



Bayesian analysis of early warning signals using a time-dependent model

Eirik Myrsvoll-Nilsen, Luc Hallali, and Martin Rypdal

Department of Mathematics and Statistics, UiT The Arctic University of Norway, 9037 Tromsø, Norway

Correspondence: Eirik Myrsvoll-Nilsen (eirik.myrsvoll-nilsen@uit.no)

Received: 14 February 2024 – Discussion started: 19 February 2024

Revised: 11 April 2025 – Accepted: 12 April 2025 – Published: 24 September 2025

Abstract. A tipping point is defined by the IPCC as a critical threshold beyond which a system reorganizes, often abruptly and/or irreversibly. Tipping points can be crossed solely by internal variation in the system or by approaching a bifurcation point where the current state loses stability, which forces the system to move to another stable state. It can be shown that before a bifurcation point is reached there are observable changes in the statistical properties of the state variable. These are known as early warning signals and include increased fluctuation and autocorrelation time. It is currently debated whether or not Dansgaard–Oeschger (DO) events, which are abrupt warmings of the North Atlantic region which occurred during the last glacial period, are preceded by early warning signals. To express the changes in statistical behavior we propose a model based on the well-known first-order autoregressive (AR) process, with modifications to the autocorrelation parameter such that it depends linearly on time. In order to estimate the time evolution of the autocorrelation parameter we adopt a hierarchical Bayesian modeling framework, from which Bayesian analysis can be performed using the methodology of integrated nested Laplace approximations. We then apply the model to segments of the oxygen isotope ratios from the Northern Greenland Ice Core Project record corresponding to 17 DO events. Statistically significant early warning signals are detected for a number of DO events, which suggests that such events could indeed exhibit signs of ongoing destabilization and may have been caused by approaching a bifurcation point. The methodology developed to perform the given early warning analyses can be applied more generally and is publicly available as the R package `INLA.ews`.

1 Introduction

If the state of a component of the climate system changes from one stable equilibrium to another, either by crossing some threshold in the form of a boundary of unstable fixed points separating two basins of attraction or by having the initial equilibrium destabilize, it is said to have crossed a tipping point. Components of the Earth system have experienced tipping points numerous times in the past, leading to abrupt transitions in the climate system. These transitions are well documented in paleoclimatic proxy records. Notably, in Greenland ice core records of oxygen isotope ratios ($\delta^{18}\text{O}$) and dust concentrations there is evidence that large and abrupt climatic transitions from Greenland stadial (GS) to Greenland interstadial (GI) conditions took place in the last glacial interval (110 000–12 000 years be-

fore 2000 CE, hereafter denoted yr b2k). These transitions are known as Dansgaard–Oeschger (DO) events (Dansgaard et al., 1984, 1993) and initialize climatic cycles where the temperature increases substantially, up to 16.5 °C for single events, over the course of a few decades. This is followed by a more gradual cooling, over centuries to millennia, returning to the GS state. A total of 17 DO events (Rasmussen et al., 2014) have been found for the past 60 kyr and they represent some of the most pronounced examples of abrupt transitions in past climate observed in paleoclimatic records.

It is widely accepted that such transitions are related to a change in the meridional overturning circulation (MOC) (Lynch-Stieglitz, 2017; Henry et al., 2016; Menviel et al., 2014, 2020; Bond et al., 1999), possibly caused by loss of sea ice in the North Atlantic. However, the physical mecha-

nisms that caused these changes in the MOC and how they triggered DO events are less understood. Some studies have found that DO events exhibit a periodicity of 1470 years (Schulz, 2002), which has made some scientists suggest that the events have been triggered by changes in the Earth system caused by quasi-periodic changes in the solar forcing (Braun et al., 2005). Others suggest that the transitions have been triggered by random fluctuations in the Earth system, without any significant changes to the underlying system caused by external forcing (Ditlevsen et al., 2007). Treating the GS and GI states as stable equilibria in a dynamical system representing the Greenland climate, and studying the statistical behavior related to the stability of the system in the period preceding DO events, can help determine whether or not they are forced or random and thus possibly constrain the number of plausible physical causes that trigger the events.

Let a time-dependent state variable $x(t)$, representing for example the $\delta^{18}\text{O}$ ratio, vary over some potential $V(x)$ with stochastic forcing corresponding to a white noise process $dB(t)$, expressed as the derivative of a Brownian motion; then the stability of the system can be modeled using the stochastic differential equation:

$$dx(t) = F(x(t))dt + \sigma dB(t). \quad (1)$$

One could interpret this equation as describing the motion of some particle in the presence of a potential $V(x)$, with drift expressed by $F(x) = -V'(x)$ and a diffusion term $\sigma dB(t)$ describing the noise that acts on the particle. Points where $dx(t)/dt = 0$ are fixed points. These are stable if a small perturbation of the state variable near the fixed point decays in time and unstable otherwise. Figure 1a illustrates an example of a potential with two valleys corresponding to stable fixed points x_1 and x_3 that are separated by an unstable fixed point x_2 . If a state variable near x_1 crosses x_2 into the basin of attraction of x_3 solely from perturbations caused by internal variation of the system, then the associated tipping point is said to be noise-induced. However, if the dynamics of the system depend on some slowly varying control parameter $\mu(t)$, then the equilibrium points may shift, vanish, or spawn as a function of $\mu(t)$. This means that an equilibrium state of a system can change over time and eventually be lost, making the system move to another equilibrium, as illustrated in Fig. 1a–c. Points in the control parameter space for which the qualitative behavior of a system changes, e.g., changes in stability or the number of fixed points, are called bifurcation points. Figure 1d illustrates these changes using a bifurcation diagram, where the stable fixed points x_1 (lower solid curve) and x_3 (upper solid curve) are separated by the unstable fixed points x_2 (middle dashed curve). Critical transitions caused by the control parameter crossing a bifurcation point are said to be bifurcation-induced tipping points.

The nature of an equilibrium can be investigated by examining the linear approximation in its nearby domain. Lin-

earizing Eq. (1) around some stable fixed point x_s yields

$$dx(t) = -\lambda(x(t) - x_s)dt + \sigma dB(t), \quad (2)$$

where $\lambda = -F'(x_s)$. This is known as the Langevin stochastic differential equation and has the solution

$$x(t) = x_0 + \int_{-\infty}^t g(t-s)dB(s), \quad (3)$$

with Green's function,

$$g(t) = \begin{cases} \exp(-\lambda t), & x \geq 0 \\ 0, & x < 0 \end{cases}. \quad (4)$$

This solution forms an Ornstein–Uhlenbeck (OU) process with variance $\text{Var}(x_t) = \sigma^2/(2\lambda)$. Under discretization this is a first-order autoregressive (AR) process,

$$x_t = \phi x_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}\left(0, \frac{1-\phi^2}{2\lambda}\sigma^2\right), \quad (5)$$

with the lag-one autocorrelation parameter $\phi = \exp(-\lambda\Delta t)$.

When the control parameter approaches a bifurcation point the restoring rate λ goes to zero and consequently the variance and autocorrelation of the state variable will increase, as can be observed in Fig. 1d. This phenomenon was first demonstrated by inspecting the power spectra of a simple physical model by Wiesenfeld (1985). The idea was later extended to a complex Earth system model by Held and Kleinen (2004) and first applied to real data by Dakos et al. (2008). These changes in statistical behavior are called early warning signals (EWSs) of the bifurcation point, or critical slowing down (Lenton et al., 2012; Dakos et al., 2008), and can be used as precursors to help determine whether or not a tipping point is imminent. EWSs are derived from the linearization of the system around its fixed points; however, even in cases where the system dynamic is far from its equilibrium the same changes in statistical behavior can be found with a delay (Ritchie and Sieber, 2016).

Recent studies have discovered that several components in the Earth system exhibit EWSs and are at risk of approaching or have already reached a tipping point. This includes the western Greenland ice sheets (Boers and Rypdal, 2021), the Atlantic meridional overturning circulation (Boers, 2021), and the Amazon rainforest (Boulton et al., 2022).

Analysis of EWSs for DO events in the Greenland ice core record has been conducted by others, e.g., Ditlevsen and Johnsen (2010), who estimated the variance and autocorrelation over a sliding window where the system was assumed to be driven by white noise. Under these assumptions they were unable to detect a statistically significant increase in EWSs, suggesting that DO events are noise-induced. However, using different model assumptions, Rypdal (2016) was able to detect statistically significant EWSs in an ensemble of




Figure 1. Panels (a)–(c) show the potential over the state variables before, at, and after the control parameter has reached the bifurcation point μ^* . Panel (a) shows the potential and fixed points x_1 , x_2 , and x_3 for some $\mu < \mu^*$, and panels (b)–(c) show the same for $\mu = \mu^*$ and $\mu > \mu^*$, respectively. When the control parameter approaches the bifurcation point μ^* , the stability of the stable fixed point x_1 decreases and eventually collapses with x_2 . In panel (b) they coincide as a single unstable saddle point, leaving x_3 as the only (stable) fixed point. (d) Bifurcation diagram describing a bifurcation-induced tipping point. The black curve represents the fixed points of the state variable x given the linearly changing control parameter $\mu \in (-1, 1)$. The solid curves represent stable fixed points x_1 and x_3 , and the dashed curve represents unstable fixed points x_2 . The red line represents a simulated process. As the control parameter μ approaches the bifurcation point μ^* the stability of x_1 decreases, which is expressed by increased variance and autocorrelation in the simulated process, causing the system to cross the tipping point x_2 prematurely.

DO events. This was achieved by analyzing individual frequency bands separately, using a fractional Gaussian noise (fGn) (Mandelbrot and Van Ness, 1968) model to describe the noise. These results were corroborated by Boers (2018), who applied a similar strategy to a higher resolved version of the Northern Greenland Ice Core Project (NGRIP) $\delta^{18}\text{O}$ dataset (Andersen et al., 2004; Gkinis et al., 2014).

Most approaches for detecting EWSs in the current literature require estimation of statistical properties such as variance and autocorrelation in a sliding window. Consequently, this requires a choice on the length of the window. Using a small window will allow the momentary state to be better depicted, but there will be fewer points used in the estimation and hence accuracy will suffer. On the other hand, if a larger window is used the estimated statistics will be more accurate but less representative of the momentary state as it represents an average over a larger timescale. The opti-

mal choice of window length should ideally represent a good trade-off between accuracy and ability to represent momentary evolution, but this can be hard to determine in practice. In this paper we circumvent this issue and present a model-based approach where such a compromise is not required. By assuming that the autocorrelation parameter is time-dependent, following a specific linear structure, it is possible to formulate this into a hierarchical Bayesian model for which well-known computational frameworks can be applied. A Bayesian approach has the additional benefit of providing uncertainty estimates in the form of posterior distributions.

The paper is structured as follows. A description of the data used in this paper is included in Sect. 2. Section 3 details our methodology, including how we treat time dependence, how to formulate our model as a hierarchical Bayesian model, and how to perform statistical inference efficiently.

Results are presented in Sect. 4 where our framework is applied first to simulated data, then to Dansgaard–Oeschger events observed in the $\delta^{18}\text{O}$ data from the NGRIP record. Our results are compared with those obtained by Ditlevsen and Johnsen (2010), Rypdal (2016), and Boers (2018). Further discussion and conclusions are provided in Sect. 5.

2 NGRIP ice core data

The $\delta^{18}\text{O}$ ratios are frequently used in paleoscience as proxies for temperature at the time of precipitation (Johnsen et al., 1992, 2001; Dansgaard et al., 1993; Andersen et al., 2004), where higher ratios signal colder climates and, conversely, warmer climates tend to result in lower ratios. We employ the $\delta^{18}\text{O}$ proxy record from the Northern Greenland Ice Core Project (NGRIP) (North Greenland Ice Core Project members, 2004; Gkinis et al., 2014; Ruth et al., 2003). There are currently two different versions of the NGRIP/GICC05 data at different resolutions. We will apply our methodology to the higher-resolution record, which is sampled every 5 cm in depth. The NGRIP $\delta^{18}\text{O}$ proxy record is defined on a temporal axis given by the Greenland Ice Core Chronology 2005 (GICC05) (Vinther et al., 2006; Rasmussen et al., 2006; Andersen et al., 2006; Svensson et al., 2008) which thus pairs the $\delta^{18}\text{O}$ measurements with a corresponding age, stretching back to 60 kyr b2k. We use segments of the $\delta^{18}\text{O}$ record corresponding to Greenland stadial phases preceding DO onsets, as given by Table 2 of Rasmussen et al. (2014). The data used in this paper can be downloaded from <https://www.iceandclimate.nbi.ku.dk/data/> (last access: 11 April 2025).

3 Methodology

During critical slowing down stationarity can no longer be assumed as we expect both the autocorrelation and variance to increase. For an AR(1) process $\mathbf{x} = (x_1, \dots, x_n)^\top$ sampled at times $t_1, \dots, t_n$, we assume that the increase in autocorrelation can be expressed by representing the lag-one autocorrelation parameter as a linear function of time:

$$\phi(t) = a + bt, \quad 0 \leq t \leq 1, \quad (6)$$

where a and b are two unknown parameters. The time-dependent AR(1) process is expressed by the difference equation given in Eq. (5) and the joint vector of variables $\mathbf{x} = (x_1, \dots, x_n)^\top$ forms a multivariate Gaussian process,

$$\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma), \quad (7)$$

where the covariance matrix is given by

$$\Sigma_{ij} = \text{Cov}(x_i, x_j), \quad (8)$$

and we assume that

$$\text{Var}(x_1) = \frac{1 - \phi^2}{2\lambda} \sigma^2. \quad (9)$$

Since the covariance matrix is always dense for $\phi \in (0, 1)$ it is computationally beneficial to instead work with the inverse covariance matrix, also known as the precision matrix $\mathbf{Q} = \Sigma^{-1}$. For consistency we hereafter use precision,

$$\kappa = \frac{2\lambda}{(1 - \phi^2)\sigma^2}, \quad (10)$$

instead of the variance as the unknown parameter of interest and denote $\kappa(t_k) = 2\lambda\phi(t_k)$ for $t_k = t_2, \dots, t_n$.

It can be shown that for a time-dependent AR(1) process the precision matrix is sparse and equal to

$$\mathbf{Q} = \begin{pmatrix} (\kappa(t_2) + \kappa(t_2)\phi(t_2)^2) & -\kappa(t_2)\phi(t_2) & & & \\ -\kappa(t_2)\phi(t_2) & (\kappa(t_2) + \kappa(t_3)\phi(t_3)^2) & -\kappa(t_3)\phi(t_3) & & \\ & \ddots & \ddots & \ddots & \\ & & -\kappa(t_{n-1})\phi(t_{n-1}) & (\kappa(t_{n-1}) + \kappa(t_n)\phi(t_n)^2) & -\kappa(t_n)\phi(t_n) \\ & & & -\kappa(t_n)\phi(t_n) & \kappa(t_n) \end{pmatrix}. \quad (11)$$

To allow for non-constant time steps $\Delta t_k = t_k - t_{k-1}$ we define

$$\phi(t_k) = e^{-\lambda(t_k)\Delta t_k/c}, \quad (12)$$

where

$$\lambda(t_k) = -\log(a + bt_k), \quad (13)$$

$c = \sum_{k=2}^n \Delta t_k / (n - 1)$, and t_k has been rescaled such that $t_k \in (0, 1)$. This modification guarantees that $\phi(t) \rightarrow 1$ as $\Delta t_k \rightarrow 0$ and $\phi(t) \rightarrow 0$ as $\Delta t_k \rightarrow \infty$. It also ensures that $\phi(t_1) = a$ and $\phi(t_n) = a + b$, which makes the interpretability of the parameters easier.

Gaussian processes with sparse precision matrices are known as Gaussian Markov random fields, and there is a wealth of efficient algorithms for fast Bayesian inference; see, e.g., Rue and Held (2005) for a comprehensive discussion on this topic. These computationally efficient properties are not shared by the fractional Gaussian noise for which both the covariance matrix and the precision matrix are dense. This means that essential matrix operations such as computing the Cholesky decomposition will have a computational cost of $\mathcal{O}(n^3)$ floating point operations (flops), as opposed to $\mathcal{O}(n)$ flops for the AR(1) process. Inference might still be possible to achieve in a reasonable amount of time if the size of the dataset remains sufficiently small. For larger datasets, however, both time and memory consumption may become an issue.

In fitting the model to data it is beneficial that the model parameters are defined on an unconstrained parameter space. We therefore introduce a suitable parameterization for the model parameters. For the precision κ we take the logarithm, $\theta_\kappa = \log \kappa$, and for a and b we use variations of the logistic transformation. Our reasoning is as follows. Assuming the lag-one autocorrelation parameter is defined on the interval $(0, 1)$, and since $t \in [0, 1]$, then the slope must be constrained by

$$|b| < 1. \quad (14)$$

An unconstrained parameterization for b thus reads

$$\theta_b = \log\left(\frac{1+b}{1-b}\right) \iff b = -1 + \frac{2}{1 + \exp(-\theta_b)},$$

$$\theta_b \in (-\infty, \infty). \quad (15)$$

The parameter space for a depends on the current state of b ,

$$0 < a + bt < 1 \iff -bt < a < 1 - bt. \quad (16)$$

Let

$$a_{\text{lower}} = -\min(b, 0) \text{ and } a_{\text{upper}} = 1 - \max(b, 0), \quad (17)$$

and then an unconstrained parameterization for a is given by

$$\theta_a = \log\left(\frac{a - a_{\text{lower}}}{a_{\text{upper}} - a}\right) \iff a = a_{\text{lower}} + \frac{a_{\text{upper}} - a_{\text{lower}}}{1 + \exp(-\theta_a)},$$

$$\theta_a \in (-\infty, \infty). \quad (18)$$

3.1 Bayesian inference

Bayesian analysis presents a powerful framework for estimating model parameters that provides uncertainty quantification and allows us to incorporate prior knowledge about the parameters. These benefits are both very valuable in making informed decisions regarding climate action. In the Bayesian paradigm the model parameters $\theta = (\kappa, b, a)$ are treated as stochastic variables for which prior knowledge is expressed using a predefined distribution $\pi(\theta)$. These are updated by new information expressed by the likelihood function of the observations $\mathbf{x} = (x_1, \dots, x_n)$, which is a Gaussian distribution here,

$$\pi(\mathbf{x} | \theta) = \mathcal{N}(\mathbf{0}, \mathbf{Q}(\theta)^{-1}), \quad (19)$$

with precision matrix $\mathbf{Q}(\theta)$ as given by Eq. (11). The updated belief is expressed by the posterior distribution $\pi(\theta | \mathbf{x})$, which is obtained using Bayes' rule,

$$\pi(\theta | \mathbf{x}) = \frac{\pi(\mathbf{x} | \theta)\pi(\theta)}{\pi(\mathbf{x})}, \quad (20)$$

where the model evidence $\pi(\mathbf{x})$ is a normalizing constant with respect to θ .

Prior selection is an essential part of any Bayesian analysis and presents a great strength of the Bayesian framework by allowing prior knowledge to be incorporated into models. Since we do not incorporate prior knowledge in this paper, and we wish to maintain objectivity, we will adopt vague prior distributions. These are distributions with large variances that express minimal information about the parameters and allow inference to be primarily driven by the data, as opposed to more informative priors which can guide the posterior to reflect prior knowledge or assumptions. Since the parameter a depends on the value of another parameter b we

assign a conditional prior, $\pi(a | b)$, such that the joint prior is expressed by

$$\pi(\kappa, b, a) = \pi(\kappa)\pi(b)\pi(a | b). \quad (21)$$

Specifically, all analyses performed in this paper assume the same set of priors, unless otherwise specified. κ is assigned a gamma distribution with shape 1 and rate 0.1. b is assigned a uniform prior on $(-1, 1)$ and $a | b$ is assigned a uniform prior on $(a_{\text{lower}}, a_{\text{upper}})$.

The main goal of this paper is to detect whether or not an early warning signal can be observed in the data. We are therefore primarily interested in the marginal posterior distribution of the slope parameter b ,

$$\pi(b | \mathbf{x}) = \int \pi(\theta | \mathbf{x}) d\kappa da. \quad (22)$$

Typically, marginal posterior distributions can be evaluated using Markov chain Monte Carlo approaches (Robert and Casella, 2004), but these are very often time-consuming and could potentially be sensitive to convergence issues. However, since our model is Gaussian with a sparse precision matrix we instead use the computationally superior alternative of integrated nested Laplace approximation (INLA) (Rue et al., 2009), which is available as an R package at <http://www.r-inla.org> (last access: 11 April 2025).

To make the methodology more accessible we have released the code associated with this model as an R package titled `INLA.ews`, which can be downloaded at <http://www.github.com/eirikmn/INLA.ews> (last access: 11 April 2025). A demonstration of this package can be found in Appendix A, and a detailed description of its features can be found in its accompanying documentation.

3.2 Incorporating forcing

Climate components may also be affected by forcing, which can be measured alongside the climate variable of interest. How the observed component responds to such forcing will be influenced by time dependence. In this subsection we adopt a similar strategy to Myrvoll-Nilsen et al. (2020) with changes to the Green's function to allow for time dependence and non-constant time steps.

Let $F(t)$ denote the known forcing component such that

$$dx(t) = -\lambda x(t) + F(t)dt + dB(t). \quad (23)$$

As shown in Gardiner (2009), the model can then be expressed as the sum of two components,

$$x(t) = v(t) + \xi(t), \quad (24)$$

where $\xi(t)$ is a time-dependent OU process and the forcing response is given by

$$v(t) = \frac{1}{\sqrt{2\lambda(t)\kappa_f}} \int_0^t e^{-\lambda(t)(t-s)} (F(s) + F_0) ds. \quad (25)$$

Here, $\lambda(t) = -\log(\phi(t))$ is the restoring rate, κ_f is an unknown scaling parameter, and F_0 is an unknown shift parameter. These parameters can be estimated using the same Bayesian framework as before, which can be computed with `INLA.ews` by specifying the `forcing` argument in the `inla.ews` function.

4 Results

4.1 Accuracy test on simulated data

To test the accuracy and robustness of the time-dependent AR(1) model we fit the model to a number of simulations. Specifically, we perform accuracy tests using a grid of $b \in [-0.8, 0.8]$ with increments of 0.1 and choose the parameter a corresponding to $\theta_a = 0$. For each b we draw $n_r = 1000$ time series of length $n = 500$ and $n = 1000$ from the time-dependent AR(1) model. The model is fitted using `R-INLA` using priors specified in Sect. 3.1. To quantify the accuracy of the model we compare the posterior marginal mean of the slope $\hat{b} = E(\pi(b | \mathbf{x}))$ to the true values b . We also compute the posterior probability of the slope being positive, denoted $P(b > 0)$. Ideally, we want $\hat{b}$ to be as close to b as possible, and $P(b > 0) > 0.5$ if $b > 0$ and, conversely, $P(b > 0) < 0.5$ if $b < 0$. We also count the number of simulations where an EWS is detected using the threshold $P(b > 0) \geq 0.95$. Since σ only scales the amplitude of the data without affecting the correlation structure we expect similar estimations for a and b regardless of the value of σ . This was confirmed by testing the model on simulations using both $\sigma = 1$ and $\sigma = 10$.

The results of the analysis are presented in Table 1 and displayed graphically as box plots in Fig. 2. Since the posterior distribution of b is skewed, especially when its absolute value approaches 1, ordinary box plots would classify a larger number of points as outliers. We instead use an adjusted box plot proposed by Hubert and Vandervieren (2008), which is better suited for skewed distributions. We obtain decent accuracy of the posterior marginal means $\hat{b}$, with a small, but consistent, underestimation which decreases as n increases. In panels (c) and (d) of Fig. 2 we observe some variation in $P(b > 0)$ for small values of $|b|$ and less so for $|b| \geq 0.2$. This behavior also improves when n increases from 500 to 1000. By counting the number of simulations with detected EWSs for $b \leq 0$ using threshold $P(b > 0) \geq 0.95$ we find 54 false positives for $n = 500$ and 56 for $n = 1000$. Out of 8000 simulations for $b \geq 0.1$ we find 669 false negatives for $n = 500$ and 321 for $n = 1000$, as reported in Table 1.

In order to assess how sensitive the model is to the choice of prior distribution we repeat the same simulation procedure with $n = 500$ using different Gaussian priors on the parameters in their internal scaling,

$$(\log(\kappa), \theta_a, \theta_b) \sim \mathcal{N}(\mathbf{0}, \sigma_\theta \mathbf{I}), \quad (26)$$

where $\sigma_\theta \in \{0.1, 1, 10\}$ and $\mathbf{I}$ is the identity matrix. The posterior marginal means and posterior probabilities $P(b > 0)$

for the different priors are compared to the default prior in Fig. 3. The results show that using the most informative prior, corresponding to $\sigma_\theta = 0.1$, will pull b too much towards the central value of $b = 0$, resulting in worse estimates for b . This pull is also reflected in the posterior probability estimates, $P(b > 0)$, where this prior performs less well compared to the others, although the effect is less strong here. Counting the number of misclassifications, we find, for σ_θ equal to 0.1, 1 and 10, 4, 48, and 67 false positives out of 9000 simulations and 1255, 657, and 631 false negatives out of 8000 simulations, respectively. This is overall quite comparable to using the default priors, which found 54 false positives and 669 false negatives. Overall, we find the model to be quite robust to the choice of prior distributions as most of the priors perform very similarly in terms of the posterior marginal mean and posterior probabilities. However, we also find that using priors that are too informative could cause the model to be overly cautious, making it less able to detect EWSs.

4.2 DO events

We apply our time-dependent AR(1) model to the high-resolution NGRIP $\delta^{18}\text{O}$ record, which is partitioned into stadial and interstadial periods as shown in Fig. 4. This version of the NGRIP record is sampled regularly every 5 cm step in depth but is non-constant in time. Having modified our model to allow for irregular time points we are able to use the raw NGRIP record without having to perform interpolation or other types of pre-processing, such as that of Boers (2018). This grants us a larger dataset for each event, which could significantly improve parameter estimation. Having implemented the model using `INLA` we are able to take advantage of this extra resolution while keeping computational time low.

Some of these datasets appear nonstationary and thus require trend estimation. Since there is no obvious choice of forcing we consider different alternatives for trend components which are then compared. The `R-INLA` framework allows us to very easily incorporate these trends into our model and estimates all model components simultaneously. First, we fit our model to the data without any additional trend, and then we assume a linear trend, followed by a second-order polynomial trend. Finally, we model the trend using a continuous second-order random walk (RW2) spline. More details on the comparison between the different trends are included in Appendix C, which also includes a plot of how well each trend fits the data.

Looking at the fits for each event we observe that most events can be fitted easily with a linear trend or even no trend, but a few events require nonlinearity. We choose the second-order polynomial trend as this gives a nice trade-off between flexibility and simplicity and appears to provide a decent fit for all events. The $\phi(t) = a + bt$ evolutions for all events using second-order polynomial detrending are included in




Figure 2. Box plots representing the results of the accuracy test for $n_r = 1000$ simulated time series for $b \in [-0.8, 0.8]$. The boxes cover the interquartile range (IQR) between the 25 % and 75 % quantiles, and the whiskers represent an adjustment to the more common boundaries of 1.5 times the IQR to better describe skewed distributions. Points that fall outside the whiskers are classified as outliers and are also included in the plot. Panels (a) and (b) show box plots of the posterior marginal mean estimated by INLA for simulations of lengths $n = 500$ and $n = 1000$, respectively. The blue line shows the true b used in the simulation. Panels (c) and (d) show box plots of the estimated posterior probability of the slope being positive against different true values used for simulations of length $n = 500$ and $n = 1000$, respectively.




Figure 3. Prior sensitivity analysis on simulated data using the default priors (black) and Gaussian priors $\mathcal{N}(0, 0.1^2)$ (red), $\mathcal{N}(0, 1)$ (orange), and $\mathcal{N}(0, 10^2)$ (yellow). Panel (a) shows the ensemble average of the posterior marginal mean estimates for b compared to the true value (blue), and panel (b) shows the ensemble average for the posterior probabilities $P(b > 0)$. The shaded gray area in both panels represents the ensemble spread of the estimates from the default prior using the 2.5 % and 97.5 % quantiles.

Table 1. Results from accuracy tests on $n_r = 1000$ simulated time-dependent AR(1) series of length n for each b ranging from -0.8 to 0.8 . The table includes the ensemble averages of the posterior marginal means $\hat{b}$ and posterior probabilities of positive slope $P(b > 0)$ for each value of b and for time series’ lengths of $n = 500$ and $n = 1000$. We also show the number of detected early warning signals using the threshold $P(b > 0) \geq 0.95$.

	$n = 500$	$n = 1000$	$n = 500$	$n = 1000$	$n = 500$	$n = 1000$
True b	$\langle \hat{b} \rangle$	$\langle \hat{b} \rangle$	$\langle P(b > 0) \rangle$	$\langle P(b > 0) \rangle$	$\# P(b > 0) > 0.95$	$\# P(b > 0) > 0.95$
-0.8	-0.781	-0.792	0	0	0	0
-0.7	-0.687	-0.694	0	0	0	0
-0.6	-0.586	-0.595	0	0	0	0
-0.5	-0.487	-0.496	0	0	0	0
-0.4	-0.393	-0.398	0	0	0	0
-0.3	-0.293	-0.299	0.001	0	0	0
-0.2	-0.194	-0.195	0.015	0.001	0	0
-0.1	-0.091	-0.096	0.16	0.074	1	1
0	0	-0.001	0.5	0.498	53	55
0.1	0.1	0.096	0.856	0.926	425	685
0.2	0.196	0.2	0.984	0.998	907	994
0.3	0.298	0.299	0.999	1	999	1000
0.4	0.393	0.396	1	1	1000	1000
0.5	0.492	0.494	1	1	1000	1000
0.6	0.588	0.594	1	1	1000	1000
0.7	0.688	0.695	1	1	1000	1000
0.8	0.783	0.793	1	1	1000	1000

Stadial and interstadial periods




Figure 4. NGRIP $\delta^{18}\text{O}$ proxy record. The time series used in our study are the parts of the curves drawn in turquoise, which are the cold stadial periods preceding the onsets of interstadial periods drawn in red. The red and turquoise vertical bars represent the start and end points of the interstadial periods, respectively.

Fig. 5, and the posterior marginal distribution of the trend, $\pi(b | \mathbf{x})$, is included in Fig. 6.

The models are fitted to the stadial periods preceding each of the 17 DO events, and the posterior probability of $\phi(t) = a + bt$ being increasing, $P(b > 0)$, is compared for all events and trend assumptions. The results are included in Table 2. Using the conventional threshold of $P(b > 0) \geq 0.95$ we are able to detect early warning signals in five events both without detrending and while using linear or second-order polynomial trends and six events using the continuous RW2 model as a trend. Averaging over the estimated $P(b > 0)$ for

all events we are not able to conclude that early warning signals have been found over the ensemble of events for any detrending model.

Having found EWSs in multiple stadial periods preceding DO events therefore indicates that DO events can exhibit evidence of ongoing destabilization, unlike the conclusion formulated in Ditlevsen and Johnsen (2010). These differences in results can be explained by the use of both a higher-resolution dataset and a methodology not involving time windows.




Figure 5. The evolution of the lag-one autocorrelation parameter $a + bt$ for each of the 17 transitions analyzed in this paper. The blue lines represent the posterior marginal means of each Greenland stadial phase, and the red shaded areas represent the 95 % credible intervals (corresponding to the region between the 2.5 % and 97.5 % quantiles of the posterior distribution). The $\delta^{18}\text{O}$ proxy measurements have been detrended using a second-order polynomial.

However, given the absence of EWSs in the ensemble of events does not support the hypothesis that all DO events exhibit signs of ongoing destabilization, and hence one cannot exclude the possibility of some events being purely noise-induced and not approaching a bifurcation point. Our results do, however, suggest that some specific transitions have undergone destabilization, which is in line with the results of Rypdal (2016) and Boers (2018), wherein significant EWSs have also been found only for some specific events. Specifically, Table 2 shows that using a square detrending our results corroborate the EWS found for events 5 and 9 by Rypdal (2016) and the one found for event 11 by Boers (2018); our results also show an EWS for event 2, similarly to these two studies.

These studies use different versions of the NGRIP record from our study and their methodologies differ from ours.

5 Conclusions

In this paper we introduce a Bayesian framework to analyze early warning signals in time series data. Specifically, we define a time-dependent AR(1) process where the lag-one autocorrelation parameter is assumed to increase linearly over

time. The slope of this parameter indicates whether or not memory is increasing and thus if early warning signals are detected. Using a Bayesian approach we automatically obtain uncertainty quantification expressed by posterior distributions and allow prior knowledge to be utilized.

To detect early warning signals of DO events we have applied our model to interstadial periods of the raw 5 cm NGRIP water isotope record. This record is sampled evenly in depth, but not in time, requiring us to make some necessary modifications to allow for non-constant time steps. Using the time-dependent AR(1) model we were unable to detect statistically significant EWSs for the ensemble of 17 DO events and only detected EWSs individually for six events using a second-order polynomial detrending. Unlike Ditlevsen and Johnsen (2010), we find evidence of EWSs in some events, corroborating Rypdal (2016) and Boers (2018). We were, however, unable to conclude that DO events are individually or generally bifurcation-induced. To better compare with other studies, we would have liked to employ a long-range-dependent process such as the fGn. However, this task is more difficult than for the AR(1) process, as necessary modifications have to be made to the model. Moreover, this would also require working with non-sparse precision ma-




Figure 6. The marginal posterior distribution of the trend parameter b . The red dotted line represents $b = 0$ and the red shaded area illustrates the density $P(b > 0)$. If the shaded area is larger than 0.95 we conclude that an EWS is detected.

trices, which are far more computationally demanding. We did attempt to implement the time-dependent fGn model presented by Ryvkin (2015), but we were unable to ensure sufficient stability. This is, however, a very interesting topic for future work.

Currently, our model can only fit an AR(1) process where the lag-one autocorrelation parameter is expressed as a linear function of time, which is not realistic. Although this is sufficient for detecting whether or not there are EWSs ex-

pressed by a linear trend, our model is unable to perform predictions or give an indication of when the tipping point could be reached. More advanced functions for the evolution of the lag-one autocorrelation parameter should be possible but would have to be implemented. One possible extension would be to include a break point such that the memory is constant for all steps before this point and starts increasing or decreasing afterwards. A simple implementation and discussion of this are included in Appendix D. Another extension

Table 2. Table comparing the probability of positive slope $P(b > 0)$ for each event given posterior distributions obtained using the time-dependent AR(1) model on preceding subintervals of lengths n . We ran the model using different trends including no trend (except for the intercept), a linear effect, a second-order polynomial, and a second-order random walk spline. Events where the posterior probability of a positive slope exceeds 0.95 are highlighted in bold. Our results are also compared with the p values obtained from Rypdal (2016) and Boers (2018).

Event	n	No trend	Linear	Square	RW2	Rypdal	Boers
1	667	0.8765	0.854	0.9146	0.9552	$p = 0.02$	–
2	3713	0.9658	0.9902	0.9728	0.9672	$p = 0.008$	$p < 0.05$
3	1370	0.5252	0.5174	0.4893	0.647	–	–
4	270	0.0684	0.078	0.084	0.1007	–	$p < 0.05$
5	962	0.9958	0.9959	0.9959	0.9893	$p = 0.13$	–
6	253	0.28	0.3174	0.2123	0.2068	–	$p < 0.05$
7	315	0.7703	0.7517	0.7132	0.79	–	–
8	345	0.8976	0.9189	0.8878	0.8524	–	–
9	492	0.9857	0.9628	0.953	0.9818	$p = 0.16$	–
10	202	0.019	0.1413	0.0732	0.1001	–	–
11	219	0.967	0.967	0.9643	0.9635	–	$p < 0.05$
12	251	0.1483	0.1319	0.1662	0	–	–
13	373	0.8872	0.8953	0.8912	0.9043	$p = 0.39$	$p < 0.05$
14	91	0.7781	0.914	0.6629	0.7808	–	$p < 0.05$
15	163	0.0227	0.0546	0.0637	0.0446	–	$p < 0.05$
16	162	0.9885	0.9915	0.9935	0.9939	–	–
17	70	0.6748	0.6366	0.6043	0.5855	–	–
Ensemble	–	0.6383	0.654	0.626	0.639	–	–

would be to formulate a model where the memory parameter follows a polynomial $\phi(t) = a + bt^c$, where the exponent term $c > 0$ is an additional hyperparameter. This would perhaps help give an indication of the rate at which the autocorrelation has increased. However, when adding more parameters one needs to be careful to avoid overfitting.

The ability to update prior beliefs in light of new evidence presents a great benefit of a Bayesian approach, and it presents an intuitive framework for iteratively updating the posterior distribution as new data become available by using the posterior distribution from previous analyses as the prior distribution in the analysis. This is of great relevance for monitoring climatic systems suspected of approaching a tipping point.

To make the methodology more accessible we have released the code associated with this model as an R package titled `INLA.ews`. This package performs all analysis and includes functions to plot and print key results from the analysis very easily. Although this paper focuses on the detection of EWSs in DO events observed in Greenland ice core records, our methodology is general and the `INLA.ews` package should be applicable to tipping points observed in other proxy records as well. We have also implemented the option of including forcing, for which the package will estimate the necessary parameters and compute the resulting forcing response. The package is demonstrated on simulated data in Appendix A.

Appendix A: Demonstration of the `INLA.ews` package

We demonstrate the `INLA.ews` package on simulated forced data with non-equidistant time steps. The time steps t_k are obtained by adding Gaussian noise such that $\tilde{t}_k = k + \xi_k$ and normalized $t_k = (\tilde{t}_k - \tilde{t}_0)/(\tilde{t}_n - \tilde{t}_0)$ such that $t_0 = 0$ and $t_n = 1$. We assume a time-dependent AR(1) process of length $n = 1000$ for the observations sampled at times $\tilde{t}_1, \dots, \tilde{t}_n$. The AR(1) process has a scale parameter $\kappa = 0.04$ and time-dependent lag-one autocorrelation $\phi(t) = a + bt_k$ given by $a = 0.3$ and $b = 0.2$. We also include a forcing $F(t)$, obtained by simulation from another AR(1) process with unit variance and the lag-one autocorrelation parameter $\tilde{\phi} = 0.95$. The forcing response is approximated by

$$v(t_k) = \frac{\sigma_f}{\sqrt{2\lambda(t_k)}} \sum_{s=t_0}^{t_k} e^{-\lambda(t_k)(t_k-t_s)} (F_0 + F(s)), \quad (\text{A1})$$

with parameters set to $\sigma_f = 0.1$ and $F_0 = 0$, and added to the simulated observations. We assign $\text{gamma}(1, 0.01)$ priors for the precision parameters κ and κ_f , uniform priors on b and a , and a Gaussian prior $\mathcal{N}(0, 10^2)$ on F_0 . For `INLA.ews`, these priors must be transformed for the unconstrained parameterization $\theta = (\log \kappa, \theta_b, \theta_a, \log \kappa_f, F_0)$ using the change-of-variable formula. The logarithm of the prior distributions is specified by creating a function as follows.

```
my.log.prior <- function(theta) {
  lprior = dgamma(exp(theta[1]), shape=1, rate=0.1) + theta[1] + #kappa
            -theta[2] -2*log(1+exp(-theta[2])) + #b
```

```

    -theta[3] -2*log(1+exp(-theta[3])) + #a
    dgamma(exp(theta[4]), shape=1, rate=0.1) + theta[4] + #kappa_f
    dnorm(theta[5], sd=10, log=TRUE) #F0
  }
  return(lprior)
}

```

It is then passed into `inla.ews` using the `log.prior` argument. The AR(1) model and forcing z sampled at time points `time` can be fitted to the data `y` with INLA using the `inla.ews` wrapper function.

```
results <- inla.ews(data=y, forcing=z,
log.prior=my.log.prior, timesteps=time)
```

The `inla.ews` function computes all posterior marginal distributions, computes summary statistics, formats the results, and returns all information as an `inla.ews` list object. Summary statistics and other important results can be extracted using the `summary` function.

```

> summary(results)

Call:
inla.ews(data = y, forcing = forcing, log.prior=my.log.prior, timesteps=time)

Time used:
  Running INLA Post processing      Total
    503.1737      151.8262      655.5909

Posterior marginal distributions for all parameters have been computed.

Summary statistics for using ar1 model (with forcing):
      mean      sd 0.025quant 0.5quant 0.975quant
a      0.3065 0.0546      0.1974      0.3072      0.4148
b      0.1929 0.0615      0.0524      0.2018      0.2878
sigma   7.0249 0.4522      6.3420      6.9522      8.0780
sigma_f  0.1000 0.0096      0.0862      0.0982      0.1229
F0      -0.0047 0.0223      -0.0514      -0.0034      0.0355

Memory evolution is sampled on an irregular grid.
Summary for first and last point in smoothed trajectory (a+b*time):
      mean      sd 0.025quant 0.5quant 0.975quant
phi0[l] 0.3065 0.0546      0.1974      0.3072      0.4148
phi0[n]  0.4980 0.0587      0.3667      0.5039      0.5945
Mean and 95% credible intervals for forced response have also been computed.

Probability of positive slope 'b' is 0.9954214

Marginal log-Likelihood: -3088.35

```

The results may be displayed graphically using the `plot` function.

```
> plot(results)
```

For this example the estimated memory evolution and forcing response are presented in Fig. A1. The estimated parameters are summarized in Table A1.

Combining forcing with irregular time steps requires more computationally intensive calculations within `rgeneric`, which increases the total computational time to around 10 min, compared to 10 s using any other model configuration. To reduce this we have implemented the model in `cgeneric`, which grants a substantial boost in speed. However, this requires pre-compiled C code using more simplistic priors for the parameters, which cannot be changed without recompiling the source code. Thus there could potentially be a small loss in accuracy of the fitted model at the cost of the improved speed. To use the `cgeneric` version of the model, set `do.cgeneric=TRUE` in the `inla.ews` function call.

Appendix B: Latent Gaussian model formulation

Section 3.1 defines our model within a Bayesian framework. However, in order for the model to be compatible with INLA we require some modifications such that it is expressed in terms of a latent Gaussian model. Latent Gaussian models

represent a subset of hierarchical Bayesian models which are defined in three stages. First, the likelihood function is specified. The likelihood is then expressed using a latent field of unobserved Gaussian variables $\mathbf{w} = (w_1, \dots, w_n)$ whose specification forms the second stage. These depend on a number of unknown hyperparameters. The final stage is to assign prior distributions to the hyperparameters.

Since our model is originally a two-stage model, a Gaussian likelihood that depends on some parameters without an intermediate latent field $\mathbf{w}$, we reshape this into three stages by defining the former likelihood as the latent field $\mathbf{w}$ and the observations $\mathbf{x}$ to be the latent field with some additional negligible noise:

$$\pi(\mathbf{x} | \mathbf{w}, \boldsymbol{\theta}) = \prod_{i=1}^n \mathcal{N}(w_i(\boldsymbol{\theta}), \sigma_x^2), \quad (\text{B1})$$

where $\sigma_x^2 \approx 0$, essentially stating that $\mathbf{x} \approx \mathbf{w}$. This trick does not change our model but creates a reformulation of the model into a latent Gaussian model where the latent field $\mathbf{w}$ is the prior of the mean of the likelihood.

The latent variables $\mathbf{w}$ follows a multivariate Gaussian process:

$$\mathbf{w} | \boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}(\boldsymbol{\theta})^{-1}), \quad (\text{B2})$$

where the precision matrix $\mathbf{Q}$ is given by Eq. (11) and $\boldsymbol{\mu}$ describes any potential trends as specified in the model. We will not discuss such trends here. When using INLA it is essential that the precision matrix is sparse in order to retain computational efficiency.

Since the Gaussian process now describes the latent variables instead of the likelihood, the parameters $\boldsymbol{\theta}$ which govern $\mathbf{w}$ will now be called hyperparameters, since they are the parameters of a prior distribution. The final step of defining the latent Gaussian model is to assign prior distributions to the hyperparameters, but since INLA prefers to work with unconstrained variables these are specified through the parameterizations derived in Sect. 3:

$$\boldsymbol{\theta} \sim \pi(\log \kappa) \pi(\theta_b) \pi(\theta_a | \theta_b). \quad (\text{B3})$$

Transforming priors chosen for (κ, b, a) to the corresponding priors chosen for $(\log \kappa, \theta_b, \theta_a)$ can be done using the change-of-variable formula.

We want to estimate the marginal posterior distribution for all hyperparameters and latent variables. These are computed by evaluating the integrals


$$\pi(w_i | \mathbf{x}) = \int \pi(w_i | \boldsymbol{\theta}, \mathbf{x}) \pi(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta}, \quad (\text{B4})$$

$$\pi(\theta_j | \mathbf{x}) = \int \pi(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta}_{-j}. \quad (\text{B5})$$

Of these we are primarily concerned with the latter, since the latent field will very be similar to the observed values $\mathbf{x}$

Table A1. Underlying values used for simulating the data, along with estimated posterior marginal means and 95 % credible intervals for all hyperparameters.

Parameter	True value	Posterior marginal mean	95 % credible interval
a	0.3	0.306	(0.197, 0.415)
b	0.2	0.193	(0.052, 0.288)
σ	5	7.025	(6.342, 8.078)
σ_f	0.1	0.1	(0.086, 0.123)
F_0	0	−0.005	(−0.051, 0.036)

**Figure A1.** Panel (a) shows the simulated data (gray) where a simulated forcing response (black) has been added. Panel (b) shows the posterior marginal mean (blue) and 95% credible intervals (red) of the forcing response. Panel (c) shows the posterior marginal mean of the lag-one autocorrelation parameter of the simulated data (gray). The fluctuations are caused by being sampled at non-constant time steps. The posterior marginal mean of the smoother evolution of $a + bt$ is included (blue), along with 95 % credible intervals (red) and the true values (white).

since $\sigma_x \approx 0$. To compute these integrals INLA uses various numerical optimization techniques to obtain an appropriate approximation. Most important is the Laplace approximation (Tierney and Kadane, 1986), which is used to approximate the joint posterior distribution,

$$\pi(\theta | \mathbf{x}) \approx \frac{\pi(\mathbf{w}, \theta, \mathbf{x})}{\pi_G(\mathbf{w} | \theta, \mathbf{x})} \Big|_{\mathbf{w}=\mathbf{w}^*(\theta)}, \quad (\text{B6})$$

where $\mathbf{w}^*(\theta)$ is the mode of the latent field $\mathbf{w}(\theta)$ and $\pi_G(\mathbf{w} | \theta, \mathbf{x})$ is the Gaussian approximation of

$$\pi(\mathbf{w} | \theta, \mathbf{y}) \propto \pi(\mathbf{w} | \theta) \pi(\mathbf{x} | \mathbf{w}, \theta). \quad (\text{B7})$$

The methodology is available as the open-source R package R-INLA, which can be downloaded at <http://www.r-inla.org> (last access: 11 April 2025).

Since there are no model components already implemented for R-INLA that meet our specifications we are required to implement the model components ourselves using the custom modeling framework of R-INLA called *rgeneric*. This adds more work and complexity in implementing our model and adds an additional barrier to further adoption of our methodology, which motivated us to create a more user-friendly R package titled *INLA.ews*, available at <http://www.github.com/eirikmn/INLA.ews> (last access: 11 April 2025).

Appendix C: Comparison of different detrending approaches

Since there is no clear choice of forcing for DO events, and not all data windows appear stationary, we assume that there is some unknown trend component reflected in the data. This trend needs to be managed or the estimates of other components will suffer. Often, this is done by first detrending the data, before the parameters of interest are estimated. This bears the risk that variation caused by the time-dependent noise component may be attributed to the trend, and it is therefore better to estimate both the trend and noise components simultaneously. This can be achieved using INLA, which supports many common model components. We perform the same analysis on the data windows preceding all 17 DO events using four different trend models.

- No trend: the data are explained using the time-dependent AR(1) noise component ε_t and an intercept β_0 only,

$$y_t \sim \beta_0 + \varepsilon_t. \quad (\text{C1})$$

We only expect this to provide accurate results for stationary data windows. The results in this paper can be recreated using the `INLA.ews` package. Let y denote the $\delta^{18}\text{O}$ ratios and `time` denote the GICC05 chronology, and then the model can be fitted by prompting the following.

```
results = inla.ews(data=y, timesteps=time, formula = y ~ 1)
```

To omit the intercept term set the formula argument to `formula = y ~ -1` instead. The `rgeneric` model component corresponding to the time-dependent AR(1) noise is added automatically. To improve numerical convergence, we perform the analysis in iterations, restarting from the previously found optima with reduced step lengths. This can be specified using the `stepsize` argument in the `inla.ews` function. The length of this argument corresponds to the number of iterations. Here we used `stepsizes = c(0.01, 0.005, 0.001)`.

- Linear trend: we incorporate an additional linear effect β_1 in the model,

$$y_t \sim \beta_0 + \beta_1 t + \varepsilon_t. \quad (\text{C2})$$

This can capture linear increases but will not be able to model any nonlinearity in the data. This model can be fitted using the following.

```
results = inla.ews(data=data.frame(y=y, trend1=time_norm),
  timesteps=time, formula = y ~ 1 + trend1)
```

Here `trend1 = time_norm` is the covariate corresponding to the normalized time steps.

```
time_norm = (time-time[1])/(time[n]-time[1])
```

- Second-order polynomial: we add another effect β_2 which allows nonlinearity to be described using a second-order polynomial trend,

$$y_t \sim \beta_0 + \beta_1 t + \beta_2 t^2 + \varepsilon_t. \quad (\text{C3})$$

This model can be fitted using the following.

```
results = inla.ews(data=data.frame(y=y, trend1=time_norm, trend2=time_norm^2),
  timesteps=time, formula = y ~ 1 + trend1 + trend2)
```

Here `trend2` specifies a linear response to the covariates defined as the square of the normalized GICC05 chronology `trend2=time_norm**2`.

- Continuous second-order random walk (CRW2): we use a random effect $f(t)$ described by a continuous second-order random walk to describe the trend,

$$y_t \sim f(t) + \varepsilon_t. \quad (\text{C4})$$

This is a continuous extension (Lindgren and Rue, 2008) of a stochastic spline model, which assumes that the second-order increments are independent Gaussian processes,

$$x_i - 2x_{i+1} + x_{i+2} \sim \mathcal{N}(0, \sigma_{\text{CRW2}}^2). \quad (\text{C5})$$

This model is able to capture more general nonlinearities compared to the second-degree polynomial trend but makes the model less interpretable. Similarly as in R-INLA, the CRW2 model is specified using the following call.

```
results = inla.ews(data=data.frame(y=y, idx=time,
  timesteps=time, formula = y ~ 1 + f(idx, model="crw2"))
```

Here `idx` specifies the time steps of the continuous RW2 trend.

In Table 2 we present the estimated posterior probability of a positive trend, $P(b > 0 | y)$, compared to the corresponding p values by Rypdal (2016) and Boers (2018). We show the fitted trends for each data interval in Fig. C1. We observe that the models tend to agree, with some exceptions where the assumed trend is unable to capture the variation of the data. Although the RW2 trend is the most flexible model it appears to exhibit irregular fluctuation for several events. The second-order polynomial trend appears to be sufficiently flexible for all events and provides a much smoother and more interpretable fit.




Figure C1. $\delta^{18}\text{O}$ proxy data from the NGRIP record (gray), with Greenland stadial phases highlighted. The posterior marginal mean (blue) and 95 % credible intervals (red) of the fitted trends are included for each event.

Appendix D: Break point in memory evolution

Early warning signals are most easily detectable shortly before a bifurcation point. If the dataset covers a much larger period, for which most of it is stationary, it could be more difficult for the time-dependent AR(1) model to detect early warning signals if they are observable only for a much smaller subset of the data. To accommodate this one could add a break point, a point in time where the lag-one autocorrelation transitions from constant to linearly increasing.

Let t_{bp} denote a break point; the lag-one autocorrelation parameter is then defined by

$$\phi(t) = \begin{cases} a, & t \leq t_{bp} \\ a + \frac{b}{1-t_{bp}}(t - t_{bp}), & t > t_{bp} \end{cases} \quad (\text{D1})$$

For stability, we constrain the break point parameter using the parameterization $\theta_{bp} = 1/(1 + \exp(-t_{bp}))$ such that $t_{bp} \in (0, 1)$. This model is demonstrated by fitting it to simulated data where $t_{bp} = 0.5$. The results are presented visually in Fig. D1.




Figure D1. Panel (a) shows simulated data (black) with a break point located at $t = 0.5$. The shaded gray area describes the standard deviation derived from using the true values. Panel (b) shows the posterior marginal means (blue) and 95 % credible intervals (red), with the true memory evolution (black).

Code and data availability. NGRIP $\delta^{18}\text{O}$ data (North Greenland Ice Core Project members, 2004; Gkinis et al., 2014) and the GICC05 chronology (Vinther et al., 2006; Rasmussen et al., 2006; Andersen et al., 2006; Svensson et al., 2008) used in this paper are available at <https://www.iceandclimate.nbi.ku.dk/data/> (Niels Bohr Institute, 2025). The code to reproduce our results is available in the `inst/include/tests` folder at <https://doi.org/10.5281/zenodo.15241983> (Myrvoll-Nilsen, 2025) or in the `INLA.ews` R package available at <http://github.com/eirikmn/INLA.ews> (last access: 11 April 2025).

Author contributions. All authors conceived and designed the study. EMN adopted the model for a Bayesian framework and wrote the code. LH and EMN carried out the examples and analysis. All authors discussed the results and drew conclusions. EMN and LH wrote the paper with input from MR.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

Acknowledgements. This project has received funding from the European Union's Horizon 2020 research and innovation program (TiPES, grant no. 820970). EMN has also received funding from the Norwegian Research Council (IKTPLUSS-IKT og digital inno-

vasjon, project no. 332901). We would like to thank Niklas Boers for helping us reproduce the results of Boers (2018), including providing code to obtain the interpolated 5-year sampled NGRIP/GICC05 dataset.

Financial support. This research has been supported by the European Horizon 2020 program (grant no. 820970) and the Norges Forskningsråd (grant no. 332901).

Review statement. This paper was edited by Jonathan Donges and reviewed by Chris Boulton and two anonymous referees.

References

- Andersen, K. K., Azuma, N., Barnola, J. M., Bigler, M., Biscaye, P., Caillon, N., Chappellaz, J., Clausen, H. B., Dahl-Jensen, D., Fischer, H., Flückiger, J., Fritzsche, D., Fujii, Y., Goto-Azuma, K., Grönvold, K., Gundestrup, N. S., Hansson, M., Huber, C., Hvidberg, C. S., Johnsen, S. J., Jonsell, U., Jouzel, J., Kipfstuhl, S., Landais, A., Leuenberger, M., Lorrain, R., Masson-Delmotte, V., Miller, H., Motoyama, H., Narita, H., Popp, T., Rasmussen, S. O., Raynaud, D., Röthlisberger, R., Ruth, U., Samyn, D., Schwander, J., Shoji, H., Siggard-Andersen, M. L., Steffensen, J. P., Stocker, T., Sveinbjörnsdóttir, A. E., Svensson, A., Takata, M., Tison, J. L., Thorsteinsson, T., Watanabe, O., Wilhelms, F., and White, J. W. C.: High-resolution record of Northern Hemisphere climate extending into the last interglacial period, *Nature*, 431, 147–151, <https://doi.org/10.1038/nature02805>, 2004.
- Andersen, K. K., Svensson, A., Johnsen, S. J., Rasmussen, S. O., Bigler, M., Röthlisberger, R., Ruth, U., Siggard-Andersen, M.-L., Steffensen, J. P., Dahl-Jensen, D., Vinther, B. M., and Clausen, H. B.: The Greenland ice core chronology 2005, 15–

- 42 ka. Part 1: constructing the time scale, *Quaternary Sci. Rev.*, 25, 3246–3257, 2006.
- Boers, N.: Early-warning signals for Dansgaard-Oeschger events in a high-resolution ice core record, *Nat. Commun.*, 9, 1–8, 2018.
- Boers, N.: Observation-based early-warning signals for a collapse of the Atlantic Meridional Overturning Circulation, *Nat. Clim. Change*, 11, 680–688, 2021.
- Boers, N. and Rypdal, M.: Critical slowing down suggests that the western Greenland Ice Sheet is close to a tipping point, *P. Natl. Acad. Sci. USA*, 118, e2024192118, <https://doi.org/10.1073/pnas.2024192118>, 2021.
- Bond, G. C., Showers, W., Elliot, M., Evans, M., Lotti, R., Hajdas, I., Bonani, G., and Johnson, S.: The North Atlantic's 1–2 kyr climate rhythm: relation to Heinrich events, Dansgaard/Oeschger cycles and the Little Ice Age, *Geophysical Monograph-American Geophysical Union*, 112, 35–58, 1999.
- Boulton, C. A., Lenton, T. M., and Boers, N.: Pronounced loss of Amazon rainforest resilience since the early 2000s, *Nat. Clim. Change*, 12, 271–278, 2022.
- Braun, H., Christl, M., Rahmstorf, S., Ganopolski, A., Mangini, A., Kubatzki, C., Roth, K., and Kromer, B.: Possible solar origin of the 1,470-year glacial climate cycle demonstrated in a coupled model, *Nature*, 438, 208–211, 2005.
- Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V., and Held, H.: Slowing down as an early warning signal for abrupt climate change, *P. Natl. Acad. Sci. USA*, 105, 14308–14312, 2008.
- Dansgaard, W., Johnsen, S., Clausen, H., Dahl-Jensen, D., Gundestrup, N., Hammer, C., and Oeschger, H.: North Atlantic Climatic Oscillations Revealed by Deep Greenland Ice Cores, 288–298, *American Geophysical Union (AGU)*, ISBN 9781118666036, <https://doi.org/10.1029/GM029p0288>, 1984.
- Dansgaard, W., Johnsen, S. J., Clausen, H. B., Dahl-Jensen, D., Gundestrup, N. S., Hammer, C. U., Hvidberg, C. S., Steffensen, J. P., Sveinbjörnsdóttir, A., Jouzel, J., and Bond, G.: Evidence for general instability of past climate from a 250-kyr ice-core record, *Nature*, 364, 218–220, <https://doi.org/10.1038/364218a0>, 1993.
- Ditlevsen, P. D. and Johnsen, S. J.: Tipping points: Early warning and wishful thinking, *Geophys. Res. Lett.*, 37, L19703, <https://doi.org/10.1029/2010GL044486>, 2010.
- Ditlevsen, P. D., Andersen, K. K., and Svensson, A.: The DO-climate events are probably noise induced: statistical investigation of the claimed 1470 years cycle, *Clim. Past*, 3, 129–134, <https://doi.org/10.5194/cp-3-129-2007>, 2007.
- Gardiner, C. W.: *Handbook of Stochastic Methods*, Springer Verlag, Berlin, 4th edn., ISBN 978-3-540-70712-7, 2009.
- Gkinis, V., Simonsen, S. B., Buchardt, S. L., White, J., and Vinther, B. M.: Water isotope diffusion rates from the North-GRIP ice core for the last 16,000 years—Glaciological and paleoclimatic implications, *Earth Planet. Sc. Lett.*, 405, 132–141, <https://doi.org/10.1016/j.epsl.2014.08.022>, 2014.
- Held, H. and Kleinen, T.: Detection of climate system bifurcations by degenerate fingerprinting, *Geophys. Res. Lett.*, 31, L23207, <https://doi.org/10.1029/2004GL020972>, 2004.
- Henry, L., McManus, J., Curry, W., Roberts, N., Piotrowski, A., and Keigwin, L.: North Atlantic ocean circulation and abrupt climate change during the last glaciation, *Science*, 353, 470–474, 2016.
- Hubert, M. and Vandervieren, E.: An adjusted boxplot for skewed distributions, *Comput. Stat. Data An.*, 52, 5186–5201, <https://doi.org/10.1016/j.csda.2007.11.008>, 2008.
- Johnsen, S. J., Clausen, H. B., Dansgaard, W., Fuhrer, K., Gundestrup, N., Hammer, C. U., Iversen, P., Jouzel, J., Stauffer, B., and Steffensen, J. P.: Irregular glacial interstadials recorded in a new Greenland ice core, *Nature*, 359, 311–313, <https://doi.org/10.1038/359311a0>, 1992.
- Johnsen, S. J., Dahl-Jensen, D., Gundestrup, N., Steffensen, J. P., Clausen, H. B., Miller, H., Masson-Delmotte, V., Sveinbjörnsdóttir, A. E., and White, J.: Oxygen isotope and palaeotemperature records from six Greenland ice-core stations: Camp Century, Dye-3, GRIP, GISP2, Renland and NorthGRIP, *J. Quaternary Sci.*, 16, 299–307, <https://doi.org/10.1002/jqs.622>, 2001.
- Lenton, T., Livina, V., Dakos, V., Van Nes, E., and Scheffer, M.: Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness, *Philos. T. Roy. Soc. A*, 370, 1185–1204, 2012.
- Lindgren, F. and Rue, H.: On the Second-Order Random Walk Model for Irregular Locations, *Scand. J. Stat.*, 35, 691–700, <https://doi.org/10.1111/j.1467-9469.2008.00610.x>, 2008.
- Lynch-Stieglitz, J.: The Atlantic meridional overturning circulation and abrupt climate change, *Annu. Rev. Marine Sci.*, 9, 83–104, 2017.
- Mandelbrot, B. B. and Van Ness, J. W.: Fractional Brownian motions, fractional noises and applications, *SIAM Review*, 10, 422–437, 1968.
- Menviel, L., Timmermann, A., Friedrich, T., and England, M. H.: Hindcasting the continuum of Dansgaard–Oeschger variability: mechanisms, patterns and timing, *Clim. Past*, 10, 63–77, <https://doi.org/10.5194/cp-10-63-2014>, 2014.
- Menviel, L. C., Skinner, L. C., Tarasov, L., and Tzedakis, P. C.: An ice–climate oscillatory framework for Dansgaard–Oeschger cycles, *Nat. Rev. Earth Environ.*, 1, 677–693, 2020.
- Myrvoll-Nilsen, E.: INLA.ews. In *Earth System Dynamics*, Zenodo [code], <https://doi.org/10.5281/zenodo.15241983>, 2025.
- Myrvoll-Nilsen, E., Sørbye, S. H., Fredriksen, H.-B., Rue, H., and Rypdal, M.: Statistical estimation of global surface temperature response to forcing under the assumption of temporal scaling, *Earth Syst. Dynam.*, 11, 329–345, <https://doi.org/10.5194/esd-11-329-2020>, 2020.
- Niels Bohr Institute: NGRIP and GICC05, Data, icesamples and software, Niels Bohr Institute [data set], <https://www.iceandclimate.nbi.ku.dk/data/>, last access: 11 April 2025.
- North Greenland Ice Core Project members: High-resolution record of Northern Hemisphere climate extending into the last interglacial period, *Nature*, 431, 147–151, <https://doi.org/10.1038/nature02805>, 2004.
- Rasmussen, S. O., Andersen, K. K., Svensson, A., Steffensen, J. P., Vinther, B. M., Clausen, H. B., Siggaard-Andersen, M.-L., Johnsen, S. J., Larsen, L. B., Dahl-Jensen, D., Bigler, M., Röthlisberger, R., Fischer, H., Goto-Azuma, K., Hansson, M. E., and Ruth, U.: A new Greenland ice core chronology for the last glacial termination, *J. Geophys. Res.-Atmos.*, 111, D06102, <https://doi.org/10.1029/2005JD006079>, 2006.
- Rasmussen, S. O., Bigler, M., Blockley, S. P., Blunier, T., Buchardt, S. L., Clausen, H. B., Cvijanovic, I., Dahl-Jensen, D., Johnsen, S. J., Fischer, H., Gkinis, V., Guillevic, M., Hoek, W. Z., Lowe, J. J., Pedro, J. B., Popp, T., Seierstad, I. K., Steffensen, J. P.,

- Svensson, A. M., Vallenga, P., Vinther, B. M., Walker, M. J. C., Wheatley, J. J., and Winstrup, M.: A stratigraphic framework for abrupt climatic changes during the Last Glacial period based on three synchronized Greenland ice-core records: refining and extending the INTIMATE event stratigraphy, *Quaternary Sci. Rev.*, 106, 14–28, 2014.
- Ritchie, P. and Sieber, J.: Early-warning indicators for rate-induced tipping, *Chaos*, 26, 093116, <https://doi.org/10.1063/1.4963012>, 2016.
- Robert, C. P. and Casella, G.: *Monte Carlo Statistical Methods*, 2nd edn., Springer, New York, <https://doi.org/10.1007/978-1-4757-4145-2>, 2004.
- Rue, H. and Held, L.: *Gaussian Markov Random Fields: Theory And Applications* (Monographs on Statistics and Applied Probability, Chapman and Hall-CRC Press, London, ISBN 1584884320, 2005.
- Rue, H., Martino, S., and Chopin, N.: Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion), *J. R. Stat. Soc. Ser. B*, 71, 319–392, 2009.
- Ruth, U., Wagenbach, D., Steffensen, J. P., and Bigler, M.: Continuous record of microparticle concentration and size distribution in the central Greenland NGRIP ice core during the last glacial period, *J. Geophys. Res.-Atmos.*, 108, 4098, <https://doi.org/10.1029/2002JD002376>, 2003.
- Rypdal, M.: Early-warning signals for the onsets of Greenland interstadials and the younger Dryas–Preboreal transition, *J. Climate*, 29, 4047–4056, <https://doi.org/10.1175/JCLI-D-15-0828.1>, 2016.
- Ryvkin, J.: Fractional Brownian Motion with Variable Hurst Parameter: Definition and Properties, *J. Theor. Probab.*, 28, 866–891, <https://doi.org/10.1007/s10959-013-0502-3>, 2015.
- Schulz, M.: On the 1470-year pacing of Dansgaard-Oeschger warm events, *Paleoceanography*, 17, 4–1, 2002.
- Svensson, A., Andersen, K. K., Bigler, M., Clausen, H. B., Dahl-Jensen, D., Davies, S. M., Johnsen, S. J., Muscheler, R., Parrenin, F., Rasmussen, S. O., Röthlisberger, R., Seierstad, I., Steffensen, J. P., and Vinther, B. M.: A 60 000 year Greenland stratigraphic ice core chronology, *Clim. Past*, 4, 47–57, <https://doi.org/10.5194/cp-4-47-2008>, 2008.
- Tierney, L. and Kadane, J. B.: Accurate approximations for posterior moments and marginal densities, *J. Am. Stat. A.*, 81, 82–86, 1986.
- Vinther, B. M., Clausen, H. B., Johnsen, S. J., Rasmussen, S. O., Andersen, K. K., Buchardt, S. L., Dahl-Jensen, D., Seierstad, I. K., Siggaard-Andersen, M.-L., Steffensen, J. P., Svensson, A., Olsen, J., and Heinemeier, J.: A synchronized dating of three Greenland ice cores throughout the Holocene, *J. Geophys. Res.-Atmos.*, 111, D13102, <https://doi.org/10.1029/2005JD006921>, 2006.
- Wiesenfeld, K.: Noisy precursors of nonlinear instabilities, *J. Stat. Phys.*, 38, 1071–1097, 1985.