Received: 30 March 2022 | Revised: 12 August 2022 | Accepted: 18 August 2022

DOI: 10.1002/env.2765

RESEARCH ARTICLE

WILEY

Functional forecasting of dissolved oxygen in high-frequency vertical lake profiles

Luke Durell[1] | **J. Thad Scott**[2] | **Douglas Nychka**[3] | **Amanda S. Hering**[1]

[1]Department of Statistical Science, Baylor University, Waco, Texas, USA

[2]Department of Biology, Baylor University, Waco, Texas, USA

[3]Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, Colorado, USA

Correspondence
Amanda S. Hering, Department of Statistical Science, Baylor University, One Bear Place 97140, Waco, TX 76798, USA.
Email: mandy_hering@baylor.edu

Funding information
National Science Foundation, Grant/Award Number: 1924146; Tarrant Regional Water District

Abstract

Predicting dissolved oxygen (DO) in lakes is important for assessing environmental conditions as well as reducing water treatment costs. High levels of DO often precede toxic algal blooms, and low DO causes carcinogenic metals to precipitate during water treatment. Typically, DO is predicted from limited data sets using hydrodynamic modeling or data-driven approaches like neural networks. However, functional data analysis (FDA) is also an appropriate modeling paradigm for measurements of DO taken vertically through the water column. In this analysis, we build FDA models for a set of profiles measured every 2 hours and forecast the entire DO percent saturation profile from 2 to 24 hours ahead. Functional smoothing and functional principal component analysis are applied first, followed by a vector autoregressive model to forecast the empirical functional principal component (FPC) scores. Rolling training windows adapt to seasonality, and multiple combinations of window sizes, model variables, and parameter specifications are compared using both functional and direct root mean squared error metrics. The FPC method outperforms a suite of comparison models, and including functional pH, temperature, and conductivity variables improves the longer forecasts. Finally, the FDA approach is useful for identifying unusual observations.

KEYWORDS

dissolved oxygen, forecasting, functional time series, vector autoregressive model, vertical profiles

1 | INTRODUCTION

Dissolved oxygen (DO) is the amount of free oxygen molecules in water, and predicting DO in lakes is necessary for both environmental management and effective water treatment. High concentrations of DO can be related to harmful algal blooms that can lead to widespread aquatic biota death (Seki et al., 1980). Conversely, aquatic ecosystems struggle to sustain life under hypoxic (scarce DO) or anoxic (zero DO) conditions (e.g., King et al., 2012; Šiljić Tomić et al., 2018). Hypoxic/anoxic conditions also lead to mobilization of reduced metals into the water column (Banks et al., 2012; Eggleton & Thomas, 2004; Riedel et al., 1997), particularly manganese and iron (Atkinson et al., 2007), resulting in higher water treatment costs and potential human health risks. If accurate forecasts of DO were available, then early warnings of either high or low DO could be issued to reservoir managers.

Forecasting DO can be complicated due to seasonality and variation across depths. We refer to a full measurement of DO from the surface to the bottom of the water body as a profile. Some water bodies, especially deep ones, are stratified by temperature year-round, where a thermocline separates the warmer surface layer from the cooler bottom layer, preventing nutrients and dissolved substances from traveling across the thermocline (Boehrer & Schultze, 2008). Other water bodies mix during the year. Thus, except for very shallow water bodies, or when the water body is mixing, surface measurements are insufficient for forecasting an entire profile.

The traditional approach to predicting water quality variables is hydrodynamic modeling, a simulation technique that involves specifying environmental conditions such as wind speed, wind direction, water inflows and outflows, approximate surface heat, and bathymetry (Hodges et al., 2000; Istvánovics & Honti, 2018; Marti et al., 2011). This approach requires complicated approximations and calibration/tuning steps for a given water body that are often not easily translated to other water bodies. The sheer number of environmental approximations required by these models can limit their accuracy (Jin et al., 2019). By contrast, forecasting DO using data-driven approaches is increasingly common. For the purpose of our application, we focus only on single-lake/reservoir studies. A summary of these single-lake, data-driven approaches to predicting DO are outlined in Table 1. The papers assigned "P" for prediction in the "Model aim" column, as opposed to "F" for forecast, are those that emphasize inter-variable relationships and estimation of in-sample values rather than modeling with the purpose of forecasting future values of DO.

In reviewing the lake DO modeling literature, both the data structure and the collection methodology vary widely. Only half of the studies in Table 1 have non-monthly measurements. All of the papers use some form of neural network (NN) model to forecast DO. Chen and Liu (2015) and Karakaya et al. (2013) include multiple regression models in addition to NNs in their studies. However, only our work and that of Saber et al. (2020) seek to reconstruct full-column water profiles. While Saber et al. (2020) offer an excellent applied example of DO forecasting, our data have the most temporally dense (measured every 2 hours) and vertically dense (measured every 0.5 m from the surface to 10 m below) full-column water profile measurements. As high-frequency vertical profile measurement instruments become cheaper, this type of data will become more common; hence, there is a pressing need for accurate examples of full-profile forecasting techniques. We seek to use functional data analysis (FDA) to simultaneously forecast the entire water quality profile with one model, which allows the DO across depths to contribute to the forecasts. While the FDA framework has been used for water quality outlier detection (Sancho et al., 2016) and identifying water body differences (Henderson, 2006), to the best of our knowledge, this is the first example of forecasting water quality profiles with FDA.

FDA is an established framework for modeling environmental data. For example, Fortuna et al. (2020) use basis function expansions to estimate ecological diversity profiles; Qu et al. (2021) apply functional multivariate analysis of variance to temperature curves and bivariate wind speed data; and Harris et al. (2022) detect change points in water vapor profiles using functional principal component analysis (FPCA). The particular family of FDA that we use in this work forecasts realizations from a functional time series (FTS) by estimating the functional principal components (FPCs) and corresponding FPC scores from the data. The scores are then forecast using a nonfunctional method, and the forecast functional observation is obtained as the product of the forecast scores and the estimated FPCs (Aue et al., 2015). Variations of this approach include the use of different score forecasting techniques (Beyaztas & Shang, 2019; Hyndman & Ullah, 2007), functional quantile forecasting (Cabrera & Schulz, 2017), aggregation of individually forecast populations

TABLE 1 Summary of literature on single-lake, data-driven prediction of DO

Authors	Sites measured	Deepest measurement	Depth freq.	Duration	Temporal freq.	Model aim	Model type
Chen and Liu (2014)	Multiple	Not reported	NA	19 mths	Not reported	P	NN
Huo et al. (2013)	Not reported	Not reported	NA	6 yrs	Monthly	P	NN
Karakaya et al. (2013)	1	1.5 m	NA	3 mths	15 min	P	NN
Ranković et al. (2010)	3	30 m	3–5 m	3 yrs	Monthly	P	NN
Ranković et al. (2011)	3	30 m	3–5 m	3 yrs	Monthly	P	NN
Saber et al. (2020)	1	116 m	5 m	5.5 yrs	6 hours	F	NN
This article	1	10 m	0.5 m	7 mths	2 hours	F	FDA

Note: Meters are designated by m. A model aim of P indicates prediction whereas an F indicates forecasting. Model type is either a neural network (NN) or a purely statistical functional data analysis (FDA) approach.

(Shang, 2020; Shang & Hyndman, 2017), sparse FPCA (Yao et al., 2005), and multivariate forecasting using vector error correction models on the scores (Gao & Shang, 2017).

The specific modeling approach that we implement is that of Aue et al. (2015), which forecasts the empirical FPC scores using a vector autoregressive (VAR) model. It has a history of adoption, modification, or comparison in numerous papers (Cabrera & Schulz, 2017; Gao & Shang, 2017; Klepsch et al., 2017; Shang, 2017, 2019). There is evidence that it is robust to model misspecification (Aue et al., 2015), which is important because the data vary widely across season, and it is difficult to capture all of the heterogeneity of the functional process. Additionally, incorporating functional exogenous variables with the Aue et al. (2015) method is straightforward. Lastly, it has readily available computational implementation in R.

The goal of this article is to present the first application of functional forecasting for vertical water quality profiles. We will estimate and forecast FPC scores with a VAR model and compare this approach to nonfunctional statistical and machine learning methods such as autoregressive (AR) and VAR models and NNs. Additionally, we will compare the forecast performance with other functional approaches, such as using FPC scores with AR models, estimating dynamic FPCs according to Hörmann et al. (2015), and utilizing functional autoregressive model of order 1 (FAR) models (Bosq, 2000). In doing so, to the best of our knowledge, we provide the first simultaneous comparison of the FPC VAR method of Aue et al. (2015) to purely statistical, purely machine learning, and functional-statistical methods. Further, this work provides a case study and forecasting benchmark for the unique and growing area of forecasting vertical lake profiles. The remainder of this article is organized as follows: In Section 2, we provide a detailed description of the unique data used in this study. Section 3 details the FPCA approach and primary forecasting strategy. Comparison models and all results are presented in Section 4 along with a discussion of outlying or poorly forecast profiles, and Section 5 offers some concluding remarks.

2 | DATA DESCRIPTION

The data were collected from Eagle Mountain Lake (EML), a reservoir in North Texas that is used for recreation and as a local water supply. EML mixes multiple times per year, resulting in similar temperatures from top to bottom during mixing. At the EML Realtime Monitoring Platform, a measurement instrument was placed at the primary controlled outflow of the reservoir. This location is both one of the deepest parts of the lake and where drinking water is drawn for treatment. The data were collected every 2 hours at 21 depths, measured every 0.5 m starting from the surface (0.0 m) and ending 10.0 m subsurface. The data set considered in this work contains 2252 vertical profiles of four variables between April 25, 2019 and October 29, 2019. To simplify computation, a small percentage of the data was imputed (less than 4%), and that process is described in the Supplementary Materials.

The variables measured are as follows: (1) Temperature (T) is measured in degrees Celsius and is expected to be warmer at the surface and cooler at the bottom, except in cases when EML is mixing. (2) Conductivity (C) measures how easily electricity moves through the water and is reported in micro Siemens per centimeter (μS/c). (3) pH (pH) indicates the alkalinity/acidity of the water from 0 (acidic) to 14 (alkaline). (4) DO %sat (D) is the percentage of DO relative to the concentration in equilibrium with the atmosphere. D can vary from 0% to 200%, where D of 100% indicates that the percentage of DO in the atmosphere is the same as the percentage of DO in the water. Values of D above 100% are caused by wind turbulence above the lake and aquatic plant photosynthesis. Values below 100% are caused by aquatic animal respiration, aquatic plants without access to light, and decomposition at the bottom of the lake. At the lake surface, D may vary substantially, but in the absence of vertical mixing, it will generally be lower and more stable at the bottom.

Figure 1 displays smoothed vertical profile measurements of these four variables across Early Summer (April, May, and June), Late Summer (July and August), and Early Fall (September and October), revealing that the profile patterns of each variable vary notably by season. The purple/dark curves are the first observations of the season, and the gold/light curves are the last observations of the season. The y-axis represents lake depth. Following convention (e.g., Saber et al., 2020), the range is reversed so that the lake surface (0) is at the top of the vertical axis. D, C, and pH display the most variability in Late Summer, and T displays the most variability in Early Fall. The apparent seasonal changes in these profiles indicate that the forecasting method should account for seasonal nonstationarity.

To investigate temporal dependence and diurnal effects in the profiles, the curves are first centered by their respective hourly mean curve (computed based on the entire data set), and then univariate autocorrelation functions (ACFs) at each depth for each variable are computed for the centered profiles (Figure 2, top). The ACFs reveal that for D and pH, the discrete measurements at the surface of the lake still contain some periodicity, even after centering by hourly mean curves.

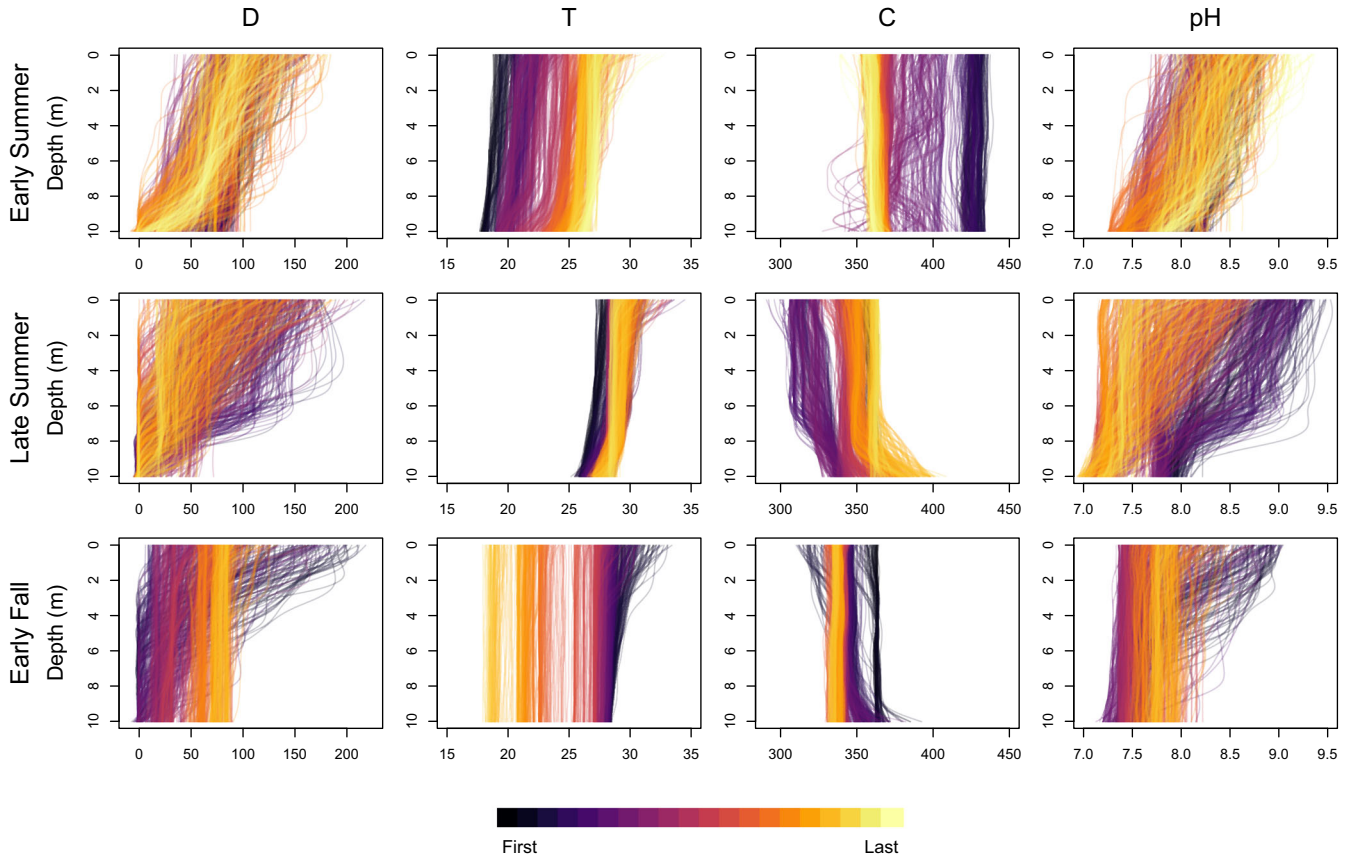

FIGURE 1 Smoothed vertical profile measurements of variables (columns) by season (rows). The colors progress from dark (first observations in the season) to light (last observations in the season).

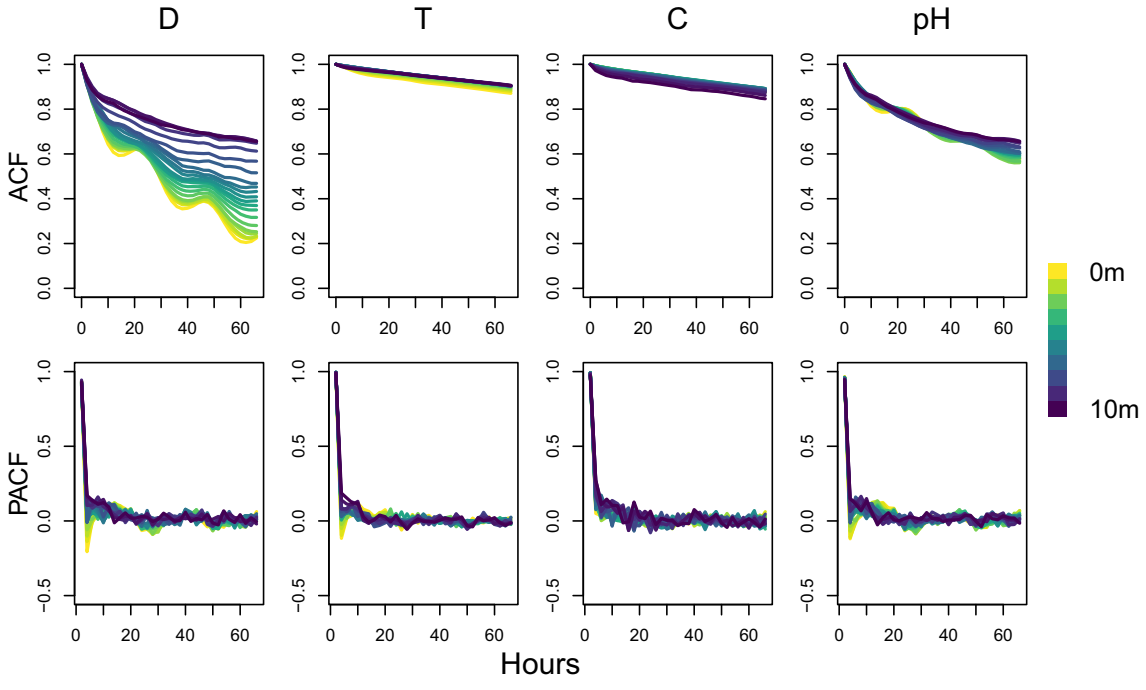

FIGURE 2 Univariate ACF (top row) and PACF (bottom row) plots of each variable (columns) by depth

This is not entirely unexpected because the detrending is performed across the entire data set and not by season. For D, the values measured at the bottom of the lake have a stronger and more persistent autocorrelation than values measured at the surface. T and C have relatively strong ACFs compared to D and pH because at 60 hours, the ACF of neither T nor C drops below 0.85. The bottom row of Figure 2 shows that for the centered data, all four variables have similar partial autocorrelation functions (PACFs) across depths. The PACF between two time points approaches zero rapidly, indicating that an autoregressive type of model would be an appropriate choice in a univariate setting.

3 | METHODOLOGY

In the following, we consider an FTS of D profiles $(W_n : n \in \mathbb{Z})$ with realizations denoted $W_n(s), n = 1, \ldots, N$, that can be evaluated for any depth, s, in the functional domain, $\mathcal{S} = [0, 10]$. In practice, we observe the discretized functional realizations of D profiles as $W_n(s_i)$ for $i = 1, \ldots, 21$. We are ultimately interested in forecasting the h-step ahead functional realization, $W_{N+h}(s)$. Our approach is to smooth the discretized curves, center the curves by their smooth hourly mean, apply FPCA to obtain a vector of FPC scores, forecast the scores h-steps ahead, reconstruct h-step ahead centered forecasts using the forecast scores, and add back the hourly mean curve.

In Section 3.1, the preliminary steps of smoothing and centering the curves are described. In Section 3.2, we describe FPCA in the population sense and show how estimated FPCs and scores can be used for curve approximation. In Section 3.3, we describe the method that we use for forecasting one-step-ahead curves without exogenous variables, which is then expanded to include exogenous variables. Section 3.4 extends the forecasting to twelve steps ahead. Finally, in Section 3.5, we summarize the steps needed to forecast h-step-ahead D profiles, including the use of a rolling window to address the nonstationarity present in the data.

3.1 | Preliminary steps

Penalized smoothing with a set of cubic B-spline basis functions approximates $W_n(s)$ with

$$W_n^K(s) = \sum_{k=1}^{K} c_{n,k} B_k(s). \tag{1}$$

Here, $B_k(s)$ is the kth B-spline basis function, and $c_{n,k}$ is the kth coefficient that weighs the influence on $W_n^K(s)$ of the kth basis function for the saturated choice of basis functions which, in our case, is $K = 23$. The smoothness of $W_n^K(s)$ is ensured by imposing a penalty on the second derivative, which is controlled by the parameter λ. This smoothing parameter is chosen by generalized cross validation within a window of observations, as described in Section 3.5.

Because of the diurnal patterns in the data, particularly for D as shown in Figure 2, we center the observations by their corresponding estimated functional hourly means, $\hat{\mu}_t(s)$ for $t = 0, 2, \ldots, 22$, prior to applying FPCA. The estimated means are constructed as $\hat{\mu}_t(s) = \frac{1}{N_t} \sum_{n \in \{n_t\}} W_n^K(s)$ where $\{n_t\}$ is the subset of indices from $\{1, \ldots, N\}$ corresponding to hour t, and N_t is the total number of elements in $\{n_t\}$. The observations of D centered by their hourly mean are denoted

$$Y_n^K(s) = W_n^K(s) - \hat{\mu}_t(s). \tag{2}$$

The goal then becomes to obtain the h-step ahead, centered forecast of D, $\hat{Y}_{N+h}^K(s)$, and then construct the final forecast as $\hat{W}_{N+h}^K(s) = \hat{Y}_{N+h}^K(s) + \hat{\mu}_t(s)$ by adding back the corresponding hourly mean.

The functional observations of the exogenous variables T, C, and pH are similarly smoothed and centered by their respective estimated hourly means. The nth hourly centered, smooth curve for the ℓth exogenous variable obtained with penalized smoothing with penalty λ_ℓ is denoted $X_{n,\ell}^K(s)$. More details concerning smoothing can be found in the Supplementary Materials.

3.2 | Functional principal component analysis

In the population sense, the covariance function, $c(s, s')$, of a mean-zero random function, $Y(s)$ with $s \in \mathcal{S}$, is related to the jth eigenfunction, or FPC, $v_j(s)$, and eigenvalue, β_j, with

$$\int_S c(s,s')v_j(s')ds' = \beta_j v_j(s), \tag{3}$$

where $s, s' \in \mathcal{S}$, and $c(s,s') = E[Y(s)Y(s')]$. The random function can be expressed with the Karhunen–Loève (KL) expansion:

$$Y(s) = \sum_{j=1}^{\infty} y_j v_j(s), \tag{4}$$

where $y_j = \langle Y, v_j \rangle = \int_S Y(s)v_j(s)ds$ is the FPC score.

Moving from the population sense to the sample case, we have an unobserved, mean-zero, stationary FTS, $Y_n(s), n = 1, \ldots, N$, assumed to be accurately approximated by $Y_n^K(s), n = 1, \ldots, N$. The sample covariance function is $\hat{c}(s,s') = \frac{1}{N}\sum_{n=1}^{N} Y_n^K(s)Y_n^K(s')$ from which q empirical FPCs (EFPCs), $\hat{v}_j(s), j = 1, \ldots, q$, and empirical eigenvalues, $\hat{\beta}_j, j = 1, \ldots, q$, are estimated. Under general dependence assumptions given in Hörmann and Kokoszka (2010), Hörmann and Kokoszka (2012) prove that the product of a constant and a given EFPC is a root-n estimator of the corresponding population FPC. Empirical scores corresponding to observation $Y_n^K(s)$ can be constructed as $y_{n,j}^e = \langle Y_n^K, \hat{v}_j \rangle, j = 1, \ldots, q$, where the e in the superscript indicates that the score is empirical. Estimation of the FPCs and scores is performed with the `pca.fd` function in the `fda` package (Ramsay et al., 2021) via a standard technique in the literature. Additional details can be found in Hörmann and Kokoszka (2012) and in Chapter 8 of Ramsay and Silverman (2005).

Based on the EFPCs and scores, a truncated KL expansion can be used to approximate $Y_n^K(s)$ as,

$$Y_n^K(s) \approx Y_n^d(s) = \sum_{j=1}^{d} y_{n,j}^e \hat{v}_j(s). \tag{5}$$

The truncation level is denoted by d where $d < q$ and is the number of EFPC curves used to approximate $Y_n^K(s)$. The analogous centered exogenous variable with truncated KL curve is denoted $X_{n,\ell}^{d_\ell}(s)$ where d_ℓ represents the ℓth exogenous variable's truncation level.

The number of EFPCs, d and d_ℓ, are often selected with data-driven approaches, such as the cumulative variance threshold or the functional final prediction error method of Aue et al. (2015). Based on our visual comparison of the reconstructed curves, $Y_n^d(s)$, and the smooth data, $Y_n^K(s)$, these automatic approaches do not perform well. Instead, we select fixed choices of d and d_ℓ, chosen as the maximum number of EFPCs such that at least two observations per parameter are available in the most parameter-dense model. A small pilot study comparing the data-driven and fixed choices of d and d_ℓ on forecasting results is outlined in the Supplementary Materials.

3.3 | Empirical score vector forecasting

We begin by considering one-step-ahead forecasting of D. FPCA is applied to the smoothed and centered observations, $Y_n^K(s), n = 1, \ldots, N$ to obtain N vectors of empirical scores, $\mathbf{y}_n^e = [y_{n,1}^e, \ldots, y_{n,d}^e]'$ and a single functional vector of EFPCs, $\hat{\mathbf{v}} = [\hat{v}_1(s), \ldots, \hat{v}_d(s)]'$. We seek to forecast the entire empirical score vector $\mathbf{y}_{N+1}^e$ using a VAR model of order p (VAR(p)). The VAR(p) model for the empirical score vectors can be written as

$$\mathbf{y}_n^e = \mathbf{c} + \mathbf{A}_1\mathbf{y}_{n-1}^e + \cdots + \mathbf{A}_p\mathbf{y}_{n-p}^e + \mathbf{w}_n, \quad n = p+1, \ldots, N, \tag{6}$$

where $\mathbf{c}$ is a $d \times 1$ vector of constants; $\mathbf{A}_1, \ldots, \mathbf{A}_p$ are $d \times d$ random matrices of coefficients; and $\mathbf{w}_n$ is a $d \times 1$ error vector of mean zero uncorrelated white noise. For a general treatment of VAR models, see Lütkepohl (2005).

The one-step-ahead score forecast can now be constructed as

$$\hat{\mathbf{y}}_{N+1}^e = \hat{\mathbf{c}} + \hat{\mathbf{A}}_1\mathbf{y}_N^e + \cdots + \hat{\mathbf{A}}_p\mathbf{y}_{N-p+1}^e, \tag{7}$$

where $\hat{\mathbf{c}}$ and $\hat{\mathbf{A}}_1, \ldots, \hat{\mathbf{A}}_p$ are ordinary least squares estimates of the corresponding parameters. Thus, the total number of parameters to estimate in a VAR(p) model is $d + pd^2$. The forecast score vector can be plugged directly into Equation (5)

to obtain a centered one-step-ahead forecast curve,

$$\hat{Y}^d_{N+1}(s) = (\hat{\mathbf{y}}^e_{N+1})'\hat{\boldsymbol{\nu}}. \tag{8}$$

The final forecast curve can then be constructed as

$$\hat{W}^d_{N+1}(s) = \hat{Y}^d_{N+1}(s) + \hat{\mu}_t(s). \tag{9}$$

The VAR(p) model allows for easy extension into a VARX(p) model that incorporates information through the score vectors obtained by applying FPCA to the exogenous functional variables. When incorporating L exogenous variables, we have L exogenous EFPC vectors, $\hat{\boldsymbol{\nu}}_1, \ldots, \hat{\boldsymbol{\nu}}_L$, of dimension $d_1, \ldots, d_L$, respectively. For the ℓth exogenous variable, we also have N empirical score vectors, $\mathbf{x}^e_{1,\ell}, \ldots, \mathbf{x}^e_{N,\ell}$, each of which are of dimension d_ℓ. The model in Equation (6) is updated to,

$$\mathbf{y}^e_n = \mathbf{c} + \mathbf{A}_1\mathbf{y}^e_{n-1} + \cdots + \mathbf{A}_p\mathbf{y}^e_{n-p} + \mathbf{B}\mathbf{X}^e_{n-1} + \mathbf{w}_n, \quad n = p+1, \ldots, N, \tag{10}$$

where the $d_\ell \times 1$ dimensional empirical score vectors $\mathbf{x}^e_{n-1,1}, \ldots, \mathbf{x}^e_{n-1,L}$ are stacked into the $\sum_{\ell=1}^L d_\ell \times 1$ dimensional vector, $\mathbf{X}^e_{n-1}$. Then, the one-step-ahead score forecast that includes exogenous variables is

$$\hat{\mathbf{y}}^e_{N+1} = \hat{\mathbf{c}} + \hat{\mathbf{A}}_1\mathbf{y}^e_N + \cdots + \hat{\mathbf{A}}_p\mathbf{y}^e_{N-p+1} + \hat{\mathbf{B}}\mathbf{X}^e_N. \tag{11}$$

The resulting forecast score vector can be directly incorporated into Equation (8) to obtain a one-step-ahead centered forecast. The only new parameter to estimate is the coefficient matrix, $\mathbf{B}$, which increases the number of parameters to estimate to $d + pd^2 + d \times (\sum_{\ell=1}^L d_\ell)$.

3.4 | Day ahead forecasting

In light of monitoring for ecosystem changes and water treatment, a longer forecast horizon than 2 hours ($h = 1$) is useful. We introduce both a direct and an iterative day-ahead ($h = 12$) forecasting model for D. The direct day-ahead VARX model is specified as,

$$\mathbf{y}^e_n = \mathbf{c} + \mathbf{A}_1\mathbf{y}^e_{n-12} + \cdots + \mathbf{A}_p\mathbf{y}^e_{n-p-11} + \mathbf{B}\mathbf{X}^e_{n-12} + \mathbf{w}_n, \quad n = p+12, \ldots, N, \tag{12}$$

with the only difference between Equations (10) and (12) being the lag in the subscripts. The day-ahead score forecast is then constructed as

$$\hat{\mathbf{y}}^e_{N+12} = \hat{\mathbf{c}} + \hat{\mathbf{A}}_1\mathbf{y}^e_N + \cdots + \hat{\mathbf{A}}_p\mathbf{y}^e_{N-p+1} + \hat{\mathbf{B}}\mathbf{X}^e_N, \tag{13}$$

which can be used in Equation (8) to obtain the direct, day-ahead, centered forecast curve, denoted $\hat{Y}^d_{N+12,dir}(s)$.

An alternative approach is an iterative day-ahead forecast that fits a VAR model with Equation (6) and computes one-step-ahead score forecasts, which are iteratively plugged back into Equation (7) until the day-ahead forecast score is obtained. The day-ahead forecast score is then used in Equation (8) to produce the iterative, day-ahead, centered forecast curve, denoted $\hat{Y}^d_{N+12,itr}(s)$. The iterative day-ahead forecast does not incorporate exogenous variables because this would require separate forecasts of each exogenous variable.

3.5 | Forecasting scheme summary

The methodology outlined in the preceding sections will be referred to as functional principal components with vector autoregressive model and exogenous variables (FPC VARX). We perform a rolling window forecasting approach where we first fix a training window size of 28, 35, or 42 days, corresponding to 4, 5, or 6 weeks. This is a vital component of our forecasting approach because both the KL expansion and the VAR(p) models assume stationarity,

which is violated by the observed seasonality in the data (see Figure 1). Because multiple years of data are not available, the seasonal components cannot be estimated, so the use of rolling windows allows us to reasonably address this nonstationarity. We fix the first forecast date to be the same for each window size and forecast horizon (h) to facilitate comparisons. After a curve is forecast, the window rolls one step ahead, incorporating the next smooth function and dropping the oldest smooth function in the window, thereby keeping N observations in the rolling window.

Steps 1–3 below are performed for all eight combinations of exogenous variables. For window size 42, which provides enough observations to fit more complex models, the steps are repeated for $p = 1$ and $p = 2$; otherwise, only $p = 1$ is used in the VAR model. When the only input variable included is D, all three model schemes (iterative day-ahead, direct day-ahead, and one-step-ahead) are used, but for all models that include exogenous variables, only the direct day-ahead forecast and the one-step-ahead forecast are computed.

In summary, the FPC VARX is implemented as follows:

1. Fix the rolling window size, N, to be 28, 35, or 42 days.
2. Let m indicate the window preceding the profile being forecast for $m = 1, \ldots, M$, for a total of M windows. For each m,
   (a) Smooth and center the N curves of D and any needed exogenous variables in window m according to Section 3.1.
   (b) Perform FPCA on D and any exogenous variables, as described in Section 3.2.
   (c) For a given value of p, fit a VAR or a VARX model to the score vectors in the training window, as described in Section 3.3.
   (d) Forecast $\hat{\mathbf{y}}^e_{N+h}$ using Equations (7), (11), or (13). Construct the centered h-step-ahead forecast as $\hat{Y}^d_{N+h}(s) = (\hat{\mathbf{y}}^e_{N+h})'\hat{\boldsymbol{\nu}}$, and obtain the final forecast $\hat{W}^d_{N+h}(s) = \hat{Y}^d_{N+h}(s) + \hat{\mu}_t(s)$.
   (e) Calculate the forecast's functional RMSE metric, computed as

$$\text{RMSE}_{f,m} = \sqrt{\langle W^K_{N+h} - \hat{W}^d_{N+h}, W^K_{N+h} - \hat{W}^d_{N+h}\rangle}, \tag{14}$$

   and what we term the direct RMSE metric,

$$\text{RMSE}_{d,m} = \sqrt{\frac{1}{21}\sum_{i=1}^{21}(W_{N+h}(s_i) - \hat{W}^d_{N+h}(s_i))^2}. \tag{15}$$

   $\text{RMSE}_{f,m}$ is a functional measure of the distance between $\hat{W}^d_{N+h}(s)$ and $W^K_{N+h}(s)$. $\text{RMSE}_{d,m}$ is a measure of the direct squared distance between the forecast curve evaluated at the twenty-one measured locations and the true observed values.
   (f) Roll the window forward by one observation, and repeat the steps (a) through (e) until the last window, M, is reached.
3. For the M forecast curves produced by each window, compute $\text{RMSE}_f = \frac{1}{M}\sum_{m=1}^{M}\text{RMSE}_{f,m}$ and $\text{RMSE}_d = \frac{1}{M}\sum_{m=1}^{M}\text{RMSE}_{d,m}$.

4 | RESULTS

After obtaining forecasts for every combination of variables, window sizes, lags, and forecast horizon, we discuss the forecast performance of FPC VARX in Section 4.1. In Section 4.2, we compare the best FPC VARX specification against a suite of other methodologies. Section 4.3 explores functional visualization and outlier detection, focusing on the performance of FPC VARX.

4.1 | Initial forecasting results

The one-step-ahead and the day-ahead results with RMSE_f are shown in Tables 2 and 3, respectively. Results for RMSE_d are very similar and are reported in the Supplementary Materials. Each table displays the seasonal RMSEs,

TABLE 2 $RMSE_f$ for FPC VARX one-step-ahead forecasting models of various p, window sizes, and exogenous variables

p	Win. size	Variable(s)	Early Summer	Late Summer	Early Fall	Overall
1	28	***D***	**10.57**	11.50	7.79	**9.85**
		DC	10.62	11.49	7.88	9.89
		DCpH	10.66	11.52	7.95	9.94
		DpH	10.62	**11.46**	7.91	9.89
		DT	10.66	11.59	**7.78**	9.90
		DTC	10.72	11.56	7.79	9.90
		DTCpH	10.77	11.47	7.87	9.90
		DTpH	10.74	11.54	7.82	9.90
1	35	***D***	**10.41**	**11.29**	7.72	**9.70**
		DC	10.47	11.31	7.80	9.76
		DCpH	10.51	11.36	7.90	9.82
		DpH	10.50	11.31	7.87	9.79
		DT	10.55	11.47	**7.68**	9.78
		DTC	10.55	11.41	**7.68**	9.76
		DTCpH	10.54	11.37	7.78	9.78
		DTpH	10.54	11.40	7.76	9.78
1	42	***D***	**10.37**	11.19	7.80	**9.69**
		DC	10.48	**11.16**	7.86	9.71
		DCpH	10.47	11.23	7.92	9.76
		DpH	10.43	11.19	7.89	9.73
		DT	10.52	11.31	7.81	9.76
		DTC	10.60	11.24	**7.77**	9.73
		DTCpH	10.50	11.28	7.87	9.77
		DTpH	10.47	11.27	7.88	9.77
2	42	***D***	**10.47**	11.22	**7.71**	**9.68**
		DC	10.58	**11.17**	7.77	9.70
		DCpH	10.59	11.22	7.85	9.75
		DpH	10.56	11.20	7.80	9.72
		DT	10.62	11.27	7.74	9.74
		DTC	10.69	11.23	7.73	9.73
		DTCpH	10.63	11.26	7.82	9.77
		DTpH	10.60	11.25	7.81	9.75

Note: The best overall model and $RMSE_f$ by season are in bold for each combination of p and window size. The best seasonal models are italicized.

corresponding to the model performance in Early Summer, Late Summer, and Early Fall. The "Overall" column is the total $RMSE_f$ across all forecasts. Prediction band computation is discussed with examples in the Supplementary Materials.

We found that for the one-step-ahead forecast, the overall best model across all window sizes and choices of p is the model with D alone, the 42-day window, and $p = 2$. This model performed similarly to the model with window size 35 and $p = 1$ with an overall $RMSE_f$ of 9.68 compared to 9.70, respectively. Using a smaller lag ($p = 1$ rather than $p = 2$) would be preferred as it reduces the number of parameters to estimate. It appears that in the Late Summer, for some model

TABLE 3 $RMSE_f$ for FPC VARX day-ahead forecasting models of various p, window sizes, and exogenous variables

p	Win. size	Fore. type	Variable(s)	Early Summer	Late Summer	Early Fall	Overall
1	28	Direct	D	22.00	23.23	16.62	20.36
			DC	21.80	23.33	**16.39**	20.27
			DCpH	**20.99**	23.11	16.50	20.08
			DpH	21.51	**22.37**	16.64	**19.92**
			DT	22.00	23.03	16.79	20.34
			DTC	21.89	23.38	16.45	20.33
			DTCpH	21.18	23.32	16.80	20.33
			DTpH	21.26	22.75	17.16	20.24
		Iterative	D	22.69	25.36	18.54	22.16
1	35	Direct	*D*	21.37	23.44	**16.40**	20.25
			DC	22.02	23.41	16.44	20.36
			DCpH	21.03	**21.87**	16.51	**19.57**
			DpH	21.73	22.38	16.67	19.97
			DT	22.91	23.04	16.93	20.56
			DTC	22.11	22.89	16.59	20.22
			DTCpH	**21.02**	21.97	16.68	19.68
			DTpH	21.68	22.36	17.29	20.20
		Iterative	D	22.08	25.10	19.13	22.18
1	42	Direct	*D*	**21.28**	23.63	16.99	20.55
			DC	21.39	23.20	17.38	20.54
			DCpH	21.31	22.36	17.29	20.14
			DpH	22.11	22.53	17.25	20.33
			DT	22.66	23.21	17.44	20.79
			DTC	21.71	22.69	**16.88**	20.18
			DTCpH	21.38	**22.34**	17.13	**20.08**
			DTpH	22.00	22.58	17.77	20.54
		Iterative	D	21.96	25.27	20.00	22.58
2	42	Direct	D	21.72	23.79	17.20	20.78
			DC	21.77	23.34	17.49	20.72
			DCpH	21.63	**22.39**	17.37	20.24
			DpH	22.57	22.60	17.39	20.49
			DT	23.22	23.29	17.69	21.02
			DTC	22.05	22.78	**17.01**	20.33
			DTCpH	21.69	22.41	17.30	**20.23**
			DTpH	22.47	22.66	18.04	20.77
		Iterative	*D*	**21.56**	25.12	18.93	22.02

Note: Both direct and iterative forecasting models are also included. The best overall model and $RMSE_f$ by season are in bold for each combination of p and window size. The best seasonal models are italicized.

combinations, the inclusion of some exogenous variables becomes helpful in reducing forecast error, particularly *C* and *pH*.

For the day-ahead forecasts, the direct model with variables *DCpH*, window size of 35 days, and $p = 1$ has the lowest RMSE_f. The iterative forecasting models do not perform as well as the direct models, and it is possible that this occurs because the exogenous variables in the direct model begin to improve the forecasts for the longer horizon. Indeed, for the direct models, different combinations of *T*, *C*, and *pH* all appear to improve forecasts for various model combinations and seasons. The iterative forecasting results are explored in more depth in the Supplementary Materials.

In terms of the best window size across forecast horizons, the 35-day window seems to provide a good balance for the application. The smaller, 28-day window may not contain enough observations for good parameter estimation, and the larger, 42-day window spans multiple months, potentially introducing nonstationarity into the training window. Based on these considerations, we fix $p = 1$ and the window size at 35 days to further investigate the FPC VARX method results for the following four model combinations: (1) one-step-ahead with *D*, (2) one-step-ahead with *DTCpH*, (3) day-ahead with *D*, and (4) day-ahead with *DTCpH*.

4.2 | Comparison methods

To evaluate the performance of the FPC VARX method, we compare it with several alternative types of models for which brief descriptions are given here:

1. *Persistence*: $W_N^K(s)$ directly forecasts $W_{N+h}^K(s)$.
2. *FPC ARX*: The EFPC scores are forecast by treating each one as a univariate time series, specifically an AR(p) with or without exogenous variables. Then, the forecast function is constructed using Equation (8) (Hyndman & Ullah, 2007).
3. *DFPCX*: Dynamic FPCs (DFPC) are derived from the eigenfunctions of the spectral density operator of $Y_1^K(s), \ldots, Y_N^K(s)$ as described in Hörmann et al. (2015) and are uncorrelated at all leads and lags, a property that regular FPCs do not possess. The goal is to obtain empirical DFPC (EDFPC) curves and scores that account for dependence in the observations. EDFPC scores are then forecast using a VAR model. Default settings from the `freqdom` package (Hörmann & Kidziński, 2017) are used.
4. *VAR*: A VAR(1) model is fit on the original 21-dimensional vector-valued observations of *D*.
5. *AR*: One univariate AR(1) model is fit for observations at each of the 21 depths of ~~the 21 values of~~ *D* that correspond to the 21 measurement depths.
6. *FAR*: The functional autoregressive model is the functional extension of the classic AR(p) model. An estimate of a functional autocovariance operator is obtained analogously to the univariate estimated autocovariance function. A useful overview is given in Kokoszka and Reimherr (2017). This method has been extended to incorporate exogenous variables (Damon & Guillas, 2005), but the software functionality does not allow lagged exogenous variables, restricting our use to the models with *D* alone.
7. *NN*: Neural network models for the 21-dimensional input vectors of *D* are fit with the `keras` package in `R` (Falbel et al., 2021). In each window, a model is fit with 750 epochs and one hidden layer for models using *D* alone and two hidden layers for models using *DTCpH*. Hyperbolic tangent activation functions are used between layers, and a linear activation function is applied for the output layer. The NN outputs a vector of 21 values of *D* corresponding to each depth. Details on the architecture selection can be found in the Supplementary Materials.

The performance of the methods is evaluated with either the RMSE_f metric, the RMSE_d metric, or both, depending on whether the method produces a forecast in the form of a smooth function, a 21-dimensional vector, or both, respectively. RMSE_f alone is computed for FPC ARX; RMSE_d alone is computed for DFPCX, VAR, AR, FAR, and NN; and both are computed for Persistence and FPC VARX. The RMSE_f and RMSE_d metrics are not directly comparable.

Results are displayed in Table 4. Within the overall column, the best performing method is clearly FPC VARX, with the lowest overall RMSE_f and RMSE_d for every horizon and variable combination except the one-step-ahead forecasts with all exogenous variables. FPC VARX reduces the overall RMSE_f over the baseline persistence method by 5% for the one-step-ahead forecasts with and without exogenous variables and by 10% and 12% for the day-ahead forecasts without and with exogenous variables, respectively. Compared to the industry-standard NN, the FPC VARX method reduces overall RMSE_d by 28% and 39% for one-step-ahead forecasts and 3% and 7% for day-ahead

TABLE 4 Comparison of eight methods with two choices of forecast horizon and two choices of variables (no exogenous variables or all exogenous variables)

Horizon	Variable(s)	Method	Metric	Early Summer	Late Summer	Early Fall	Overall
1	D	*Persistence*	$RMSE_f$	11.16	12.52	**7.58**	10.30
		FPC VARX		**10.41**	**11.29**	7.72	**9.70**
		FPC ARX		**10.41**	11.33	7.93	9.80
		Persistence	$RMSE_d$	11.92	13.08	**7.91**	10.80
		FPC VARX		**11.17**	**11.79**	8.03	**10.17**
		DFPCX		13.08	13.09	10.39	12.00
		VAR		11.47	11.99	8.15	10.36
		AR		11.74	12.54	8.64	10.83
		FAR		11.42	12.07	8.28	10.44
		NN		14.99	15.76	12.18	14.19
1	DTCpH	*Persistence*	$RMSE_f$	11.16	12.52	**7.58**	10.30
		FPC VARX		10.54	11.37	7.78	9.78
		FPC ARX		**10.46**	**11.30**	7.68	**9.70**
		Persistence	$RMSE_d$	11.92	13.08	**7.91**	10.80
		FPC VARX		**11.30**	**11.87**	8.06	**10.24**
		DFPCX		19.85	16.07	11.92	15.05
		VAR		11.76	12.10	8.22	10.48
		AR		11.87	12.58	8.41	10.78
		NN		17.22	18.45	15.04	16.87
12	D	Persistence	$RMSE_f$	26.44	25.49	17.84	22.57
		FPC VARX		21.37	**23.44**	**16.40**	**20.25**
		FPC ARX		**21.28**	23.53	16.85	20.45
		Persistence	$RMSE_d$	26.98	25.79	18.13	22.91
		FPC VARX		**21.90**	**23.70**	**16.73**	**20.58**
		DFPCX		22.20	24.12	17.04	20.94
		VAR		22.49	24.15	17.02	20.99
		AR		22.03	24.26	17.34	21.09
		FAR		22.72	25.62	20.30	22.98
		NN		22.54	24.18	17.67	21.28
12	DTCpH	Persistence	$RMSE_f$	26.44	25.49	17.84	22.57
		FPC VARX		**21.02**	21.97	**16.68**	**19.68**
		FPC ARX		21.12	**21.93**	16.78	19.72
		Persistence	$RMSE_d$	26.98	25.79	18.13	22.91
		FPC VARX		**21.58**	**22.28**	**17.03**	**20.04**
		DFPCX		35.08	27.04	21.28	26.10
		VAR		22.91	22.69	17.69	20.71
		AR		22.61	23.04	17.19	20.61
		NN		23.51	24.50	17.89	21.67

Note: All methods use a window size of 35 days and $p = 1$. Some methods are evaluated with the $RMSE_f$ metric and some with $RMSE_d$. The best overall model by season is in bold for each combination of horizon, model, and metric. The best seasonal models are italicized.

forecasts with D and $DTCpH$, respectively. The closest method to FPC VARX in terms of performance is FPC ARX. If one method and variable combination were to be chosen for each forecast horizon, it would be FPC VARX with no exogenous variables for one-step-ahead forecasting and FPC VARX with all exogenous variables for day-ahead forecasting. In the remainder of the article, these models will be referred to as the optimal $h = 1$ and $h = 12$ FPC VARX models.

The inclusion of exogenous variables affects the one-step-ahead and day-ahead models differently. For the one-step-ahead results, the inclusion of exogenous variables results in an overall performance decrease for all methods except those relying on univariate autoregressive methods: FPC ARX and AR, which improve with exogenous variables. For the day-ahead results, the inclusion of exogenous variables results in an overall performance improvement for all methods except for the DFPCX and NN methods, which worsen with exogenous variables.

4.3 | Band depth analysis

Functional boxplots (Genton & Sun, 2020) of the forecast errors are shown in Figure 3. The solid central line is the median curve, the dark gray envelope is the 50% region, the solid lines extending from the 50% region are the functional whiskers, and any dashed lines are outlier curves. Not surprisingly, the one-step-ahead errors are less variable than the day-ahead errors. The day-ahead model had its most variable forecast errors at the lake surface in the Early Fall. Several unusually shaped forecast errors are also identified.

Profiles of D with extreme values or unusual shapes are of interest for environmental and water quality monitoring. The outlying forecast errors identified in Figure 3 fall into one of two categories: either a shift in the shape or magnitude of the entire process occurs at that point in time or a "traditional" outlier occurs where the curve is distinct from both the preceding and subsequent curves. These traditional outliers may occur due to a short-duration stormy weather event or some other significant atmospheric variation (Andersen et al., 2020). Some examples of both types of outliers can be seen in the Supplementary Materials.

To understand how the FPC VARX method performs in forecasting curves of relatively high or low values of D, we order the observations by band depth, and split them into deciles from smallest (most extreme/outlying) to largest (most central) band depth. For each decile, we compute the decile's mean forecast RMSE_f. Figure 4 reveals that the extreme functional observations with the lowest values of band depth are more difficult to forecast, while the more typically shaped functions correspond to lower forecast error, most notably in Early Fall. The persistence forecasts that are shown

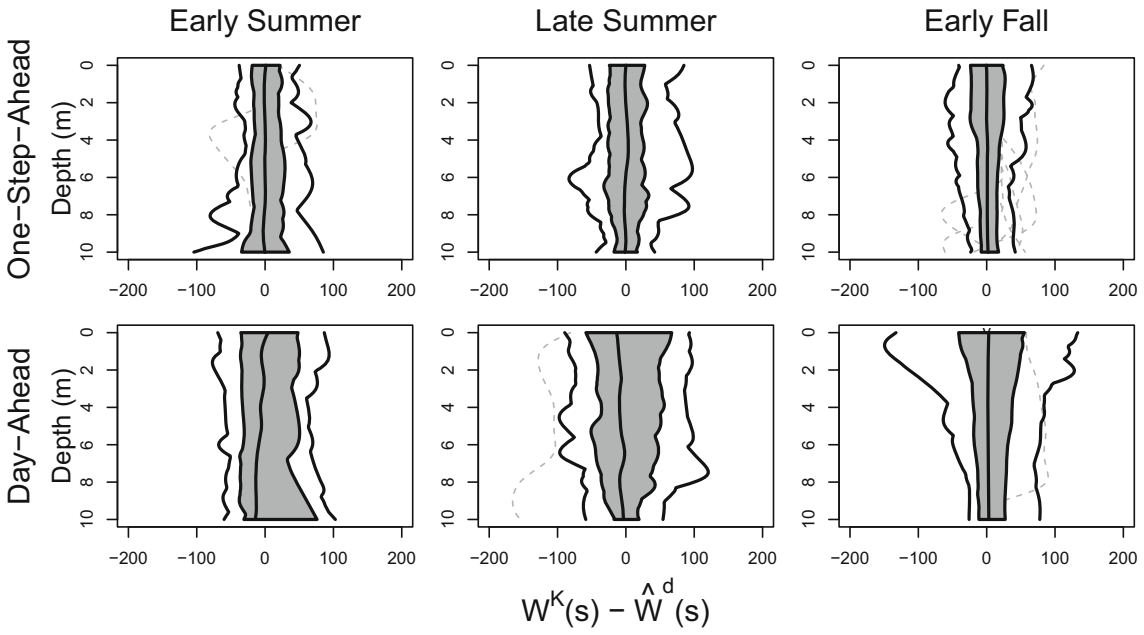

FIGURE 3 Functional boxplots for the one-step-ahead (top row) and day-ahead (bottom row) forecasting models' forecast errors by season (columns)

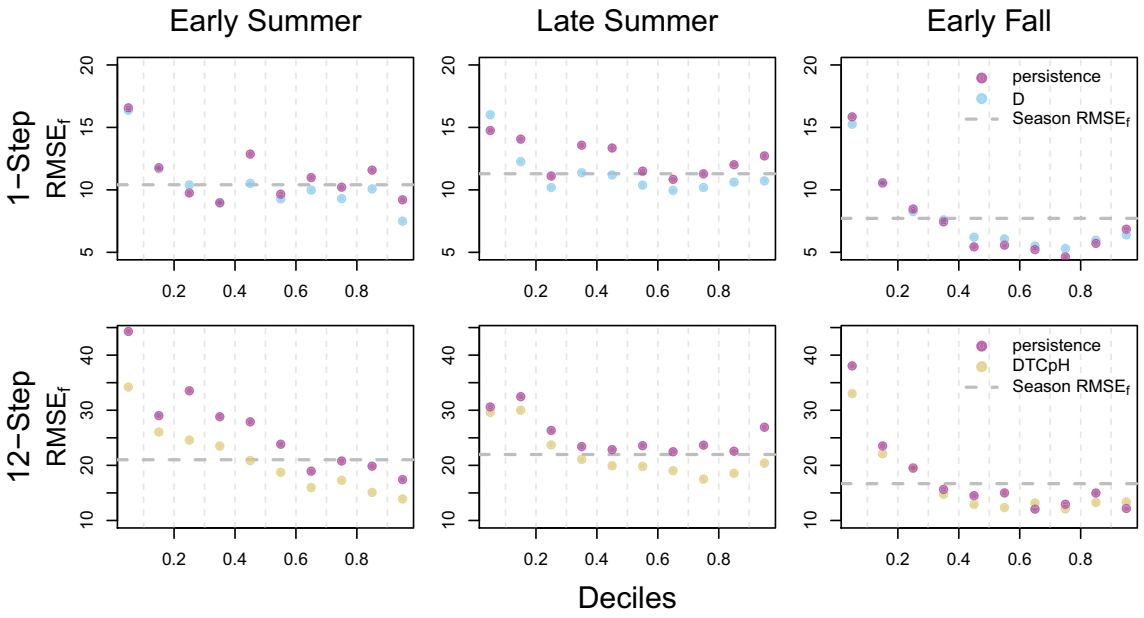

FIGURE 4 RMSE$_f$ for the one-step (top row) and day-ahead (bottom row) forecasting models, based on the deciles of a profile's band depth

for reference most nearly match the FPC VARX forecasts during the Early Fall. In the Early Summer, the day-ahead FPC VARX method performs increasingly better than persistence moving from central to smaller deciles (and more unusual curves).

5 | CONCLUSION

This work represents the first example of FDA applied to forecasting profiles of DO in a lake. We first use B-spline basis functions combined with penalized estimation to smooth the discrete observations into curves. We center the data by hourly mean curves to account for the diurnal cycle and apply FPCA to obtain score vectors and FPCs. We use a rolling window to estimate parameters in each model and to make forecasts. The FPC VARX method is the best method among those compared in almost all cases, for both one-step-ahead and day-ahead overall forecasts.

The work in this article illustrates the complexities of forecasting functional data, such as accounting for nonstationarity, implementing a smoothing step, and choosing d. One challenge we encounter is the increase in parameters to estimate due to the incorporation of exogenous variables or from increasing d. While our analysis does establish that exogenous variables improve the day-ahead forecasts, the model specifications require that d_ℓ remain so low that the exogenous KL approximations can differ from their penalized smoothed counterparts.

Overall, the FPC-based techniques outperform models that first treat the values from each depth as a vector or a univariate time series and then apply non-FTS analysis. Additionally, the comparison of FPC VARX to NN methods is valuable as NNs are the current models of choice for data-driven approaches to water quality modeling (Chen & Liu, 2014; Huo et al., 2013; Karakaya et al., 2013; Saber et al., 2020). More work could be done in this area to apply NNs or other machine learning methods directly to the FPC scores, resulting in a hybrid statistical-machine learning model.

Future research could include incorporating robust methods to address outlying curves and extreme values, possibly considering the approach of Shang and Xu (2021). Additionally, as more full-profile water quality data are collected, models could be improved in multiple ways. For example, seasonal patterns could be modeled directly, thereby reducing the need for a rolling window and increasing the sample size. Higher or variable VAR orders could be incorporated, which could improve the modeling of variable dynamics in the data. More sophisticated dynamic FPCA approaches could be explored for forecasting water quality profiles, such as Elayouty et al. (2022) who address nonstationarity in the functional covariance structure. Finally, larger sample sizes would allow for higher dimensionality used in the exogenous variable FPCA as well as comparison with more complex NNs and other parameter-rich approaches.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation Harnessing the Data Revolution: Data Science Corps project #1924146 to Amanda S. Hering, with additional support from the Tarrant Regional Water District to J. Thad Scott for the collection of the EML data. We would like to thank Dr. Nicole Wagner, Postdoctoral Research Fellow, Department of Biology, Baylor University who offered important insight into the data collection process and Dr. Alexander Aue, Professor, Department of Statistics, UC Davis who shared R code from Aue et al. (2015).

REFERENCES

Andersen, M. R., de Eyto, E., Dillane, M., Poole, R., & Jennings, E. (2020). 13 years of storms: An analysis of the effects of storms on lake physics on the Atlantic fringe of Europe. *Water*, *12*(2), 318.

Atkinson, C. A., Jolley, D. F., & Simpson, S. L. (2007). Effect of overlying water pH, dissolved oxygen, salinity and sediment disturbances on metal release and sequestration from metal contaminated marine sediments. *Chemosphere*, *69*(9), 1428–1437.

Aue, A., Norinho, D. D., & Hörmann, S. (2015). On the prediction of stationary functional time series. *Journal of the American Statistical Association*, *110*(509), 378–392.

Banks, J. L., Ross, D. J., Keough, M. J., Eyre, B. D., & Macleod, C. K. (2012). Measuring hypoxia induced metal release from highly contaminated estuarine sediments during a 40 day laboratory incubation experiment. *Science of the Total Environment*, *420*, 229–237.

Beyaztas, U., & Shang, H. L. (2019). Forecasting functional time series using weighted likelihood methodology. *Journal of Statistical Computation and Simulation*, *89*(16), 3046–3060.

Boehrer, B., & Schultze, M. (2008). Stratification of lakes. *Reviews of Geophysics*, *46*(2), 1-2.

Bosq, D. (2000). *Linear processes in function spaces: Theory and applications*. Lecture Notes in Statistics. Springer-Verlag.

Cabrera, B. L., & Schulz, F. (2017). Forecasting generalized quantiles of electricity demand: A functional data approach. *Journal of the American Statistical Association*, *112*(517), 127–136.

Chen, W.-B., & Liu, W.-C. (2014). Artificial neural network modeling of dissolved oxygen in reservoir. *Environmental Monitoring and Assessment*, *186*(2), 1203–1217.

Chen, W.-B., & Liu, W.-C. (2015). Water quality modeling in reservoirs using multivariate linear regression and two neural network models. *Advances in Artificial Neural Systems*, *2015*, 1–12.

Damon, J., & Guillas, S. (2005). Estimation and simulation of autoregressive hilbertian processes with exogenous variables. *Statistical Inference for Stochastic Processes*, *8*(2), 185–204.

Eggleton, J., & Thomas, K. V. (2004). A review of factors affecting the release and bioavailability of contaminants during sediment disturbance events. *Environment International*, *30*(7), 973–980.

Elayouty, A., Scott, M., & Miller, C. (2022). Time-varying functional principal components for non-stationary $EpCO_2$ in freshwater systems. *Journal of Agricultural, Biological and Environmental Statistics*, *27*, 506–522.

Falbel, D., Allaire, J., Chollet, F., Tang, Y., Van Der Bijl, W., Studer, M., Keydana, S., & RSudio, Google. (2021). *keras: R interface to 'Keras'*. R package version 2.4.0.

Fortuna, F., Gattone, S. A., & Di Battista, T. (2020). Functional estimation of diversity profiles. *Environmetrics*, *31*(8), e2645.

Gao, Y., & Shang, H. L. (2017). Multivariate functional time series forecasting: Application to age-specific mortality rates. *Risks*, *5*(2), 21.

Genton, M. G., & Sun, Y. (2020). *Functional data visualization*. In N. Balakrishnan, T. Colton, B. Everitt, W. Piegorsch, F. Ruggeri, & J. Teugels (Eds.), *Wiley StatsRef: Statistics reference online* (pp. 1–11). Wiley.

Harris, T., Li, B., & Tucker, J. D. (2022). Scalable multiple changepoint detection for functional data sequences. *Environmetrics*, *33*(2), e2710.

Henderson, B. (2006). Exploring between site differences in water quality trends: A functional data analysis approach. *Environmetrics*, *17*(1), 65–80.

Hodges, B. R., Imberger, J., Saggio, A., & Winters, K. B. (2000). Modeling basin-scale internal waves in a stratified lake. *Limnology and Oceanography*, *45*(7), 1603–1620.

Hörmann, S., & Kidziński, L. (2017). *freqdom: Frequency domain based analysis: Dynamic PCA*. R package version 2.0.1.

Hörmann, S., Kidziński, L., & Hallin, M. (2015). Dynamic functional principal components. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *77*(2), 319–348.

Hörmann, S., & Kokoszka, P. (2010). Weakly dependent functional data. *The Annals of Statistics*, *38*(3), 1845–1884.

Hörmann, S., & Kokoszka, P. (2012). *Functional time series*. In T. S. Rao, S. S. Rao, & C. Rao (Eds.), *Time series analysis: methods and applications* (Vol. *30*, pp. 157–186). Elsevier Science & Technology.

Huo, S., He, Z., Su, J., Xi, B., & Zhu, C. (2013). Using artificial neural network models for eutrophication prediction. *Procedia Environmental Sciences*, *18*, 310–316.

Hyndman, R. J., & Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics and Data Analysis*, *51*(10), 4942–4956.

Istvánovics, V., & Honti, M. (2018). Coupled simulation of high-frequency dynamics of dissolved oxygen and chlorophyll widens the scope of lake metabolism studies. *Limnology and Oceanography*, *63*(1), 72–90.

Jin, T., Cai, S., Jiang, D., & Liu, J. (2019). A data-driven model for real-time water quality prediction and early warning by an integration method. *Environmental Science and Pollution Research*, *26*(29), 30374–30385.

Karakaya, N., Evrendilek, F., Gungor, K., & Onal, D. (2013). Predicting diel, diurnal and nocturnal dynamics of dissolved oxygen and chlorophyll-a using regression models and neural networks. *CLEAN – Soil, Air Water*, *41*(9), 872–877.

King, A. J., Tonkin, Z., & Lieshcke, J. (2012). Short-term effects of a prolonged blackwater event on aquatic fauna in the Murray river, Australia: Considerations for future events. *Marine and Freshwater Research*, *63*(7), 576–586.

Klepsch, J., Klüppelberg, C., & Wei, T. (2017). Prediction of functional arma processes with an application to traffic data. *Econometrics and Statistics*, *1*, 128–149.

Kokoszka, P., & Reimherr, M. (2017). *Introduction to functional data analysis*. Chapman & Hall/CRC Press.

Lütkepohl, H. (2005). *Stable vector autoregressive processes*. In *New introduction to multiple time series analysis* (pp. 13–68). Springer.

Marti, C. L., Mills, R., & Imberger, J. (2011). Pathways of multiple inflows into a stratified reservoir: Thomson reservoir, Australia. *Advances in Water Resources*, *34*(5), 551–561.

Qu, Z., Dai, W., & Genton, M. G. (2021). Robust functional multivariate analysis of variance with environmental applications. *Environmetrics*, *32*(1), e2641.

Ramsay, J., Graves, S., & Hooker, G. (2021). *FDA: Functional data analysis*. R package version 5.5.0.

Ramsay, J., & Silverman, B. W. (2005). *Functional data analysis. springer series in statistics* (2nd ed.). Springer-Verlag.

Ranković, V., Radulović, J., Radojević, I., Ostojić, A., & Čomić, L. (2010). Neural network modeling of dissolved oxygen in the Gruža reservoir, Serbia. *Ecological Modelling*, *221*(8), 1239–1244.

Ranković, V., Radulović, J., Radojević, I., Ostojić, A., & Čomić, L. (2011). Prediction of dissolved oxygen in reservoirs using adaptive network-based fuzzy inference system. *Journal of Hydroinformatics*, *14*(1), 167–179.

Riedel, G. F., Sanders, J. G., & Osman, R. W. (1997). Biogeochemical control on the flux of trace elements from estuarine sediments: Water column oxygen concentrations and Benthic Infauna. *Estuarine, Coastal and Shelf Science*, *44*, 23–38.

Saber, A., James, D. E., & Hayes, D. F. (2020). Long-term forecast of water temperature and dissolved oxygen profiles in deep lakes using artificial neural networks conjugated with wavelet transform. *Limnology and Oceanography*, *65*(6), 1297–1317.

Sancho, J., Iglesias, C., Piñeiro, J., Martínez, J., Pastor, J. J., Araúo, M., & Taboada, J. (2016). Study of water quality in a Spanish river based on statistical process control and functional data analysis. *Mathematical Geosciences*, *48*(2), 163–186.

Seki, H., Takahashi, M., Hara, Y., & Ichimura, S. (1980). Dynamics of dissolved oxygen during algal bloom in lake Kasumigaura, Japan. *Water Research*, *14*(2), 179–183.

Shang, H. L. (2017). Forecasting intraday s&p 500 index returns: A functional time series approach. *Journal of Forecasting*, *36*(7), 741–755.

Shang, H. L. (2019). A robust functional time series forecasting method. *Journal of Statistical Computation and Simulation*, *89*(5), 795–814.

Shang, H. L. (2020). Dynamic principal component regression for forecasting functional time series in a group structure. *Scandinavian Actuarial Journal*, *2020*(4), 307–322.

Shang, H. L., & Hyndman, R. J. (2017). Grouped functional time series forecasting: An application to age-specific mortality rates. *Journal of Computational and Graphical Statistics*, *26*(2), 330–343.

Shang, H. L., & Xu, R. (2021). Functional time series forecasting of extreme values. *Communications in Statistics: Case Studies, Data Analysis and Applications*, *7*(2), 182–199.

Šiljić Tomić, A., Antanasijević, D., Ristić, M., Perić-Grujić, A., & Pocajt, V. (2018). A linear and non-linear polynomial neural network modeling of dissolved oxygen content in surface water: Inter- and extrapolation performance with inputs' significance analysis. *Science of the Total Environment*, *610-611*, 1038–1046.

Yao, F., Muller, H.-G., & Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, *100*(470), 577–590.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Durell, L., Scott, J. T., Nychka, D., & Hering, A. S. (2023). Functional forecasting of dissolved oxygen in high-frequency vertical lake profiles. *Environmetrics*, *34*(4), e2765. https://doi.org/10.1002/env.2765