Efficient continuous land change monitoring with pyxccd and S-CCD 2.0

Su Ye [a, b] *, Yingchu Hu [a, b], Tianjia Chu [a, b]

a *State Key Laboratory of Soil Pollution Control and Safety, Zhejiang University, Hangzhou 310058, China*

b *Zhejiang Key Laboratory of Agricultural Remote Sensing and Information Technology, Zhejiang University, Hangzhou 310058, China*

This preprint is under peer review for Science of Remote Sensing

Abstract

Continuous change detection from satellite time series requires algorithms that are robust to structural breaks, computationally efficient, and accessible for reproducible large-area processing. We present *pyxccd*, an open-source, cross-platform Python package for continuous land change monitoring based on dense satellite time series. This package provides high-performance implementations of the newly proposed S-CCD 2.0 and COLD (the latest version of CCDC), through C kernels wrapped by a user-facing Python API, with pixel- and tile-based workflows for operational mapping. Relative to earlier S-CCD implementations, S-CCD 2.0 introduces an anomaly-break hierarchical decision rule for coarse-resolution products and exposes latent state estimates for interpretable trend-seasonality decomposition. Using 6,488 independently interpreted Landsat disturbance plots, S-CCD 2.0 achieved disturbance-detection performance comparable to COLD, with maximum F1 scores of 0.653 for S-CCD 2.0 and 0.664 for COLD. S-CCD 2.0 reduced computation by 1.4-1.9x for retrospective processing and 3-6x for near-real-time updating, with larger gains as spectral dimensionality increased. Application examples across Landsat and coarse-resolution vegetation products demonstrate utilities of *pyxccd* for reproducible disturbance monitoring. This new tool lowers the technical barrier to scalable, near-real-time monitoring of land-surface change from dense satellite archives.

Keywords: Change detection, Time series analysis, Disturbance, State-space model, Near-real-time, Continuous Change Detection

1. Introduction

With the advancement of modern satellite technology, the frequency and spatial resolution of observations have been greatly improved, allowing for more granular data collection over time [1]. This has resulted in richer remotely sensed datasets capable of capturing subtle changes in environmental, climatic, and urban phenomena [2], [3], thus providing a more accurate and timely understanding of global trends. The increased density of remotely sensed data enables better tracking of short-term fluctuations as well as long-term trend analysis [4], [5], ultimately advancing our capacity to address pressing issues such as climate change [6], land use [7], and disaster monitoring [8].

Time series analysis of remotely sensed datasets is characterized by several unique features that distinguish it from other time-series application domains, such as finance and energy. First, many common time series tools, like Kalman filters [9] and Autoregressive Integrated Moving Average (ARIMA) models [10], assume that the system being modeled follows a predictable pattern with constant parameters, i.e., that the system is stationary. However, satellite-based time series are often interrupted by structural breaks caused by biophysical or biochemical processes (e.g., fires, degradation, phenological shifts, or ecological succession). When these breaks occur, statistical parameters, such as the mean, variance, circular patterns, or trends, may be significantly altered, requiring a break-aware behavior to reset the temporal and maintain model accuracy. Second, satellite-based time-series dataset is typically acquired at irregular intervals, influenced by factors such as satellite orbit schedules, cloud cover, and mission-specific objectives. This irregularity in data acquisition timing requires advanced techniques to manage and interpolate gaps in the data. In near real-time scenarios, gap filling must be

performed on-the-fly as new data arrives, which can cause detection delays when a longer dataset is needed to interpolate a missing observation. Third, the massive volumes of image-based time series datasets, often in the form of high-resolution images covering vast geographic areas over extended periods, present substantial computational challenges. Processing these billions of time series at a pixel level requires highly efficient storage solutions, optimized data processing pipelines, and sophisticated analytical methods to extract meaningful insights and support timely decision-making.

The Continuous Change Detection and Classification (CCDC) algorithm, first proposed by Zhu and Woodcock (2014), has gained widespread popularity as it continuously analyzes all clear observations in chronological order and updates the land-surface temporal model as new observations become available. Its descendants and operational implementations support continental to global land-cover and land-change products [12]–[16]. One key strength is its detecting temporal breaks as the first-class feature (i.e., break-aware), by examining spectral change magnitudes at each observation using a predefined threshold. These magnitudes are computed as the difference between the observed and modeled surface reflectance, normalized by Root Mean Square Errors (RMSE). After a break is confirmed, CCDC fits a trend-seasonal harmonic regression model to each resulting temporal segment, combining a linear trend with seasonal components represented by sine and cosine terms, and then uses the fitted multiband coefficients as features for subsequent land-cover classification. Because harmonic regression does not require equal temporal spacing, CCDC can operate directly on irregularly sampled observations without resampling or gap-filling, which is advantageous for sensors with moderate revisit frequency (e.g., Landsat and Sentinel-2) where 8-day or 16-day composites

are not consistently available. Notably, Google Earth Engine (GEE) has provided CCDC as a built-in function, enabling large-scale runs via cloud computing resources [17] and substantially reducing the engineering effort needed to process extensive image archives. However, the version of CCDC in GEE does not expose several new important refinements (e.g., mean included angle, disturbance break identification) introduced in the latest CCDC version aimed at land disturbance monitoring, named as "COLD" (COntinuous monitoring of Land Disturbance) [18]. In addition, GEE may be suboptimal for local deployments that require full on-prem control or tight integration. For example, the computation could be constrained by platform quotas and task limits, customization is bounded by the server-side API environment, and exporting large intermediate or final products can add overhead and operational friction related to authentication, networking, and pipeline orchestration.

Despite its proven effectiveness, CCDC itself is not well suited for near real-time (NRT) monitoring because it updates model coefficients by reconstructing the harmonic model from scratch. For the latest version of CCDC (i.e., COLD), Zhu et al. (2020) suggested the model updating frequency is per observation with ~10% higher F1 score than annual updating frequency, while later being changed to the 3% of the observation number that has not undergone updating with three observations as minimum for efficiency boosting purpose. This rebuilding strategy typically requires retaining and repeatedly accessing a long historical record, which increases storage demands and reduces the operational feasibility of NRT deployments. It is also computationally inefficient: re-reading archives and refitting models introduces avoidable latency and cost, undermining timely change detection. To address these limitations, Stochastic Continuous Change Detection (S-CCD) algorithm extends CCDC by reformulating

the seasonal-trend model in a state-space framework [19], where trend and seasonal terms are treated as stochastic, time-evolving states. S-CCD preserves the continuous-monitoring logic from CCDC/COLD in the sense that it still identifies temporal breaks, but differ in how they update the model: rather than refitting the harmonic model from scratch, S-CCD applies a Kalman filter to update state estimates recursively as each new observation becomes available. This short-memory formulation propagates only the current state and its uncertainty, eliminating the need to retain and revisit the full time-series archive and thereby improving both NRT feasibility and computational efficiency.

Beyond efficiency, state-space models in S-CCD allow for a dynamic representation of the system over time, with improved handling of temporal dependencies. In COLD, change detection relies on static thresholds on change magnitudes against the prediction from the fixed models, which can fail to adapt to varying patterns or trends in the data. In contrast, the state-space framework enables S-CCD to update the model dynamically based on new observations, ensuring that the model remains relevant as local fluctuations. This dynamic adaptation helps better accommodate gradual within-regime drift (e.g., phenological shifts) without triggering unnecessary model reinitialization, while still flagging true disturbances via statistically significant innovations (filter residuals). The resulting state estimates also provide an interpretable representation of evolving seasonal and trend behavior, supporting finer characterization of subtle, progressive change.

Although S-CCD was initially introduced in 2021, early applications focused primarily on forest disturbance monitoring, and its reported F1 score for a comprehensive land-cover dataset was lower than that of COLD [19]. Moreover, the initial release exposed only a low-level C

implementation and was limited to Linux, constraining accessibility, and broader adoption. Motivated by multiple mapping projects and sustained user feedback, we have substantially improved both the detection accuracy, hereafter referred to as S-CCD 2.0, and its usability through modern software engineering practices. To better support the user community, we developed *pyxccd*, an open-source, cross-platform Python package for scalable, continuous land change monitoring from satellite time series.

This manuscript makes four main contributions: (1) S-CCD 2.0 extends the original state-space formulation with an anomaly–break hierarchical decision rule and latent-state outputs; (2) *pyxccd* provides high-performance C implementations for popular CCDC-like algorithms (COLD and S-CCD 2.0) through a user-facing Python API; (3) COLD and S-CCD 2.0 are evaluated using 6,488 independently interpreted Landsat disturbance plots; and (4) pixel- and tile-level workflows demonstrate how the package supports reproducible NRT disturbance monitoring and large-area mapping. Together, these algorithmic improvements, open-source implementation, and reproducible workflows lower the technical barrier for researchers and operational agencies seeking to continuously analyze dense satellite archives for land change monitoring.

2. Algorithms implemented in *pyxccd*

2.1 COLD

After screening out the observations flagged as cloud, shadow, and snow by Fmask [20], COLD can be described as proceeding in two stages [18]: model initialization and continuous monitoring.

2.1.1 Model initialization

Model initialization aims to establish a stable "no-change" baseline trajectory from an initial set of clear observations. COLD first defines an initialization window that satisfies minimum data requirements (at least 12 clear observations, spanning at least one year, with the largest gap shorter than one year), and then applies Tmask within that window to further remove residual cloud/shadow/snow contamination missed by Fmask. It fits a harmonic regression per inputted band using the LASSO regression. The fitted model is then subjected to a stability test designed to reject a window that likely straddle an undetected change. In COLD, this test combines the absolute slope term with the larger absolute prediction bias at the start or end of the window, i.e.,

$$Magnitude_{overall} = \left|l_{slope}\right| + \max\{|l_{start}| + |l_{end}|\} < threshold \quad (1)$$

If the window fails the stability test, COLD advances the window by one clear observation (dropping the earliest and adding the next) and repeats the fitting and testing until stability is achieved. Once a stable window is found, COLD additionally applies a "look-back" step to revisit earlier clear observations that were skipped during window shifting and include those that remain consistent with the current model. Finally, the baseline model is refit for the finalized initialization window, with trend plus three seasonal harmonics:

$$\hat{\rho}_{i,t} = \mathrm{a}_{0,i} + \mathrm{c}_{1,i}t + \sum_{k=1}^{n=3}(\mathrm{a}_{k,i}\cos\left(\frac{2\pi k}{m}t\right) + \mathrm{b}_{k,i}\cos\left(\frac{2\pi k}{m}t\right)) \quad (2)$$

Where i is the ith band, t is the ordinal date, $\mathrm{a}_{k,i}$ and $\mathrm{b}_{k,i}$ are the harmonic coefficients, k indexes the harmonic frequencies (annual, semiannual, and four-month cycles), m is the number of days per year (m=365.25).

2.1.2 Continuous monitoring

After the initial baseline model is established, COLD performs break detection sequentially. For each newly acquired clear observation, COLD updates the current-segment harmonic model using all stable clear observations accumulated so far. The model is re-fit from scratch using lasso regression with penalty parameter λ (Fig. 1), which controls the L1 regularization strength (shrinkage of coefficients). A smaller λ yields weaker shrinkage and a more complex model (higher overfitting risk), while λ = 0 reduces lasso to ordinary least squares. In the *pyxccd* implementation, λ is suggested as 20 for Landsat-like reflectance scaled to [0, 10000].

Using the updated model, COLD predicts band-wise reflectance for the new date and evaluates the deviation between the observation and prediction as a multi-band change vector using a peek window (Fig.1). For the x-th observation and the i-th band, the normalized residual is

$$cm_{i,x} = \frac{\rho_{i,x} - \hat{\rho}_{i,x}}{RMSE_{i,x}} \quad (3)$$

Where $RMSE_{i,x}$ is the larger one between the temporally adjusted Root Mean Square Error (RMSE) and minimum RMSE that is lag-1 madogram for the time series [18]. The multi-band change magnitude is then defined as the squared norm of the normalized change vector:

$$CM_x = \sum_{i=1}^{k} cm_{i,x}^2 \quad (4)$$

If each $v_{i,x}$ is approximately standard normal, then CM_x follows χ^2 distribution with k degrees of freedom, which allows COLD to set a statistically interpretable threshold using a change probability P (Fig. 1) (Typically $P = 0.99$). A new observation is flagged as a break candidate when $CM_x > \chi^2_k(P)$. Using a strategy common to CCDC-like algorithms, COLD

employs a "peek window" approach to mitigate false alarms triggered by transient noise. A break is only confirmed if candidates persist for a sequence of $CONSE$ consecutive observations ($CONSE = 6$ by default).

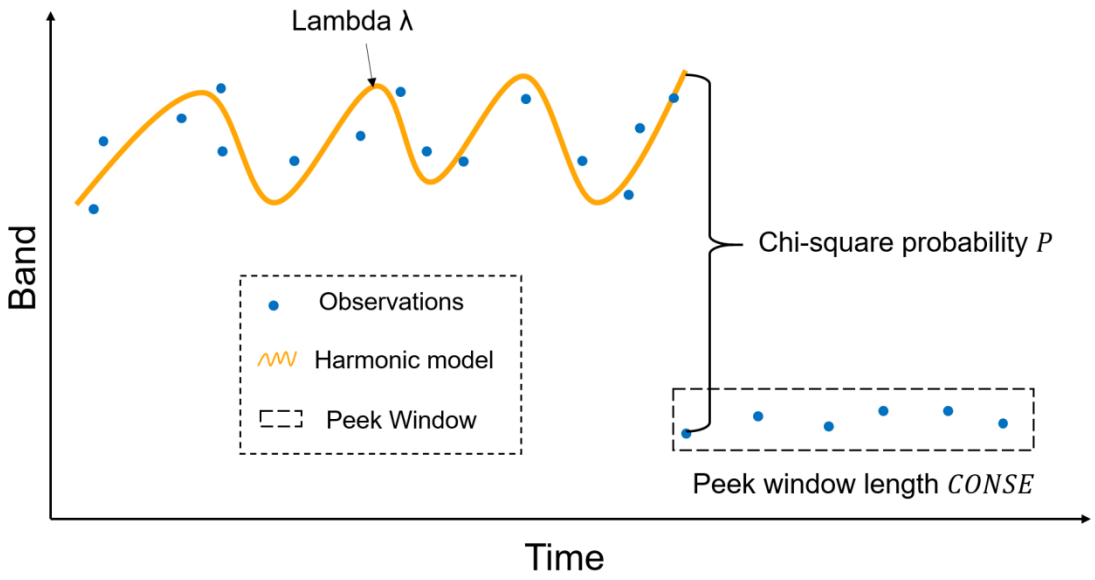

Fig.1. Conceptual figure for three critical parameters for CCDC-like algorithms.

Besides the change magnitudes, COLD applies a directional-consistency test by computing included angles between neighboring normalized change vectors over the consecutive anomalies. A break is confirmed only if the mean included angle is below a threshold (reported as 45°). Once a change is confirmed, the break is recorded and the algorithm proceeds to initialize the next temporal segment using post-break observations.

Unlike GEE-CCDC, COLD performs disturbance extraction after confirming a breakpoint, converting generic spectral breakpoints into disturbance-oriented products by excluding false positives, such as those associated with vegetation recovery. It uses a rule-based attribution system that analyzes the direction of the multiband change vector and the model slopes before and after the breakpoint. COLD identifies greener-direction breaks (NIR increases, Red and

SWIR1 decrease) as regrowth and classifies other breaks as land disturbances. For greener breaks, it differentiates regrowth from reforestation or afforestation based on trend behavior, producing three categories: disturbance, regrowth, and afforestation.

The differences between GEE-CCDC and COLD are summarized in Table 1.

Table 1. Comparison between GEE-CCDC and COLD

Steps	GEE-CCDC	COLD
Stability test	$\lvert l_{slope}\rvert + l_{start} + \lvert l_{end}\rvert < threshold$	$\lvert l_{slope}\rvert + \max\{\lvert l_{start}\rvert + \lvert l_{end}\rvert\} < threshold$
Directional consistency check	None	Mean included angle test
Model update	Update by last fitting span days multiplied by 1.33	Update at each 3% of the observation number that has not been updated with three observations as minimum
Disturbance Extraction	None	Rule-based disturbance, regrowth, and afforestation attribution

2.2 S-CCD 2.0

S-CCD 2.0 builds upon the prototype introduced by Ye et al [19] and follows the same two-step framework as COLD (Fig. 2). For retrospective scenario, S-CCD 2.0 processes historical datasets in batch mode, with observations queued and analyzed after enough time-series observations are collected. The retrospective scenario is used for backtracking historical disturbance events to understand past behavior, or to prepare state-space model for the future NRT usage. Differently, in the NRT scenario, S-CCD 2.0 continuously fed newly collected datasets into the system at a certain monitoring frequency (e.g., weekly). S-CCD 2.0 will check the current "status" of the pixel and choose to add them into the observation queue for model initialization, or detect anomalies/breaks with the current state-space model for continuous

monitoring.

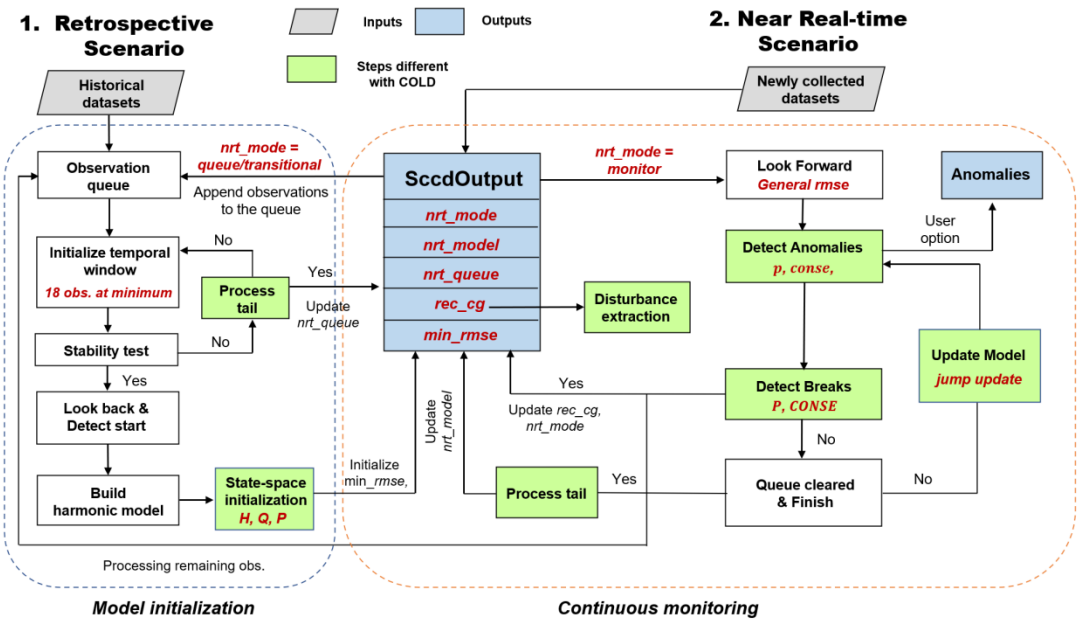

Fig. 2. The workflow of S-CCD 2.0, which is newly proposed in this study.

In the S-CCD 2.0 workflow, *SccdOutput* plays a critical role in managing and storing the results of the monitoring and detection processes. It acts as a structured record that contains all the necessary data for both model initialization and continuous monitoring. There are five key elements in *SccdOutput*.

nrt_mode: a two-digit integer describing predictability (first digit; 1 indicates no predictability) and the current monitoring status (second digit). It helps informing predictability of the state-space model, and track whether the system is in the process of initialization, monitoring real-time data, or in a transitional state.

nrt_model: a compact structure that stores the current state-space model when *nrt_mode* is monitor or transitional. It contains the updated model parameters that are used in the next monitoring cycle, ensuring that the system always has the most recent information for detecting anomalies and performing updates. See Table A1 in the appendix for details.

nrt_queue: a structured array that stores incoming observations when *nrt_mode* is queue. It allows the system to accumulate enough clear observations before reinitializing the state-space model.

rec_cg: a structured array that records temporal segments and break information detected in the past according to the selected break thresholds (P and CONSE).

min_rmse: an array containing the minimum RMSE for each band estimated during state-space model initialization. This value helps avoid unrealistically small errors for dark or shaded pixels

2.2.1 State-space initialization

State-space models, the core mathematic foundation of S-CCD, explicitly model both observation and process uncertainties, which are key components for accurate change detection:

Observation equation: $y_t = Z\alpha_t + \epsilon_t, \quad \epsilon_t \sim \text{Normal}(0, H)$ (5)

Process equation: $\alpha_{t+1} = T\alpha_t + \eta_t, \ \eta_t \sim \text{Normal}(0, Q)$ (6)

Here, y_t is the observation at time t, α_t is a state vector governing the system's evolution over time. S-CCD uses bimodal seasonal model (annual + semiannual), with six state variables by default: μ_t (trend), v_t (trend), $\gamma_{1,t}$ (annual), $\gamma^*_{1,t}$ (annual), $\gamma_{2,t}$ (semi-annual), $\gamma^*_{2,t}$ (semi-annual), where $\alpha_t = \left(\mu_t, v_t, \gamma_{1,t}, \gamma^*_{1,t}, \gamma_{2,t}, \gamma^*_{2,t}\right)^T$. To better accommodate the complex temporal variation, the users are opted to choose the trimodal (trend + annual + semiannual + four-month) in S-CCD 2.0. These states are transformed by the harmonic model from the initialization step (Eq. A1-A6 in Appendices). The system matrix Z indicates which state variables directly contribute to the observation (for six states):

$$Z = [1, 0, 1, 0, 1, 0] \quad (7)$$

Here, the states v_t, $\gamma^*_{i,t}$, $\gamma^*_{2,t}$ serve as auxiliary states that are not directly involved in observation calculation but are essential for recursive updates of the other states, therefore, their weights were assigned as zero in Z.

The observation and process noise terms ϵ_t and η_t, respectively, quantify the uncertainties in the observation and the process model. Both are independent and normally distributed, with variances H and Q, respectively. In S-CCD 2.0, H is assigned the RMSE of the initial model, and Q is defined as a six-element vector corresponding to the six state variables:

$$Q = [q, \frac{q}{slope\_scale}, q, q, q, q] \quad (8)$$

Where q is calculated as 0.25 * $\frac{LAM}{\lambda}$, indicating the influence of the parameter λ on model adjustments: Smaller values of λ increase the weight of fitting the actual observations, leading to a higher q and better fit. The slope state v_t is scaled by $slope\_scale$ (default value is 10,000,000), given that it typically has a lower numeric value than the other coefficients.

The transformation matrix T defines how the state vector evolves over time and is determined by the harmonic components. For the bimodal seasonal model, T is as follows:

$$T = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & \cos\left(\frac{2\pi}{365.25}\right) & \sin\left(\frac{2\pi}{365.25}\right) & 0 & 0 \\ 0 & 0 & -\sin\left(\frac{2\pi}{365.25}\right) & \cos\left(\frac{2\pi}{365.25}\right) & 0 & 0 \\ 0 & 0 & 0 & 0 & \cos\left(\frac{2\pi}{2*365.25}\right) & \sin\left(\frac{2\pi}{2*365.25}\right) \\ 0 & 0 & 0 & 0 & -\sin\left(\frac{2\pi}{2*365.25}\right) & \cos\left(\frac{2\pi}{2*365.25}\right) \end{bmatrix} \quad (9)$$

A crucial parameter in the state-space model is the covariance matrix P, which

represents the uncertainty of the state variables. The initial covariance P_0 reflects the uncertainty of the initial state. For the Landsat surface reflectance product, an uncertainty of 5% has been reported [21]. Accordingly, the uncertainties of the initial states are estimated as $\lambda = (\mu_0 * 5\%)^2$. The initial covariance P_0 is thus given:

$$P_0 = \begin{bmatrix} \lambda & 0 & 0 & 0 & 0 & 0 \\ 0 & \lambda/slope\_scale & 0 & 0 & 0 & 0 \\ 0 & 0 & \lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & \lambda & 0 & 0 \\ 0 & 0 & 0 & 0 & \lambda & 0 \\ 0 & 0 & 0 & 0 & 0 & \lambda \end{bmatrix} \quad (10)$$

Unlike COLD, which determines the minimum RMSE for the entire time series, S-CCD 2.0 calculates the lag-1 madogram based solely on the observations used during the initialization stage. The minimum RMSE for each inputted band is retained in the output ("*SccdOutput*") and can be used for subsequent online monitoring.

2.2.2 Detect anomalies & breaks

Once the observation y_t enter the system, S-CCD predicts the value for the band i at time t using the state vector:

$$\hat{\rho}_{i,t} = Z\alpha_{i,t} \quad (11)$$

Subsequently, it calculates the change magnitude CM_x following the same procedure of COLD following Eq. 3 and Eq. 4.

S-CCD 2.0 implements a two-level detection hierarchy, i.e., anomaly and break, to assess each observation x within its peak window. This approach differs from most CCDC-like methods, which focus exclusively on break detection (i.e., structural changes). The two-level hierarchy enables detection of subtle, short-lived signals (e.g., drought-related spectral responses) and enhances performance for noisy or coarse-resolution time series. This is

important in cases where weak anomalies provide useful information but do not require model re-initialization. In typical CCDC-like workflows, a break triggers model re-initialization, which often requires at least one year of post-change observations to stabilize the model [4]. Frequent re-initialization disrupts continuous monitoring and complicates the NRT operation, because the absence of a stable model immediately after a break reduces temporal continuity and limits the interpretability of subsequent monitoring outputs.

To address this, we introduces unified monitoring workflow with two detection levels following [8]: (1) the anomaly level, which uses a sensitive configuration with three consecutive anomalies with a change magnitude larger than 0.90 chi-square probability ($p = 0.90, conse = 3$ by default), to inclusively detect all types of spectral anomalies without triggering model re-initialization; and (2) the break level, which uses a conservative configuration ($P = 0.9999, CONSE = 8$ by default) to identify pronounced structural changes that justify model reset. We allowed users to customize p, $conse$, P, $CONSE$.

We also improve some detailed steps for S-CCD 2.0. For the current *pyxccd* version (v1.0.2), S-CCD 2.0 calculates a general RMSE, which simplifies the calculation by focusing on overall error without adjusting it for temporal changes, as opposed to the recursive temporally-adjusted RMSE used in S-CCD 1.0 [19]. The general RMSE requires much fewer variables to store for updates, while the recursive RMSE in S-CCD 1.0 necessitates storing 122 variables. The final RMSE is the larger one between the general RMSE and minimum RMSE obtained in the initialization stage. To enhance the performance for the general disturbances, S-CCD 2.0 adopted the exactly same step of directional-consistency test and disturbance extraction in COLD after anomalies and breaks are identified, as opposed to angular spread and

disturbance evidence only for forest disturbances proposed in the first version [19].

2.2.3 Update model

S-CCD leverages linear Kalman filter to update the harmonic model from the previous state-space parameter at a per-band basis, ensuring that the model is always up to date and reflective of current data. For band i, S-CCD estimates the filter state $a_{t|t,i}$ which combines observations and model prediction based upon their uncertainty measurements, i.e., the observation noise (H_i) and the covariance matrix for the states ($P_{t,i}$):

$$F_{t,i} = Z\ P_{t,i} Z^T + H_i \quad (12)$$

$$K_{t,i} = P_{t,i} Z^T \quad (13)$$

$$a_{t|t,i} = a_{t,i} + K_{t,i} F_{t,i}^{-1} v_{t,i} \quad (14)$$

Next, the model generates one-step-ahead predicts for the states and the covariance matrix for the next step:

$$a_{t+1,i} = T\ a_{t|t,i} \quad (15)$$

$$P_{t+1,i} = T\left(P_{t,i} - K_{t,i} K_{t,i}^T F_{t,i}^{-1}\right) T^T + Q_i \quad (16)$$

$a_{t+1,i}$ will be used to predict surface reflectance ($\hat{\rho}_{i,t+1}$) for $t+1$ following Eq.11. In the case of missing observations at time t, the covariance matrix could be simply updated by the transformation matrix T and Q_i:

$$P_{t+1,i} = T P_{t,i} T^T + Q_i \quad (17)$$

For m consecutive days that the observations are missing, this update process would need to be repeated m times, which could be computationally costly. To improve efficiency, we developed a "jump update" for Kalman filter to estimate missing-observation $P_{t+1,i}$ more effectively. Mathematically, processing m consecutive missing-observation days is equivalent

to perform one-time regular update for a time series that has a revisiting cycle of m days. Accordingly, S-CCD 2.0 recalculates a temporal transformation matrix for a time series with an interval of m days:

$$T^* = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & \cos\left(\frac{2\pi * m}{365.25}\right) & \sin\left(\frac{2\pi * m}{365.25}\right) & 0 & 0 \\ 0 & 0 & -\sin\left(\frac{2\pi * m}{365.25}\right) & \cos\left(\frac{2\pi * m}{365.25}\right) & 0 & 0 \\ 0 & 0 & 0 & 0 & \cos\left(\frac{2\pi * m}{2 * 365.25}\right) & \sin\left(\frac{2\pi * m}{2 * 365.25}\right) \\ 0 & 0 & 0 & 0 & -\sin\left(\frac{2\pi * m}{2 * 365.25}\right) & \cos\left(\frac{2\pi * m}{2 * 365.25}\right) \end{bmatrix} \quad (18)$$

This allows for a one-time update using T^*, instead of m times (Eq. 19):

$$P_{t+1,i} = T^* P_{t,i} T^{*T} + Q_i \quad (19)$$

Particularly, the state variables ($a_{t,i}$) are scientifically useful because they transform a noisy observation time series into an interpretable, continuously updated decomposition of ecosystem dynamics, separating components such as the seasonal baseline and the slowly varying level/trend. Fig. 3 compares the decomposition results from COLD (whose predictions rely on fixed harmonic coefficients) and S-CCD 2.0 (whose outputs are the estimated states), for the MODIS NDVI time series. Both methods recover a seasonal cycle and provide a model fit to the observations, but they differ markedly in their characterization of the long-term trajectory. COLD summarizes the record as a monotonic and linear increase, which can mask subtle curvature in the underlying trend. In contrast, S-CCD 2.0 yields a more adaptive, continuously evolving trend that highlights a gradual greening rise followed by a clear saturation/flattening ("saturation point") and subsequent slow variations. As a result, S-CCD 2.0 is particularly well suited for inspecting slowly varying trends, since it can track progressive, non-linear changes that may be smoothed over or misrepresented when the trend is constrained

to a single linear form.

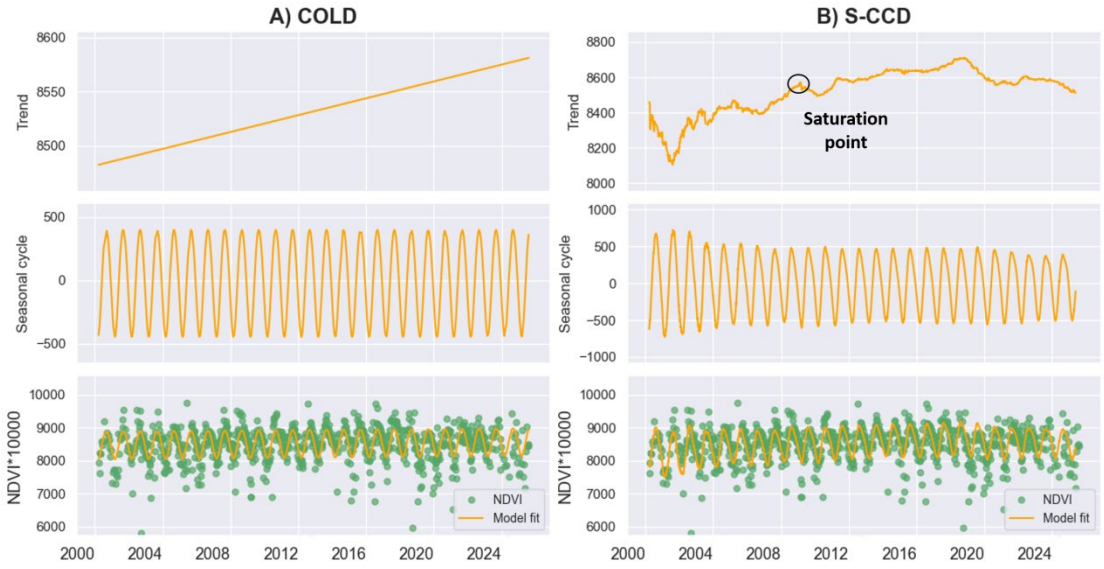

Fig. 3. MODIS-based NDVI Time-Series decomposition for analyzing greening trend over a location of Tibet, China, illustrating that "states" of S-CCD 2.0 could be used to identify slowly varying trend/seasonality.

2.2.4 Process tail

The final step of S-CCD 2.0 processes the tail end of the time series (Fig. 4), specifically the last CONSE observations, and saving all results into *SccdOutput*, for the next monitoring cycle, which includes new observations. There are three possible statuses for the tail observations (the second digit of the *nrt_mode* parameter): *monitor*, *queue*, and *transitional Mode.*

Monitor Mode: the state-space model has been initialized. We are allowed to update a dictionary-like structure *nrt_model* (see Table A1 in the appendix), retaining key parameters for future monitoring and updates (e.g., the peak window for the next update, state-space model parameters, and the sum of squared errors). The model can process newly incoming

observations with the parameters kept in *nrt_model*. S-CCD 2.0 additionally extract information into in *nrt_model* for consecutive anomalies at last (Fig. 4), as they often signify ongoing disturbances which has not satisfied the consecutive number required for identifying a break. Machine learning techniques can be further applied to these anomalies at the tail for NRT agent classification.

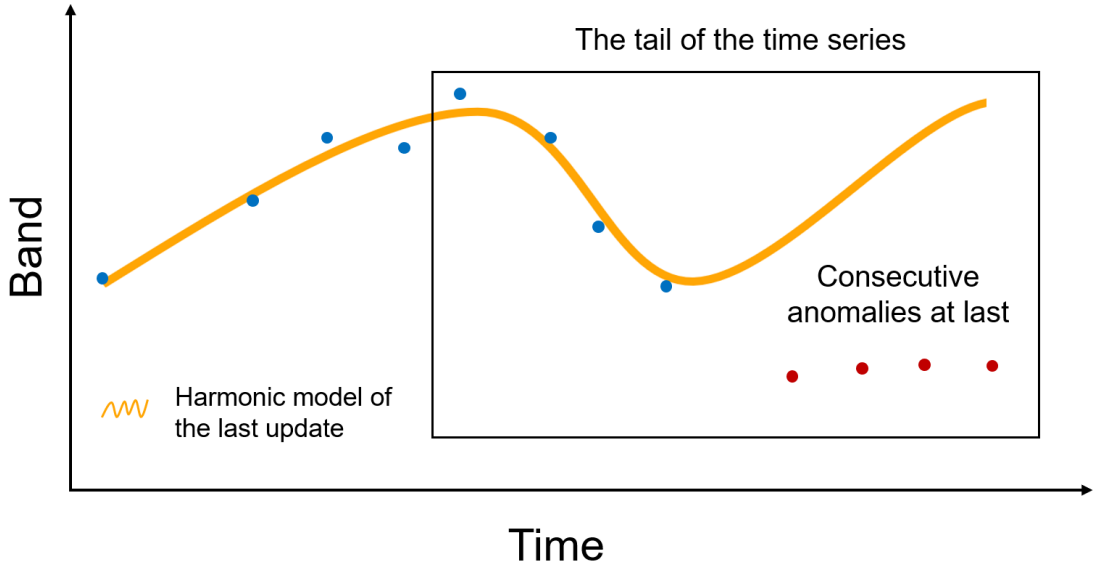

Fig. 4. S-CCD 2.0 processes the tail of the time series and reports the status for the consecutive anomalies at the last (dashed rectangle).

Queue Mode: the state-space model is not initialized yet, so S-CCD 2.0 continues to collect new clear observations into *nrt_queue*. Once new observations arrive, the system combines them with the stored observations from *nrt_queue* to perform a stability test to determine if a state-space model can be established.

Transitional Mode: designed for practical NRT monitoring tasks with weekly or bi-weekly update frequencies, the *transitional* mode is triggered right after a break is detected. Practically, users require recent anomaly information for a period to generate spatially consistent maps after

387 the break. To accommodate this, S-CCD 2.0 retains *nrt_model* for two weeks following the
388 transition from *monitor* to *queue* mode, allowing an NRT analysis on the recent anomalies even
389 they are not from the last update.

390 Additionally, the first digit of the *nrt_mode* parameter indicates the predictability, with "1"
391 indicating no predictability and "0" meaning predictability. We conducted a predictability test
392 for both *monitor* and *transitional* modes to ensure that the state-space model remains stable
393 after the last anomaly is confirmed as a disturbance (for *queue* mode, unpredictability is always
394 set). The problem arises when a disturbance process persists for several weeks or months, such
395 as construction. In such cases, ambiguous anomalies may be continuously generated, as surface
396 reflectance during an event can be unpredictable. To address this, we designed a test to assess
397 the predictability of the state-space model. The model is considered predictable if the change
398 magnitudes of half the observations in the peak window are smaller than the critical value
399 $\chi^2_k(predictability\_pcg)$. The default $predictability\_pcg$ set to is 0.90. Users can choose
400 to exclude ambiguous anomalies based on the outcome of the predictability test.

401 Overall, S-CCD 2.0 made substantial modifications on S-CCD first proposed, summarized
402 in Table 2.

403

404 **Table 2 Comparison between S-CCD** [19] **and the newly proposed S-CCD 2.0**

	S-CCD 1.0 [19]	S-CCD 2.0
State structure	Five states (no slope): $\mu_t, \gamma_{1,t}, \gamma^*_{1,t}, \gamma_{2,t}, \gamma^*_{2,t}$	Six states (or eight states as option): $\mu_t, v_t, \gamma_{1,t}, \gamma^*_{1,t}, \gamma_{2,t}, \gamma^*_{2,t}, \gamma_{3,t}, \gamma^*_{3,t}$
RMSE	Recursive temporally-adjusted RMSE	General RMSE
Detection	Only detect breaks	Detect both anomalies and breaks
Directional-consistency	Angular spread	Included angle (COLD's approach)
Disturbance extraction	Combined evidence from multiband change magnitudes	Examination on individual change magnitudes (COLD's approach)

Model updates	Covariance updates for each date regardless of missing observation	Jump covariance update for continuous missing observations.
Process tail	None	Perform additional predictability test; identify and save current mode

3. Software implementation

Pyxccd is an open-source, PyPI-installable Python package that bridges low-level computational performance with high-level scientific utility (https://github.com/Remote-Sensing-of-Land-Resource-Lab/pyxccd). It succeeds the MATLAB-based COLD implementation and serves as the primary platform for S-CCD 2.0 algorithms. Designed for both developers and researchers, *pyxccd* provides a flexible and scientifically rigorous environment for processing multi-sensor time-series datasets, moving beyond the fixed Landsat-band structure of the original CCDC framework.

The software architecture focuses on two key imperatives: computational throughput and modular flexibility.

Hybrid C-Python Core: at its core, *pyxccd* integrates C-based computational kernels to handle tasks like iterative Lasso regression and Kalman filtering, essential for the COLD and S-CCD 2.0 algorithms. Offloading these computational tasks to C ensures minimal memory overhead and high execution speed, making *pyxccd* suitable for both local workstations and distributed High-Performance Computing (HPC) clusters. Python acts as the interface layer, leveraging Cython and NumPy C-APIs for low-latency data exchange between the user interface and execution engine. This hybrid approach allows *pyxccd* to achieve much faster processing than pure Python implementations while offering an intuitive API.

Modular Functional Design: *pyxccd* employs a modular architecture that separates core detection logic from data I/O, time-series processing, and spatial orchestration. This allows specialized resource allocation, such as isolating numerical computations from data management tasks, and enables "lazy loading", where only relevant data segments are loaded into memory. The modularity also allows for easy adaptation to the evolving remote sensing landscape, supporting various satellite data formats (e.g., HLS, MODIS, Planet) without the need for a complete re-engineering of the pipeline. We presented the details for the module designs in Appendix A2.

4. Evaluation

4.1 Testing dataset

We evaluated the accuracy and computational efficiency of COLD and S-CCD 2.0 for detecting disturbance-related breaks using an identical reference dataset. Specifically, we adopted the 7,200 disturbance plots from Conterminous United States used in the COLD study by Zhu et al. (2020). Their disturbance dates were originally interpreted by Cohen et al. (2016) based on Landsat time series, high-resolution imagery and ancillary datasets. Zhu et al. (2020) subsequently refined this reference database by excluding plots with insufficient clear-sky observations or high interpretation uncertainty. Following this refinement, a total of 6488 Landsat plots were retained and used as testing samples in this study.

For the accuracy assessment, all available Landsat observations (Landsat 4, 5, 7, and 8) acquired between 1983 and 2017 were collected on a per-plot basis. For plots covered by two adjacent orbital paths, only the dominant path was retained to avoid duplication. The accuracies

of the COLD and S-CCD 2.0 were evaluated by analyzing the trade-off between omission errors and commission errors across three critical parameters (P, $CONSE$, λ). We summarized overall accuracy using the F1 score, which provides a balanced measure of omission and commission errors.

To assess computational efficiency, we controlled for differences in observation density by randomly selecting 300 temporal observations for each plot; plots with fewer than 300 observations were excluded from the analysis. Performance was evaluated across band numbers ranging from 1 to 10 for both the retrospective and NRT scenarios. In the retrospective scenario, computational performance was assessed using the two *pyxccd* functions that support user-defined band inputs: *cold_detect_flex()* and *sccd_detect_flex()*. In the NRT scenario, we compared *cold_detect_flex()* with *sccd_update_flex()* following the initial execution of *sccd_detect_flex(),* representing recursive model updating. When the specified number of bands exceeded the available Landsat bands, existing bands were replicated to maintain consistent input dimensionality. Computational efficiency was quantified as the average processing time per plot. All efficiency experiments were conducted on a desktop workstation running Windows 11 Pro 64-bit, equipped with a 13th Gen Intel® Core™ i7-13700 CPU, 64 GB RAM, and an SSD for local data storage. The experiments used Python 3.11.9, NumPy 2.3.4, and pyxccd v1.0.2. The evaluation script with the guidance for downloading testing dataset (accuracy&efficiency_test.ipynb) is available at https://github.com/Remote-Sensing-of-Land-Resource-Lab/pyxccd/tree/devel/tutorials/notebooks.

4.2 Results

Fig. 5 illustrates the comparative accuracy of COLD and S-CCD 2.0 with one varying parameter and the other two remaining default: i.e., $P = 0.99$, CONSE=*6*, and $\lambda = 20$. Both algorithms exhibit a synchronized response to the change probability threshold (Fig. 5A). The transition from a loose threshold (0.9) to a strict one (0.9999) results in an equivalent migration along the error curve for both models. While the COLD algorithm maintains a slightly lower commission rate at specific intervals, the two curves remain in proximity, suggesting that both methods are equally capable of being tuned to achieve a desired balance of detection sensitivity. The peak F1 for both algorithms are also remarkably close, with COLD achieving a maximum of 0.664 (omission: 33.9%; commission: 33.4%) and S-CCD 2.0 reaching 0.653 (omission: 37.9%; commission: 31.2%). The analysis of the peek window length (Fig. 5B) further underscores the functional similarity between the two methods. As the confirmation window expands, both CCDC-like algorithms converge toward a similar "elbow" point. This indicates that both models utilize temporal consistency in a comparable manner to filter out transient noise. Fig. 5C examines the influence of the regularization parameter λ on S-CCD 2.0 in comparison with COLD. Decreasing λ leads to higher commission errors and nonlinear omission errors in S-CCD 2.0, reflecting the coupled role of λ in controlling both model regularization and state-space responsiveness. A smaller λ weakens shrinkage and allows the model to respond more strongly to short-term spectral fluctuations, which can improve sensitivity to subtle disturbances but also increases the risk of interpreting transient noise as real change. The nonlinear behavior of omission error suggests that the influence of λ is not simply proportional to the degree of regularization; rather, it interacts with the Kalman-filter update process, and residual variance. In contrast, COLD shows a more monotonic trade-off,

where reduced regularization is associated with higher commission error and lower omission error, indicating greater model flexibility and stronger responsiveness to observations.

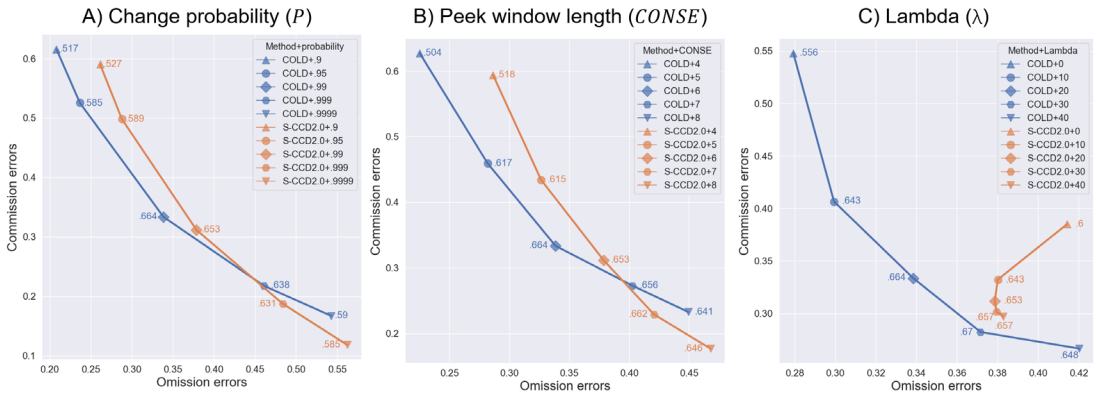

Fig. 5. Accuracy assessment results of COLD and S-CCD 2.0 based on testing dataset (n = 6,488). A) accuracies under a series of change probability thresholds as 0.9, 0.95, 0.99, 0.999 and 0.9999, under default $CONSE$ (6) and λ (20). B) accuracies for $CONSE$ from 4 to 8, under default P (0.99) and λ (20). C) accuracies for a series of λ as 0, 10, 20, 30 and 40, under default P (0.99) and $CONSE$ (6).

Fig. 6 compares the computational efficiency of COLD and S-CCD 2.0 under both retrospective and NRT scenarios as a function of the number of spectral bands. In the retrospective scenario (Fig. 6A), the running time per pixel increases approximately linearly with band number for both methods. However, S-CCD 2.0 consistently exhibits lower computational cost than COLD across all band configurations, requiring approximately two-thirds of the processing time. Moreover, the performance gap widens as the number of bands increases, indicating improved scalability of S-CCD 2.0 for multi-band retrospective processing.

In the NRT scenario (Fig. 6B), S-CCD 2.0 achieves 3–6× faster than COLD. The

computational cost of both methods increases with band number; however, the increase for S-CCD is much less pronounced than for COLD. Because the baseline processing time of S-CCD 2.0 is substantially lower, the growth trend with increasing spectral dimensionality is not visually prominent at this scale. In contrast, COLD exhibits a clear linear increase in running time. These results indicate that S-CCD 2.0 substantially reduces per-update computational overhead in NRT monitoring, and that its state-space-based design and jump-update strategy effectively mitigate the impact of increasing spectral dimensionality, making it well suited for operational NRT applications with high-dimensional spectral inputs.

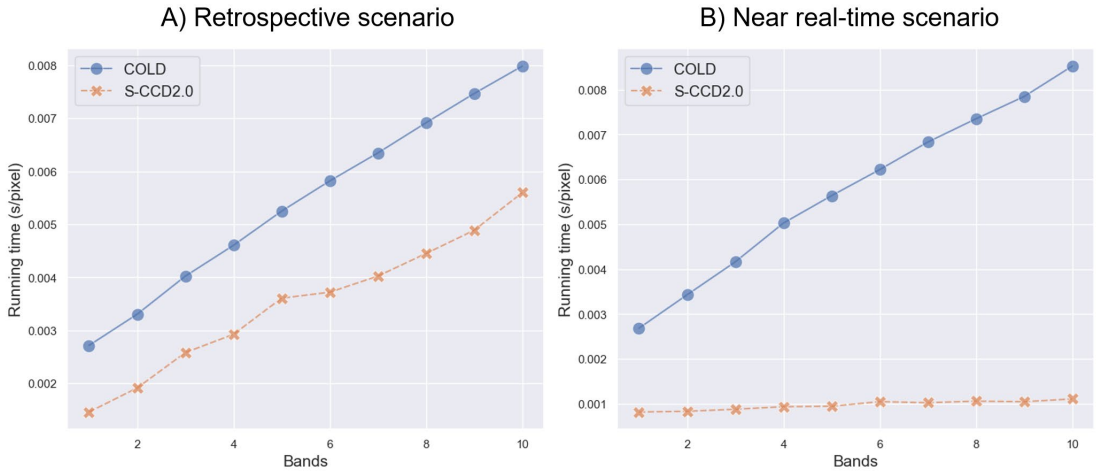

Fig. 6. Efficiency comparison between COLD (model updating at each 3% of the observation number that has not undergone updating with three observations as minimum) and S-CCD 2.0 based on testing dataset (n = 6488). The length of each time-series plot was cut into 300 temporal observations. A) Average processing time per pixel for the retrospective scenario. B) Average processing time per pixel for the NRT scenario.

5. Application

5.1 NRT detection for logging

Fig. 7 illustrates the NRT monitoring workflow of S-CCD 2.0 using *pyxccd* for a suspected logging event in Sichuan Province, China, using the SWIR2 time series. At Week 0, the observations (green dots) are well explained by the fitted seasonal model (orange curve), and no anomalous points are recorded (n_{anom}=0). Beginning at Week 1, an observation at the end of the series deviates markedly from the expected seasonal trajectory and is flagged as an anomaly (n_{anom}=1). At this time, although no break is detected, S-CCD 2.0 has outputted anomaly information at the tail, including the number of anomalies, its spectral reflectance, the current harmonic coefficients (refers to Table A1 for details). As new acquisitions are ingested (Weeks 2–3), additional high-magnitude deviations accumulate (dashed boxes), increasing the anomaly count (n_{anom}= 3 to 4), while the harmonic model is only slightly adjusted by Kalman filter. By Week 4, the sustained and repeated departures from the predicted SWIR2 signal satisfy the S-CCD 2.0 decision rule, and a structural break is declared ("Break detected", vertical line; n_{anom}= 6). The corresponding image chips (2024-03-14 to 2024-08-15) show progressive canopy disturbance and exposed ground material at the target location (blue marker), providing independent visual confirmation that the detected breakpoint (2024-04-08) corresponds to an active logging disturbance.

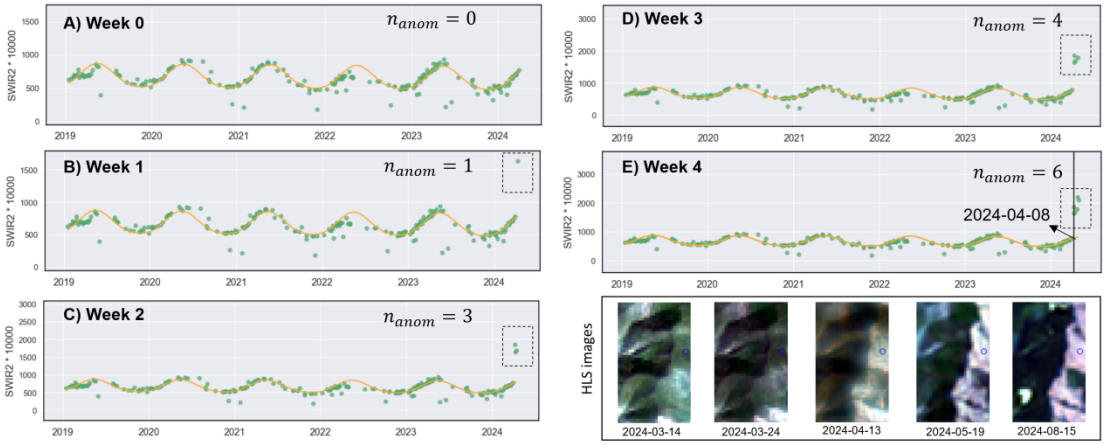

Fig. 7. An example for S-CCD 2.0 detecting a suspected logging event for an NRT scenario at a per-week updating basis in Sichuan, China. While the logging event is fully confirmed as a break in Week 4, S-CCD 2.0 starts outputting anomalies since the Week 1.

5.2 Ecosystem disturbance monitoring

Fig. 8 compares S-CCD 2.0 performance when applied to multiple coarse-resolution ecosystem products over the same location, i.e., 500-m NDVI [23], 500-m LAI [24], 0.05-degree SIF [25], and 500-m GPP [26]. For each variable, S-CCD 2.0 fits a model summed from a trimodal (trend + annual + semiannual + four-month) state-space model (orange curve) to the existing observations (green dots), and identifies departures from the expected trajectory as positive/negative anomalies and breaks. Across all products, the dominant annual phenology is captured well despite differing dynamic ranges and noise characteristics, indicating that the state-space framework can accommodate structurally different biophysical signals. Most departures in NDVI, LAI, and GPP occur as isolated anomalies (blue/red open circles) and are not sufficiently persistent to trigger a regime shift, suggesting short-lived perturbations and/or retrieval noise rather than sustained ecosystem change and hence rationale for using "anomaly-

break" detection hierarchy. In contrast, the SIF series exhibits a more coherent and sustained deviation that leads to a detected structural break (black vertical line), after which subsequent observations are evaluated relative to the updated baseline. Overall, this comparison demonstrates the transferability of S-CCD 2.0 to diverse coarse-resolution indicators and highlights the value of multi-product monitoring, where canopy-structure proxies (NDVI/LAI) and functional proxies (SIF/GPP) can show complementary sensitivity to disturbance signals.

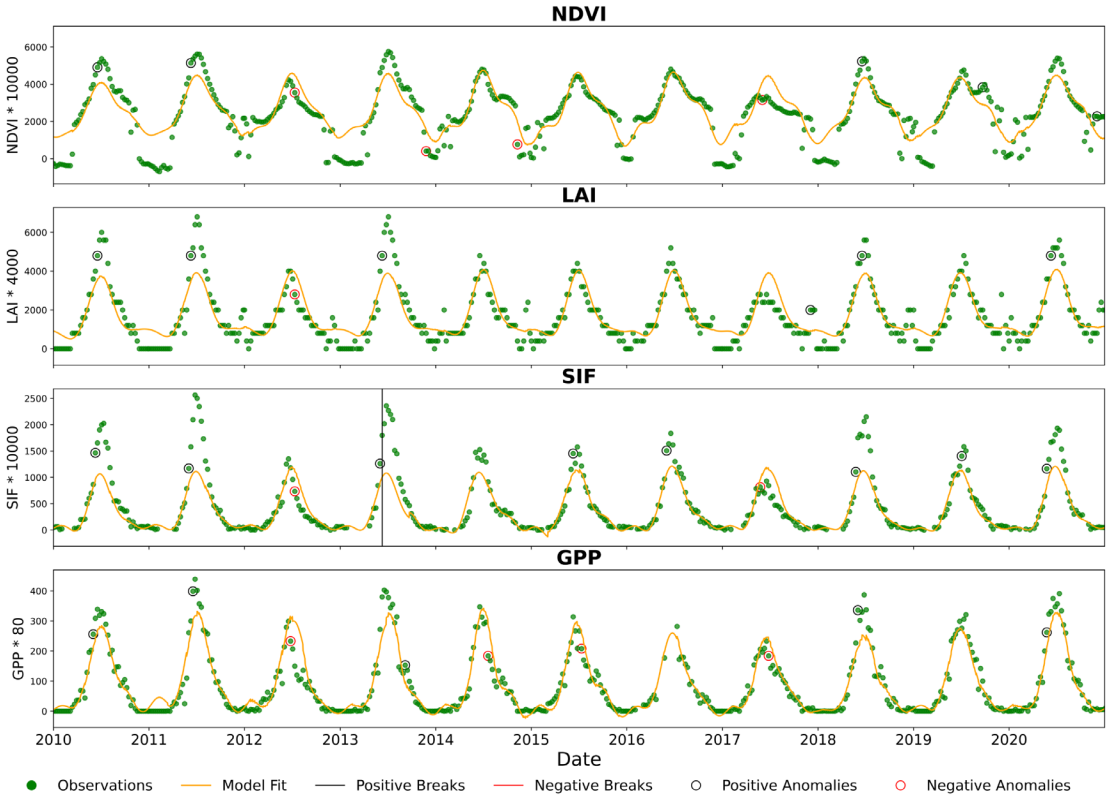

Fig. 8. Cross-product comparison of S-CCD 2.0 detecting anomalies and breaks for coarse-resolution ecosystem disturbance monitoring (NDVI, LAI, SIF, and GPP) based upon a location in Montana, United States.

5.3 Large-scale processing

Fig. 9 compares wall-to-wall disturbance mapping results produced by S-CCD 2.0 and

COLD using *pyxccd* over a large region (Hangzhou City, China), where each pixel is labeled by the recent detected change year (colors for 2019–2024) and unchanged areas are shown in black. At the landscape scale, the two methods yield highly consistent spatial patterns and similar temporal attribution of disturbance, capturing the same dominant clusters and linear features of change across the scene. The three example subsets (C-E) further illustrate this agreement at local scales: both S-CCD 2.0 and COLD delineate nearly identical disturbance footprints and assign comparable change years, and the detected patterns are visually supported by the corresponding Bing Aerial imagery (e.g., land conversion and infrastructure/clearing signatures). Differences between the two outputs are limited to scattered, isolated pixels and minor boundary variations, which are expected given mixed pixels, temporal sampling differences, and retrieval noise. Overall, these results indicate that S-CCD 2.0 achieves disturbance mapping performance that is essentially comparable to the established COLD approach in large-scale processing, while maintaining close spatial and temporal coherence with independent high-resolution reference imagery.

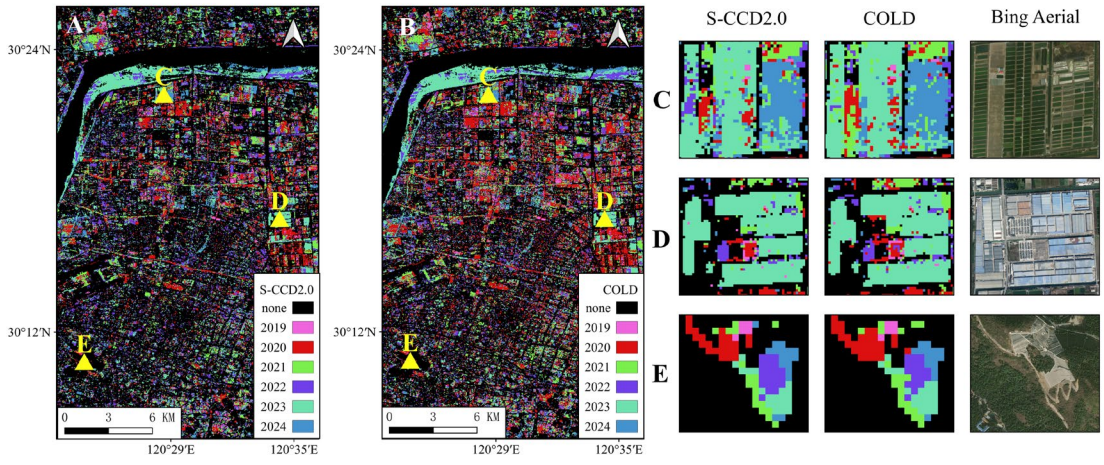

Fig. 9. Large-Scale Comparison of Disturbance Mapping Between S-CCD 2.0 and COLD with Aerial Imagery Validation

6. Conclusion

This study introduces *pyxccd*, an open-source Python package for scalable, continuous analysis from satellite-based time series dataset. It integrates high-performance C kernels with a user-friendly API, supporting COLD and S-CCD 2.0 algorithms. *Pyxccd* improves efficiency, offering 1.4–1.9× faster processing and practical utility for NRT disturbance detection and ecosystem monitoring.

Appendices

Table A1: the NRT model structure (*nrt_model*) is retained by S-CCD 2.0 to analyze the most recent anomalies and enable recursive update of the model coefficients.

Element name	Data type	Usage	Description
t_start_since1982	Short	General	The start date of the current segment, formatted as ordinal date – 723742
obs	2-d array	Peek window	The last eight multiband observations.
obs_date_since1982	1-d array	Peek window	The dates for the last eight observations
covariance	2-d array	State-space	Multiband covariance matrix (each 2-d covariance matrix is flatten to 1-d array)
nrt_coefs	2-d array	State-space	Harmonic coefficients converted from states at the last updating date (see Eq. A7-12 in the appendix)
H	1-d array	State-space	Multiband observation noise
rmse_sum	1-d array	RMSE update	Multiband sum of squared errors
num_obs	Short	RMSE update	Current number of clear observations that have been processed
anomaly_conse	Byte	Anomaly analysis	The consecutive anomaly observation number at the tail
cm_angle	Short	Anomaly analysis	The included angle for the anomalies at the tail
norm_cm	Short	Anomaly analysis	The normalized change magnitude for the anomalies at the tail

A1. Conversion between harmonic coefficients and states

Converting harmonic coefficients (a_0, c_1, a_1, b_1, a_2, b_2) to states (μ_t, v_t, $\gamma_{1,t}$, $\gamma^*_{1,t}$, $\gamma_{2,t}$, $\gamma^*_{2,t}$) at time t:

$$\mu_t = a_0 + c_1 t \quad \text{(A1)}$$

$$v_t = c_1 \quad \text{(A2)}$$

$$\gamma_{1,t} = a_1 \cos\left(\frac{2\pi}{m} t\right) + b_1 \sin\left(\frac{2\pi}{m} t\right) \quad \text{(A3)}$$

$$\gamma_{1,t}^{*} = -\mathrm{a}_1 \sin\left(\frac{2\pi}{m}t\right) + \mathrm{b}_1 \cos\left(\frac{2\pi}{m}t\right) \quad \text{(A4)}$$

$$\gamma_{2,t} = \mathrm{a}_2 \cos\left(\frac{4\pi}{m}t\right) + \mathrm{b}_2 \sin\left(\frac{4\pi}{m}t\right) \quad \text{(A5)}$$

$$\gamma_{2,t}^{*} = -\mathrm{a}_2 \sin\left(\frac{4\pi}{m}t\right) + \mathrm{b}_2 \cos\left(\frac{4\pi}{m}t\right) \quad \text{(A6)}$$

Converting states at time t to harmonic coefficients:

$$\mathrm{a}_0 = \mu_t - \mathrm{c}_1 t \quad \text{(A7)}$$

$$\mathrm{c}_1 = v_t \quad \text{(A8)}$$

$$\mathrm{a}_1 = \gamma_{1,t}\cos\left(\frac{2\pi}{m}t\right) - \gamma_{1,t}^{*}\sin\left(\frac{2\pi}{m}t\right) \quad \text{(A9)}$$

$$\mathrm{b}_1 = \gamma_{1,t}\cos\left(\frac{2\pi}{m}t\right) + \gamma_{1,t}^{*}\sin\left(\frac{2\pi}{m}t\right) \quad \text{(A10)}$$

$$\mathrm{a}_2 = \gamma_{2,t}\cos\left(\frac{4\pi}{m}t\right) - \gamma_{2,t}^{*}\sin\left(\frac{4\pi}{m}t\right) \quad \text{(A11)}$$

$$\mathrm{b}_2 = \gamma_{2,t}\cos\left(\frac{4\pi}{m}t\right) + \gamma_{2,t}^{*}\sin\left(\frac{4\pi}{m}t\right) \quad \text{(A12)}$$

A2 Key modules of pyxccd

The modules were introduced based upon *pyxccd* v1.0.2, released at Jan. 7, 2026. The package is distributed through PyPI, requires Python >=3.9, and is licensed under Apache-2.0.

(1) Change detection module

The *pyxccd.ccd* module is the computational heart of the package, containing functions to analyze spectral time-series data and detect land-surface changes.

ccd_detect(): the primary function for COLD, it analyzes historical land-surface data by accepting observation dates and spectral inputs. Users can adjust detection parameters, returning a structured array of temporal segments with metadata, including break dates, model coefficients, and spectral change magnitudes.

sccd_detect(): this function performs the S-CCD 2.0 for retrospective scenarios,

establishing the initial state-space model and returning a *SccdOutput* structure. It offers optional outputs, including a data frame of states for each band and an anomaly detection result. Users can control critical S-CCD parameters, including CONSE, P, conse, p, λ, and interval days for outputting states.

sccd_update(): this function supports NRT monitoring by updating an existing *SccdOutput* structure using only newly collected data, minimizing memory overhead and computational cost.

Flexible Variants: *Pyxccd* includes flexible variants like *cold_detect_flex()* and *sccd_detect_flex()*, which allow users to input a wide range of time-series data (e.g., NDVI, GPP, LAI) and extend the model to support complex ecosystems with multiple growing seasons through a trimodal harmonic model.

(2) Land-cover Classification Module

The *pyxccd.pyclassifier* module implements the CCDC framework to translate structural breaks and harmonic coefficients into thematic land-cover maps.

Through the *get_features()* function, the library extracts a high-dimensional feature set from the *cold_rec_cg* or *SccdOutput* structures. These features include the intercept, slope, and the sine/cosine coefficients of the harmonic cycles for each spectral band. The module is model-agnostic but includes optimized wrappers for a random-forest classifier, allowing users to train models on historical data and apply them consistently across a temporal segment until a structural break is detected.

(3) Spatial orchestration module

The *pyxccd.imagetool* module provides a high-level workflow to scale pixel-based

algorithms to large-scale geospatial datasets. First, the *prepare_ard.py* utility handles the pre-processing and ingestion of time-series images, standardizing multi-temporal image stacks and quality masks into a format optimized for time-series analysis and partitioning extensive geographic regions into manageable spatial blocks to enable efficient memory management and multi-core parallelization. Second, *tile_processing.py* serves as the primary execution driver. This script manages the distribution of data to the core *pyxccd.ccd* engines and gathers the resulting model coefficients and break records. Finally, *export_change_map.py* facilitates the synthesis of these results, mosaicking the processed tiles and translating the raw detection outputs into standardized geospatial products. This automated pipeline enables researchers to generate comprehensive raster layers with minimal manual intervention.

(4) Utilities and data integration module

The *pyxccd.common* and *pyxccd.utils* modules provide standardized Python data structures, including *anomaly_dt* and *cold_rec_cg*, for consistent data handling across algorithms. These utilities support conversion, validation, and export of algorithm outputs for external analysis tools and mapping workflows.

CRediT authorship contribution statement

S. Y.: Writing – original draft, Validation, Software, Methodology, Formal analysis, Supervision, Funding acquisition, Conceptualization. Y. H.: Writing – review & editing, Validation, Software. T. C.: Visualization, Validation.

Data availability statement

The *pyxccd* source code is publicly available through PyPI and GitHub [27]. The package version described in this manuscript is V1.0.2. The data that support the findings of this study are available from the corresponding author, S.Y., upon request.

Conflict of Interest statement

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

Funding information

The authors are grateful for the financial support provided by Major Science and Technology Collaboration Program of Zhejiang Province (2025SNJF012) and National Natural Science Foundation of China (42401393).

Reference

[1] Z. Zhu, "Change detection using landsat time series: A review of frequencies, preprocessing, algorithms, and applications," *ISPRS J. Photogramm. Remote Sens.*, vol. 130, pp. 370–384, 2017.

[2] S. Ye *et al.*, "Detecting subtle change from dense Landsat time series: Case studies of mountain pine beetle and spruce beetle disturbance," *Remote Sensing of Environment*, vol. 263. 2021. doi: 10.1016/j.rse.2021.112560.

[3] J. E. Vogelmann, A. L. Gallant, H. Shi, and Z. Zhu, "Perspectives on monitoring gradual change across the continuity of Landsat sensors using time-series data," *Remote Sens.*

Environ., vol. 185, pp. 258–270, 2016.

[4] Z. Zhu, S. Qiu, and S. Ye, "Remote sensing of land change: A multifaceted perspective," *Remote Sens. Environ.*, vol. 282, p. 113266, 2022.

[5] C. E. Woodcock, T. R. Loveland, M. Herold, and M. E. Bauer, "Transitioning from change detection to monitoring with remote sensing: A paradigm shift," *Remote Sens. Environ.*, vol. 238, p. 111558, 2020.

[6] H. West, N. Quinn, and M. Horswell, "Remote sensing for drought monitoring & impact assessment: Progress, past challenges and future opportunities," *Remote Sens. Environ.*, vol. 232, p. 111291, 2019.

[7] C. F. Brown *et al.*, "Dynamic World, Near real-time global 10 m land use land cover mapping," *Sci. Data*, vol. 9, no. 1, pp. 1–17, 2022.

[8] S. Ye, Z. Zhu, and J. W. Suh, "Leveraging past information and machine learning to accelerate land disturbance monitoring," *Remote Sens. Environ.*, vol. 305, p. 114071, 2024.

[9] G. Welch and G. Bishop, "An introduction to the Kalman filter," 1995.

[10] R. S. Tsay, *Analysis of financial time series*. John wiley & sons, 2005.

[11] Z. Zhu and C. E. Woodcock, "Continuous change detection and classification of land cover using all available Landsat data," *Remote Sens. Environ.*, vol. 144, pp. 152–171, 2014.

[12] M. A. Friedl *et al.*, "Medium spatial resolution mapping of global land cover and land cover change across multiple decades from landsat," *Front. Remote Sens.*, vol. 3, p. 894571, 2022.

[13] X. Zhang, L. Liu, X. Chen, Y. Gao, S. Xie, and J. Mi, "GLC_FCS30: Global land-cover product with fine classification system at 30 m using time-series Landsat imagery," *Earth Syst. Sci. Data Discuss.*, vol. 2020, pp. 1–31, 2020.

[14] J. F. Brown *et al.*, "Lessons learned implementing an operational continuous United States national land change monitoring capability: The Land Change Monitoring, Assessment, and Projection (LCMAP) approach," *Remote Sensing of Environment*. p. 111356, 2019. doi: 10.1016/j.rse.2019.111356.

[15] G. Z. Xian *et al.*, "Implementation of the CCDC algorithm to produce the LCMAP Collection 1.0 annual land surface change product," *Earth Syst. Sci. Data*, vol. 14, no. 1, pp. 143–162, 2022.

[16] S. Ye, Z. Zhu, and G. Cao, "Object-based continuous monitoring of land disturbances from dense Landsat time series," *Remote Sens. Environ.*, vol. 287, p. 113462, 2023.

[17] P. Arévalo, E. L. Bullock, C. E. Woodcock, and P. Olofsson, "A suite of tools for continuous land change monitoring in google earth engine," *Front. Clim.*, vol. 2, p. 576740, 2020.

[18] Z. Zhu *et al.*, "Continuous monitoring of land disturbance based on Landsat time series," *Remote Sens. Environ.*, vol. 238, p. 111116, 2020.

[19] S. Ye, J. Rogan, Z. Zhu, and J. R. Eastman, "A near-real-time approach for monitoring forest disturbance using Landsat time series: stochastic continuous change detection," *Remote Sensing of Environment*, vol. 252. 2021. doi: 10.1016/j.rse.2020.112167.

[20] Z. Zhu and C. E. Woodcock, "Object-based cloud and cloud shadow detection in Landsat imagery," *Remote Sens. Environ.*, vol. 118, pp. 83–94, 2012.

[21] M. Claverie, E. F. Vermote, B. Franch, and J. G. Masek, “Evaluation of the Landsat-5 TM and Landsat-7 ETM+ surface reflectance products,” *Remote Sens. Environ.*, vol. 169, pp. 390–403, 2015.

[22] W. B. Cohen *et al.*, “Forest disturbance across the conterminous United States from 1985–2012: The emerging dominance of forest decline,” *For. Ecol. Manage.*, vol. 360, pp. 242–252, 2016.

[23] C. Schaaf and Z. Wang, “MODIS/Terra+ aqua BRDF/Albedo nadir BRDF adjusted ref daily L3 global-500m V061,” *NASA EOSDIS L. Process. Distrib. Act. Arch. Cent. data set*, p. MCD43A4-061, 2021.

[24] R. Myneni, Y. Knyazikhin, and T. Park, “MODIS/Terra Leaf Area Index/FPAR 8-Day L4 Global 500m SIN Grid V061,” *NASA EOSDIS L. Process. Distrib. Act. Arch. Cent. data set*, p. MOD15A2H-061, 2021.

[25] X. Li and J. Xiao, “A global, 0.05-degree product of solar-induced chlorophyll fluorescence derived from OCO-2, MODIS, and reanalysis data,” *Remote Sens.*, vol. 11, no. 5, p. 517, 2019.

[26] D. A. Sims *et al.*, “A new model of gross primary productivity for North American ecosystems based solely on the enhanced vegetation index and land surface temperature from MODIS,” *Remote Sens. Environ.*, vol. 112, no. 4, pp. 1633–1646, 2008.

[27] Remote-Sensing-of-Land-Resource-Lab, 2026. pyxccd v1.0.2 [software]. GitHub/PyPI. https://github.com/Remote-Sensing-of-Land-Resource-Lab/pyxccd (accessed 19 May 2026).