Multi-Sensor Monitoring of Wetland Inundation Using a Machine Learning and Data Fusion Framework

Jenna N. Abrahamson[1], Josh M. Gray[1,2], Mirela G. Tulbure[1,2], and Erin M. Schliep[3]

[1] *Center for Geospatial Analytics, North Carolina State University, Raleigh, NC, USA*
[2] *Forestry and Environmental Resources, North Carolina State University, Raleigh, NC, USA*
[3] *Department of Statistics, North Carolina State University, Raleigh, NC, USA*

Corresponding Author Email: jnabraha@ncsu.edu

This manuscript is a non–peer-reviewed preprint submitted to **EarthArXiv** and is currently under review at *Journal of Remote Sensing*.

Multi-Sensor Monitoring of Wetland Inundation Using a Machine Learning and Data Fusion Framework

Jenna N. Abrahamson[1*], Josh M. Gray[1,2], Mirela G. Tulbure[1,2], and Erin M. Schliep[3]

[1] *Center for Geospatial Analytics, North Carolina State University, Raleigh, NC, USA*
[2] *Forestry and Environmental Resources, North Carolina State University, Raleigh, NC, USA*
[3] *Department of Statistics, North Carolina State University, Raleigh, NC, USA*
[*] *Address correspondence to: jnabraha@ncsu.edu*

Abstract

Continuous, high-resolution inundation data are needed to understand how small-scale, short-term wetland flooding influences global methane emissions and carbon cycling. Small (<1,000 m^2), variably inundated wetlands are significant methane sources, yet coarse satellite products often miss their dynamics. Integrating optical and radar imagery with resolutions <30 m offers a solution over single sensor approaches, but two needs remain: (1) understanding the relative strengths and limitations of different sensors for detecting inundation in vegetated wetlands, and (2) developing methods that leverage these sensor characteristics to transform sparse, irregular classification maps with different resolutions into complete, consistent coverage. To address this, we first evaluated sensors by classifying inundation from 2017 to 2022 across a wetland area in eastern North Carolina using Sentinel-1, Sentinel-2, and PlanetScope imagery. We used tree-based machine learning models to classify pixels as *Dry Land*, *Inundated Vegetation*, or *Open Water* and compared maps of inundation frequency produced by each model. Sentinel-2 random forest achieved the highest accuracy (94.9%) followed by PlanetScope (92.6%) and Sentinel-1 (85.9%) models. Next, we introduce a fusion framework based on Fixed-Rank Kriging, a spatiotemporal statistical model, to fuse binary classification maps into daily inundation probabilities at 6 m resolution. The fused approach balanced individual sensor variability and detected short-term inundation fluctuations missed by the Landsat Dynamic Surface Water Extent, a widely used satellite-derived data product. By characterizing sensor performance and translating that information into a multi-sensor fusion approach, this work enables high-resolution, continuous monitoring of inundation dynamics vital to predicting wetland methane emissions.

1 Introduction

Hydrology strongly influences biogeochemistry, particularly in low-lying coastal plains, where drastic changes in inundated areas can occur over short time scales. Wetlands play a fundamental role in regulating interactions between the terrestrial water and carbon cycles, which are highly sensitive

to global change [1]. However, increasing water levels and more frequent flooding can disrupt this balance and transform wetlands from carbon sinks to sources as inundation reduces oxygen availability and stimulates decomposition, leading to increased production of methane (CH_4), a potent greenhouse gas [2–4]. Wetlands are the largest natural source of CH_4, yet quantifying their emissions remains challenging due to highly variable inundation patterns, ranging from permanently flooded, low-relief zones to areas of short-term precipitation-induced pluvial flooding [5]. Discrepancies in the total area of inundation arising from inconsistent and often low-resolution (> 100 m, annual) inundation datasets, contribute to one-third of the unresolved uncertainty in global methane emissions estimates [6]. As a result, our ability to estimate methane fluxes is limited by the spatial and temporal resolution of in situ, modeled, and remotely sensed inundation data [7–10].

Using remotely sensed data to generate wetland inundation extents is an ongoing area of research [8, 10–13]. The public release of satellite data archives, combined with advances in data storage and high-performance computing, has spurred the development of satellite-derived surface water and wetland products at moderate resolutions (30 m to 100 m) on decadal time scales [14, 15]. For example, the Global Surface Water Explorer [16] and Dynamic Surface Water Extent [17, 18] products generally focus on open surface water. However, [18] improved vegetated water detection in wetlands with the DSWE product using spectral mixture models. Even so, both products struggle to detect small inundated areas (<1,000 m^2) lasting less than the average 8-16 day revisit rate provided by Landsat satellites. Similarly, the Wetland Area and Dynamics for Methane Modeling (WAD2M) product has greatly improved our ability to model seasonal and interannual variations in wetland areas at a global scale by combining multiple heterogeneous datasets [11]; yet, its 25 km spatial resolution excludes variations in small inundated areas. To overcome challenges associated with spatial resolution, the influx of commercial CubeSats such as the PlanetScope constellation launched by Planet Labs has provided new catalogs of temporally dense (i.e., near daily) and high-resolution (i.e., 3 m) images suitable for monitoring dynamic surface water on Earth [8, 19–21]. However, challenges persist in using CubeSats alone to monitor inundation across long time series due to their limited spectral resolution (particularly the absence of SWIR bands) and variable radiometric performance across sensors [8, 19].

Small-scale, short-term inundation dynamics can be better monitored by leveraging multi-sensor satellite methods to enhance the temporal density of observations. The increased frequency of images from harmonized datasets such as the Harmonized Landsat Sentinel (HLS) catalog [22] has been used to map large-scale ephemeral flood events using Landsat and Sentinel-2 (A/B) satellite missions [23]. Additionally, data collections that draw on commercial and public imagery, such as Planet Lab's Fusion product, can leverage the spatial resolution of CubeSats while calibrating and filling in gaps using publicly available imagery to achieve high spatial and temporal consistency to monitor dynamic surface water [24]. However, using optical sensors alone would likely underestimate inundation during storm events that trigger persistent cloud cover. Synthetic Aperture Radar (SAR) imagery, such as from the European Space Agency's (ESA) Sentinel-1 constellation at 10 m spatial resolution and 6-12 day revisit period, provides the ability to penetrate cloud cover and has been successfully applied in previous works, often in combination with optical data [13, 25, 26]. Although L-band SAR data are preferred for forested wetlands due to their ability to penetrate dense canopies, C-

band SAR sensors have shown some success in identifying inundated vegetation in areas with sparse, non-woody vegetation canopies [13, 27, 28] and at the time of writing, Sentinel-1 provides data at a higher spatial resolution and more consistent, global coverage than available L-band sensors.

There are many techniques for fusing data from various satellite sensors to monitor environmental changes [29, 30]. A key challenge, especially for water-related applications, is the integration of active (SAR) and passive (optical) data. Fusion methods attempting to combine these two sources can be classified into three categories: pixel level, feature level, and decision level [31]. Pixel-level fusion is where raw image values are fused at the pixel level using a specified fusion rule, a full review of which can be found in [32]. Recent advances in deep learning have focused on feature-level fusion, where features such as edges and texture are extracted from images from different sensors to produce a combined feature map for classification [31, 33]. For the sake of computational efficiency and because our training data is limited, our work focuses on decision-level fusion, where individual images are pre-processed, classified separately, and the resulting classified maps are combined [31]. Examples of post-classification fusion methods for remote sensing applications include the use of expert decision rules [34], fuzzy logic [35], Bayesian methods [36], master classifiers [37], and majority voting [38]. However, many of these approaches focus on blending images without taking into consideration the strength of individual sensors, estimating output uncertainty, or predicting how classifications might vary across unobserved times or locations.

Spatiotemporal statistical models can help us track how a process changes over space and time by considering spatial and temporal autocorrelation while allowing us to integrate data from different sources. These models have been applied in remote sensing data fusion to predict tropical ocean surface winds [39–41], global CO_2 concentrations [42, 43], and aerosol optical depths [44, 45]. Many of these studies use the Spatial or Spatiotemporal Random Effects Model (SRE), which represents the process of interest across space as a linear combination of fixed basis functions with spatially correlated random effects [45, 46]. A widely used low-rank implementation (reduced dimension) of the SRE model, specifically designed for large data sets such as those in remote sensing, is known as Fixed-Rank Kriging (FRK) [46]. In addition to modeling processes with improved scalability, FRK can fill gaps in satellite observations, provide uncertainty quantification through prediction intervals, and incorporate sensor-specific measurement error.

This study addresses two critical needs: (1) understanding the relative strengths and limitations of different sensors for detecting inundation in vegetated wetlands, environments whose spectral complexity often make them difficult to map, and (2) developing methods that leverage these sensor characteristics to transform sparse, irregular binary classification maps with varying resolutions into complete, consistent coverage. To accomplish this, we employed a two-step approach. First, we used machine learning models (random forest and extreme gradient boosting) to classify inundation across five years using Sentinel-1, Sentinel-2, and PlanetScope imagery, then evaluated the accuracy of each individual sensor. Second, we selected three key areas within our study region and implemented a data fusion framework using FRK to combine heterogeneous, binary classification maps into daily predictions of inundation probability at 6 m resolution. We selected 6 m resolution to balance high-resolution mapping requirements with computational feasibility. We then evaluated both the individual sensor maps and fused predictions using in-situ water table depth (WTD), precipitation,

and lake elevation data collected from two eddy covariance flux tower sites and a nearby USGS gauge station, comparing total inundated area, timing, and spatial patterns. Finally, we compared our fusion approach results to the Landsat Dynamic Surface Water Extent product to benchmark our performance against an operational data product.

2 Study Area

Our study area is a 5,778 km^2 area of the Albemarle-Pamlico Peninsula containing the Alligator National Wildlife Refuge (Fig. 1). The peninsula is buffered by a network of barrier islands off the coast of the Atlantic and is surrounded by the Albemarle-Pamlico estuarine system, the second largest estuary in North America [47]. The climate in this area is considered humid subtropical with a mean annual temperature of 17.9°C and mean annual precipitation of 1,496 mm/yr from 2013-2023 [48]. Nearly half of the study area is less than 1 m above mean sea level, further exacerbating the vulnerability of this area to sea level rise and tropical storm surges [47]. Portions of the area have also been drained for agriculture and forestry, resulting in a complex network of artificial ditches and canals throughout the landscape [49]. Hydrology in this area is primarily influenced by precipitation, artificial drainage manipulation, and low-relief geomorphic features. These features create a distinct hummock and hollow microtopography, where the bases of trees are often above the water table. At the same time, low-lying areas remain inundated for more than 70% of the year [47, 50, 51]. Most soils are poorly drained mucks with high organic carbon content at the surface, and the vegetation consists of mixed hardwood swamp forest in addition to emergent herbaceous vegetation and mosses [51]. According to the U.S. Fish and Wildlife Services' National Wetland Inventory (NWI), nearly half the study area is classified as wetland, 75% of which have seasonal or temporarily inundated water regime classifications [52].

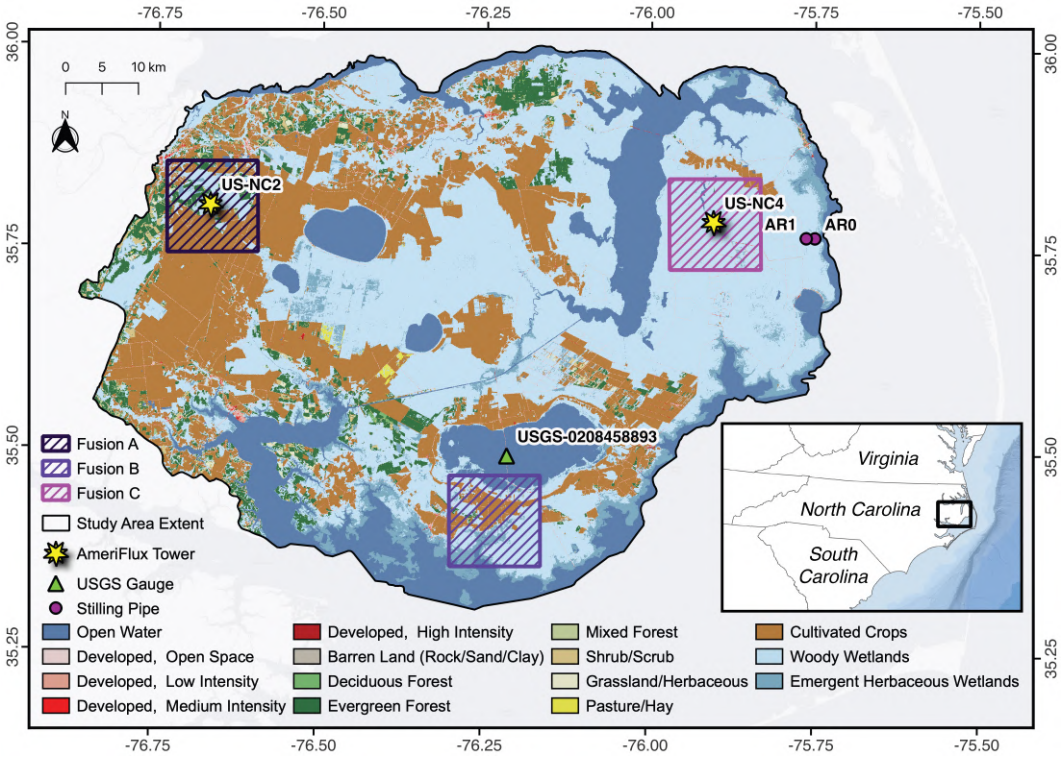

Figure 1: Map of the study area showing land cover classes from the 2019 National Land Cover Database [53], in-situ data points, and insets highlighting spatial extents where the fusion method is applied. The lower-right inset shows the study area's location within North Carolina and surrounding states.

3 Methods

3.1 Data Pre-Processing

3.1.1 Sentinel-1

We downloaded and processed 286 Sentinel-1 scenes intersecting the study area from January 2017 to December 2021 using Google Earth Engine (GEE). We began by obtaining Sentinel-1 Ground Range Detected (GRD) scenes, including VV and VH backscatter. We limited our query to images collected in interferometric wide and ascending swath mode (no descending available for this location) at 10 m resolution. The data was preprocessed automatically in GEE using tools from the European Space Agency's Sentinel-1 Toolbox ("https://github.com/senbox-orgz"). The GRD scenes are first processed to backscatter coefficient (σ°) in decibels (dB) and delivered as log-scaled ($10*\log10\sigma^\circ$). Next, the following pre-processing steps are implemented: (1) application of an orbit file for increased geometric accuracy, (2) the removal of low-intensity noise and invalid data on scene borders, (3) the removal of thermal noise to reduce discontinuities between sub-swaths, (4) radiometric calibration to backscatter intensity using sensor calibration parameters in the GRD metadata, and

(5) orthorectification to align ground range geometry to the terrain using a digital elevation model (DEM) from the Shuttle Radar Topography Mission (SRTM). Next, we implemented the Analysis-Ready Data method provided by [54] to apply additional border noise correction, speckle filtering using a multi-temporal Refined Lee filter [55], and enhanced radiometric terrain normalization. The pre-processed data were then downloaded locally for further analysis outside of GEE.

3.1.2 Sentinel-2

We downloaded all Sentinel-2 top-of-atmosphere reflectance images (L1C) between January 2017 and December 2021 that intersected the study area boundary and had $< 70\%$ cloud cover, resulting in 342 images. Batch conversion of the L1C images to L2A surface reflectance was done using the Sen2Cor software application, v2.9 [56]. Terrain correction was applied using an SRTM DEM for the study area. For cloud masking, we used the Sentinel-2 cloud probability dataset from GEE, which was created using the s2cloudless algorithm [57]. We applied strict filtering where pixels with cloud probability $> 30\%$ were excluded. Cloud shadows were masked using the Sentinel-2 L2A Scene Classification Map, where any pixel falling in the 'Cloud Shadows' category was excluded. While snow is somewhat rare in eastern North Carolina, it can still occur, therefore, pixels that had a snow probability $> 30\%$ were also masked based on the snow probability band output by Sen2Cor.

3.1.3 PlanetScope

We downloaded 4 band PlanetScope images for dates within the January 2017–December 2021 time period for which we did not already have a Sentinel-1 or Sentinel-2 image available. This gap-filling querying procedure was done to reflect typical budgetary constraints for operational monitoring systems and to demonstrate how to effectively leverage commercial data while minimizing overall data acquisition costs. PlanetScope data was queried to only download imagery with $< 40\%$ cloud cover, resulting in a total of 8,335 PlanetScope images across 531 unique dates. We used Planet's harmonization tool to radiometrically harmonize the PlanetScope imagery to Sentinel-2 and to correct for differing spectral responses and bidirectional reflectance effects [8, 19, 58]. The Usable Data Mask (UDM2) of each image was used to retain only those pixels classified as '1 - Clear' [59] and to remove any cloudy pixels.

3.2 Machine Learning Classification

3.2.1 Training & Validation

Eight different image dates were selected that had good visibility (low cloud cover) and close temporal proximity ($\leq$4 days) between the Sentinel-1, Sentinel-2, and PlanetScope image dates to generate training data from (Table 1). These dates were chosen to capture seasonal and interannual variability. Areas of *Dry Land*, *Inundated Vegetation*, and *Open Water* were identified by manually delineating and selecting polygon-based sampling regions that were representative of each type, guided by the ESA WorldCover dataset, NWI dataset (Fig. A.11 in Supplementary Material), and visual interpretation of high-resolution (3 m) PlanetScope imagery. Water was divided into *Open*

Water and *Inundated Vegetation* to account for the distinct spectral characteristics of water mixed with vegetation, enabling more accurate classification. Sampling regions were non-overlapping across image dates and were strategically delineated to ensure broad spatial representation of the study area so that our models weren't overfitting to specific areas in the study region but rather learning a broad range of possible pixels. We then performed stratified random sampling within the delineated sampling regions to reflect the relative class distributions observed in the study area based on the 2021 National Land Cover Dataset [53], targeting 50% for *Dry Land*, 30% for *Inundated Vegetation*, and 20% for *Open Water*. We had a final total of 976,000 training pixels across the three sensors used for training our models (Table 1).

Table 1: Dates of Sentinel-1, Sentinel-2, and PlanetScope image trios used to generate training points (top). The bottom portion of the table reports the number of pixels used to train each model.

Sentinel-1	**Sentinel-2**	**PlanetScope**
10/01/2017	10/03/2017	10/05/2017
11/30/2017	11/27/2017	11/27/2017
01/29/2018	01/26/2018	01/26/2018
10/14/2018	10/13/2018	10/13/2018
04/30/2019	05/01/2019	05/01/2019
08/28/2019	08/29/2019	08/29/2019
03/02/2020	03/01/2020	03/01/2020
12/08/2020	12/06/2020	12/06/2020
Number of Training Pixels		
105,600	102,400	768,000

A stratified, random probability sampling design is used to generate validation data following best practices and recommendations outlined in [60] and [61]. Validation data were derived from four trios of image dates that were excluded from model training (Table 2). These image dates were selected to represent different times of year, provide broad spatial coverage ($\geq$85%) with minimal cloud cover across images, and maintain close temporal alignment ($\leq$3 days) of Sentinel-1, Sentinel-2, and PlanetScope acquisitions. By selecting images not used to create the training set, we ensured full independence and separation between the training and validation data sets. Validation points (n=500) were randomly generated across the entire study area for each of the four dates. Each point was manually classified as *Dry Land*, *Inundated Vegetation*, or *Open Water* based on visual cues by a single interpreter using the PlanetScope image at its native 3 m resolution as the reference image. Sentinel-2 SWIR false color composites (using bands B12, B8A, B04), NWI water regime, and Google Earth imagery were used as ancillary data as needed. We implemented a stratified sampling strategy that reflected the expected class distributions of our study area, requiring a minimum of 20% representation for both *Open Water* and *Inundated Vegetation* classes. Random sampling was extended beyond $n = 500$ as needed to meet the required class proportions. As a result, each satellite image date had 300 *Dry Land* points, 200 *Inundated Vegetation* points, and 200 *Open Water* points to use for validation, resulting in a total of 2,000 validation points for each sensor.

Table 2: Dates of Sentinel-1, Sentinel-2, and PlanetScope image trios used to generate validation points based on the corresponding PlanetScope reference imagery (top). The bottom portion of the table reports the total number of validation points used to evaluate each sensor or product.

Sentinel-1	**Sentinel-2**	**PlanetScope**
04/11/2018	04/11/2018	04/11/2018
06/29/2019	06/30/2019	06/30/2019
11/14/2020	11/16/2020	11/17/2020
01/13/2021	01/15/2021	01/15/2021
Number of Validation Pixels		
2,000	2,000	2,000

3.2.2 Predictor Variables

We trained individual models for each sensor and assessed the sensitivity of our choice in ML algorithm by testing both Random Forest (RF) [62] and Extreme Gradient Boosting (XGB) ensemble decision-tree methods [63], resulting in a total of six models. Deep learning has shown in recent years to match or outperform decision tree methods, but the limited and costly generation of training data led us to test methods that could capitalize on a small amount of pixel-based training data and were robust in the literature for these specific sensors. All input features used in the models are listed in Table 3. For the Sentinel-1 model, both VV and VH backscatter (speckle-filtered) were used, as was a mix of VV and VH combinations ('VV + VH', 'VV-VH', 'VH/VV', 'VV/(VV + VH)', 'VH/(VV + VH)') following the Sentinel-1 variables considered in [26] that were shown to improve mapping of inundated vegetation. For the Sentinel-2 model, all spectral bands were used in addition to brightness, the Normalized Difference Vegetation Index (NDVI) [64], Normalized Difference Wetness Index (NDWI) [65], Modified Normalized Difference Wetness Index (mNDWI) [66], Automated Water Extraction Index (AWEIsh & AWEInsh) [67], and 3-Band (BU3) [68] spectral indices. We also included an annual BU3 index to our Sentinel-1 and Sentinel-2 models following the success [26] had by including this feature. For the PlanetScope model, we included all spectral bands, brightness, and NDWI [65].

For all models, slope and topographic wetness index [69] were included and derived based on 0.95 m resolution Digital Elevation Models (DEM) downloaded from North Carolina Spatial Data Download (https://sdd.nc.gov/). The terrain-derived features were calculated based on the DEM's native resolution and then resampled using the nearest neighbors resampling method to match each sensor's native resolution. Since inundation in our site is known to be heavily driven by precipitation events, we included 7- and 30-day accumulated precipitation from 1 km DAYMET [70] data to capture precipitation events and antecedent wetness. These data were also resampled to each sensor's native resolution using the nearest neighbors method. Percent canopy cover was downloaded from the U.S. Forest Service [71], resampled using the same method, and included as a feature to help the models distinguish between the *Open Water* and *Inundated Vegetation* classes. All models also included the mean of each feature within a 3 x 3 window to add spatial context, a method previously shown to be useful in flood mapping [23, 72]. Precipitation and canopy cover

features were excluded from the moving window transformation, as they were already at a resolution coarser than each sensor's resolution.

Table 3: Input features used in each classification model. An asterisk (*) denotes that a 3 x 3 local mean of this feature was also used as a predictor.

Sentinel-1	Sentinel-2	PlanetScope
VV*	Red*	Red*
VH*	Green*	Green*
VV+VH*	Blue*	Blue*
VV-VH*	NDVI*	NDVI*
VH/VV*	NDWI*	NDWI*
VV/(VV+VH)*	AWEInsh*	NIR*
VH/(VV+VH)*	AWEIsh*	Slope*
Slope*	NIR*	TWI*
TWI*	SWIR1*	7-Day Precip.
BU3 Annual	SWIR2*	30-Day Precip.
7-Day Precip.	Brightness*	Canopy Cover
30-Day Precip.	Slope*	
Canopy Cover	TWI*	
	BU3	
	BU3 Annual	
	7-Day Precip.	
	30-Day Precip.	
	Canopy Cover	

Models were trained using scikit-learn [73], XGBoost [63], and Dask [74] Python libraries on North Carolina State University's high-performance computing cluster. These models were tuned, trained, and evaluated using each sensor's training and validation sets. For detailed information on tuning, see the Supplementary Material (Table A.8). To assess accuracy, we quantified user's accuracy (omission error), producer's accuracy (omission error), and overall accuracy metrics in addition to the proportion of area mapped and total area mapped for each of the three classes following best practices outlined in [60]. Given the primary objective of mapping surface inundation as accurately as possible, we also computed accuracy metrics using a binary classification that grouped *Open Water* and *Inundated Vegetation* together into a single *Surface Inundation* class. Confusion between these classes is common, as heavily inundated areas may resemble *Open Water*. Therefore, both multiclass and binary metrics were calculated and analyzed across sensors. For each sensor, we selected either RF or XGB based on the highest reported accuracy for the binary *Surface Inundation* class and used that model to predict for each pixel across all five years.

3.3 Data Fusion

After generating classified maps using machine learning, we implemented a data fusion approach using FRK [46]. FRK combines all available inundation maps to estimate inundation at every location and time, filling in gaps where no sensor observed and providing a probability of inundation rather than a single 0/1 label. Time series of binary inundation maps from Sentinel-1, Sentinel-2, and

PlanetScope, each on its native grid and acquisition dates, served as inputs, with each pixel classified as 0 (*Dry*) or 1 (*Surface Inundation*). These inputs form sparse, irregularly spaced observations of inundation at different resolutions. Using the FRK model, we produced spatially and temporally complete time series of inundation maps on a uniform 6 m grid as output, assigning each pixel an inundation probability and associated uncertainty quantified through prediction interval width. In the following, we describe in detail how the FRK model is used within our machine learning–data fusion framework (Fig. 2).

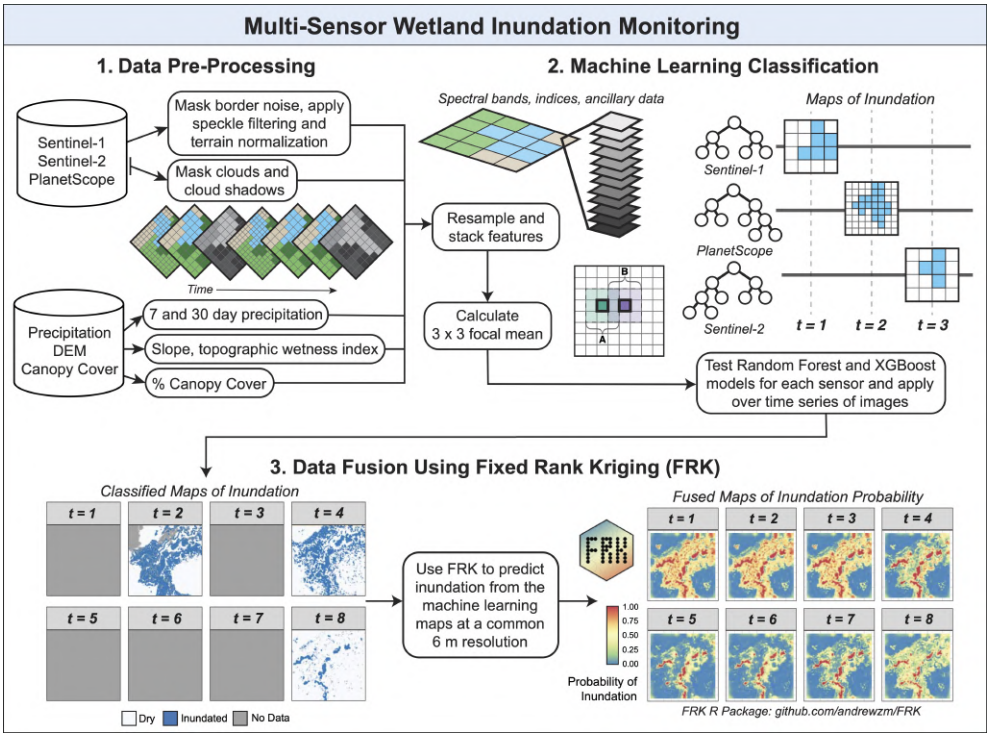

Figure 2: Workflow for the machine learning–based data fusion framework used in this study. Satellite imagery is first pre-processed, then classified using machine learning models, and the resulting inundation maps are used as inputs to the FRK data fusion model.

3.3.1 FRK Model Specification and Assumptions

FRK is uniquely suited for our problem because it accommodates multiple spatial supports (varying sensor resolutions) and is formulated to be computationally efficient with massive datasets by using a class of non-stationary covariance functions defined using a fixed set of basis functions. FRK models the conditional mean of the data as a transformation of a latent spatial process through a low-rank SRE model [46]. For our specific application, this framework operates under the assumptions that 1) the process generating our inundation observations is the same (i.e. inundation across satellite sensors is generated as a function of climate and topographic variables), 2) the observations based on the machine learning maps constitute a Bernoulli (binary) response, and 3) inundation has dependence across space and time. We used the FRK v2 R package (https://github.com/andrewzm/FRK) to

implement this approach, and we refer the reader to [75] and [76] for further details on the method and its implementation in the R package.

FRK assumes a discretized domain of interest composed of non-overlapping spatial units referred to as Basic Areal Units (BAUs) [44] to accommodate the varying spatial resolutions. These units represent the finest pixel grid at which we will make predictions with our model. For our application, we chose a resolution of 6 m to balance between our coarsest (10 m) and highest (3 m) satellite resolutions. A natural choice would be to model everything at the scale of the highest-resolution data source; however, we chose a resolution that balances the need for high-resolution predictions with the need to maintain computational efficiency. To formulate the BAUs, consider a vector of Sentinel-2 observations $\mathbf{Z}$ at m pixels $\{B_1, ..., B_m\}$ for a single point in time in our defined spatial region D (Fig. 3A). We assume each pixel in $\mathbf{Z}$ is an average of the true process $Y(\cdot)$ (Fig. 3B) plus noise, where noise includes measurement error and any sensor bias. For example, a higher resolution PlanetScope image (Fig. 3C) would have less noise than the coarser Sentinel-2 image (Fig. 3D) due to the finer spatial details more closely resembling the true process $Y(\cdot)$. Taking the vector $\mathbf{Z}$ of Sentinel-2 observations as an example, these pixels are mapped to BAUs as follows:

$$Z(B) = \frac{1}{|D \cap B|}\left\{\sum_{\mathbf{s} \in D \cap B} Y(\mathbf{s})\right\} + \epsilon(B) \tag{1}$$

where a Sentinel-2 pixel B is made up of BAUs with locations indexed by $D \cap B$, capturing the average of $Y(\cdot)$ indexed at $\mathbf{s}$ spatial location plus the measurement error term $\epsilon(B)$ (Fig. 3E.)

The FRK v2 implementation in the FRK R package uses a spatiotemporal hierarchical model defined on the BAU grid and performs kriging using tensor products of spatial and temporal basis functions, allowing nearby observations in space and time to inform predictions. It is worth noting, particularly for monitoring applications, that the temporal dependence is modeled via temporal smoothing across all time points rather than through a filtering approach that conditions each date only on past observations. Within this framework, we model binary inundation observations as Bernoulli responses with a probit link to a latent process $\boldsymbol{Y}$. We represent $\boldsymbol{Y}$ using the SRE model, which at the BAU level decomposes into fixed effects, low-rank spatial random effects, and fine-scale variation:

$$\boldsymbol{Y} = \boldsymbol{T}\boldsymbol{\alpha} + \boldsymbol{S}\boldsymbol{\eta} + \boldsymbol{\xi}, \tag{2}$$

where $\boldsymbol{Y}$ is the spatial process evaluated over N BAUs, $\boldsymbol{T}$ and $\boldsymbol{S}$ are known design matrices corresponding to spatially referenced covariates and constructed basis functions, $\boldsymbol{\alpha}$ is a vector of fixed effects, $\boldsymbol{\eta}$ is a vector of random coefficients associated with the basis functions, and $\boldsymbol{\xi}$ is a fine-scale-variation random process. Model parameters for the fixed effects and random effects were estimated by maximizing the observed-data likelihood using a Laplace approximation; full model specification and estimation details for our non-Gaussian application are provided in the Supplementary Material

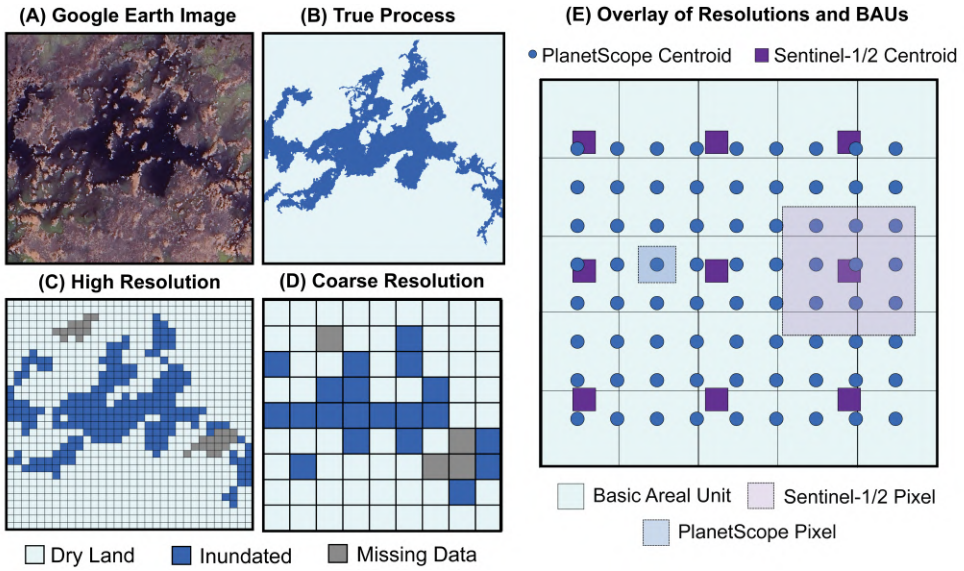

Figure 3: Comparison of (A) Google Earth satellite image of the spatial domain to (B) the true inundation process, (C) a higher resolution classification (i.e., from PlanetScope), and (D) a coarser resolution classification (i.e., from Sentinel-1/2). (E) shows an overlay of the Sentinel-1/2 and PlanetScope pixel footprints (dashed) and pixel centroids (points) to the BAU prediction grid generated using the FRK package.

(Equations 3–8) and in [76].

3.3.2 Simulation Study

We first conducted a simulation study to evaluate FRK's applicability for fusing remotely sensed classification maps. We generated a spatially auto-correlated process, $\boldsymbol{Y}$, which was subsequently used to simulate remote sensing observations of inundation, $\boldsymbol{Z}$, from data sources consisting of different spatial supports. $\boldsymbol{Y}$ was simulated using a Matern spatial covariance function and a first-order autoregressive time series covariance function across a 150 m x 150 m spatial grid for 20 points in time. A mask of potential cloud cover was also simulated using this same approach. To go from this true process $\boldsymbol{Y}$ to binary remote sensing observations at 3 m (Z_A) and 10 m (Z_B), we added Gaussian noise, split binary classes of 0 and 1 based on a mean threshold, resampled the observations to the desired resolutions, and removed pixels that intersected the cloud cover mask. An entire day of data was also removed to simulate a gap in satellite acquisition time. Although this represents a simplified data stream relative to real classification maps, it provides a useful starting point for evaluating and demonstrating FRK's suitability for this application. We specified approximately 81 bisquare spatial basis functions and 10 bisquare temporal basis functions (scale parameter = 2), yielding 810 spatiotemporal basis functions, which were used to fit and predict with the FRK model on the simulated data.

3.3.3 Implementation

After evaluating the suitability of the method through simulation, we applied it to real data across three 6 km × 6 km subsets over one year (Fig. 1). Two of these subsets coincide with eddy covariance flux tower sites where continuous WTD (and associated meteorological) data are available, and the third encompasses an area with frequent inundation changes visible in satellite imagery and located near a USGS gauge recording lake elevation. We focused the FRK application on these three subsets rather than the entire study area because (1) they are the only locations with independent ground-based hydrologic data suitable for evaluating the fused inundation products, and (2) daily 6 m spatiotemporal modeling with FRK is computationally intensive at the scale of the full 5,778 km^2 domain. Thus, this implementation is intended as a pilot study demonstrating and evaluating the method's usability, rather than producing a data product for the entire study area.

Each subset was partitioned using a spatial–temporal moving-window approach, with 2 km × 2 km spatial windows and 21-day temporal windows, retaining the middle 10 days as predictions. To optimize model performance and computational efficiency, we conducted a small grid search over (1) the number of spatial basis functions, (2) the number of spatial resolutions, (3) the number of temporal basis functions, and (4) the temporal scale. For each window, we used approximately 1,700 single-resolution bisquare spatial basis functions and 8 bisquare temporal basis functions with a scale parameter of 4, resulting in about 14,000 spatiotemporal basis functions per subset. We chose a large number of spatial basis functions to preserve fine-scale inundation patterns while restricting to a single spatial resolution to keep computation manageable. Sensor-specific measurement variances were approximated based on the classification accuracies of each machine learning model and held constant across space and time. This allowed the relative strengths of each sensor to inform the final prediction of inundation. Slope was included as a covariate because microtopography has been shown to strongly influence where water accumulates in this landscape.

Using the FRK package, we predicted daily inundation probability for each BAU at 6 m resolution over one year and quantified uncertainty with 90% credible intervals. The final predictions were spatially mosaicked by computing weighted averages in overlap areas to generate a complete prediction map for each day. We then compared the fused maps and associated time series with individual sensor outputs and the Landsat DSWE product to assess whether fusion improves monitoring of inundation dynamics. To further evaluate performance, we summarized both individual sensor outputs and fusion maps as time series of total inundated area and compared them with precipitation and WTD data from the US-NC2 and US-NC4 flux towers, as well as nearby stilling pipe records. It is important to note that each in-situ measurement represents a point location, while the satellite time series reflects broader-scale surface water patterns. This distinction is important because WTD and canal water levels can vary independently of surface water extent due to evapotranspiration, subsurface flow, and irrigation. Despite these limitations, these records provide the only available ground-based observations and thus serve as valuable references for assessing similarities in hydrologic dynamics.

We also compared our fusion and individual sensor results to the operational Landsat DSWE data product to benchmark our results against an existing surface water product. The DSWE

product provides four pixel classes at 30 m resolution: (1) Water – High Confidence, (2) Water – Moderate Confidence, (3) Partial Surface Water – Conservative, and (4) Partial Surface Water – Aggressive [17, 18], which we hereafter refer to as DSWE-1, DSWE-2, DSWE-3, and DSWE-4 respectively. We downloaded DSWE data for 2019 using the "Interpreted Layer With All Masks Applied" processing layer and restricted downloads to images with <70% cloud cover, resulting in 45 images for 2019. We then compared the total inundated area mapped in 2019 by our fusion model, the individual sensor models, and the DSWE data product by summing the inundated area over time. To provide a baseline fusion benchmark, we also applied a simple gap-filling method where the last observation was carried forward (LOCF). This approach was applied to the individual sensor maps, the combined sensor maps, and the Landsat DSWE product, and the resulting interpolated time series were compared with our FRK fusion results.

4 Results

4.1 Machine Learning

4.1.1 Model Evaluation

The results for predicting the binary classes of *Surface Inundation* versus *Dry Land* showed that RF performed best for Sentinel-2 and PlanetScope sensors, while XGB performed the best for Sentinel-1. For detailed comparison metrics, see the Supplementary Material (Tables A.9-A.10). These top-performing algorithms were applied across five years of imagery (2017–2021) to produce over 1,000 classification maps (Fig. 4). We evaluated the final models for Sentinel-1, Sentinel-2, and PlanetScope by comparing the proportion of area that was mapped (i.e. classified) as *Dry Land*, *Inundated Vegetation*, and *Open Water* (Table 4–6). Of the three models, Sentinel-1 mapped the highest proportion of *Inundated Vegetation* at 0.224, closely followed by Sentinel-2 and PlanetScope models, which had mapped area proportions of 0.218 and 0.210 for *Inundated Vegetation*. Interestingly, PlanetScope mapped the highest amount of total water (*Inundated Vegetation* + *Open Water*) at 0.442 by detecting more *Open Water* area, followed by Sentinel-2 at 0.428 and Sentinel-1 at 0.425. For *Open Water* and *Inundated Vegetation*, all models produced higher user's than producer's accuracies (Sentinel-1: +0.003, +0.072; Sentinel-2: +0.047, +0.075; PlanetScope: +0.134, +0.034), indicating that most inundated areas were correctly identified, with the models exhibiting a tendency toward under-prediction rather than over-prediction. Sentinel-1 showed the lowest accuracy in detecting the *Inundated Vegetation* class, with user's and producer's accuracies of 0.677 and 0.605, respectively, while Sentinel-2 showed the highest, with accuracies of 0.896 and 0.821, respectively. The binary accuracy metrics for these models as well as accuracy metrics by date can be found in the Supplementary Material (Table A.11 and Fig. A.12). Feature importance analysis can also be found in the Supplementary Material (Fig. A.13), and these results indicated that TWI, derived from the 5 m DEM, was the most influential predictor for the Sentinel-1 XGB model. While for Sentinel-2 and PlanetScope RF models, $AWEI_{sh}$ and NIR reflectance were the most important features, respectively.

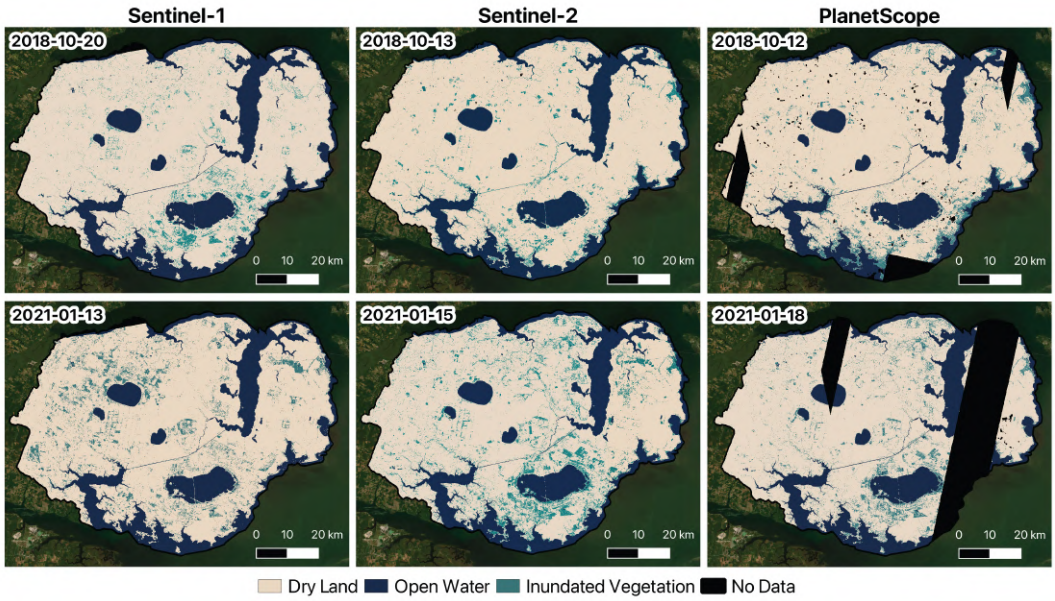

Figure 4: Maps showing example classifications for each sensor for a relatively dry date (top row) and a relatively wet date (bottom row).

Table 4: Sentinel-1 error matrix of classes expressed in terms of proportion of area and total area along with associated user's, producer's, and overall accuracies calculated using all available test pixels pooled across image dates. 95% confidence intervals are shown in parentheses.

		Mapped			
		Dry Land	Open Water	Inundated Vegetation	*Total*
Reference	Dry Land	0.513	0.001	0.086	0.600
	Open Water	0.000	0.197	0.003	0.200
	Inundated Vegetation	0.062	0.003	0.135	0.200
	Total (Area Proportion)	0.575 (±0.015)	0.201 (±0.004)	0.224 (±0.016)	1.000
	Total (Area in m^2)	115,085 (±3,090)	40,102 (±740)	44,813 (±3,134)	200,000
	User's Accuracy:	0.856 (±0.019)	0.985 (±0.012)	0.677 (±0.052)	
	Producer's Accuracy:	0.892 (±0.016)	0.982 (±0.014)	0.605 (±0.037)	
	Overall Accuracy:	0.846 (±0.015)			

Table 5: Sentinel-2 error matrix of classes expressed in terms of proportion of area and total area along with associated user's, producer's, and overall accuracies calculated using all available test pixels pooled across image dates. 95% confidence intervals are shown in parentheses.

		Mapped			
		Dry Land	Open Water	Inundated Vegetation	*Total*
Reference	Dry Land	0.560	0.001	0.039	0.600
	Open Water	0.000	0.200	0.000	0.200
	Inundated Vegetation	0.012	0.009	0.179	0.200
	Total (Area Proportion)	0.572 (±0.010)	0.210 (±0.005)	0.218 (±0.010)	1.000
	Total (Area in m^2)	114,385 (±1,1920)	41,993 (±904)	43,622 (±2,069)	200,000
	User's Accuracy:	0.933 (±0.014)	1.00 (±0.000)	0.896 (±0.032)	
	Producer's Accuracy:	0.979 (±0.008)	0.953 (±0.020)	0.821 (±0.031)	
	Overall Accuracy:	0.939 (±0.010)			

Table 6: PlanetScope error matrix of classes expressed in terms of proportion of area and total area along with associated user's, producer's, and overall accuracies calculated using all available test pixels pooled across image dates. 95% confidence intervals are shown in parentheses. Note, that PlanetScope pixels are smaller (higher resolution) than Sentinel-1 and Sentinel-2's 10 m resolution; therefore, the reported total area mapped (m^2) is smaller.

		Mapped			
		Dry Land	Open Water	Inundated Vegetation	*Total*
Reference	Dry Land	0.543	0.002	0.055	0.600
	Open Water	0.000	0.194	0.006	0.200
	Inundated Vegetation	0.016	0.036	0.148	0.200
	Total (Area Proportion)	0.558 (±0.011)	0.232 (±0.009)	0.210 (±0.014)	1.000
	Total (Area in m^2)	10,052 (±200)	4,176 (±162)	3,772 (±245)	18,000
	User's Accuracy:	0.904 (±0.016)	0.968 (±0.019)	0.740 (±0.045)	
	Producer's Accuracy:	0.972 (±0.010)	0.834 (±0.029)	0.706 (±0.037)	
	Overall Accuracy:	0.884 (±0.014)			

4.1.2 Satellite Predicted Inundation Dynamics

We summarized inundation dynamics from 2017–2021 by computing the percentage of cloud-free observations classified as inundated for each pixel (Fig. 5). We generated both a combined multi-sensor map and sensor-specific frequency maps, expressed as the fraction of observations labeled as either *Open Water* or *Inundated Vegetation*. These maps highlight zones of variable inundation identified by each sensor (Fig. 5). For example, Sentinel-1 and Sentinel-2 models map less frequent inundation along coastal areas bordering the ocean compared to the PlanetScope model (Fig. 5; Inset Y). Moreover, Sentinel-1 maps more frequent inundation in agricultural fields and open grasslands, while Sentinel-2 and PlanetScope report inundation less often in these areas (Fig. 5; Insets X, Y). Comparing these patterns show that Sentinel-2 and PlanetScope display similar patterns of mapped inundation frequency across land cover types, while Sentinel-1 tends to over map inundation in croplands and grasslands, further quantitative comparison of mapped inundation dynamics to different

land cover types can be found in the Supplementary Material (Fig. A.14). Differences in spatial resolution are also evident since the finer scale of PlanetScope captures small wetlands and ponds often missed by coarser sensors and more clearly distinguishes boundaries between inundated and dry land (Fig. 5 and Fig. A.15 in Supplementary Material). We also summarized total *Inundated Vegetation* mapped by each sensor as a time series, using only maps with at least 80% spatial coverage within a 6 km buffer around each flux tower (Fig. A.16 in Supplementary Material). Across time, inundated vegetation estimates from Sentinel-2 and PlanetScope were closely aligned, whereas Sentinel-1 indicated more frequent inundation with a less pronounced seasonal signal.

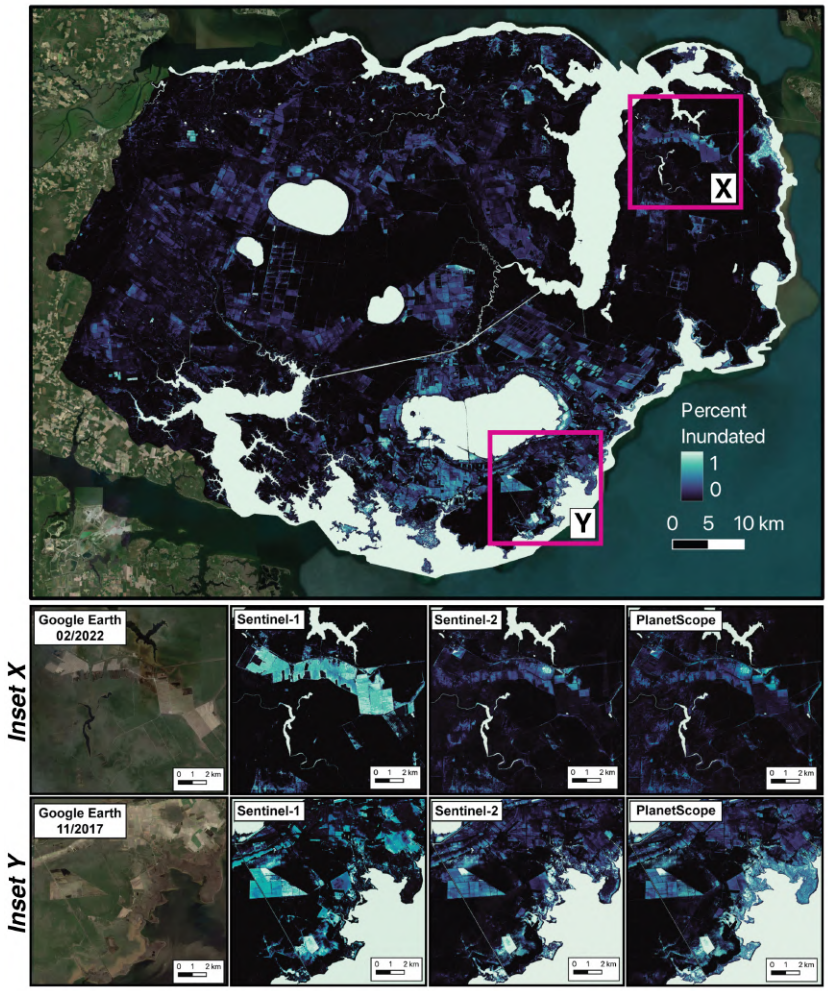

Figure 5: Percentage of time each pixel was classified as inundated from 2017–2021. Black indicates pixels that rarely or never flood, while light teal indicates areas that are permanently or frequently inundated. Insets X and Y show example areas, with inundation frequencies by sensor alongside true-color Google Earth imagery.

4.2 Data Fusion

4.2.1 Simulation Study

The simulation study produced synthetic inundation values across a 150 m x 150 m spatial grid for 20 points in time (Fig. 6). Remote sensing observations were derived from these synthetic inundation values and used to fit the FRK model with 810 basis functions, generating predictions at 6 m spatial resolution for each time point. The simulation shows how data fusion enables spatially and temporally continuous inundation predictions, even during periods of cloud cover (Fig. 6, Row 3). FRK was shown to effectively integrate observations with different spatial supports (3 m, 10 m) (Fig. 6, Row 2), producing predictions that closely matched the true process $\boldsymbol{Y}$ (Fig. 6, Row 1). We quantified uncertainty through the FRK package using Monte Carlo sampling to compute 90% prediction intervals. As expected, uncertainty increased during time steps when entire scenes were missing or partially obscured by clouds (Fig. 6, Row 4). Uncertainty also increased in areas with fine-scale inundation dynamics and at the boundaries of common transition zones between inundated and dry classes. Overall, our simulation shows that FRK effectively fuses non-uniform remote sensing classification maps to produce more complete spatiotemporal estimates for monitoring inundation dynamics.

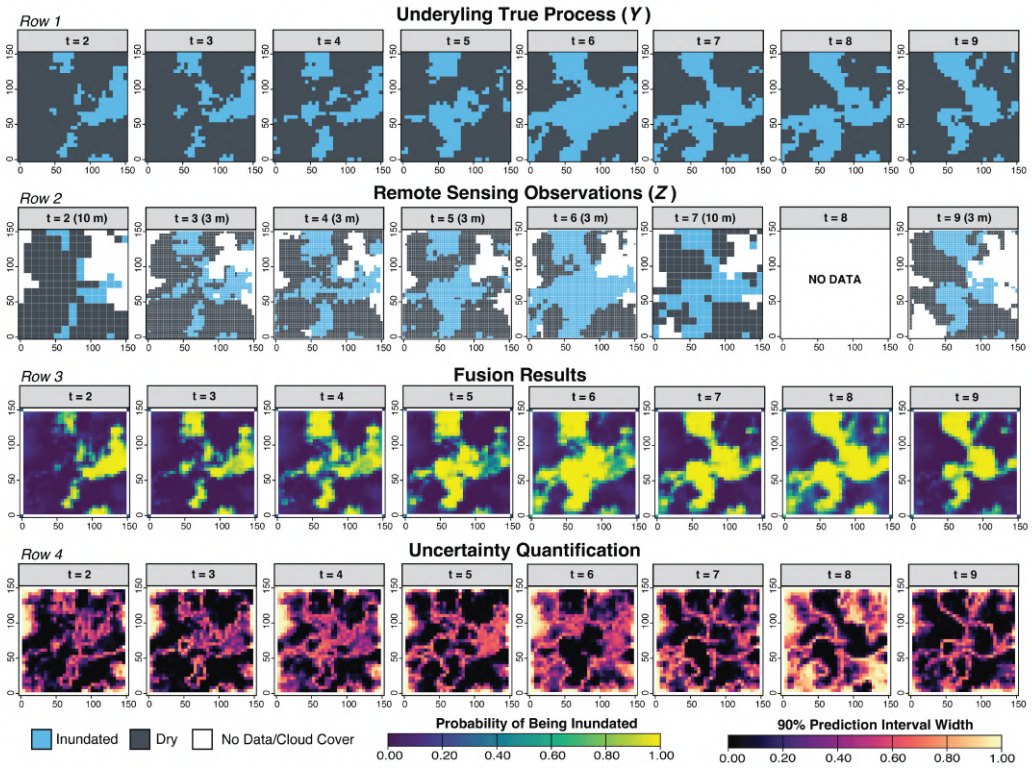

Figure 6: *Row 1* shows the true process $\boldsymbol{Y}$ we are interested in predicting, *Row 2* shows the derived remote sensing observations $\boldsymbol{Z}$ at 3 and 10 m spatial resolutions with clouds present, *Row 3* shows the resulting fusion maps depicting the probability of each pixel being inundated, and *Row 4* shows the associated uncertainty maps generated by taking the width of the 90% prediction interval.

4.2.2 Ephemeral Changes Mapped Using Fusion

Building on our simulation study, our fusion framework successfully integrated heterogeneous observations from multiple satellite sensors, filling temporal gaps and improving detection of ephemeral inundation between satellite overpasses. Using PlanetScope imagery, we highlight two examples where a small area quickly transitioned from predominantly dry conditions to being inundated, with flooding events occurring between 8/12-8/20 (left panel) and 6/17-6/21 (right panel) (Fig. 7). The fusion model captured the progressive expansion of inundated areas with high confidence (Fig. 7, Example A) by the time the reference imagery displayed the resulting inundation, while this event would go undetected for days using methods that relied on a single satellite sensor. Notably, the model maintained spatial continuity and realistic predictions even during periods of data gaps, leveraging the spatiotemporal correlations learned from available observations.

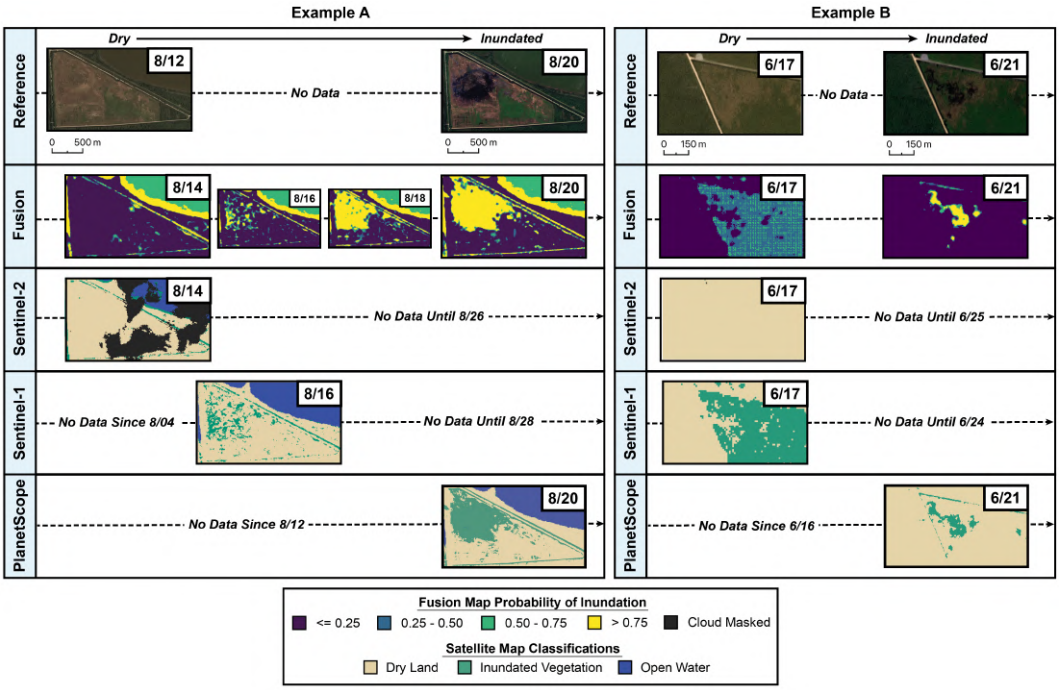

Figure 7: Examples of short-term inundation events captured by the fusion approach. **Left:** Example A, an agricultural field bordering Lake Mattamuskeet, showing the transition from dry (8/12) to inundated (8/20) and the corresponding fusion and single-sensor classifications. **Right:** Example B, a wetland near flux tower US-NC4, illustrating how the fusion model integrates differing Sentinel-1, Sentinel-2, and PlanetScope observations between 6/17 and 6/21. Only a subset of fusion maps are shown to conserve space.

The fusion framework also was shown to correctly account for sensor-specific measurement error (Fig. 7, Example B). During the 6/17-6/21 flood event, coincident Sentinel-1 and Sentinel-2 observations on 6/17 provided conflicting information: Sentinel-1 indicated inundation while Sentinel-2 classified the field as dry. This discrepancy reflects a limitation we observed with our Sentinel-1 model in that it predicts saturation, rather than true inundation in many open field and agricultural

areas. However, because Sentinel-1 exhibited lower classification accuracy than the other sensors, it was assigned a larger measurement error and thus a lower weight in the fusion process, resulting in only moderate inundation probabilities (∼50%) rather than high-confidence predictions in the fusion maps. When combined with the *Dry Land* classifications from Sentinel-2, the erroneous Sentinel-1 signal was effectively down-weighted, and subsequent observations, particularly the high-accuracy PlanetScope classification on 6/21, quickly corrected the fused inundation probabilities toward the true surface condition. This emphasizes the potential for systematic bias in inundation mapping when observations from a single sensor are used, and demonstrates our fusion framework's ability to better balance these biases and detect ephemeral inundation changes.

4.2.3 Fusion Predicted Inundation Dynamics

Applying the fusion approach across our three subsets, we found that small-scale ponds and wetland patches consistently had lower inundation probabilities than larger ponds, fields, and open areas, highlighting the influence of spatial scale on predictions. We also computed, for 2019, the fraction of days each pixel was inundated by thresholding daily inundation probabilities at 0.5, classifying pixels with predicted probability >0.5 as *Inundated* and those ≤0.5 as *Dry Land* (Fig. 8). This provides an inundation-frequency summary comparable to the maps shown in Fig. (5). From this summary, we found that fusion subset A showed variable inundation patterns in the agricultural fields and drainage canals along field edges. Fusion subset B also included many agricultural fields that experienced intermittent to more prolonged inundation, with many wetland areas along the coast experiencing semipermanent to permanent inundation dynamics. Finally, fusion subset C showed occasional inundation in the forested southwest quadrant and more sustained inundation in the southeast quadrant, encompassing the Dare County Range, an area modified by human activity for military training. The agricultural field in this area was also shown to flood frequently, and several artificial drainage canals, identifiable by their linear patterns, were shown to be regularly inundated.

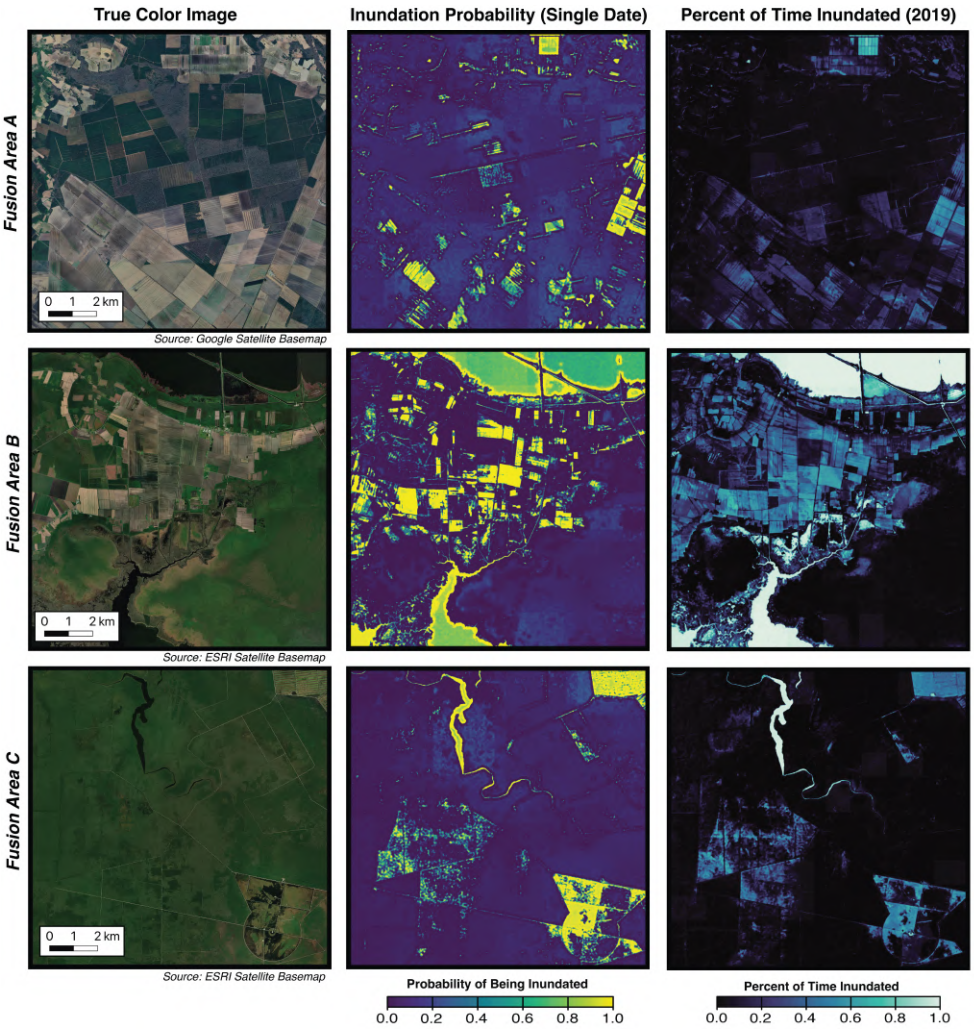

Figure 8: Maps for fusion subsets A, B, and C (see insets in Figure 1). Left column: true-color imagery (Google or Esri satellite basemaps). Middle column: example daily inundation probability maps. Right column: percentage of days inundated in 2019, derived using a 50% inundation probability threshold.

Using the same fusion subsets, we calculated the total inundated area through 2019 using varying probability thresholds to classify pixels as inundated or dry (Fig. 9). We then compared the predicted inundated area time series with WTD, precipitation, and lake elevation measurements collected at the flux towers and a nearby USGS gauge. In fusion subset A (Fig. 9), we observed high-level agreement between increased inundated area mapped and higher WTD levels from January to mid-April, and from late-October to the start of January 2020. We also observed two spikes in WTD and precipitation in mid-June and early July. These spikes are reflected in the total inundated area mapped produced during fusion, with the 50% threshold results going from <5 km^2 inundated area mapped to >20 km^2 inundated area mapped for both points in time corresponding to the dates at which an increase in WTD is observed. Moreover, between these two spikes, a gradual decrease in

WTD is observed, which is also shown in the total inundated area mapped by fusion, which at the 50% threshold drops from a maximum of approximately 24 km^2 to <2 km^2 total area mapped before spiking again to >20 km^2 in early July.

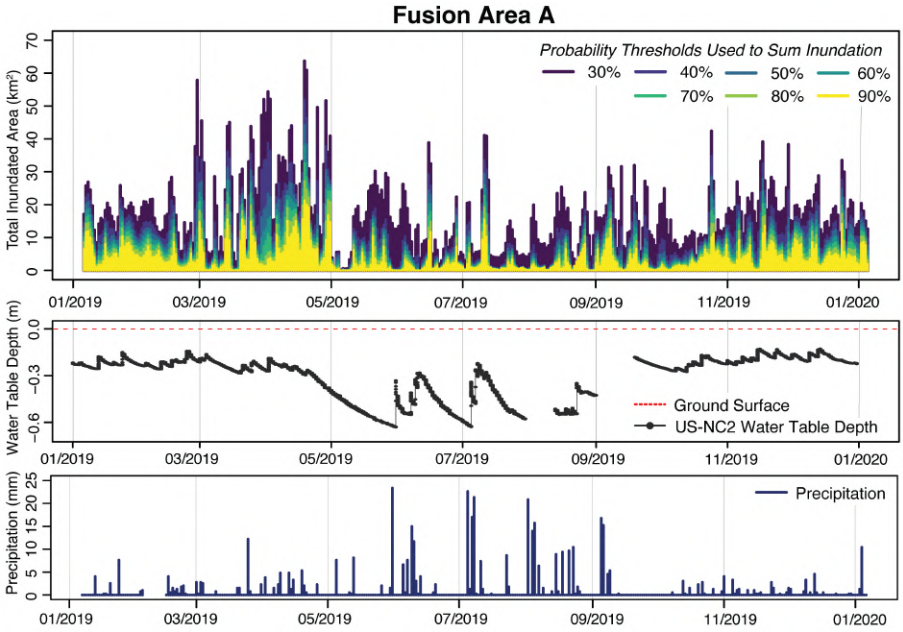

Figure 9: Summary of fusion results for subset A (Fig. 1). Top: time series of total inundated area from the fusion approach, with colors indicating different inundation probability thresholds. Middle: water table depth (WTD) relative to the ground surface at flux tower US-NC2. Bottom: precipitation measured at US-NC2.

Time series for fusion subsets B and C can be found in the Supplementary Material along with percent canopy cover maps for each subset for reference (Fig. A.17, Fig. A.18.). Fusion subset B had the highest total mapped inundation of the three subsets. Inundation levels were generally higher in winter (leaf-off months) and lower in summer, which was reflected in the water elevation of Lake Mattamuskeet recorded at a nearby USGS gauge (USGS-0208458893 in Fig. 1). In fusion subset B, we observed a spike in precipitation and WTD in early September, which was captured more subtly in the fusion results, with the 50% threshold spiking from <30 km^2 to >40 km^2 mapped during this same time. This spike was followed by a brief drawdown period leading into October, during which precipitation remained minimal. Our fusion results reflect this trend, showing a corresponding decline in total mapped inundated area with the 50% threshold going from >40 km^2 mapped at the spike to <30 km^2 after the observed decrease in lake water levels. In fusion subset C, the flux tower WTD data showed a more stable trend, oscillating around a mean value of approximately 0.7-1.0 m. We observed two notable WTD increases in late June and early September. The June increase is captured in the 30% threshold results, which go from <10 km^2 inundated area mapped to >25 km^2, while higher thresholds (40–90%) show only subtle increases. The September spike corresponds to a smaller rise in our fusion results, with our 50% threshold spiking from <5 km^2 to >13 km^2; however, the 30% threshold mirrors the sharp late-August precipitation event, spiking from <10 km^2 to >35

km^2 inundated area mapped. Further time series comparing fusion results to individual satellite sensors can be found in the Supplementary Material (Fig. A.19).

4.3 Comparisons Across Methods

4.3.1 Total Inundated Area Mapped

We benchmarked the fusion model against a simple gap-filling (LOCF) fusion alternative by comparing 2019 total inundated area from our fusion approach to gap-filled and non-gap-filled estimates derived from three data streams: 1) individual sensor maps, 2) multi-sensor maps, and 3) Landsat DSWE maps. The gap-filled Sentinel-2 data stream maps 1.2 times more inundation than the results of our fusion model using a 50% threshold. Across all non–gap-filled estimates for subsets *A*, *B*, and *C*, including individual sensor predictions and fusion results using thresholds from 50% to 90%, DSWE-1 mapped the smallest total inundated area (874.5 km^2), whereas the fusion product at the 50% probability threshold mapped the largest (22,694.5 km^2; top left panel of Fig. 10). When including gap-filled observations, DSWE-1 still mapped the lowest amount of inundation (6,212.9 km^2) while DSWE-4 mapped the most (33,427.6 km^2) (right, top panel of Fig. 10). Notably, differences between individual sensor models in both gap filled and non-gap filled approaches are shown to differ considerably, particularly in fusion subsets *A* and *B*.

When benchmarked against the operational DSWE product, which served solely as a comparison dataset rather than a validation reference, our fusion maps showed strong correspondence between DSWE confidence classes and analogous inundation-probability thresholds. The fusion approach mapped 3.7 times the inundated area of DSWE-1 at a threshold of 50%. However, when we raised the threshold to 90%, that ratio dropped to 1.5. Similarly, DSWE-4 mapped 1.5 times more inundation than the fusion model at the 50% threshold, but using a 30% threshold produced nearly identical results (1.02 times more). These comparisons illustrate how adjusting the fusion probability threshold lets us account for different confidence levels in inundation detection. Lower thresholds capture more potential inundation, aligning with products like DSWE-4 that prioritize high recall and include partial surface after pixels. In contrast, higher thresholds yield more conservative estimates, similar to DSWE-1's high-precision focus on fully inundated pixels. Notably, across subsets *A*, *B*, and *C*, our fusion model always predicts more inundation than the DSWE-1 maps, even when using a conservative 90% probability threshold. Maps showing the spatial differences in areas mapped between our fusion approach and the DSWE dataset can be found in the Supplementary Material (Fig. A.20).

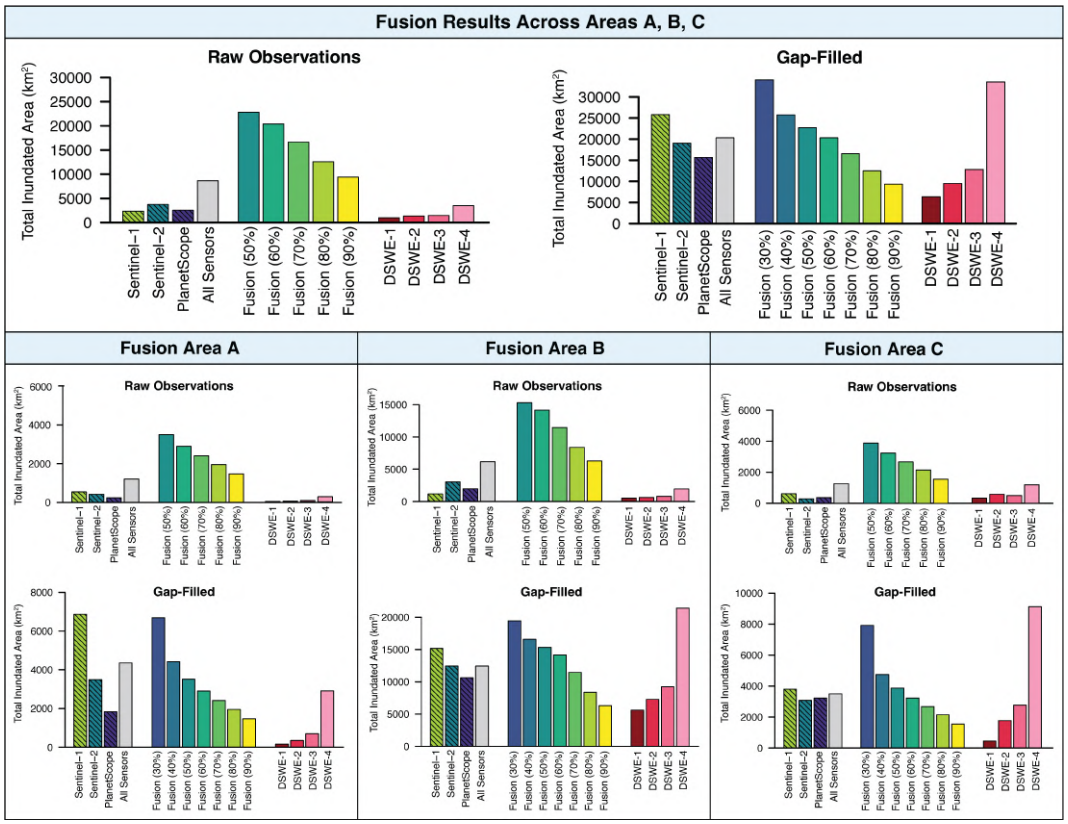

Figure 10: Bar plots showing the total summed inundation over 2019 across all fusion regions (top row), and for each fusion region (bottom panels). For each location, results are shown for both raw observations and gap-filled results.

5 Discussion

5.1 Comparing Sentinel-1, Sentinel-2, and PlanetScope Models

Accurately monitoring inundation extents at high spatial and temporal resolutions is critical for quantifying wetland carbon dynamics, but remote sensing methods face technical challenges due to observational gaps, sensor inconsistencies, and differences in resolutions across satellite platforms. In evaluating the relative strengths and limitations of different sensors and models for detecting inundation in vegetated wetlands, we identified several key findings. First, we found the performance of the machine learning models varied by sensor type: XGB was more effective with Sentinel-1 data, likely because its built-in regularization reduced overfitting to noisy features more common in SAR backscatter than in optical data [63]. For the optical data sources, RF and XGB showed comparable performance, with RF slightly outperforming XGB in most cases. However, their confidence intervals often overlapped, suggesting the two models can be used interchangeably.

After selecting the top-performing models for each sensor and applying them across the five years of imagery, we found that Sentinel-1 mapped substantially more inundation than the optical sensors

in open fields and agricultural areas. This finding may indicate sensitivity to surface moisture and saturation in our Sentinel-1 model, which may detect irrigation or wet soils rather than standing water. Additionally, we hypothesize that the inclusion of annual Sentinel-2-derived BU3 index and canopy cover features, while enhancing Sentinel-1's ability to detect inundated vegetation, may also have contributed to its tendency to over-predict inundation. Although including these features likely caused slight overestimation, for our purposes this was preferable to underestimation. For the optical sensors, Sentinel-2 achieved the highest accuracy for inundated vegetation detection, closely followed by PlanetScope, highlighting the strengths of these two sensors for wetland inundation monitoring. PlanetScope identified the highest amount of total water, yet its detection of inundated vegetation remained consistent with the other sensors. This higher total is instead explained by the PlanetScope model's ability to better capture open-water features such as small ponding and narrow, intermittent stream channels, that were omitted by the coarser-resolution sensors.

These results suggest that PlanetScope could serve as a standalone sensor for inundation monitoring in scenarios without resource constraints (e.g., where budgets allow acquisition of dense commercial satellite imagery). However, two factors highlight the value of multi-sensor integration. First, the incorporation of terrain-derived features (e.g., Slope, TWI) likely helped compensate for PlanetScope's lack of SWIR bands, which are known to improve water detection accuracy. This could be an important consideration in areas where high-resolution elevation data may not be available. Second, Sentinel-2's superior performance in classifying inundated vegetation (user's accuracy: 0.896 compared to PlanetScope's 0.740) demonstrates that inclusion of SWIR-based spectral indices enabled more accurate delineation of water in vegetated environments, which is important for wetland specific monitoring. Finally, although the Sentinel-1 model tends to overpredict inundation, its value is still in its ability to acquire data under cloud cover and increase the temporal density of observations. Moving forward, the ability for SAR data to more accurately estimate inundation, particularly inundated vegetation, will likely be improved with the release of L-band data from the NASA-ISRO SAR (NISAR) satellite mission [77].

5.2 Data Fusion with FRK

Statistical data fusion enhanced multi-sensor inundation monitoring by spatiotemporally interpolating observations and balancing sensor-specific biases. Our approach represents the first application of FRK to satellite-derived binary classification maps. Unlike traditional fusion methods that typically operate on continuous raw satellite values, our framework is applied to non-Gaussian (binary) data. The resulting daily time series of inundation probability at 6 m resolution were shown to capture hydrologic dynamics more effectively than individual sensors, particularly short-duration events such as the two inundation pulses occurring before and after July 1st in the inundation time series for subset A (Fig. 9), and the examples of ephemeral changes captured in June and August fusion maps (Fig. 7). While our FRK framework introduces temporal smoothing through its spatiotemporal covariance structure, the model retained sufficient sensitivity to detect ephemeral inundation events and appropriately weighted sensor contributions based on their classification accuracies.

Benchmarking against the LOCF fusion approach revealed distinct advantages of our FRK-

based fusion framework. The LOCF method, which assigns the most recent valid observation to data gaps, fails to account for sensor-specific biases and uncertainties. For example, using gap-filled Sentinel-1 observations alone in subset A substantially overestimated inundation extent compared to Sentinel-2 and PlanetScope (Fig. 10). While incorporating all three sensors in the LOCF approach partially mitigated individual sensor biases in the total area summed, the resulting time series would exhibit abrupt, unrealistic transitions when switching between sensors, representing artifacts of their differing sensor biases rather than true hydrologic variability. In contrast, our data fusion framework produced temporally smooth probability estimates by weighting observations according to their spatiotemporal correlation structure and measurement uncertainties. In further benchmarking this LOCF approach using the operational Landsat DSWE data product, we found that our approach consistently mapped more inundation than DSWE-1 and DSWE-2, the most widely used classes. Although DSWE-4 mapped more inundation than our fusion model in regions B and C, this is likely due to partial surface water pixels in the *Aggressive* class overestimating inundation in areas of dense tree cover, which is a known limitation of this DSWE class [18]. Overall, the substantial differences between our fusion model results and DSWE-1 and DSWE-2 suggest that studies relying on moderate-resolution surface water products for mapping wetland inundation may underestimate total inundated area. Such studies risk oversimplifying hydrologic dynamics by assuming more static inundation extents and missing seasonal fluctuations. This is corroborated by [78], who found that PlanetScope Basemaps (4.77 m resolution) identify significantly more seasonal water than Landsat DSWE, particularly for the DSWE-1 class.

In time series with inconsistent observation quality and gaps, some estimates are inherently more uncertain than others. By using a data fusion approach and quantifying our uncertainty, this approach provides a natural framework to propagate uncertainty into downstream models using Monte Carlo-based methods. This represents a substantial advantage over binary classification maps from individual sensors, which provide no inherent measure of prediction confidence. For example, we can use the inundated data generated here to estimate carbon fluxes using process-based or empirical models. This can be done through an ensemble approach, running a carbon model across multiple inundation scenarios with varying classification thresholds (e.g., 10% to 90%) to quantify the mean and variance of different emissions. Quantifying uncertainty also highlights when and where additional data, such as purchasing high-resolution PlanetScope imagery, could improve map accuracy. This is especially important not only for understanding carbon cycle dynamics but for studies that monitor long-term trends and changes like sea-level rise, saltwater intrusion, or hydrology-driven regime shifts in coastal wetland regions. More information on uncertainty quantification of our results can be found in the Supplementary Material (Fig. A.21, Fig. A.22).

5.3 Comparison to Existing Approaches

Comparing accuracy and methods across studies is difficult because most focus on either satellite classification or data fusion, with few incorporating both for high-resolution wetland inundation mapping. Hondula [8] used high-resolution PlanetScope imagery to map inundation extents in Maryland's Delmarva Peninsula and found that it predicted an order of magnitude more inundation

than the Global Surface Water Explorer [16], a product similar to DSWE. This aligns with our findings that higher spatial and temporal resolution datasets capture significantly more inundation than coarser, operational products. Vanderhoof [26] mapped open and vegetated water across the United States using Sentinel-1 and Sentinel-2 models similar to those in our study. They found strong agreement between sensor models for open water classifications, but weaker agreement when including vegetated water, consistent with our findings that Sentinel-1 maps inundation comparable to Sentinel-2 and PlanetScope in open water and river channels, but shows more discrepancies in vegetated and agricultural areas. They also noted that using the BU3 index at an annual time step may have limited their Sentinel-1 model's ability to capture seasonal or episodic variability in vegetated water, which likely contributed to similar limitations in our model.

Many studies have explored data fusion for large-scale flood mapping, including pixel-, feature-, and decision-level techniques. For comparison, we focus on those targeting wetland dynamics. Both [79] and [80] present variations of the STARFM fusion model for pixel-level fusion, or blending of NDVI and NDWI, respectively, using Landsat and MODIS data. Their post-fusion classification improved inundation mapping at higher spatial (30 m) and temporal (8–16 day) resolutions compared to single-sensor methods. While similar to our approach in enhancing resolution and frequency, STARFM blends the same-day coarse and fine-resolution imagery using rule-based methods and is not applicable to binary classification maps. In contrast, our FRK-based method uses a geostatistical model that accounts for measurement error, spatial correlation, and irregular observation spacing, while being adaptable to non-Gaussian input data [34] also used decision-level fusion with optical and radar data, applying expert rules post-classification in a Southeast Asian region that included wetlands. Similar to our findings, they reported improved mapping accuracy compared to single-sensor approaches, particularly in areas complicated by monsoons, complex topography, and dense vegetation. However, this approach, while computationally efficient, differs from ours in that it does not incorporate a geostatistical framework to generate complete spatiotemporal coverage and quantify uncertainty.

Zammit-Mangion [81] provided an overview of how FRK, along with other kriging variants, has been used to overcome data irregularities in constructing Level 3 products derived from raw remote sensing data. They demonstrated how FRK, as a spatiotemporal interpolator, was applied to predict column-averaged CO_2, highlighting its strengths in handling change-of-support issues, providing complete coverage, and quantifying uncertainty, all of which motivated the use of FRK in our study. Zhou [82] used FRK to fuse data from GOSAT and AIRS to predict near-surface CO_2, showing that FRK improved CO_2 distribution over using single data sources. This supports our findings that FRK-based fusion captured more dynamic inundation patterns and reflected certain spikes and trends in WTD and precipitation data better than if we used an individual sensor like Sentinel-2. Karapetsas [83] applied FRK to fuse soil survey data with Sentinel-2 imagery to predict soil properties, noting that while Sentinel-2 data improved results, the added computational cost was not always justified. This differs from our approach, where additional data sources were used as observations rather than covariates, but we also recognize similar computational challenges when using FRK.

5.4 Limitations & Future Directions

5.4.1 Satellite Sensors & Prediction Models

Monitoring inundation in open and vegetated areas across dense, high-resolution image time series presents significant technical challenges. While all sensors achieved high accuracy, each had limitations in capturing fine-scale inundation dynamics in forested wetlands. As stated previously, Sentinel-1 had a tendency to over-predict inundation in saturated areas like agricultural areas and open fields. In areas with sparse to moderate canopy cover shown in our inundation frequency maps (Fig. 5) , Sentinel-2 and PlanetScope detected some inundation, whereas Sentinel-1 detection was limited in areas of dense canopy cover. Future work could address this limitation by integrating with a process-based or physics-based hydrologic model to inform sub-canopy dynamics. Integrating data from new sensors, such as the L-band NISAR mission [77], may further improve the detection of sub-canopy inundation—a capability that's been previously demonstrated with ALOS PALSAR L-band data [84–86]. Additionally, Ka-band radar from the Surface Water and Ocean Topography (SWOT) mission [87] could enhance this modeling framework by incorporating water-level estimates, which have been successfully predicted in the Florida Everglades wetlands using SWOT data [88]. Future studies could also benefit from using more advanced models, such as deep learning or pre-trained foundation models; however, more work is needed to develop training and validation data sets specifically for wetlands and areas of inundated vegetation to make applying and fine-tuning these models feasible.

5.4.2 Geostatistical Data Fusion with FRK

While FRK successfully integrated heterogeneous satellite sensors into one continuous stream of predictions, we identified two notable limitations to consider in its application to remote sensing data. First, unmasked clouds had a tendency to be misclassified as water and would propagate through the spatiotemporal covariance structure, affecting neighboring predictions. This limitation was most pronounced for Sentinel-2, where the cloud mask created using the s2cloudless algorithm (with 30% probability threshold) occasionally failed. A stricter cloud probability threshold, improved cloud masking, or additional outlier filters in the framework could help improve this. Second, implementing FRK at high resolutions comes with a significant computational cost, especially when preserving fine spatial detail at high resolution due to constructing the spatial covariance matrix in the SRE model. While reducing the number of basis functions could speed up computation, this approach would have produced overly smooth spatial predictions, likely obscuring the small-scale, ephemeral inundation dynamics that were central to our analysis. To address this, we used three-week time windows and divided each fusion region into smaller spatial sections with overlapping subset boundaries. This required fitting separate models for each spatiotemporal subset, followed by post hoc integration of the model outputs. This segmented modeling process deviates from the results that would have been obtained from a single, unified model trained on the full dataset. This underscores the primary limitation we faced in applying FRK, and likely any spatial statistical method to our dataset: the need to balance computational efficiency with spatial resolution. Future work could explore a variable or multi-resolution application of FRK for wetland inundation monitoring, maintaining

high resolution in small, dynamic areas and coarser resolution in larger, more stable regions, and could also extend the current temporally smoothing formulation to a fixed-rank filtering approach to enable near–real-time applications [89].

6 Conclusion

With the ongoing rise in global temperatures coupled with shifts in hydrologic patterns, reducing uncertainties associated with understanding inundation and methane feedbacks in wetland ecosystems is crucial. Meeting this need requires methods to monitor inundation extent at high spatial and temporal resolutions using multiple satellite sensors. Our work provides both a quantitative comparison of the accuracy of different sensors and a demonstrated fusion approach to address this gap, advancing our ability to investigate this and other key science questions related to freshwater resources, biogeochemistry, and shifting water regimes in wetland ecosystems. We found that Sentinel-2 and PlanetScope satellites were the most accurate for detecting vegetated wetland inundation, whereas Sentinel-1 tended to overpredict inundation by responding to saturation in agricultural fields rather than standing water. Using the FRK model, we then generated daily predictions of inundation probability at 6 m resolution with uncertainty quantification across three fusion subsets. Our FRK-based fusion approach captured ephemeral inundation changes, balanced sensor inconsistencies, and detected substantially more inundation than DSWE's high- and moderate-confidence products. Although we likely still missed some inundation beneath dense canopies, this work pushes the limits of current satellite capabilities and offers a framework for fusing heterogeneous binary classification maps, applicable to data sources and problems beyond inundation mapping. Future research can readily build on this framework by integrating datasets that capture sub-canopy inundation, enabling more accurate and higher-resolution tracking of inundation dynamics and their implications in a changing climate.

7 Data & Code Availability

The Python and R scripts used in this study, the trained machine learning models, and the training and validation datasets will be made available at: https://github.com/jen-abrahamson/wetland_hydro_ml/ concurrent with manuscript publication.

8 Funding Sources

This work was supported by NASA through the Future Investigators in NASA Earth and Space Science and Technology (FINESST) Program Grant No. 80NSSC24K0016; and author Jenna N. Abrahamson was also supported by a National Science Foundation Graduate Research Fellowship Grant No. DGE-2137100. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

9 Authorship Contribution Statement

Jenna N. Abrahamson: Conceptualization, Data Curation, Formal Analysis, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Software, Validation, Visualization, Writing - Original Draft, Writing - Review and Editing. **Josh M. Gray**: Conceptualization, Investigation, Methodology, Supervision, Writing - Review and Editing. **Mirela G. Tulbure**: Methodology, Writing - Review and Editing. **Erin M. Schliep**: Methodology, Writing - Review and Editing.

10 Declaration of Competing Interest

The authors declare they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

11 Acknowledgments

We would like to acknowledge Dr. Justine Neville for providing the stilling pipe data used in this study. We also thank Dr. John King and Dr. Maricar Aguilos for their valuable insights on inundation dynamics in the study area and at the two flux tower sites, and Dr. Andrew Zammit-Mangion for his assistance with questions related to the FRK package. During the preparation of this work the lead author Jenna Abrahamson used ChatGPT (OpenAI) to improve prose and grammar of the manuscript. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

1. Schlesinger WH and Bernhardt ES. Biogeochemistry: An Analysis of Global Change, 3rd Edition. 3rd edition. Amsterdam Boston: Academic Press, 2013.
2. Moomaw WR, Chmura GL, Davies GT, et al. Wetlands In a Changing Climate: Science, Policy and Management. Wetlands 2018;38:183–205.
3. Mitra B, Minick K, Miao G, et al. Spectral evidence for substrate availability rather than environmental control of methane emissions from a coastal forested wetland. Agricultural and Forest Meteorology 2020;291:108062.
4. Salimi S, Almuktar SAAAN, and Scholz M. Impact of climate change on wetland ecosystems: A critical review of experimental wetlands. Journal of Environmental Management 2021;286:112160.
5. Cui S, Liu P, Guo H, et al. Wetland hydrological dynamics and methane emissions. Communications Earth & Environment 2024;5:1–17.
6. Saunois M, Stavert AR, Poulter B, et al. The Global Methane Budget 2000–2017. Earth System Science Data 2020;12:1561–623.
7. Melton JR, Wania R, Hodson EL, et al. Present state of global wetland extent and wetland methane modelling: conclusions from a model inter-comparison project (WETCHIMP). Biogeosciences 2013;10:753–88.
8. Hondula KL, DeVries B, Jones CN, and Palmer MA. Effects of Using High Resolution Satellite-Based Inundation Time Series to Estimate Methane Fluxes From Forested Wetlands. Geophysical Research Letters 2021;48:e2021GL092556.
9. Thornton BF, Wik M, and Crill PM. Double-counting challenges the accuracy of high-latitude methane inventories. Geophysical Research Letters 2016;43:12, 569–12, 577.
10. Dong B, Peng S, Liu G, et al. Underestimation of Methane Emissions From the Sudd Wetland: Unraveling the Impact of Wetland Extent Dynamics. Geophysical Research Letters 2024;51:e2024GL110690.
11. Zhang Z, Fluet-Chouinard E, Jensen K, et al. Development of the global dataset of Wetland Area and Dynamics for Methane Modeling (WAD2M). Earth System Science Data 2021;13:2001–23.
12. Du L, McCarty GW, Zhang X, et al. Mapping Forested Wetland Inundation in the Delmarva Peninsula, USA Using Deep Convolutional Neural Networks. Remote Sensing 2020;12:644.
13. Huang C, Smith LC, Kyzivat ED, Fayne J, Ming Y, and Spence C. Tracking transient boreal wetland inundation with Sentinel-1 SAR: Peace-Athabasca Delta, Alberta and Yukon Flats, Alaska. Giscience & Remote Sensing 2022;59:1767–92.
14. Woodcock CE, Loveland TR, Herold M, and Bauer ME. Transitioning from change detection to monitoring with remote sensing: A paradigm shift. Remote Sensing of Environment. Time Series Analysis with High Spatial Resolution Imagery 2020;238:111558.
15. Radeloff VC, Roy DP, Wulder MA, et al. Need and vision for global medium-resolution Landsat and Sentinel-2 data products. Remote Sensing of Environment 2024;300:113918.

16. Pekel JF, Cottam A, Gorelick N, and Belward AS. High-resolution mapping of global surface water and its long-term changes. Nature 2016;540:418–22.

17. Jones JW. Efficient Wetland Surface Water Detection and Monitoring via Landsat: Comparison with in situ Data from the Everglades Depth Estimation Network. Remote Sensing 2015;7:12503–38.

18. Jones JW. Improved Automated Detection of Subpixel-Scale Inundation—Revised Dynamic Surface Water Extent (DSWE) Partial Surface Water Tests. Remote Sensing 2019;11:374.

19. Cooley SW, Smith LC, Stepan L, and Mascaro J. Tracking Dynamic Northern Surface Water Changes with High-Frequency Planet CubeSat Imagery. Remote Sensing 2017;9:1306.

20. Perin V, Tulbure MG, Gaines MD, Reba ML, and Yaeger MA. A multi-sensor satellite imagery approach to monitor on-farm reservoirs. Remote Sensing of Environment 2022;270:112796.

21. Valman SJ, Boyd DS, Carbonneau PE, Johnson MF, and Dugdale SJ. An AI approach to operationalise global daily PlanetScope satellite imagery for river water masking. Remote Sensing of Environment 2024;301:113932.

22. Claverie M, Ju J, Masek JG, et al. The Harmonized Landsat and Sentinel-2 surface reflectance data set. Remote Sensing of Environment 2018;219:145–61.

23. Tulbure MG, Broich M, Perin V, et al. Can we detect more ephemeral floods with higher density harmonized Landsat Sentinel 2 data compared to Landsat 8 alone? ISPRS Journal of Photogrammetry and Remote Sensing 2022;185:232–46.

24. Brinkhoff J, Houborg R, and Dunn BW. Rice ponding date detection in Australia using Sentinel-2 and Planet Fusion imagery. Agricultural Water Management 2022;273:107907.

25. DeVries B, Huang C, Armston J, Huang W, Jones JW, and Lang MW. Rapid and robust monitoring of flood events using Sentinel-1 and Landsat data on the Google Earth Engine. Remote Sensing of Environment 2020;240:111664.

26. Vanderhoof MK, Alexander L, Christensen J, Solvik K, Nieuwlandt P, and Sagehorn M. High-frequency time series comparison of Sentinel-1 and Sentinel-2 satellites for mapping open and vegetated water across the United States (2017–2021). Remote Sensing of Environment 2023;288:113498.

27. Manavalan R. Review of synthetic aperture radar frequency, polarization, and incidence angle data for mapping the inundated regions. Journal of Applied Remote Sensing 2018;12.

28. Townsend PA. Relationships between forest structure and the detection of flood inundation in forested wetlands using C-band SAR. International Journal of Remote Sensing 2002;23:443–60.

29. Zhu X, Cai F, Tian J, and Williams TKA. Spatiotemporal Fusion of Multisource Remote Sensing Data: Literature Survey, Taxonomy, Principles, Applications, and Future Directions. Remote Sensing 2018;10:527.

30. Himeur Y, Rimal B, Tiwary A, and Amira A. Using artificial intelligence and data fusion for environmental monitoring: A review and future perspectives. Information Fusion 2022;86–87:44–75.

31. Zhang J. Multi-source remote sensing data fusion: status and trends. International Journal of Image and Data Fusion 2010;1:5–24.

32. Kulkarni SC and Rege PP. Pixel level fusion techniques for SAR and optical images: A review. Information Fusion 2020;59:13–29.

33. Hu B, Brown G, Stirling C, and Wang J. Land Cover Classification of Subarctic Wetlands Using Multisource Remotely Sensed Data. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 2024;17:5073–92.

34. Ahmad SK, Hossain F, Eldardiry H, and Pavelsky TM. A Fusion Approach for Water Area Classification Using Visible, Near Infrared and Synthetic Aperture Radar for South Asian Conditions. IEEE Transactions on Geoscience and Remote Sensing 2020;58:2471–80.

35. Wendl C, Le Bris A, Chehata N, Puissant A, and Postadjian T. Decision Fusion of Spot6 And Multitemporal Sentinel2 Images For Urban Area Detection. In: 2018:1734–7. DOI: 10.1109/IGARSS.2018.8517476.

36. Bioresita F, Puissant A, Stumpf A, and Malet JP. Fusion of Sentinel-1 and Sentinel-2 image time series for permanent and temporary surface water mapping. International Journal of Remote Sensing 2019;40:9026–49.

37. Shen M, Zhang F, and Wu CY. Flood Inundation Extraction Based on Decision-Level Data Fusion: A Case in Peru. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences 2022;X-3-W1-2022:133–40.

38. Shuai S, Zhang Z, Zhang T, et al. Innovative Decision Fusion for Accurate Crop/Vegetation Classification with Multiple Classifiers and Multisource Remote Sensing Data. Remote Sensing 2024;16:1579.

39. Wikle CK, Milliff RF, Nychka D, and Berliner LM. Spatiotemporal Hierarchical Bayesian Modeling Tropical Ocean Surface Winds. Journal of the American Statistical Association 2001;96:382–97.

40. Wikle CK and Berliner LM. Combining Information Across Spatial Scales. Technometrics 2005;47:80–91.

41. Foley KM and Fuentes M. A statistical framework to combine multivariate spatial data and physical models for Hurricane surface wind prediction. Journal of Agricultural, Biological, and Environmental Statistics 2008;13:37–59.

42. Nguyen H, Katzfuss M, Cressie N, and Braverman A. Spatio-Temporal Data Fusion for Very Large Remote Sensing Datasets. Technometrics 2014;56:174–85.

43. Nguyen H, Cressie N, and Braverman A. Multivariate Spatial Data Fusion for Very Large Remote Sensing Datasets. Remote Sensing 2017;9:142.

44. Nguyen H, Cressie N, and Braverman A. Spatial Statistical Data Fusion for Remote Sensing Applications. Journal of the American Statistical Association 2012;107:1004–18.

45. Cressie N, Shi T, and Kang EL. Fixed Rank Filtering for Spatio-Temporal Data. Journal of Computational and Graphical Statistics 2010;19:724–45.

46. Cressie N and Johannesson G. Fixed Rank Kriging for Very Large Spatial Data Sets. Journal of the Royal Statistical Society Series B: Statistical Methodology 2008;70:209–26.

47. Bhattachan A, Jurjonas M, Moody A, et al. Sea level rise impacts on rural coastal social-ecological systems and the implications for decision making. Environmental Science & Policy 2018;90:122–34.

48. Swanquarter Ferry Weather Station, Hyde County, NC. NOAA Online Weather Data. https://www.weather.gov/wrh/Climate?wfo=mhx. Accessed: 12 November 2024.

49. Poulter B, Goodall JL, and Halpin PN. Applications of network analysis for adaptive management of artificial drainage systems in landscapes vulnerable to sea level rise. Journal of Hydrology 2008;357:207–17.

50. Miao G, Noormets A, Domec JC, et al. The effect of water table fluctuation on soil respiration in a lower coastal plain forested wetland in the southeastern U.S. Journal of Geophysical Research: Biogeosciences 2013;118:1748–62.

51. Miao G, Noormets A, Domec JC, et al. Hydrology and microtopography control carbon dynamics in wetlands: Implications in partitioning ecosystem respiration in a coastal plain forested wetland. Agricultural and Forest Meteorology 2017;247:343–55.

52. U.S. Fish and Wildlife Service. National Wetlands Inventory. www.fws.gov/\breakprogram/national-wetlands-inventory/wetlands-mapper. Accessed: 14 November 2024.

53. Dewitz J. National Land Cover Database (NLCD) 2019 Products. 2021. DOI: 10.5066/P9KZCM54. URL: https://www.sciencebase.gov/catalog/item/5f21cef582cef313ed940043.

54. Mullissa A, Vollrath A, Odongo-Braun C, et al. Sentinel-1 SAR Backscatter Analysis Ready Data Preparation in Google Earth Engine. Remote Sensing 2021;13:1954.

55. Lee JS, Grunes M, and Grandi G de. Polarimetric SAR speckle filtering and its implication for classification. IEEE Transactions on Geoscience and Remote Sensing 1999;37:2363–73.

56. Wilm U. Sen2Cor configuration and user manual - ref. S2-PDGS-MPC-L2A-SUM-V2.9. 2020.

57. Skakun S, Wevers J, Brockmann C, et al. Cloud Mask Intercomparison eXercise (CMIX): An evaluation of cloud masking algorithms for Landsat 8 and Sentinel-2. Remote Sensing of Environment 2022;274:112990.

58. Kington J and Collison A. Scene level normalization and harmonization of planet dove imagery. Planet Labs Inc.: San Francisco, CA, USA 2022.

59. Planet Labs. Planet Imagery Product Specifications. 2023.

60. Olofsson P, Foody GM, Herold M, Stehman SV, Woodcock CE, and Wulder MA. Good practices for estimating area and assessing accuracy of land change. Remote Sensing of Environment 2014;148:42–57.

61. Stehman SV and Foody GM. Key issues in rigorous accuracy assessment of land cover products. Remote Sensing of Environment 2019;231:111199.

62. Breiman L. Random Forests. Machine Learning 2001;45:5–32.

63. Chen T and Guestrin C. XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. arXiv:1603.02754 [cs]. 2016:785–94. DOI: 10.1145/2939672.2939785. URL: http://arxiv.org/abs/1603.02754.

64. Tucker CJ. Red and photographic infrared linear combinations for monitoring vegetation. Remote Sensing of Environment 1979;8:127–50.

65. McFeeters SK. The use of the Normalized Difference Water Index (NDWI) in the delineation of open water features. International Journal of Remote Sensing 1996;17:1425–32.

66. Xu H. Modification of normalised difference water index (NDWI) to enhance open water features in remotely sensed imagery. International Journal of Remote Sensing 2006;27:3025–33.

67. Feyisa GL, Meilby H, Fensholt R, and Proud SR. Automated Water Extraction Index: A new technique for surface water mapping using Landsat imagery. Remote Sensing of Environment 2014;140:23–35.

68. Gautam V, Murugan P, and Annadurai M. A New Three Band Index for Identifying Urban Areas using Satellite Images. In: 2017.

69. Beven KJ and Kirkby MJ. A physically based, variable contributing area model of basin hydrology / Un modèle à base physique de zone d'appel variable de l'hydrologie du bassin versant. Hydrological Sciences Bulletin 1979;24:43–69.

70. Thornton M, Shrestha R, Wei Y, Thornton P, and Kao SC. Daymet: Daily Surface Weather Data on a 1-km Grid for North America, Version 4 R1. 2022. DOI: 10.3334/ORNLDAAC/2129. URL: https://daac.ornl.gov/cgi-bin/dsviewer.pl?ds_id=2129.

71. Housman I, Bender S, Schleeweis K, Heyer J, Ruefenacht B, and Megown K. National Land Cover Database Tree Canopy Cover Methods. 2023.

72. Tulbure MG, Broich M, Stehman SV, and Kommareddy A. Surface water extent dynamics from three decades of seasonally continuous Landsat time series at subcontinental scale in a semi-arid region. Remote Sensing of Environment 2016;178:142–57.

73. Pedregosa F, Varoquaux G, Gramfort A, et al. Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research 2011;12:2825–30.

74. Rocklin M. Dask: Parallel Computation with Blocked algorithms and Task Scheduling. Proceedings of the 14th Python in Science Conference 2015;130:136–43.

75. Zammit-Mangion A and Cressie N. Introduction to Fixed Rank Kriging: The R package. Journal of Statistical Software 2024.

76. Sainsbury-Dale M, Zammit-Mangion A, and Cressie N. Modeling Big, Heterogeneous, Non-Gaussian Spatial and Spatio-Temporal Data Using FRK. Journal of Statistical Software 2024;108:1–39.

77. Ferdowsi B, Bhanu M, Rao C, et al. NASA-ISRO Synthetic Aperture Radar (NISAR): The Last Steps to Launch. In: *2024 IEEE Aerospace Conference*. 2024:1–10. DOI: 10.1109/AERO58975.2024.10520949.

78. Gaines MD, Tulbure MG, Perin V, et al. Impact of Spatial Scale on Optical Earth Observation-Derived Seasonal Surface Water Extents. Geophysical Research Letters 2026;53.

79. Chen B, Chen L, Huang B, Michishita R, and Xu B. Dynamic monitoring of the Poyang Lake wetland by integrating Landsat and MODIS observations. ISPRS Journal of Photogrammetry and Remote Sensing 2018;139:75–87.

80. Ye X, Wu J, Li X, Li Y, Zhang Q, and Xu CY. Multi-source remote sensing data and image fusion technology reveal significant spatiotemporal heterogeneity of inundation dynamics in a typical large floodplain lake system. Journal of Hydrology: Regional Studies 2023;50:101541.

81. Zammit-Mangion A, Cressie N, and Shumack C. On Statistical Approaches to Generate Level 3 Products from Satellite Remote Sensing Retrievals. Remote Sensing 2018;10:155.

82. Zhou C, Shi R, and Gao W. Data fusion of CO2 retrieved from GOSAT and AIRS using regression analysis and fixed rank kriging. In: *Remote Sensing and Modeling of Ecosystems for Sustainability XII*. Vol. 9610. SPIE, 2015:315–23. DOI: 10.1117/12.2187493. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9610/96101A/Data-fusion-of-CO2-retrieved-from-GOSAT-and-AIRS-using/10.1117/12.2187493.full.

83. Karapetsas N, Alexandridis TK, Bilas G, et al. Mapping Soil Properties with Fixed Rank Kriging of Proximally Sensed Soil Data Fused with Sentinel-2 Biophysical Parameter. Remote Sensing 2022;14:1639.

84. Hübinger C, Fluet-Chouinard E, Escobar D, and Jaramillo F. Exploring the potential of using L-band InSAR for mapping flooded vegetation in tropical wetlands. Remote Sensing of Environment 2026;332:115086.

85. Oakes G, Hardy A, Bunting P, and Rosenqvist A. RadWet-L: A Novel Approach for Mapping of Inundation Dynamics of Forested Wetlands Using ALOS-2 PALSAR-2 L-Band Radar Imagery. Remote Sensing 2024;16:2078.

86. Chapman B, McDonald K, Shimada M, Rosenqvist A, Schroeder R, and Hess L. Mapping Regional Inundation with Spaceborne L-Band SAR. Remote Sensing 2015;7:5440–70.

87. Biancamaria S, Lettenmaier DP, and Pavelsky TM. The SWOT Mission and Its Capabilities for Land Hydrology. Surveys in Geophysics 2016;37. Company: Springer Distributor: Springer Institution: Springer Label: Springer number: 2 publisher: Springer Netherlands:307–37.

88. Kica S, Pavelsky TM, Fayne JV, and Williams BA. SWOT Water Surface Elevation in Herbaceous Wetlands of Florida's Everglades. Geophysical Research Letters 2025;52:e2025GL114956.

89. Cressie N, Shi T, and Kang EL. Fixed Rank Filtering for Spatio-Temporal Data. Journal of Computational and Graphical Statistics 2010;19:724–45.

90. Kristensen K, Nielsen A, Berg CW, Skaug H, and Bell B. TMB: Automatic Differentiation and Laplace Approximation. Journal of Statistical Software 2016;70. arXiv:1509.00660 [stat].

Multi-Sensor Monitoring of Wetland Inundation Using a Machine Learning–Data Fusion Framework

Supplementary Material

Jenna N. Abrahamson [1], Josh M. Gray [1,2], Mirela G. Tulbure [1,2], and Erin M. Schliep [3]

[1] Center for Geospatial Analytics, North Carolina State University, Raleigh, NC, USA
[2] Forestry & Environmental Resources, North Carolina State University, Raleigh, NC, USA
[3] Department of Statistics, North Carolina State University, Raleigh, NC, USA

National Wetland Inventory Data

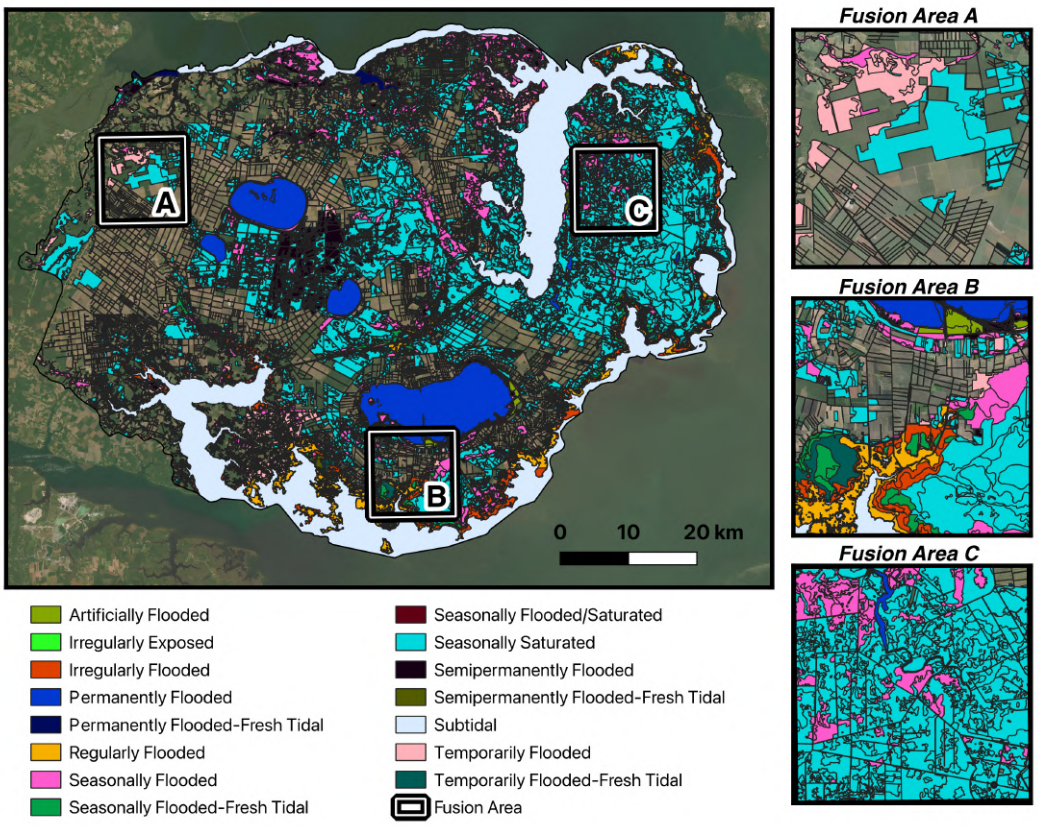

Figure A.11: Map of NWI classified water regimes for the study area with insets of the three fusion areas.

Tuning the Machine Learning Models

Both RF and XGB were extensively tuned to find the optimal hyperparameter set for each model. For RF, this included the maximum depth of each tree (`max_depth`), the number of features considered at each split (`max_features`), and the number of trees (`n_estimators`). For XGB, this included the maximum depth of each tree (`max_depth`), how much each tree contributes to the final prediction (`learning_rate`), the minimum sum of instance weights needed to create a new split (`min_child_weight`), the fraction of training data used for growing each tree (`subsample`), the fraction of features sampled for each tree (`colsample_bytree`), the minimum loss reduction required to make a split (`gamma`), and number of trees (`n_estimators`). The resulting best estimators and the range of hyperparameters tested are shown for all six trained models in Table A.8.

Table A.8: Tables showing the final, tuned hyperparameters for the machine learning models. The column name displays the name of the hyperparameter and the range of values tested using scikit learn's `GridSearchCV` function.

RF Models	**n_estimators [100, 200, 300, 500]**	**max_depth [3, 5, 7, 10, 12]**	**max_features [0.1, 0.3, 0.5, 0.7, 0.9]**
Sentinel-1	500	12	0.1
Sentinel-2	500	12	0.1
PlanetScope	200	12	0.1

XGB Models	**n_estimators [100, 200, 500]**	**max_depth [3, 5, 7, 10, 12]**	**learning_rate [0.05, 0.1, 0.15, 0.2]**
Sentinel-1	100	12	0.05
Sentinel-2	200	7	0.15
PlanetScope	200	7	0.05
min_child_weight [1, 3, 5]	**subsample [0.5, 0.7, 0.9, 1]**	**colsample_bytree [0.5, 0.7, 0.9, 1]**	**gamma [0, 2, 3, 4, 5]**
1	0.5	0.5	0
3	1	0.5	5
1	1	0.5	3

11.0.1 Spatiotemporal Hierarchical Model Used in FRK

The FRK v2 method used in the package is based on a spatiotemporal hierarchical model, which is summarized as follows:

$$Z_j|\boldsymbol{\mu}_Z,\ \psi \overset{\text{ind}}{\sim} \text{EF}(\mu_{Z_j}, \psi); \quad j = 1, ..., m, \tag{3}$$

$$\boldsymbol{\mu}_Z = \boldsymbol{C}_Z\boldsymbol{\mu}, \tag{4}$$

$$g(\boldsymbol{\mu}) = \boldsymbol{Y}, \tag{5}$$

$$\boldsymbol{Y} = \boldsymbol{T}\boldsymbol{\alpha} + \boldsymbol{S}\boldsymbol{\eta} + \boldsymbol{\xi}, \tag{6}$$

$$\boldsymbol{\eta}|\boldsymbol{\upsilon} \sim \text{Gau}(\boldsymbol{0}, \boldsymbol{Q}^{-1}), \tag{7}$$

$$\boldsymbol{\xi}|\sigma^2_\xi \sim \text{Gau}(\boldsymbol{0}, \sigma^2_\xi\boldsymbol{V}) \tag{8}$$

Equation 3 describes our data layer, which uses an exponential family probability distribution for each satellite pixel classification j of the data vector $\mathbf{Z}$, with dispersion parameter ψ. Since we are dealing with binary data of either 0 = *Not Inundated* or 1 = *Inundated*, our data layer takes the form of a binomial distribution with non-negative integers, where ψ is equal to 1 following Bin(1,π) and forming a Bernoulli distribution. Next, Equation 4 uses the matrix of weights $\boldsymbol{C}_Z$ to aggregate the BAU-level process $\boldsymbol{\mu}$ over the observation supports (i.e., the native Sentinel-1/2 or PlanetScope pixel grids). Building on the binary nature of our data, Equation 5 uses a link function $g(\cdot)$ to model $Y(\cdot)$ as a transformation of the mean process, $\mu(\cdot)$, which in our case we chose to model using a probit link function.

Equation 6 is the process model, where the latent spatial process is modeled using the SRE model. Here, $\boldsymbol{Y}$ is the spatial process evaluated over N BAUs, $\boldsymbol{T}$ and $\boldsymbol{S}$ are known design matrices corresponding to spatially referenced covariates and constructed basis functions, $\boldsymbol{\alpha}$ is a vector of fixed effects, $\boldsymbol{\eta}$ is a vector of random coefficients associated with the basis functions, and $\boldsymbol{\xi}$ is a fine-scale-variation random process. Equations 7 and 8 place distributions on the parameters $\boldsymbol{\eta}$ and $\boldsymbol{\xi}$. The parameter $\boldsymbol{\eta}$ is modeled as a mean-zero multivariate Gaussian random vector with precision matrix $\boldsymbol{Q}^{-1}$ and dependent on an unknown parameter vector $\boldsymbol{\upsilon}$. Finally, the parameter $\boldsymbol{\xi}$ is also modeled as a mean zero multivariate Gaussian random vector whose variance is modeled by $\sigma^2_\xi\boldsymbol{V}$, where $\boldsymbol{V}$ is a known, positive definite diagonal matrix that can be set according to fine-scale problem-specific information; however, in the absence of these data we set $\boldsymbol{V}$ to the identity matrix, $\boldsymbol{I}$. The parameters $\boldsymbol{\alpha}$, $\boldsymbol{\eta}$, and $\boldsymbol{\xi}$ are empirically estimated by computing the observed-data log-likelihood function along with its derivatives using a Laplace approximation via the R package TMB [90]. For further details on model fitting and estimation procedures, please refer to [76].

RF vs XGB Comparison: Accuracy Metrics Not Listed in Main Manuscript

The results for predicting the binary classes of *Surface Inundation* versus *Dry Land* showed that RF performed best for Sentinel-2 and PlanetScope sensors, while XGB performed the best for the Sentinel-1. However, the differences in accuracy metrics between the two models were generally small and often had overlapping confidence intervals. These top-performing algorithms were applied across five years of imagery (2017–2021) to produce over 1,000 classification maps. Tables A.8 and A.9 below report the accuracy metrics generated from this comparison.

Table A.9: Comparison of classification accuracy metrics of RF and XGB for combined *Surface Inundation* class calculated across time for Sentinel-1, Sentinel-2, and PlanetScope sensors. Top-performing models are shown in bold font, 95% confidence intervals are shown in parentheses.

Sensor	Model	Overall Accuracy	User's Accuracy	Producer's Accuracy
Sentinel-1	Random Forest	0.832 (±0.016)	0.875 (±0.026)	0.749 (±0.021)
	XGBoost	**0.859 (±0.015)**	**0.865 (±0.025)**	**0.800 (±0.022)**
Sentinel-2	**Random Forest**	**0.949 (±0.010)**	**0.972 (±0.012)**	**0.907 (±0.017)**
	XGBoost	0.936 (±0.011)	0.962 (±0.015)	0.888 (±0.018)
PlanetScope	**Random Forest**	**0.926 (±0.011)**	**0.959 (±0.015)**	**0.870 (±0.019)**
	XGBoost	0.922 (±0.012)	0.948 (±0.016)	0.869 (±0.019)

Table A.10: Tables showing the error matrix of classes for each satellite sensor expressed in terms of proportion of area and total area, along with associated user's, producer's, and overall accuracies calculated using all available validation pixels pooled across image dates. 95% confidence intervals are shown in parentheses. These results are for multi-class classification from the RF versus XGB model comparison. These models did not end up being used for classification.

Sentinel-1 Random Forest

		Mapped			
		Dry Land	Open Water	Inundated Vegetation	*Total*
Reference	Dry Land	0.483	0.002	0.115	0.600
	Open Water	0.001	0.196	0.003	0.200
	Inundated Vegetation	0.066	0.022	0.112	0.200
	Total (Area Proportion)	0.549 (±0.018)	0.220(±0.009)	0.230 (±0.018)	1.000
	Total (Area in m^2)	109,887 (±3,516)	44,067 (±1,773)	46,046 (±3,634)	200,000
	User's Accuracy:	0.805 (±0.021)	0.981 (±0.014)	0.558 (±0.065)	
	Producer's Accuracy:	0.879 (±0.020)	0.891 (±0.034)	0.485 (±0.040)	
	Overall Accuracy:	0.791 (±0.018)			

Sentinel-2 XGBoost

		Mapped			
		Dry Land	Open Water	Inundated Vegetation	*Total*
Reference	Dry Land	0.552	0.001	0.047	0.600
	Open Water	0.000	00.200	0.000	0.200
	Inundated Vegetation	0.014	0.015	0.171	0.200
	Total (Area Proportion)	0.565 (±0.010)	0.216 (±0.006)	0.219 (±0.012)	1.000
	Total (Area in m^2)	113,095 (±2,084)	43,176 (±1,137)	43,729 (±2,308)	200,000
	User's Accuracy:	0.919 (±0.015)	1.000 (±0.000)	0.856 (±0.037)	
	Producer's Accuracy:	0.976 (±0.009)	0.926 (±0.024)	0.783 (±0.033)	
	Overall Accuracy:	0.923 (±0.012)			

PlanetScope XGBoost

		Mapped			
		Dry Land	Open Water	Inundated Vegetation	*Total*
Reference	Dry Land	0.543	0.002	0.055	0.600
	Open Water	0.000	0.189	0.011	0.200
	Inundated Vegetation	0.019	0.041	0.139	0.200
	Total (Area Proportion)	0.562 (±0.011)	0.233 (±0.010)	0.205 (±0.014)	1.000
	Total (Area in m^2)	10,124 (±204)	4,194 (±175)	3,681 (±253)	200,000
	User's Accuracy:	0.905 (±0.016)	0.946 (±0.024)	0.700 (±0.046)	
	Producer's Accuracy:	0.965 (±0.010)	0.812 (±0.030)	0.680 (±0.038)	
	Overall Accuracy:	0.871 (±0.014)			

Binary Accuracy Metrics for Final S1, S2, and PS Models

Table A.11: Tables showing the error matrix of classes for each satellite sensor expressed in terms of proportion of area and total area along with associated user's, producer's, and overall accuracies calculated using all available validation pixels pooled across image dates. 95% confidence intervals are shown in parentheses. These results are for binary classification of *Dry Land* and *Surface Inundation* for the final models used.

Sentinel-1

		Mapped		
		Dry Land	Surface Inundation	*Total*
Reference	Dry Land	0.513	0.087	0.6
	Surface Inundation	0.054	0.346	0.4
	Total (Area Proportion)	0.568	0.432	1.0
	Total (Area in m^2)	113,515 (±3,058)	86,485 (±3,058)	
	User's Accuracy:	0.856 (±0.019)	0.865 (±0.025)	
	Producer's Accuracy:	0.905 (±0.016)	0.800 (±0.022)	
	Overall Accuracy:	0.859 (±0.015)		

Sentinel-2

		Mapped		
		Dry Land	Surface Inundation	*Total*
Reference	Dry Land	0.560	0.040	0.6
	Surface Inundation	0.011	0.389	0.4
	Total (Area Proportion)	0.571	0.429	1.0
	Total (Area in m^2)	114,299 (±1,910)	85,702 (±1,910)	
	User's Accuracy:	0.933 (±0.014)	0.972 (±0.012)	
	Producer's Accuracy:	0.980 (±0.008)	0.907 (±0.017)	
	Overall Accuracy:	0.949 (±0.010)		

PlanetScope

		Mapped		
		Dry Land	Surface Inundation	*Total*
Reference	Dry Land	0.543	0.057	0.6
	Surface Inundation	0.016	0.384	0.4
	Total (Area Proportion)	0.559	0.441	1.0
	Total (Area in m^2)	10,062 (±203)	7,938 (±203)	
	User's Accuracy:	0.904 (±0.016)	0.959 (±0.015)	
	Producer's Accuracy:	0.971 (±0.010)	0.870 (±0.019)	
	Overall Accuracy:	0.926 (±0.011)		

Accuracy Metrics for Individual Dates

We also evaluated model performance on individual dates, comparing results with DSWE for the 04/11/2018 and 11/17/2020 image dates (Fig. A.12). For context, the reported mean overall agreement rate for DSWE to in-situ data is 0.79 [17] and the mean overall accuracy for partial surface water pixels in DSWE is 0.77 [18]. For the two dates where DSWE validation was conducted, our models consistently showed higher overall and producer's accuracies than DSWE-1, DSWE-2, DSWE-3, and DSWE-4. For user's accuracy, the DSWE categories outperformed our sensor-based models. DSWE-1 and DSWE-2 both achieved perfect scores (1.0) on each date, while DSWE-3 followed closely with values ranging from 0.98 to 1.0. These values reflect high precision, which indicates they nearly always detect true surface water with few false positives. In contrast, DSWE-4 showed lower user's accuracy (0.65–0.66) but had producer's accuracy values (0.68–0.72) that were more in line with those of the DSWE-1 and DSWE-2 categories. Meanwhile, DSWE-3 had producer's accuracy values (0.76–0.77) that more closely matched those of our individual sensor models. Our Sentinel-2 and PlanetScope sensor models achieve similar precision as DSWE-1 and DSWE-2, while all of our sensor models outperform all DSWE categories in recall, capturing more true inundation events.

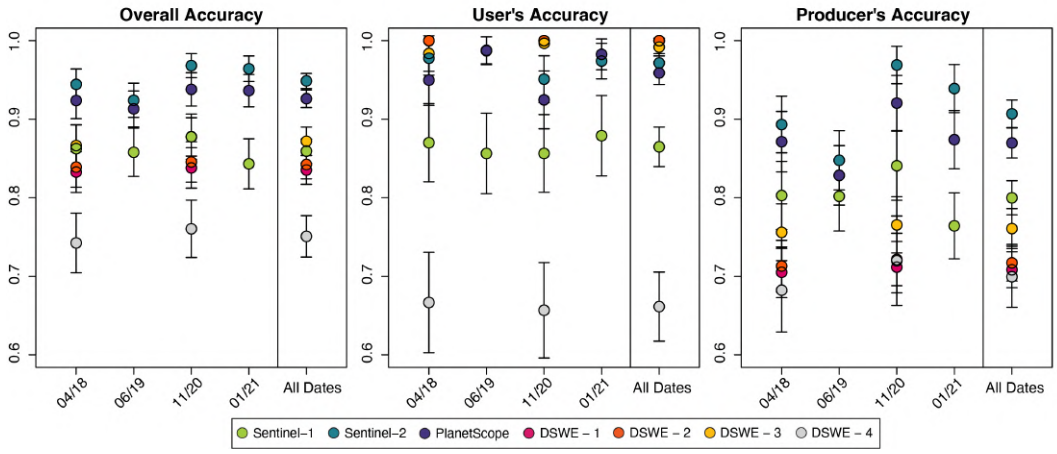

Figure A.12: Plots highlighting the differences in overall (left), producer's (middle), and user's (right) accuracy metrics for binary classification of *Dry Land* versus *Surface Inundation* across sensors and dates. Includes comparison to DSWE-1, DSWE-2, and DSWE-4 results on dates where a DSWE image was available. Error bars represent the standard errors calculated for overall, producer's and user's accuracies.

Feature Importance

When comparing the feature importance of each of the models (Fig. A.13), the Sentinel-2 model's top 10 most important predictors did not contain any terrain-derived or precipitation variables, suggesting this model can leverage its spectral resolution and availability of both the NIR and SWIR wavelengths to capture inundation. In contrast, slope, TWI, and antecedent precipitation were reported in the top 10 most important features for PlanetScope, suggesting the inclusion of these features could aid in the lack of a SWIR band on PlanetScope's CubeSats for improving inundation classifications. Similarly, TWI emerged as the most important feature in the Sentinel-1 model, with slope, canopy cover, and the BU3 annual feature also ranking among the top ten. This is consistent with findings from [26], who also identified BU3 and slope as key predictors in their Sentinel-1 model.

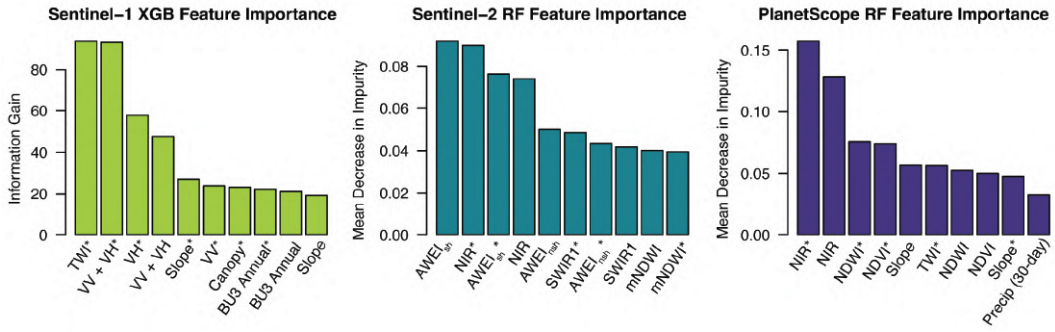

Figure A.13: Bar charts depicting the top 10 most important features as determined by information gain or mean decrease in Gini impurity for the Sentinel-1 XGB model (left), the Sentinel-2 RF model (middle), and the PlanetScope RF model (right). Features with an asterisk (*) next to them denote that a 3 x 3 focal mean was taken of that feature.

Land Cover Analysis

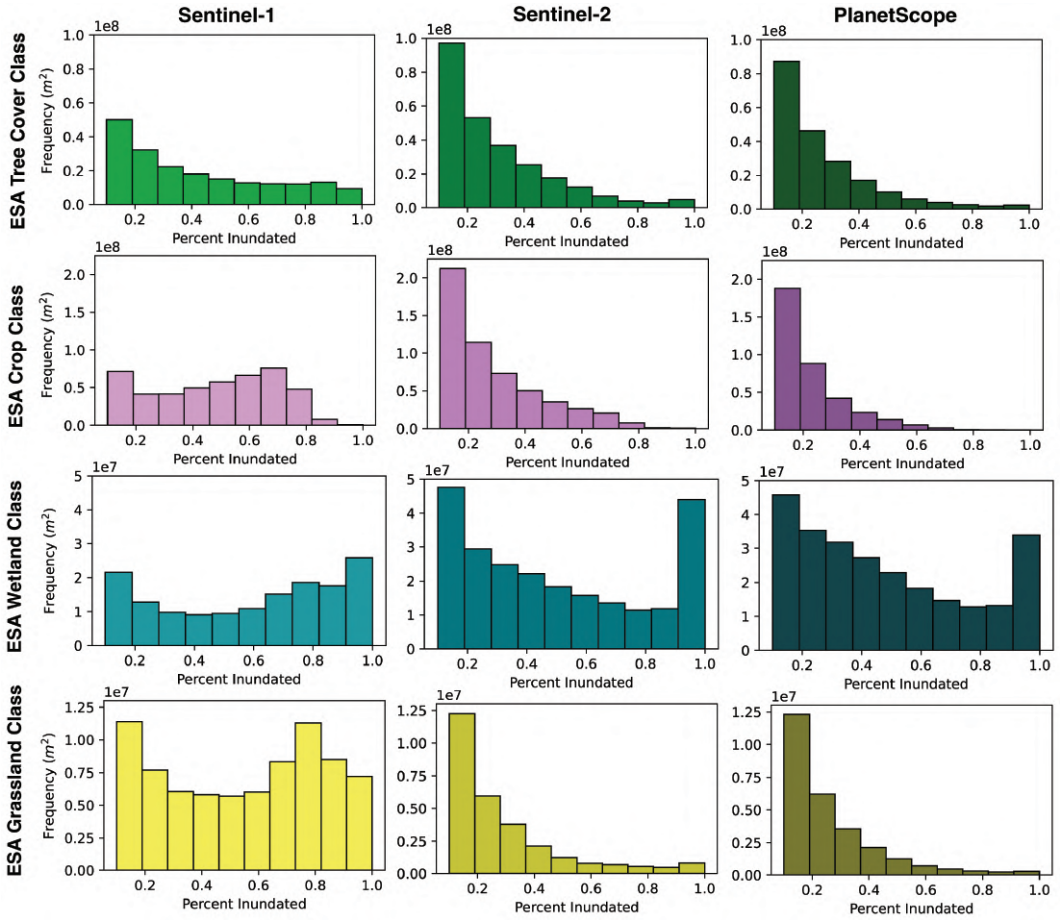

Figure A.14: Histograms showing the pixel frequency of different percent inundated values obtained from the inundation frequency maps for each sensor grouped by 2021 ESA WorldCover classifications.

Spatial Detials Mapped

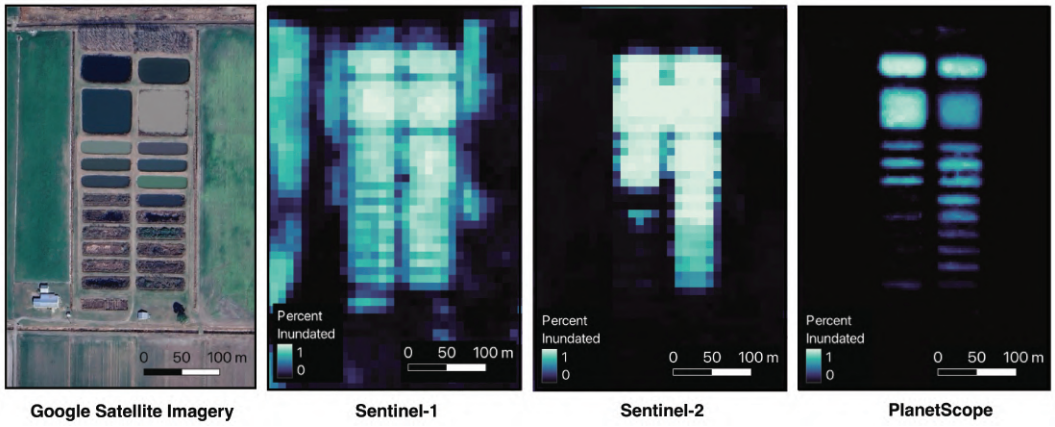

Figure A.15: Maps showing a true color Google Earth satellite image of small farm reservoirs (left) compared to the percent inundated map generated for Sentinel-1, Sentinel-2, and PlanetScope models.

Satellite Predicted Time Series

To track changes in inundation over time, we summarized total *Inundated Vegetation* as time series using maps with at least 80% spatial coverage within a 6 km buffer surrounding each flux tower (Fig. A.16). This buffer size ensured consistency with the regions selected for data fusion and allowed us to include as many observations as possible without gaps in spatial coverage limiting the analysis. The time series of mapped inundated area for US-NC4 shows a strong seasonal trend, increasing with precipitation and coinciding with leaf-off periods (late fall to winter). This pattern is evident in estimates from Sentinel-2 and PlanetScope models, which show close agreement. In contrast, the Sentinel-1 model shows a flatter trend with minimal seasonal variation, aligning more closely with data from nearby stilling wells AR0 and AR1 (Fig. 1), exhibiting a relatively stable trend. Unlike US-NC4 which consists of moderate to dense canopy cover (Fig. A.17), the US-NC2 region is dominated by agricultural fields, wetlands, and mixed vegetation. Here, the Sentinel-1 model maps the most inundation, and its temporal dynamics align more closely with Sentinel-2 and PlanetScope than they did at US-NC4. Unlike the consistent seasonal trend observed at US-NC4, inundation at US-NC2 varies more sporadically, lacking a clear seasonal signal. This is also reflected in the precipitation for this region, which experiences more frequent peaks than at US-NC4. In addition, this area has far fewer drainage canals compared to US-NC4 which may allow inundation to persist longer on the land surface. Although WTD data from the flux tower indicate that the tower itself was never inundated during the study period (i.e., WTD remained below 0), there were periods, such as the drawdown in spring 2019 and the subsequent rise in late fall, when similar trends were observed between WTD and mapped inundation dynamics.

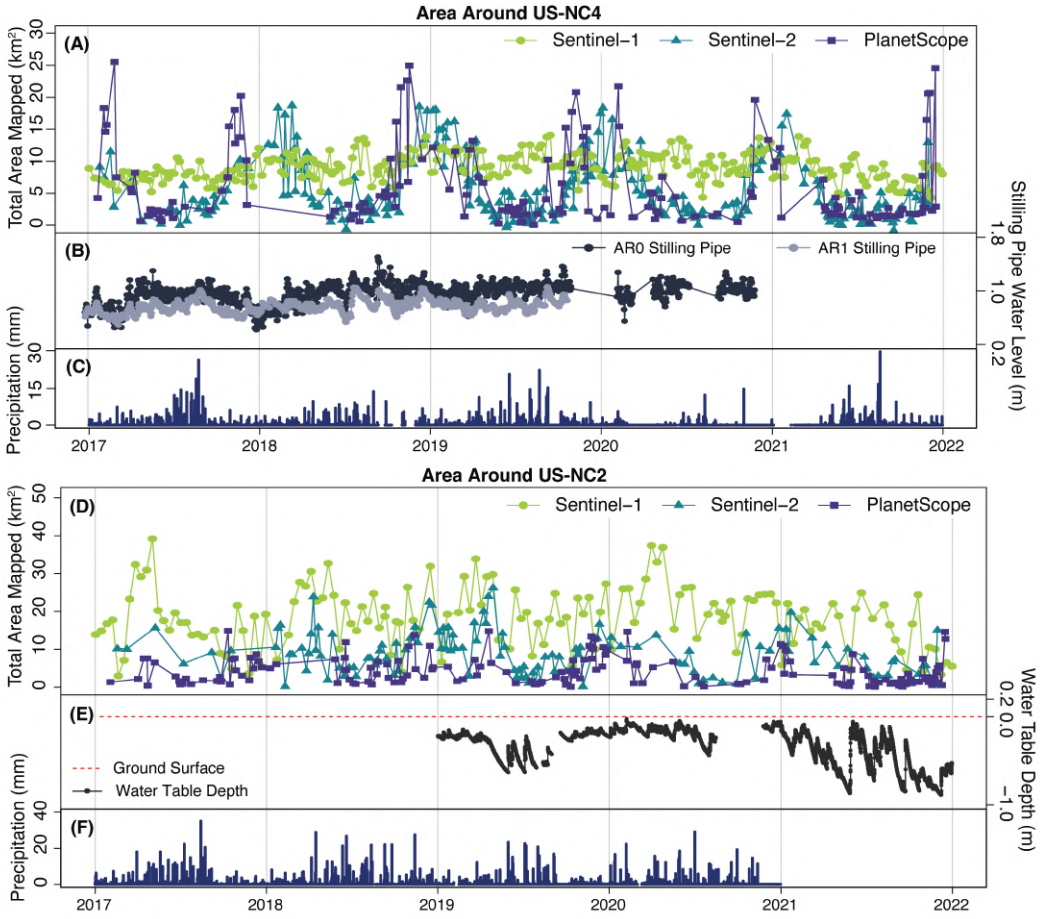

Figure A.16: (A) Time series of the total area mapped as *Inundated Vegetation* between Sentinel-1, Sentinel-2, and PlanetScope models in a 6 km buffer around US-NC4 flux tower. (B) Water levels at two stilling pipes located in the study area and near the flux tower US-NC4. (C) Precipitation measurements collected at the flux tower US-NC4. (D) Same as *(A)* but for US-NC2 flux tower. (E) WTD at US-NC2 flux tower. (F) Same as *(C)* but for US-NC2 flux tower.

Fusion Regions: Canopy Cover and Time Series

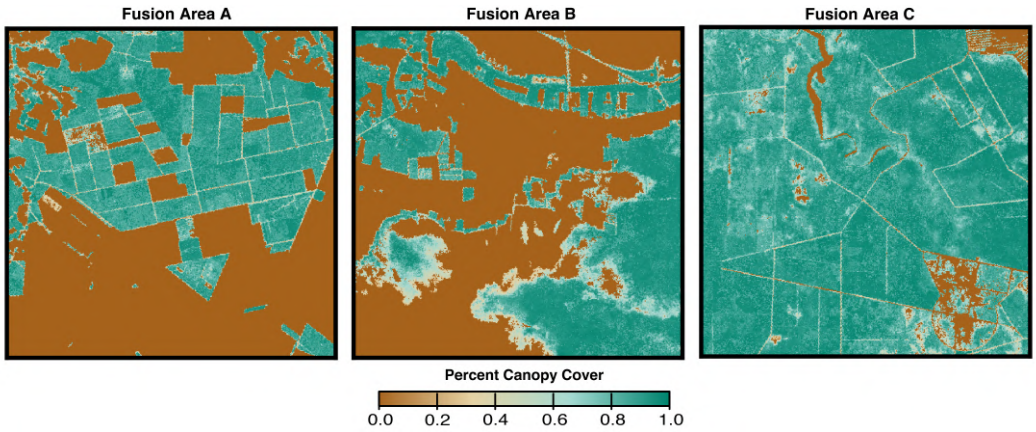

Figure A.17: Map of 2021 tree canopy cover [71] shown for each of the three fusion areas. Values ranging from 0-1 depict the percent canopy cover predicted for that pixel at 30 m resolution.

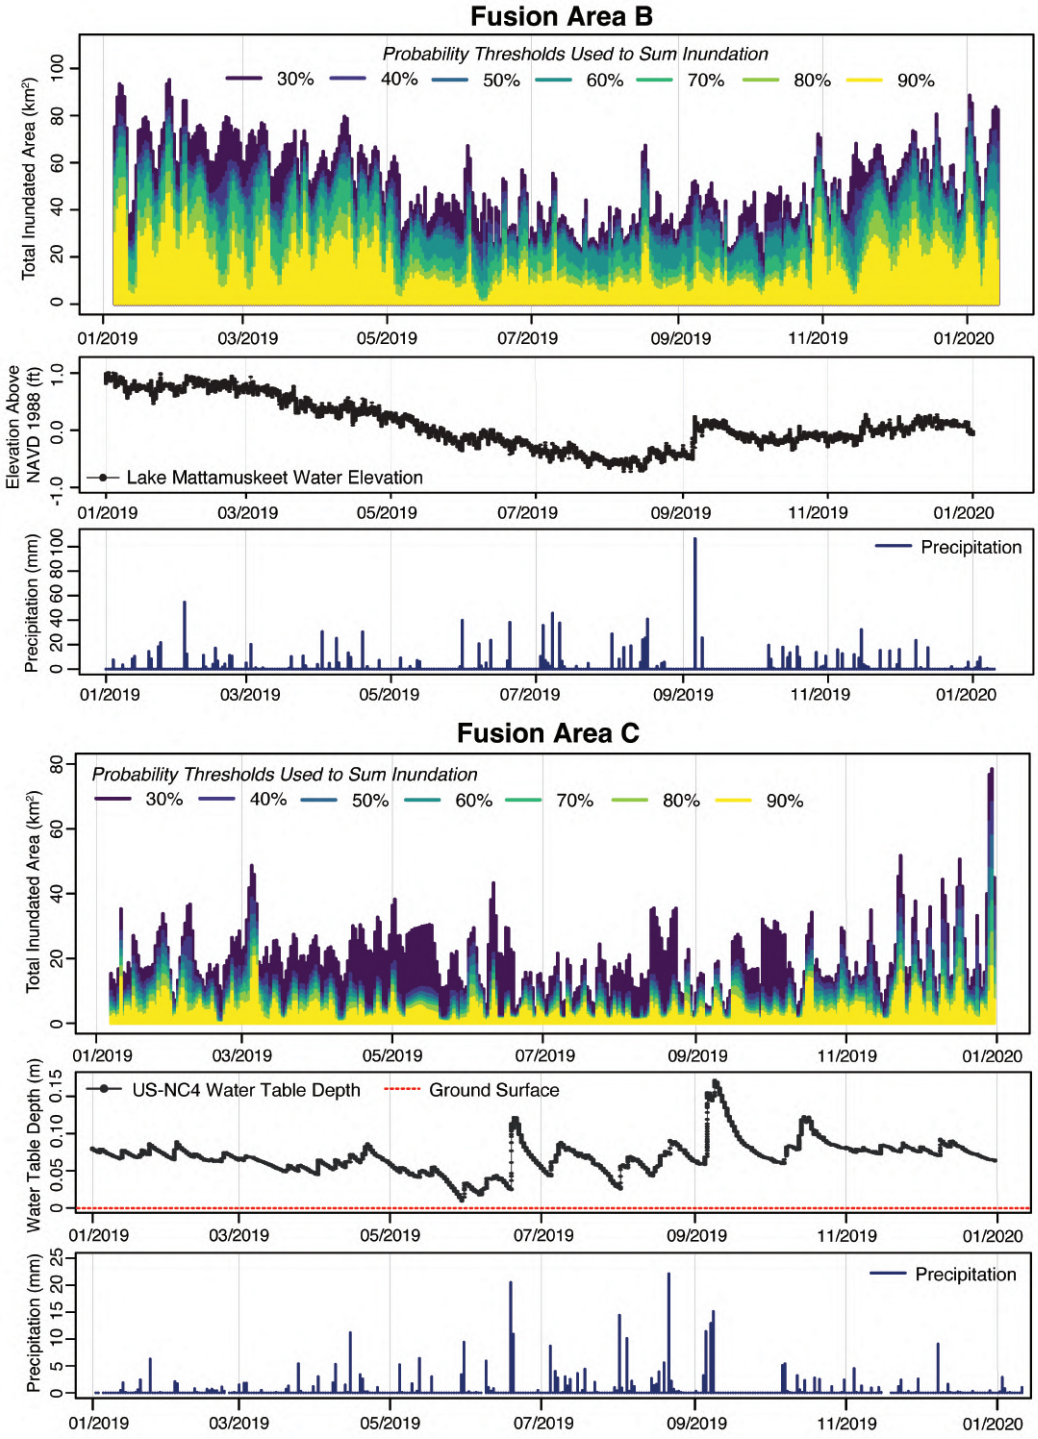

Figure A.18: Summary of fusion results for fusion areas A and B (see Figure 1). The top time series for each shows the total inundated area predicted using fusion, with each color representing a different probability threshold used to sum inundation results. For fusion area B the middle time series shows lake water elevations for nearby Lake Mattamuskeet recorded at USGS Gauge 0208458893, while for fusion area C the middle time series shows WTD measurements recorded at AmeriFlux tower US-NC4. For each area, the bottom time series shows precipitation measurements, for fusion area B these measurements are obtained from DayMet, and for fusion area C these measurements are obtained from the flux tower.

11.1 Fusion Compared to Single Sensors and Landsat DSWE

Assessing total inundated area over time revealed that the results from our fusion model captured more short-term inundation fluctuations (Fig. A.19). Across all three regions, the Sentinel-2-only data stream generally matched the broader trends in the fusion results, but there are a few interesting differences worth noting. First, in fusion region A, our fusion model captures two distinct spikes before and after July 1st, previously shown in the flux tower GWT time series for this region, and that Sentinel-2 alone did not detect. We also observe, particularly in region A, gaps in the Sentinel-2 time series due to cloud cover or limited data coverage, which prevented certain observations from being included in the time series. In contrast, the fusion model results filled these gaps, producing a continuous time series. Across all three regions, the fusion model resulted in additional inundation peaks and dips that Sentinel-2 alone did not detect.

Our fusion model also follows the high-level trends observed in the DSWE products. In fusion region A, our fusion model and the DSWE product results show greater inundation from March to May, followed by a decrease in inundation from August to mid-October. In region B, the high-level trend shows general agreement, with higher inundation in winter and lower levels in summer. However, three of the DSWE products show a spike in late August that is either not seen or observed at a significantly lower magnitude in the results from the fusion model. Similarly, in fusion region C, there is some agreement between our fusion results and DSWE on short-term periods of increased inundation, particularly for the increased levels observed in February, March, and December, but they diverge in other periods. For instance, DSWE-4 predicts a significant increase in inundation from April to late-May in region C, which is absent from DSWE-1, DWE-2, DSWE-3, and our fusion results. Region C also exhibits the largest discrepancy between the amount of inundation predicted by DSWE-4 and our fusion model results. Among the DSWE categories, only DSWE-4 approaches or exceeds the peak inundation levels predicted by our fusion model. While our results differ from individual DSWE categories in magnitude and the specific events captured, they align well with overall inundation trends when considering all categories together. This highlights both the added value of our approach and its consistency with established products.

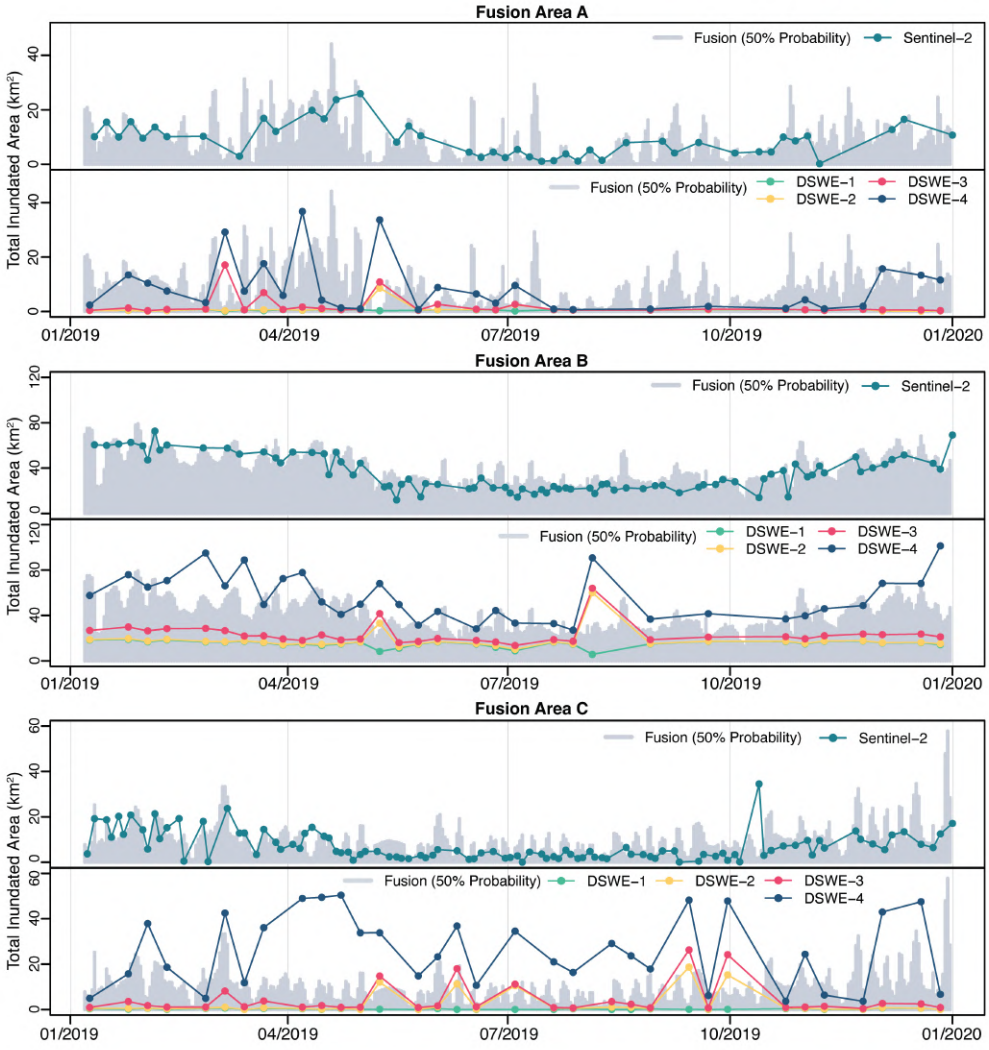

Figure A.19: Fusion results compared to a single sensor (Sentinel-2) and compared to Landsat DSWE product for fusion regions A (top), B (middle), and C (bottom). Gray shading indicates the results of the fusion model taking a 50% probability threshold, the top time series for each region displays fusion versus just Sentinel-2, and the bottom time series for each region displays colored lines indicating the results using the DSWE-1, DSWE-2, DSWE-3, and DSWE-4 categories.

Spatial Differences Between Fusion and DSWE

To assess spatial differences, we subtracted the 2019 percent inundated map generated by our fusion model (using a 50% probability threshold) from the corresponding DSWE-1 and DSWE-4 percent inundated maps. The resulting difference maps highlight areas where DSWE predicted more frequent inundation (blue), where the fusion model predicted more frequent inundation (red), and where both approaches agreed (light yellow) (Fig. A.20). It is important to note that DSWE-4 likely overestimates inundation by treating partially inundated pixels as fully inundated, whereas DSWE-1 includes only fully inundated pixels. Compared to DSWE-1, the fusion model detects substantially more inundation in small ponds and tidal wetlands, particularly in the southern portion of fusion region *B*, likely due to the incorporation of high-resolution inputs produced using PlanetScope data. Compared with NWI-designated water regimes (Fig. A.11), the fusion approach detects more inundation in the *Seasonally Flooded* class than DSWE, especially in regions *B* and *C*. In contrast, DSWE-4 maps slightly more inundation in *Seasonally Saturated* zones under dense canopy; however, this is likely due to partial surface water pixels in the *Aggressive* class overestimating inundation in areas of dense tree cover, which is a known limitation [18].

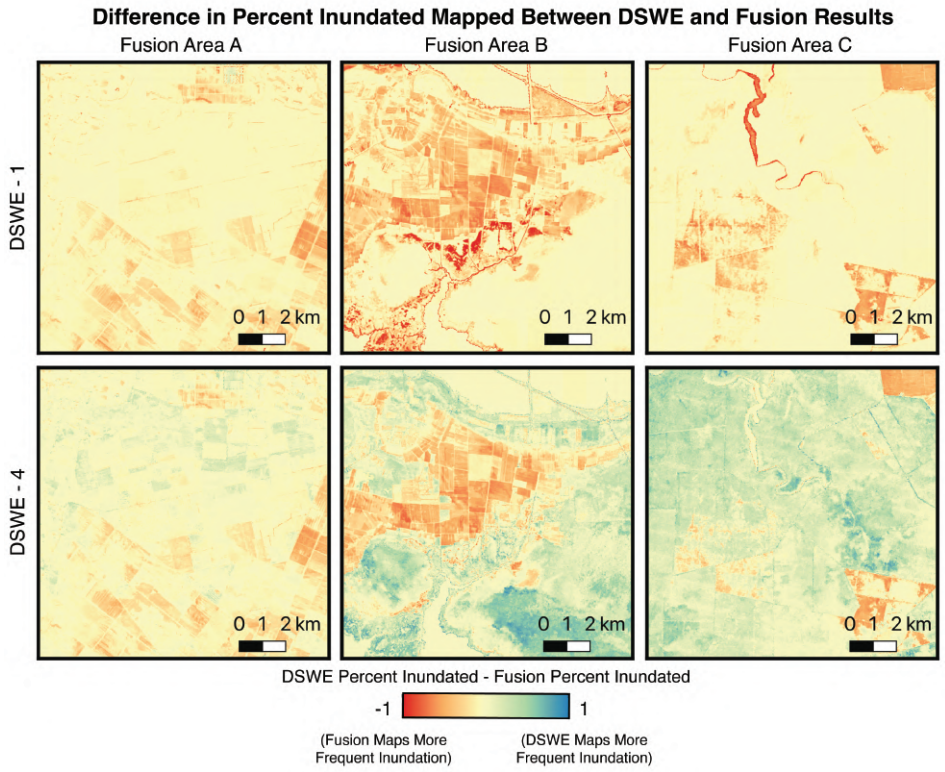

Figure A.20: Maps showing the result of taking the specified DSWE product's percent inundated map over 2019 (DSWE-1 in the top row and DSWE-4 in the bottom row) and subtracting the fusion 50% threshold percent inundated map from it. Red denotes our fusion results mapped more frequent inundation, and blue denotes DSWE mapped more frequent inundation. Light yellow indicates that both DSWE and our fusion results predicted approximately the same frequency of inundation.

11.2 Uncertainty Across Space and Time

A key advantage of using FRK for fusion is its ability to quantify prediction uncertainty. We evaluated this uncertainty both spatially and temporally, and across multiple scales. At the pixel level (Fig. A.21), gaps in satellite observations were the primary driver of variability in the width of the 90% prediction interval through time, narrowing with more inputs and widening with fewer. Notably, the model predicted increases in probability, such as those in early April, even without direct satellite observations indicating it was inundated. This is due to FRK's use of spatial dependence, where nearby inundated pixels raise the probability of inundation in a given location.

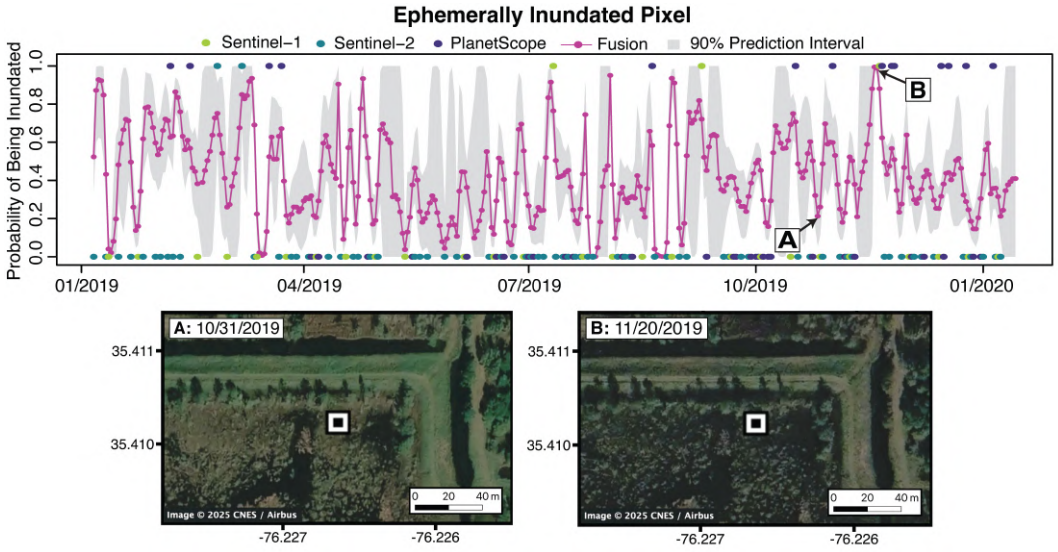

Figure A.21: Time series showing the probability of inundation and corresponding prediction interval at the individual pixel level. Satellite classification results for that pixel are also shown where 0 = *Dry Land* and 1 = *Inundated*. The location of the pixel is marked by the black and white square in the maps below, where map *A* shows a time when the pixel was dry and map *B* shows a time a month later when the pixel is inundated. These points in time are also labeled on the time series above.

We also analyzed uncertainty over time from a spatial perspective (Fig. A.22). As in the previous example, uncertainty is primarily influenced by the number of observations, evident in the December 6, 2019, maps. Here, a partial satellite observation causes uncertainty to spike across the unobserved western half of the area. Uncertainty is also spatially correlated, with higher levels observed at boundaries and transition zones between inundated and dry areas, consistent with the results of our simulation study. Certain regions, such as the southeastern and northeastern quadrants of fusion area C, consistently exhibit high uncertainty (0.8 - 1.0), even when input observations are available. Our fusion results also show a slight increase in predicted inundation probabilities from December 4th to December 5th, without any input observations. When interpreting these results, it is also important to note that the prediction intervals are influenced by the measurement variance assigned to each sensor, which is based on the classification accuracy of its machine learning model.

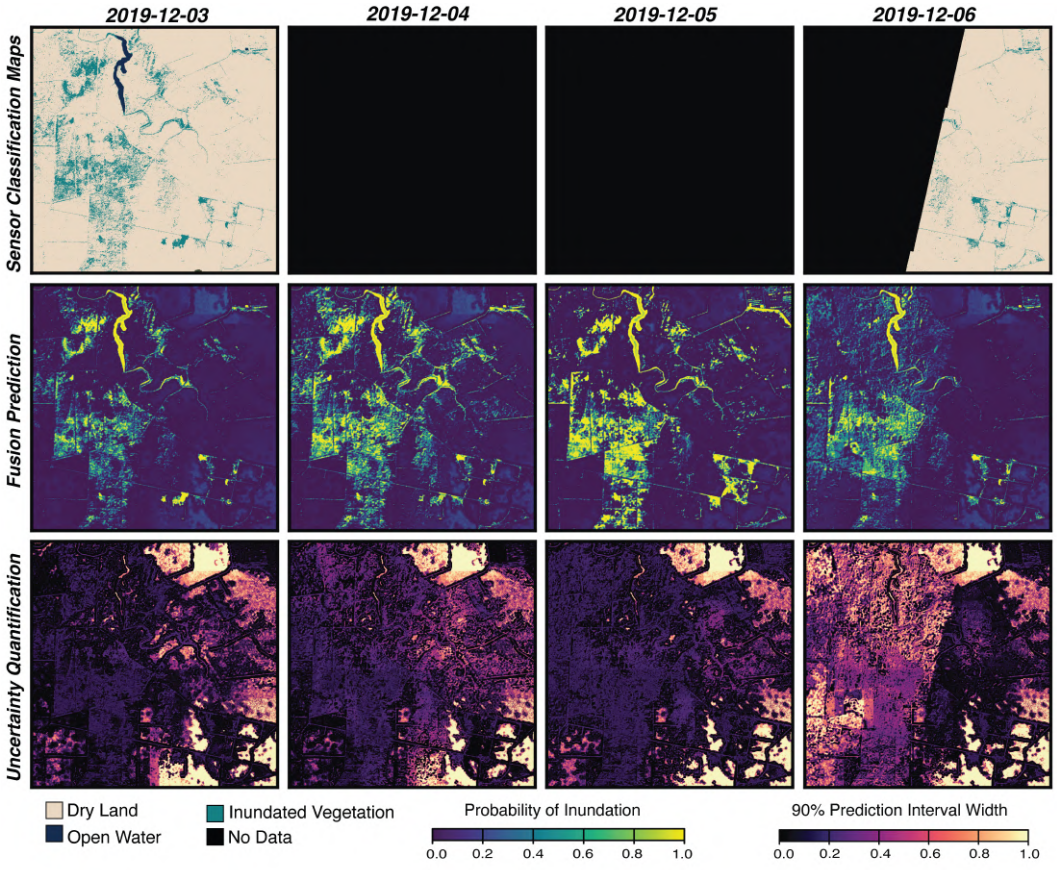

Figure A.22: Snapshots of a set of $\boldsymbol{Z}$ (PlanetScope) sensor classification observations (top row), compared to fusion predictions of inundation, or $\boldsymbol{Y}$ (middle row), as well as uncertainty quantification taking the 90% prediction interval (bottom row).

Overall, uncertainty in FRK predictions depends primarily on the number of input maps and the pixel's temporal dynamics, specifically, whether it lies near a common transition zone between dry and inundated conditions. Other factors, such as sensor measurement variance, basis function design, covariance structure, and model fit, can also impact uncertainty. For example, two areas in Fig. A.22 show persistently wide prediction intervals. This result is likely due to poor model fit caused by low observation density or noisy inputs (e.g., clouds in earlier images introducing extended effects). Additionally, uncertainty decreased and predicted inundation increased on 2019-12-04 and 2019-12-05. This effect is due to the smoothing effect of spatiotemporal basis functions, which propagated trends from surrounding observations (2019-12-03 and 2019-12-06) despite the lack of new data.