

ABSTRACT

Title of Thesis: SPATIOTEMPORAL SALINITY DYNAMICS
IN THE PATUXENT RIVER: GOVERNING
FACTORS AND RECONSTRUCTION
USING BOUNDARY CONDITIONS

Junxiao Wang, Master of Science, 2026

Thesis directed by: Assistant Professor, Jian Zhao, University of
Maryland Center for Environmental Science

Estuarine salinity integrates the effects of freshwater forcing, marine boundary conditions, mixing, and channel geometry, but these controls are often expressed unevenly in space and time. The Patuxent River, a tributary of the Chesapeake Bay, provides a setting in which these controls can be traced along a single channel axis. This thesis examines the processes driving salinity variability in this system and tests whether a machine-learning extension can improve a reduced-complexity physical framework. Chapter 2 shows that Patuxent salinity is organized by station-specific seasonal climatology rather than by a single estuary-wide signal. A reduced-complexity physical model quantifies the contributions of the two boundary conditions: freshwater discharge at the head of tide and salinity at the mouth of the Patuxent River. Counterfactual and cross-sectional diagnostics indicate that salinity in the uppermost reach and at the river mouth is each largely governed by a single driver, whereas the interaction of the two drivers controls spatiotemporal salinity variation in the mid-estuary. This intermediate region, here termed the transition zone, is where freshwater influence is filtered by channel geometry and boundary

effects. Chapter 3 reconstructs salinity variability using both the reduced-complexity physical model and a multilayer perceptron (MLP). The physical model performs well at stations where the annual salinity cycle is stable from year to year. The MLP, in contrast, improves aggregate prediction and adds the most value in the upper transition reach, particularly under high- and low-flow conditions and in representative stress-test years. Overall, the thesis shows that a reduced-complexity model can quantify the contributions of the two dominant drivers of spatiotemporal salinity change along the Patuxent River, and that an MLP offers an effective means of combining those drivers to reconstruct salinity variability. More broadly, this pairing offers a transferable strategy for estuaries where monitoring records are sparse or discontinuous and could be extended to other coastal systems and to variables beyond salinity.

SPATIOTEMPORAL SALINITY DYNAMICS IN THE PATUXENT RIVER:
GOVERNING FACTORS AND RECONSTRUCTION USING BOUNDARY
CONDITIONS

by

Junxiao Wang

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Master of Science
2026

Advisory Committee:

Dr. Hongsheng Bi, Chair

Dr. Jian Zhao

Dr. Vyacheslav Lyubchich

© Copyright by
Junxiao Wang
2026

Dedication

This thesis is dedicated to all my family and friends who support me throughout the years.

Acknowledgements

I would first like to express my sincere gratitude to the members of my thesis committee, Dr. Hongsheng Bi, Dr. Jian Zhao, and Dr. Vyacheslav Lyubchich, for their time, guidance, and thoughtful feedback throughout this research. I am especially grateful to my advisors, Dr. Bi and Dr. Zhao, for their continuous support, patience, and mentorship throughout my graduate studies. Their advice and encouragement were essential to the development and completion of this thesis. I am also deeply thankful to Dr. Lyubchich for his valuable comments, expertise, and contributions as a member of my committee.

I would also like to thank the entire UMCES community for providing a supportive and intellectually engaging environment in which to study and conduct research. I am grateful to the faculty, staff, and researchers whose assistance and encouragement made my experience at UMCES both productive and rewarding.

My sincere thanks also go to my fellow students for their companionship, encouragement, and willingness to share ideas and experiences. Their support made the challenges of graduate study easier to navigate and the journey much more enjoyable.

Finally, I would like to thank my family and friends for their unconditional support, understanding, and encouragement. Their confidence in me sustained me throughout this process, and I am deeply grateful to have had them by my side.

Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	vi
List of Figures	vii
List of Abbreviations	ix
Chapter 1: Introduction	1
Estuaries as dynamic transition environments	1
Physical controls on salinity variability	1
Salinity, monitoring records, and seasonal structure	2
The Patuxent River estuary as a study system	3
Knowledge gap	5
Modeling context	5
Research questions and objectives	7
Thesis roadmap	7
Chapter 2: Observed Salinity Structure and Physical Forcing Context in the Patuxent Estuary ..	10
Chapter introduction	10
Data sources	11
Methods	13
Day-of-year standardization and station-specific climatology	13
Annual mean anomalies from local seasonal expectation	14
Annual harmonic analysis of seasonal structure	14
Hydrologic classification and discharge-context comparison	15
Reduced-complexity predictor	16
Normalized station-transfer model	16
Counterfactual forcing and cross-sectional diagnostics	17
Results	18
Observational baseline and station coverage	18
Station-specific climatological structure	18
Annual departures from observed seasonal climatology	21
Annual harmonic departures in seasonal structure	22
Discharge background, event anomalies, and salinity departures	24
Reduced-complexity forcing predictor	26
Normalized station salinity climatology	28
Station-transfer fits	29
Counterfactual decomposition of seasonal forcing controls	33
Geometry-linked attenuation of discharge-driven salinity response	35
Effective-area sensitivity of the forcing predictor	37
Integrated result	38
Discussions	39
Contribution to estuarine-gradient interpretation	39
Seasonal and interannual departures in a sparse monitoring record	40

Freshwater forcing as a spatially filtered control.....	41
Counterfactual forcing controls and spatial separation of mechanisms.....	41
Hydrogeomorphic interpretation of the TF1.7-to-RET1.1 transition	42
Normalized fit, salinity-scale accuracy, and model-evaluation logic	43
Limitations and boundary of inference	43
Chapter Conclusion.....	44
Chapter 3: Machine-Learning Extension of the Reduced-Complexity Model	46
Chapter Introduction	46
Data and Methods	47
Raw station-day observational dataset.....	47
Direct reduced-complexity transfer baseline	47
MLP feature design.....	48
MLP architecture and training	49
Evaluation metrics and common-row comparison	49
SHAP-based model interpretation	50
Results.....	51
Direct reduced-complexity transfer captured only the broad salinity gradient.....	51
MLP improved aggregate prediction	53
MLP improvement was spatially uneven and concentrated in the upper-transition reach ...	55
Flow-stratified improvement showed hydrological-state dependence	56
Representative hydrological years showed how flow context shaped the MLP gain	58
SHAP diagnostics explained the fitted MLP as a spatial-boundary-seasonal model	61
Discussion	63
What the MLP added to the reduced-complexity baseline	63
Pooled metrics must be interpreted through the estuarine gradient.....	64
Representative years link model improvement to physical conditions.....	64
SHAP clarifies what the fitted MLP learned and where it stops	64
The transition zone remains the key diagnostic region.....	65
Implications for future model development.....	65
Chapter Conclusion.....	66
Chapter 4: General Discussion and Conclusions	67
Synthesis of the Thesis.....	67
Answers to the Research Questions.....	68
Persistent spatial salinity structure.....	68
Spatial variation in the seasonal salinity cycle	68
Interannual salinity departures.....	68
Spatial organization of physical forcing	69
Machine-learning extension of the physical baseline	70
Scientific Contributions	70
Limitations and Boundaries of Inference.....	71
Future Research	72
References	73

List of Tables

Table 2.1. Station metrics supporting the observed spatial heterogeneity result

Table 2.2. Observed day-of-year anomaly contrasts by hydrologic background and event-anomaly status

Table 2.3. Forcing construction and reduced-complexity predictor diagnostics

Table 2.4. Station salinity climatology ranges used in the normalized transfer model

Table 2.5. Station-transfer parameters for the normalized station salinity climatology

Table 2.6. Forcing-only sensitivity of the reduced-complexity predictor to effective estuary area

Table 3.1. Raw-observation support for the direct reduced-complexity application

Table 3.2. MLP model hyperparameter configuration

Table 3.3. Station-level diagnostic comparison between the direct reduced-complexity baseline and the MLP extension

Table 3.4. Largest station-flow RMSE reductions for the MLP relative to the direct reduced-complexity baseline

Table 3.5. Representative hydrological years selected from the flow background and model-performance diagnostics

List of Figures

Figure 1.1. Map of the Patuxent River Estuary with labeled stations

Figure 2.1. Observed and harmonic salinity climatology for stations

Figure 2.2. Station-specific mean state, seasonal amplitude, and fitted timing metrics

Figure 2.3. Annual mean salinity anomalies for stations

Figure 2.4. Annual harmonic mean-state, amplitude, and timing departures for stations

Figure 2.5. Annual discharge background classes and event-anomaly years over the salinity-discharge overlap

Figure 2.6. Mean observed day-of-year salinity anomaly by station and discharge background

Figure 2.7. Mean observed day-of-year salinity anomaly by station and event-anomaly status

Figure 2.8. Reduced-complexity forcing climatology on the 365-day no-leap day-of-year axis

Figure 2.9. Normalized station salinity climatology ordered from upstream to downstream

Figure 2.10. Station-transfer scatter fits

Figure 2.11. Transfer parameters by distance from the upstream retained station

Figure 2.12. Observed and fitted normalized salinity day-of-year curves

Figure 2.13. Observed and fitted salinity day-of-year curves in psu

Figure 2.14. Climatology counterfactual observed and predicted salinity curves

Figure 2.15. Climatology counterfactual sensitivity by station distance

Figure 2.16. Geometry-linked attenuation across the TF1.7-to-RET1.1 transition

Figure 2.17. Forcing-only sensitivity of the reduced-complexity predictor to smaller effective estuary area values

Figure 3.1. Observed versus predicted salinity for the direct reduced-complexity raw-observation transfer

Figure 3.2. Spatial pattern of direct reduced-complexity salinity-prediction skill

Figure 3.3. Median salinity residuals by station and month for the direct reduced-complexity transfer

Figure 3.4. Predicted versus observed salinity for the MLP random-holdout test set

Figure 3.5. Station-level MLP and direct reduced-complexity baseline metrics

Figure 3.6. Flow-stratified RMSE, MAE, and R^2 for the MLP and direct reduced-complexity baseline

Figure 3.7. Flow-stratified RMSE reduction from the direct reduced-complexity baseline to the MLP

Figure 3.8. Representative years in the annual freshwater-flow background

Figure 3.9. Observed salinity and two model predictions at TF1.7 across representative hydrological years

Figure 3.10. Station-specific MLP gains in representative hydrological years

Figure 3.11. Global mean absolute SHAP values for grouped MLP feature families

Figure 3.12. SHAP summary plot showing feature value, direction, and contribution to MLP salinity prediction

Figure 3.13. Local SHAP decomposition for a representative TF1.7 high-flow case

List of Abbreviations

AE	Effective estuary area
CB	Chesapeake Bay station prefix
CBP	Chesapeake Bay Program
CO-OPS	Center for Operational Oceanographic Products and Services
DOY	Day of year
LE	Lower-estuary station prefix
L-BFGS	Limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm
MAE	Mean absolute error
MLP	Multilayer perceptron
NOAA	National Oceanic and Atmospheric Administration
psu	Practical salinity unit
Q*	Freshwater-to-inlet-flow ratio used as the reduced-complexity forcing predictor
Q_INL	Positive inlet-flow or tidal-inflow term
Q_R	River discharge or freshwater-flow term
R ²	Coefficient of determination
RC	Reduced-complexity model or reduced-complexity baseline
ReLU	Rectified linear unit
RET	River-estuarine transition station prefix
RMSE	Root mean squared error
SD	Standard deviation
SHAP	SHapley Additive exPlanations
S_M	Downstream marine or boundary salinity

TF Tidal-fresh station prefix

USGS United States Geological Survey

Chapter 1: Introduction

Estuaries as dynamic transition environments

Estuaries are transition zones where river water and coastal ocean water meet, mix, and exchange material. They are not simply downstream extensions of rivers or protected arms of the ocean. Their physical structure emerges from the interaction of freshwater discharge, tides, density gradients, channel geometry, mixing, wind forcing, and downstream boundary conditions (Geyer 2010; MacCready and Geyer 2010). Because these controls vary across seasons and years, estuarine environments are spatially organized but rarely static.

The defining feature of most estuaries is a salinity gradient. In a typical positive estuary, salinity decreases from the ocean toward the river because freshwater input dilutes marine water landward. This horizontal gradient is more than a descriptive pattern. It contributes to density-driven circulation, stratification, salt transport, and the balance between seaward freshwater flow and landward saline influence (Pritchard 1952; Geyer 2010). For that reason, salinity is one of the most useful variables for connecting the physical state of an estuary to its broader environmental structure.

The salinity gradient also links physical forcing to ecological and biogeochemical conditions. Salinity helps define habitat zones, influences the distribution of organisms with different physiological tolerances, and covaries with water-mass history, nutrient transport, residence time, and stratification. In Chesapeake Bay and its tributaries, salinity zones are routinely used to interpret living-resource habitat and water-quality patterns (Kemp et al. 2004; Chesapeake Bay Program). Salinity therefore provides a practical and physically meaningful way to describe how an estuary is organized from upstream tidal-fresh reaches to more marine lower-estuary waters. This thesis begins from that premise. Salinity is treated not only as a water-quality variable to be summarized or predicted, but as an integrative indicator of estuarine environmental structure. The word "integrative" is important. A salinity observation reflects the combined influence of freshwater input, downstream marine conditions, mixing, transport, and station position. It does not identify all mechanisms by itself, but it records their combined expression in a variable that can be compared across stations, seasons, and years.

Physical controls on salinity variability

Freshwater discharge is one of the primary controls on estuarine salinity. Increased river flow generally pushes the salinity field downstream, strengthens seaward freshwater export, and can alter stratification and exchange flow. Low discharge can allow saline water to occupy a larger portion of the estuary. However, the response is not necessarily immediate or spatially uniform. Estuarine adjustment depends on salt storage, tidal mixing, exchange flow, geometry, and the

prior state of the system (MacCready 1999; MacCready and Geyer 2010; Geyer and MacCready 2014).

Tides, exchange flow, and mixing also shape salinity structure. Tidal energy can mix freshwater and saline water vertically and horizontally, while density-driven exchange and turbulent mixing help regulate salt transport. The interaction between tidal mixing and density-driven circulation produces different estuarine regimes, from well-mixed to partially mixed and strongly stratified systems (Geyer 2010; MacCready and Geyer 2010; Geyer and MacCready 2014; MacCready et al. 2018). These regimes matter because the same discharge event can produce different salinity responses depending on where it occurs along the estuary and how strongly the water column is mixed.

Downstream boundary conditions add another layer of variability. In tributary estuaries connected to a larger estuarine or coastal system, salinity at the mouth is not constant. It can vary with coastal ocean conditions, mainstem estuarine circulation, sea-level change, wind forcing, and regional hydrology. In Chesapeake Bay, previous work has shown that salinity variability can be linked not only to river flow but also to sea-level rise, shelf salinity, wind stress, and other basin-scale controls (Hilton et al. 2008). For a tributary such as the Patuxent River estuary, the lower boundary with Chesapeake Bay is therefore an active part of the salinity problem rather than a fixed marine endpoint.

These interacting controls make salinity a challenging but valuable response variable. A single salinity value at one station does not reveal the full estuarine state. A sequence of salinity observations ordered from upstream to downstream, however, can reveal where the gradient is steep, where seasonal amplitude is largest, and where salinity departs from local expectation during wet, dry, or event-dominated years. The scientific problem is therefore not only whether salinity changes with discharge, but how that response is organized across the estuarine gradient.

Salinity, monitoring records, and seasonal structure

Long-term monitoring records are essential for studying estuarine variability because salinity structure changes across multiple time scales. Tidal and weather events can alter salinity over hours to days. Seasonal cycles reflect recurring changes in river flow, temperature, mixing, and downstream conditions. Interannual variability reflects wet and dry years, episodic flow events, changes in boundary conditions, and longer-term environmental change. A record that is long enough to span many years but sparse within each year can still support important questions, but only if the analysis respects the sampling structure.

This issue is common in coastal and estuarine water-quality studies. Monitoring programs often provide valuable multi-decadal records, but observations may be monthly, semi-monthly, or seasonally uneven. Such records can support climatological and interannual analysis, but they should not be treated as continuous daily time series unless the sampling design supports that interpretation. Recent work on coastal water-quality time series emphasizes that low-frequency observations can quantify seasonal and longer-term signals, but the required record length and

uncertainty depend on the sampling frequency and the variability of the system (Brahmey et al. 2024).

Seasonal analysis is particularly important for salinity. The seasonal cycle is often one of the dominant repeating signals in salinity records, and harmonic approaches can provide compact summaries of seasonal mean state, amplitude, and timing (Bingham et al. 2021; Yu 2021; Brahmey et al. 2024). In an estuary, however, the seasonal cycle should not be interpreted as a single basin-wide curve unless the stations share the same salinity regime. A tidal-fresh station, a transitional station, and a lower-estuary station may all have seasonal cycles, but their mean salinity, amplitude, and response to forcing can differ strongly.

For this reason, station-specific seasonal baselines are central to this thesis. A raw salinity value has a different meaning depending on where it is observed. For example, a moderate salinity value may be unusually high at an upstream station but unusually low at a lower-estuary station. Comparing observations to a station-specific day-of-year expectation makes it possible to ask whether a station-year was saltier or fresher than normal for that location and season. This approach preserves the spatial salinity gradient while making interannual departures comparable across the retained station sequence.

The Patuxent River estuary as a study system

The Patuxent River estuary is a tributary estuary of Chesapeake Bay in Maryland. It contains tidal-fresh, riverine-estuarine transition, and lower-estuary zones, and it has been the subject of previous water-quality studies because of its long monitoring record, nutrient-management history, oxygen depletion, and connection to Chesapeake Bay (Testa et al. 2008; Dreiss et al. 2024). This combination makes the Patuxent a useful system for examining how salinity records can reveal estuarine structure.

Prior Patuxent studies provide important context for this thesis. Testa et al. examined long-term changes in water quality and productivity from 1985 to 2003 and emphasized the value of long-term monitoring, freshwater-flow variability, and exchange with Chesapeake Bay for interpreting Patuxent water-quality change. Hartzell et al. used the Patuxent salinity gradient as an organizing axis for phosphorus burial, showing how salinity-defined position can structure biogeochemical interpretation. More recent work by Dreiss et al. analyzed oxygen variability and depletion from 1985 to 2021 and described the estuary in three spatial zones: tidal fresh, riverine-estuarine transition, and lower estuary. That study also reported that dissolved oxygen and hypoxia patterns are shaped by long-term physical and biogeochemical controls, including river flow and temperature.

These studies show that the Patuxent is already recognized as a spatially structured estuarine system. However, the present thesis takes a different entry point. Rather than beginning with nutrients, productivity, or oxygen depletion, it begins with salinity itself. This choice is not meant to replace those other water-quality variables. Instead, salinity is used as the physical

organizing variable that helps define where each station sits along the river-to-bay transition and how that position changes across seasons and years.




Figure 1.1. Map of the Patuxent River Estuary with labeled stations

The monitoring framework used in this thesis is well suited to that question. This thesis uses Chesapeake Bay Program long-term surface salinity observations from Patuxent estuary stations during 1984-2023, with the station sequence ordered from upstream to downstream as TF1.5, TF1.6, TF1.7, RET1.1, LE1.1, LE1.2, LE1.3, LE1.4, and CB5.1W. Discharge context is based on the Patuxent River near Bowie, Maryland USGS gauge 01594440, and downstream tidal water-level context is based on NOAA station 8577330 at Solomons Island. These sources provide the observational basis for connecting local salinity behavior to upstream freshwater

forcing and downstream exchange context (Chesapeake Bay Program; U.S. Geological Survey; NOAA Center for Operational Oceanographic Products and Services).

Knowledge gap

The broad physical relationship between discharge and salinity is well established. Higher freshwater flow tends to freshen estuaries and shift salinity structure downstream, while lower flow can permit greater salt intrusion. Existing estuarine physics provides the conceptual basis for this expectation, and studies in Chesapeake Bay and other estuaries show that river flow, tides, mixing, wind, boundary conditions, geometry, and watershed runoff can all affect salinity variability (Pritchard 1952; MacCready 1999; Hilton et al. 2008; MacCready and Geyer 2010; Ahn, Lee, and Lyu 2020; Agonsi et al. 2026).

The unresolved issue for this thesis is more specific. For the Patuxent River estuary, it is not enough to state that salinity responds to discharge. The key question is how freshwater forcing, downstream boundary salinity, exchange-related processes, and local geometry are expressed along a structured estuarine gradient. Upper stations may remain close to fresh conditions over much of the record, limiting the absolute range of negative salinity departures. Transitional stations may show larger changes because the salinity gradient is steep and because relatively small shifts in the gradient can produce strong local effects. Lower-estuary stations may remain more closely tied to Chesapeake Bay main stem. These expectations are physically plausible, but they must be evaluated using the monitoring record rather than assumed.

This gap matters because salinity structure can be misread if the estuary is treated as one time series or one model domain with uniform behavior. A single estuary-wide seasonal curve would obscure the local meaning of salinity at each station. A single discharge-salinity relationship would overlook the fact that the same freshwater forcing can be filtered differently by station position, downstream boundary salinity, and cross-sectional geometry. A model evaluated only by pooled error could appear successful while still failing to represent the transitional reach where the gradient changes most rapidly.

The thesis therefore frames salinity variability as a spatially structured response problem. The central task is to characterize the persistent upstream-to-downstream gradient, identify how seasonal and annual departures vary by station, evaluate how freshwater and boundary forcing organize those departures, and test whether simplified and data-driven models preserve the observed structure. Modeling is useful here not because it replaces estuarine interpretation, but because its successes and failures reveal which parts of the gradient are physically easier or harder to represent.

Modeling context

Process-based hydrodynamic models are powerful tools for estuarine salinity analysis because they can represent circulation, mixing, bathymetry, forcing, and boundary conditions in physically explicit ways. They are especially valuable for scenario testing and for evaluating

mechanisms that cannot be isolated from observations alone. At the same time, they require substantial input data, calibration, and computational resources, and their results may still depend on uncertain parameterizations and boundary conditions. Studies that emulate or evaluate process-based salinity models show that predictive models can be useful, but that their value depends on how well they reproduce spatial and temporal salinity structure rather than only average accuracy (Hutton et al. 2016; He et al. 2020; Qi et al. 2022).

Reduced-complexity approaches provide a complementary diagnostic tool. A reduced-complexity model can test whether a physically interpretable forcing index captures major features of seasonal salinity variation without resolving the full hydrodynamics of the system. Recent work has shown that seasonal estuarine salinity can be predicted using simplified dimensionless forcing relationships involving river flow, tidal or inlet exchange, and marine salinity (Yavary Nia et al. 2025). In this thesis, that logic is extended by asking not only whether a simple predictor fits the seasonal salinity climatology, but also where its residuals, counterfactual sensitivities, and geometry-linked attenuation reveal limits of the simplified representation.

Machine-learning models provide another complement because they can learn nonlinear relationships among boundary conditions, spatial position, seasonal timing, hydrologic information, and salinity outcomes. Previous estuarine salinity studies have used time-series models, artificial neural networks, deep learning, and related approaches for salinity estimation, forecasting, or model emulation (Qiu and Wan 2013; Etemad-Shahidi et al. 2008; Jayasundara et al. 2020; Qi et al. 2021; Qi, He, Bai, Ding, Sandhu, Chung, et al. 2022; Saccotelli et al. 2024). Recent work has also shown that data-driven models can reproduce physically meaningful salinity dynamics at lower computational cost than full hydrodynamic simulations, although performance can weaken under extreme or unusual conditions (Gorski et al. 2024).

Recent developments in physics-informed machine learning are important for interpreting the limits of purely data-driven salinity prediction. Physics-informed neural networks embed governing-equation information or physical constraints into model training and have been proposed as a way to improve extrapolation and reduce physically unrealistic predictions (Raissi, Perdikaris, and Karniadakis 2019; Karniadakis et al. 2021). In salinity applications, physics-informed approaches and lagged-flow feature designs suggest practical pathways for improving transition-zone prediction and interannual transferability (Roh et al. 2023). These studies provide context for Chapter 3, where a compact Multi-Layer Perceptron (MLP) is used as a nonlinear extension of the reduced-complexity baseline rather than as a final operational forecasting system.

In this thesis, modeling is therefore not the central identity of the project. It is a set of tools used to evaluate and reconstruct salinity structure. The scientific object remains the Patuxent salinity variability and its forcing factors. Chapter 2 uses reduced-complexity transfer, counterfactual forcing, and cross-sectional diagnostics to define the observational and physical benchmark. Chapter 3 then asks whether embedding the reduced-complexity descriptor in a nonlinear MLP feature space improves salinity reconstruction or prediction while preserving that benchmark.

The most important modeling question is not only whether prediction error is low, but whether the model preserves the station-specific seasonal cycle, the transition-zone behavior, and the spatially uneven expression of freshwater forcing.

Research questions and objectives

The broad question of this thesis is: how does the salinity spatiotemporal structure is shaped by freshwater discharge and other physical forcings and can we use the identified driving factors to reconstruct or predict the salinity changes?

This broad question is developed through five more specific questions.

First, what persistent spatial salinity structure exists along the Patuxent River estuary? This question addresses the mean upstream-to-downstream gradient and asks whether retained monitoring stations occupy distinct salinity regimes.

Second, how does the seasonal salinity cycle vary across stations from upstream to downstream? This question addresses whether seasonal mean state, amplitude, and timing can be summarized by station-specific climatology and harmonic analysis without treating the estuary as a single basin-wide salinity curve.

Third, how do annual salinity anomalies reflect interannual variability in estuarine conditions? This question asks whether station-years are saltier or fresher than local seasonal expectation, rather than simply comparing raw annual means across the spatial gradient.

Fourth, how do freshwater discharge, downstream boundary salinity at the mouth of Patuxent River, and local geometry organize the salinity response along the gradient? This question uses hydrologic classification, reduced-complexity transfer modeling, counterfactual forcing experiments, and cross-sectional diagnostics to evaluate how freshwater and boundary salinity determine the salinity changes in space and time.

Fifth, can a compact machine-learning extension improve the direct reduced-complexity baseline while preserving the observed spatial structure? This question connects the physical benchmark from Chapter 2 to the MLP framework in Chapter 3, aiming to reconstruct the salinity changes using the factors and processes identified by the physical framework.

The objective of the thesis is to use salinity as an integrative indicator to characterize estuarine environmental variability across space and time, with particular attention to how freshwater discharge and boundary salinity organize the salinity-defined gradient of the Patuxent River estuary. The objective is deliberately bounded. The thesis does not claim that salinity alone explains all water-quality variability. Instead, it uses salinity to build an observational and modeling framework for understanding how physical forcing is expressed across the Patuxent estuarine gradient.

Thesis roadmap

Chapter 2 establishes the observational and physical benchmark of the thesis. It uses the Patuxent salinity monitoring record to construct station-specific day-of-year climatology, annual mean

anomalies, annual harmonic summaries, hydrologic discharge classes, event-anomaly comparisons, a reduced-complexity station-transfer analysis, counterfactual forcing diagnostics, and a cross-sectional attenuation diagnostic. The results show that the Patuxent estuary is not well described as a single salinity time series. Instead, salinity variability is organized around a persistent but dynamic spatial gradient, with low-salinity upper stations, a transitional reach with large seasonal and interannual variability, and lower-estuary stations closer to the Chesapeake Bay boundary.

The strongest Chapter 2 claims are the station-specific salinity gradient, the transition into TF1.7 and RET1.1, the larger anomaly variability from TF1.7 downstream, the wet-year tendency toward negative anomalies, the spatial separation of boundary and discharge counterfactual effects, the geometry-linked attenuation across the TF1.7-to-RET1.1 reach, and the concentration of reduced-complexity residuals in the transitional zone. The counterfactual analysis indicates that downstream boundary salinity primarily controls lower-estuary seasonal range, whereas fixed-discharge losses are concentrated around TF1.6, TF1.7, and RET1.1. The cross-sectional diagnostic further suggests that the TF1.7-to-RET1.1 transition is a hydrogeomorphic interval where the local expression of freshwater forcing weakens as cross-sectional area expands. Chapter 3 extends the Chapter 2 benchmark into a machine-learning-augmented reduced-complexity framework. It first applies the reduced-complexity transfer directly to raw station-day observations, where the baseline retains broad upstream-downstream salinity structure but has limited station-level temporal skill in the upstream and transition-zone stations. It then embeds the reduced-complexity descriptor S_E^* within an MLP feature space that also includes same-day downstream boundary salinity at CB5.1W, normalized station distance, and cyclic seasonal phase. The MLP is therefore framed as a nonlinear extension of the reduced-complexity baseline, not as a standalone machine-learning exercise.

The Chapter 3 results reinforce the thesis argument by showing that the nonlinear extension improves the direct reduced-complexity baseline while preserving the same spatial interpretation. Improvement is spatially uneven: downstream stations remain the highest-skill stations, while the transition zone remains the most informative test of model limits. At TF1.7, the MLP improves over the direct reduced-complexity baseline but still leaves relatively large error and compressed variability. Seasonal and annual diagnostics are consistent with stronger winter-spring and atypical-year difficulty, but these patterns are treated as diagnostic rather than causal because discharge and other dynamic forcings are not explicit MLP inputs.

Chapter 4 synthesizes what salinity reveals about the Patuxent River estuary. It will return to the central thesis framing: salinity is an integrative indicator of environmental structure, and freshwater discharge helps organize that structure through a spatially uneven response along the estuarine gradient. The broader contribution of the thesis is a defensible observational and modeling framework for interpreting salinity variability in a Chesapeake Bay tributary estuary, with explicit attention to the value and limits of station-specific climatology, reduced-complexity forcing diagnostics, counterfactual physical interpretation, and machine-learning extension.

Chapter 2: Observed Salinity Structure and Physical Forcing Context in the Patuxent Estuary

Chapter introduction

Estuaries are not spatially uniform receiving waters. They are river-ocean transition zones in which freshwater input, marine influence, tides, mixing, channel geometry, and seasonal forcing jointly organize the physical environment. The most visible expression of this organization is the salinity gradient. Along an estuary, salinity marks the transition from river-dominated upper reaches to more marine lower reaches, and this transition provides a physical template for water-quality conditions, habitat boundaries, and biogeochemical variability. In this chapter, salinity is therefore treated not simply as a response variable to be predicted, but as an observable indicator of the spatial structure of the Patuxent estuary.

The scientific problem motivating this chapter is that this salinity-defined spatial structure is dynamic rather than fixed. The location, strength, and seasonal expression of the gradient can change with freshwater forcing, downstream saline influence, tidal exchange, and episodic hydrologic events. A wet year can push fresher conditions downstream, while a dry year can allow saltier conditions to occupy a larger portion of the estuary. However, these changes are unlikely to occur uniformly along the river. Upper tidal-fresh and oligohaline stations may respond differently from transitional stations and lower-estuary stations that remain closer to the Chesapeake Bay boundary. Understanding these spatially uneven responses is necessary before later chapters can evaluate whether process-model output or data-driven models reproduce the behavior of the real system.

The Patuxent estuary is a useful setting for this question because it contains a clear upstream-to-downstream salinity gradient and a long, but irregular, monitoring record. The monitoring data are valuable because they directly measure the system over multiple decades. At the same time, they are not a dense daily time series at every station. The record is spatially and temporally uneven, with station-years differing in sample counts, seasonal coverage, and suitability for annual interpretation. A direct comparison of raw salinity values across stations would therefore mix together three different sources of variation: the persistent along-estuary salinity gradient, the expected seasonal cycle, and true interannual departures from local expectation. This chapter is designed to separate those layers as clearly as possible.

The central question of Chapter 2 is how salinity-defined spatial structure varies across stations, seasons, and years in the Patuxent estuary, and how much of that variability can be organized by freshwater-forcing context. This question is narrower than the overall thesis question, but it provides the observational foundation for the thesis. The chapter does not attempt to build the final predictive model, nor does it claim that discharge is the only driver of salinity variability. Instead, it asks what response patterns can be extracted defensibly from the observations, and

whether those patterns are consistent with an interpretation in which freshwater forcing helps reshape the estuarine salinity gradient.

This framing leads to four analytical needs. First, a station-specific seasonal baseline is required because each station occupies a different position along the salinity gradient. A salinity value of 5 psu means something different at an upper transitional station than at a lower-estuary station. Second, annual departures must be defined relative to local seasonal expectation, rather than relative to a single estuary-wide mean. Third, the seasonal cycle itself must be summarized, because interannual variability can involve not only saltier or fresher annual mean conditions but also changes in seasonal amplitude and timing. Fourth, freshwater forcing must be evaluated in a way that recognizes both annual hydrologic background and seasonal forcing balance.

The chapter therefore proceeds from empirical description toward physical interpretation. It first constructs a station-specific day-of-year climatology to describe the observed seasonal salinity structure along the estuary. It then calculates annual mean anomalies to identify years that are saltier or fresher than the local seasonal baseline. Harmonic analysis is used as a compact description of annual mean state, seasonal amplitude, and timing-related departures where sampling coverage is sufficient. Discharge context is evaluated in two complementary ways: a hydrologic classification compares salinity anomalies among dry, regular, wet, and event-anomaly years, and a reduced-complexity station-transfer analysis tests whether station climatological salinity cycles can be represented as local responses to an estuary-scale freshwater-to-inflow balance.

The chapter's main argument is that Patuxent salinity variability is spatially heterogeneous and structured around station-specific climatology. The estuary does not behave as a single salinity time series. Instead, mean state, seasonal amplitude, annual anomaly magnitude, and discharge sensitivity all vary systematically along the upstream-to-downstream station sequence. The strongest observed variability and the largest transfer-model residuals occur near the transitional portion of the gradient, while downstream stations show smaller normalized seasonal ranges and stronger influence from the downstream boundary. This spatial differentiation is the result that later chapters must preserve: any model that predicts salinity well at individual stations but fails to reproduce the structure of the gradient would not fully answer the thesis question.

Within the thesis, Chapter 2 has two roles. Scientifically, it identifies the observed salinity-gradient patterns that need explanation. Methodologically, it establishes the observational reference against which later models can be judged. The emphasis is on building a coherent chain of evidence from the raw monitoring record to a defensible interpretation of flow-regulated spatial structure.

Data sources

This chapter uses four types of information: surface salinity observations, river discharge records, tidal water-level records, and station metadata. These datasets differ in temporal resolution and purpose, so each is used only for the part of the analysis it can support. Salinity

observations provide the empirical basis for the spatial and seasonal analysis. Discharge provides the upstream freshwater-forcing context. Tidal water level provides a downstream exchange term for the reduced-complexity seasonal predictor. Station metadata provide the spatial ordering needed to interpret results along the estuarine axis.

Surface salinity observations are the primary data source. The salinity archive contains Chesapeake Bay Program long-term observations from Patuxent estuary stations during 1984-2023. Basic quality-control checks found no datetime parsing failures, no exact duplicate rows, no negative salinity values, and an observed salinity range of 0 to 19.66 psu. For the main analysis, the retained station set includes TF1.5, TF1.6, TF1.7, RET1.1, LE1.1, LE1.2, LE1.3, LE1.4, and CB5.1W. These stations provide a consistent upstream-to-downstream framework for the climatology, annual anomaly, harmonic, discharge-context, and station-transfer analyses. The stations do not all have identical sampling histories. In the retained display set, station sample totals used for climatology summaries range from 593 observations at TF1.5 to 688 observations at LE1.3. Observed standardized day-of-year coverage ranges from 284 to 293 days. Annual anomaly summaries include 39 years for every retained station. These numbers indicate that the record is long enough to support station-scale seasonal and interannual summaries, but not dense enough to treat the observations as continuous daily records.

River discharge is used to represent upstream freshwater forcing. The discharge record comes from the Patuxent River near Bowie, Maryland gauge, USGS 01594440. Discharge is used in two ways. First, daily discharge is summarized into annual hydrologic classes so that salinity anomalies can be compared across dry, regular, wet, and event-anomaly years. Second, discharge is used as the upstream freshwater term in the reduced-complexity seasonal forcing predictor. The hydrologic classification is used for interannual context, while the reduced-complexity predictor is used for average-year seasonal forcing.

Tidal water level is used in the reduced-complexity analysis to approximate downstream tidal or marine inflow. The water-level record comes from the Solomons Island, Maryland station, NOAA 8577330. The rate of water-level change is combined with an assumed effective estuarine surface area to estimate a tidal exchange term. The effective surface area is calculated in a GIS from Patuxent hydrography and watershed-boundary data obtained from the National Hydrography Dataset Plus. This term is not intended to replace a hydrodynamic model. Instead, it provides a physically interpretable seasonal-scale index that can be compared with station salinity climatology.

Station metadata include station coordinates and the upstream-to-downstream station order. This ordering is essential because the chapter's main result is spatially structured: the same freshwater-forcing context does not produce identical salinity responses at all locations. CB5.1W is also used as the downstream boundary station for normalized salinity calculations in the reduced-complexity station-transfer analysis.

Methods

The analytical workflow is organized around the scientific uncertainties introduced above. The first uncertainty is the basic spatial and seasonal form of the salinity gradient. This is addressed with a station-specific day-of-year climatology. The second uncertainty is whether individual years are saltier or fresher than expected at each location. This is addressed with annual mean anomalies defined relative to the local climatology. The third uncertainty is whether interannual changes involve the seasonal cycle itself. This is addressed with annual harmonic analysis. The fourth uncertainty is how freshwater forcing helps organize the observed variability. This is addressed with hydrologic classification and a reduced-complexity station-transfer model.

Day-of-year standardization and station-specific climatology

All seasonal analyses use a common 365-day no-leap day-of-year axis. February 29 is excluded, and dates after February 29 in leap years are shifted down by one day so that March 1 and later dates align across leap and non-leap years. This standardization is necessary because the chapter compares observations across years and fits seasonal curves on a fixed annual coordinate. It also ensures that salinity climatology and discharge-based seasonal predictors use the same day-of-year convention.

The station-specific climatology is the empirical seasonal baseline for the chapter. Because mean salinity differs strongly among stations, the climatology is calculated separately for each station rather than pooled across the estuary. Let $S_{i,y,d}$ denote observed salinity at station i , year y , and standardized day of year d . The station-specific day-of-year climatology is calculated as the mean of available observations at that station and day of year across years:

$$\bar{S}_{i,d} = \frac{1}{n_{i,d}} \sum_{y \in Y_{i,d}} S_{i,y,d}$$

where y includes all years with an observation at station i on day d .

This definition preserves the station identity of the salinity gradient. An upstream observation is compared with the upstream station's normal seasonal condition, and a downstream observation is compared with the downstream station's normal seasonal condition. Unobserved station-day combinations are not treated as new observations. The climatology is therefore an empirical baseline, not a fully imputed continuous daily dataset.

A first-order harmonic curve is also fitted to the station climatology as a compact visual and numerical summary of the dominant annual cycle. The observed station-by-day climatology remains the primary baseline for anomaly calculations, while the harmonic fit summarizes broad seasonal mean state, amplitude, and timing. This distinction is important because the observed climatology is closer to the monitoring record, whereas the harmonic fit imposes a smooth sinusoidal seasonal form.

To avoid treating the irregular monitoring record as a complete daily series, uncertainty for the climatological and harmonic summaries was evaluated with a year-block bootstrap. In each

replicate, years were resampled with replacement, station-day means were recomputed, and first-order harmonic parameters were refit. The resulting percentile intervals are used as uncertainty support for station mean state, seasonal amplitude, adjacent-station differences, and phase/timing stability.

Annual mean anomalies from local seasonal expectation

Annual mean anomalies are used to determine whether a station-year is saltier or fresher than expected after accounting for the local seasonal cycle. This step is necessary because raw annual mean salinity would largely reflect station position along the estuary. For example, lower-estuary stations are almost always saltier than upper-estuary stations, but that persistent spatial gradient is not the interannual signal of interest. The anomaly calculation removes the expected station-specific seasonal value before averaging within each station-year.

For station i and year y , the annual mean anomaly is calculated as:

$$A_{i,y} = \frac{1}{n_{i,y}} \sum_{d \in D_{i,y}} (S_{i,y,d} - \bar{S}_{i,d})$$

where $D_{i,y}$ is the set of standardized days with observations for that station-year and $n_{i,y}$ is the number of available observations. Positive values indicate saltier-than-expected conditions relative to the local seasonal baseline. Negative values indicate fresher-than-expected conditions. Because the metric is computed within station and day of year, it separates interannual departure from both the persistent spatial gradient and the mean seasonal cycle.

This anomaly metric is intentionally simple. It does not assume a particular mechanism, and it does not attempt to fill unsampled days. Its strength is interpretability: it directly answers whether the measured salinity in a given station-year was higher or lower than expected for the dates that were sampled. This makes it the chapter's clearest empirical indicator of interannual departure from local salinity structure.

Annual harmonic analysis of seasonal structure

Annual harmonic analysis is used to summarize whether interannual variability involves changes in seasonal structure, not only annual mean offset. For each station-year with sufficient sampling coverage, salinity is fitted as a first-order annual sinusoid:

$$S_{i,y,d} = \alpha_{i,y} + \beta_{i,y} \cos\left(\frac{2\pi d}{365}\right) + \gamma_{i,y} \sin\left(\frac{2\pi d}{365}\right) + \varepsilon_{i,y,d}$$

The coefficient $\alpha_{i,y}$ represents the fitted annual mean state. The sine and cosine coefficients represent the first-order seasonal component. Seasonal amplitude is calculated as:

$$R_{i,y} = \sqrt{\beta_{i,y}^2 + \gamma_{i,y}^2}$$

The phase angle can be calculated from the fitted sine and cosine terms and converted into an approximate timing metric. In this chapter, timing is reported as a descriptor but interpreted cautiously. Timing estimates are less stable when seasonal amplitude is weak or when within-

year sampling is incomplete. The primary harmonic results are therefore the annual mean-state departure and amplitude departure, because these are more directly tied to the base state and strength of the seasonal salinity cycle.

A station-year is treated as departure-ready only if it has at least 12 observations, at least 12 unique standardized day-of-year values, and observations in at least 10 unique months. This rule is designed for an irregular long-term monitoring record. It requires broad seasonal coverage while avoiding an unrealistic requirement for dense daily sampling. Across the stations, 314 of 351 station-years meet the departure-ready criterion.

The harmonic analysis is not used to replace the observed day-of-year climatology. Instead, it provides a second view of interannual variability. Annual anomalies describe whether measured salinity was generally above or below local seasonal expectation. Harmonic departures describe whether the fitted seasonal curve itself shifted in mean state or amplitude during a given year.

Hydrologic classification and discharge-context comparison

Discharge context is evaluated first through a hydrologic classification. This step asks whether station-year salinity anomalies differ systematically under different freshwater-flow backgrounds. Complete discharge years are identified using a minimum coverage rule, and annual mean discharge is used to classify years as dry, regular, or wet. Years below the lower threshold are classified as dry, years above the upper threshold are classified as wet, and intermediate years are classified as regular.

A separate event-anomaly flag is used to identify years with irregular or concentrated high-flow behavior. This flag is independent of the dry/regular/wet background class because annual mean discharge and episodic discharge structure are not the same hydrologic condition. A year can have moderate annual mean discharge but still contain unusually strong or unusually late high-flow events. The event-anomaly flag is triggered when at least two event indicators are present. These indicators include weak seasonal structure in discharge, elevated numbers of high-flow days, late-season high-flow occurrence, high-flow volume share, maximum 7-day discharge, and maximum 7-day positive discharge anomaly.

The dry, regular, and wet classes are treated as a quartile-based classification of complete annual discharge years, with the lower and upper tails defining dry and wet conditions and the middle years retained as the regular background class. Hydrologic-class and event-anomaly contrasts were evaluated with year-level resampling or permutation logic so that inference was not based on treating station-year rows as fully independent observations.

Station-year salinity anomalies are then compared across hydrologic classes. This comparison is descriptive rather than a formal causal attribution test. It asks whether dry, regular, wet, and event-anomaly years are associated with different salinity departures, while preserving station-level differences. This approach is appropriate because discharge is important, but it is not the only process shaping salinity. Downstream salinity boundary conditions, tides, mixing, wind forcing, stratification, and local geometry may also influence the observed response.

Reduced-complexity predictor

The hydrologic classification describes interannual context, but it does not provide a continuous seasonal forcing predictor. To address this limitation, a reduced-complexity predictor (Yavary Nia et al. 2025) is constructed on the same 365-day day-of-year axis as the salinity climatology. The purpose is to represent average-year seasonal conditions using a physically interpretable balance between upstream freshwater input and downstream tidal or marine inflow. For an incompressible fluid, the water volume balance is expressed as:

$$\frac{dV_E}{dt} = A_E \frac{dh_T}{dt} = Q_{INL} + Q_R$$

where V_E is the volume of water in the estuary, t is time, h_T is the tidal level, A_E is the water surface area in the estuary, assumed constant in time, and Q_{INL} and Q_R correspond to the inlet discharge and river discharge, respectively. The main analysis uses $A_E = 126,761,230 \text{ m}^2$. The positive inlet discharge can be obtained by rearranging the volume balance equation:

$$Q_{(INL>0)(t)} = \max\left(A_E \frac{dh_T}{dt} - Q_{R(t)}, 0\right)$$

Both the freshwater term and the positive inlet term are smoothed with a 30-day moving average. This smoothing focuses the reduced-complexity analysis on seasonal-scale forcing rather than individual tidal cycles or short hydrologic events. The dimensionless freshwater-to-inflow ratio is then calculated as:

$$Q_{(d)}^* = \frac{Q_{(R>0,30d)(d)}}{Q_{(INL>0,30d)(d)}}$$

This ratio is converted into an estuary-scale salinity predictor:

$$S_{(E,ss)(d)}^* = \frac{1}{1 + Q_{(d)}^*}$$

Higher values of $S_{E,ss}^*$ indicate conditions more favorable to saline influence, while lower values indicate stronger relative freshwater influence. Because the predictor is constructed as an average-year seasonal curve, it is used to explain climatological seasonal salinity structure rather than individual observation-level salinity fluctuations.

Normalized station-transfer model

To compare stations with different salinity regimes, station salinity climatology is normalized by the downstream boundary station CB5.1W:

$$S_{i(d)}^* = \frac{S_{i(d)}}{S_{CB5.1W(d)}}$$

This normalization expresses each station's climatological salinity as a fraction of the downstream boundary condition on the same day of year. It reduces the dominance of the persistent along-estuary mean gradient and focuses the analysis on how each local station cycle varies relative to the estuary-scale seasonal forcing predictor. In the resulting dataset, CB5.1W equals 1 by definition.

A station-specific transfer model is then fitted for each station:

$$S_{i(d)}^* = a_i S_{(E,ss)(d)}^* + b_i + \varepsilon_{i(d)}$$

The slope a_i is interpreted as a local response coefficient. It describes how strongly the normalized station climatology changes over the seasonal range of the estuary-scale forcing predictor. Because $S_{E,ss}^*$ varies over a narrow interval, numerical slope values can be large. Interpretation therefore focuses on relative differences among stations, sign changes along the estuary, model fit, and residual structure. The intercept b_i represents the station's background offset relative to the predictor. Model performance is summarized using R^2 , root mean squared error, mean absolute error, and mean absolute error after converting the fitted normalized values back to salinity units.

Uncertainty in the transfer parameters and residual metrics was assessed with a 30-day circular seasonal block bootstrap along the day-of-year axis. This treatment is more appropriate than ordinary independent-observation inference because both the predictor and response are smoothed seasonal climatologies with serial dependence.

This transfer model is an empirical physical-response diagnostic, not a full hydrodynamic model. It does not imply that salinity varies linearly along the estuary or that all deviations are controlled by freshwater discharge. Instead, it tests whether a simple seasonal freshwater-to-inflow balance can reproduce the timing and relative magnitude of each station's normalized climatological salinity cycle. The model is most useful for identifying where a one-predictor forcing representation works well, where it fails, and where additional processes are likely needed.

Counterfactual forcing and cross-sectional diagnostics

After fitting the normalized station-transfer model, physical counterfactuals were constructed by holding the fitted station parameters fixed and replacing one forcing component at a time with its long-term mean. The baseline prediction used the observed climatological downstream boundary salinity, freshwater discharge, and positive inlet-flow terms. The fixed_SM scenario held downstream boundary salinity constant, the fixed_Q_R scenario held freshwater discharge constant, and the fixed_Q_INL scenario held the positive inlet-flow term constant. For each scenario, predicted salinity was compared with the baseline using station-specific RMSE change and retained seasonal range.

A separate cross-sectional diagnostic was used to interpret the TF1.7-to-RET1.1 transition. Station cross-sectional areas were taken from the station geometry table, and the first-order no-mixing freshwater-advection tendency was calculated as $-(Q_R/A^*) dS/dx$ on the day-of-year climatology axis. This calculation was used only as a physical scale diagnostic. It asks how strongly the same upstream freshwater discharge would act through the local cross-section, not whether local advection alone predicts observed salinity.

Results

Observational baseline and station coverage

The cleaned salinity archive supports an analysis from 1984 to 2022 for the main observed-pattern results. The retained station set provides a continuous upstream-to-downstream sequence from TF1.5 through CB5.1W. Sample totals are similar enough to support comparison, but the observations remain sparse relative to daily sampling. Station sample totals used in the climatology summaries range from 593 observations at TF1.5 to 688 observations at LE1.3, and observed standardized day-of-year coverage ranges from 284 to 293 day-of-year values. Annual anomaly summaries include 39 years for every retained station.

These coverage results establish the appropriate level of inference for the chapter. The data are strong enough to define station-specific seasonal patterns and broad interannual departures, but they do not support treating the observed record as a continuous daily reconstruction. This is why the results emphasize climatological structure, station-year anomalies, harmonic summaries, and forcing-context diagnostics rather than event-by-event salinity dynamics.

Table 2.1. Station metrics supporting the observed spatial heterogeneity result

Station	Mean state	Seasonal amplitude	Annual anomaly SD	Median abs mean-state departure	Median abs amplitude departure
TF1.5	0.23	0.31	0.15	0.20	0.27
TF1.6	0.91	0.84	0.40	0.42	0.53
TF1.7	4.04	1.88	0.94	1.10	0.95
RET1.1	9.33	2.02	1.06	1.45	1.03
LE1.1	11.18	2.08	1.08	1.36	1.06
LE1.2	12.06	2.09	1.07	1.30	1.06
LE1.3	12.34	2.08	1.05	1.33	0.95
LE1.4	13.02	2.21	1.06	1.25	0.98
CB5.1W	13.38	2.35	1.01	1.33	0.87

Station-specific climatological structure

Salinity climatology is strongly station-specific. The harmonic mean state spans from 0.23 psu at TF1.5 to 13.38 psu at CB5.1W, showing that the stations cover a large portion of the river-to-bay salinity transition. The first-order seasonal amplitude also increases downstream, from 0.31 psu at TF1.5 to 2.35 psu at CB5.1W. The retained station sequence is therefore not characterized by one shared seasonal salinity cycle. Instead, it is organized by station-specific seasonal cycles arranged along the estuarine gradient.

Observed and harmonic salinity climatology




Figure 2.1. Observed and harmonic salinity climatology for stations. The observed station-day-of-year climatology is the empirical baseline, while the harmonic fit summarizes the dominant annual structure.

The upper stations define the low-salinity end of the system. TF1.5 has a mean state of 0.23 psu and a seasonal amplitude of 0.31 psu, while TF1.6 has a mean state of 0.91 psu and an amplitude of 0.84 psu. These stations remain close to fresh or very low-salinity conditions for much of the seasonal cycle. TF1.7 marks a clear transition, with a mean state of 4.04 psu and an amplitude of 1.88 psu. This sharp increase indicates that TF1.7 occupies a portion of the estuary where seasonal displacement of the salinity gradient becomes more visible.

The transition into TF1.7 is also supported by the bootstrap uncertainty analysis. The adjacent TF1.6-to-TF1.7 increase in harmonic mean state was 3.13 psu, with a 95% bootstrap interval of 2.77 to 3.51 psu, and the adjacent increase in seasonal amplitude was 1.04 psu, with a 95% interval of 0.77 to 1.34 psu. These intervals support describing TF1.7 as a station-scale transition in both mean salinity and seasonal expression, while avoiding interpretation as a sharp mechanistic discontinuity.

From RET1.1 through CB5.1W, mean salinity remains much higher. RET1.1 has a mean state of 9.33 psu, and the LE stations range from 11.18 to 13.02 psu. CB5.1W, used later as the downstream boundary for normalization, has a mean state of 13.38 psu. Seasonal amplitude remains relatively high across this lower-estuary group, ranging from 2.02 to 2.35 psu. This indicates that the lower estuary is not static; it also has a clear seasonal cycle. However, the main

spatial contrast shifts from low-salinity upper stations to a cluster of more saline downstream stations.

Station-specific climatology metrics




Figure 2.2. Station-specific mean state, seasonal amplitude, and fitted timing metrics. Mean state and amplitude are the main interpretive parameters.

The climatology results support the first major finding of the chapter: the Patuxent salinity gradient is a structured spatial feature rather than a set of unrelated station time series. The transition from TF1.5 and TF1.6 to TF1.7 and RET1.1 is especially important because it identifies the part of the estuary where the largest salinity gradient occurs. Later discharge and transfer-model results show that this transitional zone is also where response magnitude and residual mismatch are most pronounced.

Annual departures from observed seasonal climatology

Annual mean anomalies show that interannual salinity variability occurs around the station-specific climatological baseline. Across stations, annual anomaly standard deviation ranges from 0.15 to 1.08 psu, and all stations include both positive and negative anomaly years. This sign-changing behavior is important because it confirms that the anomaly metric is capturing departures from local expectation rather than simply reproducing the persistent upstream-to-downstream salinity gradient.




Figure 2.3. Annual mean salinity anomalies for stations. The sign-changing pattern supports interpretation as annual departures from local climatology rather than as a single monotonic trend.

Anomaly magnitude differs strongly along the estuary. TF1.5 and TF1.6 have the smallest annual anomaly standard deviations, 0.15 and 0.40 psu, respectively. TF1.7 increases sharply to 0.94 psu, and the downstream stations from RET1.1 through CB5.1W remain near one psu, ranging from 1.01 at CB5.1W to 1.08 at LE1.1. Maximum absolute annual anomalies range from 0.45 psu at TF1.5 to 2.24 psu at TF1.7. This pattern reinforces the interpretation that the transitional portion of the estuary is especially sensitive to absolute salinity terms.

Bootstrap comparison of station groups supports this spatial contrast in anomaly magnitude. The mean absolute annual anomaly for the transitional and lower stations exceeded that of the upper stations by 0.61 psu, with a 95% bootstrap interval of 0.46 to 0.75 psu. Thus, the larger anomaly variability from TF1.7 downstream is not only a visual or descriptive pattern; it is robust under year-block resampling of the observational record.

The annual anomaly result also clarifies why a station-specific baseline is necessary. If annual means had been compared directly across stations, most of the signal would have been the mean spatial gradient. By subtracting each station's own day-of-year climatology before averaging, the anomaly analysis isolates years that are relatively fresh or saline for that station. This allows upper, transitional, and lower stations to be compared on a common conceptual scale: departure from local seasonal expectation.

The anomaly patterns do not imply that every station changes in the same direction with the same magnitude each year. Rather, they show that the salinity gradient has both common forcing influences and local response structure. Some years are broadly fresher or saltier across many stations, while other years show station-specific deviations. This is consistent with the chapter's broader argument that Patuxent salinity variability is spatially heterogeneous rather than a single estuary-wide signal.

Annual harmonic departures in seasonal structure

Annual harmonic fitting shows that interannual variability is not limited to annual mean offset. Under the default coverage rule, 314 of 351 retained station-years are departure-ready. Ready-year counts range from 25 at TF1.5 to 37 at RET1.1 and the lower-estuary stations. This coverage is sufficient to compare broad patterns in mean state and amplitude departure, but timing metrics remain more sensitive to the quality and distribution of within-year sampling. Median absolute mean-state departures increase from 0.20 psu at TF1.5 and 0.42 psu at TF1.6 to 1.10 psu at TF1.7. From RET1.1 through CB5.1W, the median absolute mean-state departure generally exceeds 1.25 psu. This pattern broadly matches the annual anomaly result: upper low-salinity stations show smaller absolute departures, while transitional and lower-estuary stations show larger interannual shifts in fitted seasonal mean state.

Median absolute amplitude departures show a parallel spatial pattern. TF1.5 has a median absolute amplitude departure of 0.27 psu, TF1.6 has 0.53 psu, and TF1.7 has 0.95 psu. From RET1.1 through CB5.1W, median absolute amplitude departures range from 0.87 to 1.06 psu. These results indicate that interannual variability is not only a vertical shift of the seasonal salinity curve. The strength of the seasonal cycle itself also varies from year to year.

Timing-related metrics are retained as descriptive outputs but are not used as the central interpretation. Median absolute phase departures range from about 26 to 33 days across stations, but phase-bearing station-years are sensitive to masking rules and to the fitted amplitude. When amplitude is weak, small changes in the fitted sine and cosine coefficients can produce large apparent timing shifts. For this reason, the chapter emphasizes mean-state and amplitude departures as the robust harmonic evidence.

Annual harmonic departures




Figure 2.4. Annual harmonic mean-state, amplitude, and timing departures for stations. Mean-state and amplitude departures carry the main seasonal-structure result.

The uncertainty analysis reinforces this interpretation. Mean-state and amplitude departures retain coherent station-level patterns under the bootstrap summaries, whereas phase/timing intervals become broad where the fitted seasonal amplitude is weak or where annual coverage is limited. Phase is therefore retained as contextual evidence of seasonal-cycle change, not as a primary inferential result.

Together, the annual anomaly and harmonic results show that the salinity gradient changes in multiple ways. Some interannual variability appears as stations being saltier or fresher than expected on average. Some appears as a stronger or weaker seasonal cycle. This matters for later modeling because a model that captures annual mean conditions but misses seasonal amplitude changes would still fail to reproduce an important part of observed salinity-gradient variability.

Discharge background, event anomalies, and salinity departures

Hydrologic classification was evaluated over the 1984-2022 salinity-discharge overlap. Annual mean discharge thresholds classify years as dry below 293.89 cfs and wet above 442.95 cfs. Within the overlap period, the classification contains 9 dry years, 21 regular years, and 9 wet years. Twelve years are identified as event-anomaly years: 1989, 1996, 2003, 2005, 2006, 2008, 2009, 2011, 2012, 2014, 2018, and 2020.

The discharge thresholds are quartile-based rather than manually selected: dry and wet years represent the lower and upper 25% tails of complete annual mean discharge, while regular years form the central background class. Sensitivity checks using alternative threshold schemes showed that exact class counts can change, but the main wet-year tendency toward negative salinity anomalies remains the strongest and most physically consistent contrast.




Figure 2.5. Annual discharge background classes and event-anomaly years over the salinity-discharge overlap.

Table 2.2. Observed day-of-year anomaly contrasts by hydrologic background and event-anomaly status.

Contrast	Years	Station-years	Mean anomaly	Median anomaly
Background: dry	9	81	-0.10	0.05
Background: regular	21	189	0.20	0.03
Background: wet	9	81	-0.36	-0.25
Event anomaly: no	27	243	0.12	0.05
Event anomaly: yes	12	108	-0.26	-0.16

Observed station-year salinity anomalies differ by hydrologic context. Across 351 retained station-years, mean anomaly is -0.10 psu in dry years, 0.20 psu in regular years, and -0.36 psu in wet years. Event-anomaly years average -0.26 psu, compared with 0.12 psu in non-event years. The negative mean anomaly in wet and event-anomaly contexts is consistent with the expectation that stronger freshwater forcing shifts the system toward fresher-than-normal salinity conditions.

Year-level uncertainty estimates qualify the strength of these contrasts. Wet years were fresher than regular years by -1.04 psu, with a 95% interval of -1.89 to -0.14 psu and a year-level permutation p-value of 0.054. Dry years differed from regular years by -0.54 psu, but that contrast was more uncertain, with a 95% interval of -1.91 to 0.58 psu and a permutation p-value of 0.363. Event-anomaly years also showed negative anomalies in the unadjusted summaries, but after accounting for discharge class and station, the adjusted event effect was -0.70 psu with a 95% interval of -1.54 to 0.09 psu. Event status is therefore best interpreted as related to wet-year and high-flow structure rather than as a clearly independent driver.




Figure 2.6. Mean observed day-of-year salinity anomaly by station and discharge background.

The discharge-context response is spatially uneven. Wet-year mean anomalies are close to zero at TF1.5 (-0.03 psu) and TF1.6 (-0.10 psu), become more negative at TF1.7 (-0.29 psu), and range from -0.51 to -0.40 psu from RET1.1 through CB5.1W. Event-anomaly years show the same broad pattern: TF1.5 averages -0.04 psu, TF1.6 averages -0.15 psu, TF1.7 averages -0.34 psu, and retained downstream stations range from -0.43 to -0.22 psu.

These results support two interpretations. First, freshwater forcing helps organize salinity departures: wet and event-anomaly years are generally associated with fresher-than-expected conditions. Second, the response is not uniform along the estuary. The uppermost stations have smaller absolute anomalies, likely because their baseline salinity is already low. Transitional and downstream stations show larger negative departures under wet and event-anomaly contexts, indicating that freshwater forcing is expressed most clearly where sufficient salinity is present for a negative departure to be observed.




Figure 2.7. Mean observed day-of-year salinity anomaly by station and event-anomaly status.

The classification analysis should not be interpreted as a complete causal model. Dry years do not simply produce uniformly positive anomalies, and regular years can still include positive or negative station-level departures. This reflects the fact that annual mean discharge and event structure are coarse descriptors of freshwater forcing, while observed salinity also responds to seasonal timing, downstream boundary salinity, tidal mixing, wind forcing, and spatial position. The classification result is therefore best interpreted as evidence that freshwater context organizes part of the salinity-gradient variability, not as proof that discharge alone explains all anomalies.

Reduced-complexity forcing predictor

The reduced-complexity forcing workflow produced 350,640 30-minute records for 2001-2020. Water-level coverage was 98.88%, discharge coverage was 98.71%, and valid Q_{INL} coverage was 97.74%. The water-level record ranged from -0.908 to 1.122 m. Discharge was converted from cubic feet per second to cubic meters per second and had a median of 7.39 m³/s. The median absolute tidal exchange term, $A_E dh_T/dt$, was 2015.21 m³/s, placing the exchange term in the 10³ m³/s range.

Table 2.3. Forcing construction and reduced-complexity predictor diagnostics.

Quantity	Value
30-minute records	350,640
Water-level coverage	98.88%
Discharge coverage	98.71%
Valid Q_{INL} coverage	97.74%
Water-level range	-0.908 to 1.122 m

Quantity	Value
Median discharge	7.39 m ³ /s
Median abs(Q_{tide})	2015.21 m ³ /s
Minimum mean 30-day joint DOY coverage	0.930
$Q_R > 0, 30\text{d}$ range	7.27 to 16.12 m ³ /s
$Q_{\text{INL}} > 0, 30\text{d}$ range	908.07 to 1030.38 m ³ /s
Q^* range	0.00712 to 0.01728
$S_{E,ss}^*$ range	0.98301 to 0.99293

The average-year predictor used the same 365-day no-leap axis as the salinity climatology. The mean 30-day joint coverage by day of year was at least 0.930, and all 365 days were retained. The 30-day freshwater term ranged from 7.27 m³/s on day 248 to 16.12 m³/s on day 83. The 30-day positive inlet term ranged from 908.07 m³/s on day 21 to 1030.38 m³/s on day 207. The resulting freshwater-to-inlet ratio Q^* ranged from 0.00712 on day 248 to 0.01728 on day 83. The estuary-scale salinity predictor $S_{E,ss}^*$ ranged from 0.98301 on day 83 to 0.99293 on day 248.

Reduced-complexity forcing climatology




Figure 2.8. Reduced-complexity forcing climatology on the 365-day no-leap day-of-year axis. The panels show 30-day $Q_R > 0$, 30-day $Q_{\text{INL}} > 0$, Q^* , and $S_{E,ss}^*$.

The narrow seasonal range of $S_{E,ss}^*$ is important for interpreting the transfer-model coefficients. Because the predictor varies over less than 0.01 units in the main area scenario, fitted slopes can appear numerically large. These slopes should therefore be interpreted as response coefficients over a narrow seasonal forcing range, rather than as physical salinity gradients along distance.

Normalized station salinity climatology

Station salinity climatology and the reduced-complexity forcing predictor were aligned on the same 365-day no-leap axis. CB5.1W was used as the downstream marine boundary for normalization. Its climatological salinity ranged from 11.02 to 15.73 psu. No station had normalized salinity greater than 1.0. TF1.5 had 30 small negative harmonic-fit normalized values, and these values were set to zero for the transfer-model fit.

Table 2.4. Station salinity climatology ranges used in the normalized transfer model.

Station	Distance (km)	Salinity range (psu)	S_i^* range	Boundary station
TF1.5	0.00	-0.01 to 0.65	0.0000 to 0.0450	No
TF1.6	6.00	0.08 to 1.75	0.0067 to 0.1184	No
TF1.7	14.47	2.17 to 5.92	0.1891 to 0.3911	No
RET1.1	24.69	7.31 to 11.36	0.6576 to 0.7283	No
LE1.1	33.79	9.10 to 13.25	0.8199 to 0.8487	No
LE1.2	43.22	9.97 to 14.15	0.8905 to 0.9138	No
LE1.3	47.90	10.25 to 14.42	0.9070 to 0.9403	No
LE1.4	54.53	10.81 to 15.23	0.9678 to 0.9803	No
CB5.1W	58.78	11.02 to 15.73	1.0000 to 1.0000	Yes




Figure 2.9. Normalized station salinity climatology, $S_{i(d)}^* = \frac{S_{i(d)}}{S_{CB5.1W(d)}}$, ordered from upstream to downstream.

Normalized station salinity increased downstream from TF1.5 to CB5.1W. TF1.5 ranged from 0.0000 to 0.0450 after adjustment, TF1.6 ranged from 0.0067 to 0.1184, TF1.7 ranged from 0.1891 to 0.3911, and RET1.1 ranged from 0.6576 to 0.7283. The lower-estuary stations occupied a narrower normalized range: LE1.1 ranged from 0.8199 to 0.8487, LE1.2 from 0.8905

to 0.9138, LE1.3 from 0.9070 to 0.9403, and LE1.4 from 0.9678 to 0.9803. CB5.1W was 1.0000 by definition.

This normalized climatology shows the spatial compression of the lower-estuary salinity gradient. The upstream and transitional stations cover a large normalized range, while the lower stations are clustered near the downstream boundary. This pattern explains why lower-estuary stations can have small normalized errors but still show meaningful salinity variations in psu after reconstruction.

Station-transfer fits

The transfer model was fitted for 365 day-of-year values at each station. Slopes were positive from TF1.5 through RET1.1 and negative from LE1.1 through LE1.4. The highest internal-station slope was at TF1.7 ($a_i = 23.59$), followed by TF1.6 ($a_i = 13.15$) and RET1.1 ($a_i = 8.03$). The fitted slope changed sign between RET1.1 ($a_i = 8.03$) and LE1.1 ($a_i = -1.23$). The most negative slope occurred at LE1.3 ($a_i = -3.87$).

A 30-day circular block bootstrap supports the RET1.1-to-LE1.1 slope contrast, but the individual negative slope at LE1.1 should be interpreted cautiously. RET1.1 had a positive slope of 8.03 with a 95% bootstrap interval of 5.96 to 10.21, whereas LE1.1 had a negative point estimate of -1.23 with a 95% interval of -2.53 to 0.32. The bootstrap probability that RET1.1 was positive while LE1.1 was negative was 0.942, and the RET1.1 minus LE1.1 slope difference had a 95% interval of 6.77 to 11.88. The evidence therefore supports a strong slope contrast across this portion of the estuary, but not a stand-alone claim that the LE1.1 negative slope is statistically robust.

Table 2.5. Station-transfer parameters for $S_{i(d)}^* = a_i S_{(E,ss)(d)}^* + b_i$

Station	Dist (km)	b_i	a_i	R^2	RMSE	MAE	psu MAE
TF1.5	0.00	-5.52	5.61	0.891	0.0053	0.0041	0.06
TF1.6	6.00	-12.92	13.15	0.830	0.0162	0.0133	0.18
TF1.7	14.47	-23.00	23.59	0.817	0.0304	0.0255	0.34
RET1.1	24.69	-7.23	8.03	0.772	0.0119	0.0099	0.14
LE1.1	33.79	2.05	-1.23	0.109	0.0096	0.0085	0.11
LE1.2	43.22	3.43	-2.56	0.722	0.0043	0.0032	0.04
LE1.3	47.90	4.74	-3.87	0.810	0.0051	0.0040	0.05
LE1.4	54.53	1.92	-0.96	0.356	0.0035	0.0030	0.04
CB5.1W	58.78	1.00	0.00	NA	0.0000	0.0000	0.00




Figure 2.10. Station-transfer scatter fits. Points show day-of-year climatological S_i^* values against $S_{E,ss}^*$; lines show the station-specific linear transfer fits.

The intercepts separate the estuary into upper and lower station groups. Upper-estuary intercepts were negative: TF1.5 had $b_i = -5.52$, TF1.6 had $b_i = -12.92$, TF1.7 had $b_i = -23.00$, and RET1.1 had $b_i = -7.23$. Lower-estuary intercepts were positive: LE1.1 had $b_i = 2.05$, LE1.2 had $b_i = 3.43$, LE1.3 had $b_i = 4.74$, and LE1.4 had $b_i = 1.92$. CB5.1W had $b_i = 1.00$, slope near zero, and undefined R^2 because the normalized boundary series was constant at 1.0.

The highest internal-station R^2 values occurred at TF1.5 (0.891), TF1.6 (0.830), TF1.7 (0.817), and LE1.3 (0.810). RET1.1 had $R^2 = 0.772$, LE1.2 had $R^2 = 0.722$, LE1.4 had $R^2 = 0.356$, and LE1.1 had $R^2 = 0.109$. The low R^2 at LE1.1 does not necessarily imply a large absolute salinity error; its observed normalized range is small, so modest residuals can explain little of the limited variance.

Transfer parameters by station distance




Figure 2.11. Transfer parameters by distance from the upstream retained station. Panels show a_i , b_i , R^2 , and MAE.

Normalized MAE was largest at TF1.7 (0.0255), followed by TF1.6 (0.0133) and RET1.1 (0.0099). On the salinity scale reconstructed by multiplying fitted S_i^* by CB5.1W climatological salinity, MAE was 0.34 psu at TF1.7, 0.18 psu at TF1.6, 0.14 psu at RET1.1, 0.11 psu at LE1.1, 0.06 psu at TF1.5, 0.05 psu at LE1.3, 0.04 psu at LE1.2, and 0.04 psu at LE1.4. These values indicate that the one-predictor transfer model reproduces much of the seasonal normalized structure, but the transitional stations remain the most difficult to represent.

Bootstrap fit-metric intervals support this interpretation more strongly for residual concentration than for any single R^2 ranking. TF1.7, TF1.6, and RET1.1 were consistently among the highest-residual internal stations across normalized MAE and residual range diagnostics, indicating that the model deficits are spatially concentrated rather than evenly distributed along the retained station sequence.




Figure 2.12. Observed and fitted normalized salinity day-of-year curves. The fitted curve is $a_i S_{(E,ss)(d)}^* + b_i$.




Figure 2.13. Observed and fitted salinity day-of-year curves in psu, reconstructed as fitted $S_{l(d)}$ multiplied by $S_{CB5.1W(d)}$.

The sign change in the fitted slope is a key diagnostic result. Upstream and transitional stations have positive slopes, meaning their normalized salinity climatology increases with the estuary-scale saline-favoring predictor. Lower-estuary stations have negative slopes over the fitted seasonal range. This does not mean that freshwater directly makes the lower estuary saltier. Rather, it reflects the seasonal relationship between local normalized station salinity and the downstream boundary-normalized predictor. Because lower-estuary stations remain close to the boundary value, their normalized seasonal behavior is shaped by both local salinity changes and the seasonality of CB5.1W itself. The transfer model therefore reveals where the simple freshwater-to-inflow index aligns with local climatology and where boundary normalization and additional processes complicate interpretation.

Counterfactual decomposition of seasonal forcing controls

The station-transfer model was also used for a set of physical counterfactuals in which the fitted station parameters were held fixed while one forcing component was replaced by its long-term mean. These experiments are not a strict additive variance partition, because downstream boundary salinity, freshwater discharge, and inlet or tidal exchange are physically coupled. Instead, they ask how much of the practical climatological prediction skill and predicted salinity

range is lost when one component no longer varies seasonally within the fitted reduced-complexity structure.

Holding downstream boundary salinity fixed produced the largest change in climatology-scale performance. Across internal stations, the fixed_SM scenario increased mean RMSE by 0.92 psu relative to the baseline prediction and retained only 0.38 of the predicted baseline salinity range. The effect was strongest in the lower estuary: the largest fixed_SM increases in RMSE occurred at LE1.4 (1.53 psu), LE1.3 (1.48 psu), and LE1.2 (1.47 psu). This indicates that seasonal variation in the downstream river boundary is a dominant control on lower-estuary salinity variability.




Figure 2.14. Climatology counterfactual observed and predicted salinity curves. Observed salinity is compared with the baseline transfer prediction and the fixed_SM, fixed_Q_R, and fixed_Q_INL forcing counterfactuals with station-transfer parameters held fixed.

Holding freshwater discharge fixed produced a smaller but more spatially diagnostic effect. Across internal stations, the fixed_Q_R scenario increased mean RMSE by 0.14 psu and retained 0.72 of the predicted baseline salinity range. The largest losses were concentrated farther upstream and in the transitional section, especially TF1.7 (RMSE increase = 0.45 psu), TF1.6 (0.25 psu), and RET1.1 (0.15 psu). In contrast, holding the positive inlet-flow term fixed had little effect at the climatology scale: fixed_Q_INL increased internal-station mean RMSE by

only 0.01 psu and retained 0.97 of the baseline range.

Counterfactual sensitivity by station distance




Figure 2.15. Climatology counterfactual sensitivity by station distance for fixed downstream salinity, fixed freshwater discharge, and fixed positive inlet flow. Panels show RMSE change relative to baseline and retained predicted salinity range.

The counterfactuals therefore separate two components of the seasonal salinity structure. Downstream boundary salinity primarily controls the lower-estuary seasonal salinity range, while freshwater discharge contributes most strongly to the upper and transitional stations. The weak fixed_Q_INL response does not imply that tides or exchange are unimportant in the real estuary. Rather, it shows that, within this 30-day-smoothed climatological implementation, seasonal variation in the constructed inlet-flow term contributes less practical predictive variation than either the downstream salinity boundary or the freshwater discharge.

Geometry-linked attenuation of discharge-driven salinity response

The TF1.7-to-RET1.1 transition provides a focused test of whether the discharge-driven salinity response is filtered by local estuarine geometry. These adjacent stations lie across the upper-transition portion of the retained sequence, but their estimated cross-sectional areas differ sharply. The calculated area increases from 4,487 m² at TF1.7 to 39,914 m² at RET1.1, an 8.90-

fold downstream expansion. Because the first-order advective term scales with Q_R/A^* , this expansion should strongly reduce the local velocity associated with the same upstream freshwater discharge.

The theoretical no-mixing calculation is consistent with this expectation. The mean theoretical salinity decreases at RET1.1 is only 0.094 times the value at TF1.7, and the 95th-percentile theoretical decrease ratio is similarly small at 0.101. Expressed in the opposite direction, the theoretical discharge-driven decrease is approximately 10-fold larger at TF1.7 than at RET1.1. This close inverse correspondence between the area ratio and the theoretical salinity-decrease ratio indicates that the TF1.7-to-RET1.1 transition is not only a salinity-gradient feature, but also a geometry-linked attenuation region for local freshwater-advection effects.




Figure 2.16. Geometry-linked attenuation across the TF1.7-to-RET1.1 transition. The left panel shows RET1.1-to-TF1.7 ratios for cross-sectional area, theoretical first-order salinity decrease, and fixed_SM predicted range. The right panels show the corresponding station-distance profiles.

The fixed_SM counterfactual shows a related but less extreme attenuation. Under the scenario in which downstream marine salinity is held fixed, the predicted seasonal salinity range at RET1.1 is 0.34 times the corresponding range at TF1.7. This reduction is weaker than the approximately 10-fold theoretical reduction in the first-order advective decrease, as expected, because the station-transfer counterfactual includes seasonal structure beyond local Q_R/A^* advection. Even

so, both diagnostics point in the same direction: the non-boundary seasonal response retained under fixed downstream salinity is substantially smaller at RET1.1 than at TF1.7 (Figure 2.16). Together, the counterfactual and cross-sectional analyses suggest that the TF1.7-to-RET1.1 transition marks a shift from an upper-estuary zone where discharge-driven local advection can be large to a wider downstream reach where the same freshwater forcing is attenuated by cross-sectional expansion and increasingly embedded within exchange, storage, and boundary-condition effects. The transition is therefore better interpreted as a short hydrogeomorphic interval over which the expression of freshwater forcing changes, rather than as an isolated station anomaly.

Effective-area sensitivity of the forcing predictor

The reduced-complexity analysis also evaluates sensitivity to smaller effective estuary area values. The original sensitivity table summarizes how the forcing predictor changes with area, and the refitted station-transfer analysis extends this check by refitting the station-transfer models under the same area scenarios. These refits are used as diagnostics of scaling sensitivity rather than as a calibration of the effective area. Reducing A_E lowers the 30-day positive inlet-flow term and expands the seasonal range of $S_{E,ss}^*$. At the main area of 122.962 km², $S_{E,ss}^*$ ranged from 0.9830 to 0.9929, a range of 0.0099. At 61.481 km², the range increased to 0.0196. At 30.741 km², it increased to 0.0382. At 12.296 km², it increased to 0.0885. At 6.148 km², it increased to 0.1559.

The refitted transfer models show that reducing effective area mainly rescales the fitted slopes because the seasonal range of $S_{E,ss}^*$ expands. The RET1.1 slope decreased from 8.03 at 100% A_E to 0.51 at 5% A_E , while LE1.1 remained slightly negative across the same area scenarios. Fit metrics and residual rankings changed much less than the slope magnitudes, indicating that the area assumption strongly affects coefficient scale but does not remove the main diagnostic contrast between the transitional stations and the lower-estuary normalized behavior.

Table 2.6. Forcing-only sensitivity of $S_{E,ss}^*$ to effective estuary area.

Area scenario	Area (km ²)	$S_{E,ss}^*$ min	$S_{E,ss}^*$ max	Full range	5-95% range	30d Median $Q_{INL>0}$, (m ³ /s)
100% A_E	122.962	0.9830	0.9929	0.0099	0.0083	972.78
50% A_E	61.481	0.9663	0.9859	0.0196	0.0163	483.33
25% A_E	30.741	0.9338	0.9721	0.0382	0.0318	239.20
10% A_E	12.296	0.8432	0.9317	0.0885	0.0738	92.51
5% A_E	6.148	0.7125	0.8684	0.1559	0.1302	43.46




Figure 2.17. Forcing-only sensitivity of $S^*_{E,ss}$ to smaller effective estuary area values.

The sensitivity result demonstrates that the effective area assumption strongly controls the scale of the reduced-complexity predictor and, consequently, the numerical scale of the fitted transfer slopes. Smaller effective areas make the same water-level changes correspond to smaller inlet-flow magnitudes, increasing the relative weight of freshwater discharge in Q^* . This expands the seasonal predictor range and reduces slope magnitude when the station-transfer model is refit. For the purposes of this chapter, the main-area result is retained as the primary diagnostic, while the sensitivity analysis shows that coefficient magnitude should not be overinterpreted as a calibrated physical parameter.

Integrated result

Across climatology, annual anomalies, harmonic departures, hydrologic classification, and reduced-complexity transfer modeling, a consistent structure emerges. Patuxent salinity is spatially heterogeneous variability around station-specific climatology. Mean state and seasonal amplitude define distinct station regimes along the estuary. Annual anomalies move above and below local day-of-year baselines rather than following a single monotonic trend. Annual harmonic departures show that seasonal structure changes in mean state and amplitude, not only in annual average salinity. Discharge background and event anomalies organize part of the anomaly structure, especially through negative departures in wet and event-anomaly contexts, but the response differs by station.

The reduced-complexity transfer analysis adds a physical-response diagnostic to the observational results. A simple seasonal freshwater-to-inflow predictor can reproduce a

substantial portion of normalized climatological variation at several stations, especially TF1.5, TF1.6, TF1.7, RET1.1, LE1.2, and LE1.3. At the same time, the largest normalized errors and residual ranges occur in the upper-transitional section, especially TF1.6, TF1.7, and RET1.1. The counterfactual experiments clarify this spatial pattern: downstream boundary salinity dominates lower-estuary seasonal range, whereas fixed-discharge losses are concentrated around TF1.6, TF1.7, and RET1.1. The cross-sectional comparison between TF1.7 and RET1.1 further indicates that local geometry can attenuate the first-order discharge-driven response by about an order of magnitude across a short reach. This pattern suggests that the same part of the estuary that defines the strongest salinity transition is also the region where freshwater forcing, geometry, and exchange are most difficult to summarize with a one-predictor seasonal index. The added uncertainty analysis strengthens this integrated result. The station-scale transition into TF1.7, the larger anomaly variability from TF1.7 downstream, the wet-year tendency toward negative anomalies, and the concentration of transfer-model residuals around TF1.6-TF1.7-RET1.1 are all supported by bootstrap or year-level resampling evidence. In contrast, event-anomaly independence, phase/timing shifts, and the negative LE1.1 slope should be treated as more cautious or exploratory interpretations.

The chapter therefore establishes the observational foundation for the thesis: freshwater forcing is relevant, but its effect is expressed through a spatially structured salinity gradient rather than through a uniform estuary-wide response. The Patuxent estuary contains low-salinity upper stations, a transitional region with large seasonal and interannual variability, and lower-estuary stations closer to the downstream boundary. Later chapters should therefore be evaluated not only by overall predictive accuracy, but also by whether they reproduce this spatial organization, the station-specific seasonal cycle, discharge-related negative departures in wet and event-anomaly contexts, and the difficulty of representing the transitional salinity zone.

Discussions

Contribution to estuarine-gradient interpretation

The main contribution of this chapter is not the well-established increase in salinity from upstream to downstream. That pattern is a classical feature of positive estuaries and has been central to estuarine physics since early work on Chesapeake Bay salinity structure (Pritchard 1952). The contribution is methodological and interpretive: the Patuxent salinity gradient is used as the reference frame for interpreting variability. Station-specific climatology converts the familiar salinity gradient into an analytical baseline, so each observation is evaluated against the seasonal expectation for its own position along the estuary rather than against a pooled estuary-wide mean.

This framing extends prior Patuxent work in a specific way. Earlier studies treated the Patuxent as a spatially structured estuary, with tidal-fresh, riverine-estuarine transition, and lower-estuary regions that organize water-quality and oxygen dynamics (Testa et al. 2008; Dreiss et al. 2024).

The present chapter places salinity itself at the center of that structure. It shows that the retained monitoring stations are not repeated samples of a single salinity signal. They are positions along a dynamic river-to-bay transition whose local seasonal expectations differ.

The station-scale transition near TF1.7 and RET1.1 is therefore more than a descriptive break in the climatology. It marks the part of the retained sequence where the physical meaning of a salinity change is most sensitive to location. In estuarine physics, salinity structure reflects the interaction of freshwater input, salt intrusion, exchange, mixing, and geometry (Geyer 2010; MacCready and Geyer 2010). In the Patuxent monitoring record, this interaction appears observationally as the short reach where the mean gradient sharpens, anomaly variability increases, and reduced-complexity residuals become largest.

This interpretation links salinity to broader environmental structure without claiming more than the data show. Chesapeake Bay studies commonly use salinity zones to interpret habitat and water-quality gradients, including submerged aquatic vegetation habitat requirements and Patuxent oxygen dynamics (Kemp et al. 2004; Dreiss et al. 2024). This chapter does not test biological responses directly. Instead, it provides the salinity template around which later ecological or biogeochemical interpretations would need to be organized.

Seasonal and interannual departures in a sparse monitoring record

A second contribution is the separation of interannual variability into mean-state departures and seasonal-cycle departures. Long-term coastal monitoring records can support seasonal and multi-decadal questions, but sparse sampling limits event-scale inference (Brahmey et al. 2024). The Patuxent record has this same strength and limitation. It is long enough to define station-specific seasonal baselines and annual departures, but it should not be treated as a complete daily salinity record.

Within that boundary, the harmonic analysis adds information that annual means alone would miss. Seasonal salinity cycles are widely recognized in oceanic and coastal salinity records, and harmonic approaches are useful for summarizing mean state, amplitude, and timing (Bingham et al. 2021; Yu 2021; Brahmey et al. 2024). The Patuxent results extend this approach into a station-specific estuarine setting. They show that interannual variability includes changes in the fitted seasonal cycle, not only shifts in annual average salinity.

This distinction matters for interpretation and modeling. Seasonal amplitude can reflect changes in the annual range of freshwater influence, downstream salinity influence, mixing, or boundary conditions. A model that reproduces annual mean salinity but suppresses seasonal amplitude would therefore miss an uncertainty-supported component of the observed structure. By contrast, phase or timing estimates should remain secondary because they are more sensitive to sparse within-year sampling and weak fitted amplitude (Bingham et al. 2021; Brahmey et al. 2024).

Freshwater forcing as a spatially filtered control

The discharge-related results refine the expected freshwater-salinity relationship rather than simply confirming it. Estuarine theory predicts that increased river flow tends to freshen estuaries and displace the salinity field seaward, but the response depends on adjustment time, salt storage, tidal mixing, geometry, and boundary conditions (MacCready 1999; MacCready and Geyer 2010). The Patuxent results are consistent with this expectation because wet years tend to have negative salinity anomalies relative to local seasonal expectation. They also show that discharge acts through the spatial structure of the gradient rather than as a uniform estuary-wide forcing.

The station dependence of the anomaly response is the key point. Upper stations are expected to be strongly influenced by freshwater, yet show small negative values for salinity anomalies because their baseline salinity is already near zero salinity conditions. Transitional and lower-estuary stations have more salts available to be displaced, so the same broad freshwater context can appear as a larger absolute anomaly. The observed discharge response is therefore filtered by station position along the salinity gradient (Gibson and Najjar 2000; Monismith et al. 2002).

The incomplete mapping between hydrologic class and salinity anomaly is also informative. In Chesapeake Bay, salinity residuals have been used to identify influences beyond river flow, including sea-level-related and boundary-condition effects (Hilton et al. 2008). Recent work in other estuaries similarly shows that salinity and residence-time responses can vary by location and forcing scale (Agonsi et al. 2026). The Patuxent anomaly analysis fits this broader view. Discharge is an organizing context, but it is not a complete explanation for every station-year departure.

Event-anomaly status should be interpreted within the same boundary. It is useful as a descriptor of irregular high-flow structure, but the present evidence does not justify treating it as an independent causal factor once annual discharge background and station position are considered. With the available low-frequency salinity sampling, event structure is best used as context rather than as a standalone mechanism (Brahmey et al. 2024; MacCready 2007).

Counterfactual forcing controls and spatial separation of mechanisms

The counterfactual analysis extends the discharge interpretation by separating where different seasonal controls matter within the fitted reduced-complexity framework. Holding downstream salinity fixed caused the largest loss of climatology-scale performance in the lower estuary. This result is physically reasonable because lower-estuary stations remain closer to the Chesapeake Bay boundary and inherit much of their seasonal salinity range from marine boundary variability. In contrast, holding freshwater discharge fixed caused its largest losses at TF1.7, TF1.6, and RET1.1, placing the discharge-sensitive signal farther upstream and in the transitional section (MacCready 2004; Ralston et al. 2008).

This spatial separation is the most important contribution of the counterfactuals. It shows that boundary salinity and freshwater discharge do not simply add to one estuary-wide response.

Instead, they dominate different parts of the station sequence. The result therefore strengthens the chapter argument that salinity variability is organized by position along the gradient: lower-estuary variability is more boundary-mediated, whereas upper and transitional variability retains a stronger freshwater-discharge signature (MacCready 2004; Monismith et al. 2002; Ralston et al. 2008).

The weak fixed_Q_INL response must be interpreted carefully. The positive inlet-flow term is derived from water-level change and smoothed to a 30-day climatological scale, so it is not a complete representation of tidal exchange, wind-driven circulation, stratification, or two-layer flow. Its small counterfactual effect therefore does not mean that tides or exchange are unimportant in the real estuary. It means that this particular seasonal inlet-flow index adds little independent predictive variation after discharge, downstream salinity, and station-transfer parameters are included (Dyer 1997; Geyer and MacCready 2014; MacCready et al. 2018).

This distinction matters because reduced-complexity models are useful only when their simplifications remain visible. Recent reduced-complexity salinity modeling emphasizes that simple predictors can capture seasonal salinity variation with interpretable inputs, but they do not replace fully resolved hydrodynamics (Yavary Nia et al. 2025). The Patuxent counterfactuals follow that logic. They are diagnostic experiments that identify where the fitted framework is sensitive to each forcing component, not a complete partition of estuarine salt transport.

Hydrogeomorphic interpretation of the TF1.7-to-RET1.1 transition

The cross-sectional analysis provides a hydrogeomorphic explanation for why the discharge signal weakens across the TF1.7-to-RET1.1 transition. The same upstream discharge produces different local advective tendencies depending on the cross-sectional area through which it is distributed. In the Patuxent sequence, estimated cross-sectional area increases 8.90-fold from TF1.7 to RET1.1. Because the first-order diagnostic scales with Q_R/A^* , this expansion strongly reduces the local discharge-driven tendency for salinity decrease (Dyer 1997; Hansen and Rattray Jr 1966; MacCready 2004).

The no-mixing diagnostic supports this interpretation. The mean theoretical salinity decreases at RET1.1 is only 0.094 of the value at TF1.7, and the 95th-percentile ratio is 0.101. These values closely match the inverse of the area expansion and indicate that the TF1.7-to-RET1.1 transition is not only a salinity-gradient feature. It is also a geometry-linked attenuation point for the local expression of freshwater forcing (MacCready 2004).

The fixed_SM counterfactual shows a related but less extreme pattern. When downstream salinity is held fixed, the retained predicted seasonal range at RET1.1 is about one-third of the corresponding range at TF1.7. This attenuation is weaker than the roughly one-tenth theoretical advective ratio because the transfer counterfactual includes seasonal structure beyond local advection. The difference between these ratios is useful: local freshwater advection is strongly attenuated by the widening cross-section, while the remaining seasonal signal likely reflects exchange-mediated, storage-mediated, and boundary-mediated controls (MacCready 2004; Ralston et al. 2008).

This hydrogeomorphic interpretation extends the chapter beyond a station-by-station description. The TF1.7-to-RET1.1 reach is better understood as a short interval where the form of freshwater influence changes. Upstream, discharge can produce a larger local advective tendency because the section is narrower. Downstream, the same discharge is distributed through a wider cross-section and becomes more embedded within exchange, storage, and downstream boundary effects. Later modeling should therefore treat this reach as a process transition, not as an isolated station anomaly (Dyer 1997; Geyer and MacCready 2014).

Normalized fit, salinity-scale accuracy, and model-evaluation logic

The distinction between normalized fit and reconstructed salinity accuracy further clarifies how models should be evaluated. Normalized transfer performance asks whether local station salinity, expressed relative to the downstream boundary, follows the seasonal reduced-complexity predictor. Salinity-scale accuracy asks whether the reconstructed value is close to observed salinity in psu. These are related but not equivalent questions (Roberts et al. 2017).

This distinction matters because metric behavior depends on station position. A station with a narrower normalized range can have low R^2 even when the reconstructed salinity error is small. A station with a stronger normalized relationship can still have larger psu-scale error if it occupies a more variable part of the gradient. The chapter therefore supports a multi-diagnostic evaluation strategy using slope, intercept, R^2 , normalized residuals, and salinity-scale MAE together (Roberts et al. 2017).

The sign change in fitted slopes illustrates the same caution. Positive slopes at upper and transitional stations are consistent with increasing normalized salinity under more saline-favoring seasonal conditions. Negative or near-zero lower-estuary slopes should not be translated directly into a claim that freshwater forcing increases lower-estuary salinity. Because the normalization uses CB5.1W as the downstream boundary, lower-estuary ratios partly reflect the seasonal behavior of that boundary station itself (Monismith et al. 2002; MacCready 2004).

This evaluation logic is central for the thesis. Later models should not be judged only by overall RMSE, MAE, or visual fit. They should also be judged by whether they preserve station-specific climatology, reproduce supported wet-year freshening tendencies, distinguish boundary-mediated from discharge-mediated seasonal structure, and reduce residual structure in the transitional zone without erasing the observed gradient (Roberts et al. 2017; Gorski et al. 2024).

Limitations and boundary of inference

Several limitations define the boundary of these interpretations. First, the salinity record is long but sparse. It supports climatology, annual anomaly, harmonic, and seasonal diagnostic analyses, but it does not support event-scale claims about the detailed timing of salinity adjustment after individual flow pulses. Such claims would require denser observations, explicit lag analysis, or process-based model output (Brahmey et al. 2024; MacCready 2007).

Second, the counterfactuals are diagnostic, not a strict additive variance partition. Downstream salinity, discharge, and exchange are physically coupled in the real estuary, and the fitted station-transfer model holds station parameters fixed while one input is replaced by its mean. The counterfactual results should therefore be read as evidence about sensitivity within the fitted framework, not as a complete decomposition of causality (MacCready and Geyer 2010; MacCready et al. 2018).

Third, the cross-sectional calculation is intentionally first-order. It uses Q_R/A^* and the observed salinity-gradient scale to ask how the same discharge would be filtered by local cross-section. It does not resolve vertical shear, lateral exchange, bathymetric detail, stratification, wind forcing, or salt storage. The close area-to-response correspondence across TF1.7 and RET1.1 is physically suggestive, but it should be interpreted as a scale diagnostic rather than a full momentum or salt budget (Dyer 1997; MacCready 2004).

Fourth, the reduced-complexity predictor depends on the assumed effective estuary area. The sensitivity analysis shows that changing effective area strongly affects the numerical scale of $S_{E,ss}^*$ and the fitted slopes, even though the major residual contrast remains. This means that slope magnitudes should not be interpreted as calibrated physical parameters (Yavary Nia et al. 2025).

Finally, the analysis is based on surface salinity. Surface observations are appropriate for the station-specific climatology and anomaly framework used here, but they do not directly resolve vertical stratification or bottom-water exchange. This matters because Patuxent oxygen studies and estuarine physics both emphasize the role of two-layer circulation, stratification, and exchange in water-quality structure (Testa et al. 2008; MacCready and Geyer 2010; Dreiss et al. 2024). The chapter identifies where salinity variability is organized, but it does not fully diagnose the vertical mechanisms behind that organization.

Chapter Conclusion

Taken together, the results extend existing understanding of Patuxent and Chesapeake estuarine variability by converting a known physical expectation into a station-specific observational framework. The broad expectation that freshwater discharge freshens estuaries is not the endpoint of the chapter. The main result is that freshwater influence is expressed unevenly across a persistent but dynamic salinity gradient.

The transitional portion of the Patuxent estuary is the central diagnostic region. It marks the strongest station-scale change in climatological salinity, shows larger annual and seasonal departures than the upper low-salinity stations, concentrates the largest reduced-complexity residuals, and shows geometry-linked attenuation of discharge-driven salinity response across TF1.7 and RET1.1. These independent lines of evidence converge on the same interpretation: the transitional zone is where simplified forcing representations are most challenged and where later modeling can add the most scientific value.

The chapter therefore provides more than a descriptive summary of salinity patterns. It establishes a benchmark for interpreting and modeling Patuxent salinity as spatially structured variability around station-specific seasonal expectation. The strongest claims are the station-specific gradient, the transition into TF1.7 and RET1.1, the larger anomaly variability from TF1.7 downstream, the wet-year tendency toward negative anomalies, the spatial separation of boundary and discharge counterfactual effects, the geometry-linked attenuation across TF1.7-to-RET1.1, and the concentration of reduced-complexity residuals in the transitional zone. More cautious claims include event-anomaly independence, phase or timing interpretation, inlet-flow counterfactual meaning, and the physical meaning of individual normalized transfer slopes.

Chapter 3: Machine-Learning Extension of the Reduced-Complexity Model

Chapter Introduction

Chapter 2 showed that salinity in the Patuxent River estuary is organized by a persistent upstream-to-downstream salinity structure, with systematic differences among the upstream low-salinity reach, the transition zone, and the downstream high-salinity reach. This framing is consistent with broader estuarine studies in which salinity integrates oceanic, riverine, and mixing controls (MacCready and Geyer 2010; Geyer and MacCready 2014), and with Patuxent research that treats the salinity gradient as an organizing environmental axis (Hartzell et al. 2010). Chapter 2 also showed that this spatial structure varies with seasonal phase and interannual hydrological variability. The present chapter asks whether this reduced-complexity understanding can be extended from pattern interpretation to prediction of raw station-day observations.

The chapter is not framed as a standalone machine-learning salinity-prediction exercise. Instead, the Multi-Layer Perceptron (MLP) is used as a nonlinear extension of the reduced-complexity framework developed in Chapter 2. The reduced-complexity descriptor S_E^* is retained as a physically interpretable estuary-scale input, while downstream boundary salinity, along-channel station position, and seasonal phase provide additional information that can interact nonlinearly with that descriptor.

The central question is: how does the MLP-based extension improve the direct application of the reduced-complexity model to the raw station-day observational dataset, and under what spatial or physical conditions is this improvement most evident? This question makes the reduced-complexity model as baseline for the chapter. Overall accuracy, station-level skill, residual patterns, and seasonal or annual diagnostics are interpreted in relation to that baseline rather than as independent machine-learning performance summaries.

The chapter proceeds as follows. Data and methods section describes the raw station-day dataset, the direct reduced-complexity transfer baseline, the MLP extension, and the evaluation metrics. Results section presents the results in the order required by the chapter's central question: baseline performance first, followed by overall MLP improvement, spatial patterns of improvement, and residual or physical-condition diagnostics. Discussion section interprets how the nonlinear extension changes the reduced-complexity representation and where additional physical information may be valuable. Chapter conclusion section summarizes the conclusions.

Data and Methods

Raw station-day observational dataset

The analysis uses the raw station-day observational dataset assembled for the Patuxent River estuary. For each date with valid forcing and same-day downstream boundary salinity, salinity observations at internal stations were paired with station position and reduced-complexity descriptors. The downstream or seaward boundary condition was represented by salinity at CB5.1W, while the internal prediction stations were ordered consistently from upstream to downstream: TF1.5, TF1.6, TF1.7, RET1.1, LE1.1, LE1.2, LE1.3, and LE1.4.

After joining observations with the available forcing and boundary-salinity information, the filtered dataset contained 4,998 station-day records on 567 dates, including 4,431 internal-station records. The sampled dates ranged from 1984-02-20 to 2022-12-19. The raw $S_E^*(t)$ range was compact (0.9598-0.9982), making the prediction problem deliberately stringent: the baseline asks how much raw salinity variability can be organized by one compressed hydrological-tidal index plus station-specific linear transfer.

Table 3.1. Raw-observation support for the direct reduced-complexity application

Quantity	Value
Station-day records after boundary join	4998
Internal-station records	4431
Unique sampled dates	567
Date range	1984-02-20 to 2022-12-19
$S_E^*(t)$ range	0.9598-0.9982

Direct reduced-complexity transfer baseline

Before introducing the MLP, the reduced-complexity framework from Chapter 2 was applied directly to raw station-day observations. This direct transfer serves as the physical baseline for Chapter 3. It asks how much raw salinity variation can be organized by the same physically interpretable descriptor before nonlinear feature interactions are added.

For each date t , the downstream boundary salinity $S_M(t)$ was defined from the same-day salinity observation at CB5.1W:

$$S_M(t) = S_{\{CB5.1W\}}(t)$$

For each internal station i , observed salinity was normalized by this downstream boundary salinity:

$$S_i^*(t) = \frac{S_i(t)}{S_M(t)}.$$

This normalization expresses internal station salinity as a fraction of the downstream boundary condition. It reduces the dominance of absolute salinity differences among stations and focuses

the transfer model on how each station varies relative to the seaward end of the retained estuarine gradient.

Station-specific linear transfer relations were then fitted between the reduced-complexity descriptor and normalized salinity:

$$S_i^*(t) = a_i SE^*(t) + b_i + \varepsilon_i(t),$$

where a_i is the station-specific response coefficient, b_i is the station-specific offset, and $\varepsilon_i(t)$ is the residual. Predicted normalized salinity was converted back to salinity units by multiplying by the same-day downstream boundary salinity:

$$\hat{S}_i(t) = \hat{S}_i^*(t) \cdot S_M(t)$$

The direct reduced-complexity baseline is intentionally simple, but it is not a naive benchmark. It retains the physical logic of Chapter 2: internal salinity is interpreted relative to a downstream boundary and an estuary-scale seasonal forcing descriptor. Its value lies in interpretability. If this direct transfer performs well, then a compact reduced-complexity representation is sufficient for that part of the estuary. If it performs poorly at particular stations or flow states, those failures identify where additional nonlinear context or missing physical information is needed.

MLP feature design

The MLP extension was designed to strengthen the direct reduced-complexity baseline without discarding its physical structure. The input features were selected to represent four categories of information.

First, same-day CB5.1W salinity was used to represent the downstream or seaward boundary condition. This feature anchors the internal-station prediction to the salinity state at the lower end of the retained estuarine sequence. Because lower-estuary salinity is expected to be more strongly connected to Chesapeake Bay boundary conditions, this feature is especially important for distinguishing broad salinity level across dates.

Second, normalized station distance was used to represent position along the upstream-to-downstream salinity gradient. Distance is not a mechanistic forcing variable by itself, but it encodes the spatial coordinate along which freshwater influence, boundary influence, channel geometry, and mixing regimes vary. Including distance allows one nonlinear model to learn a continuous mapping across stations rather than fitting separate station-specific transfer relations only.

Third, S_E^* was retained as the reduced-complexity physical descriptor from Chapter 2. This variable represents a compact seasonal-scale freshwater-to-inflow balance. Its inclusion maintains a conceptual link between the MLP and the reduced-complexity framework. The MLP is therefore not a purely empirical model built only from location and observed salinity boundary conditions; it explicitly retains the physical descriptor that motivated the baseline.

Fourth, seasonal phase was represented using cyclic day-of-year terms. Day of year was mapped to sine and cosine features so that the annual cycle is treated as circular rather than linear (Qiu and Wan 2013; Qi et al. 2021; Saccotelli et al. 2024). This avoids an artificial discontinuity

between late December and early January and allows the MLP to represent smooth seasonal structure. Seasonal phase is important because salinity in the Patuxent estuary is not only spatially organized; it also has recurring seasonal variation that differs by station. Together, these features define the MLP as a nonlinear spatial-boundary-seasonal extension of the reduced-complexity baseline. The scientific role of the model is to test whether interactions among boundary salinity, station position, seasonal phase, and SE^* can improve raw-observation prediction in the locations where direct reduced-complexity transfer is incomplete.

MLP architecture and training

The MLP was implemented as a feed-forward neural network for tabular regression (Jayasundara et al. 2020; Qi et al. 2021; He et al. 2020; Qi, He, Bai, Ding, Sandhu, Zhou, et al. 2022; Qi, He, Bai, Ding, Sandhu, Chung, et al. 2022; Saccotelli et al. 2024). The final architecture used three hidden layers with sizes (128, 256, 32), ReLU activation functions, and L-BFGS optimization. L2 regularization was applied with $\alpha = 1 \times 10^{-5}$. The maximum number of training iterations was 2,000, and the convergence tolerance was 1×10^{-3} . These settings were chosen to provide sufficient nonlinear flexibility for the small-to-medium observational dataset while avoiding an unnecessarily deep architecture.

The dataset was divided by date into an 80/20 random training-test split using a fixed random seed. The random holdout evaluates interpolation skill within the sampled observational domain. It does not test extrapolation to unseen hydrological years, unusual flow regimes, or future climate conditions. That limitation is addressed in the Discussion because it matters for environmental machine-learning applications (Roberts et al. 2017).

Predictions were constrained to be non-negative values after model fitting because negative salinity values are physically unrealistic. Model performance was evaluated only on observed station-day records. No unsampled dates were added to the evaluation dataset.

Table 3.2. MLP model hyperparameter configuration

Hyperparameter	Value	Description
Hidden layer structure	(128, 256, 32)	Three hidden layers, progressively tapered
Activation function	ReLU	$f(x) = \max(0, x)$
Optimizer	L-BFGS	Quasi-Newton method, suitable for small-to-medium datasets
Regularization α	1×10^{-5}	L2 regularization coefficient
Maximum iterations	2000	Training convergence upper bound
Convergence tolerance	1×10^{-3}	Loss function change threshold

Evaluation metrics and common-row comparison

Model performance was evaluated using mean absolute error (MAE), root mean squared error (RMSE), coefficient of determination (R^2), and bias. MAE measures the average absolute prediction error. RMSE gives greater weight to large errors and is therefore useful for detecting

poor performance in transition-zone or extreme-flow cases. R^2 measures the fraction of observed variation explained within the evaluated group. Bias is defined as mean predicted salinity minus mean observed salinity; positive bias indicates overprediction, while negative bias indicates underprediction.

These metrics were interpreted together rather than individually. This is necessary because station salinity ranges differ strongly along the estuary. An upstream station may have small absolute error but low R^2 because its observed salinity range is narrow. A transition-zone station may have larger RMSE but still be the most informative diagnostic location because salinity there varies over a wider range and responds more strongly to shifts in the gradient. Therefore, the chapter emphasizes spatial and hydrological patterns of error rather than only aggregate performance.

Two comparison designs were used.

a) Station-level comparison.

This is comparison between the direct reduced-complexity baseline and the MLP extension, ordered from upstream to downstream. It asks where along the estuarine gradient the nonlinear extension adds predictive value.

b) Flow-stratified common-row comparison.

In this analysis, only station-date rows with valid predictions from both the direct baseline and the MLP were retained. This ensures that differences between models reflect model behavior rather than differences in evaluated samples.

Freshwater-flow state was defined using terciles of Q_R on the common station-date rows. Low-flow conditions corresponded to values below the 33.3rd percentile ($Q_R < 6.877 \text{ m}^3 \text{ s}^{-1}$), middle-flow conditions to values between the 33.3rd and 66.7th percentiles ($6.877 \leq Q_R < 11.808 \text{ m}^3 \text{ s}^{-1}$), and high-flow conditions to values at or above the 66.7th percentile ($Q_R \geq 11.808 \text{ m}^3 \text{ s}^{-1}$). Flow-stratified performance was evaluated for each station-flow group and then summarized across flow states. This comparison is diagnostic rather than causal. It asks whether MLP improvement is associated with freshwater-flow state, not whether discharge alone causes the improvement.

SHAP-based model interpretation

To interpret the fitted MLP, SHAP values were computed for the trained prediction function (Lundberg and Lee 2017; Lundberg et al. 2020). SHAP values provide an additive feature-attribution decomposition of each model prediction (Lundberg and Lee 2017; Lundberg et al. 2020). For a given sample x , the MLP prediction $f(x)$ can be written as:

$$f(x) = E[f(X)] + \sum_j \phi_j,$$

where $E[f(X)]$ is the expected model output over the background dataset and ϕ_j is the SHAP value for feature j . Each ϕ_j represents the contribution of a feature to moving the prediction above or below the model's expected output for that sample.

Three SHAP diagnostics were used. First, mean absolute SHAP values were used to rank global feature contributions across the test set. This identifies which variables the MLP relied on most strongly overall. Second, SHAP summary plots were used to examine the direction of feature effects, showing whether high or low feature values tend to increase or decrease predicted salinity. Third, local waterfall plots were used to decompose representative individual predictions, especially cases in the transition zone where the MLP still produced large errors. SHAP values are interpreted as diagnostics of fitted model behavior. They do not establish causal effects in the estuary. For example, a large SHAP value for station distance does not mean that distance is a physical force; rather, it means that the trained MLP used station position strongly when producing predictions. This distinction is important because the model does not explicitly include lagged discharge, wind stress, tidal phase, stratification, or antecedent salt storage.

Results

Direct reduced-complexity transfer captured only the broad salinity gradient

The direct reduced-complexity transfer retained substantial information about the broad upstream-to-downstream salinity gradient. Across all internal station-day observations, the pooled prediction R^2 was approximately 0.943, and pooled RMSE was approximately 1.33 psu. At first glance, this appears to indicate strong predictive skill. However, the pooled result must be interpreted cautiously because the station network spans a large mean salinity gradient. A model can achieve high pooled R^2 by distinguishing low-salinity upstream stations from high-salinity downstream stations, even if it explains relatively little temporal variation within individual stations.

Station-level diagnostics revealed this distinction (Figure 3.2). The direct transfer performed best in the downstream reach, where same-day boundary salinity and station position already strongly constrain local salinity. Prediction R^2 reached approximately 0.87 to 0.92 at LE1.2 through LE1.4. In contrast, station-level R^2 was much lower in the upper and transition-zone stations: approximately 0.109 at TF1.5, 0.161 at TF1.6, and 0.345 at TF1.7. The largest RMSE occurred at TF1.7, where the direct baseline had RMSE of 2.38 psu. RET1.1 and TF1.6 also retained elevated errors, with RMSE values of 1.64 and 1.25 psu, respectively.

This result shows that the direct reduced-complexity baseline is useful but incomplete. It captures the broad salinity-gradient structure, especially in the lower estuary, but it has limited station-level temporal explanatory power in the upstream and transition-zone reaches. The low R^2 values for the relationship between SE^* and normalized station salinity further indicate that one compact seasonal descriptor cannot fully organize raw station-day variability. The monthly residual structure also suggests that the remaining error is not random scatter alone; it has coherent seasonal and station-specific components, particularly around TF1.7.

These results define the baseline problem for the MLP extension. The goal is not merely to improve the pooled metric, because the pooled metric is already high. The more important question is whether nonlinear modeling improves the stations and physical conditions where direct reduced-complexity transfer leaves structured biases.




Figure 3.1. Observed versus predicted salinity for the direct reduced-complexity raw-observation transfer. Each panel reports station-level salinity prediction R^2 and RMSE.




Figure 3.2. Spatial pattern of direct reduced-complexity salinity-prediction skill. RMSE, MAE, R^2 , and bias are shown by station distance.




Figure 3.3. Median salinity residuals by station and month for the direct reduced-complexity transfer. Structured residuals indicate unresolved seasonal and station-specific behavior.

MLP improved aggregate prediction

The MLP extension improved raw-observation prediction relative to the direct reduced-complexity baseline. On the random-holdout test set, the MLP achieved MAE = 0.807 psu,

RMSE = 1.201 psu, and $R^2 = 0.956$. Compared with the direct baseline pooled RMSE of approximately 1.33 psu and pooled R^2 of approximately 0.943, this represents a moderate but meaningful aggregate improvement.

The improvement is meaningful because the MLP reduces prediction error while retaining the reduced-complexity descriptor SE^* in the feature set. However, the aggregate gain is not the primary scientific result of the chapter. The direct reduced-complexity transfer already performs well in pooled terms because it captures the broad upstream-to-downstream salinity contrast. Therefore, the remaining interpretive value lies in identifying where the MLP adds skill beyond that broad spatial structure.

The predicted-versus-observed scatter from the MLP test set shows that most predictions fall close to the 1:1 line across both low-salinity and high-salinity ranges (Figure 3.4). The model reproduces much of the observed gradient from near-fresh upper-estuary conditions to high-salinity lower-estuary conditions. However, the scatter remains broader in the intermediate salinity range. This dispersion is important because intermediate salinity overlaps with the transition-zone behavior emphasized in Chapter 2. The MLP improves the direct transfer, but it does not remove all structured error. The following analysis therefore focuses on where improvement is spatially concentrated, how it depends on flow state, and how SHAP diagnostics reveal the information used by the trained model.

Station-level comparison showed that MLP improvement was not spatially uniform (Figure 3.5). The largest absolute gain occurred at TF1.7, where RMSE decreased from 2.38 to 1.84 psu and R^2 increased from 0.345 to 0.610. TF1.6 also improved substantially, with RMSE decreasing from 1.25 to 0.92 psu and R^2 increasing from 0.161 to 0.54. TF1.5 showed a smaller absolute RMSE reduction, from 0.52 to 0.40 psu, but its R^2 increased from 0.109 to 0.47. Downstream stations improved only modestly because the direct baseline already performed well there. This pattern can be interpreted in three spatial regimes. In the upstream low-salinity reach, MLP gains were visible mainly in R^2 . Absolute salinity errors are small because the observed salinity range is narrow, but the direct baseline has limited ability to explain within-station variation. The MLP improves this component by allowing seasonal phase, boundary salinity, station position, and SE^* to interact nonlinearly. However, the physical magnitude of error in the uppermost stations remains smaller than in the transition zone.




Figure 3.4. Predicted versus observed salinity for the MLP random-holdout test set.

MLP improvement was spatially uneven and concentrated in the upper-transition reach

TF1.7 is the key station-level case. It is the station where Chapter 2 identified an important transition in mean salinity, seasonal amplitude, anomaly variability, and reduced-complexity residual structure. It is also the station where the MLP produces the largest absolute improvement over the direct baseline. This does not mean that TF1.7 variability is fully resolved. Instead, it indicates that the direct reduced-complexity transfer left substantial structured variability at the transition station and that the nonlinear extension captured part of that structure. In the downstream high-salinity reach, MLP gains were smaller. At LE1.2 through LE1.4, direct baseline R^2 was already approximately 0.87 to 0.92, and MLP R^2 increased only modestly to approximately 0.89 to 0.93. This pattern should not be read as a weakness of the MLP. Rather, it indicates that lower-estuary salinity is already strongly constrained by same-day downstream boundary salinity and station position, leaving less remaining error for nonlinear modeling to remove.

Table 3.3. Station-level diagnostic comparison between the direct reduced-complexity baseline and the MLP extension.

Station	RC RMSE	MLP RMSE	RMSE reduction	RC R ²	MLP R ²	Delta R ²
TF1.5	0.52	0.40	0.12	0.11	0.47	0.36
TF1.6	1.25	0.92	0.33	0.16	0.54	0.38
TF1.7	2.38	1.84	0.54	0.34	0.61	0.27
RET1.1	1.64	1.49	0.15	0.70	0.75	0.05
LE1.1	1.13	1.06	0.07	0.84	0.86	0.02
LE1.2	1.01	0.92	0.09	0.87	0.89	0.02
LE1.3	0.93	0.86	0.08	0.89	0.91	0.02
LE1.4	0.80	0.78	0.02	0.92	0.93	0.00

MLP versus reduced-complexity baseline station metrics




Figure 3.5. Station-level MLP and direct reduced-complexity baseline metrics. The comparison identifies where nonlinear extension most improves the direct baseline.

Flow-stratified improvement showed hydrological-state dependence

The flow-stratified comparison connected model improvement to physical conditions. Using common station-date rows, MLP improvement was strongest under high-flow conditions and weakest under middle-flow conditions. Averaged across stations, RMSE reduction was 0.34 psu

under high flow, 0.29 psu under low flow, and 0.19 psu under middle flow. This pattern indicates that the nonlinear extension added most value when conditions departed from the center of the observed freshwater-flow distribution.

The station-by-flow decomposition (Table 3.4) was more informative than the flow-state averages. The largest gain occurred at TF1.7 under high-flow conditions, where RMSE decreased by 1.34 psu and R^2 increased by 0.869. TF1.7 also showed 0.92 psu and 0.53 psu RMSE reductions under low- and middle-flow conditions, respectively. TF1.6 and RET1.1 showed smaller but consistent gains, whereas downstream stations generally showed small positive reductions. Thus, the MLP improvement was concentrated where station position and hydrological state interact most strongly.

Table 3.4. Largest station-flow RMSE reductions for the MLP relative to the direct reduced-complexity baseline.

Station	Flow state	RC RMSE	MLP RMSE	RMSE gain	Delta R^2
TF1.7	high flow	3.07	1.73	1.34	0.869
TF1.7	low flow	2.77	1.86	0.92	0.501
TF1.6	low flow	1.67	1.12	0.55	0.633
TF1.7	mid flow	2.45	1.93	0.53	0.260
RET1.1	high flow	1.89	1.42	0.48	0.196
TF1.6	mid flow	1.26	0.86	0.39	0.475
TF1.6	high flow	1.12	0.74	0.39	0.653
RET1.1	low flow	1.63	1.38	0.26	0.091




Figure 3.6. Flow-stratified RMSE, MAE, and R^2 for the MLP and direct reduced-complexity baseline.




Figure 3.7. Flow-stratified RMSE reduction from the direct reduced-complexity baseline to the MLP. Positive values indicate lower RMSE for the MLP.

Representative hydrological years showed how flow context shaped the MLP gain

To illustrate model behavior after inspecting the annual forcing and error summaries, five post hoc case years were selected from the 36 years that met minimum coverage requirements of six internal stations, 40 internal observations, and six sampled dates. The typical case, 2009, minimized the robust standardized distance from the eligible-year medians of downstream salinity, river discharge, positive inlet flow, the reduced-complexity descriptor, and station-mean baseline RMSE. The hydrological endpoints were 2003, which had the maximum mean river discharge ($23.30 \text{ m}^3 \text{ s}^{-1}$), and 1986, which had the minimum ($5.07 \text{ m}^3 \text{ s}^{-1}$). The poor-fit case, 2019, had the largest station-mean reduced-complexity baseline RMSE, whereas 2018 had the largest aggregate RC-to-MLP RMSE reduction across TF1.5, TF1.6, TF1.7, and RET1.1. These cases are illustrative diagnostics selected after model evaluation and are not independent validation years.

Across these representative years, the MLP reduced RMSE relative to the direct reduced-complexity model. As shown in table 3.5, the typical year showed a moderate aggregate gain (0.40 psu), while the wet and dry endpoints showed smaller but positive aggregate gains (0.21 and 0.18 psu, respectively). The strongest improvements occurred in the stress-test years. In 2019, aggregate RMSE decreased from 2.42 to 1.52 psu, and upstream-to-transition RMSE decreased by 1.56 psu. In 2018, aggregate RMSE decreased from 2.29 to 1.15 psu, and upstream-to-transition RMSE decreased by 1.90 psu.

The station-level pattern within these years clarifies the physical meaning of the improvement. The largest selected-year station gain occurred at TF1.7 in 2018, where RMSE decreased by 3.44 psu. Large gains also appeared at TF1.7 and RET1.1 in 2019. In contrast, downstream stations showed small gains or slightly negative changes in some representative years. This result

strengthens the interpretation that the MLP acts as a nonlinear correction where fixed transfer relationships are strained by hydrological state, station position, and event-scale departures from seasonal expectation.

Table 3.5. Representative hydrological years selected from the flow background and model-performance diagnostics.

Year	Role	Mean QR	RC RMSE	MLP RMSE	RMSE gain	Upstream gain
2009	typical year	10.85	1.51	1.11	0.40	0.63
2003	wet/high-flow year	23.30	1.11	0.90	0.21	0.14
1986	dry/low-flow year	5.07	1.05	0.87	0.18	0.23
2019	anomalous poor-fit year	16.54	2.42	1.52	0.90	1.56
2018	largest upstream MLP gain year	19.37	2.29	1.15	1.15	1.90




Figure 3.8. Representative years in the annual freshwater-flow background. Point size shows annual RMSE reduction from the direct reduced-complexity baseline to the MLP.

Representative-year predictions at the key transition station TF1.7




Figure 3.9. Observed salinity and two model predictions at TF1.7 across representative hydrological years. Pale bars show Q_R scaled within each year.




Figure 3.10. Station-specific MLP gains in representative hydrological years. Positive values indicate lower RMSE for the MLP than for the direct reduced-complexity baseline.

SHAP diagnostics explained the fitted MLP as a spatial-boundary-seasonal model

SHAP diagnostics changed the interpretation of MLP performance from a simple metrics comparison to a model-behavior explanation. Global mean absolute SHAP values showed that spatial position was the dominant contributor to the trained MLP prediction function, followed by same-day downstream boundary salinity. Seasonal phase terms had smaller but non-negligible contributions, while the reduced-complexity descriptor S_E^* acted as a compact supporting descriptor once distance, boundary salinity, and seasonal phase were already included. This ranking does not invalidate the reduced-complexity framework. Instead, it refines its role. The direct baseline defines the physically interpretable structure that motivates the chapter, while the MLP learns how that structure interacts with station position, downstream boundary condition, and seasonal phase. The SHAP results therefore place a boundary on interpretation: high MLP skill should not be read as proof that the model fully learned freshwater-driven estuarine dynamics. Much of the prediction skill reflects the strong salinity-gradient structure itself.

Local SHAP cases reinforced the same conclusion. Representative TF1.7 high-flow and low-flow cases showed that transition-zone errors can occur in both directions: the model can underpredict or overpredict when the local salinity gradient shifts more strongly than the available features diagnose. A downstream LE1.4 case, by contrast, was predicted with near-zero error because station distance and boundary salinity placed the observation in a stable high-salinity part of the learned mapping.




Figure 3.11. Global mean absolute SHAP values for grouped MLP feature families.




Figure 3.12. SHAP summary plot showing feature value, direction, and contribution to MLP salinity prediction.




Figure 3.13. Local SHAP decomposition for a representative TF1.7 high-flow case.

Discussion

What the MLP added to the reduced-complexity baseline

The main contribution of the MLP extension is not simply that a neural network can fit salinity observations. Previous estuarine and coastal studies have already shown that machine-learning models can predict salinity or emulate more complex salinity dynamics when supplied with hydrographic, boundary, and temporal information (Hutton et al. 2016; Qiu and Wan 2013; Etemad-Shahidi et al. 2008; Jayasundara et al. 2020; Qi et al. 2021; He et al. 2020; Qi, He, Bai, Ding, Sandhu, Zhou, et al. 2022; Qi, He, Bai, Ding, Sandhu, Chung, et al. 2022; Saccotelli et al. 2024). The contribution here is narrower and more diagnostic: the MLP tests how much predictive value is added when the Chapter 2 reduced-complexity salinity-gradient descriptor is placed inside a nonlinear feature space.

This distinction keeps Chapter 3 tied to the thesis argument. The direct baseline shows what can be explained by the reduced-complexity descriptor, same-day boundary salinity normalization, and station-specific transfer alone. The MLP retains this reduced-complexity information but allows it to interact with boundary salinity, station distance, and seasonal phase. It is therefore best understood as a nonlinear extension of the physical framework, not as a replacement for it (Read et al. 2019; Willard et al. 2022).

The aggregate improvement was moderate because the direct baseline already captured much of the broad spatial salinity contrast. The stronger evidence for added value came from decomposed diagnostics. The MLP improved most where the direct transfer was least complete: the upper-transition reach, high- and low-flow states, and representative years in which the direct baseline had elevated error.

Pooled metrics must be interpreted through the estuarine gradient

The results show why pooled error metrics are insufficient for this chapter. A high pooled R^2 can reflect the ability to separate upstream low-salinity stations from downstream high-salinity stations without fully explaining temporal variability within each station. This is not only a statistical issue. It follows directly from the physical structure of the Patuxent estuary, where station position and boundary salinity impose a strong salinity gradient (Roberts et al. 2017). For this reason, station-level, flow-state, and representative-year diagnostics are more scientifically informative than aggregate metrics alone. These decompositions showed that nonlinear modeling did not improve all parts of the estuary equally. It added the most value in the locations and flow contexts where the direct reduced-complexity transfer left structured residuals, especially TF1.7 and the broader upstream-to-transition reach.

Representative years link model improvement to physical conditions

The representative-year analysis provides the clearest bridge between model performance and physical interpretation. In wet/high-flow conditions, freshwater forcing compressed the salinity gradient and increased the importance of capturing flow-driven displacement of the salt field. In dry/low-flow conditions, weaker freshwater flushing allowed marine influence to extend farther upstream, increasing the sensitivity of local station salinity to boundary conditions and inlet-flow forcing. The typical year showed that the MLP advantage was not confined to extremes, while the anomalous poor-fit and high-gain years showed that the largest gains appeared when fixed transfer relationships were strained. (Gibson and Najjar 2000; Monismith et al. 2002; MacCready 2007)

This analysis should still be interpreted carefully. The representative-year and flow-stratified results show association with freshwater-flow state, not a causal decomposition of discharge effects. The current MLP does not include explicit lagged discharge, antecedent flow duration, wind, tide, stratification, or salt-storage terms. Its improvement therefore indicates that the existing features contain learnable contextual information, but the remaining transition-zone errors point to physical dynamics that are not yet explicitly represented (MacCready 2007; Ralston et al. 2008).

SHAP clarifies what the fitted MLP learned and where it stops

SHAP diagnostics showed that the fitted MLP relied primarily on station distance and same-day downstream boundary salinity. This is physically sensible because salinity in the Patuxent estuary is strongly organized by along-estuary position and by the salinity state of the downstream boundary. However, it also places a clear bound on interpretation. High model skill mainly reflects a learned nonlinear spatial-boundary-seasonal mapping, not necessarily a complete representation of freshwater-driven dynamics (Geyer and MacCready 2014; Gibson and Najjar 2000).

The smaller marginal contribution of S_E^* should not be read as a failure of the reduced-complexity framework. Instead, it indicates that within the MLP feature space, much of the information needed for prediction is already carried by distance, CB5.1W salinity, and seasonal phase. S_E^* remains valuable because it defines the direct physical baseline and maintains the conceptual link between Chapters 2 and 3. SHAP therefore helps separate prediction from explanation: it explains how the trained model used the available inputs, but it does not prove that those inputs are independent physical causes (Lundberg and Lee 2017; Lundberg et al. 2020; Aas et al. 2021).

The transition zone remains the key diagnostic region

TF1.7 is the central station linking Chapters 2 and 3. In Chapter 2, the TF1.7-to-RET1.1 marked a transition in salinity-gradient structure, anomaly variability, reduced-complexity residuals, and geometry-linked attenuation of freshwater influence. In Chapter 3, TF1.7 was also the station where the MLP produced the largest gains over the direct reduced-complexity baseline, especially under high- and low-flow conditions and in representative stress-test years. This agreement suggests that the same part of the estuary is both physically important and difficult for a compact transfer model. The transition zone is not simply noisy. It is where relatively small shifts in the salinity gradient can produce large local changes, and where missing dynamic information such as antecedent flow, lagged boundary adjustment or wind mixing is most likely to matter (MacCready 2007; Ralston et al. 2008).

Implications for future model development

The next model-development step should be targeted rather than generic. The results do not suggest that the main limitation is insufficient neural-network size. Instead, they indicate that the current feature set lacks physically interpretable dynamic information in the stations, flow states, and years where direct reduced-complexity transfer remains incomplete (Read et al. 2019; Willard et al. 2022).

Several extensions follow directly from the diagnostics. Lagged freshwater discharge should be added to represent hydrological memory, using summaries such as 7-day, 14-day, and 30-day antecedent flow means, high-flow duration, or cumulative flow anomalies. Lagged CB5.1W salinity could help represent delayed boundary adjustment. Tidal water-level derivatives, wind forcing, or flow-regime indicators could help distinguish event-scale exchange conditions that are only indirectly represented in the current model (MacCready 2007; Ralston et al. 2008).

Validation design is equally important. The random holdout test used here measures interpolation skill within the sampled observation set. It does not test whether the model transfers to unseen years, unusual hydrological regimes, or future conditions. Chronological splits, leave-one-year-out validation, and flow-regime-based validation would provide more demanding tests (Roberts et al. 2017). A transition-zone-focused validation design may also be useful because

pooled metrics can be dominated by the broader spatial gradient and may understate the importance of TF1.7 and RET1.1 errors.

Chapter Conclusion

This chapter tested whether a machine-learning extension could improve the direct transfer of reduced-complexity salinity-gradient framework to reconstruct the station-day observations. The direct reduced-complexity model captured the broad upstream-to-downstream salinity structure, with pooled RMSE of approximately 1.33 psu and pooled R^2 of approximately 0.943, but it left structured station-level residuals in the upstream and transition-zone reaches.

Embedding the reduced-complexity descriptor within a nonlinear MLP feature space improved raw-observation prediction. The MLP achieved MAE = 0.8072 psu, RMSE = 1.2012 psu, and R^2 = 0.9556 on the random-holdout test set. More importantly, the improvement was spatially and hydrologically structured. Gains were strongest around TF1.7, under high- and low-flow states, and in representative years where the direct baseline struggled.

The representative-year and SHAP analyses clarify the chapter's interpretation. The MLP should not be viewed as replacing the reduced-complexity physical framework. Rather, it acts as a nonlinear correction that is most valuable when fixed transfer relationships are strained by hydrological state, station position, and event-scale departures from the long-term averaged salinity gradient. SHAP further shows that the fitted model relies mainly on spatial position, downstream boundary salinity, and seasonal phase, which explains its predictive success while also limiting causal interpretation.

Overall, Chapter 3 shows that the reduced-complexity salinity-gradient representation can be strengthened for raw-observation prediction by embedding it within a nonlinear machine-learning framework. The most useful next steps are to add hydrological-memory and boundary-lag information, evaluate chronological and flow-regime transferability, and focus future model development on the transition zone where the physical and predictive diagnostics agree.

Chapter 4: General Discussion and Conclusions

Synthesis of the Thesis

This thesis examined how surface salinity reveals the spatial and temporal structure of the Patuxent River estuary and how that structure is organized by freshwater discharge, downstream boundary conditions, and other physical context. Salinity was treated as an integrative indicator rather than as an isolated water-quality variable. Each observation reflects the combined expression of freshwater input, bay influence, transport, mixing, geometry, and station position. The purpose of the thesis is therefore not to assign all salinity variability to a single forcing mechanism, but to establish a defensible framework for separating persistent spatial structure, recurring seasonal behavior, interannual departures, and model-representable variability.

This thesis makes three linked contributions to understanding Patuxent estuarine salinity. First, it establishes station-specific day-of-year climatology as the analytical reference needed to separate the persistent along-estuary gradient from seasonal expectation and interannual departure. This reference shows why neither a single estuary-wide time series nor a uniform discharge-salinity relation can preserve the local meaning of an observation.

Second, the analyses separate the spatial expression of freshwater and downstream-boundary sensitivity. Freshwater-related losses and residual structure are concentrated in the upper and upper-transition reaches, whereas downstream boundary salinity increasingly constrains the lower estuary. The reduced-complexity and counterfactual results therefore support a spatially filtered forcing framework: the same external forcing is expressed differently according to station position, background salinity, and geometry. These sensitivities are physically interpretable diagnostics, not an additive causal partition of salt transport.

Third, the TF1.7-to-RET1.1 transition emerges as the common physical and predictive diagnostic region of the thesis. It is where movement of the salinity gradient produces large local variability, where a fixed station-transfer relation leaves its strongest structured errors, and where the nonlinear model provides its clearest additional skill. The convergence of observational climatology, geometry-aware scaling, reduced-complexity residuals, and machine-learning gains identifies the transition zone as the priority region for future high-frequency, vertically resolved, and velocity-based observations. Because the MLP relies strongly on station position, observed downstream salinity, and seasonal phase, it should be interpreted as a diagnostic reconstruction of this structure rather than as an operational forecasting system.

Answers to the Research Questions

Persistent spatial salinity structure

The first research question asked what persistent spatial salinity structure exists along the Patuxent River estuary. The observations demonstrate a strong and repeatable upstream-to-downstream gradient, but the gradient is not spatially uniform. TF1.5 and TF1.6 define a low-salinity upper reach. TF1.7 marks a substantial increase in both climatological mean salinity and seasonal amplitude, while RET1.1 occupies the downstream side of a rapid transition into higher-salinity conditions. From the LE stations toward CB5.1W, salinity remains substantially higher and becomes more closely linked to the downstream Chesapeake Bay boundary. The stations can therefore be interpreted as occupying three broad spatial regimes: an upper low-salinity reach, an upper-to-middle transition reach, and a lower boundary-influenced reach. These regimes are not fixed water-mass boundaries, and their precise locations can shift with hydrological and marine conditions. They nevertheless provide a useful frame for interpreting how salinity variability changes along the estuary.

Spatial variation in the seasonal salinity cycle

The second question asked how the salinity seasonal cycle varies from upstream to downstream. The analysis showed that each station has its own seasonal baseline. Mean state, seasonal amplitude and phase differ substantially along the estuarine gradient. A single estuary-wide seasonal curve would consequently obscure the local manifestation of salinity observations. Station-specific climatology was essential because the same absolute salinity value can represent very different conditions at different locations. A moderate salinity observation may be anomalously high at an upstream station but anomalously low in the lower estuary. Expressing observations relative to local day-of-year expectations preserved the spatial gradient while allowing seasonal and interannual departures to be compared across stations. The harmonic analysis provided a compact summary of these seasonal cycles, but it did not replace the empirical climatology. The first-order harmonic represented broad annual mean state and amplitude effectively, whereas timing estimates were less stable when seasonal amplitude was weak or within-year coverage was incomplete. The principal seasonal result is therefore not a single estuary-wide timing pattern, but the systematic spatial differentiation of seasonal mean state and amplitude.

Interannual salinity departures

The third question asked how annual salinity anomalies reflect interannual variability. Annual departures were most meaningfully expressed relative to station-specific seasonal expectation. This approach removed much of the persistent spatial and seasonal structure before comparing individual years.

The resulting anomalies showed that the monitoring record contains genuine interannual variation rather than a fixed climatological gradient. Some station-years were consistently saltier than expected for their sampled dates, while others were fresher. The magnitude of these departures was relatively small at the uppermost stations, partly because salinity remained close to its freshwater lower bound. Larger anomaly variability occurred from TF1.7 downstream, where shifts in the salinity gradient could generate greater local changes.

Wet years tended to be fresher than regular years, and event-anomaly years also showed negative anomalies in the unadjusted summaries. These findings are consistent with freshwater forcing shifting the salinity field downstream. They do not, however, support a simple symmetric relationship in which all dry years are anomalously saline, and all wet years are uniformly fresh. Annual discharge classifications summarize hydrological background, whereas observed salinity also reflects the seasonal timing and duration of flow, downstream boundary conditions, mixing, wind, salt storage, and station position. The interannual result is therefore that freshwater context organizes part of the anomaly structure, but its consequence depends on where and when salinity is observed.

Spatial organization of physical forcing

The fourth question asked how freshwater discharge, downstream boundary salinity, and local geometry organize salinity response along the gradient. The reduced-complexity and counterfactual analyses showed that these controls are expressed differently in different parts of the estuary.

Seasonal variation in downstream boundary salinity was strongly associated with the salinity range of the lower-estuary stations within the fitted framework. Removing seasonal variation from the boundary substantially reduced the predicted range and degraded climatology-scale reconstruction, especially near the downstream end of the station sequence. In contrast, holding freshwater discharge fixed produced its largest losses around TF1.6, TF1.7, and RET1.1. This pattern is consistent with freshwater forcing having its clearest local expression in the upper and transitional portions of the salinity gradient, while lower-estuary salinity remains more closely connected to the river boundary.

The TF1.7-to-RET1.1 reach provides the clearest hydrogeomorphic example of spatial filtering. The estimated cross-sectional area expands sharply across this interval, reducing the first-order advective tendency associated with a given upstream discharge. This scaling does not establish a complete salt budget, because it does not consider vertical salt exchange. It nevertheless demonstrates that the same discharge cannot be expected to produce the same local salinity response throughout an estuary with strongly varying geometry.

Freshwater discharge should therefore be interpreted as an estuary-scale forcing whose observed local effect is mediated by the physical structure through which it acts. The thesis does not identify one universally dominant driver. Instead, it shows a spatial separation of sensitivities: stronger discharge-related sensitivity in the upper-transition reach, stronger boundary-related

sensitivity in the lower estuary, and an intervening region where geometry and exchange alter the form of the freshwater signal.

Machine-learning extension of the physical baseline

The fifth question asked whether a compact machine-learning extension could improve the direct reduced-complexity baseline while preserving the observed spatial structure. The results support a qualified affirmative answer.

The MLP reduced prediction error and added the most value where the direct reduced-complexity transfer was incomplete. Improvement was concentrated in the upper-transition reach rather than distributed uniformly across all stations. The downstream stations already had high direct-baseline skills because contemporaneous boundary salinity and station position strongly constrained their salinity. The upper and transition stations contained more unresolved within-station variability and therefore provided a more demanding test of the nonlinear extension.

The MLP did not eliminate all transition-zone error. Predicted variability remained biased at some stations, particularly TF1.7, indicating that the available inputs did not fully represent the dynamic displacement and adjustment of the salinity field. Moreover, because the model used same-day CB5.1W salinity and was evaluated through a random holdout, its performance represents interpolation or conditional reconstruction within the sampled observational domain. It does not demonstrate forecasting skill for unseen years, novel flow regimes, or future environmental conditions.

The machine-learning result is therefore not that a neural network replaced physical interpretation. Rather, nonlinear modeling identified where additional predictive information could be extracted from interactions among spatial position, boundary salinity, seasonal phase, and the reduced-complexity descriptor. The locations where the MLP improved most—and where it continued to struggle—reinforced the physical interpretation established independently in Chapter 2.

Scientific Contributions

The first contribution of this thesis is the use of station-specific seasonal climatology as a reference frame for interpreting a long but irregular estuarine monitoring record. The previously documented upstream-to-downstream salinity gradient was converted from a descriptive pattern into an analytical baseline. This allowed observations to be evaluated relative to the expected condition for their own station and season, separating persistent spatial structure from temporal departure.

The second contribution is the identification of the upper-transition reach as a common diagnostic region across multiple forms of evidence. The change into TF1.7, the TF1.7-to-RET1.1 interval, the larger anomaly variability, the concentration of reduced-complexity residuals, the geometry-linked attenuation calculation, and the spatial pattern of MLP

improvement all point to the same portion of the estuary. Their convergence supports treating this reach as a process transition rather than as an isolated station anomaly.

The third contribution is an explicitly spatial model-evaluation framework. The thesis shows why pooled metrics are insufficient in an estuary with a large mean salinity gradient. A model can achieve a high pooled R^2 by separating upstream and downstream stations while explaining relatively little temporal variation within individual stations. Station-level, flow-stratified, seasonal, and representative-year diagnostics are therefore necessary to determine whether a model reproduces the physically informative parts of the estuarine gradient.

The fourth contribution is the integration of reduced-complexity physical reasoning with machine-learning analysis without treating either approach as sufficient by itself. The physical framework supplied an interpretable baseline and identified where simplified forcing relationships failed. The MLP tested whether nonlinear feature interactions could reduce those residuals. SHAP analysis then clarified which inputs the fitted model used, separating predictive dependence from physical causality. This combined approach provides comprehensive and insightful understanding for the factors driving spatiotemporal salinity variations.

Limitations and Boundaries of Inference

Several limitations define the scope of the conclusions. First, the salinity observations form a long but temporally sparse monitoring record. The data support station-specific climatology, broad seasonal summaries, annual anomalies, and forcing-context comparisons. They do not provide continuous event-scale records at every station. The thesis therefore cannot determine the precise timing of local salinity adjustment following individual discharge pulses.

Second, the analysis is based on surface salinity. The results describe the surface expression of the Patuxent salinity gradient and do not resolve vertical stratification, bottom-water intrusion, or two-layer exchange. These processes may be particularly important in the transition and lower-estuary reaches and may explain part of the residual variability that remains unresolved by the reduced-complexity and machine-learning models.

Third, the reduced-complexity inlet term is a water-level-derived exchange proxy rather than a direct estimate of total estuarine exchange flow. Its numerical scale depends on the assumed effective estuary area, and the fitted transfer coefficients should not be interpreted as calibrated physical parameters. The counterfactuals similarly measure sensitivity within the fitted model structure. Because freshwater discharge, boundary salinity, tidal exchange, and salt storage are coupled in the real estuary, replacing one input with its mean does not provide an additive partition of physical causality.

Fourth, the cross-sectional analysis is a first-order scaling calculation. It demonstrates how widening geometry can reduce the local advective expression of a given discharge, but it does not independently resolve the full mechanism of the TF1.7-to-RET1.1 transition. Confirming the relative roles of geometry, exchange, storage, mixing, and stratification would require velocity

measurements, vertical salinity observations, salt-budget analysis, or output from a spatially resolved hydrodynamic model.

Finally, the MLP validation design measures interpolation skill within the sampled observational domain. Randomly held-out station-day records may share dates and environmental context with training observations. In addition, the use of contemporaneous downstream boundary salinity means that the model reconstructs internal salinity conditional on an observed boundary state. The current model should therefore not be interpreted as demonstrating independent forecasting or extrapolation skill.

Future Research

Future work should focus on resolving the dynamic processes that are most important in the transition zone rather than simply increasing model complexity. Higher-frequency salinity observations at TF1.7 and RET1.1 would help determine how quickly the local gradient responds to freshwater pulses and changing downstream conditions. Paired surface and bottom observations would clarify whether unresolved residuals are associated with stratification, two-layer exchange, or vertical adjustment of the salt field. Additional cross-sectional and velocity information would also allow the geometry-based scaling argument to be evaluated within a more complete salt-transport framework.

Model development should incorporate hydrological memory explicitly. Antecedent discharge summaries, cumulative flow anomalies, high-flow duration, and lagged boundary salinity could represent delayed freshwater and marine adjustment more directly than same-day or climatological forcing alone.

More demanding validation designs are also necessary. Grouped-date splits would prevent same-date observations from appearing in both training and testing sets. Chronological validation and leave-one-year-out tests would evaluate transfer to unseen hydrological years, while flow-regime validation would test whether relationships learned under typical conditions remain valid during unusually wet or dry periods. Comparisons with regularized linear models, generalized additive models, and feature-ablation experiments would clarify how much predictive value is contributed specifically by the reduced-complexity descriptor, the downstream boundary observation, seasonal phase, and nonlinear interactions.

The transition reach should remain the primary target of these extensions. Downstream stations are already strongly constrained by boundary salinity, while the uppermost stations often have a narrow absolute salinity range. TF1.7 and RET1.1 provide the most informative setting for evaluating whether new physical inputs improve the representation of gradient displacement, hydrological memory, and exchange-mediated variability.

References

- Aas, Kjersti, Martin Jullum, and Anders Loland. 2021. “Explaining Individual Predictions When Features Are Dependent: More Accurate Approximations to Shapley Values.” *Artificial Intelligence* 298: 103502. <https://doi.org/10.1016/j.artint.2021.103502>.
- Agonsi, M. C., R. M. Castelao, J. E. Sheldon, and D. Di Iorio. 2026. “Physical Controls of Salinity and Residence Time Variability in the Altamaha River Estuary.” *Estuarine, Coastal and Shelf Science* 338: 109910. <https://doi.org/10.1016/j.ecss.2026.109910>.
- Bingham, F. M., S. Brodnitz, and L. Yu. 2021. “Sea Surface Salinity Seasonal Variability in the Tropics from Satellites, Gridded In Situ Products and Mooring Observations.” *Remote Sensing* 13 (1): 110. <https://doi.org/10.3390/rs13010110>.
- Brahmey, E. I., K. J. McGlathery, and S. C. Doney. 2024. “Quantifying Seasonal to Multi-Decadal Signals in Coastal Water Quality Using High- and Low-Frequency Time Series Data.” *Cambridge Prisms: Coastal Futures 2*: e8. <https://doi.org/10.1017/cft.2024.6>.
- Chesapeake Bay Program. “CBP Water Quality Database (1984-Present).” In *Chesapeake Bay Water Quality Database*. Chesapeake Bay Program. <https://www.chesapeakebay.net/what/downloads/cbp-water-quality-database-1984-present>.
- Dreiss, A., A. R. Azarnivand, A. Hildebrand, et al. 2024. “Controls on Oxygen Variability and Depletion in the Patuxent River Estuary.” *Estuaries and Coasts* 47: 2306–23. <https://doi.org/10.1007/s12237-024-01390-3>.
- Dyer, K. R. 1997. *Estuaries: A Physical Introduction*. 2nd ed. Wiley.
- Etemad-Shahidi, A., A. Dorostkar, and W. C. Liu. 2008. “Prediction of Salinity Intrusion in Danshuei Estuarine System.” *Hydrology Research* 39 (5–6): 497–505. <https://doi.org/10.2166/nh.2008.107>.
- Geyer, W. R. 2010. “Estuarine Salinity Structure and Circulation.” In *In A. Valle-Levinson (Ed.), Contemporary Issues in Estuarine Physics (Pp. 12-26)*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511676567.003>.
- Geyer, W. R., and P. MacCready. 2014. “The Estuarine Circulation.” *Annual Review of Fluid Mechanics* 46: 175–97. <https://doi.org/10.1146/annurev-fluid-010313-141302>.
- Gibson, J. R., and R. G. Najjar. 2000. “The Response of Chesapeake Bay Salinity to Climate-Induced Changes in Streamflow.” *Limnology and Oceanography* 45 (8): 1764–72. <https://doi.org/10.4319/lo.2000.45.8.1764>.
- Gorski, G., S. E. Cook, A. M. Snyder, et al. 2024. “Deep Learning of Estuary Salinity Dynamics Is Physically Accurate at a Fraction of Hydrodynamic Model Computational Cost.” *Limnology and Oceanography* 69 (5): 1070–85. <https://doi.org/10.1002/lno.12549>.
- Hansen, Donald V., and Maurice Rattray Jr. 1966. “New Dimensions in Estuary Classification 1.” *Limnology and Oceanography* 11 (3): 319–26.
- Hartzell, J. L., T. E. Jordan, and J. C. Cornwell. 2010. “Phosphorus Burial in Sediments along the Salinity Gradient of the Patuxent River, a Subestuary of the Chesapeake Bay (USA).” *Estuaries and Coasts* 33: 92–106. <https://doi.org/10.1007/s12237-009-9232-2>.
- He, M., L. Zhong, P. Sandhu, and Y. Zhou. 2020. “Emulation of a Process-Based Salinity Generator for the Sacramento-San Joaquin Delta of California via Deep Learning.” *Water* 12 (8): 2088. <https://doi.org/10.3390/w12082088>.

Hilton, T. W., R. G. Najjar, L. Zhong, and M. Li. 2008. "Is There a Signal of Sea-Level Rise in Chesapeake Bay Salinity?." *Journal of Geophysical Research: Oceans* 113: C09002. <https://doi.org/10.1029/2007JC004247>.

Hutton, P. H., J. S. Rath, L. Chen, M. J. Unga, and S. B. Roy. 2016. "Nine Decades of Salinity Observations in the San Francisco Bay and Delta: Modeling and Trend Evaluations." *Journal of Water Resources Planning and Management* 142 (3): 04015069. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000617](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000617).

Jayasundara, N., S. Seneviratne, E. Reyes, and F. Chung. 2020. "Artificial Neural Network for Sacramento-San Joaquin Delta Flow-Salinity Relationship for CalSim 3.0." *Journal of Water Resources Planning and Management* 146 (4): 04020013. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001192](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001192).

Kemp, W. M., R. Batiuk, R. Bartleson, et al. 2004. "Habitat Requirements for Submerged Aquatic Vegetation in Chesapeake Bay: Water Quality, Light Regime, and Physical-Chemical Factors." *Estuaries* 27: 363–77. <https://doi.org/10.1007/BF02803529>.

Lundberg, S. M., G. Erion, H. Chen, et al. 2020. "From Local Explanations to Global Understanding with Explainable AI for Trees." *Nature Machine Intelligence* 2: 56–67. <https://doi.org/10.1038/s42256-019-0138-9>.

Lundberg, S. M., and S. I. Lee. 2017. "A Unified Approach to Interpreting Model Predictions." *Advances in Neural Information Processing Systems*, 30.

MacCready, P. 1999. "Estuarine Adjustment to Changes in River Flow and Tidal Mixing." *Journal of Physical Oceanography* 29: 708–26. [https://doi.org/10.1175/1520-0485\(1999\)029%253C0708:EATCIR%253E2.0.CO;2](https://doi.org/10.1175/1520-0485(1999)029%253C0708:EATCIR%253E2.0.CO;2).

MacCready, P. 2004. "Toward a Unified Theory of Tidally-Averaged Estuarine Salinity Structure." *Estuaries* 27 (4): 561–70. <https://doi.org/10.1007/BF02907644>.

MacCready, P. 2007. "Estuarine Adjustment." *Journal of Physical Oceanography* 37 (8): 2133–45. <https://doi.org/10.1175/JPO3082.1>.

MacCready, P., and W. R. Geyer. 2010. "Advances in Estuarine Physics." *Annual Review of Marine Science* 2: 35–58. <https://doi.org/10.1146/annurev-marine-120308-081015>.

MacCready, P., W. R. Geyer, and H. Burchard. 2018. "Estuarine Exchange Flow Is Related to Mixing through the Salinity Variance Budget." *Journal of Physical Oceanography* 48 (6): 1375–84. <https://doi.org/10.1175/JPO-D-17-0266.1>.

Monismith, Stephen G., Wim Kimmerer, Jon R. Burau, and Mark T. Stacey. 2002. "Structure and Flow-Induced Variability of the Subtidal Salinity Field in Northern San Francisco Bay." *Journal of Physical Oceanography* 32 (11): 3003–19.

NOAA Center for Operational Oceanographic Products and Services. "Water Level Observations for Station 8577330, Solomons Island, MD." In *NOAA Tides and Currents / CO-OPS Data API*. National Oceanic and Atmospheric Administration. <https://tidesandcurrents.noaa.gov/stationhome.html?id=8577330>.

Pritchard, D. W. 1952. "Salinity Distribution and Circulation in the Chesapeake Bay Estuarine System." *Journal of Marine Research*, 11(2). https://elischolar.library.yale.edu/journal_of_marine_research/763/.

Qi, S., Z. Bai, Z. Ding, et al. 2021. "Enhanced Artificial Neural Networks for Salinity Estimation and Forecasting in the Sacramento-San Joaquin Delta of California." *Journal of Water Resources Planning and Management* 147 (10): 04021039. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001435](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001435).

- Qi, S., M. He, Z. Bai, Z. Ding, P. Sandhu, Y. Zhou, et al. 2022. “Multi-Location Emulation of a Process-Based Salinity Model Using Machine Learning.” *Water* 14 (13): 2030. <https://doi.org/10.3390/w14132030>.
- Qi, S., M. He, Z. Bai, Z. Ding, P. Sandhu, F. Chung, et al. 2022. “Novel Salinity Modeling Using Deep Learning for the Sacramento-San Joaquin Delta of California.” *Water* 14 (22): 3628. <https://doi.org/10.3390/w14223628>.
- Qiu, C., and Y. Wan. 2013. “Time Series Modeling and Prediction of Salinity in the Caloosahatchee River Estuary.” *Water Resources Research* 49 (9): 5804–16. <https://doi.org/10.1002/wrcr.20415>.
- Ralston, David K., W. Rockwell Geyer, and James A. Lerczak. 2008. “Subtidal Salinity and Velocity in the Hudson River Estuary: Observations and Modeling.” *Journal of Physical Oceanography* 38 (4): 753–70. <https://doi.org/10.1175/2007JPO3808.1>.
- Read, Jordan S., Xiaowei Jia, Jared Willard, et al. 2019. “Process-Guided Deep Learning Predictions of Lake Water Temperature.” *Water Resources Research* 55 (11): 9173–90. <https://doi.org/10.1029/2019WR024922>.
- Roberts, D. R., V. Bahn, S. Ciuti, et al. 2017. “Cross-Validation Strategies for Data with Temporal, Spatial, Hierarchical, or Phylogenetic Structure.” *Ecography* 40: 913–29. <https://doi.org/10.1111/ecog.02881>.
- Roh, Dong Min, Minxue He, Zhaojun Bai, et al. 2023. “Physics-Informed Neural Networks-Based Salinity Modeling in the Sacramento–San Joaquin Delta of California.” *Water* 15 (13): 2320. <https://doi.org/10.3390/w15132320>.
- Saccotelli, L., G. Verri, A. De Lorenzis, et al. 2024. “Enhancing Estuary Salinity Prediction: A Machine Learning and Deep Learning Based Approach.” *Applied Computing and Geosciences* 23: 100173. <https://doi.org/10.1016/j.acags.2024.100173>.
- Strong, Aaron, John Cullen, and Sallie Chisholm. 2009. “Ocean Fertilization: Science, Policy, and Commerce.” *Oceanography* 22 (3): 236–61. <https://doi.org/10.5670/oceanog.2009.83>.
- Testa, J. M., W. M. Kemp, W. R. Boynton, and J. D. III. Hagy. 2008. “Long-Term Changes in Water Quality and Productivity in the Patuxent River Estuary: 1985 to 2003.” *Estuaries and Coasts* 31: 1021–37. <https://doi.org/10.1007/s12237-008-9095-y>.
- U.S. Geological Survey. “USGS Water Data for Monitoring Location USGS-01594440, Patuxent River near Bowie, MD: Discharge Parameter 00060.” In *National Water Information System*. U.S. Geological Survey. <https://waterdata.usgs.gov/monitoring-location/USGS-01594440/>.
- Willard, Jared, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. 2022. “Integrating Scientific Knowledge with Machine Learning for Engineering and Environmental Systems.” *ACM Computing Surveys* 55 (4): 1–37. <https://doi.org/10.1145/3514228>.
- Yavary Nia, M., A. Canestrelli, H. Klammler, M. D. Annable, D. Pinton, and J. W. Jawitz. 2025. “A Reduced-Complexity Model to Predict Seasonal Variation in Estuarine Salinity.” *Geophysical Research Letters* 52: e2025GL116332. <https://doi.org/10.1029/2025GL116332>.
- Yu, L. 2021. “Revisiting the Global Patterns of Seasonal Cycle in Sea Surface Salinity.” *Journal of Geophysical Research: Oceans* 126: e2020JC016789. <https://doi.org/10.1029/2020JC016789>.