nature communications

Article

https://doi.org/10.1038/s41467-026-78489-7

Estimating the resilience of non-stationary systems

Received: 28 April 2026

Accepted: 21 September 2026

Published online: 06 November 2026

Check for updates

Taylor Smith [1,2] ✉, Andreas Morr [3,4], Christof Schötz [2,4] & Niklas Boers [2,4]

A wide body of work has applied the concept of critical slowing down to estimate the stability of different Earth system components. Most of them – such as global vegetation – are inherently non-stationary, for example due to strong seasonal forcing that complicates the estimation of their resilience to external perturbations. Here, we introduce a method to account for non-stationarity in estimating the resilience of diverse synthetic and real-world systems via a regression-based formulation of the Langevin equation. Our method does not require extensive data pre-processing, is robust to gaps in the data record, and does not require regular time sampling. We further show that our method can incorporate time-varying data uncertainties, derive uncertainty bounds for stability estimates, and can be natively extended to examine spatial systems. Our method is a drop-in replacement for widely-used autocorrelation-based resilience estimates, and can be applied across diverse Earth system components.

There has been substantial recent research focused on understanding the stability and resilience of different Earth system components[1–7]; much of this work is based on the concept of 'critical slowing down' (CSD), which uses slowing system dynamics when approaching critical transition points to provide a warning of oncoming state changes[2,8–10]. Lag-1 autocorrelation (AC1) and variance are the most commonly used early warning indicators, though there exist several other approaches, such as spatial indicators[11], empirical and regression-based estimates[6,12], eigenvalue-based approaches[13–15], moment-based reconstructions[16,17], spectral properties[18,19], and deep learning[20,21].

The majority of CSD metrics rely on (1) evenly-sampled time series and (2) the assumption of data stationarity (i.e., that it is trend- and seasonality-free) over a chosen temporal analysis window. While many instrument records are inherently evenly sampled, this is not always the case when considering, for example, paleoclimate proxy records or using composite products based on multiple data sets. Data stationarity is often accomplished by pre-processing seasonal or otherwise non-stationary data (e.g., paleoclimate records with Milankovich cycles[2]) using statistical deseasoning and detrending techniques[22,23]. The reliance on stationary data is due to the core assumption that a system can be linearized around a stable fixed point, and hence can define an Ornstein-Uhlenbeck process[24] with a damping rate controlling how fast a given external perturbation decays back towards the current equilibrium[7]. This framing motivates the use of, e.g., AC1 and variance as early-warning signals of oncoming critical transitions, and has been widely applied across disciplines[3,7,10,14].

Existing CSD frameworks, hence, require error-prone data pre-processing to remove time-dependence (e.g., seasonality)[22,25]; significant issues with the calculation of AC1 and variance on real-world data and data with gaps have also been found in recent work[26]. Furthermore, processes independent of stability loss – such as changes in sensor signal-to-noise ratios[25,27] – can also drive changes in AC1 and variance. Current methods also struggle to integrate data uncertainty into stability estimates[28,29]; this is a critical limitation for many data sets – such as paleoclimate proxies – which have inherent and time-varying uncertainties in their measurements as well as in corresponding age estimates[30] that are often not captured using current approaches. The propagation of deseasoning and detrending errors into stability estimates is also not common practice, despite clear evidence that the choice of pre-processing strategy influences these estimates[13,22].

[1]Institute of Geosciences, Universität Potsdam, Potsdam, Germany. [2]Munich Climate Center and Earth System Modelling Group, Department of Aerospace and Geodesy, TUM School of Engineering and Design, Technical University of Munich, Munich, Germany. [3]Department of Mathematics, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany. [4]Research Domain IV–Complexity Science, Potsdam Institute for Climate Impact Research, Potsdam, Germany. ✉e-mail: tasmith@uni-potsdam.de

In this work, we introduce a regression-based approach to estimating the resilience (quantified in terms of the recovery rate λ) of a given system using the Langevin equation that can natively handle irregularly sampled data, non-stationary data, data with time-varying uncertainties, and spatially extended systems. We first motivate our approach within CSD theory, and illustrate its advantages using a synthetic time series model. We further apply our direct recovery rate estimation method to real-world data that has, until now, required several pre-processing steps, including vegetation dynamics, paleoclimate proxies, and glacier surges, and show how it can add additional context, such as uncertainty estimates, to the widely-used stability indicator λ.

Results

Theoretical framing

A general framing of the stability of a dynamical system can be motivated by assuming that a system state $x(t)$ is a stochastic Markov process characterized by a time-dependent attractor $\mu(t)$; that attractor $\mu(t)$ can be fixed, drifting, or oscillating (i.e., a limit cycle). The dynamics of that system can be described using the Langevin equation:

$$\frac{dx}{dt} = f(x,t) + \sigma(t)\eta(t) \tag{1}$$

where the deterministic function $f(x, t)$ yields the time-dependent attractor and describes the restoring force against perturbations, and $\eta(t)$ are Gaussian white noise disturbances with magnitude $\sigma(t)$ (i.e., noise intensity). As $\mu(t)$ captures the system's deterministic trajectory, evaluating $f(x, t)$ exactly on the attractor yields the attractor's trajectory gradient, here termed $\frac{d\mu}{dt}$. If we assume that the system is near the attractor $\mu(t)$, the system's fast and noisy dynamics are largely governed by the restoring forces encoded by $f(x, t)$ around $\mu(t)$. These forces can be described as approximately linear at each point in time, yielding:

$$\frac{dx}{dt} = \frac{d\mu}{dt} - \lambda(t)(x(t) - \mu(t)) + \sigma(t)\eta(t) \tag{2}$$

where $\lambda(t)$ is the restoring rate (i.e., gradient of the restoring force), or how strongly perturbations decay towards the attractor. It is this linear restoring force that is typically regarded as a measure of system resilience, with a decrease towards zero signaling a loss of resilience. Equation (2) is the basis for most common techniques of estimating λ, as removing the suspected position of the attractor from time series data yields a process that should restore towards equilibrium. In such a process, λ is directly accessible via proxy statistics such as AC1 and variance. However, the attractor location $\mu(t)$ is typically not known, and its movement $\frac{d\mu}{dt}$ is often disregarded, meaning that any attempt to remove non-stationary components by means of detrending or deseasoning can bias λ estimates in complex ways.

Here, we propose taking the functional form of the attractor as a starting point. In many natural systems (such as seasonally forced systems), the attractor's movement is well-constrained and can be functionally modeled. The instantaneous change $f(x, t)$ in the system state can thus be approximately decomposed into a part that depends on x explicitly and linearly, and a part that arises purely from external forcing and is thus a function of time t.

To analyze changes in λ through time, we construct a set of moving (temporal) windows. Within each moving window, we assume our system dynamics f to be quasi-stationary (i.e., the resilience λ and noise strength σ are constant). This means that while the system can be globally non-stationary (e.g., seasonal cycles or trends, including changes in resilience), we assume that any slow external changes (e.g., climate change) can be neglected over one temporal observation window, whose size needs to be chosen according to the speed of the system dynamics. We also assume that external perturbations (e.g., rainfall over a spatial domain) are homogeneous; this allows us to treat spatial variations in $x(t)$ as being caused by random noise and the internal dynamics and feedbacks responding to that noise, rather than by spatial variability in external perturbations. This is trivially the case for one-dimensional time series, but must be accounted for when considering spatio-temporal data.

Under the above assumptions, we can treat the restoring rate $\lambda(t)$ as approximately constant within a short temporal window. We can thus discretize Equation (2) in time using the observations $x_1, \ldots, x_n$ taken at time points $t_1, \ldots, t_n$:

$$\frac{\Delta x_i}{\Delta t_i} := \frac{x_{i+1} - x_i}{t_{i+1} - t_i} \approx \underbrace{-\lambda}_{\beta_1} x_i + \underbrace{\left[\lambda\mu(t_i) + \frac{d\mu}{dt}(t_i)\right]}_{m(t_i)} + \epsilon_i \tag{3}$$

where $\frac{\Delta x_i}{\Delta t_i}$ is the discrete rate of change of x and ϵ_i is discrete stochastic noise; we group the attractor dynamics (μ and $\frac{d\mu}{dt}$) together for clarity as they both depend on time t_i. This formulation yields a linear regression problem where λ is the negative of the slope of the regression (β_1), and both the position and movement of the attractor are contained within $m(t_i)$. The residual term ϵ_i captures any deviations from the estimated attractor μ. A diminishing contribution of the system state to the regression fit – i.e., a decreasing estimate of λ over several temporal windows – can be interpreted analogously as a decrease in system resilience. We note that the temporal scale of the recovery rate λ depends entirely on the chosen moving window; λ is estimated over a given temporal window, meaning that sub-annual temporal windows would resolve sub-annual variation in λ, and multi-annual windows would yield λ estimates averaged over that entire span.

For our analysis, we assume that the strength of process noise ($\sigma\eta$) is substantially larger than any observation or measurement noise (δ). As we estimate λ by regressing the increments $\frac{\Delta x_i}{\Delta t_i}$ on the states x_i, any measurement noise δ will influence both variables: a positive error δ_i in x_i will induce a negative error in $\frac{\Delta x_i}{\Delta t_i}$. Our λ estimates thus act as an upper bound on the true stability of the system; if we assume that $\sigma\eta \gg \delta$, this upward bias becomes negligible.

Simultaneous estimation of λ and μ. For a simple system with no expressions of seasonality or linear trends, any drift in the system collapses into the regression intercept ($\beta_0 = m(t_i)$), which contains both λ and the attractor μ. This assumes that the external forcing is negligible within the analysis window (e.g., slow climate change) and that there is no inherent non-stationarity in the system (i.e., it does not have any seasonal oscillations or trends).

Most systems on Earth are, however, not so simple – many systems are driven by oscillating forcing, such as the daily or annual cycles in temperature and sunlight that are central to their dynamics. To capture such an oscillating attractor, we need to correctly pose the linear regression problem (Equation (3)); for the case where μ contains a seasonal component with a known period (e.g., one year), we can approximate it with a Fourier series of k harmonic terms:

$$\frac{\Delta x_i}{\Delta t_i} = \underbrace{-\lambda}_{\beta_1} x_i + \underbrace{\sum_{k=1}^{n}\left[A_k \sin(k\omega t_i) + B_k \cos(k\omega t_i)\right] + \beta_0}_{m(t_i)} + \epsilon_i \tag{4}$$

where $m(t_i)$ is a time-varying intercept modeled using A_k and B_k as the harmonic fitting terms for frequency ω that constitute the linear regression factors beyond the linear restoring rate λ. As the derivative of a harmonic series yields another harmonic series, both $\lambda\mu(t_i)$ and $\frac{d\mu}{dt}(t_i)$ are absorbed into the coefficients A_k and B_k; β_0 captures any constant offset in the attractor position. We choose $n \geq 2$ to allow for asymmetric seasonality (e.g., steeper greening than browning phases in vegetation) in complex real-world data. We note that the regression

model defined in Equation (3) can also be reformulated to incorporate simpler (e.g., linear drift) or more complex components (e.g., Milankovich cycles in paleoclimate records), depending on the system under study and its assumed attractor shape. We incorporate the assumed attractor shape directly into our regression design matrix; this approach allows for flexibly modeling diverse systems, as well as combining multiple assumed forcing regimes (e.g., both harmonic and linear forcings on the system).

For any chosen attractor model, we can then capture λ via regression; in the simplest case, an Ordinary Least Squares regression can be used. For systems with known uncertainty (e.g., with independent error estimates for each sample), we can also use a weighted regression to estimate λ. In principle, other regression models can also be substituted; here we use an outlier-minimizing robust regression ("Methods"), which is flexible to both small (e.g., one-dimensional time series) and large (e.g., extensive spatio-temporal grids) data.

Relationship to widely-used CSD indicators. Several proxy statistics for estimating the restoring rate λ have been proposed in the literature, the most commonly used being AC1 and variance. Both of these metrics, however, can be polluted by signals which are unrelated to any changes in internal dynamics[25,27,28]. For an Ornstein-Uhlenbeck process driven by white noise σ, λ estimated from variance (here termed λ_{Var}) is derived from both stability and driving noise strength ($\mathrm{Var}(x) = \frac{\sigma^2}{2\lambda\Delta t}$); a change in variance can be equally caused by a loss of resilience (λ moving towards zero) or increased driving noise (increasing σ). The widely-used autocorrelation-based λ formulation (here termed λ_{AC1}) is estimated as $\mathrm{AC1} = e^{\lambda\Delta t}$; both metrics require that the system fluctuates around a constant mean (i.e., it does not have seasonal oscillations that can bias variance and autocorrelation). Further, the system must be continuously sampled at a fixed interval Δt without gaps, as λ_{AC1} relies on sequential pairs of fixed Δt representing a constant lag and λ_{Var} is influenced by both changes in data density (i.e., the sample variance can be biased by more summer than winter measurements) and data spacing (i.e., the amount of random noise between measurements scales with Δt, making λ_{Var} impossible to cleanly invert for changing Δt). In principle, both of these problems can be solved via data pre-processing schemes, although interpolation (either to fill gaps or to fit an even time sampling) will, by construction, increase autocorrelation in a time series, and the choice of deseasoning procedure can introduce complex biases in both λ_{AC1} and λ_{Var}[22]. These same stationarity and sampling limitations apply to other moving-window based λ estimation techniques, such as those using Kramers-Moyal coefficients[15] or dynamic eigenvalues[13].

Our regression-based estimate of λ (Equation (3)), however, does not require deseasoning and detrending to stationarity; any deterministic time dependence (e.g., seasonality, Equation (4)) can be modeled explicitly via covariates, allowing λ and the (time-varying) attractor μ to be estimated simultaneously. Further, by using $\frac{\Delta x_i}{\Delta t_i}$ rather than constant steps ($\Delta t_i = \Delta t$), we are not bound by fixed sampling intervals; this means that we can accommodate irregularly-sampled and incomplete time series without interpolation. We note, however, that our approach does not directly yield a continuous-time estimate of λ, as is the case for λ_{AC1} and λ_{Var}, which arise directly from the analytical solution to an Ornstein-Uhlenbeck process[24]. We rather recover a discrete-time, finite-difference estimate (λ_d), which must be converted to continuous time (λ_c) before comparison with λ_{AC1} and λ_{Var}, as well as to correct for changes in sampling density (Δt) through time ("Methods"). For all analyses presented here, we use the corrected, continuous-time, λ_c unless explicitly specified, referring to it as λ for simplicity.

If we consider a simple dynamical system with seasonality moving towards a state transition (Fig. 1), we can assess the robustness of the regression-based λ to randomly-distributed gaps and compare its performance to the common practice of first deseasoning the data and then computing λ on the nominally stationary residual time series. We model seasonality in the temporal moving windows via harmonic terms in the regression design matrix (Equation (4)); we refer to this approach generally as parametric regression.

We find that our parametric regression-based λ is robust to gaps, even when they are extremely numerous ("Methods", Fig. 1b) or cover substantial periods (e.g., a six-month winter gap, Supplementary Fig. S1). We further find that our parametric regression approach successfully accounts for seasonal oscillations, and that the computed λ on both the seasonal and a non-seasonal control model (i.e., the same model formulation with an annual seasonal amplitude of zero) is very similar (Fig. 1c). This is not the case for common deseasoning methods such as Seasonal Trend Decomposition via Loess (STL)[23], removing harmonic seasonality[22], and using the climatological mean[4], which do not necessarily preserve the full red-noise spectrum of the underlying data and hence can substantially bias the estimation of λ (Fig. 1c).

We can also natively extend our regression-based approach to the spatial context by collapsing the spatial dimension before regression ("Methods"). As in the one-dimensional case (Fig. 1), we note that our parametric regression with harmonic terms successfully mimics the non-seasonal control run, while common pre-processing schemes introduce significant bias into λ estimates (Supplementary Fig. S2). While the extension of our method to the spatial domain is of potentially great benefit, there are several assumptions that need to be made about the spatial heterogeneity of forcing that preclude its use in some cases ("Methods").

Our parametric regression approach has three further benefits beyond being robust to gaps and accounting for non-stationary attractors via regression terms: (1) we can handle uncertainty and measurement errors robustly using a weighted regression; (2) we recover uncertainty estimates of λ that explicitly account for errors in the assumed attractor shape; and (3) we implicitly account for uneven time sampling, removing the need for temporal interpolation. We use three illustrative examples to show how we can capture λ in settings that previously required extensive pre-processing and data cleaning.

Global vegetation resilience

Several studies have examined global vegetation resilience patterns[4,12,31–33]; however, estimates of vegetation resilience change vary widely with the employed method and data set. There is also substantial work showing that common pre-processing strategies bias estimates of λ[22], as well as strong evidence for gaps driving at least some of the spatial diversity in vegetation resilience trends[26]. If we re-assess one commonly-used data set (kNDVI from MODIS[34,35]), we can directly compare λ estimates over two distinct time periods covering half of the data record each (2001–2012 and 2013–2025, Fig. 2). We use both a harmonic and a linear term in our parametric regression model to account for both slow changes in the mean state – such as greening or browning – and annual seasonality.

We confirm previous results (i.e., the broad-scale patterns of λ and λ decreases[22]) with our method that is (1) robust to gaps (Fig. 1) and (2) handles seasonality and trends via parametric regression (Equation (4)) rather than statistical deseasoning techniques. We document decreasing stability in many – but not all – ecosystems, implying longer recovery times from perturbations, as has previously been reported[4,12,22]. We further note that the spatial patterns of λ derived using our method and via AC1 or variance are broadly similar (Supplementary Fig. S3), though they are strongly influenced by the choice of deseasoning procedure.

We cannot confirm or refute large-scale changes in Amazon vegetation resilience[13,31,32] that have been previously argued for based on Vegetation Optical Depth[12,31,32,36] data. We attribute the lack of spatially coherent changes in λ in the Amazon to the inherent

https://doi.org/10.1038/s41467-026-78489-7

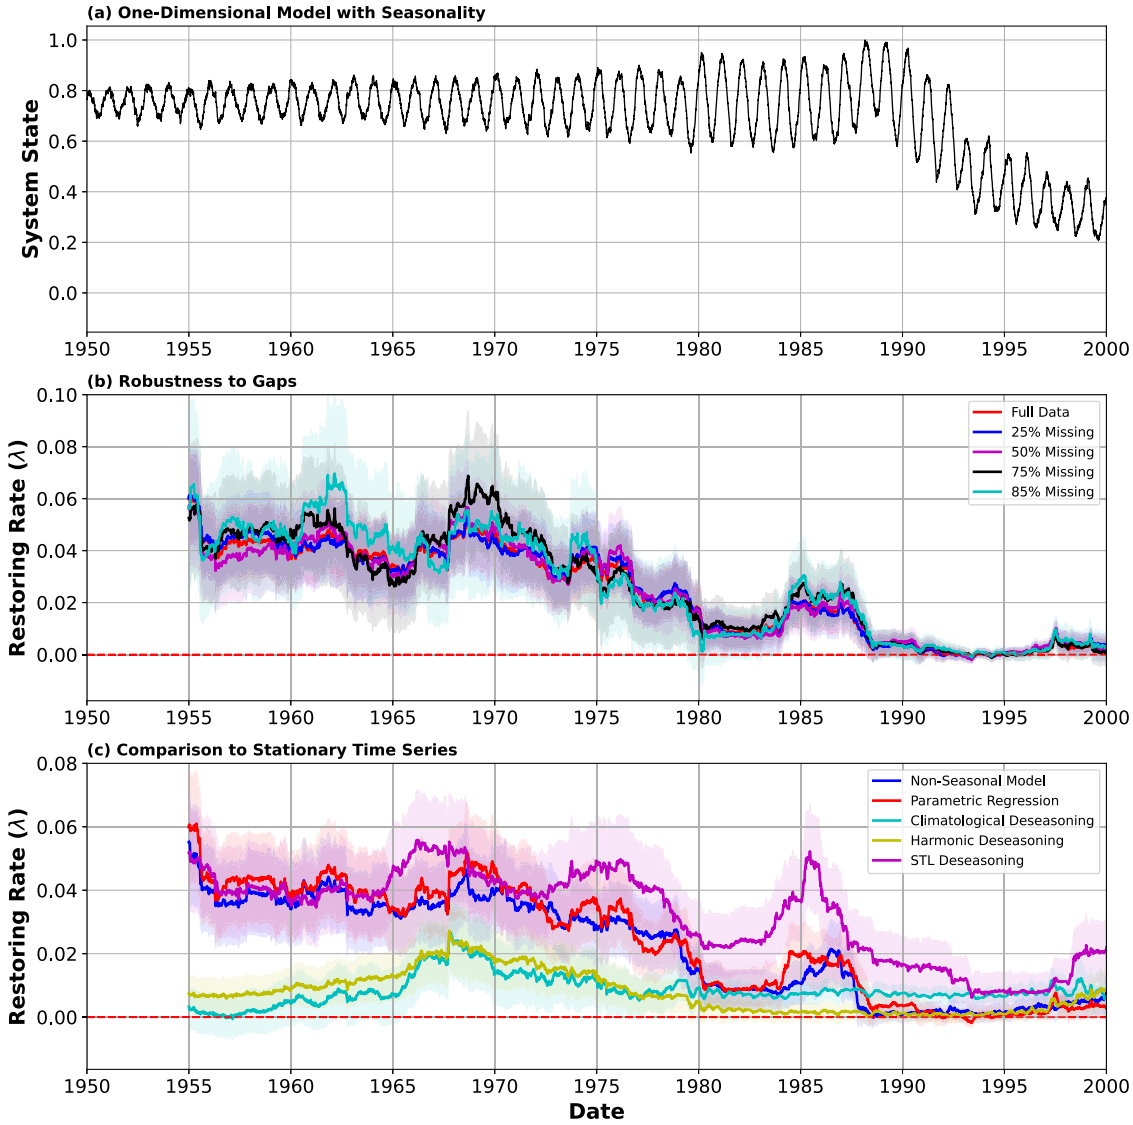


Fig. 1 | Stability of a simple time series model. a Pitchfork bifurcation model with seasonality ("Methods"). **b** Restoring rate λ estimated using parametric regression with harmonic terms on realizations with different percentages of randomly distributed gaps. **c** λ estimated on a non-seasonal control model (blue) compared to estimates on a seasonal model using parametric regression (red) as well as typical deseasoning approaches (teal, yellow, purple). Shaded bounds (**b**, **c**) cover one standard deviation uncertainty in λ ("Methods"). Note that λ estimated on our seasonal model using parametric regression yields the most similar results to the non-seasonal control λ ("Methods").

limitations of kNDVI data in dense and cloud-covered forests where optical vegetation indices tend to saturate and lose the ability to concisely capture ecosystem dynamics.

Paleoclimate proxy data

Paleoclimate proxies are a difficult data source with uneven sampling rates and diverse uncertainty estimates; previous approaches, e.g., ref. 37–39 have used a range of pre-processing techniques to make the data more tractable within the CSD framework. We use the raw NGRIP ice core data[40,41], without pre-processing or resampling, to assess changes in λ through time; this data has been used to argue for signs of CSD before Dansgaard-Oeschger (DO) events[39,42], which mark rapid shifts in Northern-Hemisphere climate with global imprints[43], although recent work including many other Greenland ice cores provides a more inconclusive picture regarding the presence of CSD prior to DO events[37]. In contrast to previous work, we incorporate both measurement error[40] and timing uncertainty into our λ estimates, and capture local changes in the mean state via a linear term in our regression ("Methods"). We further correct for uneven sampling rates in the NGRIP data to minimize biases due to variable Δt_i, especially in deeper sections of the ice core ("Methods").

We find inconclusive evidence for statistical precursor signals preceding DO events based on our bias-corrected λ_c (Fig. 3). This agrees with previous assessments, which quantified λ in terms of increasing variance and autocorrelation[37,43], and found statistically inconclusive evidence for early warning of DO events. We further note that discrete-time λ_d estimates which are not corrected for strongly variable Δt_i (Fig. 3b) are biased in the deep sections of the NGRIP data, and provide misleading signals of CSD prior to DO events due to the impacts of uneven sampling intervals ("Methods"). We note, however, that recent (past 2,000 years) λ estimates are the lowest they have been in at least 50,000 years (Supplementary Fig. S4), indicating a less stable Northern Hemisphere climate than in the past interglacial interval.

Glacier surge prediction

Glacier surging is a well-documented phenomenon in which a glacier periodically shifts from a slow quiescent state to a phase of rapid

https://doi.org/10.1038/s41467-026-78489-7

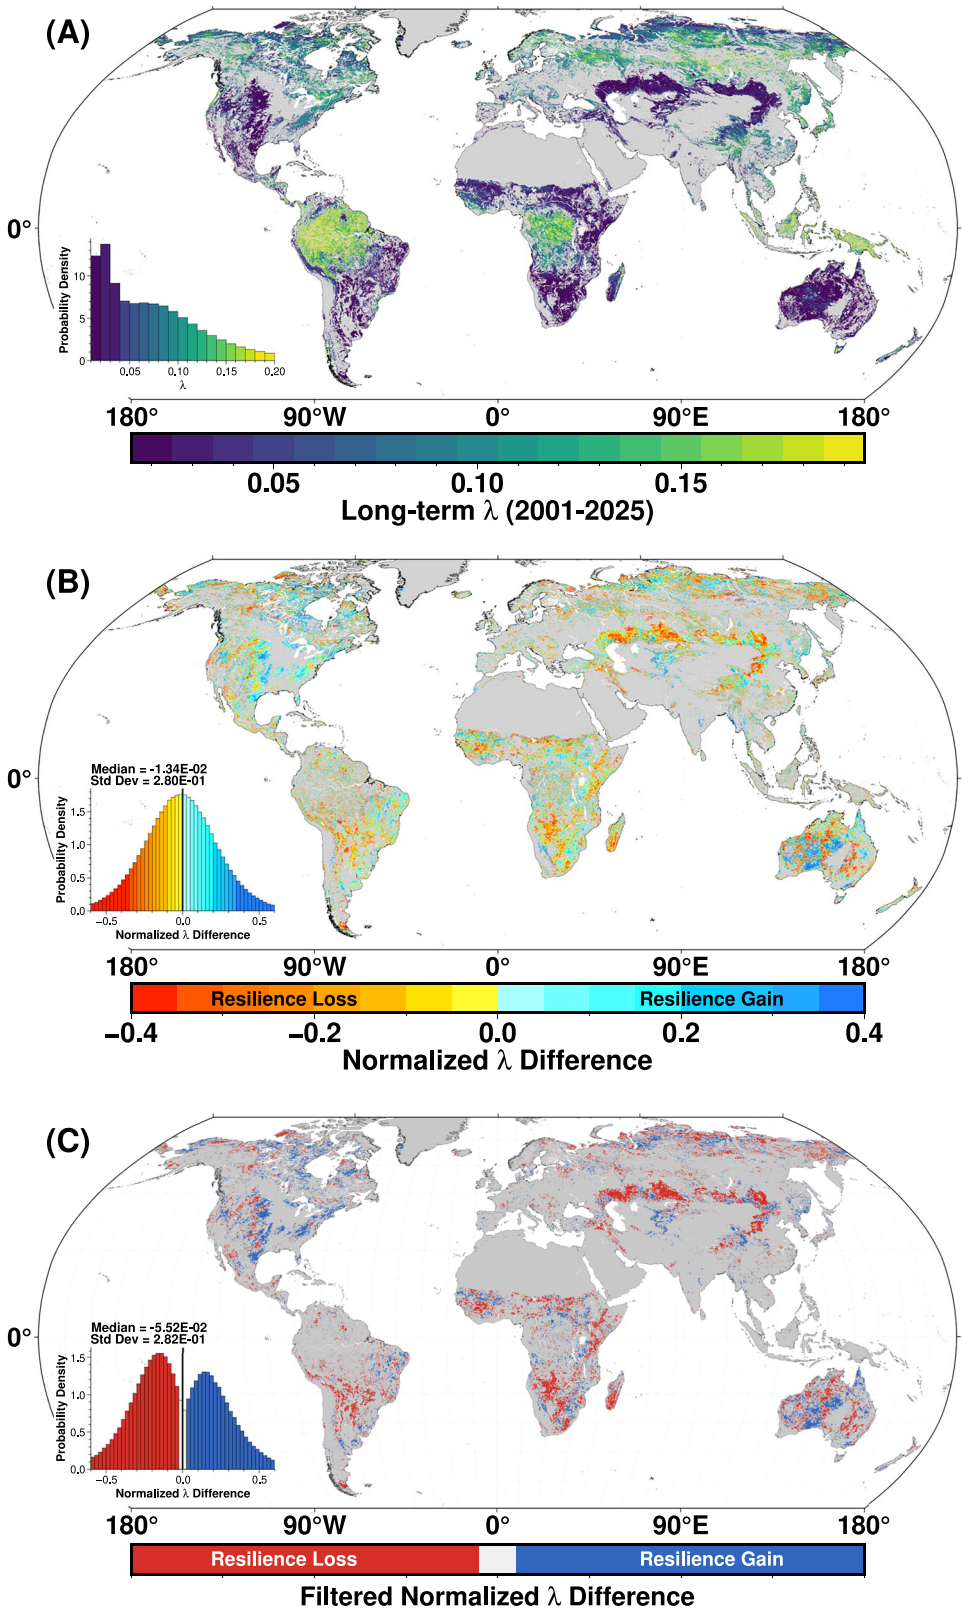


Fig. 2 | Global vegetation resilience. A Estimated λ over vegetated ecosystems (MODIS kNDVI, 2001-2025), with anthropogenic land-cover types masked based on MODIS land-cover data ("Methods"). **B** Normalized difference in λ between the period 2013–2025 and 2001–2012, showing a slight tendency towards decreasing λ globally. **C** λ differences filtered by spatial consistency ("Methods"), showing fewer coherent blocks of positive and negative λ changes globally, and a stronger tendency towards resilience loss. Insets show histograms of global values, with median and standard deviation marked for different maps.

https://doi.org/10.1038/s41467-026-78489-7

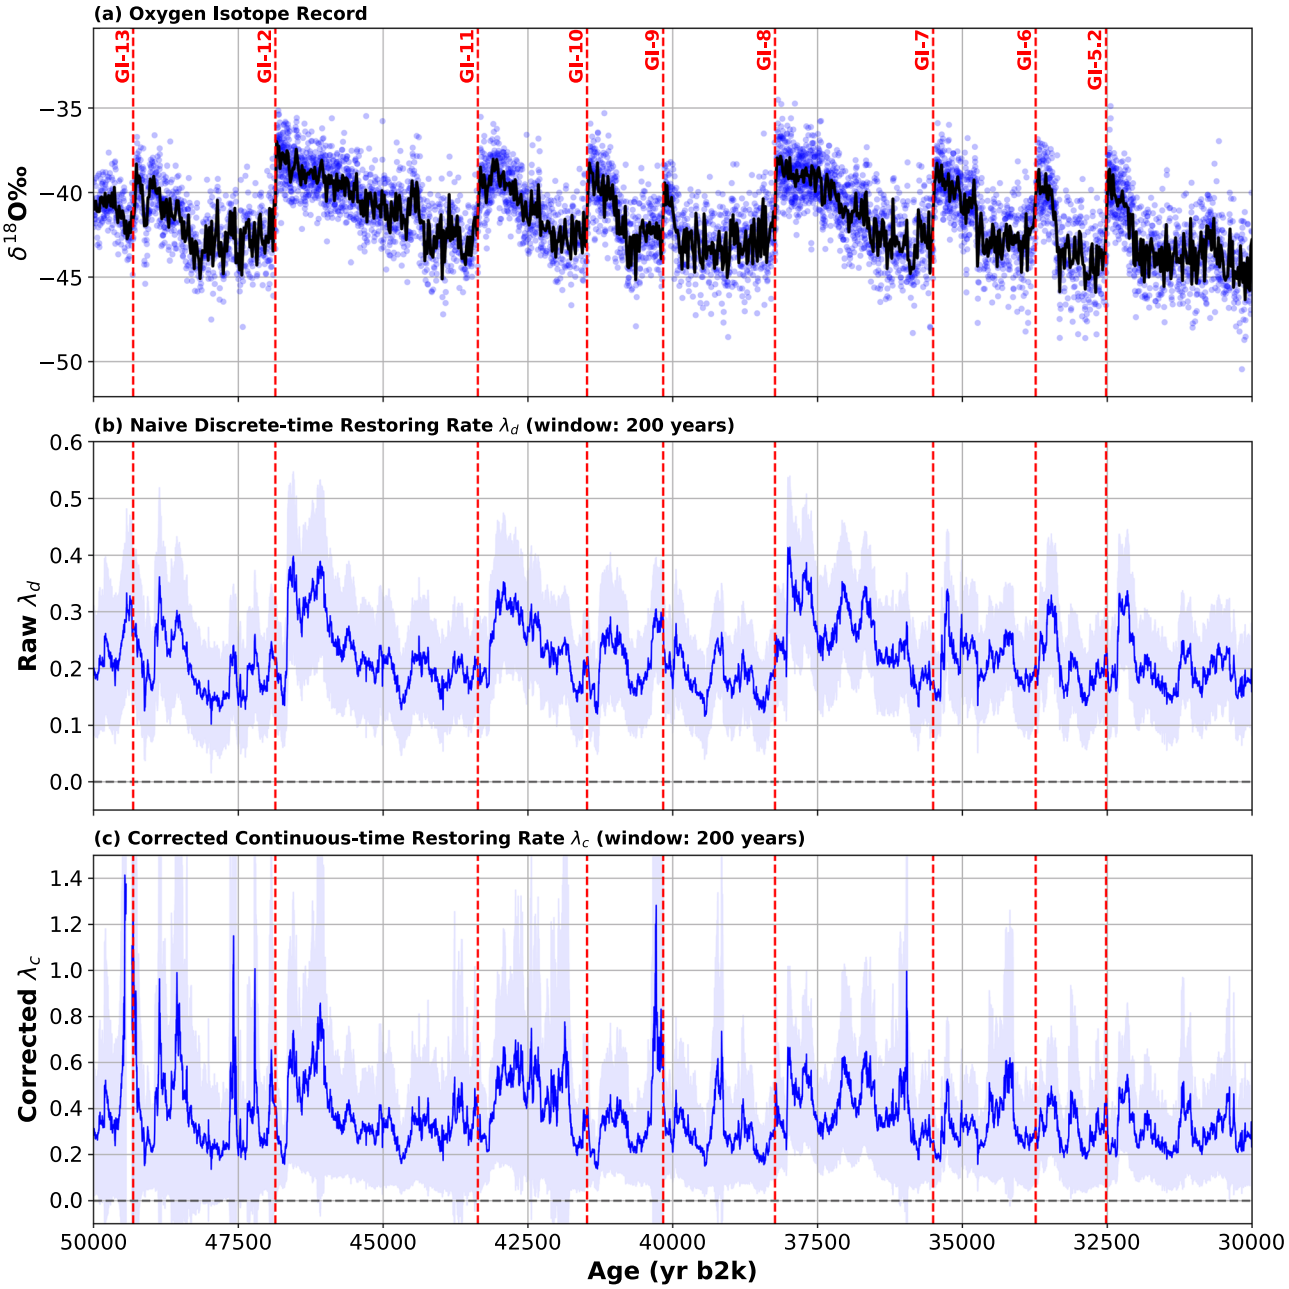


Fig. 3 | NGRIP Ice core data. a δ^{18}O isotope records from NGRIP[40]. Vertical lines mark DO event timing after[39]. Data subset to 20,000 years for clarity. The full period of the high-resolution NGRIP data can be found in Supplementary Fig. S4. **b** Changes in discrete-time λ_d calculated on the raw time series without pre-processing, incorporating measurement and age-model uncertainties via weighted least squares ("Methods"). **c** Corrected λ_c accounting for changes in sampling intervals through time ("Methods"), showing less pronounced decreases in λ before many DO events. Shaded bounds (**b**, **c**) cover one standard deviation uncertainty in λ ("Methods"). Bias-corrected λ_c estimates do not show conclusive evidence for consistent decreases in stability before DO events.

acceleration, moving large volumes of ice down-glacier[44,45]. Not all glaciers surge; whether or not there are periodic glacier surges is controlled by the interaction between downslope mass transport and basal friction forces[44,46–49]. In this sense, a glacier can be considered a dynamical system with both steady-state downhill movement and transient surges.

Recent advances in image cross-correlation have yielded globally resolved glacier velocity maps[50,51]. These velocity maps are based on spatially and temporally overlapping pixel-offsets from multiple optical and radar sensors with variable data quality, revisit times, and spatial resolutions. Glacier velocity data is hence (1) seasonal, (2) unevenly sampled, and (3) has temporally-variable noise levels. For typical (e.g., AC1-based) approaches, substantial pre-processing is necessary; removing seasonality from highly variable glacier accumulation/ablation dynamics is also difficult[13]. Properly integrating the timescale of glacier velocity measurements is also not straightforward; velocity is measured over varying time spans, meaning that many measurements overlap and have different degrees of temporal smoothing. Estimating variance or autocorrelation changes is hence difficult, as there is no fixed sampling window over which to perform the analysis.

Our regression-based approach, however, can incorporate the velocity (v) measurement span directly into the estimated $\frac{\Delta v_i}{\Delta t_i}$, and hence handle overlapping span-averaged measurements with the same

https://doi.org/10.1038/s41467-026-78489-7

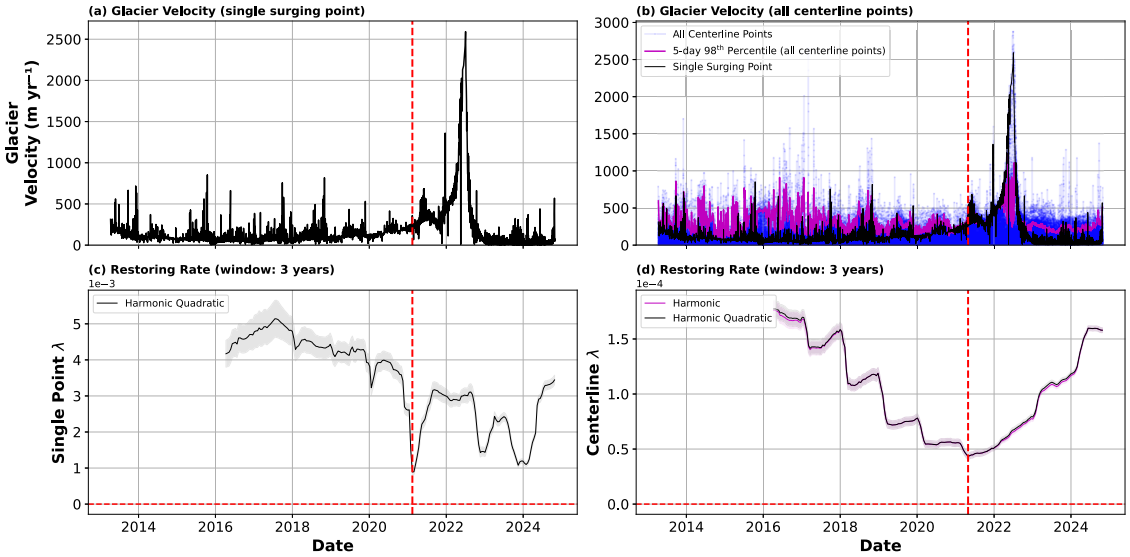


Fig. 4 | Stability of a surging glacier. a Glacier velocity sampled at a single surging point and **b** over all glacier centerline points in 500 m steps ("Methods"). **c**, **d** Restoring rate λ, showing a rapid decline in stability before the onset of a major surge event. λ is estimated using parametric regression with harmonic and quadratic terms (black), as well as only harmonic terms (purple) for comparison in the two-dimensional case. Local minima are marked with vertical dashed lines. Shaded bounds (**c**, **d**) cover one standard deviation uncertainty in λ ("Methods"). Decreases in stability before the glacier surge can be seen in both the single hand-picked location and in the set of all glacier centerline points.

rigor as instantaneous measurements (Methods). We note that the 'naive' approach of treating velocity measurements as occurring instantaneously at their span midpoints is also tractable within our framework; the midpoint approach, however, inflates the derivative $\frac{\Delta v_i}{\Delta t_i}$ for small midpoint gaps and hence amplifies high-frequency changes in the time series. To limit biases from derivative inflation, we rely on a span-aware estimate of λ ("Methods").

It is important to present a crucial detail about our glacier analysis: glacier surges inherently involve non-linear (often exponential) background acceleration. In a strict CSD sense, non-linear acceleration invalidates the basic assumption of an Ornstein-Uhlenbeck process restoring towards a stationary mean. This introduces ambiguity into the interpretation of λ, depending on how that non-linear trend is incorporated into the analysis. If we intentionally omit the non-linear term(s) from our regression design (e.g., only using harmonic terms), the linear model does not track background non-linear acceleration. λ dropping towards zero hence detects that drift away from the baseline mean velocity; that signal could be interpreted as a precursor of system change, but could no longer be interpreted as CSD. Treating the non-linear part of glacier acceleration as strictly exponential would motivate a linear analysis on log-transformed glacier velocity data, though this approach enhances the influence of large outliers. Adding higher-order terms to our parametric regression model (e.g., quadratic) is less likely to bias the estimation of λ, and allows us to again interpret changes in λ as CSD, as the higher-order polynomial absorbs local non-linear trends (Fig. 4).

There is a clear signal of stability loss (λ approaching zero from above) before the onset of the main glacier surge phase (Fig. 4); this signal is visible in both the one- and two-dimensional cases, and remains clear when the size of the moving analysis window is varied (Supplementary Fig. S5). We attribute λ decreases to modifications to bed friction as water builds up and glacier flow moves from a 'sticky' regime to a 'sliding' regime[47–49]. While a thorough examination of the precise surge initiation mechanism is beyond the scope of this study, we posit that the sharp decreases in λ before surge onset can be used to predict oncoming glacier surge activity, and λ is linked to distinct glacier physical movement regimes. We note that in the two-dimensional case, we find qualitatively the same early-warning signal when we omit the quadratic term from our parametric regression model (Fig. 4d). This behavior stems from out-of-phase non-linear acceleration throughout the glacier, which minimizes the need for a quadratic term, as opposed to in the one-dimensional case, where substantial biases to λ would be expected if the quadratic term was omitted. We further note that the higher data density and out-of-phase measurements in the two-dimensional case allow for the use of shorter temporal windows (Supplementary Fig. S5), potentially improving the operational use case for using λ changes as an early-warning signal of oncoming glacier surges.

Recent work[13] has used dynamical systems theory to predict glacier surge onset; their eigenvalue-based method, however, requires complex state-space embedding parameterization, as well as resampling glacier velocity data onto a consistent temporal grid. Our regression-based approach hence provides additional benefits, particularly in that it can incorporate velocity uncertainty information directly into the estimation of λ and that it can also integrate spatial glacier velocity data via sampling the entire glacier centerline instead of a single surging point (Fig. 4d). Treating the entire glacier as a single system, however, makes broader assumptions about the heterogeneity of external forcing (e.g., snowfall, temperature) than are required for a single time series. While the temporal pattern of inferred whole-glacier λ is suggestive of the same pre-surge dynamics, further validation against in-situ or other surge-mapping data would be needed to confirm our findings.

Discussion

The estimation of λ via parametric regression provides a few key benefits over the most common approaches to estimating λ via AC1 or variance, especially with real-world and noisy data. We can model arbitrary attractors directly in our λ estimate via regression factors, which makes it possible to concisely account for non-stationarity (e.g., seasonality, Figs. 1, 2, 4). We hence do not need to pre-process our data to remove seasonality, limiting the introduction of errors and spurious signals from the choice of deseasoning procedure. As we model both the attractor μ and λ via a single regression, we can compute one uncertainty estimate that accounts for uncertainty in both the attractor μ (e.g., due to complex seasonal cycles) and λ. This means that we recover a direct estimate of the total uncertainty in λ, which has so far been difficult with previous methods. Furthermore, by framing our

analysis around $\frac{\Delta x_i}{\Delta t_i}$, we natively account for gaps and hence do not require complete or evenly-sampled time series. This is a critical benefit for some systems, such as the vegetation (Fig. 2), paleoclimate (Fig. 3), and glacier (Fig. 4) data shown here.

Finally, by posing our estimation of λ as a regression problem, we inherit a very wide body of methodologies for dealing with noisy and complex data. For example, our method can directly integrate time-varying weights into our estimates of λ by relying on weighted least squares (Figs. 2–4), minimize the influence of outliers on our λ estimates by robust regression approaches (Methods), and handle difficult noise regimes (e.g., red noise) through appropriate regression frameworks (Methods, Supplementary Fig. S6). We can also recover error bounds via the regression residuals – as well as quantify the unexplained variance of λ – which can serve as a proxy for the external noise forcing on the system. Finally, it is possible to expand our method to the spatial domain – with caveats – by collapsing the spatial dimension and performing the same regression proposed in Equation (3) over a much denser view of a given system. It would also, in principle, be possible to extend our method to include multi-variate systems, though the temporal resolution requirements for, and interpretation of, recovered multi-variate λ estimates would be more difficult.

Over small spatial areas with homogeneous forcing (e.g., the same weather, similar vegetation mixes), using a small spatial field instead of a single time series can provide a more data-rich view of the system, and thus a more robust regression-based estimate of λ. It is not, however, well-suited to large and complex systems, particularly those where there are substantial forcing gradients (e.g., spatially varying rainfall patterns), since it is not easy to disentangle changes in the system response to forcing (λ) from spatial variability in the forcing itself. Further research is needed to develop spatially-extended regression-based λ estimates and rigorously compare them to established spatial early-warning signals[11,13].

One key limitation of our approach is the reliance on the specification of the functional form of the attractor μ. If we, for example, analyze a system with semi-annual harmonics while only including annual-scale Fourier terms in our parametric regression model, the missed harmonics will bias our estimate of λ; this can be shown concisely with a synthetic model (Supplementary Fig. S7). There is relatively less danger of biasing λ when including too many terms in the regression, as a robust regression can assign these terms very low weights (e.g., a linear term in the regression design matrix is not needed to explain the movement of a flat line); there remains, however, danger of overfitting if too many terms are included, especially higher order (e.g., quadratic) terms. If the physical process driving the system is to some degree known (e.g., vegetation follows an annual seasonal cycle with n harmonics), λ estimates via parametric regression have a strong physical basis. If the shape of the attractor μ is unknown, the danger of model misspecification and biased λ estimates is higher.

One way around this limitation is to use non-parametric regression[52–55]. Non-parametric (e.g., kernel regression) approaches avoid predefining the functional form of μ, but can introduce other issues that are not present in the parametric regression approach ("Methods", Supplementary Fig. S8). The most critical drawback relates to the separation of the attractor and the stochastic fluctuations: too small of a kernel bandwidth can result in stochastic fluctuations being treated as part of the attractor μ, and too large of a bandwidth can lead to the attractor's actual dynamics being treated as stochastic noise. Standard automated bandwidth selection methods often fail in the presence of autocorrelated noise[56,57], and a fixed kernel bandwidth assumes that sampling densities are constant throughout the time series, which is not always the case for real-world data.

In our parametric approach, λ is defined around physical constraints (e.g., a known annual periodicity), meaning that we enforce a basic signal partitioning into, for example, dynamical stability, seasonal, and residual noise components. We can therefore interpret changes in λ as independent of those specified physical constraints; without them, it is more difficult to link changes in λ to a physical process. Unfortunately, finding the correct kernel bandwidth that does not result in under- or over-fitting is not straightforward, and indeed it may change throughout a time series. For systems where the attractor is unknown or difficult to define, however, non-parametric approaches are useful, and can also be used to estimate λ within the same theoretical Langevin framework we present here ("Methods", Supplementary Fig. S8).

Our regression-based λ estimate provides a robust means of estimating the stability of diverse systems, without key limitations that require substantial pre-processing of many data sets. We further provide a straightforward means by which to capture and account for uncertainty in real-world measurements. Our method can function as a drop-in replacement for the wide body of literature relying on autocorrelation-based stability estimates, and provides a means of simplifying and minimizing data pre-processing to increase the diversity of systems that can be assessed within a critical slowing down framework.

Methods

Robust λ estimation via parametric regression

To minimize biases induced by outliers in our regression-based λ estimate (Equation (3)) – particularly in very dense and noisy time series (e.g., spatial data, paleoclimate records) – we rely on a robust regression optimized using Iteratively Reweighted Least Squares. We downweight outliers in our data using Huber's T, which linearly suppresses the influence of large residuals after an initial least squares fit. In cases where we have access to time-explicit uncertainties (i.e., Figs. 3, 4), we incorporate them by scaling both the regression factors and the dependent variable by the square root of the weights. We also tested Ordinary Least Squares and Generalized Least Squares, which account for autocorrelation (GLSAR)[6], finding that they produced similar results for our synthetic data under white noise forcing (Supplementary Fig. S9). We note that the GLSAR regression performs substantially better for data with red noise forcing (Supplementary Fig. S6); this alternative regression approach can be seamlessly used with our method. As the regression estimate of λ is auto-regressive, there will be a small sample bias of order $O(1/n)$ [58]; for the relatively large n we use here, we do not explicitly incorporate this bias. We note that our method is comparable in runtime to estimating λ from AC1, given the differences in data pre-processing involved. For the same synthetic time series (Fig. 1, daily data, 1950–2025), our regression takes 11 s to run for five-year moving windows, while first deseasoning the same time series and then estimating AC1 on the moving windows takes 10 s in total.

We capture the uncertainty in our estimated λ using the standard error (SE_λ) of the regression slope. We construct a 95% confidence interval around λ ($\lambda \pm 1.96 \cdot SE_\lambda$); if at any point that confidence interval crosses zero, we can no longer exclude the possibility that the system has become unstable. The standard error we report (confidence bounds on Figs. 1, 3, and 4 λ estimates) is expanded to account for serial correlation in the time series data (reduced degrees of freedom) using the temporal autocorrelation of the regression residuals. Code to reproduce our λ estimation procedure can be found on Zenodo[59].

Non-parametric regression

When it is difficult to define the shape of the attractor μ, parametric regression risks biasing λ estimates. One way around this is to rely on non-parametric regression[52–54], which does not fix the functional form of the attractor. Over quasi-stationary moving windows, we can use the non-parametric local attractor state $x_{eq}(t_i)$ to construct local state anomalies

$$x_{\mathrm{anom},i} = x_i - x_{\mathrm{eq}}(t_i) \tag{5}$$

and the attractor's local rate of change $\left(\frac{\Delta x_i}{\Delta t_i}\right)_{eq}(t_i)$ to construct rate anomalies

$$\left(\frac{\Delta x_i}{\Delta t_i}\right)_{anom,i} = \frac{\Delta x_i}{\Delta t_i} - \left(\frac{\Delta x_i}{\Delta t_i}\right)_{eq}(t_i) \quad (6)$$

We can then use those anomalies in a zero-intercept linear regression:

$$\left(\frac{\Delta x_i}{\Delta t_i}\right)_{anom,i} = \beta_1 x_{anom,i} + \epsilon_i \quad (7)$$

where the negative of the slope β_1 again captures λ. We can resolve this regression with the same Ordinary Least Squares, Iteratively Reweighted Least Squares, or GLSAR approaches as we use in the parametric case.

To construct the anomalies for the non-parametric regression, we use two schemes: Nadaraya-Watson Kernel Regression[53,54] and Local Linear Regression[52]. The Nadaraya-Watson approach treats $x_{eq}(t_i)$ and $\left(\frac{\Delta x_i}{\Delta t_i}\right)_{eq}(t_i)$ as locally weighted averages[53,54], whereas the local linear regression approach uses a first-order Taylor expansion to construct the local attractor[52]. For both approaches, we use a modified Gaussian kernel $K_h(t_i, t_j)$ with an h-block gap exclusion:

$$K_h(t_i, t_j) = \begin{cases} \exp\left(-\frac{(t_i - t_j)^2}{2h^2}\right) & \text{if } |t_i - t_j| > \tau \\ 0 & \text{if } |t_i - t_j| \le \tau \end{cases} \quad (8)$$

with a bandwidth h and a local exclusion zone τ (e.g., 15 days) to minimize over-tracking of high-frequency system memory by the local attractor. For the same synthetic system used in Fig. 1, we capture a reduction in λ before the critical transition using both non-parametric methods (Supplementary Fig. S8); the non-parametric λ estimates, however, are biased low, have large uncertainty bounds, and contain substantial high-frequency variability. This is largely driven by the difficulty inherent in separating the strong seasonal forcing and weaker stochastic noise with the non-parametric kernel, which leads to substantial signal leakage and biases in the estimate of λ. For the systems we study here with known forcing (e.g., seasonal sinusoids), a parametric approach yields a less biased estimate of λ.

Further, we note that the non-parametric approach is essentially a sequential, two-step estimation process (as opposed to our simultaneous, one-step parametric regression) which treats the derived anomalies $x_{anom,i}$ and $\left(\frac{\Delta x_i}{\Delta t_i}\right)_{anom,i}$ as error-free observations, and does not propagate estimation variances forward into the confidence intervals of λ. Properly accounting for that uncertainty is analytically challenging – especially when the kernel bandwidth is selected automatically – and estimating confidence intervals via bootstraps is computationally prohibitive over many moving windows. More complex one-step procedures (e.g., using varying-coefficient models) could, in principle, mitigate this bias, but would be very difficult to correctly parameterize for real-world data. Hence, we present a simplified non-parametric framework here for comparison with our parametric approach, but note that the confidence intervals of λ we present are likely too narrow (Supplementary Fig. S8).

Error propagation and uncertainty

The most common approach to estimating λ on non-stationary data is to first pre-process the data to remove time dependence (e.g., seasonality) and then estimate λ on the cleaned, nominally stationary, residuals (e.g., via AC1). Our approach of using a single parametric regression model to capture both attractor movement and λ yields two key benefits with regard to uncertainty quantification: covariance partitioning and straightforward error propagation.

An underlying and underappreciated assumption of pre-processing strategies is that the two signals of interest (e.g., seasonality and system dynamics) are orthogonal; that is, they can be perfectly decomposed. This could be the case for synthetic systems with, e.g., additive seasonality, but is unlikely to occur in real-world systems. It is often the case that the initial drift towards instability is attributed to the seasonal attractor, and the recovered deseasoned residuals are flatter than they should be. By instead performing a single robust regression, the regression coefficients for λ and μ (e.g., Equation (4)) are both tuned to minimize residual variance simultaneously. This means that if, e.g., glacier acceleration is not strictly periodic, it is correctly partitioned into the λ term rather than as part of normal seasonal glacier acceleration. We note that our parametric regression approach would be strictly equivalent to performing two linear regressions – one to remove the attractor and one to estimate λ – under the condition that Ordinary Least Squares is used for both steps (i.e., the Frisch-Waugh-Lovell theorem[60]). This equivalence is violated by most commonly used pre-processing steps, such as removing climate anomalies (e.g., long-term monthly means) or using more complex deseasoning procedures (e.g., STL).

Estimates of λ are generally taken at face value – there are no uncertainty bounds when λ is computed from AC1 or variance. Furthermore, errors due to the deseasoning and detrending process (e.g., partitioning too much signal into the seasonal term) are not propagated forward into estimates of λ_{AC1} or λ_{Var}. In our approach, we derive the standard error of λ directly from the covariance matrix of the regression, which includes the covariance between all of the components of the regression factors (e.g., seasonality) and our estimate of λ. Hence, if μ is poorly constrained in our regression (e.g., due to large gaps), we propagate that uncertainty into our derived standard error of λ; we thus account for errors in both our modeled attractor μ and λ itself. In practice, this means that poor data quality (gaps, large measurement errors) will inflate the standard error, yielding wider uncertainty bounds on λ.

Continuous-time estimation of λ

A common basis for estimating λ is to treat a system as an Ornstein-Uhlenbeck process:

$$\frac{dx}{dt} = -\lambda_c(x(t) - \mu) + \sigma_c \eta(t) \quad (9)$$

where μ is a stable equilibrium state, σ_c is continuous noise strength, and $\eta(t)$ is Gaussian white noise. In this formulation, λ_c is the true continuous-time restoring rate against perturbations to the system. Any data-driven estimate of λ_c will – by necessity – be discretized, as we recover λ_c from a time series sampled against some continuous process; we term this discrete-time estimate λ_d. In order to convert λ_d to λ_c consistently for multiple means of estimating λ (via regression, AC1, or variance), we use an analytical conversion incorporating the data sampling interval. We define the window-averaged time sampling rate as $\langle \Delta t \rangle$, and use it to convert λ_d, $\lambda_{d,AC1}$, and $\lambda_{d,Var}$ to their continuous-time equivalents.

For our regression-based λ, equating the discrete-time autocorrelation to the exact analytical solution for a single increment of Equation (2) with time step Δt yields:

$$\lambda_d = \frac{1 - e^{-\lambda_c \Delta t}}{\Delta t} \quad (10)$$

If we expand this via a Taylor series, we find a clear inherent sampling bias in the discrete-time λ_d versus the continuous-time λ_c:

$$\lambda_d = \frac{1 - \left(1 - \lambda_c \langle \Delta t \rangle + \frac{1}{2}\lambda_c^2 \langle \Delta t \rangle^2 - \mathcal{O}(\langle \Delta t \rangle^3)\right)}{\langle \Delta t \rangle} \quad (11)$$

or, simplified:

$$\lambda_d \approx \lambda_c - \frac{1}{2}\lambda_c^2\langle\Delta t\rangle \tag{12}$$

The directly-estimated λ_d can be strongly biased (underestimates resilience) for large $\langle\Delta t\rangle$ – as commonly occurs in paleoclimate data (Fig. 3) – and higher-order terms ($\mathcal{O}$) cannot be safely ignored. We thus rather invert the full analytical relationship between λ_d and λ_c as:

$$\lambda_c = \frac{-\ln(1-\lambda_d\langle\Delta t\rangle)}{\langle\Delta t\rangle} \tag{13}$$

yielding an estimate of λ_c that is less sensitive to changing sampling rates through time than the directly-estimated λ_d.

For a discretely-sampled time series, the AC1 decays exponentially as:

$$\mathrm{AC1} = e^{-\lambda_c\Delta t} \tag{14}$$

In order to convert λ_{AC1} to continuous-time, we incorporate the mean sampling rate $\langle\Delta t\rangle$ as:

$$\lambda_{c,\,\mathrm{AC1}} = \frac{-\ln(\mathrm{AC1})}{\langle\Delta t\rangle} \tag{15}$$

We note that λ_{AC1} is only strictly valid for evenly-sampled data; by correcting with $\langle\Delta t\rangle$, we can mitigate missing data biases and transform the units of $\lambda_{c,\mathrm{AC1}}$ to match those of λ_c, which are in physical-time units (e.g., restoring rate per year).

The conversion for $\lambda_{c,\mathrm{Var}}$ is slightly more complex in that it incorporates both variance in the state ($\mathrm{Var}(x)$) and the variance of the driving noise over each discrete time step. For a continuous-time Ornstein-Uhlenbeck process, the variance is related to λ as:

$$\mathrm{Var}(x) = \frac{\sigma_c^2}{2\lambda_c} \tag{16}$$

For a discrete time sampling $\langle\Delta t\rangle$, the system behaves as an AR(1) process and the discrete noise variance (σ_d^2) – obtained as residuals from an AR(1) regression (used to calculate AC1) – is an approximation of the continuous noise scaled by the time step ($\sigma_d^2 \approx \sigma_c^2\langle\Delta t\rangle$). We can use that conversion, as well as the exponential decay of the AC1 (Equation (14)), to reformulate the discrete-time relationship between variance and λ in terms of continuous variables:

$$\mathrm{Var}(x) = \frac{\sigma_d^2}{1-\mathrm{AC1}^2} = \frac{\sigma_c^2\langle\Delta t\rangle}{1-e^{-2\lambda_c\langle\Delta t\rangle}} \tag{17}$$

We can then solve for $\lambda_{c,\mathrm{Var}}$ as:

$$\lambda_{c,\,\mathrm{Var}} = \frac{-\ln\left(1-\frac{\sigma_d^2}{\mathrm{Var}(x)}\right)}{2\langle\Delta t\rangle} \tag{18}$$

We emphasize that, as with λ_{AC1}, λ_{Var} is only strictly valid for evenly-sampled data, despite the $\langle\Delta t\rangle$ correction. λ_{AC1} relies on paired neighboring observations and does not integrate the time between the observations; λ_{Var} can be biased by both changes in sample variance through time and the integration time for random noise between measurements. Unfortunately, integrating variable Δt_i into each pointwise AC1 and variance estimate is practically intractable, and as the theoretical physical restoring rate λ_c is exponential, our mean-time correction $\langle\Delta t\rangle$ is not a perfect solution. Any variability in the sampling frequency will introduce a bias into both estimators; missing data will also bias λ_{AC1} and λ_{Var} away from each other[26]. The conversion from $\lambda_{d,\mathrm{AC1}}$ and $\lambda_{d,\mathrm{Var}}$ to $\lambda_{c,\mathrm{AC1}}$ and $\lambda_{c,\mathrm{Var}}$ using $\langle\Delta t\rangle$ can account for shifts in the mean sampling rate, but cannot correct for variance in the sampling rate, which biases the $\lambda_d \mapsto \lambda_c$ conversions due to Jensen's inequality.

For varying time windows Δt_i, the discrete-time least squares regression estimates λ_d are defined as $\frac{1}{n}\sum_{i=1}^{n}\frac{1-e^{-\lambda_c\Delta t_i}}{\Delta t_i}$. As the function $s \mapsto \frac{1-e^{-\lambda_c s}}{s}$ is strictly convex for $s > 0$ and $\lambda_c > 0$, the discrete-time λ_d estimate will always be greater than or equal to $\frac{1-e^{-\lambda_c\langle\Delta t\rangle}}{\langle\Delta t\rangle}$. Our continuous-time conversion for the regression-based λ estimates hence yields an upper bound; the true λ_c will be less than or equal to our estimated λ_c, hence establishing a conservative estimate (i.e., biased away from zero) of system stability. Similarly, estimates of λ relying on autocorrelation (λ_{AC1} and λ_{Var}) will also be bounded due to Jensen's inequality. As the exponential function used to estimate $\lambda_{d,\mathrm{AC1}}$ is convex ($\mathrm{AC1} = \frac{1}{n}\sum_{i=1}^{n} e^{-\lambda_c\Delta t_i}$), our empirical AC1 estimates will be greater than or equal to $e^{-\lambda_c\langle\Delta t\rangle}$; inflation in AC1 implies a slower restoring rate and hence the true λ_c will be greater than or equal to $\lambda_{c,\mathrm{AC1}}$ and $\lambda_{c,\mathrm{Var}}$ (i.e., less conservative, as λ is biased towards zero).

In both cases, the accuracy of the λ_c estimate will be directly related to the variance of the sampling intervals; the closer the individual Δt_i values are to being strictly equal, the closer the estimators approach the true λ_c. We hence use $\langle\Delta t\rangle$ as an analytical correction for each λ_d estimate to bring them into a common physical reference frame and as close to the true continuous-time λ_c as is possible with discretely-sampled data, despite the fact that λ_{AC1} and λ_{Var} are not strictly valid for unevenly sampled data, and $\langle\Delta t\rangle$ is not a perfect solution. We note that a regression-based λ estimate maintains the core advantage over λ_{AC1} and λ_{Var} that it integrates Δt_i directly into each point-by-point estimate prior to regression for λ, which minimizes the influence of irregular time sampling and removes the core requirement that data be evenly sampled in time to estimate λ. As resampling and temporal interpolation techniques can bias AC1 and variance, a regression-based λ formulation is preferred for unevenly sampled data.

Deseasoning and detrending

For our synthetic time series (Fig. 1) and vegetation (Supplementary Fig. S2) analyses, we compare our regression-based λ to that recovered from standard data pre-processing. In short, we first deseason and detrend each time series using (1) Seasonal Trend Decomposition via Loess (STL)[23], (2) by removing both the long-term monthly mean and a five-year rolling average (here termed climatological deseasoning), and (3) using a harmonic fit for seasonality followed by a five-year rolling average[22]. All three approaches yield nominally deseasoned and detrended residual time series that can be used to compute λ via autocorrelation (λ_{AC1}) or variance (λ_{Var})[12]. For STL, we use a seasonal period of 365 days and a smoothing window of seven, as has been used in several previous publications[22].

Synthetic time series data

We generate a one-dimensional time series describing a pitchfork-bifurcation model with seasonal forcing[13] to test our regression-based λ estimate and compare it to common data analysis workflows:

$$\frac{\mathrm{d}x}{\mathrm{d}t} = p(t)\,x - x^3 + A\cos(\omega t) + \sigma\,\xi(t) \tag{19}$$

We vary the control parameter p linearly to produce the transition from a high to a low state (Fig. 1). Seasonality is controlled by the amplitude A and periodicity $\cos(\omega t)$, with additive Gaussian white noise $\sigma\xi(t)$. We use a model run with seasonal amplitude $A = 0$ to assess our parametric regression approach to handling seasonality, as well as to provide a baseline against which to compare deseasoned and detrended seasonal data. A complete description of the model parameters used here can be found in the Supplement. We note that we

find qualitatively the same results when applying our method to a simpler saddle-node bifurcation (Supplementary Fig. S10).

White noise forcing is a simplifying assumption used in many λ estimation techniques, but does not generally hold for real-world data. We test our parametric regression approach by generating versions of our synthetic model forced by (1) white, (2) red, (3) heavy-tailed (Student-t), (4) skewed (Skew-normal), and (5) multiplicative (state-dependent) noise (Supplementary Fig. S11). We find that for all noise forcings except red, our robust regression approach successfully captures changes in λ, and that those changes are qualitatively similar. Red noise forcing is more difficult for the robust regression framework to handle; we can, however, use GLSAR for estimation in our parametric regression model[6] and capture changes in λ that are comparable to those found in the same system under white noise forcing (Supplementary Fig. S6).

To test the influence of moving window size on our results, we re-analyze the data from Fig. 1 using moving windows ranging from one year to eleven years (Supplementary Fig. S12). We find qualitatively the same results, although very short (one-year) moving windows struggle to correctly partition out the seasonality in the parametric regression. Longer moving windows result in a more robust estimate of λ, with the drawback being that they could smooth over rapid changes in λ by using a long temporal baseline.

While many real-world systems satisfy our condition that process noise is larger than measurement noise ($\sigma\eta \gg \delta$), this will not always be the case. If we vary the ratio of measurement to process noise, we can test how λ responds (Supplementary Fig. S13). We find that the functional form of λ is maintained across varying noise ratios, but that the absolute value of λ changes substantially; increasing measurement noise biases λ away from zero. Even for $\sigma\eta \ll \delta$, however, we still find a negative slope in λ towards the critical transition, indicating that it is still useful as an early warning signal. It is worth noting that it is – in some cases – possible to separate process and measurement noise from the underlying dynamics[61], which could be one way to minimize biases in λ due to overprinted measurement noise.

We test the robustness of our approach to data gaps by removing a given percentage of points randomly sampled throughout the time series (Fig. 1). We further mimic, e.g., seasonal cloud or snow cover gaps by removing sequential months of data (Supplementary Fig. S1); we additionally remove a random selection of points in the time series with long gaps as well. We find that in both cases, our regression-based λ estimates are robust to gap percentages and lengths, as compared to a gap-free control run. To test whether this also holds for regular and substantial gaps (i.e., using 10-day data instead of daily data), we can add a skip to the synthetic time series (Supplementary Fig. S14). We find that even for long and consistent gaps, we still capture changes in λ concisely with our parametric regression approach. Code to reproduce our time series degradation functions and the resulting Supplementary Figures can be found on Zenodo[59].

Spatial vegetation model

To explore whether our method can be adapted to the spatial case, we use the same model presented in Smith et al.[13], which adapts a reaction-diffusion model[11] to include seasonality. The model is parameterized for vegetation V, growth rate r_v, environmental stress E, diffusion coefficient D and noise amplitude σ as:

$$\frac{\partial V}{\partial t} = r_v(x,y,t)\,V\left[1 - V\frac{(h_E^p + E^p)}{h_E^p}\right] + D\nabla^2 V + \sigma\,\eta(x,y,t) \tag{20}$$

We use spatially and temporally uncorrelated Gaussian noise for $\eta(x, y, t)$. The rate and shape of the induced critical transition are controlled by h_E, h_v, and p which modify the environmental stress E:

$$E = E_0(t)\frac{h_v}{h_v + V} \tag{21}$$

We add a seasonal component to E which is homogeneous across all grid cells with amplitude A_{E_0}:

$$E_0(t) = E_{0,\,\mathrm{base}}(t) + A_{E_0}\sin\left(\frac{2\pi t}{T_{\mathrm{year}}}\right) \tag{22}$$

The growth rate r_v is given a seasonal component which varies in space to follow, e.g., microclimates or topographic and rainfall gradients. We thus add variability to r_v pixelwise as:

$$r_v(x,y,t) = r_{v,0} + A_{r_v}(x,y)\,\sin\left(\frac{2\pi t}{T_{\mathrm{year}}} + \phi(x,y)\right) \tag{23}$$

Working with spatio-temporal data. As our method of estimating λ is regression-based, it can be simply reconfigured to analyze spatial data by collapsing the spatial dimension; in short, we can perform our regression to discover λ using the set of all data points in a given temporal window, regardless of their spatial position. The use of spatial data can dramatically increase the data density available (e.g., 50 x 50 points per time slice instead of a single point in the one-dimensional case). This increases the statistical power of the estimated λ, and can motivate shorter temporal windows due to increased data density. For seasonal data, however, a minimal one-year window is required in order to properly capture the seasonal oscillation, with more data-years being preferred to better separate the seasonal attractor μ from changes in λ.

If we consider a simple spatially-explicit vegetation model, we can calculate $\frac{\Delta x_i}{\Delta t_i}$ pixel-wise so that we do not need to assume one constant seasonality or mean state for all spatial points but can rather preserve pixel-specific equilibria. We test our method on a non-seasonal control run, a model with seasonality, and two different means of deseasoning/detrending, which are typically used[12] in the estimation of resilience on spatio-temporal data (Supplementary Fig. S2). For our parametric regression, we assume that spatial heterogeneity is fixed and can vary across space – e.g., a mountain range with pockets of high- and low-biomass areas – so that every pixel can have its own baseline equilibrium state; λ is then measured relative to that set of local equilibria. This assumption could be violated for some systems (e.g., pattern-forming vegetation), which would require a different approach to modeling the moving seasonal attractor μ through time.

Vegetation data

To analyze global vegetation resilience, we rely on the widely-used kernel Normalized Difference Vegetation Index (kNDVI)[34] based on 16-day MODIS data (MOD13A2, 2000-2025[35]). We only retain the highest-quality flagged data, and resample the native 1 km data to 5 km via a spatial mean for processing speed and to compare to previous publications[22]. We further remove outliers (kNDVI <0.05 or >1), and retain only points with at least 25% data coverage; we do not gap-fill or interpolate our data. We mask our vegetation grids using MODIS land cover data (MCD12Q1, 2001-2024[62]), first removing non-vegetated areas. In a second step, we mask out areas in which land cover has changed during the study period (e.g., forest to grassland), in order to focus our analysis on relatively stable vegetated ecosystems[22]. We split our data into two periods (2001–2012 and 2013–2025) and compute a single λ estimate for each (Fig. 2a). We further compute the difference between these time periods (Fig. 2b) and spatially filter them (Fig. 2c) using a 5 x 5 pixel window. Within each window, we mask the center pixel unless 75% of the pixels have the same difference sign and at least

50% are valid. This is done to remove regions where λ differences vary significantly in space, indicating poorly-constrained λ estimates.

In order to compare typical (i.e., AC1- and variance-based) approaches to our regression-based λ, we first deseason and detrend the vegetation data and then compute λ_{AC1} and λ_{Var} on those stationary time series. A comparison of the mean λ estimates using different approaches can be seen in Supplementary Fig. S3.

Paleoclimate data

We use NGRIP 2.5 and 5 cm sampled data[40,41] obtained from NGRIP-1 and NGRIP-2, and use the GICC05 (yr b2k) ages included with each data set. We do not resample or otherwise pre-process the data. We compute λ using a 200-year moving window, as was suggested by ref. 39. We compare our λ estimates to the timing of DO events from ref. 39, using the main onset age (i.e., the oldest date) to match the start of a DO event defined by abrupt warming at the bottom of the interstadial interval. We incorporate a simple linear drift term in our parametric regression model:

$$\frac{\Delta x_i}{\Delta t_i} = \underbrace{-\lambda}_{\beta_1} x_i + \underbrace{\lambda b}_{\beta_2} t_i + \underbrace{(\lambda a + b)}_{\beta_0} + \epsilon_i \tag{24}$$

as the NGRIP data does not resolve seasonality and is too short to incorporate, e.g., Milankovich cycles. As precise $\delta^{18}O$ estimates can sometimes yield identical neighboring measurements, we incorporate a small instrument uncertainty (± 0.1‰)[40] to maintain the stability of weighted least squares against instances of $\frac{\Delta x_i}{\Delta t_i} = 0$, especially over sections with short depth intervals where Δx_i can be dominated by instrument noise.

The ice core age model (GICC05) incorporates increasing uncertainties as deeper sections of the core are studied[30,40]. To account for this, we use the Maximum Counting Error (MCE) provided with the NGRIP data to estimate fractional timing errors $f \approx \frac{\mathrm{MCE}}{|t|}$ for each sliding window. As λ is expressed as a rate, fractional uncertainty in the time domain t (i.e., a stretch in Δt_i) is propagated directly into uncertainty in λ. We hence use the window-averaged mean fractional error $\bar{f} = \frac{\overline{\mathrm{MCE}}}{|\bar{t}|}$ and expand the uncertainty envelope of λ as $\sigma_{\lambda,\,\mathrm{time}} = |\lambda|\bar{f}$. We estimate the total standard error SE_λ by combining the statistical fitting error SE_{fit} (from the weighted least squares regression) and the age-model error as:

$$SE_\lambda = \sqrt{SE_{\mathrm{fit}}^2 + \left(|\lambda|\frac{\overline{\mathrm{MCE}}}{|\bar{t}|}\right)^2} \tag{25}$$

With this approach, we expand the λ uncertainty envelopes of deep sections of the core where chronologies are less well-constrained, taking into account the layer counting (MCE) uncertainties relative to the GICC05 ages.

Glacier velocity data

We use ITS_LIVE v2 glacier velocity data[50,51] sampled over a test glacier with known surging[44] in the Karakoram[13] (RGI2000-v7.0-G-13-05693, [71.907, 38.837])[63] (Fig. 4). We exclude data from before the Landsat 8 era in order to yield a relatively dense time series of velocity estimates. We test both a single point in a known surging region (Fig. 4a) and the collection of all glacier centerline points, sampled every 500 m (Fig. 4b). We do not resample the given surface velocity data to even time steps, as has been done in previous work[13,64], as our method natively accounts for irregular time sampling. We rely on harmonic terms in our parametric regression model to account for seasonal oscillations in glacier velocity, as well as a second-order polynomial (quadratic) trend to account for non-linear acceleration:

$$\frac{\Delta x_i}{\Delta t_i} = \underbrace{-\lambda}_{\beta_1} x_i + \underbrace{(\lambda b + 2c)}_{\beta_2} t_i + \underbrace{\lambda c}_{\beta_3} t_i^2 + \sum_{k=1}^{n}\left[A_k \sin(k\omega t_i) + B_k \cos(k\omega t_i)\right] + \underbrace{(\lambda a + b)}_{\beta_0} + \epsilon_i \tag{26}$$

We compute λ using a three-year moving window (Fig. 4).

ITS_LIVE glacier velocity data are defined over overlapping time spans, and provide velocity (v_i) estimates which are averages over those observation intervals of length S_i. Each velocity estimate v_i can hence be written as an average instantaneous velocity $u(t)$ over an observation span:

$$v_i = \frac{1}{S_i}\int_{t_i - S_i/2}^{t_i + S_i/2} u(t)\,dt \tag{27}$$

where t_i is the midpoint of the span and S_i its duration. Directly computing $\frac{\Delta v_i}{\Delta t_i}$ using fixed points (e.g., midpoints of the spans) can thus inflate derivatives when the spans overlap. We therefore incorporate the span interval S_i in our λ estimates directly; we maintain the span midpoints for gap filtering and constructing the harmonic regression terms. To integrate the variable spans S_i into our λ estimates, we define $\Delta t_{f,i}$ based on the total temporal footprint of consecutive pairs:

$$\Delta t_{f,i} = \max(t_{i+1} + S_{i+1}/2,\ t_i + S_i/2) - \min(t_{i+1} - S_{i+1}/2,\ t_i - S_i/2) \tag{28}$$

The span-aware derivative is thus:

$$\frac{\Delta v_i}{\Delta t_{f,i}} = \frac{v_{i+1} - v_i}{\Delta t_{f,i}} \tag{29}$$

and instrument error is propagated over that footprint as:

$$\sigma_{\frac{\Delta v_i}{\Delta t_{f,i}}} = \frac{\sqrt{\sigma_{v_i}^2 + \sigma_{v_{i+1}}^2}}{\Delta t_{f,i}} \tag{30}$$

where σ_{v_i} is the uncertainty provided for each span-averaged glacier velocity estimate.

The overlapping span footprint $\Delta t_{f,i}$ will always be larger than the midpoint gap between sequential observations ($\Delta t_{m,i}$), which serves to suppress or even eliminate derivative inflation for closely-spaced span midpoints. For example, two neighboring spans of 10 and 100 days with a one-day midpoint separation would yield a small $\Delta t_{m,i}$ (1 day) and a much larger span-aware $\Delta t_{f,i}$ (105 days); for the same Δv_i, the derivatives would vary greatly. Hence, the span-aware λ derived from $\frac{\Delta v_i}{\Delta t_{f,i}}$ is attenuated when compared to the naive estimate ($\frac{\Delta v_i}{\Delta t_{m,i}}$), with the factor difference being approximately $\frac{\Delta t_{f,i}}{\Delta t_{m,i}}$. Both metrics, however, can be used as statistical early-warning signals of oncoming surges; they differ in how strictly they adhere to the underlying mathematical framework of Langevin dynamics. The span-aware λ is more physically consistent, while the naive midpoint λ is conceptually simpler at the cost of inflating λ estimates in proportion to the span overlap in the underlying data. Further, they yield different outlier distributions – the span-aware λ tends to have many fewer outliers as it corrects for the derivative inflation of the midpoint-based approach. This does not make a substantial difference for the one-dimensional case, as outliers are handled cleanly by our robust linear solver. For the two-dimensional case, however, large coherent blocks of outliers (due to, e.g., many inflated derivatives over the entire glacier coherently in time) can mislead our robust solver and yield noise-dominated λ

estimates. To minimize potential biases from these inflated derivatives, we rely on a span-aware λ estimate in both the one- and two-dimensional cases (Fig. 4).

Data availability

All data is publicly available. Synthetic data can be generated via code on Zenodo: https://doi.org/10.5281/zenodo.19731233[59]. Vegetation data was accessed via Google Earth Engine[65]; the original data can be found here: https://doi.org/10.5067/MODIS/MOD13A2.061[35]. NGRIP ice core data can be found here: https://www.iceandclimate.nbi.ku.dk/data/[40,41]. ITS_LIVE glacier velocity data was accessed via the public Python API[50,51].

Code availability

Synthetic spatial data creation and glacier velocity data access codes are available here: https://doi.org/10.5281/zenodo.18031340, and one-dimensional synthetic models, regression analysis, and Figure plotting scripts are available here: https://doi.org/10.5281/zenodo.19731233.

References

1. Lenton, T. M. et al. Tipping elements in the Earth's climate system. *Proc. Natl. Acad. Sci. USA* **105**, 1786–1793 (2008).
2. Boers, N., Ghil, M. & Stocker, T. F. Theoretical and paleoclimatic evidence for abrupt transitions in the earth system. *Environ. Res. Lett.* **17**, 093006 https://doi.org/10.1088/1748-9326/ac8944 (2022).
3. Lenton, T. M. et al. Remotely sensing potential climate change tipping points across scales. *Nat. Commun.* **15**, 343 https://doi.org/10.1038/s41467-023-44609-w (2024).
4. Forzieri, G., Dakos, V., McDowell, N. G., Ramdane, A. & Cescatti, A. Emerging signals of declining forest resilience under climate change. *Nature* **608**, 534–539 (2022).
5. Boers, N. & Rypdal, M. Critical slowing down suggests that the western greenland ice sheet is close to a tipping point. *Proc. Natl. Acad. Sci. USA* **118**, e2024192118 https://doi.org/10.1073/pnas.2024192118 (2021).
6. Boers, N. Observation-based early-warning signals for a collapse of the atlantic meridional overturning circulation. *Nat. Clim. Change* **11**, 680–688 https://doi.org/10.1038/s41558-021-01097-4 (2021).
7. Boers, N. et al. Destabilization of earth system tipping elements. *Nat. Geosci.* **18**, 949–960 (2025).
8. Carpenter, S. R. & Brock, W. A. Rising variance: a leading indicator of ecological transition. *Ecol. Lett.* **9**, 311–8 (2006).
9. Dakos, V. et al. Slowing down as an early warning signal for abrupt climate change. *Proc. Natl. Acad. Sci. USA* **105**, 14308–12 https://doi.org/10.1073/pnas.0802430105 (2008).
10. Scheffer, M. et al. Early-warning signals for critical transitions. *Nature* **461**, 53–9 (2009).
11. Dakos, V., van Nes, E. H., Donangelo, R., Fort, H. & Scheffer, M. Spatial correlation as leading indicator of catastrophic shifts. *Theor. Ecol.* **3**, 163–174 https://doi.org/10.1007/s12080-009-0060-6 (2010).
12. Smith, T., Traxl, D. & Boers, N. Empirical evidence for recent global shifts in vegetation resilience. *Nat. Clim. Change* **12**, 477–484 https://doi.org/10.1038/s41558-022-01352-2 (2022).
13. Smith, T., Morr, A., Bookhagen, B. & Boers, N. Predicting instabilities in transient landforms and interconnected ecosystems. *Nat. Commun.* **17**, 1316 (2025).
14. Grziwotz, F. et al. Anticipating the occurrence and type of critical transitions. *Sci. Adv.* **9**, eabq4558 https://doi.org/10.1126/sciadv.abq4558 (2023).
15. Morr, A., Riechers, K., Gorjão, L. R. & Boers, N. Anticipating critical transitions in multidimensional systems driven by time- and state-dependent noise. *Phys. Rev. Res.* **6**, 033251 (2024).
16. Riechers, K. et al. Stable stadial and interstadial states of the last glacial's climate identified in a combined stable water isotope and dust record from greenland. *Earth Syst. Dyn.* **14**, 593–607 (2023).
17. Tabar, M. R. R. et al. Revealing higher-order interactions in high-dimensional complex systems: A data-driven approach. *Phys. Rev. X* **14**, 011050 (2024).
18. Bury, T. M., Bauch, C. T. & Anand, M. Detecting and distinguishing tipping points using spectral early warning signals. *J. R. Soc. Interface* **17**, 20200482 https://doi.org/10.1098/rsif.2020.0482 (2020).
19. Morr, A. & Boers, N. Detection of approaching critical transitions in natural systems driven by red noise. *Phys. Rev. X* **14**, 021037 https://doi.org/10.1103/PhysRevX.14.021037 (2024).
20. Bury, T. M. et al. Deep learning for early warning signals of tipping points. *Proc. Natl. Acad. Sci. USA* **118**, e2106140118 https://doi.org/10.1073/pnas.2106140118 (2021).
21. Huang, Y., Bathiany, S., Ashwin, P. & Boers, N. Deep learning for predicting rate-induced tipping. *Nat. Mach. Intell.* **6**, 1556–1565 (2024).
22. Smith, T. & Boers, N. Reliability of vegetation resilience estimates depends on biomass density. *Nat. Ecol. Evol.* **7**, 1799–1808 https://doi.org/10.1038/s41559-023-02194-7 (2023).
23. Cleveland, R. B., Cleveland, W. S., McRae, J. E. & Terpenning, I. Stl: A seasonal-trend decomposition procedure based on loess. *J. Off. Stat.* **6**, 3–73 (1990).
24. Djikstra, H. *Nonlinear Climate Dynamics* (Cambridge University Press, New York, 2013).
25. Rietkerk, M., Skiba, V., Weinans, E., Hébert, R. & Laepple, T. Ambiguity of early warning signals for climate tipping points. *Nat. Clim. Change* **6**, 1556–1565 (2025).
26. Liu, T. et al. Data gaps and outliers distort critical-slowing-down-based resilience indicators. *Sci. Adv.* **12**, https://doi.org/10.1126/sciadv.aee1916 (2026).
27. Smith, T. et al. Reliability of resilience estimation based on multi-instrument time series. *Earth Syst. Dyn.* **14**, 173–183 (2023).
28. Ben-Yami, M., Skiba, V., Bathiany, S. & Boers, N. Uncertainties in critical slowing down indicators of observation-based fingerprints of the Atlantic overturning circulation. *Nat. Commun.* **14**, 8344 (2023).
29. Ben-Yami, M., Morr, A., Bathiany, S. & Boers, N. Uncertainties too large to predict tipping times of major earth system components from historical data. *Sci. Adv.* **10**, https://doi.org/10.1126/sciadv.adl4841 (2024).
30. Boers, N., Goswami, B. & Ghil, M. A complete representation of uncertainties in layer-counted paleoclimatic archives. *Climate* **13**, 1169–1180 (2017).
31. Boulton, C. A., Lenton, T. M. & Boers, N. Pronounced loss of amazon rainforest resilience since the early 2000s. *Nat. Clim. Change* **12**, 271–278 https://doi.org/10.1038/s41558-022-01287-8 (2022).
32. Blaschke, L. L. et al. Spatial correlation increase in single-sensor satellite data reveals loss of amazon rainforest resilience. *Earth's. Future* **12**, e2023EF004040 https://doi.org/10.1029/2023EF004040 (2024).
33. Verbesselt, J. et al. Remotely sensed resilience of tropical forests. *Nat. Clim. Change* **6**, 1028–1031 https://doi.org/10.1038/nclimate3108 (2016).
34. Camps-Valls, G. et al. A unified vegetation index for quantifying the terrestrial biosphere. *Sci. Adv.***7**, https://doi.org/10.1126/sciadv.abc7447 (2021).
35. Didan, K. Modis/terra vegetation indices 16-day l3 global 1km sin grid v061. https://doi.org/10.5067/MODIS/MOD13A2.061 (2021).
36. Moesinger, L. et al. The global long-term microwave vegetation optical depth climate archive (vodca). *Earth Syst. Sci. Data* **12**, 177–196 (2020).
37. Hummel, C., Boers, N. & Rypdal, M. Inconclusive early warning signals for dansgaard-oeschger events across greenland ice cores. *Earth Syst. Dyn.* **16**, 2035–2062 (2025).
38. Boettner, C., Klinghammer, G., Boers, N., Westerhold, T. & Marwan, N. Early-warning signals for cenozoic climate transitions. *Quat. Sci. Rev.* **270**, 107177 (2021).

39. Boers, N. Early-warning signals for Dansgaard-Oeschger events in a high-resolution ice core record. *Nat. Commun.* **9**, 2556 (2018).
40. members, N. G. I. C. P. High-resolution record of northern hemisphere climate extending into the last interglacial period. *Nature* **431**, 147–151 (2004).
41. Gkinis, V., Simonsen, S., Buchardt, S., White, J. & Vinther, B. Water isotope diffusion rates from the northgrip ice core for the last 16,000 years - glaciological and paleoclimatic implications. *Earth Planet. Sci. Lett.* **405**, 132–141 (2014).
42. Mitsui, T. & Boers, N. Statistical precursor signals for dansgaard-oeschger cooling transitions. *Climate* **20**, 683–699 (2024).
43. Boers, N., Ghil, M. & Rousseau, D.-D. Ocean circulation, ice shelf, and sea ice interactions explain dansgaard-oeschger cycles. *Proc. Natl. Acad. Sci. USA* **115**, E11005–E11014 (2018).
44. Kääb, A., Bazilova, V., Leclercq, P. W., Mannerfelt, E. S. & Strozzi, T. Global clustering of recent glacier surges from radar backscatter data, 2017–2022. *J. Glaciol.* **69**, 1515–1523 (2023).
45. Guillet, G. et al. A regionally resolved inventory of high mountain asia surge-type glaciers, derived from a multi-factor remote sensing approach. *Cryosphere* **16**, 603–623 (2022).
46. Ou, H.-W. A theory of glacier dynamics and instabilities part 1: Topographically confined glaciers. *J. Glaciol.* **68**, 1–12 (2022).
47. Benn, D. I., Fowler, A. C., Hewitt, I. & Sevestre, H. A general theory of glacier surges. *J. Glaciol.* **65**, 701–716 https://doi.org/10.1017/jog.2019.62 (2019).
48. Benn, D. I., Hewitt, I. J. & Luckman, A. J. Enthalpy balance theory unifies diverse glacier surge behaviour. *Ann. Glaciol.* **63**, 88–94 https://doi.org/10.1017/aog.2023.23 (2023).
49. Thøgersen, K., Gilbert, A., Schuler, T. V. & Malthe-Sørenssen, A. Rate-and-state friction explains glacier surge propagation. *Nat. Commun.* **10**, 2823 https://doi.org/10.1038/s41467-019-10506-4 (2019).
50. Gardner, A. S. et al. Increased west antarctic and unchanged east antarctic ice discharge over the last 7 years. *Cryosphere* **12**, 521–547 (2018).
51. Gardner, A. S. et al. ITS_LIVE global glacier velocity data in near real time. *EGUsphere* **2025**, 1–29 (2025).
52. Cleveland, W. S. Robust locally weighted regression and smoothing scatterplots. *J. Am. Stat. Assoc.* **74**, 829–836 (1979).
53. Nadaraya, E. A. On estimating regression. *Theory Probab. Its Appl.* **9**, 141–142 (1964).
54. Watson, G. S. Smooth regression analysis. *Sankhyā Indian J. Stat. Ser. A* **26**, 359–372 (1964).
55. Härdle, W. *Applied Nonparametric Regression* (Cambridge University Press, 1990).
56. Hart, J. D. Kernel regression estimation with time series errors. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **53**, 173–187 (1991).
57. Opsomer, J., Wang, Y. & Yang, Y. Nonparametric regressin with correlated errors. *Stat. Sci.* **16**, 134–153 (2001).
58. Shaman, P. & Stine, R. A. The bias of autoregressive coefficient estimators. *J. Am. Stat. Assoc.* **83**, 842–848 (1988).
59. Smith, T. et al. Estimating the resilience of non-stationary systems. *Zenodo* https://doi.org/10.5281/zenodo.19731233 (2026).
60. Frisch, R. & Waugh, F. V. Partial time regressions as compared with individual trends. *Econometrica J. Economet. Soc.* **1**, 387–401 (1933).
61. Böttcher, F. et al. Reconstruction of complex dynamical systems affected by strong measurement noise. *Phys. Rev. Lett.* **97**, 090603 (2006).
62. Friedl, M. & Sulla-Menashe, D. Modis/terra+aqua land cover type yearly l3 global 500m sin grid v061 [data set]. https://doi.org/10.5067/MODIS/MCD12Q1.061 (2022).
63. RGI 7.0 Consortium. Randolph glacier inventory - a dataset of global glacier outlines, version 7.0. https://doi.org/10.5067/f6jmovy5navz (2023).
64. Charrier, L. et al. Ticoi: an operational python package to generate regular glacier velocity time series. *Cryosphere* **19**, 4555–4583 (2025).
65. Gorelick, N. et al. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sens. Environ.* **202**, 18–27 https://doi.org/10.1016/j.rse.2017.06.031 (2017).

Author contributions

T.S. conceived and designed the study, processed the data, and performed the numerical analysis. T.S. wrote the manuscript with contributions from A.M., C.S., and N.B.

Funding

T.S. acknowledges support from the DFG STRIVE project (SM 710/2-1), the ClimTip project, and the Universität Potsdam Remote Sensing Computational Cluster. This is ClimTip contribution #152; the ClimTip project has received funding from the European Union's Horizon Europe research and innovation program under grant agreement No. 101137601. N.B. acknowledges additional funding by the Volkswagen Foundation, the European Space Agency Climate Change Initiative (ESA-CCI) Tipping Elements SIRENE project (contract no. 4000146954/24/I-LR), and the Past to Future (P2F) project, which has received funding from the European Union's Horizon Europe research and innovation program under grant agreement No. 101184070. Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41467-026-78489-7.

Correspondence and requests for materials should be addressed to Taylor Smith.

Peer review information *Nature Communications* thanks M. Reza Rahimi Tabar and the other anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2026