

Water Resources Research®

RESEARCH ARTICLE

10.1029/2023WR035707

Key Points:

- Technical anomalies in water quality data from in-situ sensors are common, posing challenges in their detection that often require substantial time and effort
- We present a flexible anomaly detection framework used to automatically detect multiple anomaly types using empirical model residuals
- To facilitate real-time monitoring, we have developed a Bayesian recursive stream-network model that effectively captures spatio-temporal dependencies within the data
- The timely identification of anomalies plays a crucial role in reducing uncertainties and biases, thereby enabling a more accurate understanding and management of water resources

Correspondence to:

E. Santos-Fernandez,
santosfe@qut.edu.au;
edgar.santosfdez@gmail.com

Citation:

Santos-Fernandez, E., Ver Hoef, J. M., Peterson, E. E., McGree, J., Villa, C. A., Leigh, C., et al. (2024). Unsupervised anomaly detection in spatio-temporal stream network sensor data. *Water Resources Research*, 60, e2023WR035707. <https://doi.org/10.1029/2023WR035707>

Received 24 JUL 2023

Accepted 9 SEP 2024

Author Contributions:

Conceptualization: Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Cesar A. Villa, Catherine Leigh, Ryan Turner, Cameron Roberts, Kerrie Mengersen
Data curation: Edgar Santos-Fernandez, Cesar A. Villa, Ryan Turner, Cameron Roberts

Formal analysis: Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Cesar

Unsupervised Anomaly Detection in Spatio-Temporal Stream Network Sensor Data



Edgar Santos-Fernandez¹ , Jay M. Ver Hoef² , Erin E. Peterson¹ , James McGree¹ , Cesar A. Villa³ , Catherine Leigh⁴ , Ryan Turner^{5,6} , Cameron Roberts⁵ , and Kerrie Mengersen¹

¹School of Mathematical Sciences, Queensland University of Technology, Brisbane, QLD, Australia, ²Marine Mammal Laboratory, NOAA-NMFS Alaska Fisheries Science Center, Juneau, AK, USA, ³Science Information Services, Queensland Department of Environment and Science, Brisbane, QLD, Australia, ⁴Biosciences and Food Technology Discipline, School of Science. Royal Melbourne Institute of Technology (RMIT), Melbourne, VIC, Australia, ⁵Water Quality & Investigations, Queensland Department of Environment and Science, Brisbane, QLD, Australia, ⁶School of Earth and Environmental Sciences, University of Queensland, Brisbane, QLD, Australia

Abstract The use of in-situ digital sensors for water quality monitoring is becoming increasingly common worldwide. While these sensors provide near real-time data for science, the data are prone to technical anomalies that can undermine the trustworthiness of the data and the accuracy of statistical inferences, particularly in spatial and temporal analyses. Here we propose a framework for detecting anomalies in sensor data recorded in stream networks, which takes advantage of spatial and temporal autocorrelation to improve detection rates. The proposed framework involves the implementation of effective data imputation to handle missing data, alignment of time-series to address temporal disparities, and the identification of water quality events. We explore the effectiveness of a suite of state-of-the-art statistical methods including posterior predictive distributions, finite mixtures, and Hidden Markov Models (HMM). We showcase the practical implementation of automated anomaly detection in near-real time by employing a Bayesian recursive approach. This demonstration is conducted through a comprehensive simulation study and a practical application to a substantive case study situated in the Herbert River, located in Queensland, Australia, which flows into the Great Barrier Reef. We found that methods such as posterior predictive distributions and HMM produce the best performance in detecting multiple types of anomalies. Utilizing data from multiple sensors deployed relatively near one another enhances the ability to distinguish between water quality events and technical anomalies, thereby significantly improving the accuracy of anomaly detection. Thus, uncertainty and biases in water quality reporting, interpretation, and modeling are reduced, and the effectiveness of subsequent management actions improved.

Plain Language Summary Digital sensors are commonly used to monitor water quality in rivers and streams, providing real-time data for scientific purposes. However, these sensors are prone to technical anomalies that can affect data reliability and statistical analyses. We propose a framework for detecting anomalies in sensor data from stream networks by leveraging spatial and temporal relationships to improve detection rates. Our framework includes effective methods for handling missing data, aligning time-series, and identifying water quality events. We evaluate advanced statistical methods demonstrating the practical implementation of automated anomaly detection in near-real time. We validate our framework through simulations and a case study in the Herbert River, Queensland, Australia. Results show the effectiveness of the suggested methods in detecting various anomalies. This reduction in uncertainty and biases improves water quality reporting, interpretation, and management actions.

1. Introduction

Water-quality monitoring data are used to better understand and manage natural and anthropogenic impacts on water resources and aquatic ecosystems (Altenburger et al., 2019). In stream and river networks, the monitoring of high-frequency water-quality data often relies on digital sensors, which can be expensive, costing thousands of dollars for a single site (Queensland Government, 2022b). Consequently, monitoring efforts are typically localized, with sensors predominantly deployed at the system (Leigh, Kandanaarachchi, et al., 2019). This provides little information about water-quality dynamics in the upper reaches of the catchment, making it difficult

© 2024. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

A. Villa, Catherine Leigh, Ryan Turner, Cameron Roberts, Kerrie Mengersen

Funding acquisition: Erin E. Peterson, Kerrie Mengersen

Investigation: Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Catherine Leigh, Ryan Turner, Cameron Roberts, Kerrie Mengersen

Methodology: Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Catherine Leigh, Ryan Turner, Kerrie Mengersen

Project administration: Edgar Santos-Fernandez, Erin E. Peterson, Kerrie Mengersen

Resources: Edgar Santos-Fernandez, Erin E. Peterson, Cesar A. Villa, Ryan Turner, Cameron Roberts, Kerrie Mengersen

Software: Edgar Santos-Fernandez, Jay M. Ver Hoef

Supervision: Jay M. Ver Hoef, Erin E. Peterson, James McGree, Kerrie Mengersen

Validation: Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Catherine Leigh, Ryan Turner, Cameron Roberts, Kerrie Mengersen

Visualization: Edgar Santos-Fernandez, Jay M. Ver Hoef, Kerrie Mengersen

Writing – original draft: Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Catherine Leigh, Ryan Turner, Kerrie Mengersen

Writing – review & editing:

Edgar Santos-Fernandez, Jay M. Ver Hoef, Erin E. Peterson, James McGree, Catherine Leigh, Ryan Turner, Kerrie Mengersen

for land managers to identify pollutant sources and/or encourage practice change from private landholders. However, recent advances in sensor technology have led to lower-cost options, including new sensors currently under development (Great Barrier Reef Foundation, 2022). These cost-effective sensors offer a valuable opportunity to collect dense spatial and temporal data from stream networks in near-real time (e.g., Isaak et al., 2017).

While there are evident advantages to employing in-situ sensors for water quality monitoring, it is crucial to acknowledge the significant challenges associated with their use. In-situ sensors are susceptible to various technical issues that can be challenging to identify and mitigate (Hill & Minsker, 2010). These issues encompass calibration errors, biofouling, and battery failure, all of which can introduce inaccuracies into the collected data. We refer to these as technical anomalies, but they have also been referred to as “outliers,” “inaccuracies,” or “special causes of variation.” Some of the most common technical anomalies in water quality data include large and sudden positive or negative spikes, which are referred to as point anomalies (Chandola et al., 2009). Additionally, anomalies can manifest as periods of unusually low or high variability, drift, or shifts in the data, which are known as collective or persistent anomalies (Leigh, Alsibai, et al., 2019). If these anomalies are not properly addressed, they can have detrimental effects on statistical inference, leading to biased estimates and impacting the spatial and temporal autocorrelation structures (Barnett & Lewis, 1984; Congdon, 2019). Therefore, it is imperative to employ robust statistical methods that can effectively distinguish technical anomalies from natural water quality events. By doing so, high-frequency spatio-temporally dense data from in-situ sensors can be utilized to provide valuable insights into water quality dynamics for both management and research purposes.

Another challenge is that catchment water quality is spatially and temporally dynamic at intermediate and broad scales (Peterson et al., 2013). For example, sudden increases in water level, sediment and nutrient concentrations may be observed at multiple sites after extreme weather events (e.g., flooding, heavy rain and increased runoff), or at neighboring time lags as water moves downstream. Not surprisingly, water quality data collected in streams often exhibit spatial (Cressie et al., 2006; Garreta et al., 2010; Peterson et al., 2006) and temporal correlation (Leigh, Kandanaarachchi, et al., 2019), which tends to increase with the spatial and temporal density of data (Isaak et al., 2017; Steel et al., 2016; E. Money et al., 2009; Jackson et al., 2018; Santos-Fernandez et al., 2022).

Generally, spatial, geo-referenced and time-series observations are considered anomalous when they differ substantially from their non-anomalous neighbors (Chandola et al., 2009; Fotheringham & Rogerson, 2008). The challenge in stream systems is that data exhibit unique patterns of spatial dependency due to the branching structure of the network, longitudinal connectivity, directional flow and water volume, which must be accounted for (Cressie et al., 2006; Jackson et al., 2018; O'Donnell et al., 2014; Peterson et al., 2013).

Anomaly detection in time-series, spatial, and spatio-temporal data has been extensively studied in the literature. Numerous examples can be found in various fields such as climate data analyses (Costa & Soares, 2009; Resch et al., 2023), transportation (Djenouri et al., 2019; Shi et al., 2018), and intrusion detection (Wang et al., 2017; Zhang et al., 2020). Methods such as local outlier factor (LOF) (Bosman et al., 2017), leave-one-out cross-validation kriging methods, extreme value theory approaches (Kandanaarachchi & Hyndman, 2021), change-point detection (Tveten et al., 2022), and Hidden Markov models (HMM) (Li et al., 2017) have been proposed to identify spatial and temporal outliers. Some approaches have also been proposed to identify anomalies in data sets from stream networks where temporal dependence exist. Popular methods include autoregressive models (e.g., autoregressive integrated moving average (ARIMA) (Leigh, Alsibai, et al., 2019), vector autoregression models (VAR), machine learning methods (e.g., artificial neural networks (ANN), random forests, Long Short Term Memory) (Jones et al., 2022; Rodriguez-Perez et al., 2020).) However, little work has focused on detecting anomalous data accounting simultaneously for spatial and temporal variation in stream-network contexts, with consideration of the unique spatial relationships found in river networks.

To tackle the complex challenges in water quality monitoring, it is crucial to adopt suitable modeling frameworks that can address the following key aspects: (a) Distinguishing water quality events from technical anomalies, (b) Incorporating spatial and temporal autocorrelation to capture the underlying processes and patterns of interest more accurately, (c) Recursive update models as new data becomes available over time.

In this paper, we present a novel spatio-temporal anomaly detection approach for water quality data obtained from arrays of sensors in stream networks. We introduce a pre-processing framework involving the implementation of data imputation to efficiently handle missing data, optimal alignment of time series to address disparities in time

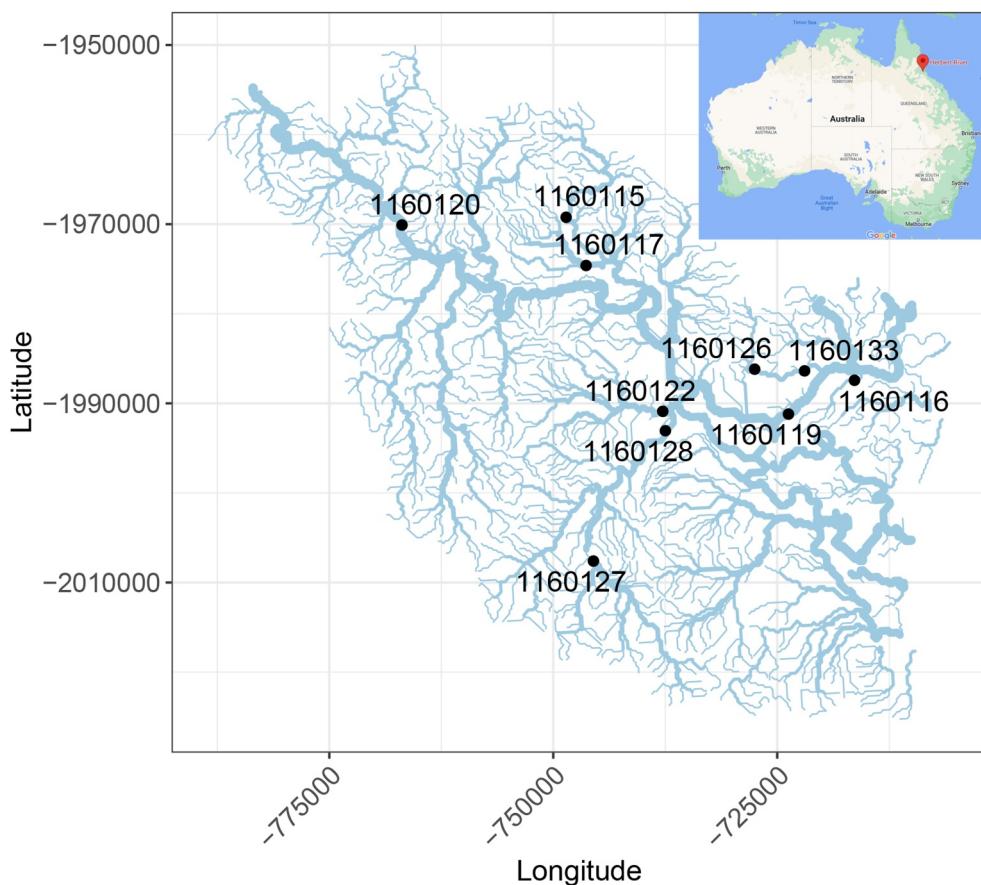


Figure 1. A sub-catchment of the Lower Herbert River in Far North Queensland, Australia. The locations correspond to 10 micro site sensors. The width of each plotted stream segment in the figure is scaled in proportion to the Strahler stream order. We used Albers Equal Area projection centered at 153° longitude, with standard parallels at -20 and -40° latitude.

series, and the identification of water quality events to facilitate the accurate detection of anomalies in spatio-temporal data. The Bayesian modeling approach provides posterior predictive distributions for each observation, which allows the agreement between models and data to be assessed using empirical residuals. Statistical tools including mixture models and HMMs are then used to detect anomalies in these residuals. The Bayesian approach is implemented in near real-time through recursive modeling, which allows the model to update automatically as new data becomes available. We use a comprehensive simulation study to assess and evaluate the ability of the methods to detect different types of anomalies. This is followed by a case-study involving a sensor network located in the Lower Herbert River in Far North Queensland, Australia, which discharges into the Great Barrier Reef Lagoon World Heritage Area.

Herein we focus on technical anomalies in sensor data while taking into account water quality events. The approach introduced in this manuscript has been carefully developed, taking into account the specific characteristics of the case study. The proposed framework aims to effectively detect anomalies and classify *water quality events*, but it also has the potential to be applied to a broad range of environmental monitoring applications. By leveraging the power of advanced statistical and machine learning techniques, this methodology offers a flexible and scalable solution for monitoring various types of water quality data, which should ultimately facilitate decision-making related to the release of trustworthy data to the public domain and in the maintenance of water in-situ sensors.

2. Motivating Data Set: Water Quality in Far North Queensland, Australia

2.1. Study Region and Sites

The Herbert River basin in far north Queensland, Australia (Figure 1) has distinct upper and lower catchments. The upper catchment, in the Southern Atherton Tableland, has an elevation over 1,000 m and an average annual



Figure 2. Opus TriOS optical water-quality sensor (model TLS02 021) produced by TriOS Mess-und Datentechnik GmbH Opus (left) and the installation on Broadwater Creek at Day Use Area/Abergowrie State Forest in the Herbert River (site 1,160,115).

rainfall just under 1200 mm; whereas the lower catchment is a large coastal delta receiving more than double the rainfall. The land use in the upper part of the catchment is dominated by grazing while the lower part of the catchment is dominated by sugarcane and conservation (Queensland Government, 2022a). The monitoring sites for this study were located in the Stone River sub-catchment of the Herbert River basin. This particular area is dominated by sugarcane and conservation as well as a small amount of forestry land use.

Water quality in this area typically changes with rainfall induced by runoff from the catchments. More specifically at the start of the wet season (October–December), monitoring will observe a *first flush* phenomenon where high concentrations of sediment or nutrients will occur with small waterway discharges. Later in the wet season, large events will also contribute to high concentrations of sediment (Queensland Government, 2022a). However, confounding occurs if many subsequent rainfall events happen, causing the catchment to become exhausted of sediment delivery potential and increasing discharge will not necessarily result in increasing concentrations.

2.2. Sensor Installation and Data Collection

We installed 10 Opus TriOS water-quality sensors (Figure 2) produced by TriOS Mess-und Datentechnik GmbH in the branching network of the Herbert River (Figure 1). These Internet of Things (IoT) devices measure among other parameters stream water level (m) and spectral total suspended solids (TSSeq) (referred as TSS from now on) (mg/L). Measurements were recorded every 30 min and uploaded via the 4G-CATM1 mobile network to a database. Each sensor was installed at the edge of the riverbank in a location best representing that part of the waterway. This positioning while given potential for edge effects in the data allows for service access during the majority of flow conditions while maintaining a safe distance from the bed to avoid issues arising from sedimentation. It is important to note that due to this fixed positioning, the probes are only able to measure when the water level is above the lens path of the probe.

2.3. Anomalies in the Data

In this study, we used time-series data for TSS and water level collected between 2021-06-15 and 2022-03-02. The data set illustrate many of the intrinsic challenges of in-situ water quality data described in the Introduction, including missing values, high variability between sites, persistent anomalies and other factors that can complicate the analysis (Figure 3). A team of water quality specialists from the Queensland Department of Environment and Science visually assessed the data and created a labeled data set. This data set contains labels of the technical anomalies in TSS following the framework developed by Leigh, Alsibai, et al. (2019). These

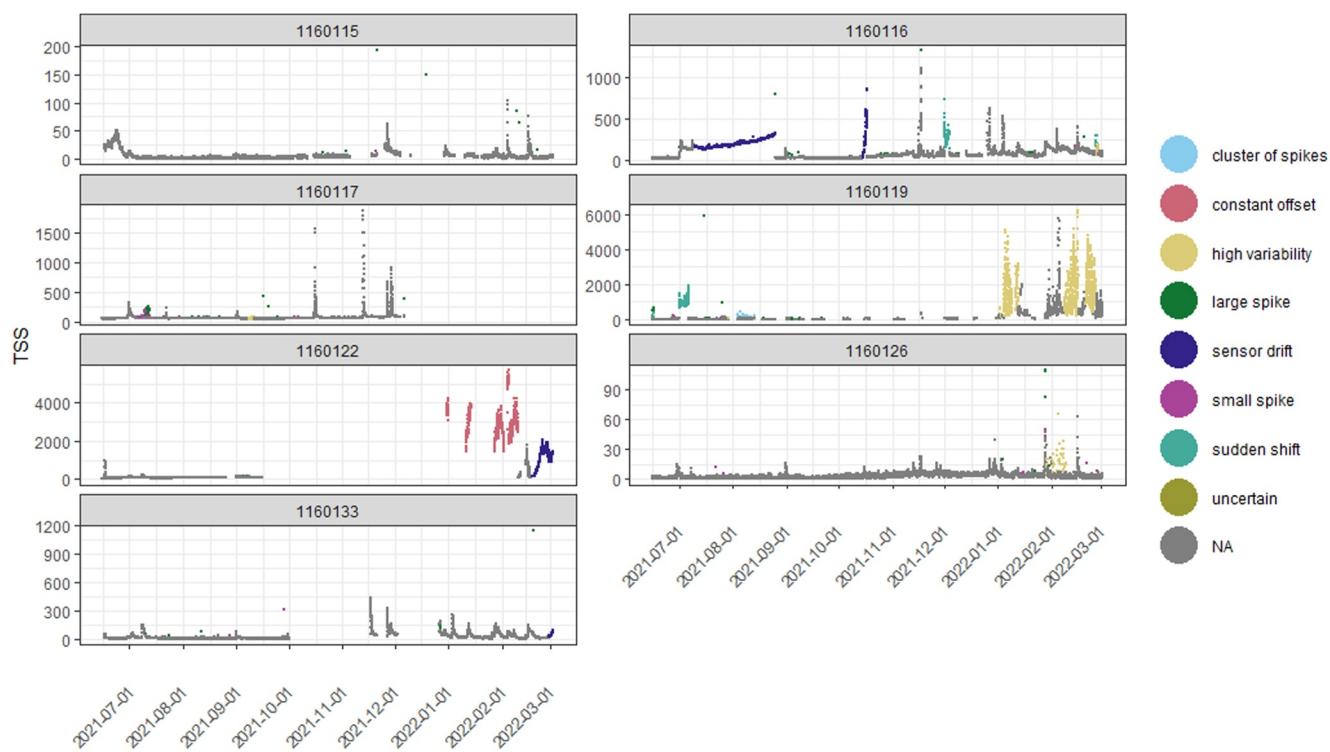


Figure 3. Time series of total suspended solids (TSS) data and anomalies. Anomalous data points are colored according to eight types. NA refers to non-anomalous data.

included comparatively large and small sudden spikes (0.131% and 0.164% of the data, respectively), clusters of spikes (0.040%), periods of anomalously low variability and/or persistent values and constant offsets (e.g., due to calibration errors (1.089%), sudden shifts (0.562%), drifts (4.262%) and periods of anomalously high variability (1.668%). The data set also includes labels of the water events associated with rainfall (ambient conditions represent 72.308%, while events and sensors out of the water are 14.852% and 12.841% respectively).

The identified sensor anomalies in the data set are generally caused by physical interference that affects the sensor readings, introducing noise and biases, thereby compromising the accuracy and reliability of the data. Such interference can manifest as spikes in the data, which may be caused by various factors, including obstructions in the probe's lens path, such as rocks, organic debris, or living organisms like fish and other fauna.

The sensor installations leverage an attached wiper unit to clean the sensor lens from bio-fouling and other debris. However, the failure or malfunction of these units can lead to the wiper blade obstructing the lens path, causing sudden shifts or high variability in the data. If the wiper unit fails or the wiper blade becomes damaged, bio-fouling can accumulate on the lens, resulting in sensor drift. This drift refers to the gradual deviation of sensor readings from their actual values over time, which typically tends to increase in a monotonic manner.

While digital sensors are generally less susceptible to electrical interference, an insufficient power supply can result in constant values being reported due to the probe relying on memory. On the other hand, the level sensors used in these installations are more analog and can be susceptible to electrical interference, often due to poor grounding connections. In such instances, high variability or negative values may be produced, depending on the specific error mechanism. Some of these anomalies are very challenging for identification using traditional methods. For instance, in location 1,160,116, there were two long periods of sensor drift (in purple). Similarly, in location 1,160,119, there was a sudden shift on the mean of the process (light green) and a period of high variability.

3. Spatio-Temporal Models

In this section, we introduce a spatio-temporal model that we use to describe stream network data. Consider S unique spatial locations on a stream network where sensors are installed. Repeated water quality measurements

(e.g., turbidity, nitrate, conductivity or water level) are collected at all spatial locations and at regular time points $t = 1, 2, \dots, T$, producing spatio-temporal stream network data. Let \mathbf{y}_t be an $S \times 1$ vector of random variables observed at time t at spatial locations $s = 1, 2, \dots, S$.

Several statistical models have been proposed to capture spatio-temporal dependence in stream network data, including those suggested by E. S. Money et al. (2009), Jackson et al. (2018), Tang and Zimmerman (2020), and Santos-Fernandez et al. (2022). In this work, we follow the approach of Santos-Fernandez et al. (2022); Santos-Fernandez et al. (2023) where the following vector auto-regressive spatial linear model is used to describe potential spatio-temporal dependencies in stream data.

$$[\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{X}_t, \mathbf{X}_{t-1}, \boldsymbol{\beta}, \boldsymbol{\Phi}_1, \boldsymbol{\Sigma}, \sigma_0^2] = \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma} + \sigma_0^2 \mathbf{I}), \quad (1)$$

$$\boldsymbol{\mu}_t = \mathbf{X}_t \boldsymbol{\beta} + \boldsymbol{\Phi}_1 (\mathbf{y}_{t-1} - \mathbf{X}_{t-1} \boldsymbol{\beta}), \quad (2)$$

where \mathbf{y}_{t-1} is the process at time $t - 1$, $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_T]$ where each \mathbf{X}_t is a $S \times p$ design matrix of p covariates at time t , and $\boldsymbol{\beta}$ is a vector of regression coefficients of length p . The total spatial variance-covariance of \mathbf{y}_t can be written as $\boldsymbol{\Sigma} + \sigma_0^2 \mathbf{I}$. Here $\sigma_0^2 \mathbf{I}$ incorporates the unstructured error term, where σ_0^2 is known as the nugget effect and \mathbf{I} is the identity matrix. The notation $\mathcal{N}(\cdot, \cdot)$ is the probability density function for the multivariate normal distribution.

In Equation 2, the mean $\boldsymbol{\mu}_t$ is expressed as a vector autoregression of order one, VAR(1), where $\boldsymbol{\Phi}_1$ is a $S \times S$ square transition matrix with diagonal parameters ϕ_s that determine the magnitude of first-order temporal autocorrelation. Note that this model can be extended to include more complex autocorrelation terms, for further details see Santos-Fernandez et al. (2022). By setting $\boldsymbol{\Sigma}$ and $\boldsymbol{\Phi}_1$ to zero and neglecting the temporal and spatial correlation, we can derive the conventional linear regression model.

Studies in the literature have investigated spatial correlation in stream data, with research conducted among others by Cressie et al. (2006); Ver Hoef et al. (2006); E. Money et al. (2009); Garreta et al. (2010); Steel et al. (2016); Jackson et al. (2018); Tang and Zimmerman (2020). These approaches allow information to be borrowed from neighboring sites, which is crucial in identifying anomalies in spatio-temporal data. We can express the spatial covariance matrix $\boldsymbol{\Sigma}$ of dimension $S \times S$ using a mixture approach (Ver Hoef & Peterson, 2010), that describes multiple patterns of spatial autocorrelation often found in data collected on streams:

$$\boldsymbol{\Sigma} = \mathbf{C}_{ED} + \mathbf{C}_{TU} + \mathbf{C}_{TD} = \sigma_e^2 \mathbf{R}_e(\alpha_e) + \sigma_u^2 \mathbf{R}_u(\alpha_u) + \sigma_d^2 \mathbf{R}_d(\alpha_d). \quad (3)$$

Here, \mathbf{C}_{ED} is a standard geostatistical covariance matrix based on Euclidean distance. The terms \mathbf{C}_{TU} and \mathbf{C}_{TD} are tail-up and tail-down covariance matrices based on stream distance.

The parameters σ_e^2 , σ_u^2 and σ_d^2 are the partial sill, and analogous parameters α_e , α_u , α_d are the spatial range parameters. Similarly, \mathbf{R}_e , \mathbf{R}_u and \mathbf{R}_d are the correlation matrices. Several variations of this model can be fit using the R statistical software (R Core Team, 2022) package SSNbayes (Santos-Fernandez et al., 2023). Here we fit a fully Bayesian model in Stan (Stan Development Team, 2018) to estimate the posterior distribution via MCMC. A reproducible example demonstrating the use of the SSNbayes package for anomaly detection is available on <https://www.kaggle.com/code/edsans/anomaly-detection-using-ssnbayes/notebook>.

Tail-up/tail-down covariance models are formulated by integrating a kernel function over a white noise process that is strictly upstream/downstream of a given site. The tail of the moving-average function points upstream in the tail-up model, and downstream in the tail-down model. In the tail-up model, autocorrelation is confined to sites that are connected by flow, while in the tail-down model, spatial dependence can still exist even between sites that are not directly connected by flow. In the tail-up and tail-down models, the distance is restricted to the stream network, which captures unique spatial relationships produced by the branching network structure, longitudinal connectivity, flow volume and flow direction of stream networks. See (Ver Hoef et al., 2006; Ver Hoef & Peterson, 2010) for more details.

The estimation of spatial covariance matrices involves estimating the partial sill and range distributions. To model these parameters, we employ flat uniform priors for both spatial components. Additionally, we set an upper bound of four times the maximum distance between observations in the river network for the range parameter. The

temporal structure is defined by the matrix Φ , and we adopt a uniform prior ranging from -1 to 1 to model the diagonal autoregressive parameter.

3.1. Bayesian Approaches for Spatio-Temporal Data

To implement the above model within a Bayesian framework we use the Bayes theorem, which involves the computation of the posterior distribution for the parameters of interest $\theta = \{\beta, \alpha, \sigma, \Phi, \sigma_0\}$:

$$f(\theta|y) = \frac{f(y|\theta)f(\theta)}{\int f(y|\theta)f(\theta)d\theta}, \quad (4)$$

where $f(\theta)$ represents the prior distribution for θ and $f(y|\theta)$ is the likelihood of the data y given θ .

Model goodness-of-fit can be assessed by comparing the data (y_t) to the predictions obtained via posterior predictive distributions, which measure the agreement between the assumed model and the observed data. The ability to identify observations that significantly deviate from the predictions is especially useful for detecting anomalies in the data (Gelman & Hill, 2006). The fitted vector autoregression spatial stream-network model is used to generate the posterior predictive distribution (PPD), (\hat{y}_t) , which is obtained from the observed response variable (y_t) and covariates for the observed and predicted data (X_t) and (\hat{X}_t) by marginalizing over the likelihood as follows:

$$p(\hat{y}_t|y_t, X_t, \hat{X}_t) = \int p(\hat{y}_t|\theta, \hat{X}_t) p(\theta|y_t, X_t) d\theta, \quad (5)$$

where θ is a vector of estimated parameters from the fitted model (i.e., posterior distributions). Note that the vector y_t may contain missing values, which may occur for various reasons. These missing values are imputed using the posterior predictive distributions integrated into our Bayesian inference approach.

The vectors X_t and \hat{X}_t are the same when we produce predictions for the observed data points. The model prediction error or empirical residual for an observation $y_{s,t}$ at location s and time t is then $e_{s,t} = y_{s,t} - \hat{y}_{s,t}$, where $\hat{y}_{s,t}$ is calculated based on a summary of location of the posterior predictive distribution such as the posterior mean.

4. Anomaly Detection Methods

Several statistical approaches can be used to detect anomalies in stream data once a spatio-temporal model is fit, and the posterior predictive distribution and the prediction errors have been computed using the estimated parameters. Here we consider four different approaches including the (a) posterior predictive approach, (b) finite mixture approach, (c) hidden Markov modeling approach, and (d) autoregressive integrated moving average method. Importantly, these approaches are all considered unsupervised, meaning that they do not require a data set containing labeled anomalies and non-anomalous/normal data. Unsupervised models provide multiple benefits, since labeled data sets are often unavailable for new data sources or monitoring sites, and are time consuming to generate manually. In addition, labeled data sets may not be transferable to new sites or time periods if they have been generated under different environmental conditions.

4.1. Method 1: Posterior Predictive Approach

The first approach we introduce is obtained via the posterior predictive distribution from Equation 5. Specifically, this approach is reasonable as it compares the predicted values from the model with the associated uncertainty to the observed data, and is based on the assumption that anomalous observations will have a low probability under the fitted model (Murphy, 2012). This method is expected to be effective in identifying spikes in the data and sudden changes in the mean of the process. The steps to implement this approach are as follows:

Steps:

1. We define anomalous observations as those that fall outside the 95% highest posterior density or posterior prediction interval (HDI). This significance level will be utilized in both the simulation and the case study. Note that tighter limits should be used if the data are expected to have a large proportion of anomalies.

2. Exclude these anomalous observations from the data.
3. In one of our approaches we iterate steps 1–2 twice. Anomalies identified in both the first and second iterations are combined, flagged, and used to compute the performance measures.

For this study, we considered up to two iterations, to ensure that we did not exclude too many observations, but other alternatives are possible (Weisberg, 2013). Our first approach (iter_1) only involves fitting the model and no data are excluded (i.e. only uses step 1); whereas in the second case (iter_2), the model is fitted again, after removing the anomalous observations (steps 1, 2, 3 and 1). Note that this is in accordance with control chart methodologies where existing outliers are removed while estimating the control limits (Montgomery, 2020). This method considers the uncertainty in the process and can produce a probabilistic score for each observation being an anomaly, which allows observations to be flagged (e.g., data not included in the 80% HDI).

Next, we concentrate on a range of statistical methods based on the residual $e_{s,t}$ obtained using the mean of the posterior distributions. However, these approaches can be extended to use the samples from the Markov chain Monte Carlo (MCMC) simulations. That is, the proposed approach would enable the implementation of fully Bayesian anomaly detection methods, utilizing all the samples drawn from the posterior distribution instead of relying on the point estimates of the residuals.

4.2. Method 2: Finite Mixture Approach

Mixtures are statistical models that are particularly useful to describe response variables from a population composed of different groups or is generated under different conditions (Gelman et al., 1995). Mixtures and other clustering-based methods have been widely used to detect anomalies (Chandola et al., 2009). In these settings, we assess heterogeneity in the residuals to identify anomalous data, especially when shifts in the mean of the process occurs. Latent observed groups in the residuals could indicate the presence of anomalous data. Anomalous data are expected to fall outside the range of the typical distribution and reside in the tails, indicating the existence of different latent classes. These anomalous groups may have different statistical properties compared to the non-anomalous data, such as a higher variance or a different statistical distribution.

To implement this approach, we propose the following steps.

1. First, we need to determine the appropriate number of clusters or components (e.g., $K = 3$). This can be done by inspecting the distribution of the residuals. For example, three components could be used if we expect anomalies above and below the mean.
2. Let $z_{s,t}$ be the cluster membership of the observed residual $e_{s,t}$ at location s and time t , and $z_{s,t} \sim \text{categorical}(1/K)$ and Z denotes the number of clusters.
3. Fit the Bayesian mixture model $e_{s,t}|z_{s,t} \sim \mathcal{N}(\mu_{z_{s,t}}, \sigma_{z_{s,t}})$, which produces a cluster membership as well as estimates of the mean and variance for each cluster.
4. Inspect the resulting clusters to identify anomalous components, which are typically characterized by observations found in the tails of the residual distribution and are likely to contain fewer observations than the main cluster(s) representing the non-anomalous data.

4.3. Method 3: Hidden Markov Model Approach

In a hidden Markov model (HMM), the response variable is determined by a hidden Markov process, where the state at time t depends only on the state at the previous time point $t - 1$ (Zucchini & MacDonald, 2009). The HMM model can be defined by $\delta = (A, B, \pi)$, where A and B are the transition probability matrix and emission probability matrix, respectively, and π is the initial state probability. The transition probability matrix gives the probability distribution of observed variables given the hidden states of the model, while the emission probability matrix describes the probability distribution of observed variables conditional on the state.

This model is often referred to as dependent or dynamical mixtures and is employed with time series data. Our proposed modeling framework involves the consideration of two or more unknown latent states (L), such as normal and anomalous, and we seek to compute the probability of each observation belonging to these states. Under this model, it is assumed that the process transitions between the latent states using a transition probability matrix.

Conditional on the state l_t , the distribution of the residual e_t is assumed to be normally distributed with a mean μ_{l_t} and a variance $\sigma_{l_t}^2$. The HMM approach is commonly used in situations where data come from multiple latent states, such as non-anomalous and anomalous states.

$$e_t|l_t = \mathcal{N}(\mu_{l_t}, \sigma_{l_t}^2), \quad (6)$$

where $\mathcal{N}(,)$ is the probability density function of the normal distribution.

To implement this approach, the following two steps are used:

1. Fit a HMM model with $L = 2$ or more components.
2. Use the component's membership to identify potential anomalies.

Our method is implemented using a Bayesian approach, which assigns each observation a probability of belonging to either the anomalous or non-anomalous cluster. This approach provides flexibility in the allocation of observations to different states. We have also incorporated the assumption of unequal variances between the clusters.

4.4. Method 4: Autoregressive Integrated Moving Average (ARIMA) (Benchmark Method)

An ARIMA approach (Leigh, Alsibai, et al., 2019) is purely temporal, meaning that it does not account for spatial correlation, and is used to benchmark the spatio-temporal methods proposed here. An appropriate ARIMA model was obtained for each time series using the `auto.arima` function in R from the `forecast` package (Hyndman & Khandakar, 2008) which finds the best model based on AIC. Observations were considered anomalous if they fell outside of the 95% prediction interval. Following Leigh, Alsibai, et al. (2019), the general ARMA model can be described as:

$$y_t = c + \phi_1 y_{t-1} + \cdots + \phi_m y_{t-m} + \theta_1 \epsilon_{t-1} + \cdots + \theta_q \epsilon_{t-q} + \epsilon_t \quad (7)$$

where c is a constant term, ϕ_1, \dots, ϕ_p are scalar autoregressive coefficients of the m terms, $\theta_1, \dots, \theta_q$ are the coefficients of the q moving average terms $\epsilon_{t-1}, \dots, \epsilon_{t-q}$ are the errors at times $t - 1$ through $t - q$.

5. Simulation Study

The methods previously presented are largely unexplored in the detection of technical anomalies in spatio-temporal stream network data. Hence, we performed a simulation study to assess the suitability of methods 1–4 under several scenarios found in real-world applications including in our case study.

To begin, we generated simulated data sets using the following steps.

1. The installation and maintenance of TSS and nitrate sensors entail a considerable financial cost, and consequently, monitoring programs in Queensland, Australia, typically have fewer than 30 sensors per catchment. Despite this limited number of sensors, it is still possible to achieve a reasonable and realistic spatial coverage that exhibits spatial correlation in the network. The use of a limited number of sensors is a common scenario in environmental monitoring, and it poses several challenges such as optimal sensor placement and data sparsity issues. Accordingly, we created artificial stream networks made up of 150 segments and $S = 30$ spatial locations representing water quality sensors. This was undertaken via using the function `createSSN()` in the `SSN` R software package (Ver Hoef et al., 2014). Figure 4 shows a simulated stream network with 30 spatial locations across the first four time points.
2. We selected a common covariance structure to model spatial dependence by means of a exponential tail-down model. To specify this structure, we used the following spatial parameters: partial sill, which represents the variance after accounting for the nugget effect, with a value of $\sigma_{TD}^2 = 3$ (this value was selected based on case study from Santos-Fernandez et al., 2022); the spatial range (defines the rate of decay in the covariance as a function of the separation distance) $\alpha_{TD} = 10$ (see Equation 3), and nugget effect $\sigma_0^2 = 0.1$ (Equation 1).
3. We considered three covariates for the mean response which were arbitrarily generated from three standard normal covariates plus an intercept ($p = 4$) at every spatial location with regression coefficients: $\{\beta_1, \beta_2, \beta_3\} = \{1, 0, -1\}$ for the slopes and the intercept $\beta_0 = 10$.

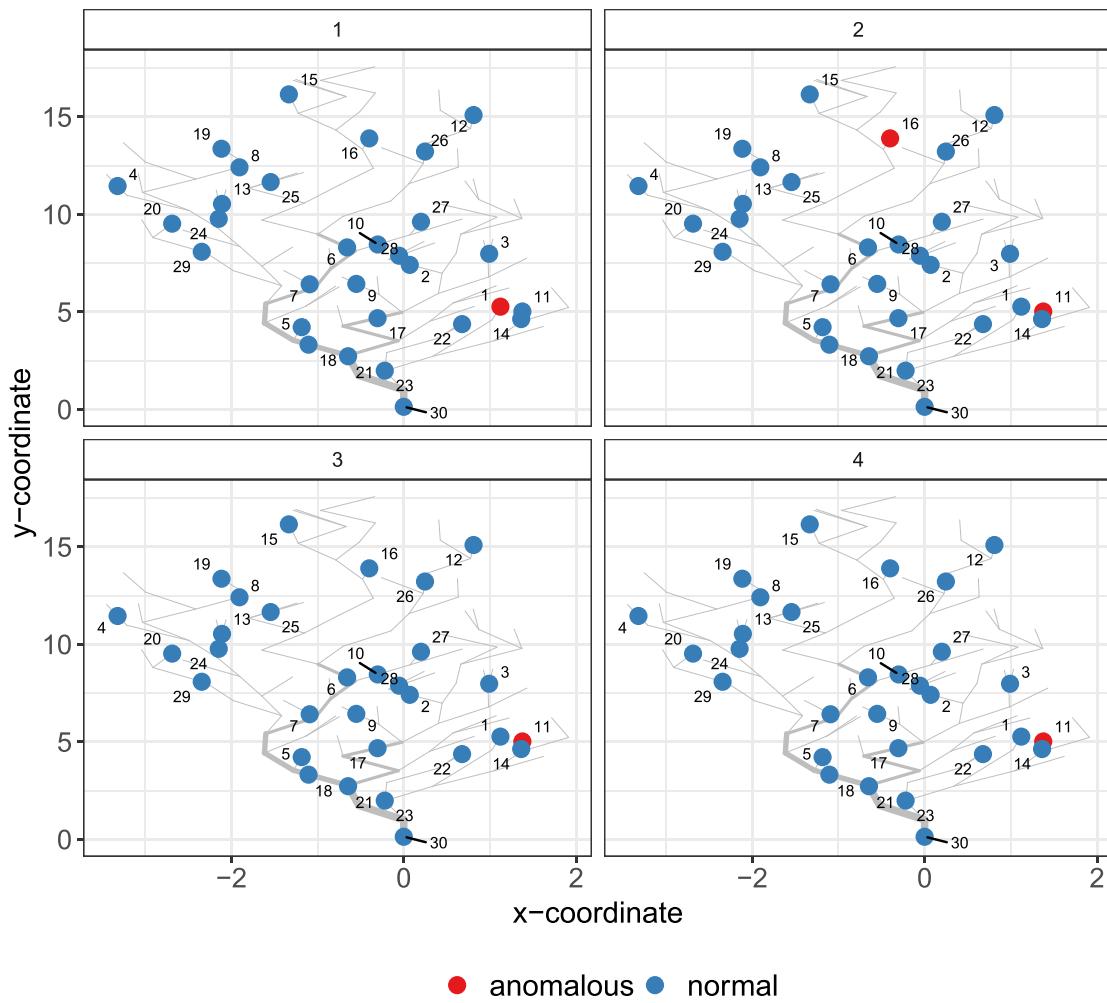


Figure 4. The artificial stream network with 30 spatial locations and observations at the first four time points. Technical water quality anomalies are shown in red and normal data are blue. The relative flow volume for each stream segment is proportional to the width of the gray line.

4. The spatio-temporal data were generated additively:
 - a. First, a response variable y was simulated from a spatial stream-network process (at time point $t = 1$) using $y = X\beta + \nu$, where $\nu \sim N(0, \Sigma)$.
 - b. Next, we defined a temporal process based on an autoregressive function of order one, AR(1), with $\phi = 0.80$ and included $T = 120$ time points. We generated the time series data using the corAR1 function from the nlme package (Pinheiro et al., 2020).
 - c. We added the temporal variation (plus some random variation $\mathcal{N}(0, 1)$) to the spatial process simulated in (a), resulting in a spatio-temporal stream network data set.
5. To create anomalies, we first defined $q_{ini} = 0.05$, which represents the *initial* probability of an observation at time t at a spatial location s being an anomaly. Within the time series, a binary indicator d representing the start of an anomaly was created at each location s and time t :

$$d_{st} \sim \text{Bern}(q_{ini}) \begin{cases} 0 & \text{non-anomalous,} \\ 1 & \text{anomalous.} \end{cases} \quad (8)$$

The binary indicator is also shown in Figure 4 in the spatial context of the network.

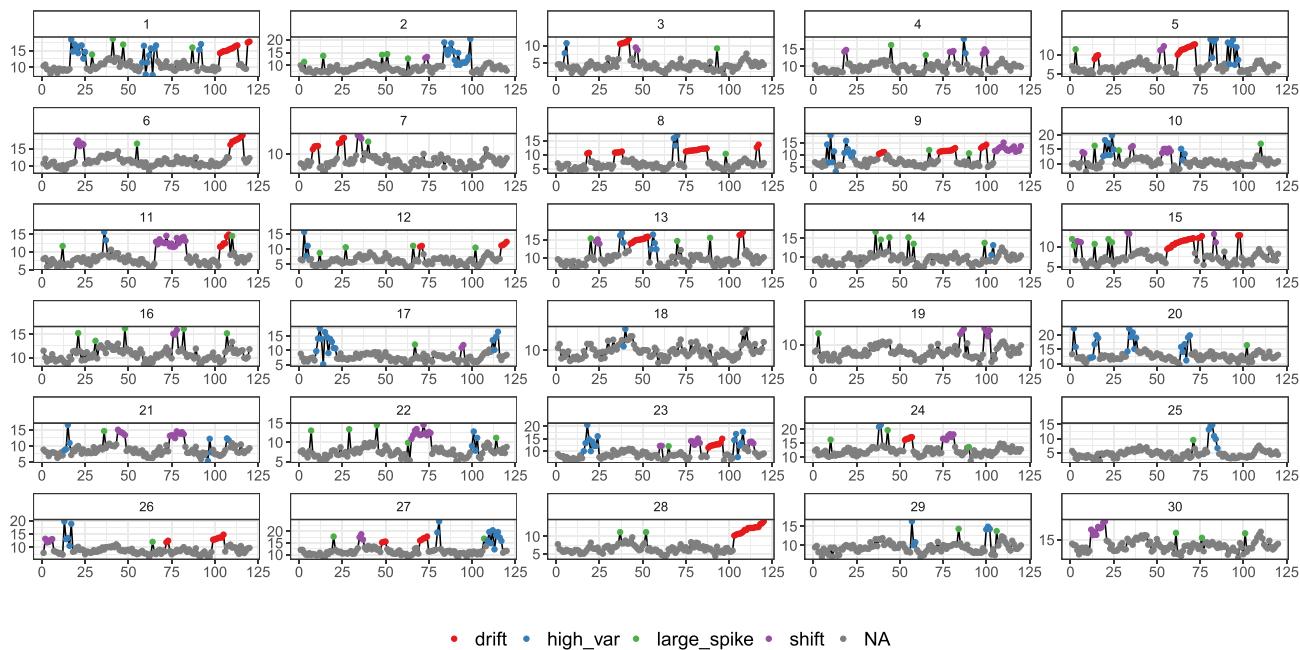


Figure 5. Time series of anomalous and normal observed data at 30 locations, across the 120 time points. Anomalous data represent drift, high variability, large spikes, and sudden shifts and non-anomalous data are labeled as NA.

6. We considered four types of anomalies according to observed data in the case study introduced in the motivating data set and reported in the literature (Leigh, Alsibai, et al., 2019). This included (a) moderate/large sudden spikes, (b) anomalously high variability, (c) a shift in the mean, and (d) drift; each with the same initial probability of occurrence.
7. For each $d_{s,t} = 1$ we simulated the persistence of the anomaly (i.e., number of consecutive observations with anomalous data) using a Poisson distribution $n_d \sim \text{Poisson}(\lambda)$ with rate 0.8, which determines subsequent anomaly indicators $d_{s,t+1}, d_{s,t+2}$, etc. Persistent anomalies occur for high variability, shift, and drift anomalies, while spikes are only composed of one observation at a single time point.
8. We defined the observed response variable, y_{obs} , as a mixture distribution:

$$y_{obs_{st}} = \begin{cases} y_{st} + d_{s,t} \mathcal{N}(5, 1) & \text{for spikes,} \\ y_{st} + d_{s,t} \mathcal{N}(1, 5) & \text{for high variability,} \\ y_{st} + d_{s,t} \mathcal{N}(5, 1) & \text{for shift,} \\ y_{st} + d_{s,t}(\mathcal{N}(5, 1)) & \text{for drift,} \end{cases} \quad (9)$$

where (\cdot) represents sort ascending to create a drifting pattern.

A time series representation of the simulated data is shown in Figure 5, where non-anomalous data are represented in gray and the different anomaly types are represented with red, purple, green and blue dots. Notice that some anomalies were one-off cases (e.g., spikes in location 28), while others persisted through time at a single location (e.g., drift in location 28). For access to the R code used to reproduce the simulations, please see the Data Availability Statement.

5.1. Simulation Results

After fitting the spatio-temporal model from Equation 1, we used approaches 1–4 (PPD, mixtures, HMMs, ARIMA) to detect the anomalies in the simulated data set. Two iterations were considered for the PDD and the finite mixtures. We simulated 100 data sets and computed the performance using a variety of model measures across the different methods.

Table 1

A Comparison of Model Performance Measures and Qualitative Factors Using ARIMA (Benchmark), PPD (Two Iterations), Finite Mixtures (Two Iterations) and the HMM Models

Approach	<i>se</i>	<i>sp</i>	<i>acc</i>	<i>acc_{adj}</i>	<i>MCC</i>	Complex
ARIMA	0.3390	0.9485	0.8690	0.6437	0.3388	Low
PPD (iter_1)	0.3593	0.9707	0.8921	0.6650	0.4272	Low
PPD (iter_2)	0.4849	0.9616	0.9002	0.7232	0.5064	Med
Finite mixtures (iter_1)	0.8896	0.6479	0.6787	0.7687	0.3638	High
Finite mixtures (iter_2)	0.8749	0.6888	0.7124	0.7818	0.3874	High
HMM	0.5587	0.9450	0.8953	0.7519	0.5197	High

Note. Performance measures include sensitivity (*se*), specificity (*sp*), accuracy (*acc*), adjusted accuracy (*acc_{adj}*), and the Matthews correlation coefficient (*MCC*). Model computational complexity is a subjective measure based on our perception. In bold, we highlight the best-performing model.

The sensitivity of the method ($se = TP/(TP + FN)$) measures the ability to detect the true positive (*TP*) anomalies when they are present. Here, false negatives (*FN*) are truly anomalous observations classified as non-anomalous. We also calculated the specificity ($sp = TN/(TN + FP)$), which measures the ability of the method to correctly identify normal (non-anomalous) data when anomalies are absent (*TN*). The false positives (*FP*) occur when non-anomalous observations are classified as anomalous. The overall accuracy of the method ($acc = (TP + TN)/(TP + TN + FN + FP)$) measures the proportion of correctly identified observations. However, the data are unbalanced (i.e., few anomalies compared to normal data) and so we also employed the adjusted or balanced accuracy (acc_{adj}) (Brodersen et al., 2010) as the midpoint between the *se* and the *sp*. Finally, we calculated the Matthews correlation coefficient ($MCC = \frac{TP \cdot TN - FP \cdot FN}{(TP + FN) \cdot (TP + FP) \cdot (TN + FP) \cdot (TN + FN)}$) (Matthews, 1975), which makes use of all the elements of the confusion matrix (containing the number of correct and incorrect predictions made by the model) to measure the quality of a binary classification.

Model overall performance measure results are provided in Table 1. All of the methods outperformed the benchmark (i.e., ARIMA) in terms of accuracy, adjusted accuracy, *MCC* and *se*. These results indicate that our spatio-temporal approaches have a higher efficacy in detecting anomalies compared to the benchmark method. The PPD method tended to have lower sensitivity than the finite mixtures and the HMM. This could be caused due to the bias in the parameters produced by the anomalies. Removing the anomalies and refitting the model, tended to produce better performance overall in terms of *MCC*. The PPD using two iterations and HMM were the best individual performing method in terms of *MCC*. A boxplot comparing the performance of the six methods based on the *MCC* is presented in Figure 6.

Persistent anomalies, which occur at multiple time points, were challenging for all of the methods to detect because they affect the estimated parameters (i.e., temporal correlation), causing bias and distorting the predictions. However, these models captured well the beginning and the end of the anomalous period, which may be useful for manual quality coding or labeling the whole period as abnormal by water quality specialists.

We calculated performance measures by method and anomaly type to determine whether the methods had a stronger ability to detect some anomaly types more than others (See Table 2). Large spikes were identified by all the approaches with high sensitivity, especially by the mixtures ($se > 0.99$). The PPD seemed to have less propensity for false positives ($1 - sp$) and the highest *MCC*. Anomalies that involved high variability in the data were better identified using HMM based on *MCC* and adjusted accuracy. When shifts in the mean occurred, the HMM also tended to outperform the rest of the methods. Drift was by far the hardest type of anomaly to identify. Apart from the finite mixtures, the rest of the methods had a low probability of detecting these subtle patterns in the data (low *se*). Nevertheless, most of the models tended to detect the beginning of the drift period, which could be used to automatically flag these patterns in the data set and then manually adjust them. In terms of model computational complexity, the PPD (iter_1) method is relatively low in complexity since it involves fitting the spatio-temporal model only once. The PPD (iter_2) model, on the other hand, has a medium complexity as it requires fitting the models twice after excluding anomalies in the first iteration. Mixtures and the HMM exhibit high complexity due to their reliance on residual computation in the spatio-temporal model and the inherent

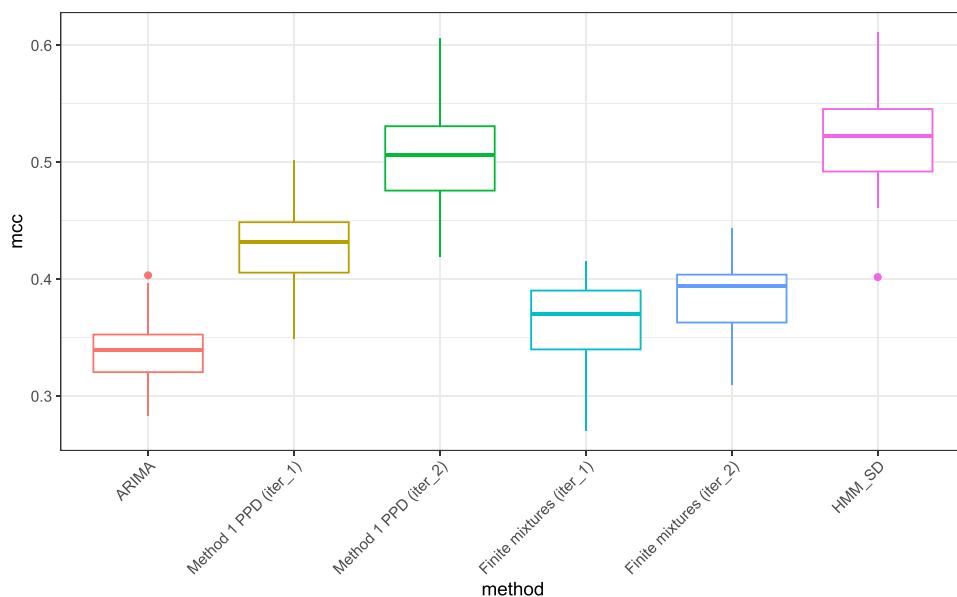


Figure 6. Comparison of the Matthews correlation coefficient (MCC) in 100 simulated data sets.

challenges associated with handling identifiability issues. This qualitative aspect is critical from both the implementation and operational perspectives. It guides the decision-making process toward simpler models when all methods (PPD, mixtures, and HMM) perform well.

The posterior distributions of the regression coefficients obtained from the fitted model indicated that all the regression parameters were slightly biased, apart from β_1 which captured the true value of 1 (Figure A1 in Appendix A), demonstrating the negative impact of anomalies in the model. The temporal autoregression parameter, ϕ , was slightly overestimated by the model (Figure A2 in Appendix A). This was also true of the spatial autoregression parameters; the model underestimated the partial sill, $\sigma_{TD}^2 = 3$, and overestimated the nugget effect, $\sigma_{NUG}^2 = 0.1$ and the range parameter, $\alpha_{TD} = 10$ (Figure A3 in Appendix A).

6. Case Study: Analysis and Results

In this section, we delve into the analysis of the motivating data set introduced in Section 2.

6.1. Analysis of Water Level Using IMPALE - A Data Pre-Processing Framework

First, we will implement a pre-processing step to enhance the data quality and improve overall performance of the anomaly detection approaches. The objective of this process is to discern events that can subsequently serve as covariates in the statistical model, fill in missing covariate data gaps, and evaluate the similarity among sites.

In this study, we introduce the IMPALE framework, a three-step approach designed specifically for water level analysis. IMPALE involves the following steps: IMPute and fill the gaps, ALign the time series, and assess the relationship and spatial dependence among sensors, and carry out Events identification.

1. **Imputation:** water level sensors often get out of the water producing unreliable level data. Note that this might or might not cause anomalies in the other parameters such as TSS as the sensors might be independent. We start by performing multivariate time series imputation using the R package *mtsdi* (Junger & de Leon, 2018). The results from the imputation are shown in Figure 7. An alternative imputation using Bayesian inference can be implemented using package *SSNbayes* (Santos-Fernandez et al., 2023), which allows capturing spatial dependence and propagating the uncertainty of the prediction on water level into the TSS estimation.
2. **Alignment:** optimal match and alignment of the time series was obtained using Dynamic Time Warping (DTW) which minimizes the sum of absolute differences among time series (Senin, 2008). This approach

Table 2

Performance Measures by Anomaly Type, With Models (ARIMA, PPD With Two Iterations, Finite Mixtures With Two Iterations and the HMM Models) Listed in Increasing Level of Complexity

Anomaly	Approach	se	sp	acc	acc_adj	MCC
large_spike	ARIMA	0.8077	0.9489	0.9459	0.8783	0.4360
large_spike	Method 1 PPD (iter_1)	0.8183	0.9709	0.9675	0.8946	0.5493
large_spike	Method 1 PPD (iter_2)	0.9315	0.9616	0.9609	0.9466	0.5595
large_spike	Finite mixtures (iter_1)	0.9939	0.6468	0.6544	0.8203	0.1951
large_spike	Finite mixtures (iter_2)	0.9888	0.6885	0.6951	0.8386	0.2120
large_spike	HMM	0.9479	0.9454	0.9454	0.9466	0.5027
high_var	ARIMA	0.4258	0.9489	0.9276	0.6874	0.2941
high_var	Method 1 PPD (iter_1)	0.4534	0.9709	0.9503	0.7122	0.3920
high_var	Method 1 PPD (iter_2)	0.5590	0.9616	0.9455	0.7603	0.4271
high_var	Finite mixtures (iter_1)	0.7004	0.6468	0.6488	0.6736	0.1401
high_var	Finite mixtures (iter_2)	0.7826	0.6885	0.6920	0.7355	0.1943
high_var	HMM	0.6489	0.9454	0.9336	0.7971	0.4309
Shift	ARIMA	0.2081	0.9489	0.9142	0.5785	0.1383
Shift	Method 1 PPD (iter_1)	0.2335	0.9709	0.9372	0.6022	0.2178
Shift	Method 1 PPD (iter_2)	0.3824	0.9616	0.9349	0.6720	0.3114
Shift	Finite mixtures (iter_1)	0.9173	0.6468	0.6590	0.7820	0.2404
Shift	Finite mixtures (iter_2)	0.8260	0.6885	0.6946	0.7572	0.2249
Shift	HMM	0.4766	0.9454	0.9237	0.7110	0.3334
Drift	ARIMA	0.1592	0.9489	0.9194	0.5541	0.0879
Drift	Method 1 PPD (iter_1)	0.1578	0.9709	0.9413	0.5643	0.1308
Drift	Method 1 PPD (iter_2)	0.2795	0.9616	0.9366	0.6206	0.2090
Drift	Finite mixtures (iter_1)	0.9909	0.6468	0.6593	0.8189	0.2450
Drift	Finite mixtures (iter_2)	0.9704	0.6885	0.6988	0.8294	0.2595
Drift	HMM	0.3450	0.9454	0.9234	0.6452	0.2185

Note. The best performing method in terms of MCC are highlighted in bold.

allows clustering sites based on similarity or time series distance (Aghabozorgi et al., 2015). Figure 8 shows the clusters formed performing DTW (two clusters in blue and green) and dendograms with the clustering structure among sites. For the rest of this work and for the sake of simplicity we excluded three sites with very different behavior and use the seven sites in cluster one. This ensures good spatial association between sites and will help borrow strength from the neighbors when detecting anomalies. In practice, these three sites can be analyzed as a separate cluster or individually. Not done here for the sake of space. These sites are geospatially distinct from the rest, with 1,160,120 serving as the reference upstream site for the Herbert River, and 1,160,127/1,160,128 located along the southern Stone River catchment. Due to the period of observation analyzed, it is probable that a significant portion of the rainfall occurred in the lower catchment, leading to a distinct distribution of rainfall events throughout the time series. Therefore, it is not unexpected to observe these sites clustering separately from the remaining sites.

3. **Events identification:** Water quality and discharge events tend to produce higher levels of turbidity, nutrients and contaminants. It is critical to identify these water quality events as they help differentiating sensor anomalies (which do not occur across multiple sites at the same periods of time). We used a Multivariate Hidden Markov Model (Zucchini & MacDonald, 2009) to estimate the latent state of the catchment. This approach produces “ambient” and “event” labels for the catchment for each time points. We show in Figure 9 the separation of ambient (black) and event (blue). The vertical gray lines are the true water events produced by water quality specialists assumed the ground truth. Note that on site 1,160,117, no ground truth was produced during 2022, as no data was available. However, we show the labels for the imputed data from the MHMM.

This method was also effective in a site like 1,160,119 that is affected tidal influence. Overall, we found that the MHMM is effective in the identification of events. The overall accuracy was 87.89%. In total, 80% of the data points pertaining to events where correctly classified (sensitivity) and the specificity was 89.6% with a $MCC = 62.62$. The Brier score, as defined by Brier (1950), was calculated as $\frac{1}{N} \sum_{i=1}^N (f_i - o_i)^2$, where f_i represents the predicted and o_i is the observed outcome. In this MHMM analysis, the Brier score was 0.121.

6.2. Analysis of Total Suspended Solids (TSS)

We developed a recursive online method for fitting vector autoregressive Bayesian spatio-temporal regression models. Details of the general Bayesian recursive model can be found in Särkkä (2013).

We denote the total suspended solids as $TSSeq$, which is modeled as a function of the water level (*level*), a binary variable indicating event versus ambient conditions obtained in previous section (*event*) and precipitation data (*prec*). This can be represented more formally using mathematical notation as:

$$TSSeq = \beta_0 + \beta_1 \cdot level + \beta_2 \cdot event + \beta_3 \cdot prec \quad (10)$$

where β_0 is an intercept and β_1, β_2 , and β_3 are the coefficients associated with each predictor variable. Note that Equation 10 corresponds to $X\beta$ in Equation 2.

The data has a temporal resolution of one observation every 30 min. For ease of processing, we organized the data into monthly batches. This means that we took 1 month of data, fitted the model, estimated the parameters, and identified anomalies. In each batch of data, we derive posterior predictive distributions for the spatial and temporal parameters (σ_{td} , α_{td} , and ϕ_s), along with the regression coefficients (β). The posterior distributions obtained from each batch of data then served as priors for subsequent batches, allowing for the recursive updating of the model's parameters.

Batches with anomalous behavior did not inform the estimation of parameters. Specifically, if the i th batch contained anomalous data, the prior distributions for the subsequent batch were formed from the last non-anomalous batch. This approach ensures that the estimated parameters are less influenced by anomalous data and reduces the bias that may arise from such data.

For more effective online updating and to account for the dynamic change in the environment, models can be fitted to smaller chunks of data, such as hourly, daily, and weekly intervals. The computational time for fitting these models for all batches of data was smaller than the frequency of the data (<30 min).

We employed the PPD approach, which demonstrated suitable performance in our simulation study and does not require fitting additional methods such as mixtures and HMMs. Although we only utilized one iteration to avoid recomputing the models, we acknowledge that using two iterations may lead to further improvements in performance. However, we maintain that our chosen approach provides a suitable level of accuracy and practicality for the purposes of this study.

For each observation, we computed the posterior predictive distribution, and identified any observation falling outside the 95% highest density interval (HDI) as an anomaly. The identified anomalies were highlighted in red in Figure 10, with the (true) anomalous sections in the data represented with gray vertical lines. As previously mentioned, the HDI is a credible interval used to measure the uncertainty of the posterior distribution and the likelihood of observing new data. The 95% HDI, in particular, represents the range of values that contain 95% of the probability mass for the posterior distribution.

Our model demonstrates suitable performance in detecting drift (in site 1,160,116), sudden shifts, and periods of high variability (in 1,160,119). Additionally, it accurately identified the constant offset present in 1,160,122. Overall, using this approach, we were able to identify and highlight the anomalous observations with high accuracy.

Table 3 presents the performance measures of a spatio-temporal anomaly detection method for the seven types of anomalies in TSS data. The performance measures include sensitivity (*se*), specificity (*sp*), accuracy (*acc*), adjusted accuracy (*acc_{adj}*), Matthews correlation coefficient (*MCC*), and Brier score.

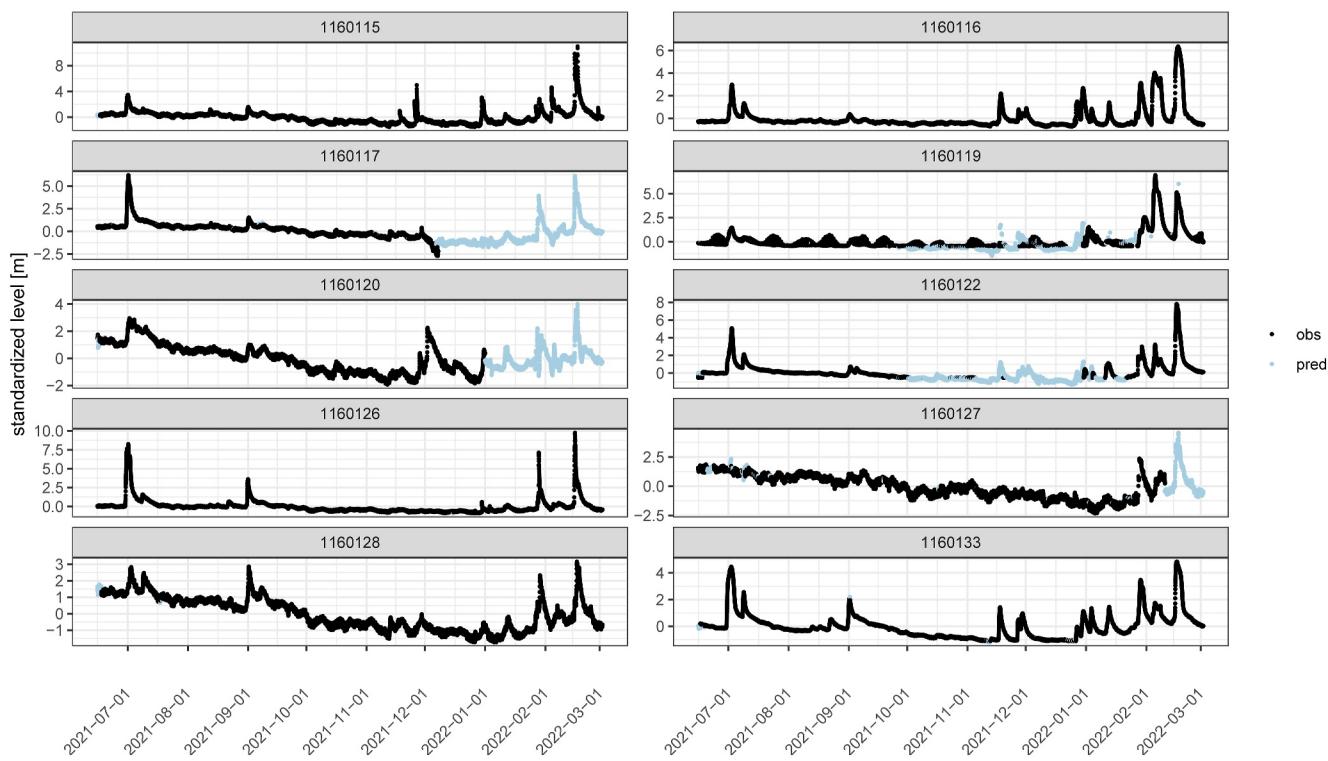


Figure 7. Time series of standardized water level data. Standardization was performed by subtracting the mean and dividing by the standard deviation for each site. Imputed values are represented in light blue.

The results show that the PPD method performs well in detecting most types of anomalies. Across all anomaly types, 74% of anomalous observations (in the high-frequency data) were correctly identified and the accuracy and adjusted accuracy were 83.6% and 79.2%. Note that the overall $sp = 84\%$ (not reported for individual types of anomaly as they are all the same).

Analyzing by type of anomaly, the highest sensitivity is observed for the constant offset ($se = 90\%$). The highest overall accuracy and adjusted accuracy are observed for the constant offset type of anomaly, while the highest MCC is observed for the sensor drift. The Brier score, which measures the accuracy of the method's probability estimates, is low for all types of anomalies, indicating that the PPD method provides reliable probability estimates. The results in the table indicate that the PPD method is effective in detecting various types of anomalies in TSS data, making it a valuable tool for online monitoring of water quality in spatio-temporal stream networks.

We conducted similar analyses on the same catchment using TSS data obtained from low-cost sensors, which posed greater challenges for the identification of anomalies. The results of these analyses are presented in Appendix B.

7. Discussion and Conclusions

Water quality data from in-situ sensors can be challenging to analyze due to the presence of technical anomalies as well as water quality events and spatio-temporal dependencies. Here, we have introduced and successfully demonstrated a general suite of unsupervised statistical methods to detect technical anomalies in data arrays from in-situ sensors used for water-quality monitoring in stream networks.

Our results show that these tools are suitable for detecting challenging technical anomalies in high-frequency data produced by sensor technology. The four approaches discussed were more effective in the identification of a range of anomalies, compared to the ARIMA, which can only describe temporal correlation in the data at individual sites. These spatially aware models capture events happening simultaneously across multiple locations, allow more effective data imputation of time series and help identify spatial anomalies. Since these models were set within a

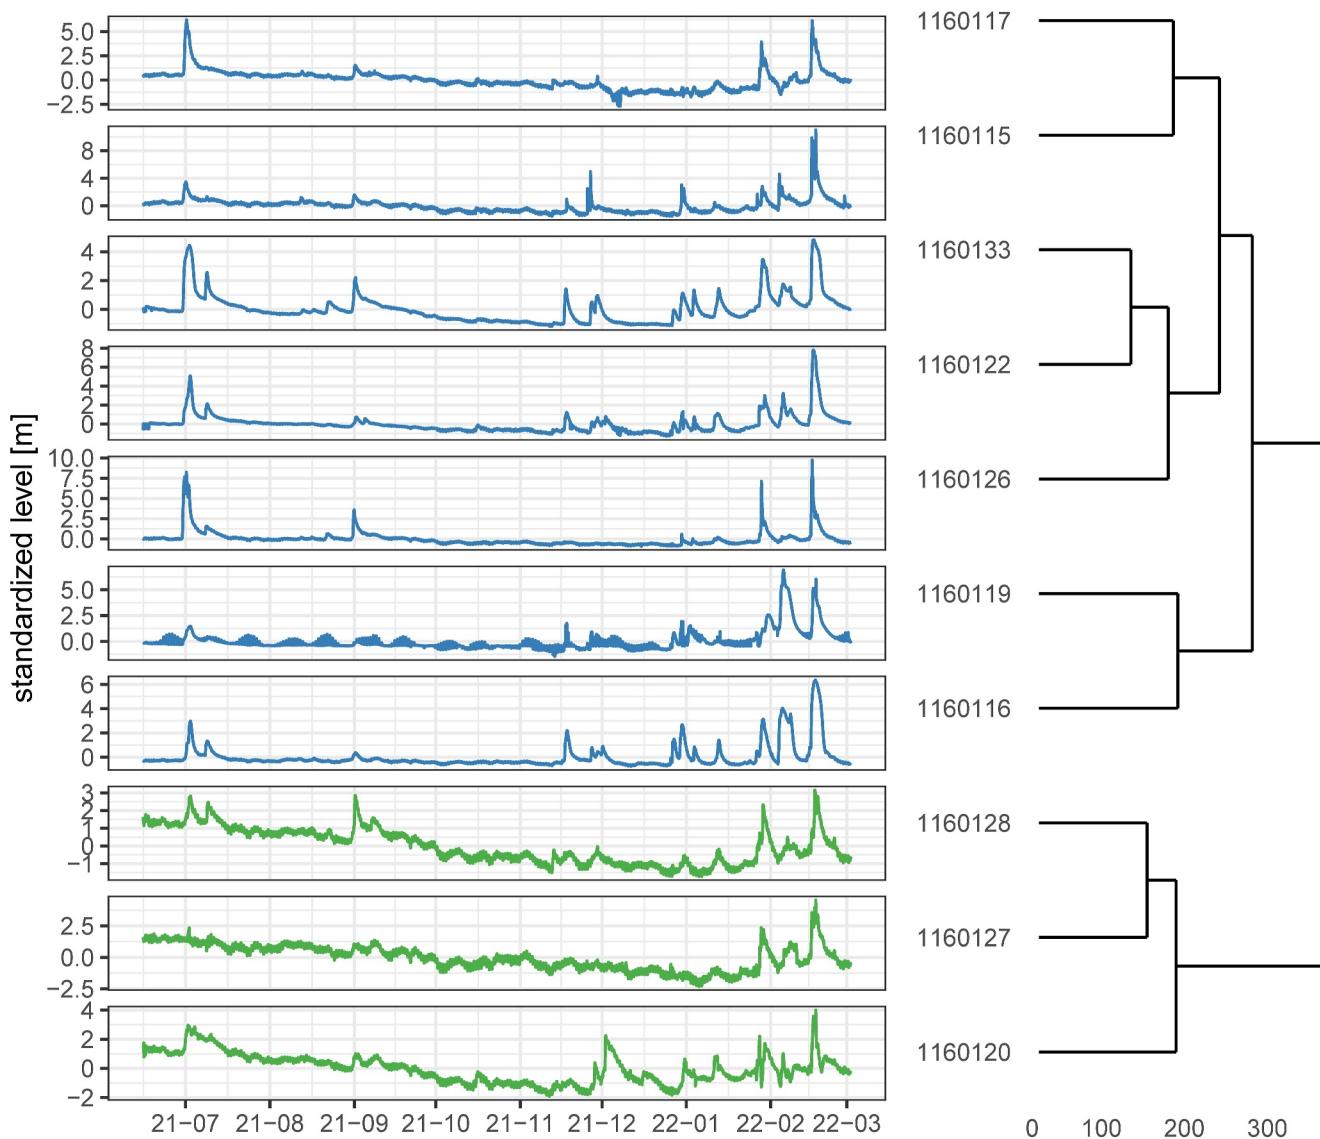


Figure 8. Time series of standardized water level data (left) and dendrograms with the clustering structure among sites (right).

Bayesian framework, we can obtain probabilistic estimates, anomalous scores and quantify the uncertainty in the data.

The findings of this study were reinforced through a comprehensive simulation study that incorporated various data sets and types of anomalies. Nonetheless, future investigations could take into account scenarios where the model assumptions are violated, such as deviations in the temporal structure and non-linear associations between the response variable and covariates. These circumstances can help to provide a more complete understanding of the model's behavior and its robustness under different scenarios of data generation. Therefore, additional research in these areas could provide valuable insights and enhance the overall applicability of the proposed methodology.

We chose to implement a spatio-temporal stream-network model (Santos-Fernandez et al., 2022) because the unique characteristics of streams (e.g., branching network topology, flow connectivity and directionality), together with high-frequency sampling tended to produce data with spatial and temporal dependencies. However, analyzing the empirical residuals from the model fit via finite mixtures and HMMs offers a flexible set of solutions to distinguish anomalous from normal observations, regardless of the initial approach used. For example, the spatio-temporal model could be replaced with other approaches that produce predictions, such as Gaussian

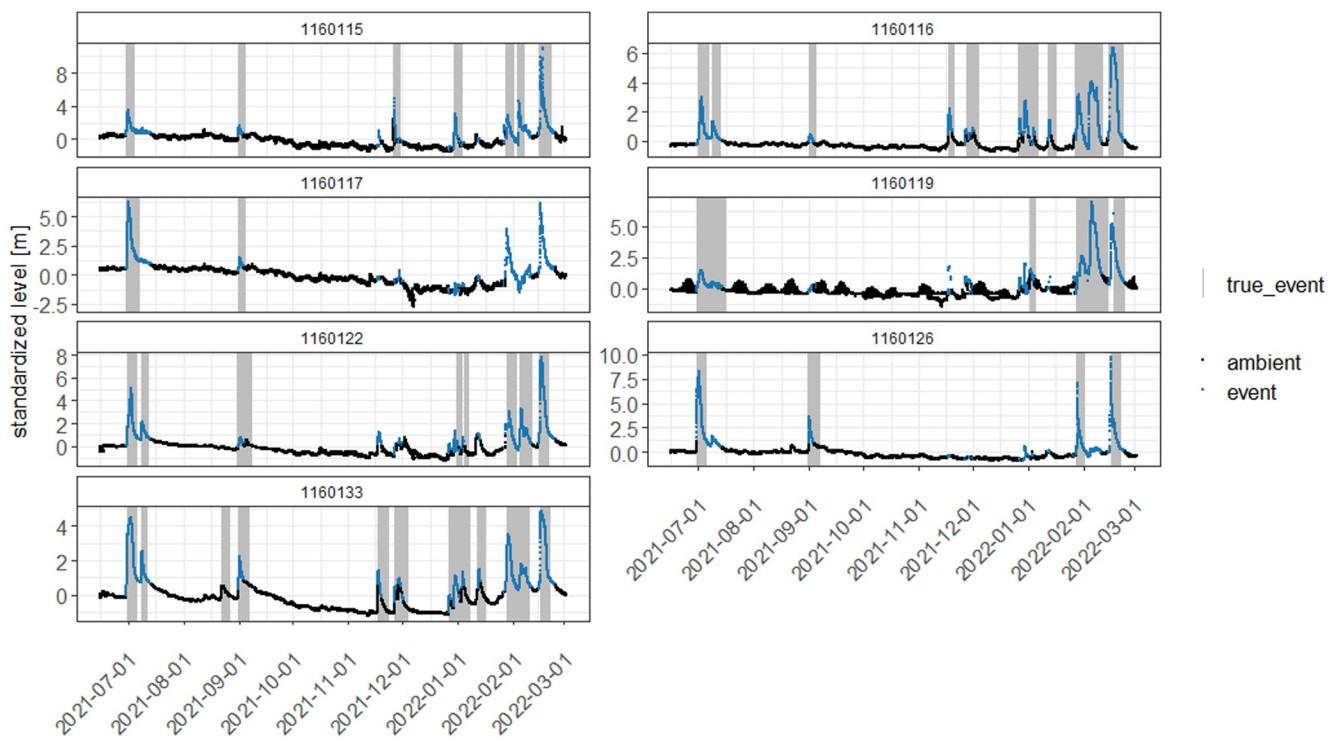


Figure 9. Ambient and events identified in the time series of standardized water level data using a Multivariate Hidden Markov Model (MHMM).

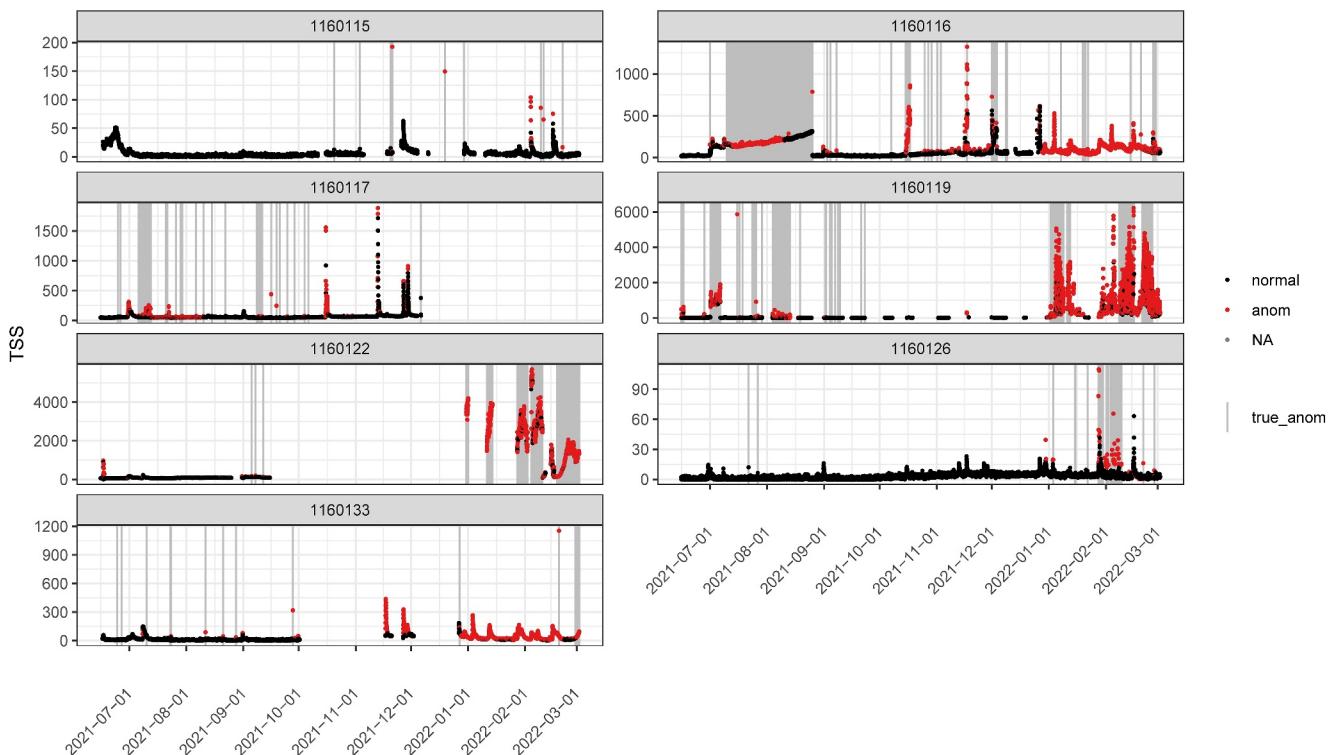


Figure 10. Anomalies identified in the turbidity (TSS) data. The gray vertical lines and shaded areas delineate labeled anomalies (ground truth). Points in red represent the predicted anomalies from the models.

Table 3

Performance Measures in the Detection of Anomalies in Turbidity Data by Type of Anomaly in the Labeled Data Set and Overall

Anomaly type	Method	se	sp	acc	acc_adj	MCC	Brier
Small spike	ppd	0.4530		0.8453	0.6496	0.0381	0.1547
Large spike	ppd	0.6774		0.8459	0.7618	0.0595	0.1541
Cluster of spikes	ppd	0.8276		0.8461	0.8369	0.0428	0.1539
Sudden shift	ppd	0.5940		0.8443	0.7201	0.1024	0.1557
High variability	ppd	0.7965		0.8451	0.8213	0.2474	0.1549
Constant offset	ppd	0.9017		0.8469	0.8739	0.2361	0.1531
Sensor drift	ppd	0.7074		0.8389	0.7768	0.3186	0.1611
All	ppd	0.7397	0.8461	0.8363	0.7929	0.4187	0.1637

additive models (GAMS) (Wood, 2017), time series models (Hyndman & Athanasopoulos, 2018), or machine learning algorithms (Chandola et al., 2009).

Physical models are constructed based on well-established physical principles that govern processes within river networks. These models are capable of simulating water quality parameters across diverse environmental conditions, offering a dynamic representation of system behavior. By comparing observed data with model predictions, physical models can detect discrepancies, potentially identifying anomalies in the data set. An extensive body of literature has introduced various physical methods for identifying anomalous data and change points in time series analysis (Housh & Ohar, 2017; Zégre et al., 2010). Since these models can capture intricate interactions and dependencies within the river system, they can complement the capabilities of statistical methods. However, these physical models can be very difficult to develop, require many parameters (which are often unidentifiable from available data) and the models can very computationally expensive to form predictions. Future research endeavors could explore the development of efficient hybrid approaches, merging physical modeling with Bayesian statistical methods, to achieve enhanced performance and gain deeper insights into water quality dynamics.

Water quality parameters often are non-normally distributed and can be affected by censorship or zero-inflation, posing a challenge for statistical analysis. To address this issue, current practice involves using log transformations to normalize the data. Alternatively, our modeling framework can be adapted to follow the Generalized Linear Model (GLM) philosophy, which provides a flexible approach to analyzing non-normal data and can account for the various distributional characteristics of water quality parameters.

From a practical point of view, using methods such as mixtures and HMMs involves challenging issues such as the selection of the number of components. This, however, could be approached using common model-selection measures such as Akaike information criterion (AIC), Bayesian information criterion (BIC) and widely applicable information criterion (WAIC) to compare competing models. Other natural extensions of these methods include multivariate approaches to explain the complex relationships between multiple correlated parameters such as nitrate, turbidity, and conductivity (Kermorvant et al., 2021). The tools can also be easily extended to build statistical process control approaches (e.g., multivariate cumulative sum (MCUSUM) or multivariate exponentially weighted moving average (MEWMA) control charts (Montgomery, 2020; Santos-Fernández, 2013). Similarly, the PPD Bayesian method can be extended by constructing multivariate credible intervals (Besag et al., 1995). We are currently exploring multiple variations of these models.

The motivation for developing these methods was primarily for automating the Quality Assurance and Control (QA/QC) processes for high-frequency in-situ data, but they could also be applied retrospectively. For example, anomalies could be identified and removed during the modeling phase to obtain less biased estimates of the parameters of interest (e.g., regression coefficients). This is important in cases where data are released for public use, especially where the presence of anomalies due to sensor issues could reduce confidence in the quality of the data.

This work opens up new avenues of research, especially around the detection of anomalies in space-time data in streams. The outcomes of this research are paramount for the delivery and presentation of environmental near

real-time data to the public. As more high-frequency monitoring data are made available to the general community, timely detection of the anomalies/errors will be imperative for gaining and maintaining trust and transparency in the measured data. This is particularly important when data are used to identify sources of water quality pollutants and demonstrate to landholders the benefits associated with changes in land management practices. Future research should explore the integration of the statistical models here described with other statistical quality control tools for high-frequency data produced by arrays of sensors.

Appendix A: Other Results

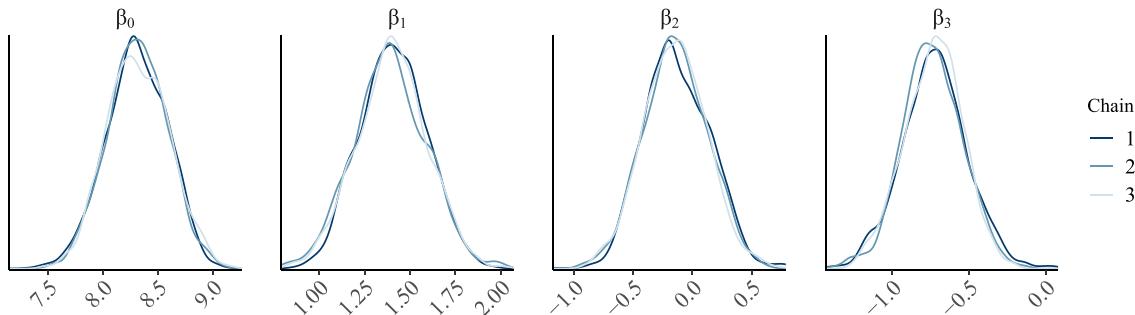


Figure A1. Posterior distribution of the regression coefficients. True values $\beta_0 = 10$, $\beta_1 = 1$, $\beta_2 = 0$ and $\beta_3 = -1$.

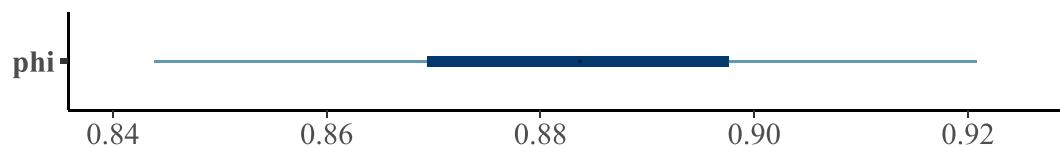


Figure A2. Posterior boxplot of the autoregression parameter. The true value is $\phi = 0.80$.

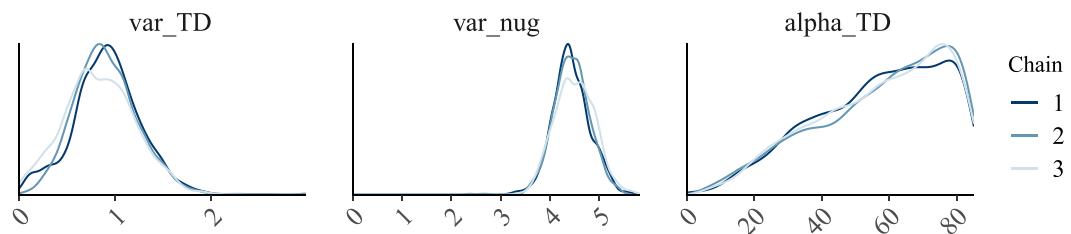


Figure A3. Posterior distribution of the spatial parameters. True values: $\sigma_{TD}^2 = 3$, $\alpha_{TD} = 10$, $\sigma_{NUG}^2 = 0.1$.

Appendix B: Extended Analysis Using Low-Cost Sensors

In this section we extend the analyses from Section 6 using data from low-cost sensors. We installed five low-cost BlueSpot water-quality sensors (model TLS02 021; Figure B1) produced by IntelliDesign in the branching network of the Herbert River. These Internet of Things (IoT) devices measure stream water level (m), specific conductivity ($\mu\text{S}/\text{cm}$) and turbidity (NTU). Measurements were recorded every 15 min and uploaded via the 4G mobile network to a database. Each BlueSpot was installed at the edge of the riverbank in a location best representing that part of the waterway. The BlueSpot telemetry unit was attached to a star picket anchored to the ground using cables. The probe itself was suspended in the water column approximately 30 cm above the riverbed or riverbank depending on where it was installed. We used time-series data for turbidity and water level collected between 16-Feb-2021 and 04-Mar-2021 in this study.

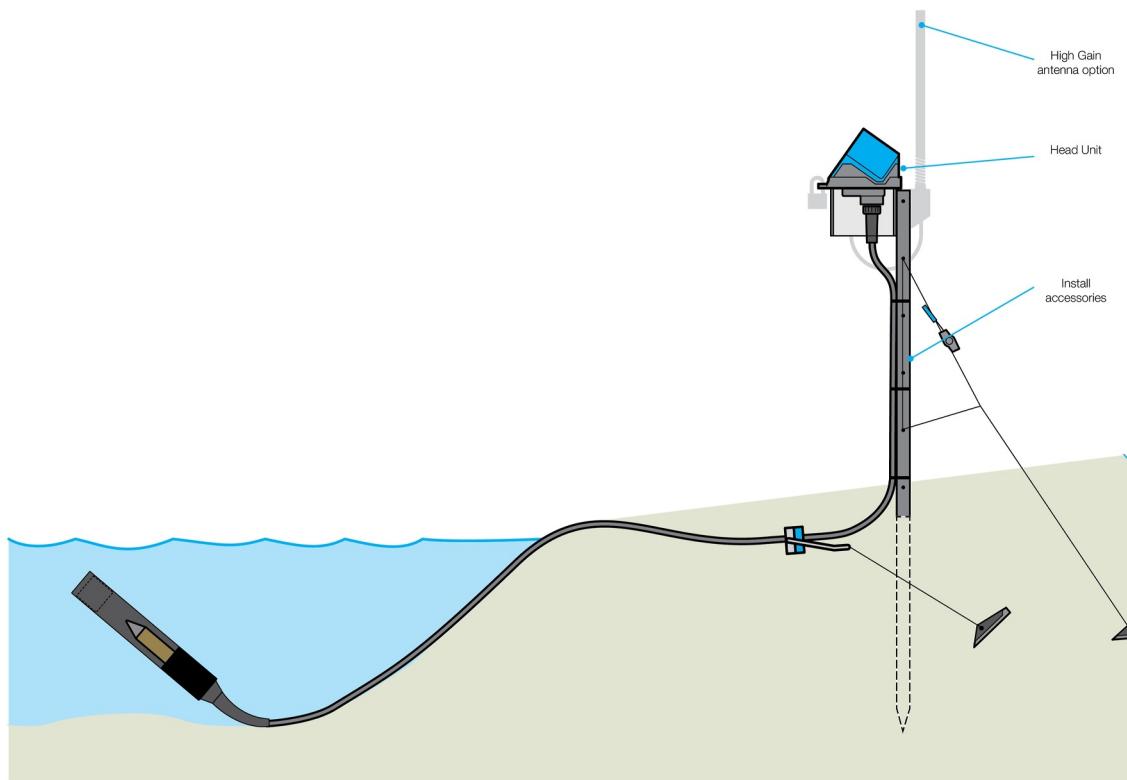


Figure B1. Diagram of the BlueSpots used in the Herbert River.

B1. Results

We used the best-performing methods from the simulation study (PPD and HMM) to identify anomalies in water level and turbidity measurements. We first fit a space-time model introduced in Equation 1 to identify anomalies in water level using a temporal and two spatial covariates:

1. Precipitation (in mm) obtained from a monitoring station in the Herbert River. <https://water-monitoring.information.qld.gov.au/>;
2. Proportion of stream with land uses where fertilizer is likely to be used; and
3. Proportion of natural forest cover in the catchment.

The spatial covariates were sourced from the Australian National Environmental Stream Attributes data set (Stein et al., 2012).

Next we fit a spatio-temporal stream-network model for (log) turbidity using the same three covariates described above, plus the (log) water level. All the methods were then used to predict anomalies.

For illustration, let us consider the turbidity from Bluespot sensor 1, where some anomalous behavior occurred (Figure B2). In green, we represent the *data points (with vertical lines)* and *periods (with rectangles)* where true technical (sensor) anomalies are present in the labeled data set. In red we represented the predicted anomalies using the PPD method. Notice the two sudden spikes on the 20-Feb-2021 at 19:45 and 21-Feb-2021 at 09:15. There is also a period of high variability (28-Feb-2021 to 03-Mar-2021) followed by an ongoing drift (3-Mar-2021 16:00) (Figure B2).

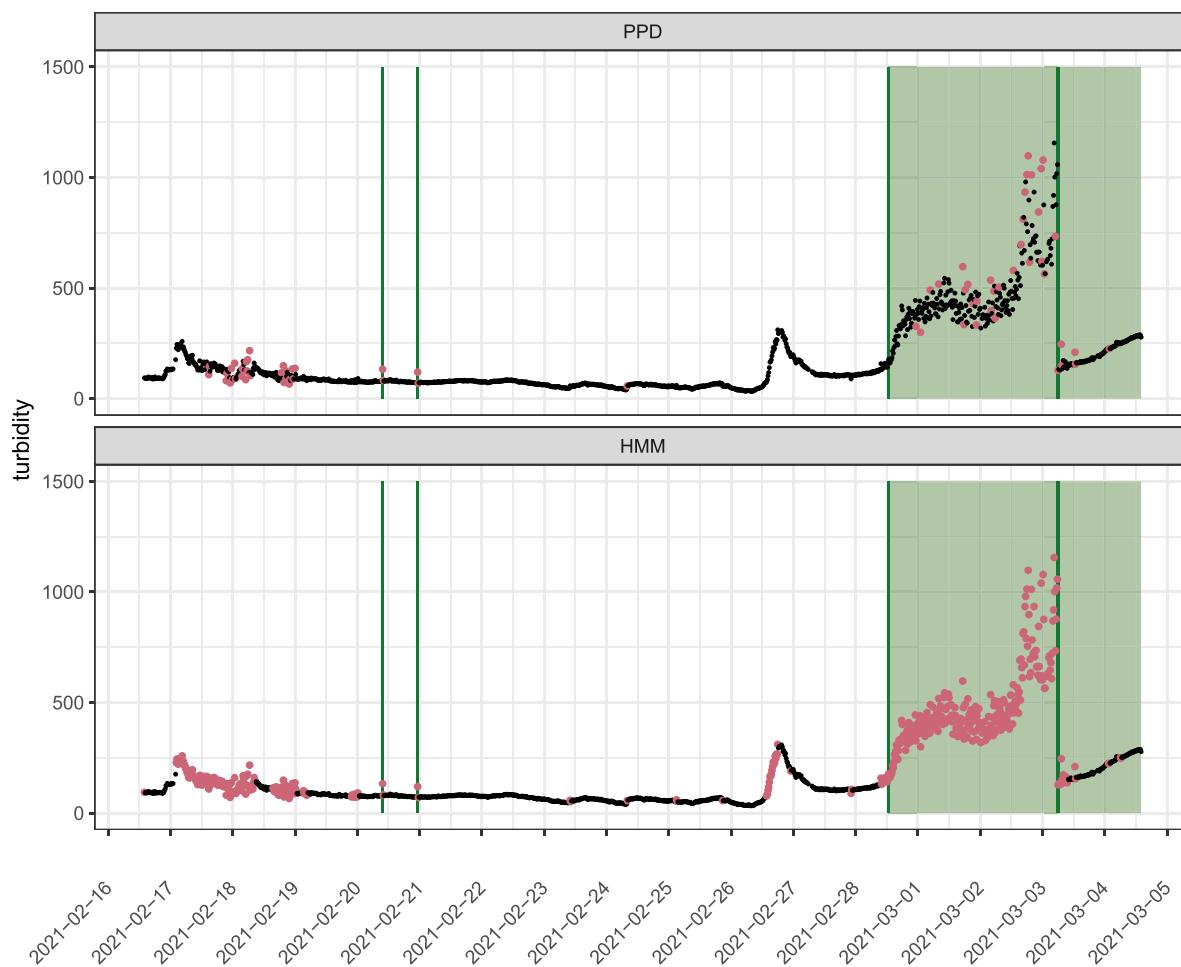


Figure B2. Time series of turbidity data and anomalies from Bluespot sensor 1 using (a) PPD and (b) HMM. The green vertical lines and shaded areas delineate labeled anomalies. Points in red represent the predicted anomalies from the models.

The two approaches successfully identified the four anomalies in the data, including the start of the drift event. The PPD model identified the high variability event, in addition to the two sudden spikes and the beginning of the drift (Figure B2a).

However, the HMM method was more sensitive to changes in the process than the other two models (Figure B2c). Interestingly, the HMM models also identified a true water quality event around 26-Feb-2021, which was observed across the other Bluespot sensors.

Data Availability Statement

Software: The spatio-temporal stream network models were fit using the R package SSNbayes (Santos-Fernandez et al., 2023). R code used to simulate spatio-temporal stream network data and to incorporate anomalies can be found at (Santos-Fernandez, 2023). Data: The Herbert river data set from the case study has been deposited in Zenodo (Santos-Fernandez, 2024b). It can also be accessed through the R package herbert (Santos-Fernandez, 2024a). See <https://github.com/EdgarSantos-Fernandez/herbert>.

Acknowledgments

This research was supported by the Australian Research Council (ARC) Linkage Project LP180101151: “Revolutionising water-quality monitoring in the information age.” We thank the Water Quality and Investigations group from the Queensland Department of Environment and Science for the installation of the water quality sensors and for producing the labeled data sets used in the case study, specially Matthew Sinclair and Ben Ferguson. We also thank Puwasala Gamakumara for developing some of the R scripts used to access the in-situ data from low-cost sensors. The authors are grateful to the six reviewers and Associate Editor for their insightful comments. Open access publishing facilitated by Queensland University of Technology, as part of the Wiley - Queensland University of Technology agreement via the Council of Australian University Librarians.

References

Aghabozorgi, S., Shirkhorshidi, A. S., & Wah, T. Y. (2015). Time-series clustering—a decade review. *Information Systems*, 53, 16–38. <https://doi.org/10.1016/j.is.2015.04.007>

Altenburger, R., Brack, W., Burgess, R. M., Busch, W., Escher, B. I., Focks, A., et al. (2019). Future water quality monitoring: Improving the balance between exposure and toxicity assessments of real-world pollutant mixtures. *Environmental Sciences Europe*, 31(1), 1–17. <https://doi.org/10.1186/s12302-019-0193-1>

Barnett, V., & Lewis, T. (1984). *Outliers in statistical data*. Wiley Series in Probability and Mathematical Statistics. Applied Probability and Statistics.

Besag, J., Green, P., Higdon, D., & Mengerson, K. (1995). Bayesian computation and stochastic systems. *Statistical Science*, 10(1), 3–41. <https://doi.org/10.1214/ss/1177010123>

Bosman, H. H., Iacca, G., Tejada, A., Wörtche, H. J., & Liotta, A. (2017). Spatial anomaly detection in sensor networks using neighborhood information. *Information Fusion*, 33, 41–56. <https://doi.org/10.1016/j.inffus.2016.04.007>

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1), 1–3. [https://doi.org/10.1175/1520-0493\(1950\)078<0001:yofei>2.0.co;2](https://doi.org/10.1175/1520-0493(1950)078<0001:yofei>2.0.co;2)

Brodersen, K. H., Ong, C. S., Stephan, K. E., & Buhmann, J. M. (2010). The balanced accuracy and its posterior distribution. In *2010 20th International Conference on Pattern Recognition* (pp. 3121–3124).

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58. <https://doi.org/10.1145/1541880.1541882>

Congdon, P. D. (2019). *Bayesian hierarchical models: With applications using R*. CRC Press.

Costa, A. C., & Soares, A. (2009). Homogenization of climate data: Review and new perspectives using geostatistics. *Mathematical Geosciences*, 41(3), 291–305. <https://doi.org/10.1007/s11004-008-9203-3>

Cressie, N., Frey, J., Harch, B., & Smith, M. (2006). Spatial prediction on a river network. *Journal of Agricultural, Biological, and Environmental Statistics*, 11(2), 127–150. <https://doi.org/10.1198/108571106x110649>

Djenouri, Y., Belhadi, A., Lin, J. C.-W., & Cano, A. (2019). Adapted k-nearest neighbors for detecting anomalies on spatio-temporal traffic flow. *IEEE Access*, 7, 10015–10027. <https://doi.org/10.1109/access.2019.2891933>

Fotheringham, A. S., & Rogerson, P. A. (2008). *The SAGE handbook of spatial analysis*. Sage.

Garreta, V., Monestiez, P., & Ver Hoef, J. M. (2010). Spatial modelling and prediction on river networks: Up model, down model or hybrid? *Environmetrics*, 21(5), 439–456. <https://doi.org/10.1002/env.995>

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (1995). *Bayesian data analysis*. Chapman and Hall/CRC.

Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.

Great Barrier Reef Foundation. (2022). Great barrier Reef foundation. Retrieved from <https://www.barrierreef.org/what-we-do/reef-trust-partnership/water-quality-improvement/innovation-and-systems-change>

Hill, D. J., & Minsker, B. S. (2010). Anomaly detection in streaming environmental sensor data: A data-driven modeling approach. *Environmental Modelling and Software*, 25(9), 1014–1022. <https://doi.org/10.1016/j.envsoft.2009.08.010>

Housh, M., & Ohar, Z. (2017). Integrating physically based simulators with event detection systems: Multi-site detection approach. *Water Research*, 110, 180–191. <https://doi.org/10.1016/j.watres.2016.12.003>

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and practice*. OTexts.

Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, 26(3), 1–22. <https://doi.org/10.18637/jss.v027.i03>

Isaak, D. J., Wenger, S. J., Peterson, E. E., Ver Hoef, J. M., Nagel, D. E., Luce, C. H., et al. (2017). The NorWeST summer stream temperature model and scenarios for the western US: A crowd-sourced database and new geospatial tools foster a user community and predict broad climate warming of rivers and streams. *Water Resources Research*, 53(11), 9181–9205. <https://doi.org/10.1002/2017WR020969>

Jackson, F. L., Fryer, R. J., Hannah, D. M., Millar, C. P., & Malcolm, I. A. (2018). A spatio-temporal statistical model of maximum daily river temperatures to inform the management of Scotland’s Atlantic salmon rivers under climate change. *Science of the Total Environment*, 612, 1543–1558. <https://doi.org/10.1016/j.scitotenv.2017.09.010>

Jones, A. S., Jones, T. L., & Horsburgh, J. S. (2022). Toward automating post processing of aquatic sensor data. *Environmental Modelling & Software*, 151, 105364. <https://doi.org/10.1016/j.envsoft.2022.105364>

Junger, W., & de Leon, A. P. (2018). mtsdi: Multivariate time series data imputation. Computer software manual. [https://CRAN.R-project.org/package=mtsdi\(Rpackageversion0.3.5](https://CRAN.R-project.org/package=mtsdi(Rpackageversion0.3.5)

Kandanaarachchi, S., & Hyndman, R. J. (2021). *Leave-one-out kernel density estimates for outlier detection (Tech. Rep.)*. Monash University, Department of Econometrics and Business Statistics.

Kermorvant, C., Liquet, B., Litt, G., Mengerson, K., Peterson, E., Hyndman, R., et al. (2021). Understanding links between water-quality variables and nitrate concentration in freshwater streams using high-frequency sensor data. *arXiv preprint arXiv:2106.01719*.

Leigh, C., Alsabai, O., Hyndman, R. J., Kandanaarachchi, S., King, O. C., McGree, J. M., et al. (2019). A framework for automated anomaly detection in high frequency water-quality data from in situ sensors. *The Science of the Total Environment*, 664, 885–898. <https://doi.org/10.1016/j.scitotenv.2019.02.085>

Leigh, C., Kandanaarachchi, S., McGree, J., Hyndman, R. J., Kandanaarachchi, S., King, O. C., & Peterson, E. E. (2019). Predicting sediment and nutrient concentrations from high-frequency water-quality data. *PLoS One*, 14(8), e0215503. <https://doi.org/10.1371/journal.pone.0215503>

Li, J., Pedrycz, W., & Jamal, I. (2017). Multivariate time series anomaly detection: A framework of hidden Markov models. *Applied Soft Computing*, 60, 229–240. <https://doi.org/10.1016/j.asoc.2017.06.035>

Matthews, B. W. (1975). Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2), 442–451. [https://doi.org/10.1016/0005-2795\(75\)90109-9](https://doi.org/10.1016/0005-2795(75)90109-9)

Money, E., Carter, G. P., & Serre, M. L. (2009a). Using river distances in the space/time estimation of dissolved oxygen along two impaired river networks in New Jersey. *Water Research*, 43(7), 1948–1958. <https://doi.org/10.1016/j.watres.2009.01.034>

Money, E. S., Carter, G. P., & Serre, M. L. (2009b). Modern space/time geostatistics using river distances: Data integration of turbidity and E. coli measurements to assess fecal contamination along the Raritan River in New Jersey. *Environmental Science and Technology*, 43(10), 3736–3742. <https://doi.org/10.1021/es803236j>

Montgomery, D. C. (2020). *Introduction to statistical quality control*. John Wiley and Sons.

Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. MIT press.

O’Donnell, D., Rushworth, A., Bowman, A. W., Scott, E. M., & Hallard, M. (2014). Flexible regression models over river networks. *Journal of the Royal Statistical Society - Series C: Applied Statistics*, 63(1), 47–63. <https://doi.org/10.1111/rssc.12024>

Peterson, E. E., Merton, A. A., Theobald, D. M., & Urquhart, N. S. (2006). Patterns of spatial autocorrelation in stream water chemistry. *Environmental Monitoring and Assessment*, 121(1), 571–596. <https://doi.org/10.1007/s10661-005-9156-7>

Peterson, E. E., Ver Hoef, J. M., Isaak, D. J., Falke, J. A., Fortin, M.-J., Jordan, C. E., et al. (2013). Modelling dendritic ecological networks in space: An integrated network perspective. *Ecology Letters*, 16(5), 707–719. <https://doi.org/10.1111/ele.12084>

Pinheiro, J., Bates, D., DebRoy, S., & Sarkar, D., & R Core Team. (2020). nlme: Linear and nonlinear mixed effects models. Computer software manual. <https://CRAN.R-project.org/package=nlme>

Queensland Government. (2022a). Modelling reductions of pollutant loads due to improved management practices in the Great Barrier Reef catchments. Retrieved from https://www.reefplan.qld.gov.au/__data/assets/pdf_file/0019/46180/wet-tropics-catchment-modelling-report.pdf

Queensland Government. (2022b). Small business innovation research (SBIR) pilot: Reef water quality monitoring challenge overview. Retrieved from <https://openinnovation.com.au/static/downloads/sbir/SBIR%20Pilot%20-%20Reef%20Water%20Quality%20Monitoring%20-%20Challenge%20Overview.pdf>

R Core Team. (2022). R: A language and environment for statistical computing. Computer software manual. <https://www.R-project.org/>

Resch, G., Koch, R., Marty, C., Chimani, B., Begert, M., Buchmann, M., et al. (2023). A quantile-based approach to improve homogenization of snow depth time series. *International Journal of Climatology*, 43(1), 157–173. <https://doi.org/10.1002/joc.7742>

Rodriguez-Perez, J., Leigh, C., Liquet, B., Kermorvant, C., Peterson, E., Sous, D., & Mengersen, K. (2020). Detecting technical anomalies in high-frequency water-quality data using artificial neural networks. *Environmental Science and Technology*, 54(21), 13719–13730. <https://doi.org/10.1021/acs.est.0c04069>

Santos-Fernández, E. (2013). *Multivariate statistical quality control using R* (Vol. 14). Springer. Retrieved from <http://www.springer.com/statistics/computational+statistics/book/978-1-4614-5452-6>

Santos-Fernandez, E. (2023). Anomaly detection in spatio-temporal sensor data (0.0.1) [Dataset]. Zenodo. <https://doi.org/10.5281/zenodo.8146087>

Santos-Fernandez, E. (2024a). herbert: Herbert river water quality dataset. Computer software manual. <https://github.com/EdgarSantos-Fernandez/herbert>

Santos-Fernandez, E. (2024b). Herbert river water quality dataset. Herbert: Version: 0.0.1 [Dataset]. Zenodo. <https://doi.org/10.5281/zenodo.13883326>

Santos-Fernandez, E., Hoef, J. M. V., McGree, J., Isaak, D. J., Mengersen, K., & Peterson, E. E. (2023). SSNbayes: An R package for Bayesian spatio-temporal modelling on stream networks [Software]. *The R Journal*, 15(3), 26–58. <https://doi.org/10.32614/RJ-2023-061>

Santos-Fernandez, E., Ver Hoef, J. M., Peterson, E. E., McGree, J., Isaak, D. J., & Mengersen, K. (2022). Bayesian spatio-temporal models for stream networks. *Computational Statistics and Data Analysis*, 170, 107446. <https://doi.org/10.1016/j.csda.2022.107446>

Särkkä, S. (2013). *Bayesian filtering and smoothing* (Vol. 3). Cambridge University Press.

Senin, P. (2008). *Dynamic time warping algorithm review* (Vol. 855(1–23), p. 40). Information and Computer Science Department University of Hawaii at Manoa Honolulu.

Shi, Y., Deng, M., Yang, X., & Gong, J. (2018). Detecting anomalies in spatio-temporal flow data by constructing dynamic neighbourhoods. *Computers, Environment and Urban Systems*, 67, 80–96. <https://doi.org/10.1016/j.compenvurbsys.2017.08.010>

Stan Development Team. (2018). RStan: The R interface to Stan. Retrieved from <http://mc-stan.org/Rpackageversion2.18.2>

Steel, E. A., Sowder, C., & Peterson, E. E. (2016). Spatial and temporal variation of water temperature regimes on the snoqualmie river network. *Journal of the American Water Resources Association*, 52(3), 769–787. <https://doi.org/10.1111/1752-1688.12423>

Stein, J., Hutchison, M., & Stein, J. (2012). *National environmental stream attributes v1. 1.5. 5* (1). Australian National University.

Tang, J., & Zimmerman, D. (2020). Space-time covariance models on networks with an application on streams. *arXiv preprint arXiv:2009.14745*.

Tweten, M., Eckley, I. A., & Fearnhead, P. (2022). Scalable change-point and anomaly detection in cross-correlated data with an application to condition monitoring. *Annals of Applied Statistics*, 16(2), 721–743. <https://doi.org/10.1214/21-aos1508>

Ver Hoef, J. M., Peterson, E., Clifford, D., & Shah, R. (2014). SSN: An R package for spatial statistical modeling on stream networks. *Journal of Statistical Software*, 56(3), 1–45. <https://doi.org/10.18637/jss.v056.i03>

Ver Hoef, J. M., Peterson, E., & Theobald, D. (2006). Spatial statistical models that use flow and stream distance. *Environmental and Ecological Statistics*, 13(4), 449–464. <https://doi.org/10.1007/s10651-006-0022-8>

Ver Hoef, J. M., & Peterson, E. E. (2010). A moving average approach for spatial statistical models of stream networks. *Journal of the American Statistical Association*, 105(489), 6–18. <https://doi.org/10.1198/jasa.2009.ap08248>

Wang, W., Sheng, Y., Wang, J., Zeng, X., Ye, X., Huang, Y., & Zhu, M. (2017). Hast-ids: Learning hierarchical spatial-temporal features using deep neural networks to improve intrusion detection. *IEEE Access*, 6, 1792–1806. <https://doi.org/10.1109/access.2017.2780250>

Weisberg, S. (2013). *Applied linear regression*. John Wiley and Sons.

Wood, S. N. (2017). *Generalized additive models: An introduction with R*. CRC Press.

Zégre, N., Skaugset, A. E., Som, N. A., McDonnell, J. J., & Ganio, L. M. (2010). In lieu of the paired catchment approach: Hydrologic model change detection at the catchment scale. *Water Resources Research*, 46(11). <https://doi.org/10.1029/2009wr008601>

Zhang, J., Ling, Y., Fu, X., Yang, X., Xiong, G., & Zhang, R. (2020). Model of the intrusion detection system based on the integration of spatial-temporal features. *Computers and Security*, 89, 101681. <https://doi.org/10.1016/j.cose.2019.101681>

Zucchini, W., & MacDonald, I. L. (2009). *Hidden Markov models for time series: An introduction using R*. Chapman and Hall/CRC.