DWR-1288

Environmental Modelling & Software 93 (2017) 193 208

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

A hybrid empirical-Bayesian artificial neural network model of salinity in the San Francisco Bay-Delta estuary

John S. Rath [a], Paul H. Hutton [b], Limin Chen [a], Sujoy B. Roy [a,*]

[a] Tetra Tech, Lafayette, CA, United States
[b] Metropolitan Water District of Southern California, Sacramento, CA, United States

ARTICLE INFO

Article history:
Received 1 September 2016
Accepted 22 March 2017

Keywords:
Salinity
Artificial neural network
Structural validation
Bayesian parameter estimation

ABSTRACT

This paper reports the refinement of a published empirical model of salinity in the San Francisco Bay Delta estuary by integration with a Bayesian artificial neural network (ANN) model and incorporation of additional inputs. Performance goals established for the resulting hybrid model are based on the quality of fit to observed data (replicative and predictive validation) as well as sensitivity when compared with *a priori* knowledge of system behavior (structural validation). ANN model parameters were con strained to provide plausible sensitivity to coastal water level, a key input introduced in the hybrid formulation. In addition to representing observed data better than the underlying empirical model while meeting structural validation goals, the hybrid model allows for characterization of prediction uncer tainty. This work demonstrates a real world application of a general approach integration of a pre existing model with a Bayesian ANN constrained by knowledge of system behavior that has broad application for environmental modeling.

1. Introduction

Artificial neural networks (ANNs) have seen growing application in the modeling of water resources and water quality processes over the past two decades. Reviews by Maier and Dandy (1996), Maier et al. (2010), and Wu et al. (2014) have documented more than 300 publications on the field, and readers are referred to these papers for a broad discussion of potential applications. ANN models are particularly attractive because they can be developed using standardized approaches even where mechanistic representations of a system are unavailable, prohibitively expensive computation ally, or are too complex to formulate. As empirical formulations, ANN models are based on data that can be directly observed or simulated, as opposed to abstract processes that are embodied in mechanistic models. To develop such models, sufficient data must be available on the dependent variable of interest as well as inde pendent variables known to affect the dependent variable, and appropriate software must be utilized to encapsulate these re lationships through a data fitting (or calibration) process. This data fitting process is commonly referred to as "training", during which the adjustable parameters of an ANN model are estimated using error minimization algorithms. Data are generally partitioned for training, with a fraction set aside for model validation. A successful model fit matches both the training and validation data. These steps are common to most modeling efforts, whether ANN based or mechanistically based, and have been termed "replicative" and "predictive" validity, respectively (Wu et al., 2014).

With the growing maturity of ANN applications in the literature, it has become clear that the "black box" model relationship be tween inputs and outputs embodied in ANNs may not adequately represent the physical system being modeled (Jain et al., 2004; Kingston et al., 2005a). Thus, a trained and validated ANN model may fit the aggregate response to multiple inputs well, even though the sensitivity to a specific input is not physically meaningful, or in some cases, not physically plausible. The condition of representing inputs and outputs in a manner that is physically plausible, given an *a priori* understanding of a system, is termed "structural" validity (Wu et al., 2014). For a model to be successful under a wide range of future conditions, both predictive and structural validity are necessary prerequisites and have been the focus of a small subset of recent ANN applications (Jain et al., 2004; Kingston et al., 2005a; Jain and Kumar, 2009). Assessment of predictive validity is often the main criterion in the development of ANN models, as noted in a review of recent literature on water quality related models by Wu

* Corresponding author.
E-mail address: sujoy.roy@tetratech.com (S.B. Roy).

http://dx.doi.org/10.1016/j.envsoft.2017.03.022

et al. (2014). Of 99 ANN models evaluated by the authors, only 15 considered structural validity as a criterion.

The focus of this work is on the development of predictively and structurally valid ANN models for representing salinity behavior in the San Francisco Bay Delta estuary, with performance better than existing models. The San Francisco Bay Delta is the largest fresh water estuary on the Pacific Coast of the United States. The man agement (and therefore accurate prediction) of salinity in this region is of great consequence; it is tied to the health of a highly diverse ecosystem in the estuary and to the management of freshwater withdrawals that support the needs or more than 20 million people and 3 million acres of irrigated land in California, thus supporting one of the largest economies in the world (Luoma et al., 2015; Delta Plan, 2013). Salinity behavior in an environment such as the San Francisco Bay Delta estuary is highly dynamic, with freshwater flow and tidal influences having time varying behavior on scales of hours to months (Schoellhamer, 2000; Monismith et al., 2002). A variety of empirical and mechanistic modeling tools have been developed to predict salinity behavior in the es tuary for management decision making. Despite the availability of these tools, there remains an ongoing need for robust modeling approaches that can be applied with rapid turnaround times; the ANN modeling framework presented here is intended to meet this specific need. This work is motivated by the general utility of ANNs in modeling complex problems in the water resources domain, as well as the specific utility of ANN applications related to salinity intrusion in the San Francisco Bay Delta estuary (Finch and Sandhu, 1995; Seneviratne et al., 2008; Rajkumar and Johnson, 2001; Chung and Seneviratne, 2009) and other estuaries (Maier and Dandy, 1996; Bowden et al., 2002, 2005; Huang and Foo, 2002; Conrads et al., 2006).

The model structure presented in this paper is a hybrid approach that integrates neural networks within an existing empirical modeling framework. Hybrid modeling approaches may achieve better results than using a single modeling approach (Maier et al., 2010) and can be used to provide some degree of structural validity in a modeling framework. Hybrid systems are an active area of research and Maier et al. (2010) see this as a matu ration of the ANN methodology where their strengths are best used in conjunction with existing modeling approaches in the water resources domain. In this work, we utilized Bayesian infer ence—the application of Bayes' rule to obtain probabilistic esti mates of parameters in statistical models—to calculate hybrid ANN model parameters; this approach allowed predictions to account for uncertainty in the model parameters and to improve structural validity by formally incorporating prior knowledge into the model through the use of prior distributions. This approach has been used with success elsewhere (Kingston et al., 2005b; Gelman et al., 2013; Humphrey et al., 2016). The consideration of structural validity in this work is of benefit by allowing for more robust model pre dictions; it also serves a broader purpose of addressing well founded skepticism of ANN models based on their "black box" character. This approach has the potential to improve overall model performance (e.g. quality of fits, incorporation of prediction un certainty, and consideration of additional inputs in a structurally valid framework) compared to existing modeling frameworks for salinity in the San Francisco Bay Delta estuary. Specific elements of this real world application also hold potential in other environ mental domains where ANN models are not targeted as a replacement for a pre existing empirical or mechanistic model, but rather are used in tandem to enhance the performance of the resulting hybrid.

2. Background

Salinity management in estuaries is a subject of interest in ba sins where a significant amount of freshwater is diverted (Alber, 2002). Examples of similar well studied systems occur globally and include the Murray Darling River basin in Australia (Murray Darling Basin Ministerial Council, 1999), estuaries in the South eastern United States along the Gulf of Mexico and Atlantic coast (Sklar and Browder, 1998; Reinert and Peterson, 2008), the Mekong River Delta in Vietnam (Dat et al., 2011), and the Guadalquivir River estuary in Spain (Fernández Delgado et al., 2007).

Freshwater flow through the northern reach of the San Francisco Bay Delta estuary, specifically the Suisun Bay and the western Delta of the Sacramento and San Joaquin Rivers (Fig. 1), is a function of seasonally and annually varying runoff as well as anthropogenic influences such as upstream reservoir releases and freshwater di versions. Over the past two decades, the balance between envi ronmental and water export needs has been managed through the location of the low salinity zone, which is correlated with various aquatic species distributions (Jassby et al., 1995). Operationally, the low salinity zone is defined as the position of the 2 parts per thousand bottom salinity isohaline, termed X2. Under current regulations, it is interpolated as an equivalent surface salinity from fixed monitoring stations near the surface and reported as a dis tance from Golden Gate Bridge in kilometers (Fig. 1). Besides X2 position, which is largely driven by habitat considerations, salinity compliance is also maintained at several discrete locations to pro tect municipal and agricultural beneficial uses (CSWRCB, 2006). For example, agricultural salinity standards are maintained in the western Delta at Emmaton and Jersey Point (Fig. 1).

The basic conceptual model of freshwater saltwater mixing in estuaries with seasonally varying flow patterns such as the San Francisco Bay Delta estuary is as follows: freshwater flows repel salinity downstream (seaward) across a mixing zone (with longi tudinal and vertical gradients) and saltwater intrudes upstream (landward) during periods of low freshwater flow. The extent of the salinity gradients varies with tides on hourly to daily time scales and varies with freshwater flows on daily to seasonal time scales. Salinity management in the estuary is primarily concerned with daily, fortnightly and seasonal variability; thus, salinity variation at sub daily time scales is beyond the scope of this work.

Over the past two decades, a variety of modeling frameworks have been developed and applied to the prediction of salinity intrusion in the San Francisco Bay Delta estuary. These frameworks range from simple empirical statistical models to complex three dimensional hydrodynamic models. Several simple empirical models are available to predict X2 position and/or salinity at fixed locations in the estuary as a function of flow on a daily time scale (Denton, 1993; Jassby et al., 1995; Monismith et al., 2002; MacWilliams et al., 2015; Hutton et al., 2015). The Delta Salinity Gradient (DSG) model proposed by Hutton et al. (2015) integrates the approaches used by Monismith et al. (2002) for X2 position prediction and Denton (1993) for fixed location salinity prediction as follows:

$$X2(t) \quad \Phi_1 * G(t)^{\Phi_2} \tag{1}$$

where G(t) is antecedent outflow and Φ_1 and Φ_2 are empirically determined constants. Antecedent outflow (Denton, 1993) is defined by the following routing function similar to one proposed by Harder (1977):

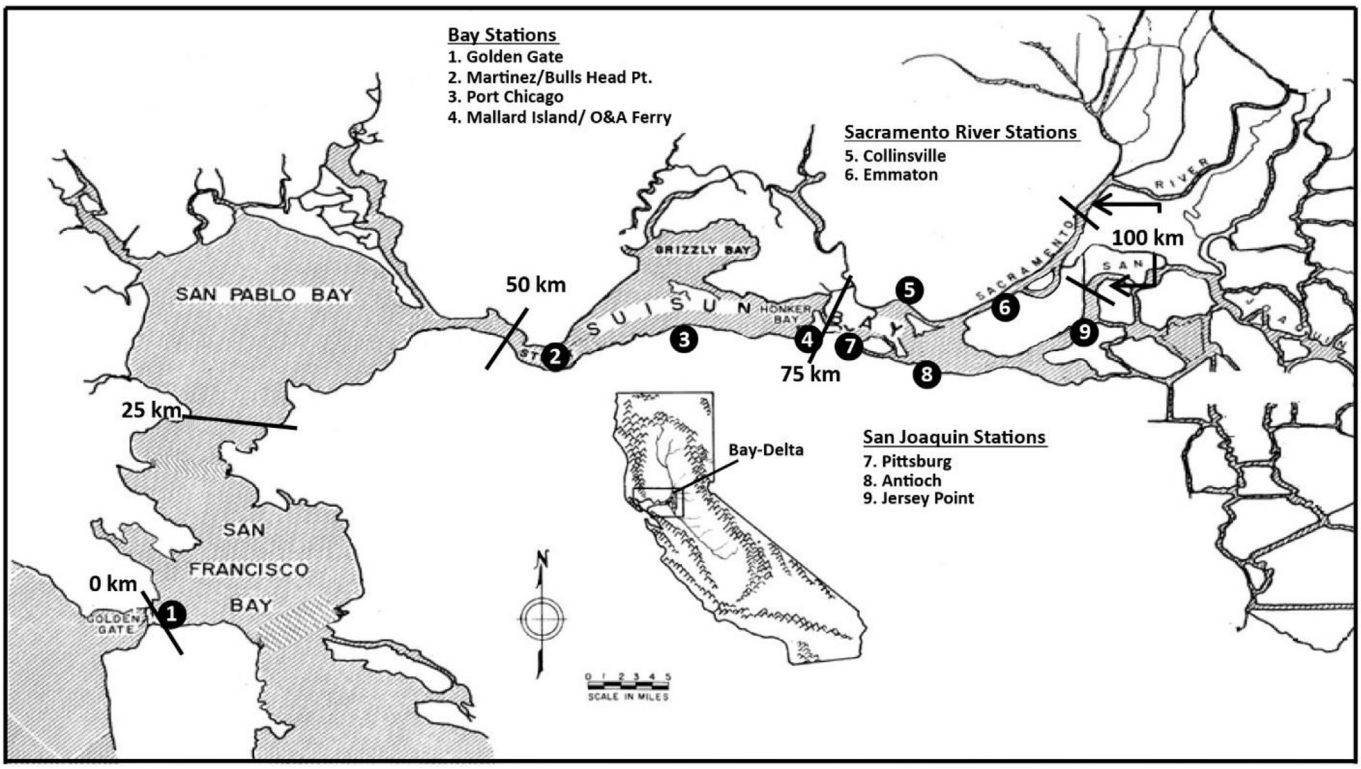

Fig. 1. Key salinity stations are identified in Suisun Bay and the western Delta. Salinity data from these and other locations were used for the ANN training and to compute X2, the position of two parts per thousand bottom salinity in the estuary. X2 position is reported as the distance in kilometers from Golden Gate. The Sacramento and San Joaquin River branches are identified on the map.

$$\frac{\partial G}{\partial t} \quad \frac{(Q(t) \quad G(t))*G(t)}{\beta} \tag{2}$$

where Q is Delta outflow and β is an empirically determined con stant. Denton (1993) observed that the term β/G(t) is a time con stant governing the rate at which G approaches steady state. Salinity is estimated at individual locations through the following relationship:

$$S(X,t) \quad (S_o \quad S_b)* \exp\left[\tau* \left(\frac{X}{X2(t)}\right)^{-\frac{1}{\Phi_2}}\right] + S_b \tag{3}$$

where S is the salinity at a fixed location (X) in the estuary, S_o and S_b are representative downstream and upstream riverine boundary salinities, $\tau \quad \ln\left[\frac{S_i - S_b}{S_o - S_b}\right]$, and S_i is a salinity index associated with the X2 isohaline. Salinity is reported in units of specific conductance or electrical conductivity (EC), standardized to 25° C and reported as milliSiemens/cm or mS/cm. Hutton et al. (2015) assumed $S_i \quad 2.64$ mS/cm, based on field observations of surface water specific conductance and consistent with X2 regulatory compli ance. The authors also assumed the representative upstream boundary salinity S_o varies as a function of X2 position:

$$S_o \quad \hat{S} + \left(S_i \quad \hat{S}\right)*\exp\left(\quad \gamma* X2^{\delta}\right) \tag{4}$$

where $\hat{S}$ is the salinity of seawater and γ and δ are empirically determined constants. Equation (1) through (4) can be used to determine salinity at any longitudinal distance from Golden Gate (X) given one dependent variable (Q) and five empirical constants (β, Φ_1, Φ_2, γ, δ). Best fit parameter values reported by Hutton et al. (2015) represent the full range of flows in the historical record. However, this model does not include a tidal term necessary to accurately represent salinity variation over spring neap tidal cycles in the estuary.

In addition to simple empirical statistical models, several nu merical (one , two and three dimensional) models have been developed to predict salinity intrusion in the estuary. These mechanistically based models have been successfully used in several applications, ranging from operations and facility planning (CDWR, 2013) to scientific exploration of fundamental estuarine mechanics (Cheng et al., 1993; Gross et al., 1999, 2007; Chua and Fringer, 2011; MacWilliams et al., 2015). Although theoretically rigorous and capable of providing insight into the basic physics of freshwater saltwater mixing in the estuary, the computational demands of these tools—especially two and three dimensional models—limit application for studies that require consideration of long hydrologic sequences and/or multiple scenarios. Impor tantly, to date, these mechanistically based models have not been calibrated over the full record of hydrologic data and are therefore of arguable reliability in simulating extremely low flow conditions that occurred in the early part the 20th century.

ANN models have also been developed and applied to the pre diction of salinity intrusion in the San Francisco Bay Delta estuary over the past two decades. An initial effort developed models to predict salinity at four discrete locations in the estuary, with training and validation based on weekly averages of observed data from 1971 to 1991 (Finch and Sandhu, 1995); a subsequent effort was undertaken by Rajkumar and Johnson (2001) using data for three years (1996–1998). Ongoing development and application of ANNs has focused on emulating data generated from a one

dimensional hydrodynamic model of the Delta (Wilbur and Munevar, 2001; Mierzwa, 2002; Senevirante et al., 2008; Chung and Seneviratne, 2009) and embedding the resulting ANN models in a water resources allocation model (Draper et al., 2004) – thereby allowing consideration of estuarine salinity constraints with minimal computational burden (Wilbur and Munevar, 2001). Importantly, these prior modeling efforts are based either on data over very limited periods or on numeric values generated from a hydrodynamic model; none have been subjected to rigorous eval uation of input sensitivity.

3. Methods

Given the flexibility and range of modeling approaches that constitute the ANN methodology at present, the need for clarity and consistency has been emphasized in previous reviews of the liter ature (Maier et al., 2010; Wu et al., 2014). To allow the specific methods in this work to be adequately evaluated and reproduced, we follow the protocol proposed by Wu et al. (2014) with the following steps: model input selection, data splitting, model ar chitecture and structure selection, and model calibration and validation.

ANN model development in this study focused on a hybrid approach, where ANNs were combined with the empirical DSG model (Hutton et al., 2015) described previously in Equation (1) through (4). As part of the hybrid model training, additional con straints on model weights were implemented through the use of a Bayesian approach (Kingston et al., 2005b). Performance of the trained models was compared against the performance of the DSG model to evaluate the incremental benefits of incorporating ANNs in the existing salinity modeling framework. Key details associated with the model development approach are presented below.

3.1. Model inputs

Selection of model inputs was guided by an underlying assumption that the variables influencing salinity are reasonably well known *a priori* and that the training approach finds the rela tionship that minimizes error between the modeled and observed salinity. The estuary's salinity gradient is governed by the opposing forces of freshwater and tidal flows. Each of these drivers can be represented by one or more specific inputs. Various combinations of inputs were explored with regard to their ability to fit observed salinity data at multiple locations in the estuary (Chen and Roy, 2013; Chen et al., 2015).

Selection of model inputs was based on comparison of model performance at key locations along the salinity gradient (see Fig. 1) and in consideration of model parsimony. The following inputs were utilized in the final models: daily freshwater flow to the es tuary, daily mean coastal water level, and the daily tidal range (difference between maximum and minimum water level). Note that the latter two inputs are not part of the DSG model and are incorporated through the ANN. Available data span a 91 year period on a daily time step from Water Year (WY) 1922–2012, where water years in California run from October 1 to September 30. Data representing freshwater flow (i.e., net Delta outflow) were obtained from the DAYFLOW program (CDWR, 2015). Data repre senting water level and tidal range at Golden Gate were obtained from the National Oceanic and Atmospheric Administration (NOAA, 2015). Salinity data were based on a composite of legacy grab sample and modern continuous sensor data collected across a network of stations in the estuary. The compilation, cleaning, and data quality assurance procedures associated with unifying the legacy and modern salinity data are reported elsewhere (Hutton et al., 2015). Because sub daily salinity data were frequently unavailable in the early part of the record, averages were computed over 24 h rather than a tidal day (25 h). Monismith et al. (2002) reported that the errors associated with this approximation were "very slight". In addition to the location specific salinity data, data representing the position of the 2 parts per thousand isohaline (X2) were also used in ANN training (Hutton et al., 2015). X2 is an interpolated value based on salinity measurements at fixed loca tions and is reported in kilometers from Golden Gate. Because the actual shape of the salinity profile between two observation sta tions (typically separated by 10 km) is not known, the interpolation method introduces some uncertainty into the estimate. Fig. 2 il lustrates how X2 position was interpolated from observed salinity data for two different measurement dates. This figure highlights that the interpolation procedure depends on measurement loca tions that vary over the 91 year period of record. Due to the length of the record, most of the sampling locations were not operational over the entire period, and in many instances the sampling loca tions were characterized by extensive periods of missing data.

Given the above model inputs, the following salinity responses are considered hydrodynamically plausible and are the basis of the structural validation of the trained ANNs: salinity decreases monotonically with distance from Golden Gate; salinity decreases with increasing freshwater flow; and salinity increases with increasing coastal water levels. No *a priori* response to the tidal range input is assumed.

3.2. Data splitting

ANN models are typically fit using a subset of the available data; this data splitting step is performed to assess model ability to generalize through comparison of predictions with the remaining "testing data" that were excluded from the fitting process. The hybrid ANN approach adopted a Bayesian technique to help prevent over fitting; no early stopping procedure was used. In this approach, a random data subset (approximately 15% or 5000 days of data) was used for training and the remaining data were used for testing. Once selected, the training data were fixed throughout the entire model calibration process. The use of larger training data sets did not improve model predictive or structural validity and suffered the drawback of longer fitting times.

3.3. Model architecture and structure

The architecture and structure of ANN models are inspired by biological neural networks such as the central nervous system of animals (in particular the brain). In an ANN model, numerous ab stract "neurons" and the strength of connections between them, collectively referred to as "weights and biases", determine the functional relationship between input and output data. Network architecture is the term encompassing important details about the structure of the network of neurons in an ANN model, such as number of neurons and the ways in which they are connected. The multi layer perceptron (MLP)—a feedforward ANN model that maps sets of input data onto a set of outputs—is by far the most popular network architecture used in water resource applications of ANN models to date, representing more than 90% of the peer reviewed applications related to water resources generally (Maier et al., 2010) and two thirds of recent publications specifically related to water quality (Wu et al., 2014). For this reason, the feedforward MLP network was selected for this work. As used here, the MLP uses a layer of neurons to represent the inputs, a layer of neurons to represent the outputs, and a single "hidden" layer of neurons to connect the input and output layers. The neurons are connected by excitation functions whose weights and biases are estimated during training. The number of neurons in the hidden

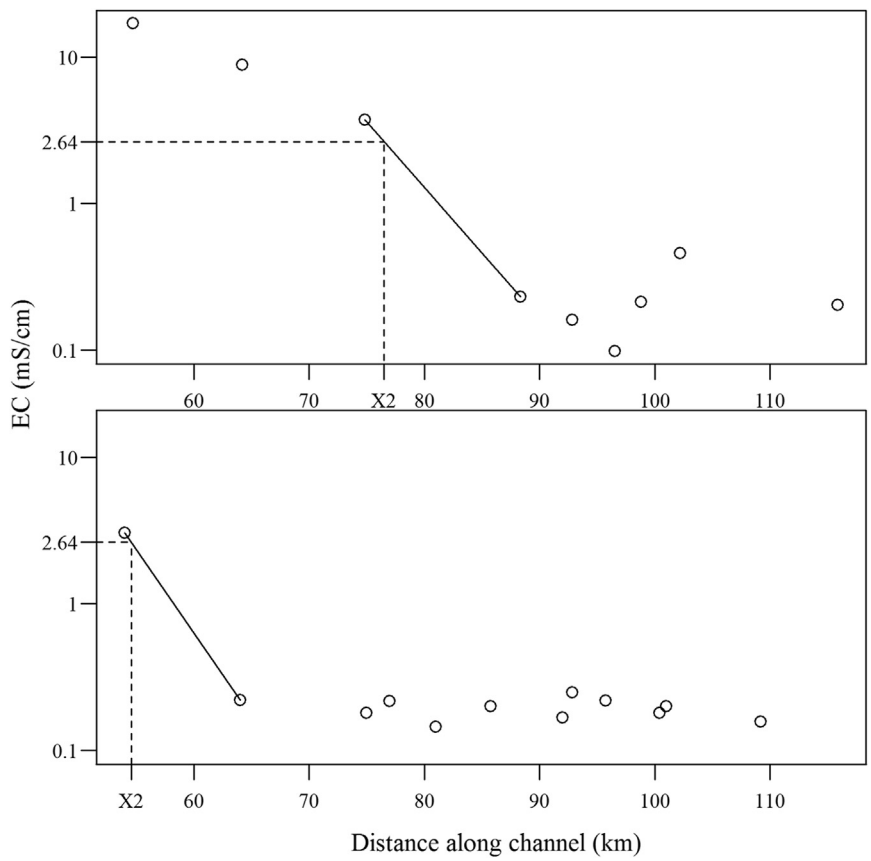

Fig. 2. The observed longitudinal salinity gradient (represented as specific conductance or EC) is shown for two example days from the 91-year data record: a moderate flow period on 7/26/1932 (top) and a high flow period on 4/10/1996 (bottom). These examples illustrate that the number and location of measurement sites (indicated by the x-axis value) vary over the full period of record. The X2 position, which is defined as the distance (in kilometers) where surface salinity is 2.64 mS/cm (indicated by the dashed lines), was estimated by interpolation (log salinity over linear distance) using the two stations bounding the 2.64 mS/cm salinity value.

layer of a MLP determines the potential complexity of the model. The MLP within the hybrid ANN model uses a hidden layer with three neurons, and increasing the number of hidden neurons above three did not substantially improve the model's predictive ability.

The hybrid ANN model exploits existing knowledge of the estuary's salinity dynamics by combining an ANN with the previously introduced DSG model (Hutton et al., 2015). This approach incorporates the physical intuition and empirical results from previous studies showing that X2 behavior is largely explained by the recent time history of net Delta outflow. For each daily time step, X2 position was modeled as the sum of the DSG model value (Eq. (1)) and a correction to that value (η) from the ANN:

$$X2 \quad \Phi_1 \cdot G^{\Phi_2} + \eta \tag{5}$$

For prediction of salinity at fixed locations along the estuary, we used the modeled X2 from Eq. (5) and the DSG formulation of the salinity gradient (Eqs. (3) and (4)); the parameters γ and δ were estimated via nonlinear least squares using the observed station salinity data and the ANN predicted X2 position.

An advantage of the hybrid approach is that it uses the DSG model to fit the station level salinity data prior to training and is thus relatively robust to extensive spatial and temporal data gaps (spanning years) that occurred in the data record. Indeed, as shown in Fig. 2 above, available data on different days may be present at different stations, and the DSG model may be fit to the gradient in each case. For this reason, the hybrid model could be developed using the full 91 year period of record, thereby exposing it to a wider range of flow and salinity conditions. The wider range of data used for training is beneficial for empirical tools such as ANNs that are poor at extrapolating beyond the training range.

3.4. Model calibration and validation

The calibration approach used in developing the hybrid ANN model is similar to that of Kingston et al. (2005b), who demonstrated the utility of Markov Chain Monte Carlo (MCMC) algorithms in Bayesian training of hydrological ANN models. Bayesian inference is based on making probabilistic statements about the parameters of interest via the application of Bayes' rule:

$$p(\theta|y,x) \quad \frac{p(\theta)p(y|\theta,x)}{\int p(\theta)p(y|\theta,x)\ d\theta}, \tag{6}$$

where θ is the set of all the model parameters being estimated (θ represents ANN weights and biases as well as the DSG parameters Φ_1 and Φ_2 in this work), y is the set of model outputs, and x is the set of model inputs. In the Bayesian paradigm, $p(\theta)$ is the prior distribution, a probability density representing the modeler's beliefs about the model parameters before the data are evaluated. Our beliefs concerning X2 response to coastal water level are reflected in our prior distributions for ANN weights and biases that are constrained to provide a positive response to increasing coastal

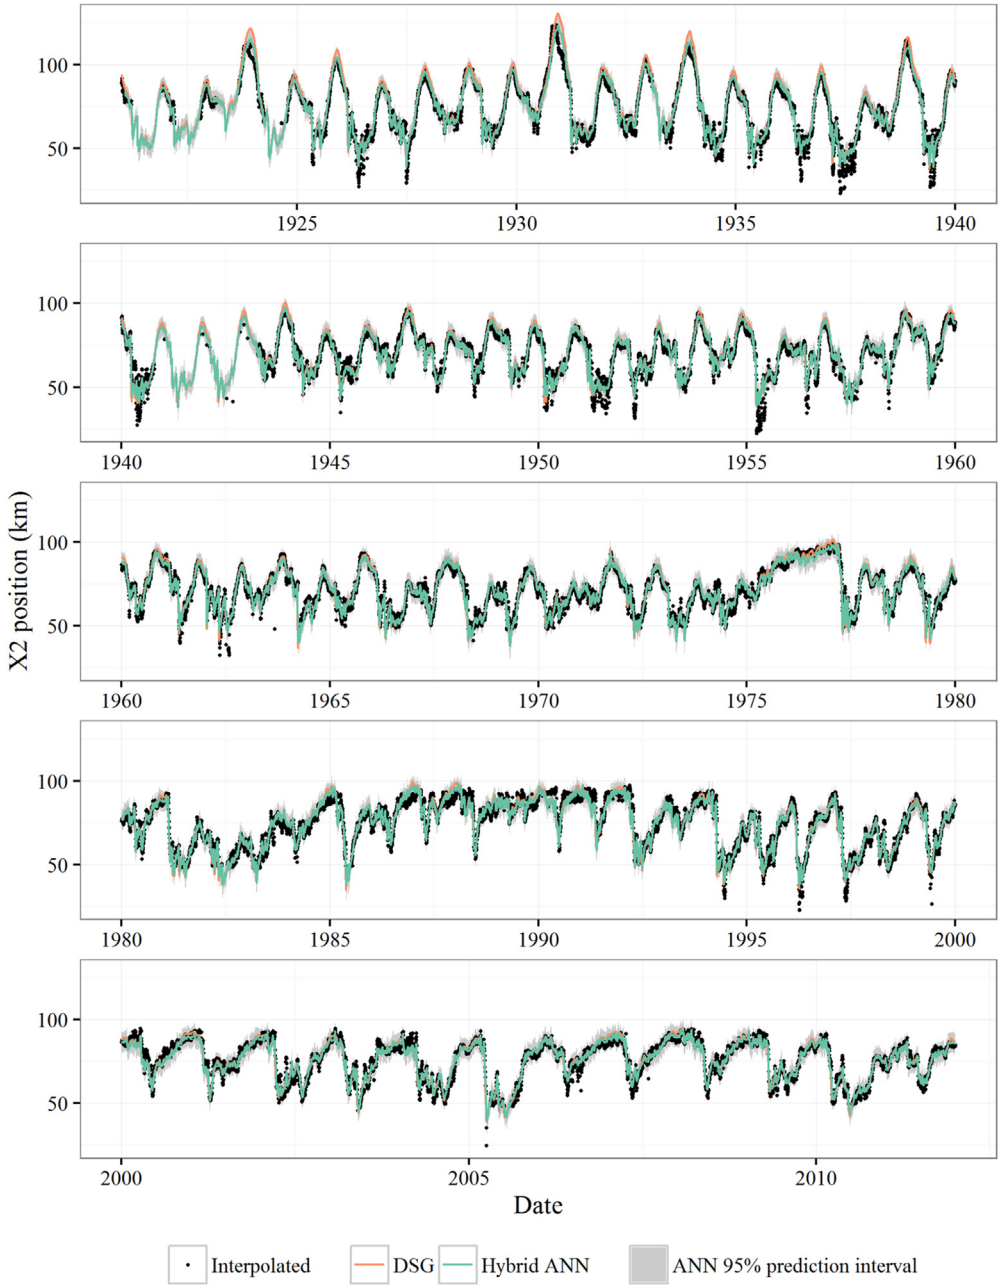

Fig. 3. The 91-year time series of observed X2 values is compared with predicted values from the DSG and hybrid ANN models. Observed X2 values are interpolated from daily salinity measurements.

Table 1
Evaluation of model fits to interpolated X2 from observed salinity data.

River Branch	Count	Mean Residual (km)		Coefficient of Determination (r^2)		Standard Error (km)	
		DSG	Hybrid ANN	DSG	Hybrid ANN	DSG	Hybrid ANN
Sacramento	30,753	0.76	<0.01	0.93	0.95	3.63	3.22
San Joaquin	30,224	−0.31	−0.07	0.93	0.94	4.03	3.73

water levels. The prior is combined with the likelihood, $p(y|\theta, x)$—a probability model for the data if the parameters were known—to obtain a probability density for the model parameters called the posterior distribution: $p(\theta|y, x)$.

The denominator of Eq. (6) is a normalizing constant depending only on the model inputs and outputs, which ensures that the posterior is a proper density function integrating to unity (Gelman et al., 2013). In general, this quantity must be approximated unless

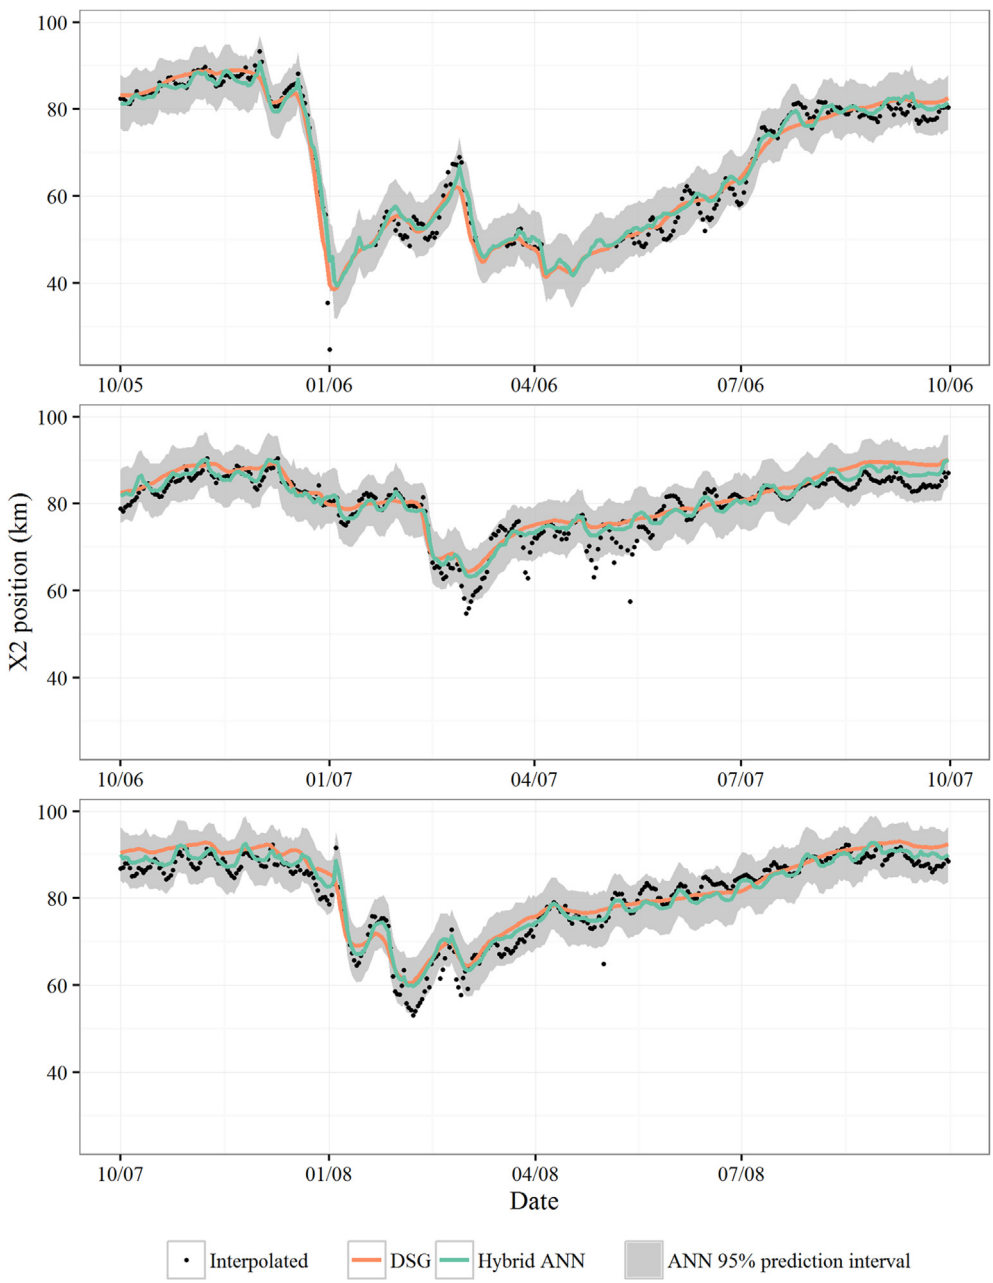

Fig. 4. A three-year (WY 2006 2008) time series of observed X2 values is compared with predicted values from the DSG and hybrid ANN models. The hybrid ANN model captures some of the spring-neap tidal variation of the salinity field.

limited, mathematically convenient pairs of $p(\theta)$ and $p(y|\theta, x)$ (known as conjugate priors) are used. Numerous methods exist for approximating posterior distributions (Gelman et al., 2013), but Kingston et al. (2005b) found success in using MCMC for Bayesian training of hydrological ANNs and we followed a modified version of that approach here. MCMC algorithms approximate posterior distributions by repeatedly constructing random proposals of the model parameters. At each iteration, specific rules are used to accept or discard the latest proposal such that the iterations eventually converge to the desired posterior distribution. The specific MCMC algorithm used here is known as the No U Turn Sampler (Hoffman and Gelman, 2014), which is an extension of the Hamiltonian Monte Carlo (HMC), also known as Hybrid Monte Carlo. HMC has been used in Bayesian ANN models previously (Neal, 1995). Kingston et al. (2005b) note the benefits of HMC but mention that it can be difficult to implement. Since the time of that publication, open source software has been developed to ease the implementation burden associated with HMC. We used the R interface for the C++ library Stan software to perform our MCMC analysis (Stan Development Team, 2016).

As in Kingston et al. (2005b), we used an independent and identically distributed $p(y|\theta, x)$ for each day (indexed by i) of data. In this work we use a student's t distribution for $p(y|\theta, x)$. The no tation $t_\nu(\mu, \sigma)$ represents the student's t distribution with degrees of freedom ν, mean μ, and scale parameter σ. We use ν 7 which results in higher probabilities for values far from the mean

Table 2
Evaluation of model fits to salinity observations at key monitoring stations.

Location[a]	Distance (km)[b]	Count	Approximate Period of Record	Mean Residual (mS/cm)		Coefficient of Determination (r^2)		Standard Error (mS/cm)	
				DSG	Hybrid ANN	DSG	Hybrid ANN	DSG	Hybrid ANN
Bay & Sacramento River Stations									
Martinez	52.6	1672	1946 to 1971	1.86	1.61	0.85	0.86	2.85	2.71
Martinez	54.0	5688	1995 to 2012	−0.34	−0.50	0.93	0.95	2.16	1.86
Bulls Head Point	54.7	1199	1926 to 1941	1.47	0.48	0.91	0.93	2.77	2.43
Martinez	55.0	9894	1965 to 1996	−0.22	−0.64	0.90	0.91	2.73	2.55
Port Chicago	64.0	14,205	1966 to 2012	0.51	0.07	0.90	0.92	2.25	2.06
Port Chicago	66.0	2219	1946 to 1971	−0.69	−1.24	0.93	0.94	1.78	1.62
O & A Ferry	74.8	2674	1921 to 1957	0.38	−0.45	0.96	0.97	1.31	1.25
Mallard Island	75.0	11,258	1976 to 2012	−0.02	−0.45	0.87	0.90	1.64	1.44
Collinsville	81.0	15,200	1966 to 2012	0.14	−0.18	0.83	0.86	1.22	1.10
Collinsville	81.8	3296	1921 to 1969	0.24	−0.34	0.94	0.95	1.15	1.12
Emmaton	92.0	16,294	1964 to 2012	0.06	−0.08	0.70	0.74	0.54	0.50
Emmaton	92.9	1937	1922 to 1971	0.21	−0.17	0.93	0.93	0.80	0.78
San Joaquin River Stations									
Pittsburg	77.0	15,810	1965 to 2012	0.13	0.20	0.86	0.89	1.52	1.35
Antioch	85.8	14,972	1967 to 2012	0.08	0.08	0.80	0.84	0.85	0.77
Antioch	88.4	4180	1921 to 1971	−0.16	−0.26	0.93	0.94	0.74	0.69
Jersey Point	95.8	15,977	1964 to 2012	−0.06	−0.09	0.65	0.68	0.38	0.37
Jersey Point	98.9	1179	1921 to 1971	−0.12	−0.20	0.90	0.91	0.51	0.48

[a] The same location may appear more than once when reported as different sampling locations. Salinity calculations are performed using the distance shown in the table; salinity calculations for the Bay stations utilize the Sacramento River reach DSG and ANN models. General locations of these stations are shown in Fig. 1.
[b] Measured from Golden Gate Bridge.

compared to the normal distribution and thus provides some degree of robustness against outliers (Gelman et al., 2013). The hybrid model (Eq. (5)) can be re written as $p(y|\theta, x)$ in Bayes' rule:

$$y(i) \sim t_7\left(\Phi_1 \cdot G(i)^{\Phi_2} + \eta(x(i), U, V), \sigma\right), \tag{7}$$

where the output (y) is X2 position, the input (x) represents the flow and water level input data, U denotes the matrix of input weights and biases, V denotes the matrix of hidden weights and biases, and G, Φ_1, Φ_2, and η are as previously defined. Our prior distribution, $p(\theta)$, should reflect the fact that the DSG model already fits the data quite well. Therefore, the prior distributions for the DSG parameters (Φ_1, Φ_2) are chosen to be normal distributions centered on values from previous work (Hutton et al., 2015) considering data from a similar time period:

$$\Phi_1 \sim N(4.25,\ 0.05),$$

$$\Phi_2 \sim N(\ \ .186, .005),$$

where $N(\mu, \sigma)$ indicates the normal distribution with mean μ and standard deviation σ.

As noted in the previous description of model architecture, ANN models have several weights and biases connecting the inputs to the hidden neurons and then to the output neuron. Depending on the ordering and sign of the hidden neurons, training may result in non unique U and V matrices for the same output y. To minimize this fitting challenge, we required V to be ordered (smallest to largest weights) and positive throughout the training, consistent with best practices in Bayesian inference (Gelman et al., 2013). We also restricted the row of U corresponding to the coastal water level input to be positive. This restriction, in combination with the constraints on V described above, enforced a positive response for X2 with increasing coastal water levels, thereby meeting our goal for ANN structural validity.

Network weights and biases were sampled and then re scaled to have near unit scale, following methods to improve the efficiency of the MCMC algorithm in terms of number of effective posterior samples per unit of computer time (Gelman et al., 2013). The DSG model is already known to predict X2 position reasonably well, and we expect the η term predicted by the ANN in Eq. (6) to be small; our prior distributions for the weights and biases of the ANN model reflect this information. For the weights and biases of all inputs except coastal water levels, we imposed a prior distribution with mean zero (where the tildes indicate this prior is for the values before re scaling):

$$\tilde{U} \sim N(0, 1),$$

$$\tilde{V} \sim N(0, 1).$$

The weights and biases that correspond to the coastal water level inputs instead have a half normal prior distribution (positive side of a standard normal distribution with mean 0, and standard deviation of 1).

3.5. Model outputs and prediction uncertainty

The MCMC algorithm outlined above produces 2000 samples of the hybrid ANN model parameters U, V, Φ_1, and Φ_2. Applying each of these 2000 samples to the input data yields a distribution of X2 values that reflect uncertainty in the model parameters. The time series of the posterior mean is obtained by averaging over the X2 distribution on each day. Uncertainty limits around the mean are computed by adding random draws corresponding to the 2000 MCMC samples of the residual variance, σ, and computing sample percentiles on each day. For example, the 2.5% and 97.5% percentiles determine a 95% prediction limit. This method of quantifying pre diction uncertainty has been used previously in other studies using Bayesian ANNs in water resources applications (Kingston et al., 2005a,b; Humphrey et al., 2016).

4. Results

ANN model performance was evaluated by comparing model

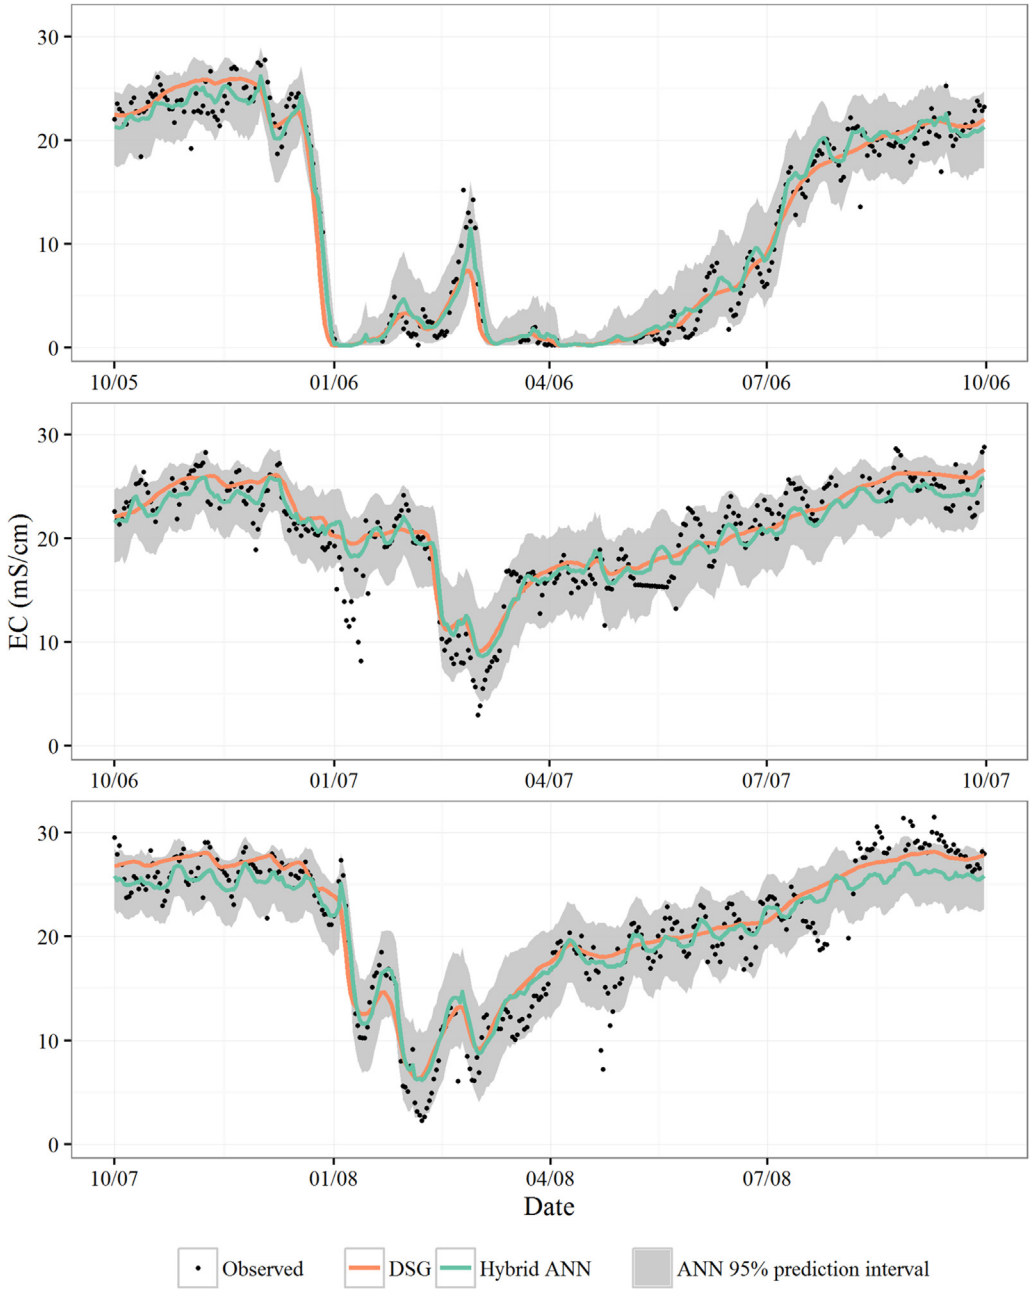

Fig. 5. A three-year (WY 2006 2008) time series of observed salinity values at Martinez is compared with predicted values from the DSG and hybrid ANN models. This station is located 54 km upstream of Golden Gate (see Fig. 1). The hybrid ANN model captures some of the spring-neap tidal variation of the salinity field.

predictions to observed data, including X2 isohaline data and salinity at discrete monitoring sites. Similarly, ANN model predictions were compared with DSG model predictions. Hybrid model response to freshwater flow and coastal water level input was explored through sensitivity analysis and, the latter through a case study application. In this work, sensitivity analysis refers to the process of identifying model response to a single response, with all other variables kept constant. Separate ANN models were developed for the Sacramento and San Joaquin River branches because channel characteristics lead to differences in salinity intrusion.

4.1. Hybrid ANN model

Hybrid ANN model estimates of daily X2 position are compared in Fig. 3 with daily DSG model estimates and interpolated values obtained from salinity observations over the full period of record. The hybrid model visually appears to capture the seasonal variability of X2, with some values falling outside the modeled range during periods of very high and low salinity (corresponding to low and high freshwater flows, respectively). Statistics comparing model performance are shown in Table 1. The hybrid model developed for the Sacramento River branch provides a superior fit to interpolated X2 data (relative to the DSG model) with a standard error of 3.22 km and r^2 0.95; the DSG model fits the interpolated data with a standard error of 3.63 km and r^2 0.93. The hybrid model appears to provide less biased predictions, with a mean residual of <0.01 km (0.003 km) compared to 0.76 km for the DSG model. For the San Joaquin River branch, the hybrid model provides

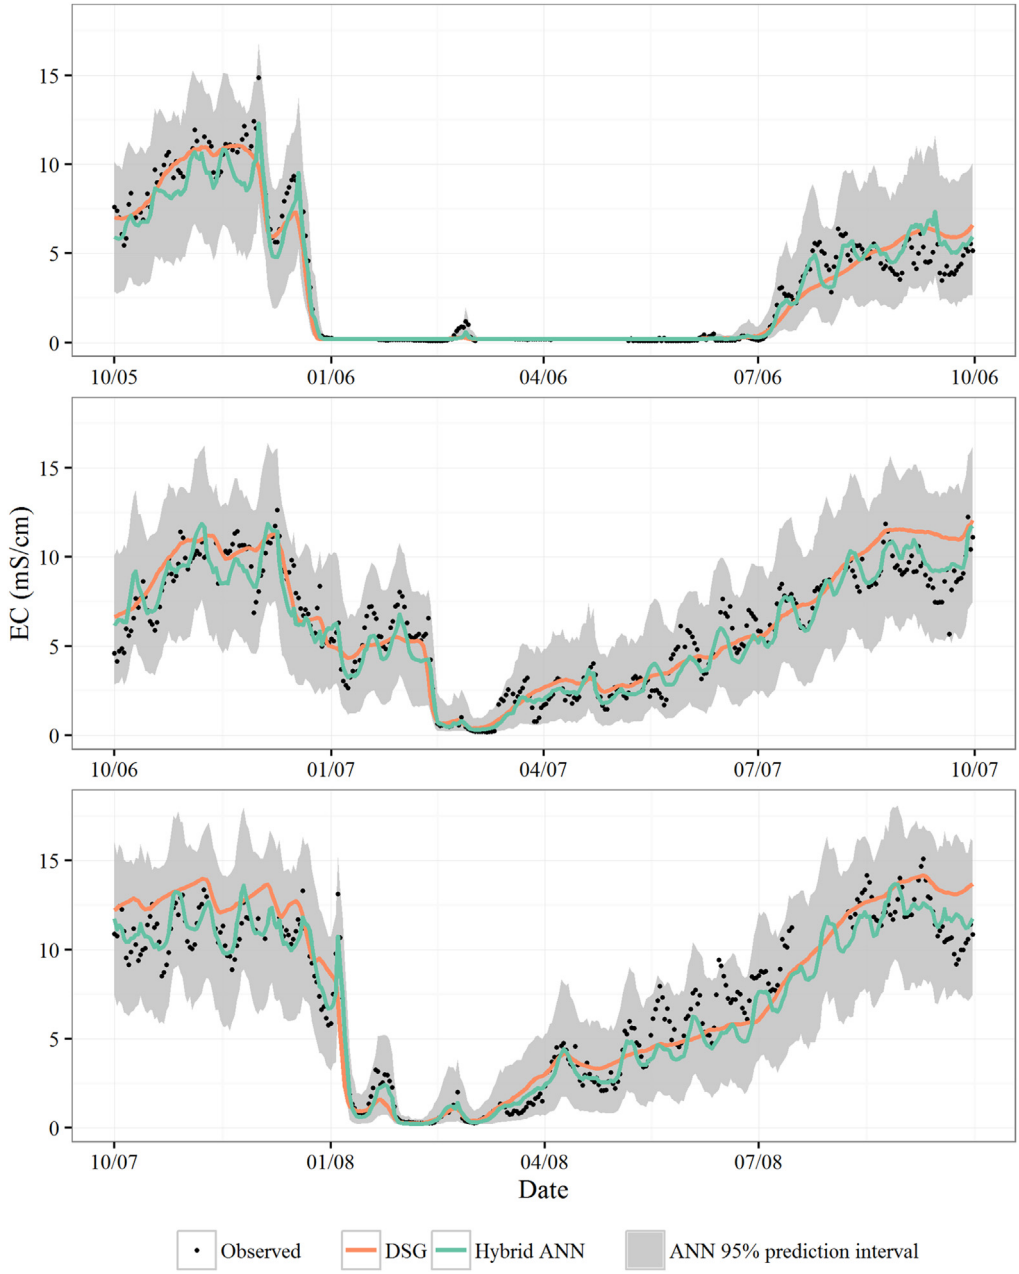

Fig. 6. A three-year (WY 2006 2008) time series of observed salinity values at Mallard Island is compared with predicted values from the DSG and hybrid ANN models. This station (X 75 km) is located further upstream of Golden Gate than Martinez and thus shows lower salinity levels. The hybrid ANN model captures some of the spring-neap tidal variation of the salinity field.

similar improvements relative to the DSG model (see Table 1).

To illustrate how simulated and observed salinity data respond to the estuary's tidal signal, Fig. 4 compares daily hybrid ANN model estimates with DSG model estimates and interpolated values over a three year subset of the data record (water years 2006–2008). The hybrid model appears to reasonably capture the spring neap tidal signal seen in the interpolated X2 data, thereby providing superior performance relative to the DSG model. This is a consequence of incorporating additional inputs (coastal water level and tidal range) that are not part of the original DSG model.

Table 2 provides a statistical comparison of model performance relative to observed salinity, measured as specific conductance, at several key monitoring sites in the estuary. The hybrid ANN models for the Sacramento and San Joaquin River branches consistently provide superior performance (relative to the DSG models) as measured by r^2 and standard error. The hybrid models typically show somewhat greater mean residuals in an absolute sense, suggesting greater (albeit small) prediction bias relative to the DSG models.

Fig. 5 through 7 compare hybrid model estimates of daily salinity, measured as specific conductance, with daily DSG model estimates and observations at three discrete locations in the estuary (Martinez, Mallard Island and Collinsville) for water years 2006–2008. These plots show that the general salinity behavior is reasonably represented through both modeling approaches, although the hybrid model captures the spring neap tidal variation

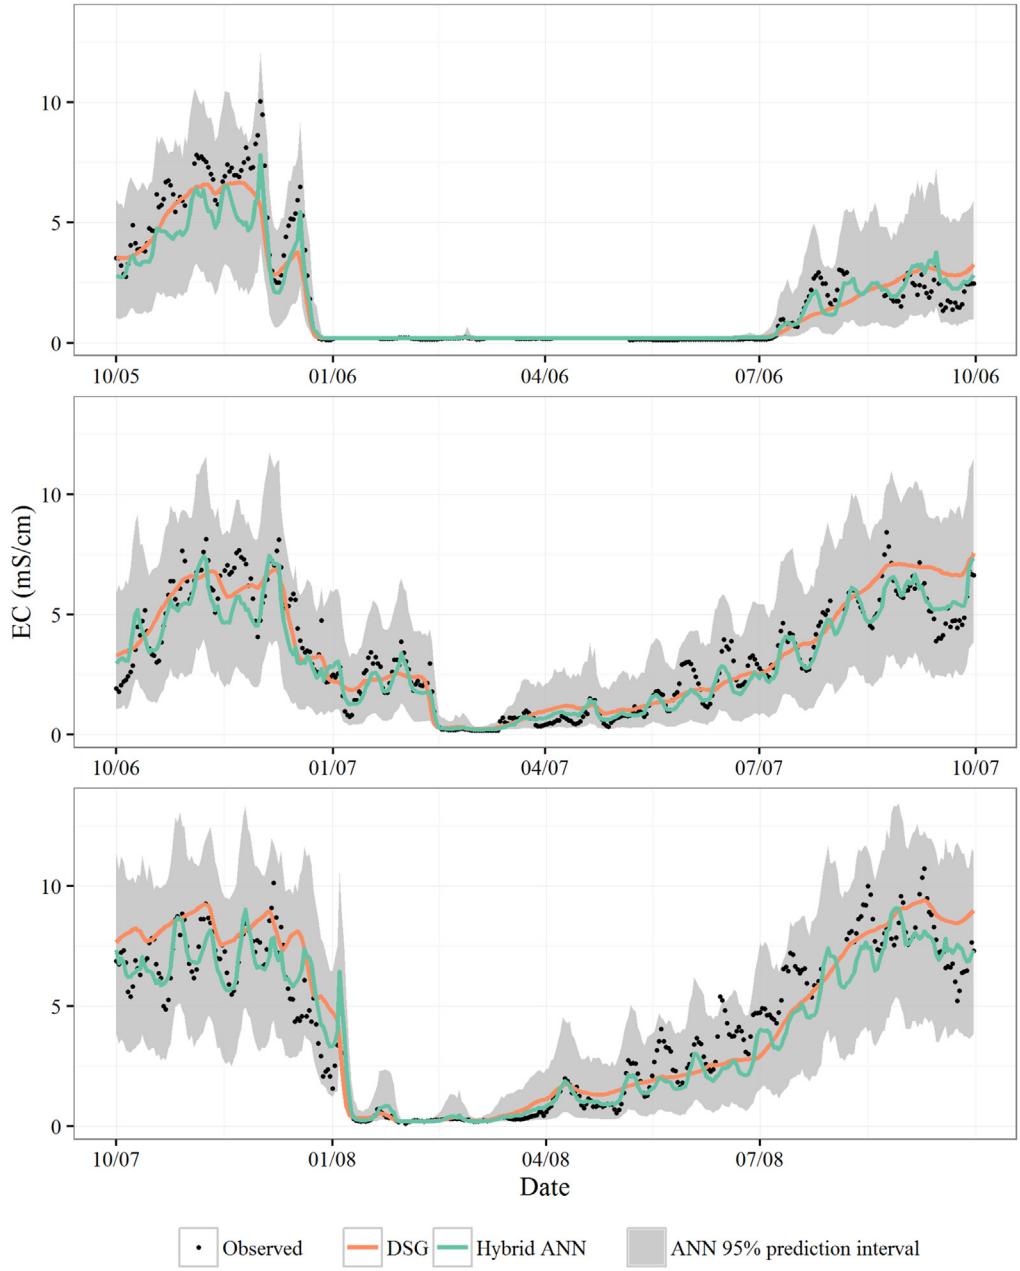

Fig. 7. A three-year (WY 2006 2008) time series of observed salinity values at Collinsville is compared with predicted values from the DSG and hybrid ANN models. This station (X 81 km) is located further upstream of Golden Gate than Martinez and Mallard Island and thus shows lower salinity levels. The hybrid ANN model captures some of the spring-neap tidal variation of the salinity field.

while the DSG model does not. The fortnightly variation in the hybrid model estimates typically follows the observed salinity signal; however, the magnitude of the modeled signal is sometimes under predicted compared to the observed data. As observed in the X2 comparison provided in Fig. 4, accounting for tidal response is a clear benefit of the use of the ANN modeling approach.

A consequence of the salinity gradient structure assumed within the hybrid ANN formulation is that it forces salinity values at up stream locations to converge to a constant freshwater value defined by S_b in Eq. (3). In reality, the relationship between flow and salinity becomes increasingly complex with distance upstream of the Sac ramento and San Joaquin River confluence. Thus, the assumed hybrid formulation does not fully capture the salinity variation at the upstream end of the gradient. This limitation was addressed through the development of targeted location specific data driven models; these models generally provide improved predictive vali dation at key upstream locations. Details are beyond the scope of this paper and are described in Chen et al. (2015).

4.2. Sensitivity analysis: Delta outflow and coastal water level

The sensitivity of the hybrid ANN model was explored by per turbing Delta outflow and coastal water level over the entire period of record (water years 1922–2012). Fig. 8a shows the response of X2 to a constant increase of antecedent Delta outflow by 5000 cfs; this outflow was previously defined by Eq. (2). The DSG predicted

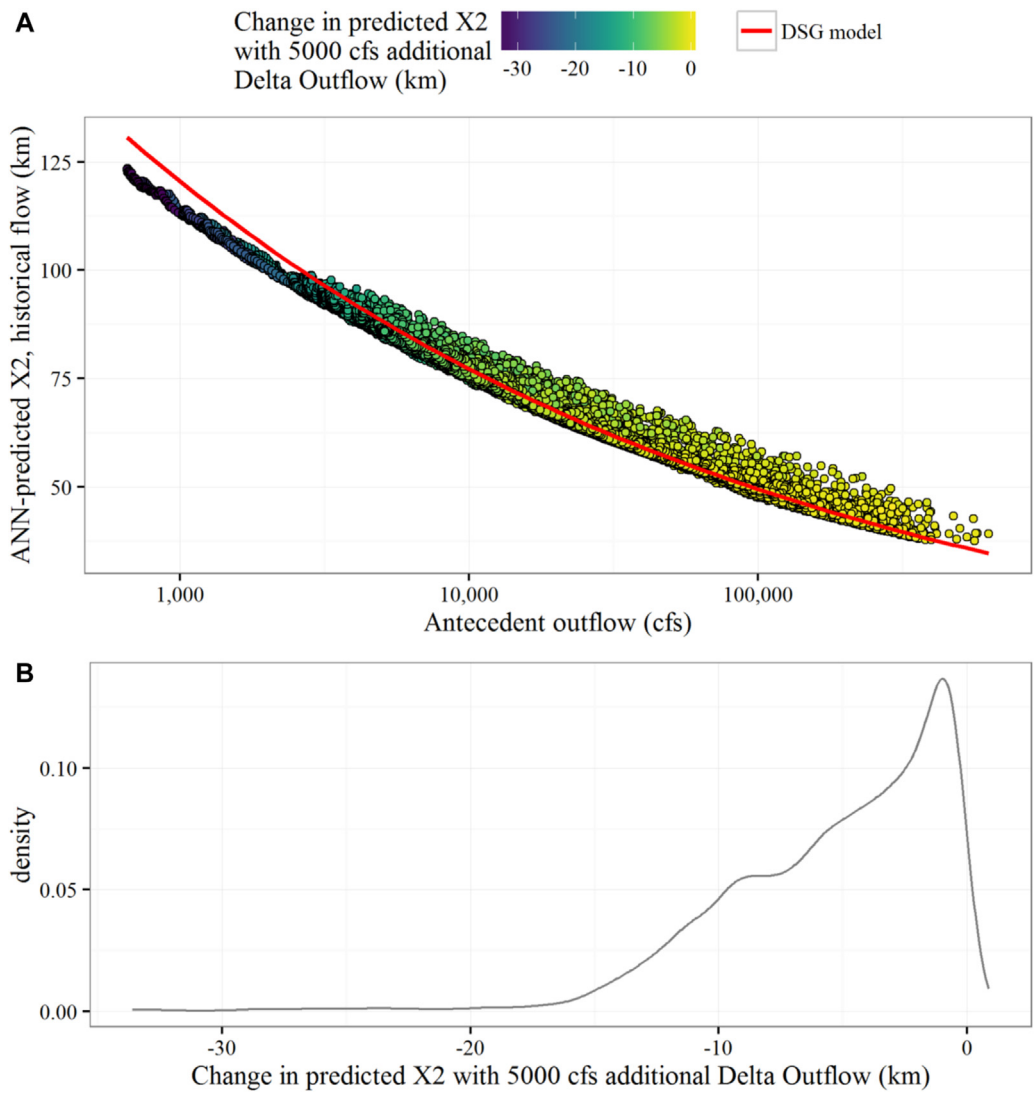

Fig. 8. The top panel (a) compares predicted X2 position as a function of antecedent outflow for the DSG and hybrid ANN models over the 91-year period of record (WY 1922 2012). The hybrid model X2 predictions are color-coded to reflect sensitivity to a constant 5000 cfs increase in Delta outflow. The bottom panel (b) shows the predicted X2 change as a probability density function over the same period of record. The X2 change for a given flow is much larger for lower flows and minimal for higher flows. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

X2 response to antecedent outflow is also provided in Fig. 8a for comparison. This exercise shows a decrease in X2 with larger de creases occurring at lower flows, and minimal change at high outflows, as indicated by the symbols in Fig. 8a and the distribution in Fig. 8b. This is consistent with our understanding of estuarine salinity, where a constant flow perturbation is much more conse quential at low flows. Fig. 9a shows model predicted X2 response to water level perturbation of 0.15 cm as a function of antecedent outflow. The hybrid model exhibits a variable response to the water level perturbation, with the greatest changes (up to approximately 3 km) observed over the intermediate flow range. Fig. 9b summa rizes the sensitivity analysis results as a frequency distribution. This sensitivity analysis demonstrates that the hybrid model response to a water level increase is always positive, following the constraints imposed as part of the hybrid ANN training process. Furthermore, this sensitivity analysis demonstrates that the water level increase is of the same magnitude as the standard error of the X2 prediction reported in Table 1. Hybrid model sensitivity to coastal water level is explored further in a case study as discussed below.

4.3. Case study: coastal water level fixed to 1920 level

A potential application of the hybrid ANN model is to isolate the historical effect of sea level rise on X2 position in the San Francisco Bay Delta estuary. A case study, summarized here, was defined to provide such an estimate. While coastal water level is subject to annual and decadal variation, a long term rising trend of 1.9 mm/yr has been documented in the estuary (Ryan and Noble, 2007). This case study does not consider future sea level rise, which may occur at a greater rate than observed to date. ANN models trained on historical data would likely not be valid for an assessment of future sea level rise, as such an analysis would entail significant extrap olation outside the model's training range. For this case study, we de trended the historical Golden Gate mean water level time series using linear regression to derive a 1920 level input time series and maintained other model inputs at baseline conditions (with a water level difference of +18.3 cm between the two scenarios in 2012). Fig. 10 summarizes the case study results by comparing hybrid model predictions of X2 assuming historical and 1920 level coastal

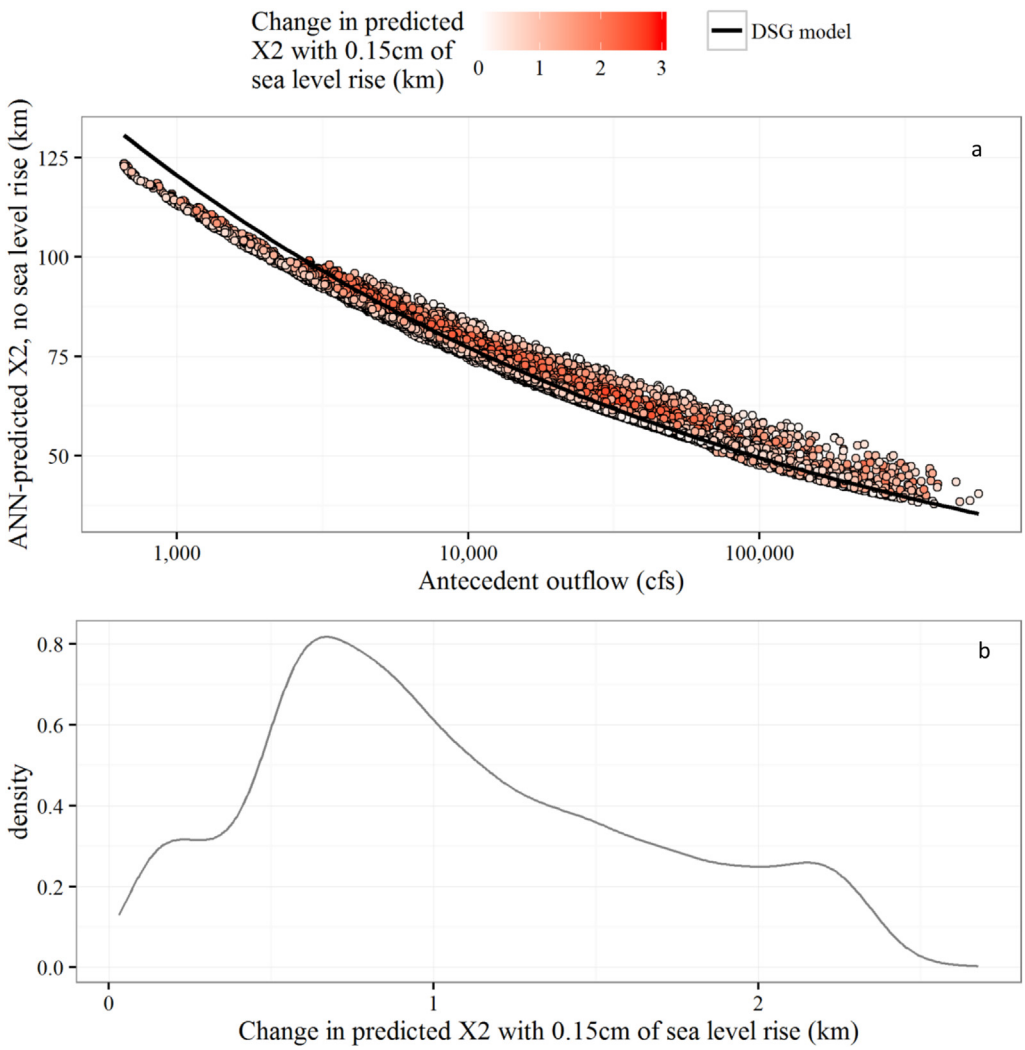

Fig. 9. The top panel (a) compares predicted X2 position as a function of antecedent outflow for the DSG and hybrid ANN models over the 91-year period of record (WY 1922 2012). The hybrid model X2 predictions are color-coded to reflect sensitivity to a constant 0.15 cm increase in coastal water level. The bottom panel (b) shows the predicted X2 change as a probability density function over the same period of record and demonstrates that the hybrid model response to a water level increase is always positive, following the constraints imposed as part of the training process. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

water levels for the hydrologic period spanning water years 2006–2008. The comparison is shown in absolute terms in Fig. 10a and in relative terms in Fig. 10b. This case study suggests that historical sea level rise has resulted in slightly greater salinity intrusion into the estuary, with an approximate upstream shift in X2 of 1 km relative to 1920 conditions.

5. Discussion

This work explored the benefit of a hybrid empirical ANN modeling approach to predict salinity in the San Francisco Bay Delta estuary. The existing empirical model (Hutton et al., 2015) is based on freshwater outflow as the only input, and while it performs well at predicting salinity in the estuary, the goal of the present work was to further improve performance by inclusion of additional inputs. The ANN models were thus used in conjunction with the empirical model, with coastal water level and tidal range also incorporated as model inputs. The ANN models were trained with an extensive data set spanning nine decades including observed salinity, flow, and coastal water level data. Over the last nine decades, the estuary has experienced significant hydrologic variability in conjunction with changes in land use, water management facilities and regulations, and coastal water level; these perturbations are implicitly represented in the flow, salinity and water level data and are therefore embodied in the trained ANN models. In addition to the focus on sensitivity analysis and prediction uncertainty, use of the long term observed dataset differentiates our work from previous efforts to model salinity in the estuary with ANNs, as these efforts utilized synthetic data generated from a water quality model of the Delta (Wilbur and Munevar, 2001; Mierzwa, 2002; Senevirante et al., 2008) or utilized observed data from a limited period of record (Finch and Sandhu, 1995; Rajkumar and Johnson, 2001).

The hybrid ANN approach provides improvements over the empirical model (Hutton et al., 2015), including the ability to capture the effect of the spring neap tidal cycle and historical sea level rise on X2 position and salinity at site specific locations. By incorporating a Bayesian approach, the hybrid model also provides a measure of prediction uncertainty in the calculations. A case study was presented to evaluate how salinity in the estuary has changed

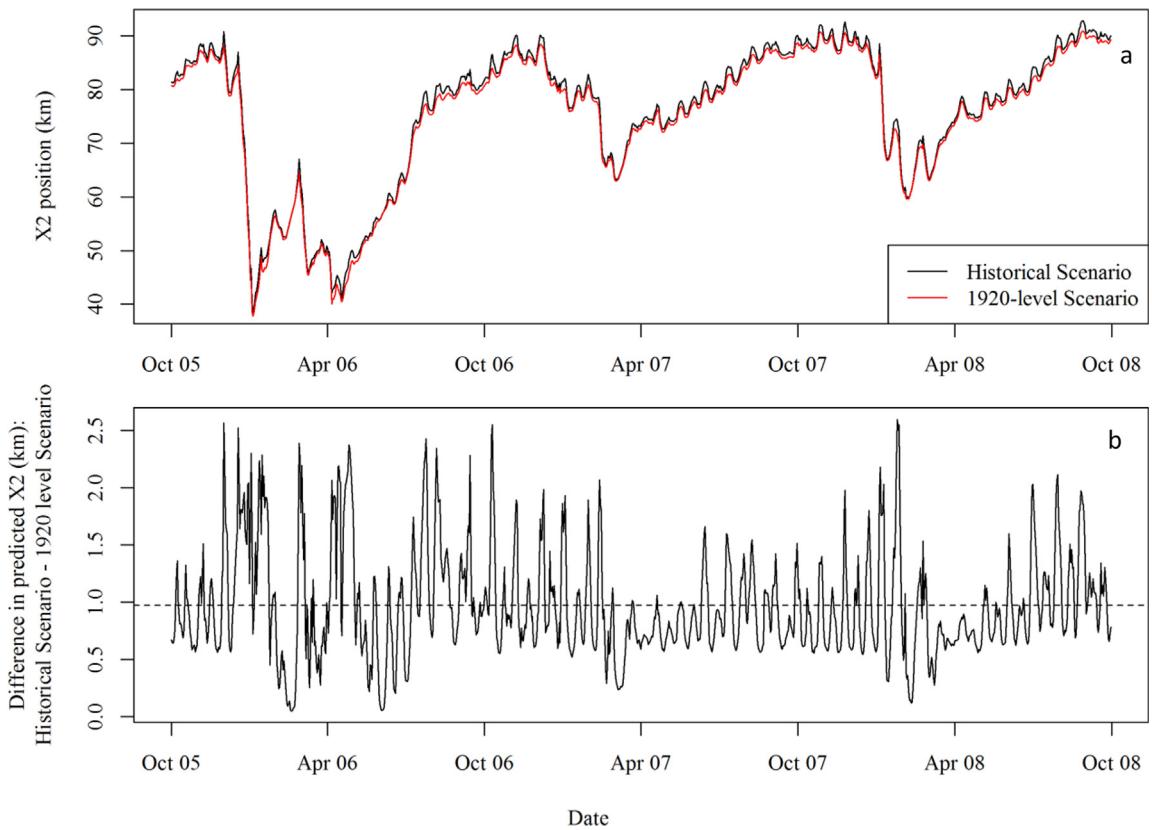

Fig. 10. The top panel (a) compares hybrid ANN X2 predictions under historical and 1920-level scenarios for a three-year (WY 2006-08) time series. The bottom panel (b) shows the difference of the two scenarios; the dashed line represents the average difference over the period shown. The assumed coastal water level differential between the historical and 1920-level scenarios is +18.3 cm.

as a result of sea level rise over the last nine decades; this study suggests that historical sea level rise has resulted in greater salinity intrusion into the estuary, with an approximate upstream shift in X2 of 1 km relative to 1920 conditions. Given the great economic value associated with efficient water management in the San Francisco Bay Delta estuary, the improvement in model performance demonstrated in this work (as measured by predictive and structural validation goals) and the incorporation of uncertainty justifies further consideration of hybrid ANN modeling approaches in future planning and operational efforts.

Although the hybrid approach provided improved predictive model performance relative to the empirical DSG model, unexplained residual error remains. For example, as reported in Table 1, the standard error of X2 estimates for the Sacramento River branch was 3.63 km for the DSG model and 3.22 km for the hybrid model. Potential causes of the residual error may include noise associated with measurement error, error associated with interpolation of X2 position between fixed monitoring locations, historical changes in the estuary's bathymetry, and complex hydrodynamics (particularly at the high and low flow extremes) that are insufficiently represented in the model formulation and training.

Potential model improvements were explored as part of this work (Chen et al., 2015; Chen and Roy, 2013) but were not reported here for brevity. Follow up work in some of these areas is recommended for applications specific to the San Francisco Bay Delta estuary. One area for potential improvement relates to the representation of hydrodynamic complexities in the upstream reaches of the estuary. The hybrid ANN model presented in this paper, typical of empirical approaches in the estuary, assumes that a single lumped measure of freshwater flow (i.e., Delta outflow) sufficiently captures the response of salinity to flow changes. We found that targeted data driven ANN models provided superior predictive validity at key upstream locations; however, this work (which was beyond the scope of this paper) was not critiqued relative to structural validity. Given the importance of complying with salinity regulations in the upstream reaches of the estuary (CSWRCB, 2006), we recommend that future efforts explore improved methods of representing salinity intrusion in this locale. A second area for potential improvement relates to the representation of tidal conditions in the estuary. Utilizing astronomical predictions of coastal water levels (rather than observed water levels) as a model input may offer significant advantages when using the tool in a forecasting application. We recommend that future efforts explore alternate model inputs, including astronomical tidal predictions along with measures of atmospheric conditions (e.g. barometric pressure and wind) that impact water levels.

This work contributes to the growing body of literature on the application of artificial neural networks for modeling dynamic environmental systems. More specifically, it illustrates the use of ANNs for a complex, data rich, real world problem, where application of a simple black box model would likely be rejected by the stakeholder community without sufficient understanding of model sensitivity. Furthermore, as with many important problems, a pre existing empirical or mechanistic model may be available and provide a reasonable representation of the underlying system. In such cases, a hybrid approach may be adopted to build upon the foundation of the pre existing model, with ANN features being used to enhance certain elements of performance that are of

interest. Such an approach does not require "reinventing the wheel" and provides a minimum level of performance consistent with the pre existing model. This outcome cannot easily be guaranteed for a fully independent ANN model, and the hybrid approach may provide a more resource efficient path for model development. The present application also incorporated recent advances in Bayesian ANNs to better focus the application. Bayesian priors were used to define constraints in model response based on *a priori* understanding of the system and to ensure that certain sensitivity responses are always hydrodynamically plausible. The Bayesian approach also had the advantage of allowing quantification of prediction uncertainty that was reported with the model outputs. Taken together, the hybrid Bayesian ANN approach may provide a template for model development where the following criteria are met: existence of a credible model; need for improving the existing model without modifying its basic structure; prior knowledge of certain model behavior that can be encapsulated in priors associated with model weights and biases; and the need for quantification of prediction uncertainty.

Acknowledgements

This work was made possible through funding from the Metropolitan Water District of Southern California (Agreement number143878).

Appendix A. Supplementary data

Supplementary data related to this article can be found at http://dx.doi.org/10.1016/j.envsoft.2017.03.022.

References

Alber, M., 2002. A conceptual model of estuarine freshwater inflow management. Estuaries 25 (6), 1246 1261.
Bowden, G.J., Maier, H.R., Dandy, G.C., 2002. Optimal division of data for neural network models in water resources applications. Water Resour. Res. 38 (2), 2 1.
Bowden, G.J., Maier, H.R., Dandy, G.C., 2005. Input determination for neural network models in water resources applications. Part 2. Case study: forecasting salinity in a river. J. Hydrol. 301 (1), 93 107.
California Department of Water Resources (CDWR), 2013. DSM2 Documentation. On the Internet at: http://baydeltaoffice.water.ca.gov/modeling/deltamodeling/models/dsm2/toc.html (Accessed 25 March 2016).
California Department of Water Resources (CDWR) (2015). DAYFLOW model data available online at: http://www.water.ca.gov/dayflow/, (Accessed 24 March 2016).
California State Water Resources Control Board (CSWRCB), 2006. Water Quality Control Plan for the San Francisco Bay/Sacramento-San Joaquin Delta Estuary. Division of Water Rights. On the Internet at: http://www.waterboards.ca.gov/waterrights/water_issues/programs/bay_delta/wq_control_plans/2006wqcp/ (Accessed 24 March 2016).
Chen, L., Rath, J.S., Roy, S.B., 2015. Modeling Salinity in Suisun Bay and the Western Delta Using Artificial Neural Networks, Report Prepared for Metropolitan Water District of Southern California, April.
Chen, L., Roy, S.B., 2013. Modeling Salinity in Suisun Bay and the Western Delta Using Artificial Neural Networks, Report Prepared for Metropolitan Water District of Southern California, April.
Cheng, R.T., Casulli, V., Gartner, J.W., 1993. Tidal, residual, intertidal mudflat (TRIM) model and its applications to San Francisco Bay, California. Estuar. Coast. Shelf Sci. 36 (3), 235 280.
Chua, V.P., Fringer, O.B., 2011. Sensitivity analysis of three-dimensional salinity simulations in North San Francisco Bay using the unstructured-grid SUNTANS model. Ocean. Model. 39 (3), 332 350.
Chung, F., Seneviratne, S., 2009. Developing artificial neural networks to represent salinity intrusions in the Delta. World Environ. Water Resour. Congr. http://dx.doi.org/10.1061/41036(342)483.
Conrads, P.A., Roehl, E.A., Daamen, R.C., Kitchens, W.M., 2006. Simulation of Water Levels and Salinity in the Rivers and Tidal Marshes in the Vicinity of the Savannah National Wildlife Refuge, Coastal South Carolina and Georgia. U. S. Geological Survey. On the Internet at: http://pubs.usgs.gov/sir/2006/5187/ (Accessed 29 March 2016).
Dat, T.Q., Likitdecharote, K., Srisatit, T., Trung, N.H., 2011. Modeling the Influence of river discharge and sea level rise on salinity intrusion in Mekong Delta. Environ. Support. Food Energy Secur. Crisis Oppor. 35, 685 701.
Delta Plan (2013). Published by the Delta Stewardship Council, available online at: http://deltacouncil.ca.gov/delta-plan-0, (Accessed 29 March 2016.
Denton, R.A., 1993. Accounting for antecedent conditions in seawater intrusion modeling applications for the San Francisco Bay-Delta. Hydraulic Engineering 93(1): 448 453,. In: Proceedings of ASCE National Conference on Hydraulic Engineering (San Francisco).
Draper, A.J., Munévar, A., Arora, S.K., Reyes, E., Parker, N.L., Chung, F.I., Peterson, L.E., 2004. CalSim: generalized model for reservoir system analysis. J. Water Resour. Plan. Manag. 130 (6), 480 489.
Fernández-Delgado, C., Baldó, F., Vilas, C., García-González, D., Cuesta, J.A., González-Ortegón, E., Drake, P., 2007. Effects of the river discharge management on the nursery function of the Guadalquivir river estuary (SW Spain). Hydrobiologia 587 (1), 125 136.
Finch, R., Sandhu, N., 1995. Artificial Neural Networks with Application to the Sacramento San Joaquin Delta. California Department of Water Resources Delta Modeling Section, Division of Planning.
Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. Bayesian Data Analysis. CRC Press.
Gross, E.S., Koseff, J.R., Monismith, S.G., 1999. Three-dimensional salinity simulations of South San Francisco Bay. J. Hydraul. Eng. 125 (11), 1199 1209.
Gross, E.S., Nidzieko, N.J., MacWilliams, M.L., Stacey, M.T., 2007. Parameterization of estuarine mixing processes in the San Francisco Estuary based on analysis of three-dimensional hydrodynamic simulations. J. Estuar. Coast. Model. ASCE 322 338.
Harder, J.A., 1977. Predicting estuarine salinity from river inflows. Journal of the Hydraulics Division, ASCE 103, 877 888.
Hoffman, M.D., Gelman, A., 2014. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. J. Mach. Learn. Res. 15 (Apr), 1593 1623.
Huang, W., Foo, S., 2002. Neural network modeling of salinity variation in Apalachicola River. Water Res. 36 (1), 356 362.
Humphrey, G.B., Gibbs, M.S., Dandy, G.C., Maier, H.R., 2016. A hybrid approach to monthly streamflow forecasting: integrating hydrological model outputs into a Bayesian artificial neural network. J. Hydrol. 540, 623 640.
Hutton, P.H., Rath, J.S., Chen, L., Ungs, M.J., Roy, S.B., 2015. Nine decades of salinity observations in the San Francisco Bay and Delta: modeling and trend evaluation. ASCE J. Water Resour. Plan. Manag. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000617,04015069.
Jain, A., Sudheer, K.P., Srinivasulu, S., 2004. Identification of physical processes inherent in artificial neural network rainfall runoff models. Hydrol. Process. 18 (3), 571e581.
Jain, A., Kumar, S., 2009. Dissection of trained neural network hydrologic models for knowledge extraction. Water Resour. Res. 45 (7).
Jassby, A.D., Kimmerer, W.J., Monismith, S.G., Armor, C., Cloern, J.E., Powell, T.M., Schubel, J.R., Vendlinski, T.J., 1995. Isohaline position as a habitat indicator for estuarine populations. Ecol. Appl. 5, 272 289.
Kingston, G.B., Maier, H.R., Lambert, M.F., 2005a. Calibration and validation of neural networks to ensure physically plausible hydrological modeling. J. Hydrol. 314 (1e4), 158e176.
Kingston, G.B., Lambert, M.F., Maier, H.R., 2005b. Bayesian training of artificial neural networks used for water resources modeling. Water Resour. Res. 41 (12).
Luoma, S.N., Dahm, C.N., Healey, M., Moore, J.N., 2015. Challenges facing the Sacramento San Joaquin Delta: complex, chaotic, or simply cantankerous? San Franc. Estuary Watershed Sci. 13 (3).
MacWilliams, M.L., Bever, A.J., Gross, E.S., Ketefian, G.S., Kimmerer, W.J., Center, R.T., 2015. Three-dimensional modeling of hydrodynamics and salinity in the San Francisco Estuary: an evaluation of model accuracy, X2, and the low salinity zone. San Franc. Estuary Watershed Sci. 13 (1).
Maier, H.R., Dandy, G.C., 1996. The use of artificial neural networks for the prediction of water quality parameters. Water Resour. Res. 32 (4), 1013 1022.
Maier, H.R., Jain, A., Dandy, G.C., Sudheer, K.P., 2010. Methods used for the development of neural networks for the prediction of water resource variables in river systems: current status and future directions. Environ. Model. Softw. 25, 891 909.
Mierzwa, M., 2002. Chapter 4: CALSIM versus DSM2 ANN and G-model Comparisons. Methodology for Flow and Salinity Estimates in the Sacramento San Joaquin Delta and Suisun Marsh, 23rd Annual Progress Report. On the Internet at: http://baydeltaoffice.water.ca.gov/modeling/deltamodeling/delta/reports/annrpt/2002/2002Ch4.pdf (Accessed 24 March 2016).
Monismith, S.G., Kimmerer, W., Burau, J.R., Stacey, M.T., 2002. Structure and flow-induced variability of the subtidal salinity field in northern San Francisco Bay. J. Phys. Oceanogr. 32 (11), 3003 3019.
Murray Darling Basin Ministerial Council, 1999. The Salinity Audit of the Murray Darling Basin: a 100 Year Perspective. On the Internet at: http://www.mdba.gov.au/sites/default/files/archived/mdbc-salinity-reports/439_Salinity_audit_of_MDB_100_year_perspective.pdf (Accessed June 14, 2015).
National Oceanic and Atmospheric Administration (NOAA) (2015). San Francisco water height data available online at: https://tidesandcurrents.noaa.gov/stationhome.html?id 9414290, (Accessed 14 June 2015).
Neal, R.M., 1995. Bayesian Learning for Neural Networks (PhD thesis). Department of Computer Science, University of Toronto.
Rajkumar, T., Johnson, M.L., 2001. Prediction of salinity in San Francisco Bay Delta using neural network. In: In Systems, Man, and Cybernetics, vol. 1. IEEE, pp. 329 334, 2001 IEEE International Conference on.
Reinert, T.R., Peterson, J.T., 2008. Modeling the effects of potential salinity shifts on the recovery of striped bass in the Savannah River Estuary, Georgia South Carolina, United States. Environ. Manag. 41 (5), 753 765.

Ryan, H.F., Noble, M.A., 2007. Sea level fluctuations in central California at subtidal to decadal and longer time scales with implications for San Francisco Bay, California. Estuar. Coast. Shelf Sci. 73 (3), 538 550.

Schoellhamer, D.H., 2000. Influence of salinity, bottom topography, and tides on locations of estuarine turbidity maxima in Northern San Francisco Bay. Proc. Mar. Sci. 3, 343 357.

Senevirante, S., Wu, S., Liang, Y., 2008. Chapter 3: Impacts of Sea Level Rise and Amplitude Change on Delta Operations. Methodology for flow and salinity estimates in the Sacramento San Joaquin Delta and Suisun Marsh. 29th Annual Progress Report. June.

Sklar, F.H., Browder, J.A., 1998. Coastal environmental impacts brought about by alterations to freshwater flow in the Gulf of Mexico. Environ. Manag. 22 (4), 547 562.

Stan Development Team, 2016. RStan: the R Interface to Stan. Version 2.9.0. Available online at: http://mc-stan.org/.

Wilbur, R., Munevar, A., 2001. Chapter 7: Integration of CALSIM and Artificial Neural Networks Models for Sacramento San Joaquin Delta Flow-salinity Relationships. Methodology for flow and salinity estimates in the Sacramento San Joaquin Delta and Suisun Marsh. 22nd Annual Progress Report. On the Internet at: http://baydeltaoffice.water.ca.gov/modeling/deltamodeling/delta/reports/annrpt/2001/2001Ch7.pdf (Accessed March 25, 2016).

Wu, W., Dandy, G.C., Maier, H.R., 2014. Protocol for developing ANN models and its application to the assessment of the quality of the ANN model development process in drinking water quality modelling. Environ. Model. Softw. 54, 108 127.