1-month-ahead predictions versus raw well observations (n = 5,635 pairs at 97 wells); (b) spatial RMSE distribution at grid level (background) with per-well RMSE overlaid (dots).

4.5.6 Seasonal error characteristics

The error growth behavior of ConvLSTM was further analyzed by stratifying the autoregressive RMSE by target season (DJF, MAM, JJA, and SON) across lead times of 1–12 months (Fig. 10). This decomposition revealed that model skill varied systematically with the season being predicted, reflecting differences in the predictability of GWL dynamics across the hydrological year.

At the 1-month horizon, ConvLSTM achieved the lowest RMSE for spring targets (MAM: 0.389 m) and the highest for autumn targets (SON: 0.430 m). The spring advantage reflects the relatively smooth, recharge-dominated GWL rise during March–May, which is predictably followed by antecedent winter precipitation. Autumn targets are more challenging because they occur during the transition from the ET-driven drawdown phase to the onset of winter recharge, a period of competing forces that produces greater uncertainty in the GWL trajectory than spring targets.

As the lead time increased, all seasons showed gradual RMSE growth, but the rate of degradation varied: winter targets (DJF) exhibited the steepest error growth (RMSE increased from 0.420 to 0.505 m over 12 leads), whereas summer targets (JJA) showed the most moderate growth (0.405 to 0.491 m). ConvLSTM consistently outperformed persistence at all leads and seasons, with the skill advantage widening most dramatically in winter and spring, where persistence degraded rapidly (DJF persistence RMSE reached 0.757 m at a 3-month lead versus 0.445 m for ConvLSTM).

The seasonal error structure has practical implications for operational forecasting: spring salinization risk forecasts (the March–June critical window identified in Section 4.3.2) benefit from the period of highest model skill, with MAM RMSE remaining below 0.49 m even at the 12-month horizon. This fortunate alignment between the seasons of greatest management relevance and highest forecast accuracy enhances the practical utility of the coupled risk forecasting system.
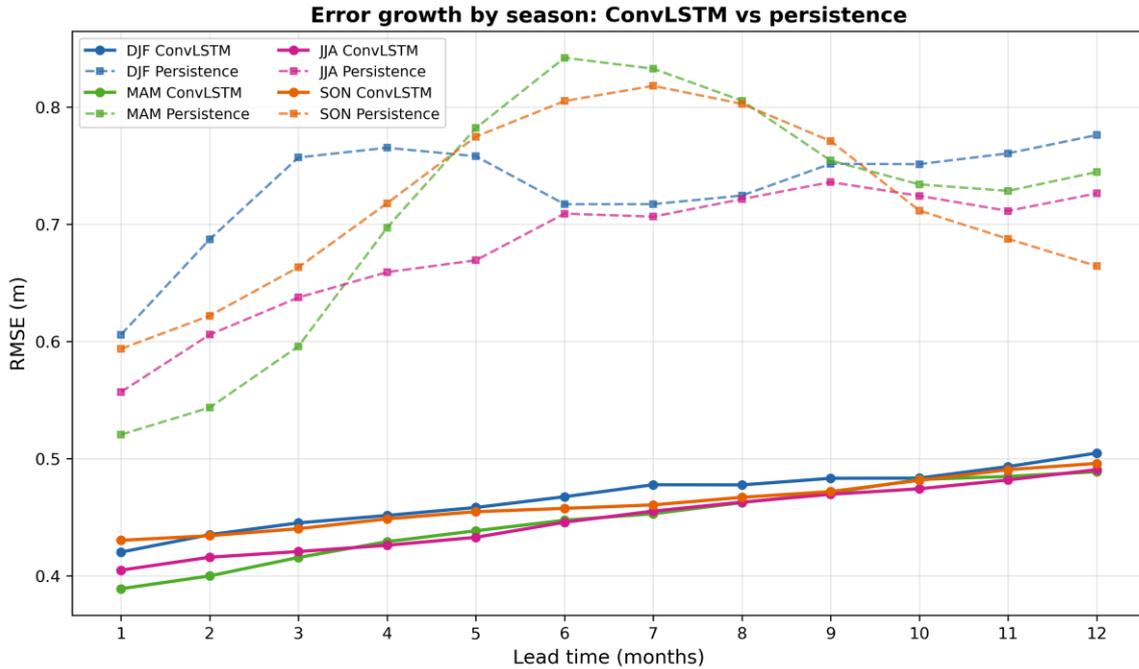
[Figure 10] Error growth curves stratified by target season (DJF, MAM, JJA, and SON) for both ConvLSTM (solid lines with circles) and persistence (dashed lines with squares). ConvLSTM outperformed persistence at all leads and seasons, with the largest advantage in winter and spring.

4.6 Coupled salinization risk forecasts

4.6.1 Framework architecture

The coupled forecasting system operates in two stages. First, the trained ConvLSTM generates spatially explicit GWL forecasts using autoregressive multi-step inference: given the trailing 12 months of observed GWL and climate forcing (precipitation and temperature), the model predicts GWL for the next month. This prediction is then fed back as input for subsequent months, allowing forecasts to extend to arbitrary horizons (tested up to 24 years, 2026–2050, using FORESEE RCP 4.5 climate projections as future forcing). Uncertainty increases with the horizon; results beyond the validated 1– 12 month range are illustrative of spatial patterns rather than exact magnitudes of the relationships.

Second, the predicted GWL fields are passed through the dynamic risk model (Eq. 9) to compute $R_{forecast}(x, y, t+h)$ at each forecast lead time h. This coupling preserves the full spatial resolution (1 km) and captures the nonlinear relationship between GWL and risk, which cannot be obtained by directly forecasting R.

4.6.2 Forecast capabilities

The dashboard implementation allows for the interactive exploration of forecast scenarios through the Forecast Horizon sidebar control, which specifies the target year (2026–2050). For each selected horizon, the system displays the following information:

4. **Predicted GWL maps** with a monthly time slider, enabling visual inspection of the spatial patterns and temporal evolution of the forecasted water table.
5. **Predicted critical depth maps** showing how the dynamic critical depth evolves under forecasted GWL conditions.
6. **Predicted salinization risk maps**, presenting the coupled risk index $R_{forecast}$ with the same interactive click-to-timeseries functionality as the historical analysis.

The autoregressive inference scheme introduces cumulative uncertainty at longer lead times because the error of each prediction step propagates into subsequent inputs. This is a known limitation of autoregressive deep learning forecasts and motivates the use of ensemble approaches (e.g., Monte Carlo dropout) for uncertainty quantification in operational deployments.

4.6.3 Validation context

The multi-step evaluation (Table 6) demonstrates that ConvLSTM outperforms persistence at all tested lead times, with 27–40 % RMSE reduction through 12-month horizons. At the 1-month horizon, the coupled risk forecast directly inherits the ConvLSTM's skill (NSE = 0.839). At longer horizons, autoregressive error accumulation causes gradual NSE degradation (from 0.839 to 0.776 at 12 months); however, the spatial patterns of risk (hotspots vs. low-risk areas) remain robust because they are controlled primarily by the static landscape configuration (topography, soil type, and proximity to discharge zones) rather than by the precise GWL value.

4.7 Summary of key findings

The principal findings of this analysis are as follows:

(i) Salinization risk is phase-controlled, peaking in April rather than at peak ET demand. The domain-averaged risk index reached its maximum in April (R = 0.338), when the water table was shallowest, and not in July–August, when ET peaked. This 3–4-month phase offset between peak risk and peak ET defines a "window of vulnerability" during which the water table remains within reach of evaporative demand, while atmospheric forcing intensifies. This finding has direct implications for the timing of salt-management interventions.

(ii) Nearly 30 % of the observation period experienced above-critical-depth conditions on average, but 19.2 % of the domain was in the "high" or "critical" risk category. The bimodal spatial distribution (40.6 % of the domain at "no risk" and 9.1 % at "critical") indicates that salinization risk in the TIKEVIR region is concentrated in well-defined floodplain discharge zones rather than being distributed uniformly. Spatial targeting is valuable for prioritizing monitoring and intervention.

(iii) ConvLSTM outperformed the persistence and linear baselines at all tested lead times, enabling credible coupled risk forecasting. The ConvLSTM achieves 27–40 % RMSE reduction over persistence across 1–12 month horizons (Tables 4 and 6), with a per-cell linear AR baseline confirming that the nonlinear spatiotemporal representations account for an additional 16 % RMSE reduction beyond linear climate–GWL relationships. Monte Carlo dropout analysis confirmed a narrow prediction uncertainty (mean risk credible interval width = 0.014), and seasonal error decomposition showed that the spring critical window (MAM) coincided with peak model skill (1-month RMSE = 0.389 m). Among the model configurations (Table 5), cyclical temporal encoding and increased capacity contributed the most to performance gains at longer lead times. Point-scale validation against raw monitoring well observations (97 wells, 5,635 observation pairs) confirmed that ConvLSTM did not introduce errors beyond the inherent grid representation gap (well-level RMSE: ConvLSTM 1.958 m vs. SGS 1.982 m; mean bias -0.23 m), verifying that the model learned genuine hydrogeological patterns rather than SGS interpolation artifacts.

Two additional noteworthy findings are: (iv) the secular water table deepening of 0.051 m per decade (p = 0.032), while producing a weak overall reduction in risk, simultaneously masks local intensification in persistently shallow areas; and (v) the 2022 extreme drought paradoxically produced the lowest risk year on record (peak monthly R = 0.090), demonstrating the non-monotonic relationship between

drought severity and salinization risk: extreme droughts suppress risk by pushing the water table below the capillary zone.

The results are discussed in the context of the existing literature, methodological limitations, and practical implications in Section 5 below.

5 Discussion

5.1 The dynamic nature of the critical groundwater depth

The results demonstrate that the salinization risk in the TIKEVIR region is fundamentally governed by the seasonal phase relationship between groundwater level fluctuations and atmospheric evaporative demand rather than by static soil-dependent thresholds. The 3–4-month phase offset between the GWL minimum (shallowest in April, Sect. 4.2.1) and peak ET demand (July–August) defines a phase-controlled "window of vulnerability" that, to the author's knowledge, has not been explicitly quantified for any comparable continental lowland system in the literature. This phase lag is consistent with the climate-driven periodicities identified by Garamhegyi et al. (2018) and the 7–10 month precipitation–GWL lag times documented by Illyés et al. (2023), which together indicate that the seasonal water table trajectory is largely predetermined by antecedent winter–spring recharge conditions.

A static extinction depth parameterization, as employed in the classical MODFLOW EVT and ETS packages (McDonald and Harbaugh, 1988; Banta, 2000), would miss this temporal structure. If the time-averaged GWL field was used against a fixed d_{crit} threshold, the fraction of the domain classified as "at risk" would correspond to the 29.8 % long-term exceedance frequency, but the critical information about *when* during the year this risk materializes would be lost. The segmented-linear ETS formulation allows depth-dependent flux attenuation but lacks temporal coupling between ET forcing and GWL response; it treats each stress period independently without accounting for the hydrologic memory that determines whether the water table will be above d_{crit} when ET demand arrives. The dynamic formulation presented here, in which d_{crit} varies as $d_{crit} = k_d \cdot (0.8 + 0.4 \cdot z_{GWL}/(z_{GWL} + 3))$ and is evaluated against the full monthly GWL time series, explicitly captures this coupling.

The dynamic d_{crit} formulation should be understood in the context of the established Hungarian "critical groundwater level" (kritikus talajvízszint) tradition. Szabolcs et al. (1969) originally defined the critical level as the depth above which groundwater of known salinity and ion composition causes harmful salt accumulation, and mapped it as a static, soil- and salinity-dependent parameter for irrigation planning across the Tisza irrigation systems. Bakacsi et al. (2019) subsequently produced a national-scale assessment for the broader Hungarian Great Plain, reporting critical depths of 2–4 m, depending on soil type, exchangeable sodium content, and groundwater chemistry. The time-averaged d_{crit} range in this study (2.10–2.83 m for $k_d = 2.5$ m) falls within this empirically established envelope, providing indirect validation, despite the absence of direct evaporative flux measurements. The conceptual advance here is not the absolute threshold value, which is broadly consistent with the existing empirical tradition, but rather the dynamic, time-varying evaluation that enables phase-controlled risk assessment.

The critical salinization window spans March through June, a period during which the water table remains within capillary reach of the soil surface while atmospheric demand accelerates from its winter minimum. During this window, capillary fluxes sustain upward moisture movement into the unsaturated zone, where rising ET concentrates dissolved salts at the evaporative front (Brunner et al., 2008). In the TIKEVIR context, the salt source is provided by ascending saline groundwater from deeper Pannonian formations, as documented by Mádl-Szőnyi and Tóth (2009) for the discharge zones of the Great Hungarian Plain. The dynamic critical depth mechanism thus provides the transport pathway that links the deep salt reservoir to the surface expression of salinization, a connection that is temporally modulated by the seasonal GWL cycle rather than constant in time.

The spatial variability of the GWL seasonal amplitude across the domain (SD of monthly climatology ranges from 0.2 m in clay-dominated floodplains to 0.8 m in sandy interfluves) suggests that the duration of the critical window also varies. Sandy aquifers, with their higher hydraulic conductivity and faster drainage response, likely exhibit a compressed overlap window in which the water table

transitions rapidly from above to below d_{crit}. Conversely, clay-dominated systems buffer the seasonal signal and sustain above-critical depth conditions for a longer period, extending the capillary flux window but with lower flux rates per unit time. This spatial heterogeneity has direct implications for monitoring design: high-frequency (weekly) observations are needed in sandy areas to capture brief risk episodes, whereas monthly monitoring may suffice in clay zones, where risk evolves more gradually.

It is worth emphasizing that the GWL-only risk formulation (Eq. 5) does not neglect ET; rather, it captures the ET effect implicitly through the water table position. The seasonal GWL cycle is a product of the recharge–ET balance: the spring water table high reflects winter–spring recharge accumulation during the low-ET season, whereas the autumn low reflects the integrated ET deficit of the growing season. Therefore, the GWL trajectory at any point in time encodes the cumulative ET signal. A sensitivity analysis was performed using an ET-weighted variant $R_{ET} = R \times f_{ET}$, where f_{ET} is the normalized Hargreaves reference ET0 climatology derived from FORESEE daily temperature extremes (Sect. Section 4.3.2 and Fig. 11), confirm this interpretation quantitatively: the ET-weighted index shifts the peak risk month from April (month 4) to June (month 6), aligning with the period when both residual water table elevation and atmospheric demand are simultaneously elevated. Crucially, the spatial correlation between time-averaged R and R_{ET} was r = 0.999, demonstrating that ET weighting modulates timing but does not alter the spatial ranking of risk hotspots. The April risk peak identified by the GWL-only index thus represents a conservative upper bound on the timing of maximum salinization potential: the actual salt flux peaks approximately two months later, but the *transport capacity* peaks in spring. This is the period during which the water table is above d_{crit}.

5.2 Salinization risk patterns and their drivers

The time-averaged risk map (Fig. 5a) revealed a strongly bimodal spatial structure, with 40.6 % of the 4,041 km² domain classified as "no risk" (R < 0.1) and 9.1 % in the "critical" category (R >= 0.7), while 25.7 % occupies the "low" transitional zone (Table 2). The highest persistent risk values were concentrated in the low-lying discharge zones along the Körös and Berettyó floodplains and in the eastern and southern margins of the study area, coinciding with areas mapped as solonchak and solonetz soils in the AGROTOPO database. This spatial agreement between the model-derived risk hotspots and independently mapped salt-affected soils provides a measure of validation for the dynamic risk framework, although the causal direction is circular to some degree, as shallow water tables both drive salinization and are a consequence of the landscape position that favors salt accumulation. An independent correlation analysis against the ESDAC topsoil clay content (Ballabio et al., 2016) confirmed a weak but statistically significant positive association (Spearman's $\rho = 0.150$, $p < 10^{-21}$; Fig. 6), consistent with the expectation that clay-rich floodplain soils sustain stronger capillary fluxes and coincide with shallow water-table levels. Notably, the CaCO3 content showed no correlation with the dynamic risk index ($\rho = -0.010$, $p = 0.52$), indicating that the index captures active salinization susceptibility rather than legacy carbonate accumulation, a distinction that validates the dynamic GWL-driven formulation over static soil property approaches.

Temporal trend analysis revealed a weak, non-significant decline in domain-averaged annual peak R (-0.011 per decade, p = 0.127) over the 1951–2025 record (Fig. 5d). This trend is consistent with the statistically significant secular water table deepening of 0.051 m per decade (p = 0.032, Fig. 4b), which Fehér and Rakonczai (2019) and Tran et al. (2022) independently documented across the broader Great Hungarian Plain as a consequence of climate change and land-use intensification. As the water table gradually deepens, fewer grid cells experience above-critical-depth conditions during the spring recharge peak. The two effects that govern the long-term risk trajectory, a deepening water table (which suppresses risk by pulling GWL below d_{crit}) and intensifying ET under climate warming (which amplifies capillary flux during the narrowing windows when the water table *is* above d_{crit}), appear to

partially cancel in the domain-averaged signal, explaining the non-significance of the risk trend despite the significant GWL trend. This partial cancellation masks local divergence: in areas where the water table has permanently receded below d_{crit}, risk has effectively ceased; in persistently shallow discharge zones, risk may intensify as ET rates increase without a compensating deepening trend.

The extreme droughts of 2022 and 2024 provide instructive natural experiments for this study. The year 2022 produced the lowest domain-averaged risk in the entire record (peak R = 0.090, annual mean R = 0.048), a counterintuitive result that arose because the severe drought drove the water table rapidly below d_{crit} across the entire domain, effectively shutting off the capillary transport pathway. The 2024 drought, while also severe, produced a moderately higher risk (peak R = 0.237 in February) because brief winter recharge events temporarily lifted the water table above d_{crit} before the continuing dry conditions pushed it back down. These contrasting responses demonstrate that the relationship between drought severity and salinization risk is non-monotonic: extreme droughts suppress the risk by desiccating the capillary zone, whereas moderate droughts or drought-recovery transitions, in which the water table oscillates around d_{crit}, produce the highest risk episodes. This non-monotonicity has implications for post-drought management, as the most dangerous period for salinization may be the recharge phase following drought termination, when the water table rises back into the capillary zone while accumulated salts remain close to the surface.

5.3 The irrigation–salinization feedback under climate warming

The results of this study, particularly the ConvLSTM-derived forecasts under the FORESEE RCP 4.5 projections to 2050, raise a fundamental question for agricultural water management in the TIKEVIR region: does increasing irrigation demand under a warming climate create a self-reinforcing salinization feedback loop?

This mechanism operates through three coupled pathways. First, rising temperatures increase the atmospheric evaporative demand, extending the growing season and amplifying crop water stress. In the TIKEVIR region, where rainfed agriculture dominates, this translates to increasing pressure for supplemental irrigation, a trend that has already been documented across the Great Hungarian Plain (Mezősi et al., 2016). Second, irrigation raises the local water table by increasing recharge through the unsaturated zone, bringing the GWL closer to or above the critical depth. The sensitivity analysis (Table 3) demonstrates that even modest changes in the baseline scaling parameter (from $k_d = 1.5$ to $k_d = 3.5$ m) shift the domain-averaged risk from R = 0.062 to R = 0.471, which demonstrates a 660 % increase. Analogously, irrigation-induced water table rise of comparable magnitude would push currently "safe" transitional zones (the 25.7 % of the domain in the "low" risk category) into the "moderate" or "high" risk classes. Third, the elevated water table sustains capillary flux to the surface, where the irrigation water itself, even if of acceptable quality, concentrates salts through evapotranspiration, progressively degrading soil quality.

The paradox is that climate warming simultaneously makes irrigation more necessary and more dangerous. The ConvLSTM forecasts, driven by RCP 4.5 temperature and precipitation projections, indicated continued water table deepening under the "no irrigation" baseline scenario, which should, in principle, reduce the salinization risk. This is the trajectory captured by the secular trend of 0.051 m/decade, as documented in Section 4.2.2. Under this scenario, the domain transitions toward a drier equilibrium, in which fewer areas experience above-critical-depth conditions. The 2022 drought, which produced the lowest risk year on record, offers an extreme preview of this trend.

The critical question is what happens when irrigation is superimposed on this drying baseline. Irrigation acts as a localized reversal of this deepening trend, creating a wetter-than-natural water table regime in agricultural areas where salt management is critical. In the TIKEVIR context, where 19.2 % of the

domain already falls in the "high" or "critical" risk categories under historical conditions (Table 2), even a modest irrigation-induced GWL rise could push additional areas above the critical depth during the March–June critical window. The timing of this irrigation is particularly unfavorable, as it coincides exactly with the period when the water table is at its shallowest and salinization risk peaks (Fig. 5b).

This irrigation–salinization feedback creates a management dilemma with no simple resolution. Restricting irrigation to protect against salinization reduces agricultural productivity in warming climates. Maintaining or increasing irrigation to sustain yields accelerates salt accumulation in vulnerable soils, potentially rendering them unproductive within a few decades. The dynamic risk framework developed in this study provides a quantitative basis for navigating this trade-off: by identifying the spatial zones where irrigation-induced salinization risk is highest (the "high" and "critical" areas in Fig. 5a), water managers can target irrigation investments toward areas where the water table is sufficiently deep that irrigation recharge will not lift it above d_{crit}, while deploying salt-tolerant varieties and drainage improvements in the vulnerable zones.

The ConvLSTM forecasting system offers additional value. By providing spatially explicit GWL forecasts at a monthly resolution, the coupled framework can simulate "what if" irrigation scenarios: what happens to the risk map if irrigation raises the water table by 0.5 m in a given sub-region? The sensitivity analysis quantified the response of R to changes in k_d (Fig. 7), and the autoregressive ConvLSTM architecture could, in principle, be extended to include irrigation application rates as an additional forcing variable. This scenario modelling capability transforms the tool from a passive forecasting system into an active decision-support instrument for irrigation planning under salinization constraints.

5.4 Performance and limitations of the ConvLSTM forecasting approach

At the 1-month horizon, ConvLSTM achieved an RMSE of 0.416 m, MAE of 0.319 m, and NSE of 0.839 for GWL forecasting over the 2020–2024 evaluation period, representing a 27 % RMSE reduction over persistence (RMSE = 0.573 m, NSE = 0.694; Table 4). Unlike kriging-based approaches that smooth fine-scale spatial variability, the GWL fields used here were generated via Sequential Gaussian Simulation (SGS) with 16 conditional realizations per time step (Sect. 2.2.2.2), which preserves the full spectrum of spatial variability in the observed well data. This methodological choice increased the absolute RMSE values compared to the kriged fields but provided a more stringent and realistic evaluation. Because the SGS fields honor the observed well values exactly (conditional simulation) and reproduce the target spatial covariance structure, the learned spatiotemporal patterns reflect genuine hydrogeological variability rather than interpolation artifacts, and the stochastic component contributes short-range variability that is consistent with the variogram model. This was confirmed by point-scale validation at 97 monitoring wells with 5,635 observation pairs over the evaluation period (Sect. 4.5.5, Fig. 12), ConvLSTM achieved a well-level RMSE of 1.958 m, which was nearly identical to the SGS grid values themselves (1.982 m) and slightly better than persistence (1.990 m). The small mean bias (-0.23 m for ConvLSTM, -0.20 m for SGS) confirms that the ~2 m well-level RMSE is dominated by the inherent point-to-grid representativeness mismatch at 1 km resolution, not by systematic artifact learning. The ConvLSTM's clear advantage over persistence at even the 1-month horizon confirms that the model has learned meaningful spatiotemporal relationships rather than benefiting from the artificial smoothness of deterministic interpolation. The advantage strengthens at longer horizons: at 3 months, 34 % RMSE reduction (0.433 vs. 0.654 m); at 6 months, 40 % (0.455 vs. 0.755 m); and even at 12 months, 32 % (0.495 vs. 0.732 m; Table 6). The bounded error growth through the full annual cycle indicates that ConvLSTM captures the seasonal GWL dynamics robustly, without the catastrophic divergence sometimes observed in autoregressive deep learning forecasts.

A sensitivity analysis across four model configurations (Table 5) revealed that the most impactful improvements came from cyclical temporal encodings and increased model capacity: the addition of sin/cos month channels, larger hidden dimensions ([32, 16] vs. [16, 8]), and learning rate scheduling (Config D vs. Config C) produced the largest gains, particularly at longer lead times (12-month NSE: 0.776 vs. 0.744). Static spatial channels (DEM, clay content, and AWC) and an extended training period (1971 vs. 2002) provided more modest improvements, indicating that ConvLSTM already captured much of the spatial heterogeneity through the GWL input channel itself.

A per-cell linear autoregressive baseline incorporating lagged GWL, precipitation, and temperature (Table 4) achieved an RMSE of 0.495 m and NSE of 0.771, confirming that climate–GWL relationships at each grid cell were partially linear, but that the ConvLSTM's spatiotemporal convolutions captured an additional 16 % RMSE reduction through nonlinear interactions and cross-cell information transfer. Monte Carlo dropout analysis (50 stochastic forward passes; Fig. 9) yields a mean GWL prediction standard deviation of only 0.036 m – roughly one-tenth of the prediction RMSE – indicating that the model's epistemic uncertainty is small and that the risk classification is robust. The propagated risk credible interval width (P90 – P10 = 0.014 on average) confirms that risk categories are well-constrained across the domain, with appreciable uncertainty only in the transitional zone near d_{crit}, where the nonlinear risk function is most sensitive to GWL perturbations. A seasonal decomposition of error growth (Fig. 10) revealed that the ConvLSTM's skill advantage over persistence was the largest in winter (DJF) and spring (MAM), and that the March–June critical salinization window coincided with the period of highest forecast accuracy (MAM 1-month RMSE = 0.389 m). This fortunate alignment between peak management relevance and forecast skill enhances the practical value of the coupled system.

These metrics compare favorably with published benchmarks: Wunsch et al. (2021) reported NSE values of 0.7–0.9 for 1-month LSTM forecasts at individual monitoring wells in Germany, while Sun et al. (2019) achieved $R^2 > 0.8$ for GRACE-based spatiotemporal GWL predictions. However, the comparison is not fully commensurate, as the present model predicts a full spatial field (4,041 cells at 1 km resolution) on stochastically simulated GWL grids rather than individual well time series, introducing additional complexity from spatial simulation uncertainty and spatial heterogeneity in aquifer properties.

The spatial distribution of the forecast error (Fig. 8a) revealed systematic patterns that informed the reliability of the coupled risk forecasts. Performance is strongest in the central portions of the domain, where monitoring well density is highest and the GWL signal is dominated by climate-driven recharge and ET, which are the processes explicitly represented in the model inputs. The largest errors (RMSE > 1.0 m) occurred along the domain boundaries and near the regulated river channels, where the lateral exchange with surface water introduced rapid GWL fluctuations that were not captured by precipitation and temperature forcing. Anthropogenic abstractions, which were not included as model inputs, contributed additional unmodelled variance in areas with irrigation pumping. These spatial patterns indicate that the coupled risk forecast is most reliable in the interior floodplain areas where salinization risk is also highest, a fortunate coincidence that increases the practical utility of the system.

The model's ability to forecast extreme events is a critical consideration for operational deployment. The training period (1971–2019) does not include the 2022 and 2024 droughts, which fall entirely within the evaluation set (2020–2024), providing a genuine test of the model's ability to generalize to extreme events not encountered during training. The domain-averaged time series comparison (Fig. 8c) indicates that the model tracks interannual variability well, including the seasonal amplitude variations associated with wet and dry years. However, a fundamental limitation of data-driven approaches remains: ConvLSTM extrapolates poorly beyond the range of conditions encountered during training.

If future water table deepening exceeds the historical range, as projected under high-emission scenarios, the model's predictions may underestimate the depth and duration of below-critical-depth episodes. Ensemble approaches combining the ConvLSTM with physics-based constraints (Nearing et al., 2021; Reichstein et al., 2019) represent a promising direction for improving robustness under non-stationary conditions.

For the coupled risk application, the nonlinear risk function (Eq. 9) partially mitigates GWL forecast errors by compressing them into a bounded [0, 1] risk interval. A GWL forecast error in an area where the water table is well above d_{crit} (R near 1.0) or well below it (R near 0.0) has negligible effect on the risk classification; the same error in a transitional zone near d_{crit} can shift the classification by one or more categories. This suggests that validation efforts should focus on the transitional zones rather than on domain-averaged metrics.

5.5 Implications for water management in the TIKEVIR region

The coupled forecasting system provides actionable forecasts at 3–6-month lead times (Table 6), with ConvLSTM maintaining NSE > 0.80 through 6 months and achieving 34–40 % RMSE reduction over persistence, enabling spatially targeted management responses during the March–June critical window. This lead time is sufficient for practical interventions, including the adjustment of irrigation scheduling to avoid spring applications in areas where the water table is approaching d_{crit}, activation of subsurface drainage infrastructure in chronic hotspots, and selection of salt-tolerant crop varieties for the upcoming growing season. Water management practices in the TIKEVIR region currently lack predictive tools for salinization risk, and the framework developed here represents a concrete step toward anticipatory water management.

The strong spatial concentration of risk (19.2 % of the domain in "high" or "critical" categories, Fig. 5a) argues for spatially targeted rather than uniform intervention strategies. In a resource-constrained setting, directing drainage investments, soil monitoring, and irrigation restrictions to the 776 km² identified as persistently at risk represents a far more efficient allocation than domain-wide programs. The risk map provides a rational basis for spatial prioritization, whereas the temporal forecast component identifies the years and months that require the most urgent attention.

The irrigation–salinization feedback discussed in Section 5.3 underscores the need for the forecasting tool to be embedded within an integrated water-allocation framework. The TIKEVIR Water Management Cooperative represents an existing institutional structure into which the dynamic risk index can be integrated as an additional decision layer, informing the trade-offs between agricultural water use, ecological flow requirements, and long-term soil quality preservation. The interactive Streamlit dashboard implementation (Section 3.7) serves as a prototype for this integration, enabling stakeholders to explore forecast scenarios, adjust risk parameters, and evaluate the consequences of management decisions in near-real time.

5.6 Limitations and future directions

Several data limitations affect this analysis. The gridded GWL product was derived from the Sequential Gaussian Simulation (SGS) of monitoring well observations, which preserves spatial variability but introduces stochastic uncertainty that propagates into both the risk index and ConvLSTM training data, particularly in areas with sparse well coverage along the domain periphery. The absence of spatially distributed groundwater electrical conductivity (EC) observations at most monitoring locations precluded the inclusion of an explicit soil susceptibility factor, meaning that the risk index captures salinization potential based on water table position alone, without accounting for local variations in groundwater chemistry or salt content. The 0.1° resolution of the FORESEE climate grid is too coarse

to resolve local topographic or land-use effects on ET, and the static CORINE land cover classification does not capture interannual crop rotation effects on crop coefficients. Each of these limitations could be addressed with targeted data collection: densification of the monitoring network in data-sparse peripheral zones, EC sensors at existing wells, and integration of Sentinel-2-derived crop type maps for dynamic K_c estimation in the future.

Methodologically, the empirical GWL-dependent formulation for d_{crit}, although physically motivated, has not been validated against field measurements of evaporative flux extinction. The resulting d_{crit} range (2.10–2.83 m for $k_d = 2.5$ m) is, however, consistent with the empirically mapped critical groundwater depths of 2–4 m reported by Bakacsi et al. (2019) for comparable soil and salinity conditions across the broader Hungarian Great Plain, and with the texture-based extinction depths of Shah et al. (2007), who demonstrated that extinction depth depends on soil texture through the van Genuchten parameters in a strongly non-linear manner. The sensitivity analysis (Table 3) confirmed that the choice of k_d parameterization was the primary source of methodological uncertainty. The autoregressive multi-step forecasting strategy accumulates errors at longer lead times, which is a well-known limitation of recursive deep learning forecasts, although the Monte Carlo dropout analysis (Sect. 4.5.4) demonstrates that the resulting risk classifications remained robust (mean CI width = 0.014). The absence of groundwater abstraction data as a model input represents perhaps the most consequential limitation for operational deployment, as pumping creates localized drawdowns that ConvLSTM cannot explain from meteorological forcing alone.

Future work should pursue the following four directions. First, the incorporation of Sentinel-2-derived bare soil salinity indices and MODIS-based actual ET products would enable independent validation of the risk index against remotely sensed surface indicators. Second, coupling with EURO-CORDEX climate projections under multiple RCP/SSP pathways would enable long-term (2050/2100) salinization risk scenarios, quantifying the irrigation–salinization feedback under various emission trajectories. Third, extending to a physics-informed neural network (PINN) architecture that constrains GWL predictions with the Boussinesq equation or mass balance considerations would improve robustness during out-of-distribution events. Fourth, the operational deployment of the dashboard within the TIKEVIR Water Management Cooperative during the 2026–2027 growing seasons would provide real-world validation and iterative refinement of the forecasting framework.

6 Conclusions

This study introduces a dynamic, time-varying formulation of the critical groundwater table depth for assessing salinization risk. This approach replaces the traditional static extinction depth parameterization in groundwater models with a groundwater level (GWL)-dependent formulation that effectively captures the phase relationship between seasonal evapotranspiration demand and groundwater level fluctuations. Applied to a 4,041 km² area of the TIKEVIR region in southeastern Hungary, utilizing 74 years of gridded groundwater monitoring data (1951–2025) and FORESEE climate projections, the framework identified a critical salinization window centered between March and June. This window is driven by a 3–4 month phase offset between the water table maximum in April and the evapotranspiration peak in July–August. A ConvLSTM neural network, trained on the 1971–2019 spatiotemporal record, achieved a root mean square error (RMSE) of 0.416 m and a Nash-Sutcliffe efficiency (NSE) of 0.839 for one-month-ahead GWL forecasting over the independent 2020–2024 evaluation period, outperforming persistence by 27–40% RMSE at 1–12 month lead times. Monte Carlo dropout analysis confirmed narrow prediction credible intervals, with a mean risk credible interval width of 0.014. Point-scale validation at 97 monitoring wells confirmed that the model does

not learn small-scale geostatistical (SGS) interpolation artifacts, with a well-level RMSE of 1.958 m for ConvLSTM compared to 1.982 m for SGS, and a mean bias of -0.23 m.

The primary conclusion of this study is that the risk of salinization in shallow-water-table lowland systems cannot be adequately characterized by static, soil-dependent thresholds. Instead, it is influenced by the transient interaction between the rising phase of the seasonal water table cycle and the increasing demand for evapotranspiration. In the TIKEVIR region, this interaction results in peak risk occurring in April, approximately 3–4 months prior to the maximum evapotranspiration, a temporal pattern not easily captured when extinction depth is parameterized as a fixed cell property. This dynamic approach may be applicable to other continental or semi-arid lowlands where the groundwater level and evapotranspiration seasonal cycles are phase-offset, such as similar environments in the Pannonian Basin, North European Plain, and irrigated lowlands of Central Asia and northern China. However, site-specific calibration of the kd parameter and validation against local soil salinity observations would be necessary.

The integrated forecasting framework offers actionable lead times of up to six months (NSE = 0.807, 40% RMSE reduction over persistence) for identifying emerging salinization hotspots at a 1 km grid-cell resolution. This enables spatially targeted management strategies, including drainage management, irrigation scheduling, and the selection of salt-tolerant crops. The 19.2% of the study area identified as chronically at high or critical risk constitutes a priority zone for infrastructure investment, while the remaining area benefits from anomaly-based early warning during years when groundwater levels deviate from climatological norms. The analysis also highlights a potentially self-reinforcing irrigation–salinization feedback: as climate warming increases irrigation demand, the consequent rise in the water table may counteract the secular deepening of the water table and reactivate capillary salt transport in areas that would otherwise transition toward safety under a drying baseline.

As climate projections suggest continued warming and increasing evapotranspiration demand across the Carpathian Basin, the frequency and spatial extent of episodes exceeding critical depth will depend on the balance between declining natural recharge and increasing anthropogenic recharge from irrigation practices. The dynamic risk framework and coupled forecasting system developed in this study provide a quantitative basis for navigating these trade-offs.

Code and data availability

The CriticalGWL framework's source code, including ConvLSTM model, training scripts, risk index modules, and Streamlit dashboard, is archived at https://doi.org/10.5281/zenodo.18770854 (Fehér 2026).The trained ConvLSTM model weights and normalization parameters, gridded GWL product (`gwl_interpolated.nc`, 894 monthly time steps, 1951–2025, 100 × 65 cells at 1 km resolution), and aligned training data cube (including FORESEE RCP 4.5 climate projections to 2050) were deposited alongside the code.

The FORESEE-HUN meteorological database is publicly available from the Department of Meteorology at the Eötvös Loránd University (ELTE), Budapest (https://nimfea.geo.elte.hu/FORESEE/). Groundwater level observations from the national monitoring network were obtained from the General Directorate of Water Management (OVF) upon request from the authors. The ESDAC topsoil physical properties (Ballabio et al., 2016) and EU-SoilHydroGrids (Tóth et al., 2017) are available from the European Soil Data Center (ESDAC;

https://esdac.jrc.ec.europa.eu). The AGROTOPO soil database is maintained by the Institute for Soil Sciences at the HUN-REN Center for Agricultural Research.

Author contributions

ZZF designed the study, developed the dynamic risk index formulation, implemented the ConvLSTM forecasting framework and gap-filling/SGS interpolation pipeline, performed all analyses, and wrote the manuscript.

Competing interests

The authors declare no conflict of interest.

Acknowledgements

References

Allen, R.G., Pereira, L.S., Raes, D., and Smith, M.: Crop evapotranspiration – Guidelines for computing crop water requirements, FAO Irrigation and Drainage Paper 56, Food and Agriculture Organization of the United Nations, Rome, 300 pp., 1998.

Bakacsi, Zs., Tóth, T., Makó, A., Barna, Gy., Laborczi, A., Szabó, J., Szatmári, G., and Pásztor, L.: National level assessment of soil salinization and structural degradation risks under irrigation, Hung. Geogr. Bull., 68, 141–156, https://doi.org/10.15201/hungeobull.68.2.3, 2019.(Bakacsi et al., 2019)

Ballabio, C., Panagos, P., and Montanarella, L.: Mapping topsoil physical properties at European scale using the LUCAS database, Geoderma, 261, 110–123, https://doi.org/10.1016/j.geoderma.2015.07.006, 2016.(Ballabio et al., 2015)

Banta, E.R.: MODFLOW-2000, the U.S. Geological Survey modular ground-water model – Documentation of packages for simulating evapotranspiration with a segmented function (ETS1) and drains with return flow (DRT1), U.S. Geological Survey Open-File Report 00-466, 2000.(Banta, 2000)

Brunner, P., Li, H.T., Kinzelbach, W., and Li, W.P.: Extracting phreatic evaporation from remotely sensed maps of evapotranspiration, Water Resour. Res., 44, W08428, https://doi.org/10.1029/2007WR006063, 2008.(Brunner et al., 2008)

Deutsch, C.V. and Journel, A.G.: GSLIB: Geostatistical Software Library and User's Guide, 2nd edn., Oxford University Press, New York, 369 pp., 1998.

Fehér, Z.Zs. and Rakonczai, J. Analyzing the sensitivity of Hungarian landscapes based on climate change-induced shallow groundwater fluctuations, Hung. Geogr. Bull., 68(4), 355–372, https://doi.org/10.15201/hungeobull.68.4.3, 2019.(Fehér & Rakonczai, 2019)

Fehér, Z.Zs.: CriticalGWL: Dynamic critical groundwater depth and salinization risk forecasting framework (v1.0), Zenodo [code], https://doi.org/10.5281/zenodo.18770854, 2026.

Gai, L., Shin, Y., Gusyev, M., and Hao, Z.: Graph neural networks for simulating spring discharge and capturing spatial dependence of groundwater propagation, Water Resour. Res., 59, e2022WR033814, 2023.

Gal, Y. and Ghahramani, Z.: Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In: Proceedings of the 33rd International Conference on Machine Learning, JMLR W&CP, vol. 48, 1050–1059, 2016.(Gal & Ghahramani, 2015)

Garamhegyi, T., Kovács, J., Pongrácz, R., Tóth, T., and Hatvani, I.G.: Investigation of the climate-driven periodicity of shallow groundwater level fluctuations in a Central-Eastern European agricultural region, Hydrogeol. J., 26, 677–688, https://doi.org/10.1007/s10040-017-1665-2, 2018.(Garamhegyi et al., 2017)

Hoyer, S. and Hamman, J.: xarray: N-D labeled arrays and datasets in Python, J. Open Res. Softw., 5, 10, https://doi.org/10.5334/jors.148, 2017.(Hoyer & Hamman, 2017)

Huber, P.J.: Robust regression: Asymptotics, conjectures, and Monte Carlo, Ann. Stat., 1, 799–821, https://doi.org/10.1214/aos/1176342503, 1973.(Huber, 1973)

Illyés, Cs., Tóth, M., and Szűcs, P.: Calculating the connection between precipitation and shallow groundwater, Multidiszciplináris Tudományok, 13(3), 103–112, https://doi.org/10.35925/j.multi.2023.3.29, 2023.(Illyés et al., 2023)

Kern, A., Dobor, L., Hollós, R., Marjanovic, H., Torma, Cs.Zs., Kis, A., Fodor, N., and Barcza, Z.: Seamlessly combined historical and projected daily meteorological datasets for impact studies in Central Europe: the FORESEE v4.0 and the FORESEE-HUN v1.0, Climate Services, 33, 100443, https://doi.org/10.1016/j.cliser.2023.100443, 2024.(Kern et al., 2023)

Kingma, D.P. and Ba, J.: Adam: A method for stochastic optimization, in: Proceedings of the 3rd International Conference on Learning Representations (ICLR), San Diego, 2015.

Major, P.: Síkvidéki erdők hatása a vízháztartásra [Effect of lowland forests on the water balance], Hidrológiai Közlöny, 82, 319–324, 2002.

McDonald, M. G. and Harbaugh, A. W.: A modular three-dimensional finite-difference ground-water flow model, Techniques of Water-Resources Investigations, Book 6, Chap. A1, U.S. Geological Survey (1988).

Mezősi, G., Blinova, I., Meyer, B.C., and Blanka, V.: The changing potential of irrigated agriculture in the Great Hungarian Plain, Geographica Pannonica, 20, 218–226, 2016.

Moriasi, D.N., Arnold, J.G., Van Liew, M.W., Bingner, R.L., Harmel, R.D., and Veith, T.L.: Model evaluation guidelines for systematic quantification of accuracy in watershed simulations, Trans. ASABE, 50, 885–900, https://doi.org/10.13031/2013.23153, 2007.(Moriasi et al., 2007)

Mádl-Szőnyi, J. and Tóth, J.: A hydrogeological type section for the Duna-Tisza Interfluve, Hungary, Hydrogeol. J., 17, 961–980, https://doi.org/10.1007/s10040-008-0421-z, 2009.(Mádl-Szőnyi & Tóth, 2009)

Müller, S., Schüler, L., and Zech, A.: GeoStat-Framework/GSTools v1.4: a comprehensive Python package for geostatistical modelling, Geosci. Model Dev., 15, 3161–3185, https://doi.org/10.5194/gmd-15-3161-2022, 2022.(Müller et al., 2021)

Nearing, G.S., Kratzert, F., Sampson, A.K., Pelissier, C.S., Klotz, D., Frame, J.M., Priber, C., and Gupta, H.V.: What role does hydrological science play in the age of machine learning?, Water Resour. Res., 57, e2020WR028091, https://doi.org/10.1029/2020WR028091, 2021.(Nearing et al., 2021)

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S.: PyTorch: An imperative style, high-performance deep learning library, in Advances in Neural Information Processing Systems, vol. 32, 8024–8035, 2019.(Paszke et al., 2019)

Pálfai, I.: Az Alföld aszályossága [Drought conditions of the Hungarian Great Plain], Hidrológiai Közlöny, 74, 89–97, 1994.

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., and Prabhat: Deep learning and process understanding for data-driven Earth system science, Nature, 566, 195–204, https://doi.org/10.1038/s41586-019-0912-1, 2019.(Reichstein et al., 2025)

Shah, N., Nachabe, M., and Ross, M.: Extinction depth and evapotranspiration from ground water under selected land covers, Ground Water, 45, 329–338, https://doi.org/10.1111/j.1745-6584.2007.00302.x, 2007.(Shah et al., 2007)

Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., and Woo, W.C.: Convolutional LSTM network: a machine learning approach for precipitation nowcasting, in: Advances in Neural Information Processing Systems, vol. 28, 802–810, 2015.(Shi et al., 2015)

Soylu, M.E., Istanbulluoglu, E., Lenters, J.D., and Wang, T.: Quantifying the impact of groundwater depth on evapotranspiration in a semi-arid grassland region, Hydrol. Earth Syst. Sci., 15, 787–806, https://doi.org/10.5194/hess-15-787-2011, 2011.(Soylu et al., 2010)

Streamlit Inc.: Streamlit – A faster way to build and share data apps, https://streamlit.io, last access: 15 January 2025, 2024.(Kamboj et al., 2024)

Sun, A. Y., Scanlon, B. R., Zhang, Z., Walling, D., Bhanja, S. N., Mukherjee, A., and Zhong, Z.: Combining physically based modeling and deep learning for fusing GRACE satellite data: can we learn from mismatch?, Water Resour. Res., 55, 1179–1195, https://doi.org/10.1029/2018WR023333, 2019.(Sun et al., 2019)

Szabolcs, I. (Ed.): European Solonetz Soils and Their Reclamation, Akadémiai Kiadó, Budapest, 1971.

Szabolcs, I., Darab, K., and Várallyay, Gy.: A tiszai öntözőrendszerek és a Magyar Alföld talajainak termékenysége. II. A talajvíz "kritikus" mélysége a kiskörei öntözőrendszer által érintett területen [The Tisza irrigation systems and the fertility of the soils in the Hungarian Lowland. II. The "critical depth" of the water table in the area belonging to the irrigation system of Kisköre], Agrokémia és Talajtan, 18(2), 211–220, 1969.

Szabó, B., Mészáros, J., Laborczi, A., Takács, K., Szatmári, G., Bakacsi, Zs., Makó, A., and Pásztor, L.: From EU-SoilHydroGrids to HU-SoilHydroGrids: A leap forward in soil hydraulic mapping, Sci. Total Environ., 921, 171258, https://doi.org/10.1016/j.scitotenv.2024.171258, 2024.(Szabó et al., 2024)

Szalai, S., Szinell, C., and Zoboki, J.: Drought monitoring in Hungary, in: Early Warning Systems for Drought Preparedness and Drought Management, World Meteorological Organization, Geneva, 182–199, 2004.

Tran, H.Q., Fehér, Z.Zs., Túri, N., and Rakonczai, J.: Climate change as an environmental threat on the central plains of the Carpathian Basin based on regional water balances, Geographica Pannonica, 26(3), 184–199, https://doi.org/10.5937/gp26-37271, 2022.(Quang et al., 2022)

Tóth, B., Weynants, M., Nemes, A., Makó, A., Bilas, G., and Tóth, G.: New generation of hydraulic pedotransfer functions for Europe, Eur. J. Soil Sci., 68, 613–624, https://doi.org/10.1111/ejss.12456, 2017.(Tóth et al., 2014)

Tóth, T. and Várallyay, Gy.: Salt accumulation in Hungarian soils, in: Proceedings of the International Symposium on Solonetz Soils, Osijek, 203–210, 1980.

Várallyay, Gy., Szűcs, L., Murányi, A., Rajkai, K., and Zilahy, P.: Magyarország termőhelyi adottságait meghatározó talajtani tényezők 1:100 000 méretarányú térképe [Map of soil factors determining the site conditions of Hungary at 1:100 000 scale], Agrokémia és Talajtan, 28, 363–384, 1979.

Várallyay, Gy.: A dunavölgyi talajok sófelhalmozódási folyamatai [Salinization processes of the Duna Valley soils], Agrokémia és Talajtan, 16, 327–349, 1967.

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., et al.: The FAIR Guiding Principles for scientific data management and stewardship, Sci. Data, 3, 160018, https://doi.org/10.1038/sdata.2016.18, 2016.

Wunsch, A., Liesch, T., and Broda, S.: Groundwater level forecasting with artificial neural networks: a comparison of long short-term memory (LSTM), convolutional neural networks (CNNs), and non-linear autoregressive networks with exogenous input (NARX), Hydrol. Earth Syst. Sci., 25, 1671–1687, https://doi.org/10.5194/hess-25-1671-2021, 2021.

References

Bakacsi, Z., Tóth, T., Makó, A., Barna, G., Laborczi, A., Szabó, J., Szatmári, G., & Pásztor, L. (2019). National level assessment of soil salinization and structural degradation risks under irrigation. *Hungarian Geographical Bulletin*, *68*(2), 141–156. https://doi.org/10.15201/hungeobull.68.2.3

Ballabio, C., Panagos, P., & Monatanarella, L. (2015). Mapping topsoil physical properties at European scale using the LUCAS database. *Geoderma*, *261*, 110–123. https://doi.org/10.1016/j.geoderma.2015.07.006

Banta, E. R. (2000). *MODFLOW-2000, the U.S. Geological Survey Modular Ground-Water Model; documentation of packages for simulating evapotranspiration with a segmented function (ETS1) and drains with return flow (DRT1)*. Us Geological Survey. https://doi.org/10.3133/ofr00466

Brunner, P., Li, H. T., Kinzelbach, W., Li, W. P., & Dong, X. G. (2008). Extracting phreatic evaporation from remotely sensed maps of evapotranspiration. *Water Resources Research*, *44*(8). https://doi.org/10.1029/2007wr006063

Fehér, Z. Z., & Rakonczai, J. (2019). Analysing the sensitivity of Hungarian landscapes based on climate change induced shallow groundwater fluctuation. *Hungarian Geographical Bulletin*, *68*(4), 355–372. https://doi.org/10.15201/hungeobull.68.4.3

Gal, Y., & Ghahramani, Z. (2015). *Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning*. https://doi.org/10.48550/arxiv.1506.02142

Garamhegyi, T., Kovács, J., Pongrácz, R., Tanos, P., & Hatvani, I. G. (2017). Investigation of the climate-driven periodicity of shallow groundwater level fluctuations in a Central-Eastern European agricultural region. *Hydrogeology Journal*, *26*(3), 677–688. https://doi.org/10.1007/s10040-017-1665-2

Hoyer, S., & Hamman, J. (2017). xarray: N-D labeled Arrays and Datasets in Python. *Journal of Open Research Software*, *5*(1), 10. https://doi.org/10.5334/jors.148

Huber, P. J. (1973). Robust Regression: Asymptotics, Conjectures and Monte Carlo. *The Annals of Statistics*, *1*(5). https://doi.org/10.1214/aos/1176342503

Ilyés, C., Tóth, M., & Szűcs, P. (2023). Calculating the connection between precipitation and shallow groundwater. *Multidiszciplináris Tudományok*, *13*(3), 292–299. https://doi.org/10.35925/j.multi.2023.3.29

Kamboj, M., Goyal, C., & Ratra, N. (2024). Unleashing Potential: Snowflake's Streamlit Strategy for GenAI Solutions Leveraging External Network Access and OpenAI. *INTERANTIONAL JOURNAL OF SCIENTIFIC RESEARCH IN ENGINEERING AND MANAGEMENT*, *08*(03), 1–5. https://doi.org/10.55041/ijsrem29339

Kern, A., Dobor, L., Hollós, R., Marjanović, H., Torma, C. Z., Kis, A., Fodor, N., & Barcza, Z. (2023). Seamlessly combined historical and projected daily meteorological datasets for impact studies in Central Europe: The FORESEE v4.0 and the FORESEE-HUN v1.0. *Climate Services*, *33*, 100443. https://doi.org/10.1016/j.cliser.2023.100443

Mádl-Szőnyi, J., & Tóth, J. (2009). A hydrogeological type section for the Duna-Tisza Interfluve, Hungary. *Hydrogeology Journal*, *17*(4), 961–980. https://doi.org/10.1007/s10040-008-0421-z

Moriasi, D., Arnold, J., Liew, M., Bingner, R., Harmel, R., & Veith, T. (2007). Model Evaluation Guidelines for Systematic Quantification of Accuracy in Watershed Simulations. *Transactions of the ASABE*, *50*(3), 885–900. https://doi.org/10.13031/2013.23153

Müller, S., Schüler, L., Zech, A., & Heße, F. (2021). *GSTools v1.3: A toolbox for geostatistical modelling in Python*. Copernicus Gmbh. https://doi.org/10.5194/gmd-2021-301

Nearing, G. S., Kratzert, F., Sampson, A. K., Pelissier, C. S., Klotz, D., Frame, J. M., Prieto, C., & Gupta, H. V. (2021). What Role Does Hydrological Science Play in the Age of Machine Learning? *Water Resources Research*, *57*(3). https://doi.org/10.1029/2020wr028091

Paszke, A., Tejani, A., Lerer, A., Liu, F., Devito, Z., Bai, J., Antiga, L., Killeen, T., Massa, F., Raison, M., Köpf, A., Steiner, B., Yang, E. S., Gross, S., Lin, Z., Bradbury, J., Desmaison, A., Chanan, G., Chilamkurthy, S., … Chintala, S. (2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. Cornell University. https://doi.org/10.48550/arxiv.1912.01703

Quang, T., Fehér, Z., Túri, N., & Rakonczai, J. (2022). Climate change as an environmental threat on the central plains of the Carpathian basin based on regional water balances. *Geographica Pannonica*, *26*(3), 184–199. https://doi.org/10.5937/gp26-37271

Reichstein, M., Baghirov, Z., Jung, M., & Kraft, B. (2025). *Deep learning and Process Understanding for Data-Driven Earth System Science*. https://doi.org/10.5194/egusphere-egu24-15874

Shah, N., Nachabe, M., & Ross, M. (2007). Extinction Depth and Evapotranspiration from Ground Water under Selected Land Covers. *Groundwater*, *45*(3), 329–338. https://doi.org/10.1111/j.1745-6584.2007.00302.x

Shi, X., Yeung, D., Chen, Z., Wang, H., Wong, W. C., & Woo, W.-C. (2015). *Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting*. Cornell University. https://doi.org/10.48550/arxiv.1506.04214

Soylu, M. E., Istanbulluoglu, E., Lenters, J. D., & Wang, T. (2010). *Quantifying the impact of groundwater depth on evapotranspiration in a semi-arid grassland region*. Copernicus Gmbh. https://doi.org/10.5194/hessd-7-6887-2010

Sun, A. Y., Scanlon, B. R., Zhang, Z., Walling, D., Bhanja, S. N., Mukherjee, A., & Zhong, Z. (2019). Combining Physically Based Modeling and Deep Learning for Fusing GRACE Satellite Data: Can We Learn From Mismatch? *Water Resources Research*, *55*(2), 1179–1195. https://doi.org/10.1029/2018wr023333

Szabó, B., Mészáros, J., Laborczi, A., Takács, K., Szatmári, G., Bakacsi, Z., Makó, A., & Pásztor, L. (2024). From EU-SoilHydroGrids to HU-SoilHydroGrids: A leap forward in soil hydraulic mapping. *Science of The Total Environment*, *921*, 171258. https://doi.org/10.1016/j.scitotenv.2024.171258

Tóth, B., Weynants, M., Nemes, A., Makó, A., Bilas, G., & Tóth, G. (2014). New generation of hydraulic pedotransfer functions for Europe. *European Journal of Soil Science*, *66*(1), 226–238. https://doi.org/10.1111/ejss.12192