

Review

Machine Learning for Water and Environmental Modeling in the Sacramento–San Joaquin Delta of California: A 30-Year Retrospective and Future Prospects

Minxue He * and Prabhjot Sandhu

California Department of Water Resources, Sacramento, CA 95814, USA

* Correspondence: kevin.he@water.ca.gov

Abstract

The Sacramento–San Joaquin Delta (Delta) is a critical hub for California’s water supply and ecological health, necessitating rapid and accurate modeling for decision-support. This paper reviews three decades of machine learning (ML) applications in the Delta, tracing the evolution from early artificial neural networks (ANNs) developed in the 1990s for salinity estimation to modern integration within operational planning tools like CalSim. Recent advancements highlight a shift toward deep learning and hybrid physics-informed models, expanding applications to water quality forecasting, ecosystem monitoring, toxigenic and disaster risk assessment, and infrastructure resilience. Despite these successes, widespread operational deployment remains constrained by critical gaps in model interoperability, generalizability, interpretability, and operationalization. To address these challenges, we propose a strategic roadmap emphasizing standardized ML development protocols, multi-modal data fusion, digital twins, generative artificial intelligence (AI), explainable AI, and AI agents. This roadmap aims to transform isolated research ML models into a cohesive, transparent, and adaptive decision-support framework for sustainable water and environmental management in the Delta.

Keywords: Sacramento–San Joaquin Delta; machine learning; water and environmental modeling; decision support; artificial intelligence

1. Introduction

1.1. Background

California, the world’s fourth-largest economy [1,2], is home to nearly 40 million people and supports one of the most productive agricultural regions in the world [3,4]. Sustaining this population and economy depends critically on reliable water supply. However, California’s hydrology is fundamentally mismatched with its human water demand. Most precipitation falls in northern mountain ranges during winter, while peak demand occurs in the southern part of the state during the dry summer and fall [5]. To bridge this spatial and temporal gap, California has developed an extensive water storage and conveyance system anchored by the State Water Project (SWP) managed by the state and the federally managed Central Valley Project (CVP). Collectively, these two projects comprise over 1000 km of canals, aqueducts, tunnels, and pipelines, as well as major dams and reservoirs, pumping plants, and hydropower generation facilities [6,7].

Academic Editor: Jianjun Ni

Received: 10 June 2026

Revised: 17 July 2026

Accepted: 20 July 2026

Published: 21 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

The fulcrum of this complex system is the Sacramento–San Joaquin Delta (hereafter “Delta”), a mosaic of islands intersected by approximately 1100 km of tidal channels (Figure 1). The Delta receives freshwater inflows from the Sacramento and San Joaquin Rivers, which are diverted for municipal, agricultural, and environmental uses or managed to repel seawater intrusion from the San Francisco Bay. Ecologically, the Delta is a global biodiversity hotspot, supporting over 750 plant and animal species, including endangered species such as Delta Smelt [8,9]. Socioeconomically, its watershed supplies water to nearly two-thirds of California’s population and over 15,000 km² of farmland [10]. From a regulatory perspective, state and federal regulations govern flows, water quality, and biological protections to balance water supply reliability with ecological sustainability [11–13]. Consequently, operations of the SWP and CVP are shaped by these tightly coupled physical, ecological, and socioeconomic, and regulatory constraints.




Figure 1. Location map of the San Francisco Bay and the Sacramento–San Joaquin Delta (Delta; outlined by a dashed polygon) in California. The intake locations of the State Water Project (SWP) and Central Valley Project (CVP) are also shown.

Despite its central importance, the Delta faces a suite of interconnected challenges. Water supply reliability is increasingly constrained by climate change-induced droughts, upstream water diversions, and aging infrastructure [14,15]. Water quality degradation from agricultural runoff, urban pollutants, and industrial discharges introduces excessive nutrients, pesticides, and other contaminants, impairing habitat quality and fostering harmful algal blooms [16–20]. Concurrently, habitat loss, invasive species, and altered flow regimes have collectively contributed to declining ecosystem health and biodiversity in the Delta [21–23]. Sea level rise-driven salinity intrusion further complicates these challenges, affecting agricultural productivity, water quality, and ecosystem health, particularly in low-lying areas vulnerable to inundation [24]. Compounding these pressures, land subsidence driven by microbial oxidation and compaction of organic soils threatens the integrity of the Delta’s levee network [25].

Addressing these coupled challenges requires accurate, timely, and computationally efficient decision support tools, motivating the development and continuous evolution of numerical modeling frameworks for the Delta.

1.2. Modeling for Decision-Support in the Delta

Computer-based numerical modeling, together with monitoring data, forms the backbone of decision-support for Delta water and environmental management by addressing monitoring sparsity through interpolation and enabling scenario analysis, forecasting, and evaluation of trade-offs under changing hydrologic, climatic, structural, operational, and regulatory conditions.

Specifically, management of water and environmental resources in the Delta involves a complex network of federal, state, and local agencies, as well as academic institutions and private-sector partners, each playing distinct but interconnected roles in model development, deployment, and use. The California Department of Water Resources (DWR) operates the SWP, while the U.S. Bureau of Reclamation manages the CVP, and both rely heavily on modeling tools to support planning and real-time operations. Other federal agencies, including the U.S. Geological Survey (USGS), U.S. Fish and Wildlife Service (USFWS), U.S. Army Corps of Engineers (USACE), and National Marine Fisheries Service (NMFS), contribute to model development, scientific evaluation, and compliance with flow and environmental regulations. In parallel, state agencies such as the California Department of Fish and Wildlife (CDFW), the State Water Resources Control Board (SWRCB), and the Delta Stewardship Council develop and apply models to inform decision-making on, for instance, permitting, water quality standards, and ecosystem management decisions. Academic researchers and consulting firms further support model advancement, calibration, and application. This distributed ecosystem of actors underscores the importance of cross-agency collaboration, as models are often co-developed, shared, and iteratively refined to ensure consistency, transparency, and credibility across planning, regulatory, and operational contexts.

Early modeling efforts relied on empirical and mass-balance approaches (defined in Table A1 of Appendix A). For example, the Bulletin 27 Salinity Intrusion Model [26] estimated freshwater flow requirements to repel seawater intrusion into the Delta, while DAYFLOW [27] applied a mass-balance framework to quantify daily Delta inflows, exports, and outflows (Figure 2). Although these early modeling tools lacked detailed spatial resolution and process representation, they remain valuable for their transparency and computational efficiency.

Building on this foundation, more sophisticated operations and planning models were developed to simulate the complex interactions of the SWP and CVP. DWRSIM [28,29] was developed to simulate coordinated SWP and CVP operations. It was later replaced by CalSim [30], CalSim II [31,32], and more recently CalSim 3.0 [33] which serves as the primary operations model for statewide water planning (Figure 2). While these models do not explicitly resolve detailed Delta hydrodynamics and water quality processes, they are indispensable for long-term water supply planning in a changing climate.

Process-based models were also developed to simulate Delta hydrodynamics and water quality more explicitly (defined in Table A1 of Appendix A). The Bay-Delta Salinity model (TVRK) was developed and became operational in 1979 [34], integrating many of the desirable capabilities of the existing empirical flow-salinity models used across different Delta sub-regions while addressing several of their limitations. This effort was followed by the development of the Fischer Delta Model (FDM), delivered to DWR in 1984. FDM consists of two primary modules: DELFLO, which computes Delta hydrodynamics, and DELSAL, which simulates the salinity mixing using the Lagrangian method [35]. FDM later evolved into the DWR Delta Simulation Model (DWRDSM, also referred to as

DSM1), incorporating enhanced representation of Delta geometry, agricultural divisions, and gate operations [36] (Figure 2).

DWRDSM was subsequently replaced by the Delta Simulation Model 2 (DSM2) [37], a one-dimensional (1-D) network model capable of simulating flow, stage, velocity, water temperature, salinity, and constituent transport throughout the Delta (Figure 2). Unlike DWRSIM, DSM2 integrates the Four-Point Model [38–40] for hydrodynamic simulation and the Branched Lagrangian Transport Model (BLTM) [41,42] for water quality constituent transport. Since its release, DSM2 has served as a primary tool for real-time Delta operations and planning assessments, offering a balance between process representation and computational efficiency. More recently, three-dimensional (3-D) models such as SCHISM [43–45] have been applied in the Delta to resolve 3-D processes, including tidal mixing, salinity stratification, and localized ecological dynamics (Figure 2). These models provide higher physical fidelity and are well suited for evaluating climate change, sea level rise, and ecosystem restoration scenarios, though they require greater data and computational resources.




Figure 2. Timeline of selected operational models used to support water and environmental decision-making in the Delta.

Despite the high physical fidelity of multi-dimensional models, their integration into routine operations and long-term assessments is frequently hindered by prohibitive computational overhead. While 3-D models like SCHISM provide unparalleled detail, simulating a single water year can require 2–4 days of time on high-performance computing clusters [43–45], depending on model configuration, computational resources, and the number of cores allocated per simulation. This computational bottleneck creates a significant gap in decision-support; it precludes real-time forecasting during emergency events (e.g., levee failures) and limits the ability of planners to run the massive ensembles required for robust climate adaptation and uncertainty analysis.

Furthermore, the high-dimensional output of these models often requires manual, expert-intensive post-processing, introducing operational latency that is incompatible with the “seconds-to-minutes” requirement of real-time decision-support systems. Machine learning (ML) addresses these limitations by serving as a high-speed surrogate model. Once trained on the outputs of process-based models or historical data, ML architectures can perform rapid inference, often within milliseconds to seconds depending on the model architecture and hardware configuration. This computational efficiency provides an opportunity to emulate complex Delta system responses through data-driven surrogate frameworks, enabling rapid scenario screening and supporting real-time or near-real-time decision-support applications. However, unlike the 1-D DSM2 and 3-D SCHISM models, which are process-based, ML models are data-driven in nature. Therefore, caution should be exercised when interpreting ML model results; this will be discussed later in the paper. Figure 3 presents a comparison of the ML model with the 1-D and 3-D process-based models. Additional details are provided in Table A1 of Appendix A.




Figure 3. Comparison of an example high-fidelity three-dimensional process-based model, an example one-dimensional process-based model, and a typical machine learning model. Approach denotes the methodology used to represent system behavior; structure refers to the internal organization of equations or learned relationships; runtime is the computational time required to generate outputs from a calibrated or trained model; and interpretability indicates how readily model behavior and input–output relationships can be understood.

Over the past three decades, ML adoption in Delta water and environment modeling has evolved from early artificial neural networks (ANNs) to modern deep learning architectures. Research and applications now span a vast array of ecological, physical, and structural disciplines across local, state, and federal agencies, academia, and the private sector. These studies are listed in Table A2 of Appendix A. The table summarizes the target variables, model architectures, author affiliations, adoption status, and whether the associated data and code are open source. For adoption status, we distinguish three tiers: Research, referring to published demonstrations or proof-of-concept studies not yet integrated into an agency workflow; Pilot, referring to implementations that run within an operational tool or workflow but are explicitly designated as experimental, undergoing testing, or awaiting formal acceptance; and Operational, referring to models that are routinely used to inform planning, regulatory, or real-time management decisions and are maintained as part of an agency’s production toolset.

Across these tiers, ML has been applied to a diverse set of problems. For example, ML models are integrated into water resources planning models CalSim II and CalSim 3.0 [46,47] turbidity forecasting [48–50], suspended sediment characterization [51], and satellite-based mapping of dissolved organic matter [52]. Recent advancements also extend into toxigenic risk assessment, utilizing ML to predict the probability of harmful algal blooms [53,54] and to perform high-accuracy conversions of salinity measurements into individual ion concentrations [55–57].

In addition, ML has catalyzed innovations in ecological protection and infrastructural design [58] in the Delta. Applications range from characterizing fine-scale fish migration strategies through computational fluid dynamics to forecasting population-scale entrainment risks at major water diversion facilities [59]. ML also serves as a critical surrogate for compound coastal flooding assessments under future climate change [60,61] and provides unsupervised algorithms for mapping the internal stability of aging levee systems [62]. In emergency management, ML-based recommendation engines now provide optimized recovery plans for multi-island breach events in seconds, enabling a proactive and data-centric approach to Delta water quality and disaster response [63].

Furthermore, management of the Delta was historically constrained by a data-poor environment, where sparse observations and coarse grab samples limited the pace of scientific understanding [64]. Over the past four decades, advances in high-frequency telemetry, multi-spectral satellite imagery, and genomic sequencing [65–71] have transformed the system into a far more data-rich environment. However, this growth has been uneven, and important monitoring gaps remain, particularly for aquatic vegetation, water temperature, groundwater, and sediment dynamics in the Delta [72]. Expanding remote-sensing capabilities offer opportunities to address some of these gaps. Yet the rapid increase in data volume has created a new bottleneck: synthesizing disparate datasets from different sources and integrating them with analytical tools to produce actionable science remains challenging.

Synthesis provides a structured way to take stock of current knowledge, articulate that understanding clearly to both internal and external audiences, and identify where existing data and scientific interpretations either reveal new insights or expose gaps and inconsistencies that warrant further refinement [73]. There have been recent fruitful synthesis efforts on water temperature data [74], climate change effects on habitat suitability [75], pelagic fishes [76,77], and food webs [78] to help answer management-relevant research questions in the Delta. However, scaling retrospective synthesis efforts to meet operational demands remains challenging, because conventional process-based models and manual statistical frameworks are often computationally intensive, structurally rigid, and difficult to adapt for efficiently leveraging large, heterogeneous, and increasingly high-frequency environmental data streams [79,80]. Consequently, the traditional analytical frameworks are often too slow or too rigid to integrate these massive, multi-modal data streams into real-time management [81]. In this context, ML is not merely an alternative surrogate tool, but increasingly serves as essential infrastructure, providing the computational backbone needed to ingest vast, disparate datasets—from in situ monitoring networks to satellite imagery—and to identify complex, non-linear relationships that elude conventional rigid analytical frameworks [82–84]. By serving as a high-speed engine for data fusion and pattern recognition, ML can translate raw digital abundance into actionable knowledge—such as real-time salinity forecasts—needed to manage Delta’s competing demands for water reliability and ecological health in an era of rapid change.

1.3. Study Scope

While prior efforts have synthesized the state of the science of Delta modeling with respect to multi-dimensional process-based models [85], the rapid expansion of data-driven techniques over the past three decades necessitates a dedicated assessment of ML. The present study aims to complement that foundational work [85] by focusing specifically on ML applications and addressing the need for data synthesis mentioned above. By providing a comprehensive three-decade retrospective and a forward-looking roadmap, this paper aims to update the state of the science and practice of Delta modeling, illustrating how ML can bridge the gap between increasing data availability and actionable management in Delta.

Specifically, this study traces the evolution of ML techniques used to address Delta water and environmental management challenges, highlighting key milestones, successes, and limitations. Building on this foundation, the study further identifies emerging trends and future directions for ML applications in the Delta, emphasizing how these advances can support the evolving needs of water supply operations and ecosystem management. By offering both a retrospective overview and forward-looking perspectives, the study aspires to inform a broad audience, including modelers, researchers, water managers, policymakers, and other stakeholders on the development and deployment of ML for sustainable Delta water and environmental management. It should be noted that, although this review covers a broad spectrum of ML applications in Delta water and envi-

ronmental modeling, the depth of discussion varies across topics. Salinity prediction and CalSim/DSM2 surrogate modeling receive comparatively greater attention because they represent the earliest and most extensively developed ML applications in the Delta and are therefore discussed in greater detail in Section 2. More than three decades ago, the need for rapid salinity prediction and computationally efficient surrogate models for planning and operations provided the primary motivation for applying ML to Delta water management. Consequently, these applications have the longest history, the largest body of published literature, and the highest level of operational maturity. The emphasis throughout this review therefore reflects the historical evolution and maturity of the field rather than the relative importance of different water and environmental management challenges. It should also be noted that ChatGPT 5 and Gemini 3 were used to create selected icons in Figures 3, 6, 7, 9 and A2. The lead author reviewed these icons and is solely responsible for their content.

The remainder of the paper is organized as follows. Section 2 reviews ML applications in the Delta. Section 3 explores prospects and potential research directions. Section 4 presents concluding remarks.

2. Three-Decade Retrospective

This section presents a comprehensive retrospective of ML applications in the Delta, drawing on a structured literature search supplemented by the authors' domain expertise. Peer-reviewed publications were primarily identified using the Web of Science database with combinations of keywords including "Sacramento–San Joaquin Delta," "machine learning," "deep learning," "artificial intelligence," "neural network," and related terms. The search covered publications from 1980 through 2025. Retrieved publications were screened for relevance and included if they developed, applied, or evaluated ML methods for Delta-related water resources, hydrodynamic, water quality, ecological, infrastructure, or environmental applications. To improve the comprehensiveness of the review, the database search was further supplemented by the authors' knowledge of the field and a review of references to identify additional Delta-specific studies, technical reports, and conference proceedings that may not have been captured through database keyword searches. Table A2 in Appendix A provides an inventory of the studies identified and included in this review.

2.1. Establishing the Foundation (1990s)

2.1.1. Motivation

ML applications in the Delta emerged in the early 1990s, motivated by the need for fast and reliable simulation of salinity at key Delta locations, including regulatory compliance stations and export intakes of SWP, CVP, and other water conveyance systems [86]. At that time, the standard practice was to use a computer module called Minimum Delta Outflow (MDO), a sub-routine of the planning model DWRSIM, to calculate the Delta outflows required to meet predefined salinity targets. These estimates were based on export rates and the operational status (open or closed) of the Delta Cross Channel (DCC) gates. A key limitation of the MDO was its inability to incorporate antecedent flow conditions, which influences current salinity levels. Consequently, its performance was suboptimal in certain years and months [87].

To address this limitation, the "G-model" was introduced to correlate salinity with both current and antecedent net Delta outflow (NDO) [88]. NDO is a calculated water balance term representing total Delta inflows minus outflows, including exports, channel depletion, diversions, and consumptive use. Incorporating antecedent flow conditions improved estimation of flow-salinity relationships. Despite the advances, the G-model ex-

clusively relied on NDO and could not distinguish between different flow patterns that produced identical NDO values. In reality, varying combinations of inflows and exports with the same NDO could produce markedly different salinity patterns. Additionally, the G-model did not account for the effects of gate operations, which also play an important role in shaping salinity conditions within Delta.

Attempts have also been made to develop statistical models for estimating Delta salinity in the 1980s, although these efforts did not yield satisfactory results [89]. Moreover, classical time-series analysis techniques have been explored for this purpose in the early 1990s [90]. However, such methods typically require transforming the data into a stationary time series format, which ultimately undermines the practicality of these models as viable substitutes for the MDO module.

Process-based hydrodynamics and water quality models, such as DWRDSM and its successor version DSM2, can incorporate a wide range of inputs such as exports, inflows, and gate operations. These models also account for antecedent flow conditions and generate realistic salinity patterns across the Delta. However, their computational demands are much higher, often by several orders of magnitude compared to those of the MDO module [91]. As a result, these models are generally not considered optimal replacements for MDO in operational settings.

Artificial Neural Networks (ANNs) are flexible function approximators widely used to map nonlinear relationships between input factors (e.g., flow) and output variables (e.g., salinity) [92]. With a sufficient number of hidden neuron nodes, ANNs can approximate functions with finite discontinuities to a desired level of precision. Consequently, ANNs do not require predefined assumptions about the form of the input-output relationships, such as prescribing a specific transfer function. Once trained, ANNs offer rapid computation and reliable predictive performance, making them a compelling alternative to the MDO module.

2.1.2. Development of Early ANN Models

The initial generation of ANNs for Delta applications was implemented using the open-source Stuttgart Neural Network Simulator (SNNS) software [93] in the mid-1990s [87]. A conventional multi-layer perceptron (MLP) architecture [94–96], a type of ANN model, was adopted due to its structural simplicity and relatively efficient training performance. An MLP generally consists of three main components: an input layer, one or more hidden layers, and an output layer. Each layer contains processing units (neurons) that are interconnected to neurons in neighboring layers. Training an MLP network includes forward propagation, loss (or error) evaluation, back propagation, and iterative optimization (Figure 4).

During training, input features (e.g., flows and gate positions) are propagated through successive network layers, where weighted connections, bias terms, and activation functions enable the ANN to learn complex nonlinear relationships. This forward propagation produces model predictions, which are compared with observed values using a loss function. The resulting prediction error is backpropagated through the network to compute gradients, and model parameters (weights and biases) are iteratively updated using optimization algorithms (e.g., gradient descent [97]) until convergence or a predefined stopping criterion is reached (Figure 4). A more detailed description of the training process is provided in Appendix A.

For the early MLP implementations in Delta modeling, multiple architectural setups were tested. This included varying both the depth of the network and the numbers of neurons within each hidden layer. Such trial-and-error design was essential for determining the optimal network structure. Model performance was primarily assessed using mean squared error (MSE) [98] as the loss function. MSE quantifies the average squared differences between model-simulated and observed values. Parameter optimization was

commonly carried out using the scaled conjugate gradient method [99]. This method was selected due to its faster convergence relative to the standard gradient descent techniques.




Figure 4. Architecture of early artificial neural network (ANN) models, their training process as well as applications in the Delta.

2.1.3. Early Applications

Foundational Work

Building on the motivations described above, the earliest ML applications in the Delta (before 2000) explored ANNs, specifically MLPs, as nonlinear, data-driven alternatives to traditional regression techniques and simplified conceptual models for representing Delta hydrodynamics and water quality. These pilot studies established the foundational architecture still used in Delta modeling today. They are highlighted in Figure 4 and are discussed as follows.

(a) Flow and Salinity Dynamics

Initial ANN applications focused on modeling the relationship between Delta inflows including Sacramento River flow, exports, gate operations, and salinity (typically expressed as electrical conductivity or total dissolved solids) and chloride concentrations at key boundary and interior locations [86,91,100]. A pivotal methodological advancement was the shift from using NDO as a single aggregated predictor (as briefly noted in Section 2.1.1) to a multi-input, memory-based approach. By training ANNs on component flows (e.g., Sacramento River, San Joaquin River, exports, channel depletions, eastside streams) and gate positions, combined with explicit temporal memory through lagged inputs, model performance was significantly improved. These applications demonstrated that ANNs could outperform linear and quasi-nonlinear regression models, particularly under low-flow, high-salinity conditions where system behavior is strongly nonlinear.

(b) Carriage Water Estimation

One of the most conceptually significant ANN applications in the 1990s was the investigation of carriage water, also referred to as marginal export cost (MEC). Carriage water represents the additional Sacramento River inflow required to maintain a constant Delta salinity level when exports are increased. To ensure the models could independently evaluate operational impacts, synthetic DSM-simulated data were augmented with historical data to serve as the training dataset using “inverted” gate positions to decouple historical correlations between gate operations and flow regimes [101]. The findings indicated that carriage water is a significant and highly variable quantity that depends on the controlling salinity station, antecedent flow conditions, and the nature of the export increase.

By categorizing costs into transient and continuous impulses, the studies revealed that sudden perturbations generally require higher water costs than sustained increases [91,102]. General results showed that MEC is lowest in the western Delta and increases

toward the interior and export pumps, fluctuating notably depending on which station is controlling. Furthermore, the ANN framework confirmed a temporal “lag” effect, where the salinity response to flow conditions is typically spread over two to three months due to the Delta system’s inherent memory. This physically consistent methodology provided a more robust tool for operational planning compared to existing modules.

(c) Total Trihalomethane (TTHM) Formation Modeling

In parallel with flow-salinity and carriage water applications, ANNs were explored for modeling TTHM formation and bromine incorporation factors in Delta source waters [103]. These models incorporated water quality variables such as bromide concentration, dissolved organic carbon, ultraviolet absorbance, chlorine dose, ammonia, temperature, reaction time, and pH.

ANN-based TTHM formulations improved predictive performance relative to traditional log-linear empirical equations, particularly across different THM species. Sensitivity analyses showed that the ANN models captured nonlinear interactions among precursor concentrations and treatment conditions, offering a more flexible framework for evaluating disinfection by-product formation relevant to drinking water supply.

Significance of Pioneering Efforts

Although exploratory by modern standards, these early ANN applications were highly influential. They demonstrated—well ahead of broader adoption in Delta modeling—that ML could serve as a fast, nonlinear surrogate for regression and process-based conceptual models, enabling scenario analysis, planning studies, and operational decision support in the Delta.

Several foundational contributions stand out. First, these studies established the feasibility of using data-driven models to emulate sophisticated hydrodynamic and water quality models while reducing computational cost by orders of magnitude. Second, they introduced the idea that ML models could be embedded within planning and operations frameworks (e.g., DWRSIM) rather than used solely as standalone predictive tools. Third, the explicit use of memory and multiple flow components anticipated many features of modern sequence-based and physics-informed ML approaches.

Collectively, this body of work laid the conceptual groundwork for later integration of ML into Delta modeling. They showed that data-driven approaches could capture nonlinear relationships in Delta hydrodynamics, water quality, and water operations, supporting their use in operational decision-making and long-term planning. At the same time, they spurred further research to refine model structures, expand predictor sets, and improve training strategies, advancing both scientific understanding and practical applications of ML in the Delta.

Limitations of Pioneering Efforts

Despite their innovations, these early ML applications faced several limitations. Foremost among these was computing power. Training ANN models in the 1990s relied mostly on gradient-descent-based optimization and required substantial computational time, often constraining network complexity, memory length, and the scope of sensitivity analyses that could be performed.

Data availability also posed challenges. While daily records spanning approximately two decades were considered extensive at the time, these datasets were relatively short by modern standards and limited the ability to explore rare or extreme conditions. In addition, strong correlations among flow variables and operational decisions made it difficult to disentangle causal relationships, raising concerns about interpretability and robustness under non-stationary conditions.

Finally, these early applications depended on feed-forward architectures with manually engineered lag structures, reflecting both software limitations and limited theoretical guidance on network design. Uncertainty quantification, formal validation under changing conditions, and explicit integration of physical constraints were largely absent, highlighting important gaps in the research that require further development and refinement.

2.2. Enhancements and Integration into Planning Tools (2000–2010)

Between 2000 and 2010, ANNs transitioned from experimental tools to core operational components of Delta water management, yet the predominant model architecture remained the MLP. By mimicking the high-fidelity outputs of process-based models (e.g., DSM2) at a fraction of the computational cost, these models became essential for planning and operational decision-making. This section reviews key studies that document the architectural refinement of these networks from this era, their validation against traditional methods, and their integration into CalSim for flow-salinity estimation.

2.2.1. ANN Architecture and Input Refinement

Initial efforts in the early 2000s focused on refining the internal structure of ANN models to better mimic the complex hydrodynamics of the DSM2 model. A foundational study [104] described significant improvements in model accuracy achieved through input perturbation, where data from multiple DSM2 runs was manipulated to isolate the effects of the DCC and remove correlations between Sacramento River flows and gate positions. This study also detailed how increasing the number of internal hidden nodes from 8 to 32, supported by this expanded training dataset, allowed the network to better approximate salinity peaks. A concurrent research effort by [105] explored the use of hourly data and the Levenberg–Marquardt backpropagation algorithm to capture sub-daily salinity fluctuations. Their work highlighted the importance of tidal flux and even experimented with using moon phase illumination as a proxy for tidal effects to improve prediction accuracy. In a follow-up research study, [106] introduced Genetic Algorithm (GA) optimization to automate the selection of training data and the number of hidden neurons, demonstrating that GA-optimized networks could further reduce errors when compared to traditional neural network predictions.

2.2.2. Implementation and Integration into CalSim

The transition of ANNs from standalone predictive tools to integrated components of operational modeling was a critical advancement. Reference [107] detailed the technical framework for integrating ANNs into the CalSim model to replace the less sophisticated G-model. They addressed the primary challenge of using non-linear ANNs within CalSim's linear programming solver by developing a linearization technique that calculated a feasible "slope and intercept" for constant salinity between specific export points. This allowed CalSim to dynamically adjust upstream reservoir releases and Delta export limits while satisfying salinity standards. The ANN's improved accuracy was then verified against DSM2 model results by [108] using a "full-circle" methodology. The full-cycle analysis, as illustrated in Figure 5, begins by running CalSim to generate Delta boundary inflows, exports, and operational conditions (Step 1). DSM2 is then run using the CalSim outputs from Step 1 to simulate flow and salinity in the Delta (Step 2). In the next step (Step 3), ANN models are trained using the DSM2 simulations generated in Step 2. CalSim is then rerun with the trained ANN embedded to produce salinity estimates at key Delta compliance locations, along with updated Delta inflows, exports, and operations (Step 4). These updated inflows, exports, and operational conditions are subsequently used to drive DSM2 to simulate salinity at the same compliance locations (Step 5). In Step 6, the salinity values obtained from Step 4 and Step 5 are compared. If the results agree favora-

bly, the ANN model is considered validated; otherwise, the ANN model is retrained using DSM2 simulations from Step 5, and the process is repeated.




Figure 5. CalSim-DSM2-ANN Full Circle Analysis Schematic.

A comparative study [109] further utilized this full-cycle analysis to benchmark the ANN model against the G-model, finding the G-model systematically underestimated salinity, resulting in standard violations. Reference [110] further refined this ANN implementation through the work of the CalSim ANN Refinement Team, focusing on systematic improvements to model stability and the selection of more robust input parameters to ensure the model performed reliably across varied hydrological conditions.

2.2.3. Advanced Modeling of Tidal Influence

By the late 2000s, modeling efforts advanced to incorporate physical drivers of salinity intrusion that were previously approximated. Reference [111] introduced an enhanced ANN that explicitly incorporated tidal range as an input variable. This development allowed the model to represent the marginal salinity response to unit flow changes with much greater fidelity, particularly in its ability to mimic both continuous and transient impulses that earlier ANN models struggled to capture. This work represented a culmination of a decade of refinement, providing a robust, tidal-aware ANN that could faithfully serve as a rapid, high-performance surrogate for DSM2 within the CalSim II model.

2.2.4. Synthesis and Implications for Delta Water Management

Taken together, these studies mark the transition of ANN models from research tools to operational decision-support components in Delta water management. Early refinement efforts [104,107,108,110] are still pilot in nature, but they improved predictive skill and robustness of the ANN models, while comparative analyses with DSM2 [109] built confidence in ANN behavior relative to trusted process-based mechanistic models. Subsequent work addressing tidal influence [111] established a clear pathway for operational integration of ANNs into CalSim II.

This body of work laid the technical groundwork for later developments in surrogate modeling, hybrid physics–machine learning approaches, and modern efforts to incorporate advanced data-driven tools into Delta water operations and planning decision making.

2.3. Expanding the Landscape (2010–Present)

Since 2010, the application of ML in water and environmental modeling for the Delta has diversified substantially across modeling targets, modeling approaches, and participating institutions.

This expansion reflects the increasingly complex and interconnected challenges facing the Delta system. A wide range of target variables are being modeled, driven by the need to better understand and manage water supply reliability, water quality degradation, ecosystem stressors, and infrastructure vulnerability. Methodologically, this period has been characterized by rapid advancement in ML techniques, particularly in the field of deep learning (Figure A1 of Appendix A), as well as hybrid process-based and ML modeling. Studies have explored a broad suite of algorithms, enabled by increased computational power and the growing availability of high-resolution observational and simulation datasets. These developments have enhanced model accuracy, scalability, and applicability to complex, nonlinear Delta processes. Institutionally, ML research and implementation in the Delta have expanded from state agencies to include federal partners, academic institutions, and private-sector organizations. This growing multi-sector collaboration has accelerated innovation, facilitated data sharing, and increased investment in operational ML tools for decision-support.

Against this backdrop, this subsection reviews recent ML applications organized around four thematic areas: (1) hybrid modeling and deep learning, (2) water operations and water budget, (3) water quality and ecosystem health, and (4) infrastructure resilience and disaster risk (Figure 6). Most applications within these themes remain research-oriented, with a few exceptions at the pilot or operational stage which will be highlighted in detail below.




Figure 6. Growth of machine learning development and applications in the Delta since 2010. ET_0 refers to reference evapotranspiration; CO_2 and CH_4 indicate carbon dioxide and methane emissions, respectively; DOM designates dissolved organic matter.

2.3.1. Hybrid Modeling and Deep Learning

Pioneering Efforts

While purely data-driven models are exceptionally flexible and capable of identifying complex patterns in historical data, they often operate as “black boxes” (Figure 3) that may produce physically impossible results (e.g., negative salinity) because they lack an inherent representation of the physical laws. Hybrid approaches bridge this gap by integrating the predictive power of ML with the theoretical rigor of process-based models. These hybrid frameworks ensure that the resulting simulations not only achieve high statistical accuracy but also remain consistent with the physical laws of mass balance and fluid dynamics, making them more reliable for long-term planning under changing climatic conditions.

One prominent hybrid strategy involves using the output of a process-based model as an input feature for a machine learning model, a technique explored by [112]. This study focused on emulating salinity at Martinez, the downstream boundary of the Delta. To capture the long-term cumulative effects of freshwater flow on salinity, they utilized the results of the Martinez Boundary Salinity Generator (MBSG)—a process-based empirical relationship—as an auxiliary input feature for Deep Learning models like Long Short-Term Memory (LSTM) networks. By feeding the MBSG’s salinity estimates into the neural network, the ML model is “pre-conditioned” with physical information about antecedent flow conditions, which outperformed the benchmark model across all five metrics, particularly during extreme wet, dry, and critically dry years.

A more direct integration of physics occurs through the embedding of physical governing equations into the loss function of the ML model, as demonstrated by [113] in their work on Physics-Informed Neural Networks (PINNs). Unlike standard neural networks that minimize only the error between predicted and observed data, PINNs incorporate a “physics loss” term derived from the one-dimensional advection-dispersion equation. This additional constraint penalizes the model if its predictions violate the fundamental transport mechanics of the Delta. This approach is particularly effective in data-sparse scenarios, as the physical equations guide the model toward realistic solutions even when training data is limited, ensuring that the simulated salinity gradients across the Delta remain physically plausible.

Lastly, a third hybrid method uses process-based model results to fit or constrain station salinity data before training the ML component, as detailed by [114]. This study developed a hybrid empirical-Bayesian ANN by first utilizing a refined version of the empirical G-model called the Delta Salinity Gradient (DSG) model to establish a baseline physical relationship for each salinity station. A Bayesian ANN was then integrated to reconcile the residuals and refine the predictions. By using the process-based model to provide the fundamental structural behavior of the system, [114] were able to ensure “structural validation”—meaning the model responds correctly to changes in coastal water levels and flows—while the ANN component provided the necessary flexibility to capture site-specific nuances (e.g., spring-neap tidal variations and prediction uncertainty) that the empirical model might otherwise overlook.

Deep Learning Architectures

The evolution of salinity modeling in the Delta has recently shifted from traditional, shallow ANNs toward sophisticated deep learning architectures. While traditional ANNs—such as those used in CalSim II and CalSim 3.0—typically consist of a few fully connected layers, deep learning models leverage multiple hidden layers and specialized structural units to capture higher-order nonlinearities and temporal dependencies. This transition has been largely facilitated by the dramatic advancement in computational

power, particularly the availability of Graphics Processing Units (GPUs) and specialized deep learning frameworks, which make training complex models with millions of parameters feasible within reasonable timeframes. Unlike the manually engineered feature extraction used in earlier models (e.g., using fixed 11-day averages as inputs), deep learning architectures can automatically identify relevant spatial and temporal patterns directly from raw high-frequency data.

A significant advancement in this field is the move from Single-Task Learning (STL) to Multi-Task Learning (MTL), as detailed by [115] and [116]. Traditional models required a separate ANN for each salinity station, but the MTL architecture uses shared hidden layers to jointly estimate salinity at multiple locations across the Delta. Qi et al. [115] found that this approach not only reduces training and inference time by 90% but also improves predictive accuracy by allowing the model to learn shared hydrodynamic features that influence the entire Delta. Furthermore, [115] introduced trainable convolution layers to replace predetermined input preprocessing, allowing the model to optimize its own “memory” of antecedent flow conditions rather than relying on static, human-defined averages.

Recent research has also integrated specialized recurrent and residual units to better handle the long-term memory requirements of Delta salinity. Reference [112] demonstrated the effectiveness of LSTM networks and Convolutional Neural Networks (CNNs) for emulating boundary salinity at Martinez, showing that these models can better capture the “long-memory” effects of freshwater outflow compared to standard MLPs. Building on this, [117] introduced Residual LSTM (Res-LSTM) and Residual GRU (Res-GRU) architectures. By incorporating residual connections, these models mitigate the vanishing gradient problem in deep networks, allowing for more stable training and superior performance in capturing both tidal fluctuations and seasonal trends.

The most recent frontier in Delta salinity modeling involves the application of transfer learning and data augmentation to address persistent challenges related to data scarcity. As demonstrated by [118], deep learning models can be pre-trained on expansive synthetic datasets generated by process-based hydrodynamic models, which provide a physically consistent representation of salinity variability in the Delta. These models are subsequently fine-tuned using limited historical observations from individual monitoring stations to capture local dynamics. This hybrid framework enables deep learning architectures to maintain high predictive skill at data-sparse locations, effectively bridging the gap between traditional mechanistic modeling and modern data-driven approaches.

2.3.2. Water Operations and Water Budget

Further Refinement and Implementation into CalSim

Building on earlier ANN implementations in CalSim II, [47] developed a new generation of flow–salinity ANN models fully integrated into the CalSim 3.0 operational planning framework. A primary motivation for redevelopment was that the CalSim II ANNs—trained on a relatively short 16-year period (water years 1974–1991)—performed well during calibration but exhibited degraded skill under independent validation and were incompatible with the enhanced spatial representation and updated Delta processes introduced in CalSim 3.0. In contrast, the CalSim 3.0 ANNs were trained using a substantially longer historical record (1940–2015), explicitly including extreme hydrologic conditions such as the 2011–2015 drought, thereby improving robustness across a wider range of operational and climatic regimes. Validation against DSM2 demonstrated markedly improved agreement at key compliance locations, particularly Rock Slough and Jersey Point, which are critical for SWP and CVP operational decisions [47].

The CalSim 3.0 ANN architecture also differs from that used in CalSim II. While CalSim II relied on relatively compact feedforward networks with limited memory and fewer input nodes, the CalSim 3.0 ANN employs a deeper feedforward architecture with

multiple hidden layers and extensive antecedent memory, resulting in 126 input nodes that capture both short-term dynamics and longer-term hydrologic conditioning. Inputs were expanded to include not only major inflows, exports, Delta Cross Channel operations, and consumptive use, but also tidal energy metrics and upstream salinity, reflecting the recognized importance of tidal influence and estuarine memory in Delta salinity dynamics. In addition, CalSim 3.0 replaced the simplified Consumptive Use representation with the Delta Channel Depletion model, providing spatially refined and physically consistent depletion estimates that substantially improved ANN training fidelity and operational realism [47].

From an implementation perspective, these enhancements significantly strengthened the role of ANN models as fast, reliable surrogates for DSM2 within a statewide planning context. The improved training strategy, expanded input space, and deeper network architecture allow the CalSim 3.0 ANN to more faithfully emulate DSM2 salinity responses while remaining computationally efficient for long-term scenario analysis. This represents a critical step beyond earlier ANN applications, demonstrating that ML-based emulators can be robustly embedded within operational planning models to support salinity compliance evaluation, export management, and systemwide trade-off analyses under both historical and future hydrologic conditions [47].

More recent work has extended this paradigm beyond emulation of a single, fixed Delta configuration toward explicitly representing how changes in Delta geometry, restoration actions, and sea level rise propagate through statewide operations. Reference [119] advanced this evolution by coupling high-resolution, multi-dimensional hydrodynamic modeling with modern ML techniques to assess the long-term water supply implications of future Delta conditions within the CalSim framework.

In this pilot study [119], recurrent ANN surrogate models based on Gated Recurrent Units (GRUs) were developed to capture the temporal dynamics of flow and salinity transport. A key methodological innovation was the use of transfer learning, whereby models were first pre-trained using a century-long dataset generated from the 1D DSM2 model and subsequently fine-tuned using results from high-resolution 3D (SCHISM) and 2D (RMA2D) simulations. To enhance sensitivity to scenario-specific changes, contrastive learning was employed to explicitly train the models to reproduce differences between baseline and alternative conditions. Model training efficiency was achieved through a Design of Experiments framework combining k-means clustering and stochastic perturbations, allowing a broad range of hydrologic and operational variability to be represented using a limited set of multi-dimensional simulations. The trained ANN surrogates were then embedded into CalSim using a custom Java-based interface to TensorFlow models, enabling rapid estimation of Delta outflow requirements for salinity compliance under scenarios including Suisun Marsh restoration and sea level rise. This integration allowed the study to evaluate how localized changes in Delta mixing influence upstream reservoir storage and south-of-Delta deliveries within a statewide operations context. Results indicated that restoration scenarios that reduce tidal mixing can yield moderate improvements in water supply reliability, whereas sea level rise substantially increases mixing and imposes large water supply penalties.

Importantly, [119] emphasized that the CalSim implementation of these advanced ML surrogates is pilot and experimental in nature. While round-trip validation against multi-dimensional models demonstrated that the approach reliably reproduces relative scenario impacts, some discrepancies in absolute salinity levels were attributed to differences in model assumptions (e.g., channel depletion and boundary condition treatments). As a result, further testing, refinement, and coordination with CalSim development teams are required before these methods can be incorporated into production versions of CalSim. Nevertheless, this pilot study represents a significant methodological advance,

illustrating a viable pathway for integrating physics-informed, high-resolution machine-learning surrogates into statewide planning models to evaluate future Delta management strategies under climate change and sea level rise.

Volumetric Fingerprinting

In addition to direct salinity modeling, the use of ANNs has expanded to include the emulation of process-based source water fingerprinting. This approach allows for the simulation of the underlying hydrodynamic routing that drives salinity transport, rather than just the concentration itself. Reference [120] demonstrated this by developing ANNs to replicate the volumetric contributions (i.e., fingerprints) of various source waters—such as the Sacramento River and San Joaquin River—at 17 Delta locations.

This methodology effectively decouples the transport physics from constituent concentrations; the ANN-derived fingerprints can be converted into salinity values by simply multiplying the volumetric fractions by the observed boundary salinity levels. A major advantage of this approach is its robustness; unlike direct salinity ANNs that require re-training if boundary concentration patterns shift, the fingerprinting ANN remains valid as long as the underlying hydrodynamics of the system are consistent. The study reported high precision across diverse hydrologic conditions, with correlation coefficients generally exceeding 0.90, marking a critical milestone in the development of rapid, process-informed surrogates for complex estuarine modeling.

Reference Evapotranspiration

Evapotranspiration (ET) is a fundamental component of the hydrologic cycle and a critical driver of agricultural water demand, particularly in intensively managed regions such as the Delta. In practice, actual ET is commonly estimated from reference evapotranspiration (ET_o), which represents the potential evapotranspiration rate over a standardized, well-watered reference surface. Traditionally, ET_o has been calculated using empirical equations that rely on multiple meteorological inputs, often requiring sensor-dense and costly monitoring infrastructure. In recent years, ML approaches have emerged as an alternative for estimating reference evapotranspiration. In California, Ahmadi et al. [121] developed an ML-based framework, termed SolarET, that uses solar radiation as the sole input variable, thereby bypassing the need for expensive, multi-sensor reference stations. The model was trained using data from 131 stations distributed across the state, including several located within the Delta, enabling robust spatial generalization and regional applicability. By providing reliable, site-specific ET_o estimates in environments where maintaining traditional reference surfaces (e.g., well-watered grass) is challenging, SolarET offers a valuable tool for localized water management. Improved estimation of reference evapotranspiration directly enhances the accuracy of consumptive use calculations—one of the most uncertain yet critical components of the Delta water budget.

Similarly, the shift toward deep “global learning” architectures—where models are trained on data from diverse locations and conditions to learn broadly applicable relationships—has enhanced the ability to forecast monthly reference evapotranspiration up to one year across the Central Valley and the Delta, capturing broad regional climatic signals [122]. By training on a diverse array of monitoring sites, including those directly within the Delta, these models achieve a level of generalizability that ensures high-accuracy forecasting even in data-sparse zones. For Delta water managers, the ability to predict future atmospheric water demand is essential for long-term resource planning. These advanced forecasting tools provide the necessary data to refine consumptive use projections, ultimately supporting more informed operational decisions as they manage the complex interplay between upstream supply and downstream demand [122].

2.3.3. Water Quality and Ecosystem Health

Ion Constituents

Reliable prediction of major ionic constituents—such as chloride, bromide, and sulfate—is critical for effective water quality management and ecosystem protection in the Delta. Historically, these relationships have been described using empirical regression models that link Electrical Conductivity (EC), commonly used as a surrogate for salinity, to individual ion concentrations. While such formulations [123,124] have been widely applied to support water operations in the Delta, their rigid structures and simplified assumptions about source contributions limit their ability to represent the system's strongly nonlinear and time-varying mixing behavior. Because Delta salinity arises from the combined influence of seawater intrusion, riverine inputs, and agricultural return flows under fluctuating tidal and hydrologic regimes, identical EC values can be associated with substantially different ionic compositions.

ML offers a flexible, data-driven alternative that can capture strong nonlinear relationships among large observational datasets without depending on pre-defined functional forms. By incorporating a wide range of predictors such as EC, water year type, and season, ML models can effectively and dynamically characterize constituent behavior across both spatial and temporal scales.

A series of pilot investigations by Namadi et al. [55–57] progressively strengthened the application of ML to ion–salinity modeling in the Delta. The initial work [55] showed that four ML approaches, most notably Random Forest models, substantially exceeded the performance of traditional regression equations in predicting 12 ionic constituents from EC at South Delta monitoring sites. Building on this foundation, a subsequent study [56] expanded the framework across the entire Delta using approximately six decades of observations, demonstrating that ANNs delivered the greatest predictive accuracy. That effort was complemented by the development of an interactive web-based dashboard to support real-time simulation and visualization. The most recent study [57] applied ANN-based models at three major water diversion facilities, yielding coefficients of determination exceeding 0.96 and reductions in mean absolute error of up to 88% compared with earlier techniques. This study further incorporated model interpretability tools and hybrid ML–hydrodynamic dashboards, enhancing both transparency and spatial representation.

Collectively, this body of work illustrates a transition from initial feasibility demonstrations to operational deployment, highlighting the role of ML in augmenting conventional modeling frameworks by representing nonlinear processes, enhancing predictive performance, and delivering interpretable, user-oriented tools to support water quality management in the Delta.

Water Temperature

Expanding beyond salinity modeling, recent work has applied ANN models to simulate daily water temperature at three critical boundary locations of the Delta: the Sacramento River at Freeport, the San Joaquin River at Vernalis, and Martinez. This effort was primarily driven by the need to extend the simulation period of the DSM2 to cover water years 1922–2015, thereby aligning with the long-term planning horizons used in the CalSim3 model. A longer historical or model simulated data record is essential because it captures a wider range of hydroclimatic variability, including extreme wet and dry conditions, which supports more robust long-term water resources planning. By utilizing ANNs as surrogate models, He et al. [125] derived historical water temperature boundaries for periods where direct observations were unavailable.

The methodology employed a three-layer MLP architecture consisting of an input layer, a hidden layer with five neurons, and an output layer. A separate MLP was devel-

oped for each of the three boundary locations, utilizing a sigmoid activation function to capture non-linear relationships between inputs and temperature outputs. The models were trained using historical daily observations from 1990 to 2013. Validation results indicated that the ANN models effectively mimicked observed variation patterns and magnitudes, achieving correlation coefficients consistently above 0.9 and high Nash-Sutcliffe efficiency values at all three locations. Notably, the downstream boundary at Martinez exhibited the most robust performance metrics, confirming the utility of the MLP structure for reliable long-term temperature boundary generation.

Suspended Sediment and Turbidity

Machine learning applications have expanded into more complex water quality parameters such as turbidity and suspended sediment concentration (SSC). These parameters are vital for managing the habitat of endangered species like the Delta smelt, which is often salvaged at export facilities when turbidity levels exceed a critical threshold. The Delta Turbidity ANN Model (DASM-T) was developed to provide a rapid alternative to computationally intensive mechanistic models like DSM2 and RMA-2 [48–50]. Chen and Roy [48,49] focused on validating MLP and Nonlinear Autoregressive with Exogenous Inputs (NARX) architectures, demonstrating their ability to accurately emulate DSM2-simulated turbidity at key locations across the Delta. Reference [50] expanded this work by utilizing feedforward ANNs to evaluate operational management scenarios, specifically investigating how modifying Old and Middle River (OMR) flows through export curtailments could control turbidity levels at interior locations during “first flush” events.

The architectural shift toward recurrent neural networks has further enhanced the simulation of sediment transport dynamics. While the DASM-T turbidity models primarily utilized standard feedforward MLP structures, Kim et al. [51] introduced LSTM networks to capture the temporal dependencies of SSC. Comparative evaluations revealed that the LSTM models often achieved superior predictive accuracy compared to the process-based DSM2-General Transport Model (DSM2-GTM), yielding significantly higher Nash-Sutcliffe Efficiency (NSE) values and lower root-mean-square error to standard deviation ratios (RSR).

A notable similarity across these studies is their shared goal of creating rapid “surrogate” models that can be integrated into broader planning tools like CalSim to bypass the computational overhead of mechanistic models. However, they differ in their handling of physical constraints and additional variables. The DASM-T turbidity models primarily rely on flow and turbidity boundary conditions from the Sacramento River, San Joaquin River, and Eastside streams. In contrast, [51] explored “hybrid” modeling by incorporating DSM2-GTM simulations as input features to provide physical guidance to the LSTM’s “black box” structure, alongside hydro-meteorological data such as precipitation and wind speed. While precipitation improved performance in some cases, wind speed was found to have a negligible overall effect. Ultimately, these studies demonstrate a progression from early emulation of turbidity to the development of sophisticated deep learning models capable of capturing sediment fluctuations with improved fidelity relative to traditional process-based counterparts in selected applications. This highlights that deep learning does not necessarily replace process-based models but can complement and enhance them by improving predictive skill, computational efficiency, and the representation of unresolved processes. In particular, physics-informed deep learning approaches can incorporate governing equations or physical constraints while also learning from data, enabling them to identify systematic model biases and potentially refine or improve process-based model formulations.

Dissolved Organic Matter

Avouris et al. [52] established a significant advancement in Delta monitoring by developing an automated, ‘sensor-to-satellite’ paradigm for modeling dissolved organic matter (DOM). By integrating high-resolution Sentinel-2 imagery with a ‘ground-to-space’ network of nine continuous monitoring stations, the researchers compiled 982 match-up samples. Utilizing a Random Forest (RF) model, the study translated remote sensing reflectance into fluorescent dissolved organic matter (fDOM) concentrations with high accuracy ($R^2 = 0.76$). This framework demonstrated that fDOM serves as a robust and cost-effective surrogate for dissolved organic carbon, allowing for the mapping of carbon ‘hotspots’ across complex tidal and seasonal cycles. While the model captured broad variability, performance was highest in winter and spring and slightly limited in narrow channels by adjacency effects from levee structures. Ultimately, this research provides a scalable, replicable framework for dynamic water quality management in complex estuarine systems.

Ecological Monitoring, Entrainment Risk Forecasting, and Behavioral Strategy Discovery

ML has also been applied to understand the ecological mechanics and risks facing fish populations within the Delta. This research has evolved from decoding the fine-scale behavioral strategies of individual fish to developing population-level risk forecasts and habitat occupancy models for endangered species.

A significant advancement in understanding individual-level behavior was introduced by [59], who developed a framework merging Computational Fluid Dynamics (CFD) with a LSTM network to reveal the movement rules of migrating juvenile Chinook salmon. By using an unsteady CFD model to simulate the dynamic fluid environment, the researchers decomposed observed fish motion into physical drag and active locomotion. They then trained LSTM to predict this active locomotion. This approach demonstrated that migratory movements can be accurately forecasted by relating environmental conditions to behavioral responses, revealing that fish actively navigate complex tidal environments rather than simply drifting with the current.

Building upon the ability to predict fish movement based on environmental cues, researchers have scaled these techniques to forecast entrainment risks at major water diversion facilities. Tillotson et al. [126] developed a machine learning framework using Quantile Regression Forest (QRF) to predict the weekly entrainment “loss”—a regulatory measure derived from salvage—of juvenile winter-run Chinook Salmon and Central Valley Steelhead. The model utilizes environmental and operational drivers such as the prior week’s loss, water exports, OMR flows, and water temperature. Although point-prediction accuracy was modest, the use of quantile intervals provided a reliable tool for managers to consider “worst-case” scenarios. Specifically, the 75th predicted quantile offered a proactive benchmark to help adjust export volumes and avoid exceeding Incidental Take Limits.

Finally, as the focus shifted from operational risk to broader conservation monitoring, machine learning has been adapted to address the challenge of detecting rare and endangered species in their natural habitats. Tillotson et al. [127] applied tree-based machine learning, specifically a Boosted Regression Trees (BRT), to evaluate the habitat characteristics associated with the detection of the endangered Delta Smelt (*Hypomesus transpacificus*). Because direct detections of Delta Smelt are now extremely rare, initial models using only target-species data exhibited limited predictive ability. To circumvent this, the study utilized a ‘proxy species’ approach, demonstrating that the catch per unit volume of the broader small pelagic fish community was the strongest predictor of Delta Smelt presence. Furthermore, by modeling combined detections of Delta Smelt and their more common relative, the Wakasagi (*H. nipponensis*), the researchers achieved significantly higher predictive power and reliability, confirming that these species utilize similar habitats and may serve as effective surrogates for conservation monitoring.

Collectively, these studies illustrate a sophisticated shift in Delta management, where ML provides actionable insights ranging from the fine-scale behavioral mechanics of migration to population-scale entrainment risks and habitat occupancy.

Algal Bloom Detection and Risk Assessment

ML has also emerged as a powerful tool in modeling the occurrence and risk of algal blooms in the Delta. Recent research has bifurcated into two primary modeling objectives: distinguishing the fundamental presence of phytoplankton blooms from non-bloom conditions and classifying the specific risks associated with toxigenic harmful algal blooms (HABs). Addressing the drivers of general phytoplankton productivity, Framsted et al. [54] utilized Random Forest (RF) modeling to evaluate the factors associated with blooms, defined by chlorophyll-*a* (*chl-a*) concentrations exceeding 16 µg/L. Their analysis across the south Delta revealed that these events are primarily driven by a combination of antecedent turbidity and water temperature, regional differences, and San Joaquin River inflows. While the potentially toxigenic cyanobacterium *Microcystis* was consistently observed during the summer, the study found it was not necessarily associated with elevated *chl-a* or high toxin concentrations. This suggests that traditional bloom metrics like *chl-a* primarily capture beneficial productivity in this region and may not always align with toxigenic risk.

In contrast, Saha et al. [53] developed a predictive framework specifically focused on the risk levels of toxigenic *Microcystis* blooms, categorizing events into “Low Risk” and “Risky” based on cell concentrations. By evaluating multiple machine learning architectures—including logistic regression, decision trees, and artificial neural networks, they identified Extreme Gradient Boosting (XGBoost) as the superior model for risk classification, achieving an accuracy of 0.88 on a scale from 0 to 1, where 1.0 indicates perfect classification accuracy. This research pinpointed water temperature, pH, and dissolved oxygen as the most critical environmental predictors for HAB risk. To facilitate proactive management, the researchers integrated these top-performing models into a web-based interactive dashboard, allowing water managers to simulate how shifts in environmental variables impact the probability of a “Risky” bloom event in real time.

Carbon and Methane Fluxes

Brereton et al. [128] developed the Regional Carbon and Climate Analytics Tool (RCCAT), an open-source machine learning framework designed to predict CO₂ and CH₄ fluxes in non-tidal wetlands. Recognizing the dual role of wetlands in carbon sequestration and methane emissions, the study evaluates a suite of models—including linear regression, gradient boosting methods, and recurrent neural networks—trained on over a decade of eddy covariance measurements from the Delta.

Among the tested approaches, LSTM models demonstrated the strongest overall performance, although improvements over simpler methods were modest, highlighting trade-offs between model complexity and predictive gain. The framework integrates widely available predictor variables, including meteorological and hydrological data and Landsat-based vegetation indices, to estimate net CO₂ exchange and methane emissions.

Using a leave-one-site-out (LOSO) validation strategy, the authors show that RCCAT can generalize to unmonitored sites, supporting its use as a scalable, data-driven complement to traditional process-based models. The framework enables regional-scale assessments of wetland carbon dynamics, offering useful insights for climate mitigation and restoration planning in the Delta.

2.3.4. Infrastructure Resilience and Disaster Risk

Fish Barrier Design

In addition to simulating or predicting traditional water quality variables, ML has recently been applied to ecological engineering in the Delta to improve the survival of out-migrating juvenile salmonids. Swyers et al. [58] developed a pilot Barrier Optimization Tool (BOT) that utilizes the GA approach to iteratively evolve the design of behavioral fish barriers at three key river junctions in the Delta. In this tool, a “chromosome” is defined as the fundamental digital data structure that encodes the unique physical traits of a candidate barrier solution. By executing thousands of simulations across varying water years and flow conditions, the GA evolves “chromosomes” representing barrier parameters like shape, length, and location to converge on designs that maximize through-Delta survival. The significance of the GA evolving these chromosomes lies in its ability to systematically navigate a massive design space to find high-performance solutions that would be impossible to identify through costly and time-consuming manual field trials.

The application of the tool demonstrated notable improvements over designs based solely on professional expert judgment, with BOT-optimized barriers increasing simulated winter-run Chinook salmon survival by 1.0% to 6.3%. The study also found that operational timing was a critical factor, as earlier implementation periods captured a larger proportion of the migrating population. More importantly, Swyers et al. [58] indicated that the BOT-identified Georgiana Slough barrier design was selected by DWR for implementation as the Georgiana Slough salmon migratory barrier. This reflects the practical utility of the tool in supporting real-world design decisions. All in all, this work illustrates a shift toward using high-fidelity, data-driven ML tools to optimize complex infrastructure design for fisheries management in the Delta.

Levee Stability Assessment

Extending the scope of machine learning from ecological protection to structural resilience, recent research has also applied these computational techniques to assess the health and internal stability of the Delta’s aging levee systems. Russo and Athanasopoulos-Zekkos [62] investigated the application of unsupervised machine learning algorithms to characterize the internal structure of sandy levees on Grand Island using non-invasive geophysical data. By integrating high-resolution surface data from unmanned aerial vehicle-mounted LiDAR with subsurface parameters—specifically shear wave velocity and apparent resistivity—the study established a multidimensional framework for detecting potential internal failure mechanisms. This approach is particularly valuable for the Delta, where internal deterioration often precedes visible surface indicators like cracks or landslides, necessitating a move beyond traditional visual inspections toward advanced subsurface mapping.

The study specifically explored the efficacy of K-Means clustering and Density-Based Spatial Clustering of Applications with Noise (DBSCAN) in identifying natural soil groupings within geophysical datasets. A key finding was that the choice of feature engineering—such as raw, standardized, or normalized data—significantly impacted the optimal model parameters and the resulting cluster accuracy. While K-Means effectively incorporated all data points into clusters, it tended to underestimate the number of natural soil layers when guided by traditional metrics like the Elbow or Silhouette rules. In contrast, DBSCAN more accurately identified the number of natural clusters based on depth but frequently classified valid data as noise. The authors further proposed a hybrid approach using DBSCAN to determine the initial cluster count for K-Means. This strategy is particularly promising as it overcomes the mathematical limitations of each algorithm in isolation, though its application in anomaly detection is still in a preliminary phase. All

in all, this proof-of-concept study provided a promising pathway, but further work is required to transition from basic clustering to standardized, semi-unsupervised anomaly detection systems that can proactively identify failure signatures across the Delta's vast and deteriorating levee network.

Levee Breach Emergence Responses

While ML applications in the Delta have addressed proactive levee health monitoring [62], the Delta Emergency Response Tool (Delta-ERT) extended these capabilities into the realm of post-disaster mitigation [63]. Specifically, the Delta-ERT utilizes a machine learning algorithm designed to suggest optimal recovery strategies—defined as specific combinations of temporary barriers—that minimize the duration of export disruptions caused by salinity intrusion following a levee breach event in the Delta.

The development of this tool involved addressing a vast and highly dimensional problem space consisting of 81 years of varying hydrologic conditions, 53 potential island breach locations, and 21 possible temporary barrier configurations. To make this computationally tractable, researchers utilized k-means clustering to subsample representative hydrologic states based on initial salinity positions (X2) and inflow volumes from the Sacramento and San Joaquin rivers. A similar subsampling strategy was applied to island breach events by grouping them into eight geographic corridors based on their hydrodynamic impacts, ensuring that the training set spanned the full range of potential disaster scenarios.

The resulting ML Ranker built using gradient-boosted decision trees functions as a rapid surrogate for complex process-based models like DSM2. By leveraging a training set of millions of simulations—executed efficiently through a specialized Python 3.12 harness and cloud computing environment—the algorithm can instantly list barrier combinations predicted to minimize export disruption for any unique breach event without requiring new simulations. This enables emergency managers to rapidly transition from identifying a structural failure to implementing a high-performance recovery plan, tailored specifically to the prevailing hydrologic and breach conditions of the event.

Flood Risk Evaluation

ML has also been applied to model potential flooding in the San Francisco Bay, including the western Delta, through ML surrogate modeling of compound water levels. Wang et al. [60] introduced a hybrid statistical–dynamical framework in which Gaussian Process Regression (GPR) was used as a surrogate for a high-fidelity hydrodynamic model (Delft3D-FM) to efficiently predict total water levels driven by tides, waves, winds, sea level pressure, monthly mean sea level anomalies, and Sacramento–San Joaquin River discharge. GPR models were trained on several hundred targeted hydrodynamic simulations selected from a much larger stochastic ensemble of forcing combinations, enabling rapid and spatially explicit prediction of water levels along the Bay centerline, including near the Delta mouth. This surrogate approach preserved the accuracy of physics-based simulations while reducing computational cost by orders of magnitude, allowing the authors to quantify extreme water levels, flood frequency, and the relative contribution of individual drivers over centennial timescales.

Building on this framework, Wang et al. [61] applied the same GPR-based surrogate modeling approach to assess compound flooding under future climate change scenarios. The trained ML surrogates were embedded within a stochastic climate–hydrology–hydrodynamics pipeline to evaluate how sea-level rise and increased river discharge jointly affect extreme water levels across the San Francisco Bay-Delta. Results showed that GPR-enabled simulations could resolve spatial gradients in compound flooding, highlighting the growing influence of river discharge on extreme water levels in the North Bay and western Delta under high sea-level rise scenarios, while tides and sea-level rise dominate

farther seaward. Together, these studies demonstrate how Gaussian Process-based ML surrogates can bridge statistical and process-based modeling, enabling computationally tractable, physically consistent assessment of compound flooding risks in complex estuarine–delta systems such as the Delta.

2.4. Critical Gaps

While the past three decades have demonstrated the immense potential of ML in water and environmental modeling for the Delta, the transition from research prototypes and isolated implementations to trusted, Delta-wide operational tools remain uneven. To date, only a limited number of ML applications, most notably the ANN-based salinity emulator embedded within CalSim, have been fully integrated into routine operational workflows (ANNs in CalSim), with most efforts remaining confined to case-specific demonstrations or proof-of-concept studies. This retrospective review reveals four critical gaps that currently hinder the widespread adoption of ML in management workflows. These are illustrated in Figure 7 and described as follows. Section 3 discusses potential strategies for addressing these gaps.

It should be pointed out that these gaps combine generic challenges that affect ML applications in essentially all environmental domains with Delta-specific drivers rooted in California’s multi-agency governance, regulatory complexity, and the Delta’s unique physical and ecological setting. For each gap, we therefore distinguish the portion that reflects the global state of ML practice from the portion that is sharpened or uniquely shaped by the Delta context. This distinction matters because the two portions call for different remedies: generic gaps benefit from adopting tools and norms developed by the broader ML community, whereas Delta-specific gaps require coordinated institutional action among California water agencies, federal partners, and the research community. It should also be noted that although these four critical gaps are presented separately for clarity, they are closely related and often reinforce one another. Table A3 in Appendix A summarizes how each gap affects the principal ML applications reviewed in this paper, the evidence supporting its identification, and priority research directions for future development. This synthesis provides a direct connection between the retrospective review presented earlier in this section and the future research roadmap proposed herein.




Figure 7. Gaps between current and state-of-the-science ML applications in the Delta.

Gap 1: Interoperability

ML development in the Delta is fragmented across agencies and platforms. Different teams use disparate datasets, preprocessing steps, and software tools, so models and data cannot be easily shared or combined. In practice, this lack of common standards and interoperable infrastructure prevents building a cohesive modeling ecosystem and makes benchmarking nearly impossible.

This gap is particularly pronounced in the Delta, where a fragmented landscape of state, federal, academic, and consulting actors has led to diverse data platforms, modeling traditions, and workflows. While the underlying technical challenges—such as heterogeneous data schemas and the absence of common application programming interfaces (APIs)—are broadly applicable across domains, their impact is amplified in the Delta due to this institutional fragmentation. As a result, interoperability barriers are more severe, limiting model reuse, integration, and reproducible evaluation.

Addressing this gap requires moving beyond ad hoc integration efforts toward a coordinated framework for ML development and deployment in the Delta. This includes standardized ML development protocols, shared data schemas, and interoperable modeling infrastructures to enable consistent data exchange, model reuse, and reproducible benchmarking across institutions. These elements collectively form the foundation of a unified modeling ecosystem.

Gap 2: Generalizability

Most Delta ML models are trained using historical observations, limiting their validity under novel or shifting conditions. Climate change and evolving operations introduce unprecedented scenarios (e.g., mega-floods, sea level rise) that lie outside the training record. As a result, purely data-driven models often lose predictive skill when the system departs from past regimes. This lack of robustness undermines confidence in long-term planning and constrains model utility for future conditions.

While this limitation reflects a well-known and broadly studied challenge in ML—namely distribution shift and extrapolation beyond the training domain—it is particularly acute in the Delta. The Delta system is highly non-stationary due to accelerating climate change, sea level rise, shifting hydrologic extremes, and ongoing or potential structural and operational changes (e.g., habitat restoration, levee subsidence, and conveyance re-configuration). These factors jointly create a rapidly evolving physical and operational landscape that is only partially represented in historical datasets, thereby amplifying the risk of model degradation under future conditions.

Addressing this gap requires approaches that extend beyond purely empirical learning, including hybrid physics–ML frameworks, scenario-informed training strategies, and explicit integration of process-based constraints. Improving model robustness and generalizability across diverse hydrodynamics, water quality, and ecological regimes is therefore essential.

Gap 3: Interpretability

Many ML algorithms operate as opaque “black boxes,” providing limited insight into why a particular prediction was made. In a highly regulated system like the Delta, decisions must be scientifically defensible; opaque models undermine stakeholder trust and make regulators hesitant to rely on ML outputs. Without interpretable reasoning, managers are reluctant to use ML forecasts for high-stakes decisions.

While interpretability is a widely recognized challenge in ML more generally, its implications are particularly pronounced in the Delta. This is due to the region’s highly regulated, multi-agency governance structure, frequent legal scrutiny, and adversarial decision environment, where model-based conclusions may be subject to regulatory review, stakeholder challenge, or litigation. In this setting, predictive accuracy alone is insufficient; models must also provide traceable and defensible reasoning grounded in established process understanding.

Addressing this gap requires approaches that enhance transparency without sacrificing predictive performance, including physics-informed ML, hybrid mechanistic–data-driven models, and explainable AI (XAI) [129] techniques that link model outputs to known system drivers and domain knowledge. Ensuring that ML outputs are interpreta-

ble, traceable, and consistent with established scientific understanding is therefore essential for building trust and enabling operational adoption.

Gap 4: Operationalization

Finally, many ML models are not designed for practical deployment. A model can achieve very high accuracy in a controlled or retrospective setting but remains impractical for real-time or operational use if it cannot run within existing operational tools, ingest real-time inputs, or quantify uncertainty. This “accuracy vs. actionability” gap means that most prototypes never get deployed into operations. For example, if an ML salinity model cannot feed into the water planning model CalSim or requires data that is not available in real time, it has little value for operators.

An additional limitation is that many existing ML applications provide deterministic predictions without estimates of predictive confidence or uncertainty. This is particularly important for Delta water and environmental management, where operational decisions must account for uncertainty and risk. Future ML systems should therefore incorporate reliable uncertainty estimates to better support risk-informed and adaptive decision making.

While this challenge reflects a broader issue in machine learning related to deployment maturity, software integration, and real-time data engineering, it is sharpened in the Delta due to the presence of deeply embedded operational modeling infrastructure and formalized statutory compliance requirements. As a result, operational constraints are not simply technical afterthoughts but fundamental design requirements that shape whether ML models can be adopted in practice.

Addressing this gap requires closer coupling between model development and operational workflows, including standardized interfaces with existing hydrodynamic, water quality, and planning models, and robust pipelines for real-time data assimilation and uncertainty quantification.

2.5. When ML Is Not the Ideal Tool

Although the preceding sections have emphasized the demonstrated successes of ML in Delta modeling, ML is not universally appropriate. In the Delta context, purely data-driven approaches are often poorly suited to extrapolation far beyond the historical training envelope, particularly under non-stationary climate conditions, sea-level rise, major infrastructure modifications, or novel flow operations for which little or no observational precedent exists. For example, predicting hydrodynamic and salinity responses under extreme sea-level rise scenarios, or future climate regimes substantially different from the historical record may require physically based process models to preserve governing system dynamics and mass-balance relationships. Similarly, ML approaches may be less reliable for rare, high-consequence events with limited training data, such as extreme levee failures, compound flood events, harmful algal bloom outbreaks under unprecedented conditions, or multi-year drought sequences more severe than historical extremes.

Purely data-driven methods may also be insufficient in regulatory and environmental decision-making contexts where transparent and auditable causal mechanisms are required. For instance, biological opinions, water quality compliance evaluations, or endangered-species management decisions often require mechanistic understanding and traceable cause-and-effect relationships that are difficult to establish using black-box ML models alone. In addition, some Delta applications remain relatively data-poor, particularly for ecological response variables with sparse spatial and temporal observations, making it difficult to train robust ML models without substantial uncertainty or overfitting risk. In such cases, the cost of developing, validating, and maintaining ML systems may exceed the benefits relative to simpler empirical approaches or established process-based models.

Finally, ML may provide limited added value in applications where existing process-based models already achieve acceptable accuracy, interpretability, and operational performance. For example, well-established hydrodynamic and water quality models used for tidal flow, salinity transport, or gate operations may already satisfy operational latency and regulatory requirements while providing physically interpretable outputs. Recognizing these limitations is important to ensuring that ML is deployed in areas where it provides clear advantages, while hybrid and physics-based approaches continue to serve as the foundation for applications where physical realism, extrapolation capability, interpretability, and regulatory defensibility remain essential.

3. Prospects

The gaps identified above underscore the need for Delta modeling to move beyond the ad hoc application of new algorithms. Instead, advancing the use of ML in the Delta will require the development of a standardized, transparent, and multi-modal modeling paradigm. This section outlines a strategic roadmap for transforming ML from a series of isolated research efforts into a cohesive operational toolkit. By adopting rigorous ML developmental protocols and proactively integrating emerging technologies, the Delta community can build a ML-informed management framework that supports balanced decision-making across water supply, water quality, ecosystem health, and infrastructure resilience. In this paper, we use the term protocol to refer to a documented, repeatable procedure for conducting specific ML development activities (e.g., data preparation, model training, validation, documentation, deployment) within a project or organization.

3.1. Implementation of ML Protocols

An essential future prospect is the adoption of standardized ML protocols by the Delta modeling community, including relevant local, state, federal agencies, academia, and the private sector (as highlighted in Section 1.2) to enable more robust, transparent, and operational ML applications that support water and environmental decision-making in the Delta. This directly addresses Gap 1 discussed in the previous section.

As discussed in He et al. [130], while ML applications have become increasingly diverse—ranging from hydrodynamic emulations to ecological risk assessment—the lack of unified ML development protocols presents a significant barrier to their long-term utility. Establishing a common procedural framework is not merely a technical preference but a necessity for ensuring that these data-driven tools can be peer-reviewed, reproduced, and ultimately trusted by the multi-agency stakeholders responsible for Delta governance.

The protocols proposed by He et al. [130] establish a structured lifecycle for ML model development, encompassing four critical phases (Figure A2 of Appendix A). The process begins with problem definition, in which clear management objectives are articulated and performance criteria—such as acceptable thresholds for accuracy, interpretability, and operational relevance—are predefined to ensure that the model directly addresses specific operational and regulatory needs. This is followed by data preparation, which emphasizes standardized data collection, rigorous quality control, and harmonization—an essential step given the heterogeneity and complexity of Delta datasets. The model development phase advocates systematic progression from simple, interpretable models to more sophisticated hybrid or ensemble approaches, with evaluation guided by the predefined performance criteria to determine when sufficient model skill and usability have been achieved. Within this evaluation stage, standardized performance benchmarking—defined as the application of a consistent set of evaluation datasets, metrics, and diagnostic tests (e.g., flow, stage, and salinity error statistics at compliance locations)—is used to enable consistent comparison of model performance and utility across different modeling approaches (e.g., ML-based and process-based). This staged strategy ensures that domain

knowledge is explicitly incorporated, thereby enhancing physical realism and model robustness. Finally, model deployment focuses on comprehensive documentation, hands-on training, open-source practices to promote transparency, reproducibility, and cross-agency collaboration.

Importantly, the framework also highlights the development of user-friendly interactive ML-based tools, enabling practitioners with limited ML expertise to readily apply these models in operational decision-making. For example, interactive ML-based dashboards for Delta salinity [118] and ion constituent simulations [55–57] allow users to specify input conditions, select locations of interest, and choose preferred ML algorithms. Model predictions are then generated and visualized in near real time, facilitating rapid scenario exploration and adaptive management. Figure 8 illustrates an example dashboard that uses ML models to emulate DSM2 in generating salinity across the Delta (https://dwrbdodash.azurewebsites.net/dashboard_v20250918; accessed on 1 April 2026). Users can select the year of interest (Step 1), scale up or down input variables (Step 2), choose a location (Step 3), and click ‘Compute’ (Step 4). The results are generated and displayed within seconds, with an option for users to download them.




Figure 8. ML-Based DSM2 salinity emulation dashboard.

Adopting these standardized steps offers several transformative benefits for Delta water and environmental management. By providing a roadmap for quality assurance, these protocols mitigate common risks such as overfitting and ensure that models perform reliably across varying conditions. Furthermore, standardization facilitates easier comparison and collaboration between different research and application efforts, while interoperability ensures that new ML tools can seamlessly integrate with existing decision-support systems. Following these structured steps will allow for better peer review and cross-agency collaboration, ensuring that future ML applications are as rigorous and reliable as the process-based models they are designed to complement or emulate. Ultimately, this standardization is vital for transforming ML from a series of isolated research exercises into a cohesive, operational toolkit.

3.2. Roadmap of Integrating Emerging Technologies and Multi-Modal Data Sources

In addition to standardizing modeling procedures, another critical future direction is to evaluate the potential of emerging technologies—including XAI [129], Generative AI (GenAI) [131], Digital Twins [132], AI agents [133], mega-learners (stacked ensemble

[134]), and multi-modal modeling [135]—along with the integration of diverse data streams including remote sensing data. As the Delta faces escalating challenges related to water supply reliability, water quality degradation, ecosystem stress, and infrastructure vulnerability, traditional data-driven approaches must evolve to incorporate these advanced computational paradigms and information-rich datasets. By embracing these cutting-edge tools, the Delta community can transition from reactive monitoring toward a more predictive, adaptive, and transparent management framework that better balances water supply reliability with ecosystem health. A conceptual roadmap along these lines is depicted in Figure 9 and discussed in detail below.




Figure 9. Near-future outlook of ML applications in the Delta.

3.2.1. Multi-Modal Data Fusion (Addressing Gap 1: Interoperability)

The first pillar of the roadmap is establishing an integrated data infrastructure: a centralized data hub that harmonizes heterogeneous Delta datasets (remote sensing, in situ sensors, grab samples, etc.) into a common repository. By adopting open data principles (e.g., the FAIR principles [136]) and incorporating emerging protocols, the hub ensures consistent formats and discoverability across datasets [137]. More broadly, the development and adoption of open, community-defined standards would further strengthen these capabilities. In fact, the State maintains an open data portal to host water-related datasets statewide, following California’s Open and Transparent Water Data Act [138]. While the portal contains much of the available Delta data, other important data sources (such as satellite products) need to be accessed from external platforms. Establishing a new centralized data hub or augmenting the current open data portal ensures that all models and tools can access consistent inputs—a remedy for the current fragmentation of data sources. In practice, this involves adopting open standards (per California’s Open Data Act) and cloud-based platforms that automatically ingest satellite imagery (e.g., optical and multispectral imagery of Delta turbidity), sensor telemetry (e.g., continuous flow and salinity observations), and biological indicators (e.g., eDNA sampling that provide direct insight into species distributions and ecosystem conditions) into a shared database.

Multi-modal learning frameworks then leverage advances in deep learning to learn joint representations across these data streams, capturing correlations and interactions that are not apparent when modalities are analyzed in isolation. In Earth observation and environmental science, multimodal graph neural networks and representation learning methods have recently been shown to improve robustness and prediction accuracy by fusing optical, radar, vector, and in situ data into integrated latent spaces that reflect the multi-faceted nature of ecological systems [139–141].

For the Delta, multi-modal data fusion models can combine visual observations of fish behavior (e.g., video from monitoring platforms), chemical and physical water quality measurements (salinity, ion concentrations), and remote sensing indicators to build richer ecosystem state representations than traditional uni-modal models. This multimodal fusion enhances both predictive performance and ecological interpretability, enabling more

nuanced assessment of habitat conditions and stressor impacts. For instance, while current fish entrainment models rely on sparse, in situ turbidity sensors that may miss localized plumes, a multi-modal architecture can integrate satellite imagery (e.g., Sentinel-2) to derive a spatial turbidity field via empirical algorithms to convert reflectance to turbidity. This spatial data can serve as a dynamic input to ML models, allowing the models to adjust their entrainment risk weightings based on the movement of a sediment plume across the Delta. By fusing the temporal density of sensor telemetry with the spatial coverage of remote sensing, we move from point predictions to a spatially aware decision-support system that more accurately reflects the heterogeneous nature of the Delta's aquatic environment.

By enforcing common data definitions and sharing a centralized data hub, this data fusion framework makes outputs interoperable across agencies and tools, directly addressing the interoperability gap.

3.2.2. Delta Digital Twin and Computational Engine (Addressing Gap 2: Generalizability)

Beyond data infrastructure, Digital Twins represent a powerful platform for near-real-time, scenario-based decision support [142,143]. Digital Twins are high-fidelity, dynamic virtual representations of physical systems that are continuously updated with real-time data to reflect the current state of their real-world counterparts [144,145]. The second step of the roadmap involves constructing a high-fidelity digital twin of the Delta—a continuously updated, virtual replica of the hydrodynamic and ecological system. This allows managers to run instantaneous “what-if” simulations, assessing how a specific change in, for example, water export volumes or the installation of a temporary drought barrier would impact salinity, turbidity, and fish migration in real-time. By providing a safe, virtual playground for experimentation, the Delta Digital Twin reduces the risk of unintended consequences in one of the most highly contested water systems in the world.

Complementing the Delta Digital Twin, the roadmap envisions an advanced computational engine that uses ensemble modeling and generative AI to boost ML model robustness and generalizability. Ensemble ML techniques (also referred to as “mega-learners”) can be employed to synthesize insights across multiple multi-task (e.g., targeting Delta flow, salinity, water temperature, ion levels, etc., simultaneously) ML models, reducing bias inherent to any single algorithm and improving robustness across diverse flow regimes and ecological conditions. Ensemble strategies have become increasingly valuable in environmental applications, where complex interactions and non-linear responses are common. Previous ML studies in the Delta have developed multiple ML algorithms that were developed to simulate ion concentrations [55–57], salinity [112,116–118], reference evapotranspiration [121,122], and harmful algal bloom risks [53]. However, in most cases, the outputs of individual algorithms were presented separately to demonstrate model diversity, rather than being formally integrated into a unified ensemble framework. Moving forward, these model outputs could be strategically weighted, and these algorithms can be combined within a mega-learning architecture to explicitly leverage the strength of each algorithm.

In addition to ensemble modeling, GenAI offers an emerging set of tools for scenario augmentation and knowledge synthesis. GenAI involves a class of ML models capable of generating new, synthetic data or content that mimics the statistical properties of their training set. In the Delta context, GenAI holds significant promise for data augmentation, generating high-fidelity synthetic time-series for data-sparse periods or to simulate rare, extreme events that are underrepresented in historical records [146,147]. For instance, a persistent bottleneck in Delta ecological modeling is the class-imbalance problem, where the extreme rarity of endangered species—such as the Delta Smelt—results in datasets dominated by zeros, often causing traditional ML models to fail or bias toward absence

predictions. Rather than relying on simple oversampling, GenAI can be used to synthesize high-stakes environmental scenarios that mirror the rare conditions under which endangered fish are historically detected; however, because generative models trained on imbalanced data may inherit those biases, their use typically involves conditioning or physics-informed constraints to ensure realistic representation of rare events. For example, process-based models can be run under perturbed inputs representing rare or unprecedented conditions to generate physically consistent synthetic datasets, which can then be used to train or guide GenAI to better capture these low-probability but high-impact scenarios. This hybrid workflow shifts GenAI from a purely generative tool toward a complementary component in uncertainty-aware modeling frameworks, improving the ability of decision-support systems to represent rare but consequential environmental conditions even in data-scarce regimes.

Moreover, GenAI also shows potential for document synthesis and policy support. Large language models, a type of GenAI, can summarize extensive technical and policy documents, helping decision makers efficiently distill findings and emerging themes. For instance, a California state agency recently used NotebookLM to synthesize publicly available reports and studies on climate stressors and hazards affecting the Delta, informing internal and external discussions on emerging climate change impacts on the Delta [148]. Such applications demonstrated how GenAI can support evidence-based policymaking when used with appropriate transparency and verification.

Ultimately, the Delta digital twin infrastructure integrates real-time multi-modal data fusion with simulations or forecasting from the advanced computational engine. It enables the development of highly generalizable models capable of navigating unprecedented scenarios that typically confound ML models that are purely trained on historical data.

3.2.3. Explainable Artificial Intelligence (Addressing Gap 3: Interpretability)

To promote trust and accountability in decision-making contexts, XAI should be integrated into both individual and ensemble (mega-learner) ML models. This integration enables stakeholders to understand why a model produces certain outputs, whether for salinity intrusion forecasts, turbidity predictions, or ecological response indicators. Explainability is essential in regulated water systems such as the Delta where transparency underpins both scientific defensibility and stakeholder confidence.

For Delta, a recent study demonstrated the feasibility of incorporating physical interpretability into ML frameworks. Specifically, Namadi et al. [57] assessed the sensitivity of several machine learning models in simulating ionic constituents to different input features. They also developed an interactive dashboard to illustrate model interpretability (https://dwrdashionsensitivity.azurewebsites.net/Sensitivity_IonStudy; accessed on 1 April 2026). Roh et al. [113] embedded salinity transport equations into an ANN network, allowing the model to learn physical constraints while maintaining predictive flexibility.

In addition to physics-informed approaches, post hoc XAI frameworks—such as SHAP (SHapley Additive exPlanations) [149] and LIME (Local Interpretable Model-agnostic Explanations) [150] can be applied to quantify the influence of specific environmental drivers, such as flows, water temperature, and tidal conditions on model predictions. However, the degree of explainability achievable depends on model complexity, and highly flexible models such as deep neural networks may remain only partially interpretable at the process level. In practice, the goal is often sufficient and consistent explanation of model behavior to support stakeholder confidence in Delta decision-making, particularly by identifying dominant drivers and ensuring that model responses align with established domain knowledge.

3.2.4. Decision Support (Addressing Gap 4: Operationalization)

The final pillar of the roadmap focuses on decision support. AI agents could potentially serve as a bridge between ML outputs and operational decision-support by integrating model predictions into decision-support tools, control systems, and human-in-the-loop interfaces. While still an emerging area, AI agents represent a more autonomous paradigm in which ML-driven systems may eventually be able to perceive dynamic conditions, reason across objectives, and suggest actions within regulatory constraints.

In the Delta context, such capabilities could eventually support adaptive real-time operations, including the coordinated operations of the Delta Cross Channel gates and project exports. During high-flow ‘first flush’ events, an AI agent could potentially synthesize real-time acoustic fish tracking and turbidity sensor streams to assist operators in evaluating high-frequency adjustments, such as modifying exports or gate positions in coordination with tidal cycles to reduce fish entrainment while maintaining water supply reliability. This shift toward agent-based management enables an agility that traditional rule-curve-based governance cannot match, particularly in an environment where conditions can change within a single tidal cycle. However, these applications remain largely conceptual and represent a long-term research direction rather than operational capability. Given the high-consequence nature of Delta water management, any future implementation would require rigorous validation, uncertainty quantification, robustness testing, regulatory review, and appropriate human oversight before being considered for operational use (see Appendix B for a detailed discussion of operational deployment considerations). As such, near-term applications are more likely to augment human decision making through decision-support systems rather than enable full autonomous operations.

Crucially, an iterative feedback loop is established: actions taken (and their outcomes) are logged back into the Delta Digital Twin infrastructure, continuously refining ML model accuracy. Over time, the whole framework depicted in Figure 9 becomes embedded in management practice—essentially eliminating the silo between accurate model predictions and actual operations. By automating routine responses, alerting humans for oversight, and integrating model results into agency workflows, this decision–action layer ensures that “better information” truly leads to “better decisions,” addressing the operationalization gap.

3.2.5. Roadmap Timeline

The roadmap illustrates a cohesive Delta ML/AI ecosystem. The unified data platform enables cross-disciplinary interoperability; the digital twin and ensemble engine extend models beyond historical confines; embedded physics and XAI ensure transparent, scientifically grounded outputs; and ML/AI-driven decision-support connects model outcomes to management actions. Together, this roadmap aims to ensure that future ML applications in the Delta are interoperable, robust to change, interpretable to stakeholders, and fully integrated into adaptive management practices.

To realize this integrated technological ecosystem in practice, a phased implementation strategy is advisable. In the near term (0–5 years), focus could be placed on strengthening data infrastructure, improving interoperability across datasets and modeling platforms, deploying ensemble and XAI-augmented forecasting models, and piloting modular multi-modal fusion techniques using existing datasets. In the mid-term (5–10 years), digital twin prototypes and GenAI-augmented scenario analysis frameworks, and more advanced multi-modal modeling systems can be developed and evaluated through iterative stakeholder engagement. Over longer time horizons (10+ years), supervised AI agents and fully interoperable decision support ecosystems incorporating real-time data assimilation and multi-modal model integration may become feasible as governance frameworks and computational infrastructure mature. Aligning these developments with on-

going collaborative modeling and synthesis initiatives within the Delta science community will help ensure that technological innovations translate into tangible improvements in both scientific understanding and water management outcomes.

Implementation of such a roadmap would likely require coordinated participation across the broader Delta modeling community, including local, state, and federal agencies, academia, and the private sector involved in data collection, model development, and decision-making. A centralized coordinating entity or governance framework could play an important role in facilitating communication across organizations, promoting interoperability standards, supporting shared cyberinfrastructure, and coordinating long-term implementation activities. At the same time, responsibilities related to data provision, model development, model application, and operational decision-support are inherently distributed across organizations operating within different mandates and workflows. Consequently, these roles are often overlapping rather than cleanly separable into fixed categories of “developer,” “user,” or “data provider.” For example, agencies may simultaneously function as data providers, model developers, and end users depending on the application context. Given this distributed structure, the objective of the proposed roadmap is not to prescribe institutional responsibilities, but rather to highlight the importance of sustained interagency coordination, shared technical standards, open scientific collaboration, and community-driven governance mechanisms to support the long-term evolution of ML/AI capabilities in the Delta.

4. Summary and Conclusions

This review differs from previous reviews in terms of its geographic focus and methodological scope. Specifically, this review provides a Delta-specific assessment of the evolution of ML applications over the past three decades and outlines future perspectives for implementing ML to support Delta operations and decision-making. Unlike broader reviews that focus on general ML methodologies or water resources applications, this review synthesizes the historical development, current status, and future opportunities of ML within the unique context of the Delta. A comparison of this review with other recent relevant reviews is provided in Table A4 in Appendix A. The remainder of this section summarizes the retrospective analysis (Section 4.1), presents a roadmap for future ML/AI applications in the Delta (Section 4.2), and discusses the importance of collaboration and partnership for successful implementation (Section 4.3).

4.1. Summary of the Retrospective Review

This paper has provided a comprehensive, three-decade retrospective on the evolution of ML applications within the Delta. The journey began in the 1990s with pioneering efforts to utilize ANNs as computationally efficient surrogates for flow-salinity relationships and carriage water estimation. Over the subsequent decade, these tools transitioned from experimental research into core operational components, notably through their integration into the CalSim planning framework. In the most recent era (2010–Present), the landscape has expanded significantly to leverage deep learning architectures and hybrid physics-informed models. These modern approaches have successfully moved beyond traditional salinity management to address a broad spectrum of complex challenges, including fish migration patterns, toxigenic risk from harmful algal blooms, and the structural resilience of the Delta’s levee network. Collectively, this evolution demonstrates that ML can effectively compress the complex physics of the Delta into agile frameworks that enable real-time optimization and large-scale scenario screening.

4.2. Roadmap for Future Delta ML/AI Applications

The roadmap for future ML applications in the Delta focuses on transforming these techniques into a cohesive, standardized, and transparent operational toolkit. Key strategic directions include:

(a) Implementation of Standardized Protocols

Adopting unified ML procedural frameworks for the entire model lifecycle—from problem definition and data preparation to model development and deployment—is essential for ensuring that ML/AI tools in the Delta are reproducible, scientifically defensible, peer-reviewed, and trusted by regulatory and governance agencies.

(b) Multi-Modal Data and Infrastructure

Developing an integrated, interoperable, and centralized Delta data hub is critical for harmonizing diverse data streams, including remote sensing and in situ data. This will support multi-modal modeling approaches capable of providing a more holistic understanding of the Delta's coupled physical, biogeochemical, and ecological processes.

(c) Integration of Emerging Technologies

The Delta community should proactively adopt cutting-edge paradigms such as Explainable AI (XAI) to enhance model interpretability and scientific defensibility, Generative AI for rare-event augmentation and knowledge synthesis, and Delta Digital Twins for near-real-time, “what-if” scenario analysis under changing climatic, structural, operational, and regulatory conditions.

(d) Transition to AI Agents

The long-term evolution of Delta ML/AI systems may involve AI agents capable of supporting adaptive and multi-objective water management by rapidly evaluating operational alternatives and recommending control strategies that balance water reliability with ecological protection in real time.

It is worth noting that several emerging technologies—including XAI, digital twins, generative AI, and AI agents—offer promising opportunities to complement existing ML workflows. However, their readiness levels differ substantially. While XAI and standardized ML workflows are already applicable for many operational needs, technologies such as digital twins, GenAI, and AI agents should currently be viewed as longer-term research opportunities that require rigorous technical validation, governance frameworks, and operational testing before they can be deployed in mission-critical water and environmental management applications in the Delta.

4.3. The Importance of Collaboration and Partnership

The future success of ML in shaping Delta management will depend on strong collaboration and sustained cross-agency partnerships. Establishing standardized modeling protocols is not merely a technical step but a necessity for facilitating easier comparison and cooperation between different research efforts. Furthermore, embracing open-source practices and interactive, user-friendly tools will promote transparency and ensure that ML-driven insights are accessible to practitioners with varied levels of expertise. By working together through shared data infrastructure and unified developmental standards, the Delta modeling community can ensure that these powerful technologies are applied both rigorously and reliably. Such collective action is vital for building a predictive and adaptive management framework capable of navigating the escalating challenges of climate change and regulatory complexity and uncertainty while maintaining the delicate balance between human needs and ecosystem sustainability.

The next thirty years of Delta management will be shaped not just by the volume of available data, but increasingly by the ability to synthesize that information into timely,

effective, and reliable decisions. To achieve this, the Delta modeling community should pivot from isolated ML algorithmic experimentation toward a culture of shared digital infrastructure and standardized validation. Establishing open-source benchmarking datasets and explainable ML protocols is a requirement for trust. By laying these collaborative foundations today, we can transform ML from a tool of limited implementation into a reliable operational framework that guides decisions to support both Delta water supply reliability and ecosystem health in an era of unprecedented changes.

Author Contributions: Conceptualization, M.H. and P.S.; methodology, M.H. and P.S.; validation, M.H. and P.S.; formal analysis, M.H. and P.S.; investigation, M.H. and P.S.; resources, M.H. and P.S.; writing—original draft preparation, M.H.; writing—review and editing, M.H. and P.S.; visualization, M.H.; supervision, P.S.; project administration, P.S. All authors have read and agreed to the published version of the manuscript.

Funding: This study received no external funding.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The authors would like to express their sincere gratitude to their past and present colleagues, particularly the late Ralph Finch, for their vision, ingenuity, diligence, and unwavering support in advancing the development and application of machine learning in Delta modeling over the past three decades. Specifically, the authors would like to thank their colleagues Lucy Andrews, Bryan Nguyen, and Andrew Schwarz for taking the time to review an earlier version of the paper and provide insightful comments that significantly improved its quality. The authors are also grateful to their colleagues Zachary Roy, Malinda Wimalaratne, and Nicole Osorio for their support in creating Figure 5. In addition, the authors are grateful to their colleagues, Erik Reyes and Louise Conrad, for their support during this work. The authors also sincerely thank the four anonymous reviewers and the editor for their thoughtful and constructive comments, which have significantly improved the quality of this work. Generative AI tools were used to create selected icons in Figures 3, 6, 7, 9 and A2. The lead author reviewed these icons and is solely responsible for their content. The views expressed in this paper are those of the authors and do not reflect the views of their employer.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
ANN	Artificial Neural Network
API	Application Programming Interface
BOT	Barrier Optimization Tool
BRT	Boosted Regression Trees
CalSim	California Water Resources Planning Model
CalSim II	California Water Resources Planning Model Version 2
CalSim 3.0	California Water Resources Planning Model Version 3
CatBoost	Categorical Boosting
CFD	Computational Fluid Dynamics
CH ₄	Methane
CNN	Convolutional Neural Network
CO ₂	Carbon Dioxide
CVP	Central Valley Project
DASM-T	Delta Turbidity ANN Model
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DCC	Delta Cross Channel

DCD	Delta Channel Depletion
Delft3D-FM	Delft3D Flexible Mesh Suite
Delta-ERT	Delta Emergency Response Tool
D-Linear	Decomposition Linear
DNN	Deep Neural Network
DOM	Dissolved Organic Matter
DSG	Delta Salinity Gradient Model
DSM2	Delta Simulation Model 2
DT	Decision Tree
DWR	California Department of Water Resources
DWRDSM	DWR Delta Simulation Model
DWRSIM	Original California Water Resources Planning Model
EC	Electrical conductivity
eDNA	Environmental Deoxyribonucleic Acid
ET	Evapotranspiration
ETo	Reference Evapotranspiration
fDOM	Fluorescent Dissolved Organic Matter
FDM	Fischer Delta Model
GA	Genetic Algorithm
GAM	Generalized Additive Model
GB	Gradient Boosting
GenAI	Generative Artificial Intelligence
GPR	Gaussian Process Regression
GPU	Graphics Processing Unit
GRU	Gated Recurrent Unit
GTM	General Transport Model
HAB	Harmful Algal Bloom
LightGBM	Light Gradient Boosting Machine
LIME	Local Interpretable Model-agnostic Explanations
LOSO	Leave-One-Site-Out
LP	Linear Programming
LSTM	Long Short-Term Memory
MBSG	Martinez Boundary Salinity Generator
MEC	Marginal Export Cost
MDO	Minimum Delta Outflow
ML	Machine Learning
MLP	Multi-Layer Perceptron
MSE	Mean Squared Error
MTL	Multi-Task Learning
NARX	Nonlinear Autoregressive with Exogenous Inputs
N-BEATS	Neural Basis Expansion Analysis for Interpretable Time Series
NDO	Net Delta Outflow
N-HiTS	Neural Hierarchical Interpolation for Time Series Forecasting
N-Linear	Normalized Linear
NSE	Nash-Sutcliffe Efficiency
OMR	Old and Middle River
PINN	Physics-Informed Neural Network
QRF	Quantile Regression Forest
RCCAT	Regional Carbon and Climate Analytics Tool
Res-GRU	Residual Gated Recurrent Unit
Res-LSTM	Residual Long Short-Term Memory
ResNet	Residual Network
RF	Random Forest
RMA-2	Resource Management Associates 2D Model
RSR	Standard Deviation Ratio
RT	Regression Tree
SCHISM	Semi-implicit Cross-scale Hydrosience Integrated System Model
SHAP	SHapley Additive exPlanations
SNNS	Neural Network Simulator

SSC	Suspended Sediment Concentration
STL	Single-Task Learning
SVM	Support Vector Machine
SWP	State Water Project
TCN	Temporal Convolutional Network
TFT	Temporal Fusion Transformer
TiDE	Time-series Dense Encoder
TTHM	Total Trihalomethane
XAI	Explainable Artificial Intelligence
XGBoost	Extreme Gradient Boosting

Appendix A. Supplementary Tables and Figures

Table A1. Comparison of empirical and mass-balance, process-based, and machine learning modeling approaches in the Sacramento–San Joaquin Delta of California.

Aspect	Empirical/Mass-Balance Models	Process-Based Models	Machine Learning Models
Definition	Formulated using observed input-output relationships and/or simplified representation of conservation principles	Physically based models that simulate hydrodynamics, transport, and biogeochemical processes using governing equations	Data-driven models that learn nonlinear mappings between inputs and outputs from data without explicit physical equations
Model Structure	Simple; often algebraic or regression-based	Mechanistic; equation-based; spatially and temporally explicit	Flexible architectures
Key Principle	Mass conservation or empirical relationships	Physical laws	Pattern recognition; function approximation
Interpretability	High	Moderate to High	Low to moderate
Data Requirement	Low to moderate	High	High
Runtime	Low	High	Low (once trained)
Strengths	Simple; fast; useful for screening and first-order estimates	Physically consistent; widely used in regulatory and decision-support contexts	Captures nonlinear relationships; emulates or complements process-based models
Limitations	Limited physical realism; weak extrapolation	Parameter-heavy; computationally expensive; may suffer from equifinality	Limited physical interpretability; risk of overfitting; weak extrapolation
Typical Delta Applications	Flow-salinity and salinity-ion relationships; simple mass balance (e.g., DAYFLOW)	Hydrodynamics and transport simulations; ecosystem process modeling (e.g., DSM2)	Flow and salinity modeling; surrogate modeling; pattern detection; hybrid ML-physics frameworks (e.g., CalSim ANN)

Table A1 summarizes the principal modeling approaches discussed throughout this review, including empirical, process-based, and machine learning (ML) methods that have been applied to water and environmental modeling in the Sacramento–San Joaquin Delta (Delta). The table provides a high-level comparison of each modeling framework, including its underlying methodology, computational characteristics, strengths, limitations, and representative Delta applications. Because the reviewed ML methods span a wide range of architectures—from conventional feedforward neural networks to recent hybrid and physics-informed models—a concise comparison of their respective strengths, limitations, and suitable application scenarios is provided below to complement the discussion presented in the main text.

The ML methods reviewed in this paper have evolved from relatively simple feed-forward Artificial Neural Networks (ANNs) to increasingly sophisticated deep learning and physics-informed frameworks, with each approach offering distinct advantages and

limitations. Conventional ANNs, particularly Multi-Layer Perceptrons (MLPs), have proven highly effective as computationally efficient surrogate models for operational planning because they require relatively modest computational resources, are straightforward to implement, and perform well when representative training datasets are available. However, they rely on manually selected antecedent inputs to represent temporal memory and generally have limited ability to capture long-term temporal dependencies. Recurrent neural network architectures, such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models, address this limitation by explicitly learning temporal relationships from sequential data, making them particularly suitable for forecasting variables such as salinity, water temperature, and suspended sediment that exhibit strong seasonal and hydrodynamic memory. These models, however, typically require substantially larger training datasets, greater computational resources, and more extensive hyperparameter tuning, while their internal decision-making processes remain relatively difficult to interpret.

Tree-based ensemble methods, including Random Forest (RF) and Extreme Gradient Boosting (XGBoost), provide an attractive alternative for tabular environmental datasets because they are generally robust to noisy observations, require relatively little data pre-processing, naturally provide measures of feature importance, and often perform well with moderate-sized datasets. Their principal limitation is that they do not explicitly represent temporal dependencies unless lagged variables or other temporal features are incorporated during feature engineering. More recently, hybrid modeling approaches and Physics-Informed Neural Networks (PINNs) have sought to combine the complementary strengths of process-based and data-driven models by incorporating governing equations, process-model outputs, or physically meaningful constraints into the learning process. These approaches generally improve physical consistency, extrapolation capability, and robustness under data-sparse conditions, making them particularly attractive for climate change assessments and scenario analysis. However, they are also more computationally demanding, require expertise in both numerical modeling and ML, and remain largely at the research or pilot stage within Delta applications.

Overall, the studies reviewed in this paper suggest that no single ML approach is universally superior. Rather, the selection of an appropriate model should be guided by the intended application, data availability, computational constraints, and the desired balance among predictive accuracy, interpretability, computational efficiency, and physical realism. Conventional ANNs remain well suited for rapid surrogate modeling and operational planning, tree-based methods are particularly effective for structured environmental datasets where interpretability is important, recurrent deep learning architectures are preferable when long-term temporal dependencies dominate system behavior, and hybrid or physics-informed approaches offer the greatest promise for applications requiring physically consistent predictions under changing environmental conditions.

Table A2 provides a comprehensive inventory of ML applications reported for the Delta through 2025. The inventory was compiled from the systematic literature search described in Section 2 and supplemented by the authors' domain knowledge to capture relevant Delta-specific journal articles, conference proceedings, technical reports, and other publicly available studies. For each study, the table summarizes the publication year, authors, institutional affiliation, target variable(s), ML architecture(s), data and code availability, and operational adoption status. Together, these attributes document the evolution of ML applications in the Delta over the past three decades—from early ANN-based surrogate models for salinity prediction in the 1990s to recent applications involving deep learning, physics-informed ML, ensemble learning, and transformer-based architectures across a much broader range of water and environmental variables. The inventory also illustrates increasing collaboration among state agencies, federal agencies, academia, and private or-

ganizations, gradual improvements in data and code accessibility, and the transition of selected ML applications from research and pilot studies toward operational deployment.

Table A3 provides a synthesis of the four critical gaps identified in this review, linking each gap to representative ML applications in the Delta, the supporting evidence summarized from the literature, and corresponding priority research pathways. Although the four gaps—interoperability, generalizability, interpretability, and operationalization—are discussed separately in the main text (Section 2.4), they are closely interconnected and often influence one another. The table is intended to provide readers with a concise summary of how the retrospective findings presented throughout the review translate into a forward-looking research roadmap for advancing ML in Delta water and environmental management.




Figure A1. Hierarchical relationships among Artificial Intelligence (AI), Machine Learning (ML), Deep Learning (DL), and Generative AI (GenAI). ML is a subset of AI, DL is a subset of ML, and GenAI typically builds on DL methods to generate new content.

Figure A1 illustrates the hierarchical relationships among artificial intelligence (AI), machine learning (ML), deep learning (DL), and generative artificial intelligence (GenAI). AI is the broadest field encompassing computational methods that emulate aspects of human intelligence. ML is a subset of AI that enables models to learn patterns from data without explicit programming. DL is a specialized subset of ML that employs multi-layer neural networks to automatically learn hierarchical feature representations from large datasets. GenAI typically builds upon advanced DL architectures to generate new content, such as text, images, code, or synthetic data.

Table A2. List of machine learning studies in the Sacramento–San Joaquin Delta of California from 1995 to 2025 *.

Year	Author(s)	Affiliation(s)	Target Variable(s)	Architecture(s)	Data/Code Open Source?	Adoption Status
1995	Sandhu and Finch [86]	State Agency	Salinity (EC, Chloride)	MLP	No	Pilot
1995	Hutton and Sandhu [103]	State Agency	TTHM	MLP	No	Pilot
1995	Sandhu and Finch [87]	State Agency	Salinity (TDS)	MLP	No	Pilot
1996	Sandhu [100]	State Agency	Salinity (TDS)	MLP	No	Pilot
1997	Sandhu and Wilson [101]	State Agency	Salinity (TDS)	MLP	No	Pilot
1998	Wilson [102]	State Agency	Salinity (EC)	MLP	No	Pilot
1999	Chung et al. [91]	State Agency	Salinity (EC)	MLP	No	Pilot
2000	Pranger [104]	State Agency	Salinity (EC)	MLP	No	Pilot
2001	Wilbur and Munevar [107]	State Agency	Salinity (EC)	MLP	No	Pilot
2001	Hutton and Seneviratne [108]	State Agency	Salinity (EC)	MLP	No	Pilot
2001	Rajkumar and Johnson [105]	Private and Academia	Salinity (EC)	MLP	No	Research
2002	Rajkumar and Thompson [106]	Federal Agency and Private	Salinity (EC)	GA-MLP	No	Research
2002	Mierzwa [109]	State Agency	Salinity (EC)	MLP	No	Pilot
2002	Senevirante [110]	State Agency	Salinity (EC)	MLP	No	Pilot
2007	Seneviratne and Wu [111]	State Agency	Salinity (EC)	MLP	No	Operational
2014	Chen and Roy [48–50]	Local Agency and Private	Turbidity	MLP, NARX	No	Pilot
2017	Rath et al. [114]	Private	Salinity (EC) and X2	Hybrid Empirical-Bayesian MLP	Partial (Data)	Research
2018	Chen et al. [120]	Private	Volumetric Fingerprints	MLP	Partial (Data)	Research
2020	Jayasundara et al. [47]	State Agency	Salinity (EC)	MLP	Partial (Data)	Operational
2020	He et al. [112]	State Agency	Salinity (EC)	MLP, LSTM	Yes	Pilot
2021	Qi et al. [115]	State Agency and Academia	Salinity (EC)	CNN-MLP	Yes	Pilot
2021	Olivetti et al. [59]	Federal Agency and Academia	Migratory Behavior of Chinook Salmon	LSTM	Yes	Research
2021	RMA [63]	State Agency and Private	Barrier Strategies	LightGBM	No	Pilot
2022	Tillotson et al. [126]	Local Agency and Private	Salmon/Steelhead Entrainment	QRF	No	Research
2022	Qi et al. (a) [116]	State Agency and Academia	Salinity (EC)	MLP, LSTM, GRU, ResNet	Yes	Pilot
2022	Namadi et al. [55]	State Agency	Ion Constituents	GAM, RT, RF, MLP	Yes	Pilot
2022	Kim et al. [51]	State Agency	Suspended Sediment Concentration	LSTM	Yes	Pilot

2022	Qi et al. (b) [117]	State Agency and Academia	Salinity (EC)	MLP, LSTM, GRU, ResNet, Res-LSTM, Res-GRU	Yes	Pilot
2022	He et al. [125]	State Agency	Water Temperature	MLP	Yes	Pilot
2023	Roh et al. [113]	State Agency and Academia	Salinity (EC)	PINN	Yes	Pilot
2023	Qi et al. [118]	State Agency and Academia	Salinity (EC)	MLP, LSTM, GRU, ResNet, Res-LSTM, Res-GRU	Yes	Pilot
2023	Namadi et al. [56]	State Agency	Ion Constituents	RT, RF, GB MLP	Yes	Pilot
2023	Swyers et al. [58]	Federal Agency	Barrier Configurations	GA	No	Operational
2024	Ahmadi et al. [121]	Academia	Reference Evapotranspiration	CatBoost, DNN, LightGBM, RF, XGBoost	Yes	Research
2024	Russo and Athanasopoulos-Zekkos [62]	Academia	Levee Subsurface	K-means Clustering, DBSCAN Clustering	No	Research
2024	Saha et al. [53]	State Agency	Algal Bloom	RF, DT, XGBoost, MLP	Yes	Pilot
2024	Wang et al. [60]	Federal Agency and Academia	Water Level	GPR	Yes	Research
2025	Avouris et al. [52]	Federal Agency and Academia	DO	RF	Yes	Research
2025	Ahmadi et al. [122]	State Agency and Academia	Reference Evapotranspiration	LightGBM, LSTM, N-BEATS, TCN, Transformer, TFT, N-HiTS, TiDE, N-Linear, D-Linear	Yes	Research
2025	Brereton et al. [128]	State Agency and Academia	CO ₂ and CH ₄ Emissions	RF, XGBoost, LightGBM, SVM, GRU, LSTM	Yes	Research
2025	Framsted et al. [54]	Federal Agency	Algal Bloom	RF	Yes	Pilot
2025	Namadi et al. [57]	State Agency	Ion Constituents	MLP	Yes	Pilot
2025	Tillotson et al. [127]	Private	<i>Hypomesus transpacificus</i>	BRT	No	Research
2025	Wang et al. [61]	Federal Agency and Academia	Water Level	GPR	Yes	Research

Note: * All acronyms used in this table are defined in the Abbreviations section preceding Appendix A.

Table A3. Synthesis of critical gaps, supporting evidence, and future research priorities for advancing machine learning applications in the Sacramento–San Joaquin Delta.

Critical Gap	Representative Delta Applications	Evidence from This Review	Priority Research Pathways
Interoperability	DSM2 surrogate modeling [47]	Most ML models have been developed for individual applications and remain loosely coupled with process-based models and operational decision-support systems. Only a limited number of studies have demonstrated integrated frameworks,	Develop standardized model interfaces, hybrid physics–ML frameworks, interoperable data pipelines, cloud-based modeling services, and digital twin capabilities to enable seamless integration across models and agencies.

		such as DSM2 surrogate models embedded within CalSim using ANNs.	
Generalizability	Salinity forecasting [115–118], harmful algal bloom (HAB) prediction, ion constituent prediction [55,56], suspended sediment prediction [51], evapotranspiration forecasting [122]	Many models are trained for specific locations, hydrologic periods, or historical conditions, limiting their applicability under climate change, sea-level rise, and changing management conditions. Recent advances in transfer learning and physics-informed neural networks (PINNs) provide promising pathways but remain limited in operational Delta applications.	Develop transfer learning [151], domain adaptation, synthetic training datasets, continual learning frameworks, and domain-specific foundation models to improve robustness across locations, time periods, and future scenarios.
Interpretability	Salinity forecasting [113], ion constituent prediction [57]	Many early ANN and deep learning models operate as black-box approaches. Only a limited number of recent studies have incorporated explainable AI (e.g., SHAP and feature importance analysis), sensitivity analysis, or physics-informed constraints to improve transparency and scientific understanding.	Expand explainable AI, causal inference, physics-informed learning, sensitivity analysis, and uncertainty quantification approaches to improve model transparency, reliability, and stakeholder confidence.
Operationalization	CalSim 3.0 ANN [47], ion constituent forecasting dashboard [57], fish types barrier design [58]	Only a small number of ML models have progressed to pilot or operational implementation. Most remain research prototypes due to challenges associated with model maintenance, governance, validation, software integration, reproducibility, and uncertainty assessment.	Establish standardized ML lifecycle management, continuous performance monitoring, model governance frameworks, uncertainty-aware forecasting, automated retraining, reinforcement learning–based adaptive decision-support and optimization [152–154], cloud-based deployment, and operational evaluation protocols.

Table A4. Comparison with Recent Reviews on Machine Learning Applications in Water Resources.

Study	Scope and Review Approach	Main Contribution	Distinction/Gap Addressed by This Study
Saha and Pal (2024) [155]	State-of-the-art review of ML and emerging remote sensing across hydrological applications	Links ML with remote sensing for hydrological processes, extremes, groundwater, and water quality	Limited treatment of operational deployment, governance, and validation protocols
Tripathy and Mishra (2024) [156]	Conceptual and methodological review of DL in hydrology and water resources	Synthesizes DL architectures, physics-guided learning, explainability, applications, and research directions	Emphasizes methods more than end-to-end operational workflows in regulated systems
Ahmed et al. (2024) [157]	Broad review of ML applications in water-resources management	Covers water supply, water treatment, water drainage, water quality, hydropower, and marine energy	Limited focus on agency adoption, model governance, and integration with established decision-support systems

Zhi et al. (2024) [158]	Focused review of DL for water-quality prediction	Synthesizes DL methods for water-quality variables and environmental processes	Does not address the wider portfolio of hydrodynamic, ecological, infrastructure, and planning applications
Willard et al. (2024) [159]	Survey of ML time-series prediction at unmonitored sites	Organizes transfer learning and spatial-generalization methods for data-limited locations	Transferability focus is narrower than operational integration, governance, and multi-objective management
Nie et al. (2025) [160]	Bibliometric review of ML/DL research in hydrology	Quantifies publication trends, influential topics, and the evolution of application areas	Provides limited implementation-level guidance for validation, deployment, and institutional adoption
Lokman et al. (2025) [161]	Focused review of ML and statistical methods for water-quality forecasting and classification, with emphasis on hybrid models and Malaysia	Compares model performance, strengths, and limitations; integrates residual analysis, principle component analysis, feature importance, standards, and real-world applications	Specialized in water quality; limited coverage of hydrodynamics, allocation, ecological and infrastructure applications, and transition into agency operations
Torres González et al. (2026) [162]	PRISMA 2020 systematic review (2010–2025; 18 included studies) of predictive ML for streamflow, drought, floods, and water management	Compares deep, ensemble, shallow, and ML–physical hybrid models; evaluates validation, reproducibility, uncertainty, and anthropogenic factors	Broad cross-study evidence synthesis but not a long-term, place-based account of model evolution, process-model integration, governance, and operational deployment
This study	Thirty-year, place-based retrospective and forward roadmap for ML in the Sacramento–San Joaquin Delta	Connects historical model evolution with current hydrodynamic, water-quality, ecological, infrastructure, and planning applications	Bridges research and operations through protocols for robustness, uncertainty, interpretability, process-model integration, deployment, governance, and human oversight

ANNs, one of the most widely used ML models reviewed in this paper, are trained through an iterative optimization process. During training, input features (e.g., flows, gate position) are fed into the input layer, where each neuron corresponds to a specific feature. Each neuron in the input layer represents an input feature. Signals are then transmitted to the next layer, where incoming values are multiplied by connection weights and summed together. A bias term may be added to this sum. The aggregated value is then passed through an activation function (e.g., sigmoid) to introduce nonlinearity, allowing the network to capture complex relationships. The outputs of the activation function serve as inputs for neurons in the next layer. This process continues sequentially until the final layer produces a prediction. This process is referred to as forward propagation. Once predictions are generated, a loss function is applied to measure discrepancies between predicted outputs and observed targets. Backpropagation then computes partial derivatives (i.e., gradients) of the loss with respect to model parameters including weights and biases. These gradients are calculated sequentially from the output layer back toward the input layer. The parameters are subsequently adjusted to minimize the loss via optimization algorithms (e.g., gradient descent [97]). This cycle of forward propagation, loss calculation, back propagation, and parameter optimization process are repeated until convergence or until a predefined stopping condition is reached (Figure 4).




Figure A2. Lifecycle of a typical machine learning model.

Appendix B. Operational Considerations: Validation, Generalizability, Uncertainty, and Transferability

As ML transitions from research to operational deployment, considerations beyond predictive accuracy become increasingly important. The studies reviewed in this paper demonstrate that rigorous model validation has become standard practice for many operational applications in the Delta. Early ANN surrogate models were evaluated using independent testing datasets and comprehensive “full-circle” comparisons against DSM2 simulations before integration into CalSim, while more recent studies have adopted independent validation periods, cross-validation, and benchmarking against established process-based models. These approaches provide confidence that ML models can reproduce historical system behavior under conditions represented in the training data.

However, operational deployment requires demonstrating model robustness under conditions that differ from the historical training period. Future Delta ML applications would therefore benefit from a more comprehensive validation framework that extends

beyond conventional train–test evaluation. Recommended validation practices include out-of-distribution validation, in which models are evaluated under hydrologic, meteorological, or water-quality conditions that fall outside the range of the training dataset; extreme-event testing, using major droughts, floods, heat waves, harmful algal bloom outbreaks, or other rare events to assess model performance during high-impact conditions; rolling-origin validation, which repeatedly trains models using historical data and evaluates them on future time periods to better represent operational forecasting; and independent-event validation, where entire hydrologic events or management scenarios are withheld during model development and used only for final evaluation. Together, these approaches provide a more realistic assessment of model reliability under operational conditions than random data partitioning alone.

Climate change introduces an additional challenge because future environmental conditions may differ substantially from those represented in historical observations. Consequently, climate-change stress testing should become an integral component of operational ML model evaluation. This may include applying trained models to climate projection datasets generated from downscaled global climate models, evaluating robustness under synthetic extreme scenarios, or coupling ML models with process-based simulations representing future hydrologic, sea-level rise, or operational conditions. Such analyses help identify the limits of model applicability and determine when retraining, transfer learning, or hybrid physics-informed approaches may be required.

Uncertainty quantification has likewise received increasing attention but remains an active area of research. Most existing ML applications in the Delta focus primarily on deterministic predictions, whereas many water management decisions—including water quality compliance, water supply reliability, endangered species protection, flood risk management, and infrastructure resilience—are inherently risk-based and require information on forecast confidence rather than a single predicted value. Emerging approaches such as Bayesian neural networks, empirical Bayesian frameworks, deep ensembles, quantile regression, quantile regression forests, conformal prediction, Monte Carlo simulations, probabilistic forecasting, and model interpretability techniques have demonstrated considerable potential for characterizing prediction uncertainty while improving model transparency and robustness. Nevertheless, comprehensive uncertainty quantification and propagation through operational decision-support systems remain relatively limited. Future operational implementations should more systematically quantify uncertainties arising from observations, input forcing, model structure, parameter estimation, and future climate scenarios, and should provide probabilistic predictions or confidence intervals whenever feasible. Integrating uncertainty quantification into operational ML systems will enhance their reliability, transparency, and value for adaptive water and environmental management under changing climatic and hydrologic conditions.

Finally, although this paper focuses on applications within the Delta, many of the methodologies reviewed are broadly transferable to other regulated river basins and estuarine systems. Techniques such as surrogate modeling, deep learning, transfer learning [151], physics-informed ML, and reinforcement learning [152] are largely model-agnostic and can be adapted wherever sufficient observational or simulation data are available. Among these, transfer learning has emerged as a particularly promising strategy for improving model generalizability under data-sparse conditions. Rather than training a model entirely from scratch, transfer learning leverages knowledge learned from related tasks or data-rich domains and adapts it to a new target application through fine-tuning. Recent studies have demonstrated that pre-training deep learning models using large synthetic datasets generated by high-fidelity process-based models, followed by fine-tuning with limited field observations, can substantially improve predictive skill while reducing data requirements [118,119]. Such hybrid approaches offer an effective pathway for ex-

tending machine learning models to new monitoring locations, future climate conditions, and emerging environmental variables where historical observations remain limited.

However, direct transfer of trained models between regions is generally inappropriate because hydrologic processes, infrastructure, operational rules, monitoring networks, and environmental conditions vary substantially among watersheds. Consequently, scenario transferability assessments should evaluate the extent to which model performance is maintained when applied to new geographic regions, operational settings, or future management scenarios. Such assessments may include cross-basin validation, domain adaptation, transfer learning, benchmarking against local process-based models, and comparisons across multiple climate and management scenarios. In practice, transferability is better viewed as the transfer of modeling frameworks, methodologies, and learned representations rather than the direct application of site-specific trained models.

Beyond improving predictive models, reinforcement learning [152] offers a promising pathway toward operational ML applications by enabling adaptive decision-making, dynamic control, and optimization under changing hydrologic and management conditions. Reinforcement learning has attracted growing interest in reservoir operations, water allocation, irrigation scheduling, and adaptive water management under uncertainty. In particular, multi-agent reinforcement learning [153] enables multiple autonomous decision-makers to cooperatively or competitively optimize system-wide performance. Representative studies have demonstrated multi-agent Q-learning-based approaches for urban water allocation [154], in which individual water users or management districts are represented as intelligent agents that collectively learn adaptive allocation strategies while balancing competing objectives. Although no published applications of multi-agent reinforcement learning have yet been reported for the Delta, the underlying concepts are well aligned with the distributed and multi-objective nature of Delta water management. Future research could integrate multi-agent reinforcement learning with fast surrogate models of DSM2, SCHISM, or CalSim3 to support adaptive reservoir operations, Delta exports, environmental flow management, and coordinated water allocation under changing hydrologic, climatic, and regulatory conditions. Together, advances in transfer learning and reinforcement learning broaden the machine learning toolbox available for future water-resources decision-support systems. Whereas transfer learning focuses on improving predictive model generalization across domains, reinforcement learning focuses on adaptive sequential decision-making under uncertainty.

References

1. International Monetary Fund (IMF). World Economic Outlook Databases. October 2025. Available online: <https://www.imf.org/en/publications/sprolls/world-economic-outlook-databases> (accessed on 1 January 2026).
2. U.S. Bureau of Economic Analysis (BEA). Gross Domestic Product by State. 2025. Available online: <https://www.bea.gov/data/gdp/gdp-state> (accessed on 1 January 2026).
3. Public Policy Institute of California (PPIC). California's Population: Just over 39 Million and Still the Most Populous State (California Demographic Facts). 2024. Available online: <https://www.ppic.org/publication/californias-population/> (accessed on 1 January 2026).
4. Luoma, S.N.; Dahm, C.N.; Healey, M.; Moore, J.N. Challenges facing the Sacramento–San Joaquin Delta: Complex, chaotic, or simply cantankerous? *San Franc. Estuary Watershed Sci.* **2015**, *13*. <https://doi.org/10.15447/sfew.2015v13iss3art7>.
5. Dettinger, M.D.; Ralph, F.M.; Das, T.; Neiman, P.J.; Cayan, D.R. Atmospheric Rivers, Floods and the Water Resources of California. *Water* **2011**, *3*, 445–478. <https://doi.org/10.3390/w3020445>.
6. Becker, L.; Yeh, W.; Fults, D.; Sparks, D. Operations models for central valley project. *J. Water Resour. Plan. Manag. Div. Am. Soc. Civ. Eng.* **1976**, *102*, 101–115.
7. Sabet, M.H.; Coe, J.Q. Models for water and power scheduling for the California State Water Project. *J. Am. Water Resour.* **1986**, *22*, 587–596.
8. Myers, N.; Mittermeier, R.A.; Mittermeier, C.G.; Da Fonseca, G.A.; Kent, J. Biodiversity hotspots for conservation priorities. *Nature* **2000**, *403*, 853.

9. Moyle, P.B.; Brown, L.R.; Durand, J.R.; Hobbs, J.A. Delta smelt: Life history and decline of a once-abundant species in the San Francisco Estuary. *San Franc. Estuary Watershed Sci.* **2016**, *14*. <https://doi.org/10.15447/sfew.2016v14iss2art6>.
10. Healey, M.; Dettinger, M.; Norgaard, R. Perspectives on Bay–Delta Science and Policy. *San Franc. Estuary Watershed Sci.* **2016**, *14*. <https://doi.org/10.15447/sfew.2016v14iss4art6>.
11. CSWRCB. *Water Right Decision 1641*; CSWRCB: Sacramento, CA, USA, 1999; p. 225.
12. USFWS. *Formal Endangered Species Act Consultation on the Proposed Coordinated Operations of the Central Valley Project (CVP) and State Water Project (SWP)*; USFWS: Sacramento, CA, USA, 2008; p. 410.
13. DSC. *The Delta Plan*; Delta Stewardship Council: Sacramento, CA, USA, 2013.
14. Hanak, E. *Managing California's Water: From Conflict to Reconciliation*; Public Policy Institute of California: San Francisco, CA, USA, 2011.
15. Lund, J.R. California's agricultural and urban water supply reliability and the Sacramento–San Joaquin Delta. *San Franc. Estuary Watershed Sci.* **2016**, *14*. <https://doi.org/10.15447/sfew.2016v14iss3art6>.
16. Jassby, A.D.; Cloern, J.E.; Cole, B.E. Annual primary production: Patterns and mechanisms of change in a nutrient-rich tidal ecosystem. *Limnol. Oceanogr.* **2002**, *47*, 698–712.
17. Jassby, A. Phytoplankton in the upper San Francisco Estuary: Recent biomass trends, their causes, and their trophic significance. *San Franc. Estuary Watershed Sci.* **2008**, *6*. <https://doi.org/10.15447/sfew.2008v6iss1art2>.
18. Beck, M.W.; Jabusch, T.W.; Trowbridge, P.R.; Senn, D.B. Four decades of water quality change in the upper San Francisco Estuary. *Estuar. Coast. Shelf Sci.* **2018**, *212*, 11–22.
19. Kudela, R.M.; Howard, M.D.; Monismith, S.; Paerl, H.W. Status, trends, and drivers of harmful algal blooms along the freshwater-to-marine gradient in the San Francisco Bay–Delta system. *San Franc. Estuary Watershed Sci.* **2023**, *20*. <https://doi.org/10.15447/sfew.2023v20iss4art6>.
20. Bouma–Gregson, K.; Bosworth, D.H.; Flynn, T.M.; Maguire, A.; Rinde, J.; Hartman, R. Delta Blue(green)s: The Effect of Drought and Drought-Management Actions on Microcystis in the Sacramento–San Joaquin Delta. *San Franc. Estuary Watershed Sci.* **2024**, *22*. <https://doi.org/10.15447/sfew.2024v22iss1art2>.
21. Quiñones, R.M.; Moyle, P.B. Climate change vulnerability of freshwater fishes of the San Francisco Bay area. *San Franc. Estuary Watershed Sci.* **2014**, *12*. <https://doi.org/10.15447/sfew.2014v12iss3art3>.
22. Moyle, P.B.; Lund, J.R.; Bennett, W.A.; Fleenor, W.E. Habitat variability and complexity in the upper San Francisco Estuary. *San Franc. Estuary and Watershed Sci.* **2010**, *8*. <https://doi.org/10.15447/sfew.2010v8iss3art1>.
23. Herbold, B.; Bush, E.; Castillo, G.; Colombano, D.; Hartman, R.; Lehman, P.; Mahardja, B.; Sommer, T. Climate change impacts on San Francisco estuary aquatic ecosystems: A review. *San Franc. Estuary Watershed Sci.* **2022**, *20*. <https://doi.org/10.15447/sfew.2022v20iss2art1>.
24. , J.E.; Knowles, N.; Brown, L.R.; Cayan, D.; Dettinger, M.D.; Morgan, T.L.; Schoellhamer, D.H.; Stacey, M.T.; Van der Wegen, M.; Wagner, R.W.; et al. Projected evolution of California's San Francisco Bay-Delta-River system in a century of climate change. *PLoS ONE* **2011**, *6*, e24465.
25. Mount, J.; Twiss, R. Subsidence, sea level rise, and seismicity in the Sacramento–San Joaquin Delta. *San Franc. Estuary Watershed Sci.* **2005**, *3*.
26. California Department of Water Resources (DWR) *Variation and Control of Salinity*; California Division of Water Resources Bulletin: Sacramento, CA, USA, 1931.
27. California Department of Water Resources (DWR). Dayflow Documentation, Sacramento, CA, USA. 1978. Available online: <https://water.ca.gov/Programs/Integrated-Science-and-Engineering/Compliance-Monitoring-And-Assessment/Dayflow-Data> (accessed on 1 October 2025).
28. Barnes, G.W., Jr.; Chung, F.I. Operational planning for California water system. *J. Water Resour. Plan. Manag.* **1986**, *112*, 71–86.
29. Chung, F.I.; Archer, M.C.; DeVries, J.J. Network flow algorithm applied to California aqueduct simulation. *J. Water Resour. Plan. Manag.* **1989**, *115*, 131–147.
30. Munévar, A.; Chung, F.I. Modeling California's Water Resource Systems with CALSIM. In *Proceedings of the 29th Annual Water Resources Planning and Management Conference*; American Society of Civil Engineers (ASCE): Tempe, AZ, USA, 1999.
31. DWR; USBR. *Benchmark Studies Assumption*; California Department of Water Resources (DWR): Sacramento, CA, USA; U.S. Department of the Interior, Bureau of Reclamation (USBR): Sacramento, CA, USA, 2002.
32. Draper, A.J.; Munévar, A.; Arora, S.K.; Reyes, E.; Parker, N.L.; Chung, F.I.; Peterson, L.E. CalSim: Generalized model for reservoir system analysis. *J. Water Resour. Plan. Manag.* **2004**, *130*, 480–489.

33. DWR and USBR. *A Water Resources System Planning Model for State Water Project (SWP) & Central Valley Project (CVP)*; CalSim 3.0 Draft Report; Sacramento, CA, USA, 2017.
34. California Department of Water Resources (DWR). Salinity Predictions. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: 1st Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1979.
35. California Department of Water Resources (DWR). The Fischer Delta Model. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Tenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1988.
36. California Department of Water Resources (DWR). Calibration and Verification of DWRDSM. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twelfth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1991.
37. Nader-Tehrani, P.; Rajbhandari, H.; Mayr, S. New Model Development: DSM2. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Fifteenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1994.
38. DeLong, L.L. Mass conservation: 1-D open channel flow equations. *J. Hydraul. Eng.* **1989**, *115*, 263–269.
39. Nader-Tehrani, P. Application of Four Point model to the Sacramento-San Joaquin Delta. In *National Conference on Hydrologic Engineering*; Proceedings: American Society of Civil Engineering (ASCE): San Francisco, CA, USA, 1993; pp. 442–447.
40. DeLong, L.L.; Thompson, D.B.; Lee, J.K. *The Computer Program FourPt (Version 95.01)—A Model for Simulating One-Dimensional, Unsteady, Open-Channel Flow*; U.S. Geological Survey Water-Resources Investigations Report 97-4016; US Geological Survey: Sacramento, CA, USA, 1997; 69p.
41. Jobson, H.E.; Schoellhamer, D.H. *Users Manual for a Branched Lagrangian Transport Model*; U.S. Geological Survey Water-Resources Investigations 87-4163; US Geological Survey: Sacramento, CA, USA, 1986; 80p.
42. Jobson, H.E. *Enhancements to the Branched Lagrangian Transport Modeling System*; USGS Water Resources Investigation Report 97-4050; US Geological Survey: Sacramento, CA, USA, 1997.
43. Zhang, Y.J.; Ateljevich, E.; Yu, H.C.; Wu, C.H.; Yu, J.C. A new vertical coordinate system for a 3D unstructured-grid model. *Ocean Model.* **2015**, *85*, 16–31.
44. Ateljevich, E.; Nam, K.; Liu, L.; Saha, S.; Wang, R.; Zhang, J. Bay-Delta SCHISM model developments and applications. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh. 36th Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2007.
45. Zhang, Y.J.; Ye, F.; Stanev, E.V.; Grashorn, S. Seamless cross-scale modeling with SCHISM. *Ocean Model.* **2016**, *102*, 64–81.
46. Chung, F.I.; Seneviratne, S.A. Developing artificial neural networks to represent salinity intrusions in the Delta. In *Proceedings of the World Environmental and Water Resources Congress 2009: Great Rivers*, Kansas City, MO, USA, 17–21 May 2009; pp. 1–10.
47. Jayasundara, N.C.; Seneviratne, S.A.; Reyes, E.; Chung, F.I. Artificial neural network for Sacramento-San Joaquin Delta flow-salinity relationship for CalSim 3.0. *J. Water Resour. Plan. Manag.* **2020**, *146*, 04020015.
48. Chen, L.; Roy, S.B. *DASM-T: Delta ANN Simulation Model for Turbidity, Phase 1 Results*; Report prepared for the Metropolitan Water District of Southern California; Metropolitan Water District of Southern California: Los Angeles, CA, USA, 2012.
49. Chen, L.; Roy, S.B. *Delta Turbidity ANN Model (DASM-T) Development Using DSM2: Phase 2 Results*; Report prepared for the Metropolitan Water District of Southern California; Metropolitan Water District of Southern California: Los Angeles, CA, USA, 2013.
50. Chen, L.; Roy, S.B. *Delta Turbidity ANN Model (DASM-T) Development Using DSM2: Phase 3 Results*; Final Report to Metropolitan Water District of Southern California; Metropolitan Water District of Southern California: Los Angeles, CA, USA, 2014.
51. Kim, H.S.; He, M.; Sandhu, P. Suspended sediment concentration estimation in the Sacramento-San Joaquin Delta of California using long short-term memory networks. *Hydrol. Process.* **2022**, *36*, e14694.
52. Avouris, D.M.; Hestir, E.L.; Fleck, J.; Hansen, J.A.; Bergamaschi, B.A. An integrated sensor network and data driven approach to satellite remote sensing of dissolved organic matter. *Earth Space Sci.* **2025**, *12*, e2024EA004048.
53. Saha, G.K.; Namadi, P.; He, M.; Sandhu, P. Machine Learning-Based Harmful Algal Blooms (HABs) Modeling in the Sacramento-San Joaquin Delta. In *Proceedings of the 2024 Bay-Delta Science Conference*, Sacramento, CA, USA, 30 September–2 October 2024.
54. Framsted, N.; Sturgeon, C.; Kraus, T.; Bouma-Gregson, K. *Assessing Drivers of Harmful and Beneficial Algal Blooms Along the Riverine to Freshwater Tidal Continuum in a Modified Delta*; USGS: Sacramento, CA, USA, 2025.

55. Namadi, P.; He, M.; Sandhu, P. Salinity-constituent conversion in South Sacramento-San Joaquin Delta of California via machine learning. *Earth Sci. Inform.* **2022**, *15*, 1749–1764.
56. Namadi, P.; He, M.; Sandhu, P. Modeling ion constituents in the Sacramento-San Joaquin Delta using multiple machine learning approaches. *J. Hydroinform.* **2023**, *25*, 2541–2560.
57. Namadi, P.; He, M.; Sandhu, P. Advancing Ion Constituent Simulations in California's Sacramento–San Joaquin Delta Using Machine Learning Tools. *Water* **2025**, *17*, 1511. <https://doi.org/10.3390/w17101511>.
58. Swyers, N.M.; Blake, A.R.; Stumpner, P.; Burau, J.R.; Burdick, S.M.; Anwar, M.S. *A Machine Learning Tool for Design of Behavioral Fish Barriers in the Sacramento-San Joaquin River Delta*; US Geological Survey: Sacramento, CA, USA, 2024.
59. Olivetti, S.; Gil, M.A.; Sridharan, V.K.; Hein, A.M. Merging computational fluid dynamics and machine learning to reveal animal migration strategies. *Methods Ecol. Evol.* **2021**, *12*, 1186–1200.
60. Wang, Z.; Leung, M.; Mukhopadhyay, S.; Sunkara, S.V.; Steinschneider, S.; Herman, J.; Abellera, M.; Kucharski, J.; Nederhoff, K.; Ruggiero, P. A hybrid statistical–dynamical framework for compound coastal flooding analysis. *Environ. Res. Lett.* **2024**, *20*, 014005.
61. Wang, Z.; Leung, M.; Mukhopadhyay, S.; Sunkara, S.V.; Steinschneider, S.; Herman, J.; Abellera, M.; Kucharski, J.; Ruggiero, P. Compound coastal flooding in San Francisco Bay under climate change. *npj Nat. Hazards* **2025**, *2*, 3.
62. Russo, B.M.; Athanasopoulos-Zekkos, A. Exploration of feature engineering techniques and unsupervised machine learning clustering algorithms for geophysical data on levees. In *Geo-Congress 2024*; American Society of Civil Engineers: Reston, VA, USA, 2024; pp. 454–463.
63. RMA. *Delta Emergency Response Tool (Delta-ERT): Machine Learning Technical Documentation*; Prepared for California Department of Water Resources; California Department of Water Resources: Sacramento, CA, USA, 2021; p. 6.
64. Lund, J.; Hanak, E.; Fleenor, W.; Bennett, W.; Howitt, R. *Comparing Futures for the Sacramento-San Joaquin Delta*; University of California Press: Berkeley, CA, USA, 2010.
65. Fichot, C.G.; Downing, B.D.; Bergamaschi, B.A.; Windham-Myers, L.; Marvin-DiPasquale, M.; Thompson, D.R.; Gierach, M.M. High-resolution remote sensing of water quality in the San Francisco Bay–Delta Estuary. *Environ. Sci. Technol.* **2016**, *50*, 573–583.
66. Schreier, B.M.; Baerwald, M.R.; Conrad, J.L.; Schumer, G.; May, B. Examination of predation on early life stage Delta Smelt in the San Francisco estuary using DNA diet analysis. *Trans. Am. Fish. Soc.* **2016**, *145*, 723–733.
67. Kraus, T.E.C.; Bergamaschi, B.A.; Downing, B.D. *An Introduction to High-Frequency Nutrient and Biogeochemical Monitoring for the Sacramento–San Joaquin Delta, Northern California*; U.S. Geological Survey Scientific Investigations Report 2017–5071; US Geological Survey: Sacramento, CA, USA, 2017; 41p. <https://doi.org/10.3133/sir20175071>.
68. Downing, B.D.; Bergamaschi, B.A.; Kraus, T.E.C. *Synthesis of Data from High-Frequency Nutrient and Associated Biogeochemical Monitoring for the Sacramento–San Joaquin Delta, Northern California*; U.S. Geological Survey Scientific Investigations Report 2017–5066; US Geological Survey: Sacramento, CA, USA, 2017; 28p. <https://doi.org/10.3133/sir20175066>.
69. Bergamaschi, B.A.; Downing, B.D.; Kraus, T.E.C.; Pellerin, B.A. *Designing a High-Frequency Nutrient and Biogeochemical Monitoring Network for the Sacramento–San Joaquin Delta, Northern California*; U.S. Geological Survey Scientific Investigations Report 2017–5058; US Geological Survey: Sacramento, CA, USA, 2017; 40p. <https://doi.org/10.3133/sir20175058>.
70. Nagarajan, R.P.; Bedwell, M.; Holmes, A.E.; Sanches, T.; Acuña, S.; Baerwald, M.; Barnes, M.A.; Blankenship, S.; Connon, R.E.; Deiner, K.; et al. Environmental DNA methods for ecological monitoring and biodiversity assessment in estuaries. *Estuaries Coasts* **2022**, *45*, 2254–2273.
71. Nagarajan, R.P.; Sanders, L.; Kolm, N.; Perez, A.; Senegal, T.; Mahardja, B.; Baerwald, M.R.; Schreier, A.D. CRISPR-based environmental DNA detection for a rare endangered estuarine species. *Environ. DNA* **2024**, *6*, e506.
72. Climate Change MAST (Management, Analysis and Synthesis Team). *Synthesis of Data and Studies Related to the Effects of Climate Change on the Ecosystem and Biota of the Upper San Francisco Estuary*. IEP Technical Report 99. 2022. Available online: <https://iep.ca.gov/Publications/Technical-Publications> (accessed on 1 January 2026).
73. Sommer, T.; Conrad, J.L.; Culbertson, S. Data to Decisions: How to Make Science More Relevant for Management of the San Francisco Estuary. *San Franc. Estuary Watershed Sci.* **2023**, *21*. <https://doi.org/10.15447/sfew.2023v21iss1art1>.
74. Bashevkin, S.M.; Mahardja, B.; Brown, L.R. Warming in the upper San Francisco Estuary: Patterns of water temperature change from five decades of data. *Limnol. Oceanogr.* **2022**, *67*, 1065–1080.
75. Dettinger, M.; Anderson, J.; Anderson, M.; Brown, L.R.; Cayan, D.; Maurer, E. Climate change and the Delta. *San Franc. Estuary Watershed Sci.* **2016**, *14*. <https://doi.org/10.15447/sfew.2016v14iss3art5>.
76. Mac Nally, R.; Thomson, J.R.; Kimmerer, W.J.; Feyrer, F.; Newman, K.B.; Sih, A.; Bennett, W.A.; Brown, L.; Fleishman, E.; Culbertson, S.D.; et al. Analysis of pelagic species decline in the upper San Francisco Estuary using multivariate autoregressive modeling (MAR). *Ecol. Appl.* **2010**, *20*, 1417–1430.

77. Thomson, J.R.; Kimmerer, W.J.; Brown, L.R.; Newman, K.B.; Nally, R.M.; Bennett, W.A.; Feyrer, F.; Fleishman, E. Bayesian change point analysis of abundance trends for pelagic fishes in the upper San Francisco Estuary. *Ecol. Appl.* **2010**, *20*, 1431–1448.
78. Rogers, T.L.; Bashevkin, S.M.; Burdi, C.E.; Colombano, D.D.; Dudley, P.N.; Mahardja, B.; Mitchell, L.; Perry, S.; Saffarinia, P. Evaluating top-down, bottom-up, and environmental drivers of pelagic food web dynamics along an estuarine gradient. *Ecology* **2024**, *105*, e4274.
79. Hampton, S.E.; Strasser, C.A.; Tewksbury, J.J.; Gram, W.K.; Budden, A.E.; Batcheller, A.L.; Duke, C.S.; Porter, J.H. Big data and the future of ecology. *Front. Ecol. Environ.* **2013**, *11*, 156–162.
80. Clark, M.P.; Nijssen, B.; Lundquist, J.D.; Kavetski, D.; Rupp, D.E.; Woods, R.A.; Freer, J.E.; Gutmann, E.D.; Wood, A.W.; Brekke, L.D.; et al. A unified approach for process-based hydrologic modeling: 1. Modeling concept. *Water Resour. Res.* **2015**, *51*, 2498–2514.
81. Colombano, D.D.; Rudnick, J.; Rowlands, N.A. Five Perspectives to Advance Science-Informed Decision-Making in the Era of Climate. *San Franc. Estuary Watershed Sci.* **2025**, *23*, 3.
82. Shen, C. A transdisciplinary review of deep learning research and its relevance for water resources scientists. *Water Resour. Res.* **2018**, *54*, 8558–8593.
83. Reichstein, M.; Camps-Valls, G.; Stevens, B.; Jung, M.; Denzler, J.; Carvalhais, N.; Prabhat, F. Deep learning and process understanding for data-driven Earth system science. *Nature* **2019**, *566*, 195–204.
84. Nearing, G.S.; Kratzert, F.; Sampson, A.K.; Pelissier, C.S.; Klotz, D.; Frame, J.M.; Prieto, C.; Gupta, H.V. What role does hydrological science play in the age of machine learning? *Water Resour. Res.* **2021**, *57*, e2020WR028091.
85. MacWilliams, M.L.; Ateljevich, E.S.; Monismith, S.G.; Enright, C. An overview of multi-dimensional models of the Sacramento–San Joaquin Delta. *San Franc. Estuary Watershed Sci.* **2016**, *14*. <https://doi.org/10.15447/sfews.2016v14iss4art2>.
86. Sandhu, P.; Finch, R. Artificial Neural Networks and Their Applications. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Sixteenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1995.
87. Sandhu, P.; Finch, R. Application of artificial neural networks to the Sacramento-San Joaquin Delta. In *Proceedings of the 1995 4th International Conference on Estuarine and Coastal Modeling*, San Diego, CA, USA, 26–28 October 1995; pp. 490–503.
88. Denton, R.A. Accounting for antecedent conditions in seawater intrusion modeling—Applications for the San Francisco Bay-Delta. In *Hydraulic Engineering*; CSWRCB: Sacramento, CA, USA, 1993; pp. 448–453.
89. Winkler, E.D. *Preliminary Development of a Statistical Approach to Salinity Modeling in the Western Sacramento-San Joaquin Delta and Suisun Bay*; State of California, Resources Agency, Department of Water Resources, Central District: Sacramento, CA, USA, 1985.
90. Shumway, R.H.; Azari, R.S. *Statistical Analysis of Historical Delta Water Quality Data*; Report prepared for Delta Modeling Section, Division of Planning, California Department of Water Resources under Contract No. B-59259; California Department of Water Resources: Sacramento, CA, USA, 1993.
91. Chung, F.; Sandhu, N.; Wilson, D.; Finch, R. *Modeling Flow Salinity Relationships in the Sacramento-San Joaquin Delta Using Artificial Neural Networks*; Technical report OSP-99-1; Department of Water Resources Office of SWP Planning: Sacramento, CA, USA, 1999.
92. Kohonen, T. An introduction to neural computing. *Neural Netw.* **1988**, *1*, 3–16.
93. Zell, A.; Mache, N.; Huebner, R.; Mamier, G.; Vogt, M.; Schmalzl, M.; Herrmann, K.U. SNNS (stuttgart neural network simulator). In *Neural Network Simulation Environments*; Springer: Boston, MA, USA, 1994; pp. 165–186.
94. Rosenblatt, F. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychol. Rev.* **1958**, *65*, 386.
95. Minsky, M.; Papert, S. *An Introduction to Computational Geometry*; MIT: Cambridge, MA, USA, 1969; Volume 479, p. 104.
96. Baum, E.B. On the capabilities of multilayer perceptrons. *J. Complex.* **1988**, *4*, 193–215.
97. Amari, S.I. Backpropagation and stochastic gradient descent method. *Neurocomputing* **1993**, *5*, 185–196.
98. Prasad, N.N.; Rao, J.N. The estimation of the mean squared error of small-area estimators. *J. Am. Stat. Assoc.* **1990**, *85*, 163–171.
99. Møller, M.F. A scaled conjugate gradient algorithm for fast supervised learning. *Neural Netw.* **1993**, *6*, 525–533.
100. Sandhu, P. Emulation of DWRDSM Using Artificial Neural Networks and Estimation of Sacramento River Flow. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Seventeenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1996.
101. Sandhu, P.; Wilson, D. Marginal Export Cost and MDO Replacement. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Eighteenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1997.

102. Wilson, D. Artificial Neural Networks and MEC Estimates. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Nineteenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1998.
103. Hutton, P.; Sandhu, P. Disinfection By-product Formation. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Sixteenth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 1995.
104. Pranger, T. Artificial Neural Network Development. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twenty-First Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2000.
105. Rajkumar, T.; Johnson, M.L. Prediction of salinity in San Francisco Bay-Delta using neural network. In *2001 IEEE International Conference on Systems, Man and Cybernetics. e-Systems and e-Man for Cybernetics in Cyberspace (Cat. No. 01CH37236)*; IEEE: New York, NY, USA, 2001; Volume 1, pp. 329–334.
106. Rajkumar, T.; Thompson, D. Optimization of a model to predict salinity intrusion in San Francisco Bay Estuary using a genetic algorithm. Integrated assessment and decision support. In *Proceedings of the First Biennial Meeting of the International Environmental Modelling and Software Society*, Lugano, Switzerland, 24–27 June 2002; Volume II, pp. 148–153.
107. Wilbur, R.; Munevar, A. Integration of CALSIM and Artificial Neural Networks Models for Sacramento-San Joaquin Delta Flow-Salinity Relationships. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twenty-Second Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2001.
108. Hutton, P.; Seneviratne, S. An Initial Assessment of Delta Carriage Water Requirements Using a New CALSIM Flow-Salinity Routine. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twenty-Second Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2001.
109. Mierzwa, M. CALSIM versus DSM2 ANN and G-model Comparisons. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twenty-Third Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2002.
110. Seneviratne, S. Short-Term Improvements to Artificial Neural Network Implementation in CALSIM. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twenty-Third Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2002.
111. Seneviratne, S.; Wu, S. Enhanced Development of Flow-Salinity Relationships in the Delta Using Artificial Neural Networks: Incorporating Tidal Influence. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: Twenty-Eighth Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2007.
112. He, M.; Zhong, L.; Sandhu, P.; Zhou, Y. Emulation of a Process-Based Salinity Generator for the Sacramento–San Joaquin Delta of California via Deep Learning. *Water* **2020**, *12*, 2088. <https://doi.org/10.3390/w12082088>.
113. Roh, D.M.; He, M.; Bai, Z.; Sandhu, P.; Chung, F.; Ding, Z.; Qi, S.; Zhou, Y.; Hoang, R.; Namadi, P.; et al. Physics-Informed Neural Networks-Based Salinity Modeling in the Sacramento–San Joaquin Delta of California. *Water* **2023**, *15*, 2320. <https://doi.org/10.3390/w15132320>.
114. Rath, J.S.; Hutton, P.H.; Chen, L.; Roy, S.B. A hybrid empirical-Bayesian artificial neural network model of salinity in the San Francisco Bay-Delta estuary. *Environ. Model. Softw.* **2017**, *93*, 193–208.
115. Qi, S.; Bai, Z.; Ding, Z.; Jayasundara, N.; He, M.; Sandhu, P.; Seneviratne, S.; Kadir, T. Enhanced artificial neural networks for salinity estimation and forecasting in the Sacramento-San Joaquin delta of California. *J. Water Resour. Plan. Manag.* **2021**, *147*, 04021069.
116. Qi, S.; He, M.; Bai, Z.; Ding, Z.; Sandhu, P.; Zhou, Y.; Namadi, P.; Tom, B.; Hoang, R.; Anderson, J. Multi-Location Emulation of a Process-Based Salinity Model Using Machine Learning. *Water* **2022**, *14*, 2030. <https://doi.org/10.3390/w14132030>.
117. Qi, S.; He, M.; Bai, Z.; Ding, Z.; Sandhu, P.; Chung, F.; Namadi, P.; Zhou, Y.; Hoang, R.; Tom, B.; et al. Novel Salinity Modeling Using Deep Learning for the Sacramento–San Joaquin Delta of California. *Water* **2022**, *14*, 3628. <https://doi.org/10.3390/w14223628>.
118. Qi, S.; He, M.; Hoang, R.; Zhou, Y.; Namadi, P.; Tom, B.; Sandhu, P.; Bai, Z.; Chung, F.; Ding, Z.; et al. Salinity Modeling Using Deep Learning with Data Augmentation and Transfer Learning. *Water* **2023**, *15*, 2482. <https://doi.org/10.3390/w15132482>.
119. Ateljevich, E.; DeGeorge, J. *Analysis of Delta Salinity During Extended Drought–Pilot Project*; Final Project Report for Delta Stewardship Council; Delta Stewardship Council: Sacramento, CA, USA, 2025.

120. Chen, L.; Roy, S.B.; Hutton, P.H. Emulation of a process-based estuarine hydrodynamic model. *Hydrol. Sci. J.* **2018**, *63*, 783–802.
121. Ahmadi, A.; Kazemi, M.H.; Daccache, A.; Snyder, R.L. SolarET: A generalizable machine learning approach to estimate reference evapotranspiration from solar radiation. *Agric. Water Manag.* **2024**, *295*, 108779.
122. Ahmadi, A.; Daccache, A.; He, M.; Namadi, P.; Bafti, A.G.; Sandhu, P.; Bai, Z.; Snyder, R.L.; Kadir, T. Enhancing the accuracy and generalizability of reference evapotranspiration forecasting in California using deep global learning. *J. Hydrol. Reg. Stud.* **2025**, *59*, 102339.
123. Denton, R. *Delta Salinity Constituent Analysis*; Richard Denton & Associates: Concord, CA, USA, 2015.
124. Hutton, P.; Sinha, A.; Roy, S. 2022 Simplified Approach for Estimating Salinity Constituent Concentrations in the San Francisco Estuary & Sacramento-San Joaquin River Delta. User Guide. July 2022. Available online: <https://rtdf.info/> (accessed on 1 March 2026).
125. He, M.; Zhou, Y.; Kim, H.; Nader-Tehrani, P.; Sandhu, P. DSM2 Water Temperature Modeling Input Extension: 1922–2015. In *Methodology for Flow and Salinity Estimates in the Sacramento-San Joaquin Delta and Suisun Marsh: 43rd Annual Progress Report to the State Water Resources Control Board*; California Department of Water Resources: Sacramento, CA, USA, 2022.
126. Tillotson, M.D.; Hassrick, J.; Collins, A.L.; Phillis, C. Machine learning forecasts to reduce risk of entrainment loss of endangered salmonids at large-scale water diversions in the Sacramento–San Joaquin Delta, California. *San Franc. Estuary Watershed Sci.* **2022**, *20*. <https://doi.org/10.15447/sfews.2022v20iss2art3>.
127. Tillotson, M.; Lee, C.Y.; Garcia, C. Association of delta smelt with other small pelagic fishes suggests the potential for competition and proxy monitoring. *Endanger. Species Res.* **2025**, *56*, 117–134.
128. Brereton, A.; Mekonnen, Z.A.; Arora, B.; Riley, W.J.; Yuan, K.; Xu, Y.; Zhang, Y.; Zhu, Q.; Anthony, T.L.; Paytan, A. Development of a model framework for terrestrial carbon flux prediction: The Regional Carbon and Climate Analytics Tool (RCCAT) applied to non-tidal wetlands. *Geosci. Model Dev.* **2025**, *18*, 8157–8173.
129. Doshi-Velez, F.; Kim, B. Towards a rigorous science of interpretable machine learning. *arXiv* **2017**, arXiv:1702.08608.
130. He, M.; Sandhu, P.; Namadi, P.; Reyes, E.; Guivetchi, K.; Chung, F. Protocols for Water and Environmental Modeling Using Machine Learning in California. *Hydrology* **2025**, *12*, 59. <https://doi.org/10.3390/hydrology12030059>.
131. Epstein, Z.; Hertzmann, A.; Investigators of Human Creativity; Akten, M.; Farid, H.; Fjeld, J.; Frank, M.R.; Groh, M.; Herman, L.; Leach, N.; et al. Art and the science of generative AI. *Science* **2023**, *380*, 1110–1111.
132. Grieves, M.W. Product lifecycle management: The new paradigm for enterprises. *Int. J. Prod. Dev.* **2005**, *2*, 71–84.
133. Park, J.S.; O'Brien, J.; Cai, C.J.; Morris, M.R.; Liang, P.; Bernstein, M.S. Generative agents: Interactive simulacra of human behavior. In Proceedings of the 36th Annual Acm Symposium on User Interface Software and Technology, San Francisco, CA, USA, 29 October 2023–1 November 2023; pp. 1–22.
134. van der Laan, M.J.; Polley, E.C.; Hubbard, A.E. Super learner. *Stat. Appl. Genet. Mol. Biol.* **2007**, *6*, 25.
135. Baltrušaitis, T.; Ahuja, C.; Morency, L.P. Multimodal machine learning: A survey and taxonomy. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *41*, 423–443.
136. Crystal-Ornelas, R.; Varadharajan, C.; O’Ryan, D.; Beilsmith, K.; Bond-Lamberty, B.; Boye, K.; Burrus, M.; Cholia, S.; Christianson, D.S.; Crow, M.; et al. Enabling FAIR data in Earth and environmental science with community-centric (meta) data reporting formats. *Sci. Data* **2022**, *9*, 700.
137. Martinez, E.S.; Tejada-Gutiérrez, E.; Sorribas, A.; Mateo-Fornes, J.; Solsona, F.; Defacio, R.; Alves, R. Multimodal data integration to model, predict, and understand changes in plant biodiversity: A systematic review. *Ecol. Inform.* **2025**, *92*, 103485.
138. Cantor, A.; Kiparsky, M.; Kennedy, R.; Hubbard, S.; Bales, R.; Pecharroman, L.C.; Guivetchi, K.; McCreedy, C.; Darling, G. 2018 *Data for Water Decision Making: Informing the Implementation of California’s Open and Transparent Water Data Act Through Research and Engagement*; UC Berkeley Center for Law, Energy & the Environment: Berkeley, CA, USA, 2018.
139. Aboagye-Ntow, S.; Huang, X. A survey of multimodal data fusion in Earth observation-remote sensing. *Preprints.org* **2025**, <https://doi.org/10.20944/preprints202510.0743.v1>.
140. Li, J.; Hong, D.; Gao, L.; Yao, J.; Zheng, K.; Zhang, B.; Chanussot, J. Deep learning in multimodal remote sensing data fusion: A comprehensive review. *Int. J. Appl. Earth Obs. Geoinf.* **2022**, *112*, 102926.
141. Lin, W.; Li, T.; Li, X. Deep learning-based object detection for environmental monitoring using big data. *Front. Environ. Sci.* **2025**, *13*, 1566224.
142. Li, X.; Feng, M.; Ran, Y.; Su, Y.; Liu, F.; Huang, C.; Shen, H.; Xiao, Q.; Su, J.; Yuan, S.; et al. Big Data in Earth system science and progress towards a digital twin. *Nat. Rev. Earth Environ.* **2023**, *4*, 319–332.
143. Zhao, T.; Song, C.; Yu, J.; Xing, L.; Xu, F.; Li, W.; Wang, Z. Leveraging Immersive Digital Twins and AI-Driven Decision Support Systems for Sustainable Water Reserves Management: A Conceptual Framework. *Sustainability* **2025**, *17*, 3754.

144. Bam, P.G.; Rezaei, N.; Roubanis, A.; Austin, D.; Austin, E.; Tarroja, B.; Takacs, I.; Villez, K.; Rosso, D. Digital Twin Applications in the Water Sector: A Review. *Water* **2025**, *17*, 2957.
145. Durden, J.M. Environmental management using a digital twin. *Environ. Sci. Policy* **2025**, *164*, 104018.
146. Ren, Y.; Zhang, T.; Dong, X.; Li, W.; Wang, Z.; He, J.; Zhang, H.; Jiao, L. WaterGPT: Training a large language model to become a hydrology expert. *Water* **2024**, *16*, 3075.
147. He, G.; Luo, J.; Luo, F.; Yang, X.; Jing, X.; Cui, H. IWMS-LLM: An intelligent water resources management system based on large language models. *J. Hydroinform.* **2025**, *27*, 1685–1702.
148. Google. *Overview of Climate Science in the Sacramento-San Joaquin Delta*; AI Generated Summar of Selected Documents for the Delta Stewardship Council; Delta Stewardship Council: Sacramento, CA, USA, 2025.
149. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 4766–4777.
150. Ribeiro, M.T.; Singh, S.; Guestrin, C. “Why should I trust you?” Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; pp. 1135–1144.
151. Pan, S.J.; Yang, Q. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **2009**, *22*, 1345–1359.
152. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*; MIT Press: Cambridge, MA, USA, 1998.
153. Albrecht, S.V.; Christianos, F.; Schäfer, L. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*; MIT Press: Cambridge, MA, USA, 2024.
154. Ni, J.; Liu, M.; Ren, L.; Yang, S.X. A multiagent Q-learning-based optimal allocation approach for urban water resource management system. *IEEE Trans. Autom. Sci. Eng.* **2013**, *11*, 204–214.
155. Saha, A.; Pal, S.C. Application of machine learning and emerging remote sensing techniques in hydrology: A state-of-the-art review and current research trends. *J. Hydrol.* **2024**, *632*, 130907.
156. Tripathy, K.P.; Mishra, A.K. Deep learning in hydrology and water resources disciplines: Concepts, methods, applications, and research directions. *J. Hydrol.* **2024**, *628*, 130458.
157. Ahmed, A.A.; Sayed, S.; Abdoulhalik, A.; Moutari, S.; Oyedele, L. Applications of machine learning to water resources management: A review of present status and future opportunities. *J. Clean. Prod.* **2024**, *441*, 140715.
158. Zhi, W.; Appling, A.P.; Golden, H.E.; Podgorski, J.; Li, L. Deep learning for water quality. *Nat. Water* **2024**, *2*, 228–241.
159. Willard, J.D.; Varadharajan, C.; Jia, X.; Kumar, V. Time series predictions in unmonitored sites: A survey of machine learning techniques in water resources. *Environ. Data Sci.* **2025**, *4*, e7.
160. Nie, Y.; Yu, K.H.; Wang, Y.; Liu, P. Applications of machine learning and deep learning in hydrology from a bibliometric perspective: A comprehensive review. *Discov. Artif. Intell.* **2025**, *5*, 242.
161. Lokman, A.; Ismail, W.Z.; Aziz, N.A. A review of water quality forecasting and classification using machine learning models and statistical analysis. *Water* **2025**, *17*, 2243.
162. Torres González, M.A.; Ceballos Pérez, S.G.; Lara Figueroa, H.N.; Ávila Camacho, F.J.; Moreno Villalba, L.M.; Carrillo, J.M.; Meléndez Ramírez, A. Machine learning and predictive models for water management: A systematic review. *Front. Water* **2026**, *8*, 1756052.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.