

StratoBayes: A Bayesian method for automated stratigraphic correlation and age modelling

Kilian Eichenseer¹, Matthias Sinnesael^{1,2}, Martin R. Smith¹ and Andrew R. Millard³

¹Department of Earth Sciences, Durham University, South Road, DH1 3LE, Durham, United Kingdom

²Geology, School of Natural Sciences, Trinity College Dublin, Dublin 2, Ireland

³Department of Archaeology, Durham University, South Road, DH1 3LE, Durham, United Kingdom

Corresponding author: Kilian Eichenseer - kilian.eichenseer@gmail.com

Abstract

Stratigraphic correlation and age modelling are fundamental to reconstructing Earth's history, biological evolution, and palaeoclimate, and underpin the exploration for subsurface resources. Correlations are produced by integrating diverse stratigraphic data across multiple sites, typically by visual inspection. Here, we introduce 'StratoBayes', a Bayesian statistical framework that combines stratigraphic correlation and depositional age estimation of stratigraphic horizons, i.e. age modelling. Our method aligns quantitative signals from two or more sites by shifting and scaling, allowing for sedimentation rate changes between stratigraphic partitions. The likelihood of an alignment is evaluated by how well the adjusted signals conform to a shared smooth trend, represented by a cubic spline. Tie points or independent age constraints, such as radiometric dates or biostratigraphic markers, can be integrated within this framework, providing age estimates for all sites. Our approach identifies multiple alignments where distinct alternatives exist, estimates their relative probabilities, and quantifies the uncertainty associated with correlations and age estimates. We apply StratoBayes to a lower Cambrian dataset comprising a combination of $\delta^{13}\text{C}$ records, radiometric dates and astrochronology from four sites in Morocco and Siberia. The results demonstrate its capacity to quantify existing alignments, and provide the first precise age estimate for the evolutionary appearance of trilobites in Siberia, one of the hallmarks of the Cambrian Explosion. Beyond this application, StratoBayes offers a generalisable framework for probabilistic stratigraphic correlation, with potential to improve age models across a range of proxy records and time intervals.

1 Introduction

Stratigraphic correlation works on the basis that rocks that were deposited under similar conditions or at the same time tend to share characteristics that allow for their attribution to a stratigraphic or temporal horizon. For example, insofar as temporal changes in the global $\delta^{13}\text{C}$ composition of seawater are reflected in marine sedimentary rocks, matching trends of changing $\delta^{13}\text{C}$ in rock sections from different locations can be used to place those sections on a relative time scale (Cramer and Jarvis, 2020; Saltzman et al., 2012). Quantitative signals such as isotopic compositions, elemental concentrations or geophysical well-log data present a particular challenge: in aligning those signals by eye, the stratigrapher has to make a large number of intuitive decisions about which peaks and troughs are likely to line up. Trying to integrate all the stratigraphic evidence from multiple sites often results in more than one potential alignment solution and differing interpretations between different workers (Bowyer et al., 2022, 2023; Landing and Kruse, 2017; Smith et al., 2016).

Computer algorithms have been designed to address the problems arising from visual correlation (Agterberg, 1990; Lisiecki and Lisiecki, 2002; Rudman and Lankston, 1973). Algorithms designed for aligning quantitative signals from two or more sites typically use a point-based approach, aligning each point of site A with zero, one or multiple points from site B. This approach proposes variable sedimentation rates between points. This flexibility in principle allows the most precise alignments, though potentially at the cost of overfitting. Point-based algorithms commonly use dynamic time warping (DTW), a technique that finds the optimal match between two time-series data by adjusting their alignment (Sakoe and Chiba, 1978). For a selection of recent approaches using dynamic time warping for stratigraphic alignment, see Wheeler and Hale (2014); Hay et al. (2019); Baville et al. (2022); Sylvester (2023); and Hagen et al. (2024). The limitations of DTW-based approaches are that they commonly require known section tops and bottoms (Sylvester, 2023); and they are generally deterministic, providing only a single solution without any indication of uncertainty or alternative alignments (but see Al Ibrahim, 2022; Hay et al., 2019). The integration of additional stratigraphic information besides the quantitative signals tends to be difficult, requiring extra steps outside of the core DTW-algorithm (e.g. Hagen and Creveling, 2024).

Probabilistic approaches overcome some of these limitations by estimating the probabilities of different outcomes, rather than producing deterministic predictions. An effective probabilistic approach is offered by the Bayesian framework, which integrates multiple sources of uncertainty by combining prior knowledge, encapsulated mathematically as a prior probability distribution, with a custom likelihood function that is used to evaluate the likelihood of observed data. Given an appropriate prior and likelihood

function it is straightforward to integrate different types of stratigraphic information. Bayesian approaches are commonly employed in age-depth models that interpolate between absolute age constraints or tie points; examples include *Behron* (Haslett and Parnell, 2008) and *Oxcal* (Ramsey, 1995). This approach can be extended by incorporating prior expectations on hiatuses, sedimentation rates, and rate variability, including external information such as astrochronological data (e.g. Blaauw and Christen, 2011; Trayler et al., 2024).

Recent Bayesian methods have attempted to combine stratigraphic correlation and age modelling. Lee et al. (2022) have implemented a Bayesian method that uses Gaussian process regression to match Cenozoic oxygen isotope data from one site to an oxygen isotope stack, while simultaneously integrating age estimates from radiocarbon dates to produce probabilistic age-depth models (i.e. the BIGMACS model). This method improves upon earlier approaches by specifying uncertainty for tie points and integrates prior knowledge on Cenozoic sedimentation rates with absolute age information from the aligned site. However, age uncertainties from the reference site are not included, and varying sampling resolution or large sedimentation rate changes may violate model assumptions and impede the broader adoption of this method in its current form (Middleton et al., 2024). Edmonson and Dyer (2024) have developed a different Bayesian method based on Gaussian process regression that works without prior knowledge of sedimentation rates, but requires minimum and maximum age estimates for all sections, and the absence of an explicit prior on sedimentation rates may risk overfitting. Here, we introduce a versatile Bayesian method for stratigraphic correlation and age modelling that can align quantitative signals from two or more sites without the need to specify tie points or top and bottom ages, and with no restrictions on sampling frequencies. Possible sedimentation rates can be specified by the user as priors, and the likelihood encompasses the alignment of the signals and, optionally, additional age constraints such as dated horizons. The method requires only vague prior knowledge on the ages and on the degree of overlap of the sections, along with order-of-magnitude estimates of sedimentation rates; it is not necessary to specify matching section tops or bottoms. The model is able to integrate radiometric dates from different sites, meaning that ages from well-dated sites can inform age estimates at sites with little or no age information. Age estimates with uncertainty can thus be obtained for any point within any site, and alternative alignments can be identified. Additional stratigraphic knowledge, such as hiatuses or tie points, can be readily incorporated.

Our Bayesian model works by evaluating the fit of a single cubic spline (Heaton et al., 2020) to the combined quantitative signal of all sites. If more than one type of signal is included, e.g. $\delta^{13}\text{C}$ and $\delta^{18}\text{O}$, a different spline is constructed for each signal type, and their joint likelihood is used to evaluate the alignment. Different alignments are generated by shifting the sites relative to each other, and by scaling segments of the sites using different sedimentation rates. Markov chain Monte Carlo methods are used to

obtain the posterior distributions of the unknown model parameters. Our method is implemented as an R package, ‘StratoBayes’.

To demonstrate the potential of this method, we apply it to artificial stratigraphic data and to a real case study using lower Cambrian $\delta^{13}\text{C}$ records from Morocco (Magaritz et al., 1991; Maloof et al., 2005, 2010; Tucker, 1986) and Siberia (Kouchinsky et al., 2007). Integrating radiometric dates (Landing et al., 1998, 2021; Maloof et al., 2010), we provide age estimates for the studied sections of an interval spanning several lower Cambrian carbon isotope excursions, and compare our algorithm-derived correlation with recent stratigraphic models relying on visual expert-based interpretations (Bowyer et al., 2022, 2023). Our solution also provides a fully quantitative age estimate for the appearance of the first Siberian trilobites, which are thought to be the world’s oldest trilobites (Landing et al., 2021).

2 Bayesian stratigraphic model

StratoBayes generates and evaluates alignments of quantitative stratigraphic signals. A signal consists of, for example, geochemical or geophysical measurements that vary across height or depth (Fig. 1a), obtained from a contiguous sedimentary sequence, which may be interrupted by hiatuses at known horizons. Alignments are generated by shifting the sites containing the signals either (a) against a fixed reference site, or (b) against each other on an absolute age scale. Additionally, the sites are scaled (“stretched” or “squeezed”) assuming different sedimentation rates. The fit of different alignments, corresponding to different shifts and sedimentation rates, is evaluated in the Bayesian framework.

Statistical analysis in the Bayesian framework starts by formulating a probabilistic model that includes known data y and unknown model parameters θ . Instead of trying to identify a single estimate for θ , Bayesian inference involves estimating probability distributions for the model parameters, termed “posterior probability distributions”. Posterior distributions are obtained by combining prior knowledge of the parameters with the data via a likelihood function. Bayes’ theorem states that the probability of the parameters given the data, $p(\theta|y)$, i.e. the posterior probability, is proportional to the probability of the data given the model parameters (i.e. the likelihood), $p(y|\theta)$, times the prior probability of the model parameters, $p(\theta)$:

$$p(\theta|y) \propto p(y|\theta)p(\theta) \quad (1)$$

In our case, this approach requires specifying prior probability distributions for the unknown model parameters that control the shifting and scaling (Fig. 1b), and optionally for the duration of pre-determined hiatuses. Our model assumes that the measurements in each sedimentary sequence are samples (with noise)

from a common underlying signal, whose value can be modelled by a smooth curve described by a cubic B-spline. Our likelihood function quantifies how well a cubic B-spline fitted to a given alignment explains the observed data (Fig. 1c). Additional likelihood components can integrate absolute age constraints such as radiometric dates or other tie points, e.g. index fossils. Using Bayes' theorem, the priors are combined with the likelihood to obtain the posterior probability for any alignment.

We obtain probability distributions for the parameters of the model by running a Markov chain Monte Carlo (MCMC) simulation. This involves repeatedly generating parameter values over a large number of iterations. To ensure thorough exploration of the parameter space, we employ parallel tempering, i.e. we run multiple chains in parallel, flattening the likelihood of the tempered (hot) chains, which can therefore move between different posterior modes; swaps between chains are proposed at every iteration. For the posterior estimates, we retain samples only from the primary (cold) chain. An initial portion of the samples is discarded (burn-in) to remove dependency on starting values, and only every n^{th} iteration is recorded to reduce autocorrelation. Details on the MCMC implementation are provided in Appendix A.

In the following, we will assume that measurements were taken on a height scale (increasing from the bottom to the top), but depth-scale measurements can be used interchangeably by inverting their sign.




Figure 1: Schematic of the alignment algorithm. a) Input data: Quantitative stratigraphic measurements (e.g. geochemical data) from two sites recorded along their section height (here given in meters). b) Priors must be placed on the shift parameter α and on the relative sedimentation rate v . Here, α determines the reference height (at Site 1) corresponding to the top of the height range of Site 2, and v corresponds to the sedimentation rate of Site 2 relative to Site 1. The vertical, dashed lines denote the α

and ν values, 12.3 m and 3.0, respectively, that were used in the creation of the data of site 2. c) An alignment corresponding to a single sample from the posterior. The blue dashed line indicates the position of the top of the data from Site 2 at the reference height scale (α ; median: 12.5 m). The relative sedimentation rate ν has been estimated at a median of 2.8, corresponding to a shortening of the dataset from Site 2 relative to the reference site (indicated by the dashed and solid light brown line). Note that the posterior estimates of α and ν are similar, although not identical to the values used in creating the data (see Sect. 3). The curved grey line shows the cubic B-spline corresponding to the alignment.

2.1 Evaluating alignments with a cubic B-spline

Identifying good alignment positions requires evaluating and comparing different potential alignments. In the Bayesian framework, the measure used for this evaluation is the likelihood. We derive the likelihood of an alignment from its fit to a single cubic B-spline (Eilers and Marx, 1996), fitted to the measurements from all sites, including the reference site (see Fig. 1c).

We model each measured value y_i as normally distributed:

$$y_i \sim \text{Normal}(\mu_i, \sigma), \quad (2)$$

where μ_i is the mean, and the standard deviation σ represents the scatter around the spline. μ_i is given by the spline function

$$\mu_i = \sum_{j=1}^{k+2} \beta_j B_j(h_i) \quad (3)$$

Here, μ can be interpreted as an underlying common signal of which the observations from each site, including the reference site, are noisy realisations. k denotes the number of internal knots of the spline, with more knots implying that the spline can potentially capture higher-frequency variations. β_j is the spline coefficient associated with the j -th basis function, and $B_j(h_i)$ is the j -th B-spline basis function evaluated at a reference height h_i . A roughness penalty controlled by a smoothing parameter λ is incorporated in the prior on $\boldsymbol{\beta}$, such that higher values of λ serve to favour smoother splines (Appendix A). The number of knots and the roughness penalty each influence spline flexibility in different ways: increasing k provides a finer resolution for fitting local features, whereas increasing λ penalizes abrupt changes and yields smoother fits. The knots for the spline can be distributed across the reference height range that the converted measurement heights occupy for a specific combination of shift parameters (α) and scale parameters (ν ,

i.e. relative sedimentation rates). Our current model implementation uses evenly spaced knots, but knot placement could also follow, for example, the density of measurements. Alternatively, the knots can be fixed at specific heights on the reference scale, in which case combinations of α and ν that result in converted measurement heights falling outside the knot range cannot be evaluated.

The likelihood of an alignment, given β , σ and λ , is determined by the residual deviations of the y_i values from the corresponding μ_i values. The overall likelihood for n data points is obtained by taking the product over all individual likelihoods for each pair of y_i and μ_i :

$$L(\mathbf{y}|\beta, \sigma, \lambda) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} \times e^{\left(-\frac{(y_i-\mu_i)^2}{2\sigma^2}\right)} \quad (4)$$

We thus assume that the deviations of the data from the spline are independently and identically distributed according to a normal distribution with mean 0 and standard deviation σ .

Our model allows for using more than one type of measurement simultaneously. In this case, a separate spline is fitted to all data, from all sites, for each type of measurement. The product of all likelihoods from all measurement types gives the overall likelihood.

2.2 Alignment and partitioning

In order to generate alignments of stratigraphic signals from different sites, one site is picked as a fixed reference site. The other sites are shifted and stretched (or squeezed) relative to the fixed reference site r . This requires specifying a shift parameter (height) α_s , which anchors an arbitrary, specified height of site s to a height in the reference site r . Here, we anchor the top of site s , so we set $\alpha_s = \alpha_{top,s}$ meaning $\alpha_{top,s}$ will be the height at site r that aligns with the top of site s . To stretch or squeeze site s , a relative sedimentation rate ν_s can be specified, where ν_s is defined relative to the reference site. For any height $h_{x,s}$ at site s , the corresponding height in the reference site r can then be calculated as

$$h_r = \alpha_{top,s} - \frac{1}{\nu_s} \times (h_{top,s} - h_{x,s}), \quad (5)$$

where $h_{top,s}$ is the height of the top of site s . Although we here chose the top of site r as the reference horizon α for simplicity, any horizon at site r can be used as α . A $\nu_s < 1$ implies that site s has a lower sedimentation rate than site r , and consequently, s has to be stretched to match r . A $\nu_s > 1$, i.e. a higher sedimentation rate at site s will lead to s being squeezed to match r .

The model described here is simple in that the same v is applied to all measurements of the same site. In this scenario, any site may be used as the reference site. Below, we introduce more complex models with more than one sedimentation rate per site, and with hiatuses. With these models, it is practical to select the site with the most sedimentation rate changes and hiatuses as the reference site. This reduces the number of unknown parameters in the model, making it easier to obtain a representative sample from the posterior.

2.2.1 Multiple sedimentation rates per site

Instead of having one sedimentation rate per site, sites can be partitioned, reflecting for example lithological units, with each partition being modelled with a distinct sedimentation rate:

$$h_r = \alpha_{top,s} - \sum_l^{n_{p,s}-1} \left(\frac{1}{v_p} \times (h_{p,s} - h_{p+1,s}) \right) - \frac{1}{v_{n_{p,s}}} \times (h_{n_{p,s}} - h_{x,s}), \quad p = 1 \dots n_{p,s} \quad (6)$$

Here, $n_{p,s}$ is the number of partitions encountered from h_{top} to $h_{x,s}$, $h_{p,s}$ is the top height of partition p at site s , and $h_{p+1,s}$ is the top height of the partition below partition p at site s . If $h_{x,s}$ falls in the first partition from the top, the calculation simplifies to the equivalent of Equation 5, with h_{p_s} , the top height of the first partition being also the top height of site s . The relative sedimentation rates of partitions, v_p , can differ for each partition in each site, or partitions in different positions within a site or across sites may share sedimentation rates.

2.2.2 Site-specific sedimentation rate multipliers

The sedimentation rate model above can be further expanded by adding an overall site-specific sedimentation rate multiplier ζ_s :

$$h_r = \alpha_{top,s} - \sum_l^{n_{p,s}-1} \left(\frac{1}{\zeta_s v_p} \times (h_{p,s} - h_{p+1,s}) \right) - \frac{1}{\zeta_s v_{n_{p,s}}} \times (h_{n_{p,s}} - h_{x,s}), \quad p = 1 \dots n_{p,s} \quad (7)$$

This may be useful in scenarios where sedimentation rates systematically differ between sites, perhaps due to varying distances from a sediment source, but where the sedimentation rate ratios of different partitions are assumed to be constant across sites.

2.2.3 Hiatuses

Known hiatuses (also referred to as unconformities or stratigraphic gaps) can be included at specific pre-defined locations in a site. Expanding Equation 5 to include gaps of height δ , we obtain

$$h_r = \alpha_{top,s} - \frac{1}{v_s} \times (h_{top,s} - h_{x,s}) - \sum_g^{n_{G_s}} \delta_g, \quad g = 1 \dots n_{G_s} \quad (8)$$

where n_{G_s} is the number of gaps encountered from $h_{top,s}$ until height $h_{x,s}$. In a correlation on an absolute age scale (Sect. 2.2.5), hiatuses would instead be expressed as durations, not heights.

2.2.4 Tie points

Tie points define specific heights within an aligned site and assign a probability distribution to indicate to which horizon these heights correspond on the reference scale. For example, a tie point might be a lithological boundary, a biostratigraphic horizon, or a radiometric date. If tie points are specified, the likelihood of an alignment is expanded to include not only the fit of the signal data to the spline, but also the positions of the ties on the reference height scale relative to the specified probability distribution.

For example, a point in an aligned section which is tied by observation to the reference section at a position m_t with a normally distributed uncertainty with standard deviation s_t that ends up being shifted to a reference height h_t (computed from the relevant α and v parameters) contributes a likelihood of

$$L(m_t|h_t, s_t) = \frac{1}{\sqrt{2\pi s_t^2}} \times e^{\left(-\frac{(m_t-h_t)^2}{2s_t^2}\right)} \quad (9)$$

to the overall likelihood of the model.

2.2.5 Age-scale alignment

Data on an (absolute) age scale can be aligned using the methods introduced above by using ages instead of heights. However, height-scale data can be aligned on an age-scale if absolute age constraints (specified as ties) are provided from at least one site. In this case, all sites will be shifted to align on a common age scale, i.e., there is no reference site.

Analogous to the heights in the reference height scale in Equation 5, ages (a) can be calculated as:

$$a = \alpha_{top,s} + \frac{1}{v_s} \times (h_{top,s} - h_{x,s}) \quad (10)$$

Here, $\alpha_{top,s}$ is the top age (minimum age), rather than top height (maximum height), of site s . Sedimentation rates v_s need to be expressed on the common age scale, rather than relative to a reference site. Equations 6–8 can be modified accordingly for an analysis on the age scale.

It should be noted that due to sedimentation rates being fixed for an entire site or within partitions, our current model implementation does not necessarily result in increasing age uncertainty away from absolute age constraints. Potential sedimentation rate changes within sites or partitions could lead to our model underestimating age uncertainty with growing stratigraphic distance from absolute age constraints (see De Vleeschouwer and Parnell, 2014).

2.3 Priors

The Bayesian framework requires priors to be placed on all unknown model parameters. In our model, these include the alignment parameters (e.g. α , ν), the smoothing parameter λ , the residual standard deviation σ (if it is not fixed), and the spline coefficients β . The priors on the alignment parameters determine the range of possible alignments and need to be chosen with care. For the other parameters, weakly informative priors with minimal influence on the analysis are preferred (Appendix A). In addition to those priors, we penalise a lack of overlap by specifying a prior probability of data points from different sites overlapping each other.

2.3.1 Alignment parameters

The priors on the alignment parameters should reflect the stratigraphic knowledge on the input data. The user may specify different types of prior distributions (e.g., normal, uniform, exponential) for the alignment parameters during model setup.

- α determines the reference site (site r) height or age that a specific position within the aligned site (site s) corresponds to. In the absence of prior knowledge on how the sites are likely to align, a uniform prior can be placed on α . For example, if α refers to the top of site s , a uniform prior on α with min and max equal to the height or age range of site r implies that the top of site s will be placed within the height range of site r .
- ν is either a relative (*height scale alignment*) or an absolute (*age scale alignment*) sedimentation rate. In our model implementation, priors are placed on the natural logarithm of ν , $\ln(\nu)$, rather than on ν directly. Specifying rate parameters on the logarithmic scale ensures that their priors are symmetric: a doubling or halving of a rate has equivalent distances on the logarithmic scale. If the sedimentation rate is relative, $\ln(\nu) < 0$ (i.e. $\nu < 1$) results in “stretching”, and $\ln(\nu) > 0$ (i.e. $\nu > 1$) results in “squeezing” of site s relative to site r . In the absence of strong prior knowledge about the relative sedimentation rate, a normal prior on $\ln(\nu)$ with a mean of 0 places equal prior probability on “stretching” or “squeezing” of site s relative to site r . The standard deviation requires at least a broad guess of the potential magnitude of sedimentation rate

differences. For example, a standard deviation of $\frac{\ln(4)}{1.96}$ places 95% of prior probability on $\frac{1}{4} < \nu < 4$ for $\ln(\nu) \sim \text{Normal}\left(0, \frac{\ln(4)}{1.96}\right)$. If ν is an absolute sedimentation rate, the range of plausible prior sedimentation rates may be estimated from the absolute age constraints.

- ζ_s is a multiplier applied to all relative or absolute sedimentation rates ν corresponding to a single site s . As with ν , $\ln(\zeta_s) < 0$ (i.e. $\zeta_s < 1$) causes additional “stretching”, and $\ln(\zeta_s) > 0$ (i.e. $\zeta_s > 1$) causes additional “squeezing” of site s .
- δ is the reference height range or duration of a hiatus. An exponential prior may be useful when little is known about the extent of the hiatus, placing higher probabilities on short extents. The rate needs to be chosen to make sense in the context of the height of the sections, or of the anticipated age range of the sites.

2.3.2 Penalising a lack of overlap

Individual splines fitted to data from each site separately can almost always follow the data more closely than a single spline fitted to aligned data from all sites. Given enough knots, alignments in which the data do not overlap, or only overlap little, will thus generally result in a higher likelihood than alignments with a partial or full overlap. This means that if the priors allow non-overlapping alignments, those will generally be preferred in the model inference. To counteract this tendency, we impose a prior on the overlap of each individual data point from all sites that penalises non-overlap with data from other sites.

The prior on overlap for data point i from site s is

$$P(i_s) = e^{(-\sqrt{S-1} + \sqrt{S_{\text{overlap},s,i}}) \times c_{\text{overlap}}}, \quad (11)$$

where S is the number of sites in the analysis, $S_{\text{overlap},s,i}$ is the number of other sites overlapping the reference height h_r or age a of point i_s , and c_{overlap} is a constant. This formulation implies that the penalty for a point i_s that overlaps all other sites is 0, and the penalty is strongest (most negative) if i_s overlaps no other sites. To work effectively, the penalty needs to be stronger for data sets with little noise (low residual σ), to offset the larger likelihood differences resulting from fitting a spline with low σ . A range of c_{overlap} values may work in practice. A formulation that we have found works well in many scenarios sets

$$c_{\text{overlap}} = c \times \frac{1}{S} \sum_{s=1}^S \left(\frac{\sigma_{y,s}}{\sigma_s} \right)^q \quad (12)$$

where c is a constant determining the strength of the overlap penalty (set to a default of $c = \frac{1}{4}$), $q = 1$ if σ is fixed, and $q = \frac{1}{2}$ if σ is variable (i.e. estimated in the model inference). Here, $\sigma_{y,s}$ is the standard deviation of all data y from site s , and σ_s is the residual standard deviation of a Bayesian spline fitted to the data y from site s , using the same priors as for the overall model inference.

3 Model illustration

We illustrate the performance of our stratigraphic alignment method with a simple, artificial dataset (Fig 2a). We generated measurements from a reference site (Site_{ref}) using a sine wave covering 3.5 periods, where each period corresponds to 2π radians. To generate the signal data, we intercepted this sine wave at heights h with 250 evenly spaced points per period, i.e. the number of data points (n) is $3.5 \times 250 = 875$. Each signal value y_i was generated with random white noise $\sigma = \frac{1}{5}$ added, such that

$$y_i \sim \text{Normal}\left(\eta_i \sin\left(h_i - \frac{1}{2}\pi\right), \sigma\right), \quad i = 1 \dots n \quad (13)$$

The factor η_i modulates the amplitude of the sine wave at each height h_i . It was set to $\eta = 1$ for the heights ranging from -0.5π to 5π , and to $\eta = 0.75$ from heights 5π to 6.5π , which reduces the amplitude beginning in the middle of the third period of the sine wave. The aligned signal was simulated as above, but from a sine wave covering one period, sampling 250 data points, again with random noise using $\sigma = \frac{1}{5}$ and $\eta = 1$. To simulate a sedimentation rate twice as high as at the reference site, we multiplied the heights of Site_{align} by 2. The heights of Site_{align} were then shifted to start at 0.

The aligned signal should thus match either the first or the second, but not the third period of the reference signal. To align the two sites, we used a simple model with a site-specific shift α , referring to the top of Site_{align} and relative sedimentation rate ν as in Equation 5. From the data generation, we know that the posterior of ν should be ≈ 2 , with $\ln(\nu) \approx 0.69$, and α (defined as the reference height corresponding to the top height of Site_{align}) should be $\approx 2\pi$ (top of first period) or $\approx 4\pi$ (top of second period).

To minimise the influence of the priors, we used a uniform prior on α that extends well beyond the alignment positions known from generating the data, and a broad normal prior on $\ln(\nu)$ that encompasses the known sedimentation rate $\nu = 2$ (Fig. 2b):

$$P(\alpha) \sim \text{Uniform}(-\pi, 8\pi) \quad (14)$$

$$P(\ln(v)) \sim \text{Normal}(0,1) \quad (15)$$

These priors place 95% of prior probability for the relative sedimentation rate of $\text{Site}_{\text{align}}$ between 0.14 and 7.1, and place the top of $\text{Site}_{\text{align}}$ anywhere from half a period below the start of the first period ($-\pi$) up to one period above the third period (8π). For the cubic spline, we specify 20 evenly spaced knots, which is more than enough to approximate the three periods of the sine wave.

We estimated the posterior of the model with three independent runs, each with 16 chains and 60,000 iterations. The first 10,000 samples were discarded as burn-in, and every 25th iteration was recorded, resulting a total of 6000 samples after burn-in across all three independent model runs.

The results show that the analysis identified both matching alignments, corresponding to the first and second period of the reference site (Fig. 2b). The posterior probability for ($\text{Site}_{\text{align}}$) matching period 1 is 50.1%, and 49.9% for matching period 2. A density plot of the posterior of α and $\ln(v)$ shows that α has a bimodal posterior, corresponding to the two alignments (Fig. 2c). The trace plots indicate good mixing of the chains (Fig. 2d), suggesting that the posterior estimates are robust.

It is notable that the model estimate for the relative sedimentation rate v is lower at 1.90 (95% credible interval: 1.82 to 1.99) than the value used for the data generation (2.00). Reported values, here and throughout, represent the posterior median, with 95% credible intervals – given in brackets – referring to the interval between the 2.5% and 97.5% points of the posterior distribution. This deviation of the posterior from the known sedimentation rate estimate arises because the priors favour greater overlap (see Sect. 2.3.2). The posterior alignment tends to “compress” the data from $\text{Site}_{\text{align}}$ slightly less than expected, leading to an increased overlap of points (see also Fig. 5b).




Figure 2: Model illustration using artificial data. a) Input data: Quantitative stratigraphic data from two sites. The blue line indicates the range in which $Site_{ref}$ was created with $\eta = 1$, and the purple line above indicates the range for which η was set to 0.75 to lower the amplitude. b) Two alignments identified by the inference, with ($Site_{align}$; blue squares) matching the first or second period of ($Site_{ref}$; red points). The alignments shown here correspond to two distinct samples from the posterior; other samples will result in slightly different positions of ($Site_{align}$). The curved dark lines show the cubic spline corresponding to each alignment. c) Posterior densities of α and $\ln(v)$. The two modes of α correspond to the two distinct alignments in b). The dotted lines indicate the v values with which

(Site_{align}) was simulated, and the two plausible α values. d) Trace plots of α and $\ln(v)$. The three distinct colours correspond to the three independent model runs. For visual clarity, only 75 selected samples are shown from each run.

4 Case study: Lower Cambrian $\delta^{13}\text{C}$ records

To demonstrate the utility of this method, we use it to align stable carbon isotope records ($\delta^{13}\text{C}$) from lower Cambrian marine shelf carbonates (Fig. 3). We integrate a combination of radiometric dates, $\delta^{13}\text{C}$ and astrochronological information from four sites to obtain age estimates for the sampled intervals from all sites, and use this age model for dating the first documented occurrence of Siberian trilobites.

4.1 Data

We selected three records from the Anti-Atlas mountains in Southern Morocco, corresponding to the Oued Sdas, the Tiout and the Talat n' Yissi sections, which were part of West-Gondwana during the early Cambrian (Magaritz et al., 1991; Maloof et al., 2005, 2010; Tucker, 1986). Oued Sdas and Tiout harbour multiple precise U-Pb radiometric ages (Landing et al., 2021; Maloof et al., 2010). Talat n' Yissi has no radiometric dates, but a radiometric date exists from the stratigraphically equivalent Lemdad syncline (Landing et al., 1998) that has been correlated biostratigraphically to Talat n'Yissi with the *Antatlasia gutta-pluviae* zone (Maloof et al., 2005); we include this date in the analysis. We will align these sites with each other, and with a $\delta^{13}\text{C}$ record from the Sukharikha section from the northwestern Siberian platform (Kouchinsky et al., 2007), corresponding to the palaeocontinent Siberia. There are no radiometric dates available for the Siberian section for this stratigraphic interval. Data that was inferred to be below the lower leg of the prominent “5p” excursion (lowest peak in Fig. 3a and d) was excluded to simplify the correlation, reducing the number of modelled sedimentation rates unconstrained by radiometric dates. This cropping of data affects the Oued Sdas and Sukharikha sections; Fig. 3 shows all data that was included in the analysis. $\delta^{13}\text{C}$ values were used as reported by the authors of the respective publications without any scaling or other adjustments.




Figure 3: Cambrian $\delta^{13}\text{C}$ data and radiometric dates from Morocco (a - c) and Siberia (d). Different colours, in conjunction with different symbols, delineate different lithological units or formations. Circles indicate the position of radiometric dates, with mean age and 2 standard deviations denoting the uncertainty. Stars denote the positions where the oldest trilobite remains are found in Morocco (a) and Siberia (d). The dashed line in (d) indicates a hiatus. α indicates the reference horizon chosen for specifying the prior on the shift parameter α for each site.

4.2 Model specification

To align the four sites on the age scale, we specify an α parameter on the absolute age scale (Ma) for each site, and use absolute, rather than relative sedimentation rates (expressed in m Myr^{-1}). We encapsulate variation in sedimentation rates (ν) by partitioning sites into members, formations or lithological units, leading to multiple sedimentation rates per site. As there are few radiometric dates to constrain sedimentation rates, partitions shared between the Moroccan sites are set to have the same relative sedimentation rate across sites. To account for potentially faster or slower sedimentation rates at different sites, a site-specific sedimentation rate multiplier ζ is added for Oued Sdas and Talat n'Yissi that is multiplied with the ν from those sites. The ν for a partition applies to all sites at which this partition occurs; for Tiout, they are used unaltered, and no ζ is needed for Sukharikha as there are no shared partitions with other sites. We partition the Moroccan data based on the lithostratigraphy from Maloof et al. (2005). We divide the Adoudounian Tifnout Member into a lower part (Tifnout l.), and an upper stromatolitic part (Tifnout stromatolite), as preliminary results suggested pronounced sedimentation variability between those parts. We subdivide the Lie de Vin Formation into three members; the Igoudine Formation is subdivided into two members. The Amouslek and Isaafen formations are not subdivided. The Sukharikha section is

divided into two formations, which we assign separate sedimentation rates. At the boundary, a substantial hiatus is evinced by the truncation of the “7p” $\delta^{13}\text{C}$ peak (Kouchinsky et al., 2007). We include the duration of this hiatus (δ) as an additional unknown parameter in the model.

The model requires priors to be specified for each of its 18 alignment parameters: Four α , eleven ν , two ζ and one δ (Fig. 4). These priors are broadly guided by the radiometric dates and by previous work (Bowyer et al., 2023; Landing et al., 2021; Sinnesael et al., 2024). The α for the Tiout and Sukharikha sites are placed at the height positions of the first trilobite fossil remains found at Tiout (Sinnesael et al., 2024), and the first appearance of Siberian trilobites correlated to Sukharikha (Landing et al., 2021; Varlamov et al., 2008). Here, we place normal distributed priors with mean age 520 Ma and a wide standard deviation of 2 Myr on the α parameters at Tiout and Sukharikha. This prior reflects the notion that first appearance dates of trilobites may be broadly similar at ≈ 520 Ma, but not necessarily identical, and the data is allowed to determine the exact age of each α . The α priors for Oued Sdas and Talat n’Yissi are placed at the position of the lowest or the only available radiometric date, respectively, consisting of normal distributions with mean age equal to the mean age estimate of the radiometric data and a wide standard deviation of 2 Myr.

For the sedimentation rates, priors informed by an astrochronology of the Tiout section (Sinnesael et al., 2024) are used for the following five stratigraphic partitions: The lower, middle and upper members of the Lie de Vin Formation, and for the lower and upper (Tiout Member) members of the Igoudine Formation. Those priors are chosen such that the 95 percentile interval of ν spans the minimum and maximum of the astrochronological sedimentation rate estimates when using an uncertainty of ± 1 short eccentricity cycle for each partition, with an estimated duration of short (≈ 100 kyr) eccentricity cycles ranging from 92.5 to 100.5 kyr (two standard deviations, following Lantink et al., 2022).

To specify priors for the remaining Moroccan partitions (lower part of Tifnout Fm., Tifnout stromatolite, Amouslek Fm., and Isaafen Fm.), sedimentation rates between the radiometric dates from Oued Sdas and Tiout are calculated using the mean ages of the dates. The prior on $\ln(\nu)$ is defined as a normal distribution with a mean of 5.39, corresponding to the mean of the empirical sedimentation rates from Oued Sdas and Tiout, calculated on the logarithmic scale. A wide standard deviation of 0.75 is set, resulting in the 95 percentile interval of ν spanning 50.3 to 951 m Myr^{-1} . This interval significantly exceeds the range of sedimentation rates inferred from the radiometric dates at Oued Sdas and Tiout, 147 to 314 m Myr^{-1} , allowing for the possibility of lower or higher sedimentation rates in some partitions.

Prior sedimentation rate estimates for the Siberian formations are estimated in the absence of radiometric dates, very broadly based on global correlations by Bowyer et al. (2023). These correlations suggest average sedimentation rates on the order of 20 to 30 m Myr^{-1} ; we place a normal prior on $\ln(\nu)$ with a mean of

3.30 and a standard deviation of 0.75, resulting in a 95 percentile interval of ν spanning 6.23 to 117.9 mMyr⁻¹, which allows for the possibility of significantly different sedimentation rates from those inferred by Bowyer et al. (2023).

Finally, a prior needs to be placed on the duration of the hiatus δ between the Sukharikha and the Krasnoporog formations. Kouchinsky et al. (2007) do not give an indication of the potential duration of this hiatus, but if the under- and overlying $\delta^{13}\text{C}$ peaks are correlated as indicated by previous work (Bowyer et al., 2022; Landing et al., 2021), a relatively short hiatus of ≈ 1 Myr is likely. To express considerable uncertainty about the duration of the hiatus, we place an exponential prior on δ with a rate of 1, which places 95% of prior probability on the duration being < 3 Myr, with 5% probability accounting for the possibility of a longer gap.

The cubic spline comprises 40 evenly spaced knots, allowing it to closely follow trends in the $\delta^{13}\text{C}$ records while keeping the MCMC runtime manageable, as a higher knot count increases computational cost. For the smoothing parameter λ , we applied a gamma prior with StratoBayes' default values of $a_\lambda = 1$ and $b_\lambda = 1000$. We fixed σ , which is the residual standard deviation of the overall spline, at 0.66, which is the average residual standard deviation of individual cubic splines fitted to each $\delta^{13}\text{C}$ record from the four respective sites. These individual splines were constructed with 40 knots evenly spaced across the height range of each respective site and fitted with Gibbs sampling using 2000 iterations, discarding 25% of samples as burn-in. The same default λ priors as described above were applied, while the prior for these splines' standard deviations was specified as a gamma prior on the precision τ , with $a_\tau = b_\tau = 0.01$ (see Appendix A for details).




Figure 4: Priors on the 18 alignment parameters for the Cambrian model. Prior probability density is shown (a) for four α parameters corresponding to one site each (priors for Tiout and Sukharikha in grey are identical), (b) for six v (sedimentation rate) parameters with little prior knowledge, (c) for five v parameters from Morocco with tight priors based on astrochronology, (d) for ζ parameters (site-specific sedimentation rate multipliers) for Oued Sdas and Talat n'Yssi (identical), and (e) for the duration of the hiatus between the Sukharikha and the Krasnoporog formations. The width of the red bar in (b) visualises the range of sedimentation rates spanned by (c). Panel (f) visualises two alignments generated by randomly drawing parameter values from their respective priors, to give an indication of the broad range of alignments that the priors on the alignment parameters allow; colours correspond to the four sites (see Fig. 6). Panels (b), (c), and (d) are depicted with a logarithmic x-axis as the priors were specified on $\ln(v)$ and $\ln(\zeta)$.

4.3 Parameter estimation

This model is more complex than our earlier examples, and hence requires longer runs with more chains. We conducted four independent model runs, each with 750,000 iterations and 24 chains. The runs were executed in parallel using four workers on a desktop computer (Intel i7-10700 CPU, 8 cores / 16 threads, 40 GB RAM) and completed within 5 days. The first 150,000 iterations were discarded as burn-in. From

the remaining 600,000, every 50th iteration was retained, resulting in 12,000 samples per run and 48,000 samples in total.

Inspection of trace plots of the model runs indicates stationarity and good mixing of the chains with the exception of infrequent visits of secondary posterior modes (Appendix B, Fig. B1). The potential scale reduction factor (using eq.4 in Vats and Knudson, 2021) is between 1.00 and 1.05 for all alignment parameters, suggesting approximate convergence of the MCMC. The multivariate effective sample size (Vats et al., 2019) of the 48,000 samples is 4161.

4.4 Results

To identify distinctly different alignments in the posterior, a hierarchical density-based cluster analysis (Campello et al., 2015) was conducted using the inferred ages of all partition boundaries of the four sites (Fig. 4a,b). We specified 1% of samples (480) as the minimum number of points per cluster, resulting in three distinct clusters with 93%, 2.8% and 2.6% of posterior samples, respectively, and 1.5% of samples not being assigned to any cluster. These alignment clusters also differ in the prior probabilities and likelihoods associated with individual posterior samples. On average, samples from alignment 1 tend to exhibit a lower degree of overlap, but a higher likelihood (Fig. 4c), indicating a better fit to the data.




Figure 5: (a, b) 2D density plots of the inferred top ages of the four sites, representing some of the ages used for obtaining alignment clusters from posterior samples. (c) 2D density plot of the $\ln$ prior probability of overlap against the overall $\ln$ likelihood. Areas with more opaque shadings correspond to a higher density of individual posterior samples. Colours correspond to alignment clusters: alignment 1 - violet; alignment 2 - blue; alignment 3 - green; outlier samples not assigned to any cluster - yellow.

Using samples from the posterior of the model parameters, alignments can be generated. Fig. 6 visualises three alignments drawn from the three alignment clusters identified in the posterior. For each alignment cluster, the iteration with partition boundary ages that are, on average, closest to the median ages of the partition boundaries within that cluster is selected for displaying. All three alignments exhibit a good match between the long-term trends of the $\delta^{13}\text{C}$ curves from the four sites and the common spline curve, although many shorter-term deviations are visible (Fig. 6a-c). The spline curve notably follows the more densely sampled sites (Oued Sdas, Talat n'Yissi) more so than the thinly sampled sites (Tiout, Sukharikha), resulting in greater deviations of the latter two sites.

The posterior age estimates for the stratigraphic positions of the radiometric dates broadly match the age estimates that were used as inputs in the analysis (Fig. 6d). The deviations are greatest for the Talat n'Yissi date (Ta_1), which has large uncertainty and therefore less influence on the analysis, and the second date from Oued Sdas (Ou_2). The first appearances of trilobites are visualised alongside the dates in Fig. 6d, and are dated to 519.46 Ma (519.25 to 519.68 Ma) at Tiout. The age estimate for the first Siberian trilobites differs considerably between the different alignment solutions: For the most likely alignment 1, the age estimate is 520.79 Ma (520.98 to 520.61 Ma), and for alignment 2 the estimate is somewhat higher at 521.05 Ma (521.19 to 520.91 Ma). Alignment 3 suggests a significantly later appearance of Siberian trilobites at 519.98 Ma (520.15 to 519.84 Ma). All three alignments place the appearance of the first Siberian trilobites before their appearance at Tiout, with the temporal gap (computed directly from the posterior distribution) being estimated at 1.33 Myr (1.09 to 1.54 Myr) for alignment 1, 1.71 Myr (1.54 to 1.87 Myr) for alignment 2, and 0.63 Myr (0.53 to 0.74 Myr) for alignment 3.




Figure 6: Three possible alignments identified by the inference with Cambrian data. (a) Exemplary sample from the cluster of the most likely alignment (93% of posterior samples). (b, c) Exemplary samples from a second and third identified alignment cluster (2.8% and 2.6% of posterior samples, respectively). Each shown alignment corresponds to a single sample from the posterior; other samples will result in slightly different alignments. 1.5% of samples were not assigned to any cluster (see Fig. 5). The curved dark lines show the cubic B-splines corresponding to each visualised sample. The coloured bars to the right of each alignment show the median duration of the stratigraphic partitions under each respective alignment cluster, based on the median ages of partition boundaries, with colours repeating the colour scheme of Fig. 2. (d) Posterior density of the inferred ages corresponding to the radiometric dates to the left (3 from Tiout, 4 from Oued Sdas, and 1 from Talat n'Yssi) and the first occurrences of trilobites at Tiout (Ti tril.) and Sukharikha (Sh tril.) to the right, in colours. All samples from all alignment clusters were included. Greater width corresponds to higher posterior density; all densities are scaled to have the same maximum for better visibility. Densities representing the uncertainties of radiometric dates based on their mean and standard deviation are shown in grey (left). The faint yellow shading to the right shows the prior density on α , i.e. the first appearance of trilobites at Tiout and Siberia based on a mean age of (520 Ma) and a standard deviation of 2 Myr (identical for

Tiout and Siberia). Colours and shapes of the points correspond to the four sites: Tiout - brown circles; Oued Sdas - pink squares; Talat n'Yissi - green diamonds; Sukharikha - blue triangles.

The posterior of the model runs allows the construction of age models that span the entire height of each site (Fig. 7). As sedimentation rates are constrained to be constant within the pre-defined partitions, sedimentation rate changes are visible as inflections at the boundaries of these partitions. Age uncertainties are relatively low at Tiout and most of Oued Sdas, which are relatively well constrained by radiometric dates in the top (Tiout) and middle (Oued Sdas) parts of the sections, as well as by astronomical priors on sedimentation rates. Uncertainty noticeably increases towards the top and bottom of Oued Sdas. The lowest partition of Oued Sdas is constrained only by its match to the lower part of the Sukharikha Fm., their age estimates are thus varying considerably between different alignments (Fig. 6). Differences in the positioning of the $\delta^{13}\text{C}$ curves between alignments are greatest at Talat n'Yissi and the Siberian Krasnoporog Fm. (Fig. 6), which results in large uncertainties in the age models (Fig. 7c, d).




Figure 7: Age-depth model for each of the four sites. The solid lines indicate the median posterior ages corresponding to the respective heights; the shaded interval denotes the 95% credible interval of posterior ages. Colours correspond to the colours of partitions introduced in Fig. 3. Circles indicate the mean age estimates of radiometric dates, with vertical lines spanning two standard deviations around the mean of these age estimates. Stars denote the first appearances of trilobites in Morocco and Siberia. See Fig. B2 for separate visualisations of age-depth models for different alignment solutions.

5 Discussion

5.1 Lower Cambrian stratigraphy

We used StratoBayes to correlate and date four lower Cambrian carbonate sections using $\delta^{13}\text{C}$ records, radiometric dates and astrochronological sedimentation rate estimates. From a large space of possible alignment configurations (Fig. 4), the software identified alignment solutions that visibly match the large-scale features in the $\delta^{13}\text{C}$ records from multiple sites, while simultaneously achieving an approximate fit to the radiometric dates (Fig. 6).

The most likely alignment solution from the posterior, alignment 1 (probability = 93%), results in a correlation of the three Moroccan sites that has much in common with that proposed by Maloof et al. (2005). In our model, we used common sedimentation rates for the stratigraphic partitions (members, formations) shared between the sites, whilst allowing sedimentation rates to systematically differ from the reference sedimentation rates at Tiout by adding a site-specific multiplier. This multiplier, ζ , is 1.02 (95% credible interval: 0.97 to 1.08) for Oued Sdas, meaning the model estimates very similar sedimentation rates for Tiout and Oued Sdas (Fig. 6a), consistent with their close geographical proximity. Sedimentation rates for the shared partitions at Talat n'Yissi are lower by a factor of 0.86 (0.76 to 0.96), which would be consistent with a moderately lower accommodation space at Talat n'Yissi relative to Tiout and Oued Sdas (as suggested by Fig. 3B in Maloof et al., 2005). We deliberately chose broad priors that did not explicitly enforce a relationship between sedimentation rates and palaeogeography; nonetheless, the model identified a geologically plausible solution. In contrast, the higher $\zeta_{\text{Talat n'Yissi}}$ of alignment 2 (probability = 2.8%, 1.07 to 1.37) and alignment 3 (probability = 2.6%, 2.07 to 2.45) are harder to reconcile with the palaeogeographic context.

Alignments 2 and 3 also suggest different sedimentation rates between Tiout and Oued Sdas, with a higher value of $\zeta_{\text{Oued Sdas}}$ (1.13 to 1.26) being estimated by alignment 2, and a lower value of $\zeta_{\text{Oued Sdas}}$ (0.83 to

0.88) by alignment 3. The most consistent lithostratigraphic alignment between Tiout and Oued Sdas is achieved by alignment 1, meaning that the age estimates for partition boundaries (based on members or formations) are most similar (Fig. 6). For the more distant Talat n'Yissi, age estimates of partition boundaries differ to varying degrees across all three alignments.

Breaking down the posterior probability into individual components – likelihood (fit of $\delta^{13}\text{C}$ measurements to the spline, fit of age estimates to the radiometric dates) and prior probability from the overlap penalty – reveals that samples from alignment 1 have a higher likelihood, on average (Fig. 5c). In contrast, alignments 2 and 3 have a greater number of overlapping $\delta^{13}\text{C}$ points, which results in higher overlap prior probabilities (Fig. 5c). The overlap prior reflects the prior belief that substantial parts of the sections involved in the correlation should be overlapping. However, the weight of that prior is somewhat arbitrary and reflects the technical requirement to facilitate overlap despite non-overlap allowing for closer fit to the spline, similar to the role of the “edge value” in some DTW implementations (Hay et al., 2019). A lower prior weight on overlap would thus have caused alignments 2 and 3 to receive lower posterior probabilities relative to alignment 1. Taken together, the evidence from above leads us to strongly favour alignment 1, and we will focus further discussion on that most likely alignment solution.

A radiometric date of 517.0 Ma (± 2 SD: 515.5 – 518.5 Ma) has been recovered from the Lemdad Syncline in the Atlas mountains (Landing et al., 1998), and has been correlated biostratigraphically to a horizon in the lower Isaafen Fm. at Talat n'Yissi (Maloof et al., 2005). In our alignment 1, this horizon has a posterior age estimate of 519 Ma (519.2 to 518.8 Ma) – ≈ 2 Myr older than the mean of the radiometric date. This date has informed the age estimates for Talat n'Yissi in Maloof et al. (2005) and Maloof et al. (2010), whereas alignment 1 produces age estimates close to those of Bowyer et al. (2022) and Bowyer et al. (2023). Age estimates deviating from radiometric dates are not necessarily incorrect: Although radiometric dates are sometimes treated as “absolute truth” within the stratigraphic community, they are the result of various sources of technical uncertainties (Condon et al., 2024) and geological interpretations like the actual zircon crystallisation versus eruption age (Keller et al., 2018). This is illustrated by the recalculation of the radiometric date from Landing et al. (1998) to 515.56 Ma (± 2 SD: 514.40 – 516.72 Ma) in the Geological Time Scale 2012 (Schmitz et al., 2012).

The two radiometric dates measured at Tiout at the bottom of and within the Amouslek Formation suggest a sedimentation rate of 146 m Myr^{-1} (± 2 SD: 78.7 to 613 m Myr^{-1}) for the Amouslek formation. However, the posterior estimates for the sedimentation rate in the Amouslek formation are poorly constrained and high compared to the sedimentation rates of all other partitions, at 3030 m Myr^{-1} (800 to 17,300 m Myr^{-1}). It appears that the model has overestimated the Amouslek sedimentation rate in aligning the $\delta^{13}\text{C}$ record of

the overlying Isaafen formation with a part of the Siberian Krasnoporog formation which has similar $\delta^{13}\text{C}$ values (Fig. 6a). The alignments of Bowyer et al. (2022) imply significant sedimentation rate changes within the Krasnoporog formation, allowing the $\delta^{13}\text{C}$ records to be better reconciled with the radiometric dates. We didn't allow for sedimentation rate changes within the Krasnoporog formation because the stratigraphic log of Kouchinsky et al. (2007) indicates a uniform facies. Additional sedimentation rate changes might lead to a closer alignment with the radiometric dates, at the cost of greater model complexity.

The alignment of the Siberian Sukharikha section with the Moroccan sites is relatively precise in the lower half of the records: The prominent positive $\delta^{13}\text{C}$ excursions interpreted as the “5p” and “6p” excursions have a similar magnitude both at Oued Sdas and Sukharikha, and are readily aligned visually (Bowyer et al., 2022) and by our model (Fig. 6). Our model aligns the main 6p peak of Sukharikha with the first subpeak of the second large excursion at Oued Sdas, as in model C in Bowyer et al. (2022). The lesser, positive excursion below the hiatus at the top of the Sukharikha formation lines up with the positive excursion in the lower Lie-de-Vin formation, representing the “II” peak as in model C in Bowyer et al. (2022). The upper parts of the Moroccan records and the Siberian Krasnoporog formation appear to be aligned primarily by matching the prominent positive excursion interpreted as excursion “IV” (Bowyer et al., 2022; Kouchinsky et al., 2007). The “III” peak below is only weakly expressed at Oued Sdas, leading to uncertainty in the alignment with the corresponding part of the Krasnoporog formation, and in the inferred duration of the hiatus even within alignment solution 1 (Fig. B3a-c). Similarly, considerable uncertainty exists in how the top of Talat n'Yissi corresponds to the Krasnoporog formation. This is evident from variations between samples in alignment solution 1 (Fig. B3a-c) and in the wide credible intervals of those parts of the age models (Fig. 7). The relatively small magnitude of $\delta^{13}\text{C}$ changes limits the model's ability to identify a definitive alignment solution for that part of the record.

Our estimate for the Moroccan first appearance of trilobites at Tiout from alignment 1, 519.47 Ma (519.68 to 519.26 Ma), is slightly younger and somewhat less precise than the recent, astrochronological estimate of 519.62 Ma (95% highest posterior distribution: 519.70 to 519.54 Ma) by Sinnesael et al. (2024). We attribute this difference to our model simultaneously combining different data types from multiple sites. Additionally, Sinnesael et al. (2024) allowed sedimentation rates to vary between cycles, whereas our model assumed a single sedimentation rate per member. In our alignment 1 solution, the highest $\delta^{13}\text{C}$ values of Tiout correlate to shortly after the peak of the IV $\delta^{13}\text{C}$ excursion. This correlation suggests that the actual peak of the excursion at Tiout has not been sampled by Magaritz et al. (1991) and Tucker (1986), which may result in misalignments when correlating the record to other sections. Further $\delta^{13}\text{C}$ samples from the lower Igoudine and upper Lie-de-Vin formation at Tiout are required to improve correlation with other sections, including the correlation presented herein.

Our model successfully reconstructs the first appearance of trilobites at Tiout, within error, despite using a simpler astrochronology and enforcing a less variable sedimentation rate history than Sinnesael et al. (2024). It also provides the first fully quantitative estimate for the first appearance of trilobites in Siberia based on chemostratigraphic correlation and the Moroccan radiometric dates and astrochronology, at 520.79 Ma (520.98 to 520.61 Ma). This refines earlier estimates of ≈ 521 Ma (Landing et al., 2021), and quantifies the temporal gap between the appearance of trilobites in Siberia and Morocco as 1.33 Myr (1.09 to 1.54 Myr). We do not suggest that these estimates are definitive; indeed, we anticipate that the incorporation of additional $\delta^{13}\text{C}$ data from Tiout, the inclusion of astrochronological estimates of individual short eccentricity cycles, and the relaxation of the assumption of constant sedimentation rates within partitions may update the estimate. A high-resolution temporal sequence of trilobite first occurrence dates could be used to delineate trilobite evolutionary rates and dispersal; to evaluate evolutionary hypotheses on the origins and biomineralisation of trilobites (Holmes and Budd, 2022; Paterson et al., 2019); and to inform the definition of the base of the Cambrian Series 2 (Zhang et al., 2017).

5.2 Statistical alignment and age modelling

5.2.1 Advantages of Bayesian stratigraphic alignment

As shown above, our algorithm can identify the correct alignment positions in scenarios with one (Fig. 1) or more than one (Fig. 2) known solution. In scenarios where more than one distinctly different alignment is identified, the probability of each solution, given the specified data and model, is identified. This can be used to evaluate the likelihood of competing models for the alignment of stratigraphic records found by visual (e.g. Bowyer et al., 2023; Landing and Kruse, 2017) or algorithmic (e.g. Hay et al., 2019) correlation. The requirement to specify priors for the alignment parameters can be leveraged to provide information beyond that which is contained in the signals: for example, information on sedimentation rates may be expressed in the prior.

Because our model can integrate absolute age constraints such as radiometric dates, a user is able to correlate stratigraphic records and construct probabilistic age models in a single step. In our Cambrian example, the posterior alignment and the posterior age model are thus influenced by the priors, the quantitative signals and the radiometric dates. In contrast, age models constructed in a separate step after identifying alignments do not reflect uncertainty arising during the alignment stage (Hagen and Creveling, 2024).

In our integrated approach, discrepancies between radiometric dates and signal alignment are resolved probabilistically, with the model weighting the available evidence based on its likelihood and prior

information. This means that posterior age estimates may diverge from the age information provided by radiometric dates, as seen with the Ou_2 date in Fig. 6d. This is not necessarily a deficiency of the model; rather, it indicates that the priors and non-radiometric data provide sufficiently strong evidence to suggest that the actual age of the horizon associated with the radiometric date falls toward the tails of its confidence interval, or that the radiometric uncertainty may be underestimated. Some degree of discrepancy is expected when integrating multiple data types rather than relying on a single proxy (see also Lee et al., 2022).

If, on the other hand, the user wishes to increase the influence of radiometric dates on the posterior age estimates, this can potentially be achieved by introducing additional sedimentation rate changes to allow more flexible alignment of the proxy signals, reducing the weight of the proxy signal records – such as by imposing a larger σ for the cubic spline – or by weakening priors.

5.2.2 Model choice and priors

Stratigraphic alignment using algorithms has the advantage of removing some of the inherent subjectivity of visual alignment (Sylvester, 2023). Yet, somewhat subjective decisions are still explicitly or implicitly made with every alignment algorithm. In the case of DTW, subjectivity is introduced e.g. with restrictions on the warping path (i.e. relative sedimentation rates, Sakoe and Chiba, 1978), with the amount of overlap required between sections (Hay et al., 2019), or with the choice of an exponent controlling the weight of outlier values (Wheeler and Hale, 2014). All of those settings can alter the outcome of DTW-based alignments. Likewise, our Bayesian approach comes with a number of subjective choices. The appropriate model structure can be readily determined when the data-generating process is known (Sect. 3), but has to be carefully considered and potentially revised when dealing with complex real-world data (Sect. 4). Lithological data may guide the partitioning of data and can inform somewhat objective choices of horizons with likely sedimentation rate changes (Sect. 4.2), but such information may not be readily available with some datasets, such as with well logs.

Besides the model structure, StratoBayes requires the user to specify priors for several model parameters: relative or absolute sedimentation rates (v, ζ), the shifts of sections relative to one another (α), the duration of hiatuses (δ), the degree of smoothing of the spline (λ), the extent to which overlap of signal points should be favoured (C_{overlap}), and optionally the residual standard deviation of the spline (σ). Although the choice of any of those parameters has the potential to affect posterior alignments and age models, they also offer a chance to explicitly include geological information that could otherwise only be incorporated by discarding or modifying alignment solutions after the algorithmic alignment.

While it is relatively straightforward to express prior beliefs on the alignment parameters α , ν , ζ and δ , it is hard to specify suitable priors for λ , σ and $C_{overlap}$, as they do not correspond to measures used by geologists. The default priors on λ , σ and $C_{overlap}$ in the StratoBayes software were chosen iteratively by working with various test data sets. Users should avoid fine-tuning these priors directly on the data sets to which they intend to apply StratoBayes, as this could introduce unintended circularity. Instead, analogous independent data sets could be used to identify suitable priors for λ , σ and $C_{overlap}$. For example, priors on λ and σ for correlating $\delta^{13}\text{C}$ curves could be meaningfully specified from pre-existing reconstructed $\delta^{13}\text{C}$ composite curves.

5.2.3 Challenges with the proxy and sedimentary record

Chemostratigraphy, and, more broadly, correlating geological sections based on proxy data relies on the proxies accurately reflecting a common, underlying signal. Several processes may disrupt this assumption. For example, $\delta^{13}\text{C}$ recorded in carbonates differs between different depositional environments, water depths, and grain types (Geyman and Maloof, 2021), while the $\delta^{13}\text{C}$ recorded in restricted basins may be offset significantly relative to contemporary carbonates elsewhere (Uhlein et al., 2019). Where known, such offsets could be accounted for by subtracting or adding the estimated offset relative to global values. Alternatively, anticipated offsets could be modelled as additional unknown variables, as in Edmonson and Dyer (2024). This approach will likely require substantial prior knowledge on the potential magnitude and direction of offsets; otherwise, the combination of variation along the height or time axis and along the proxy value axis may result in a large range of mathematically feasible alignments.

A more fundamental problem is posed when similar patterns in a proxy curve are asynchronous in different sections: Shifting and stretching proxy data from multiple sites may result in a strongly correlated composite curve, but this correlation does not prove that the patterns or excursions observed at different sites were in fact synchronous (Blaauw, 2012). Unless supported by independent evidence such as precise radiometric dates, relative age estimates derived from proxy correlations (e.g. $\delta^{13}\text{C}$) are conditional on the assumption of synchronicity.

Several challenges arise from the variability of sedimentation and the incompleteness of the sedimentary record. Sediment accumulation rates vary with measurement scale (Sadler, 1981): closer spacing between measurements allows more variability to be identified, with actual sedimentation rate histories displaying fractal properties (Miall, 2015). This implies that depositional ages tend to vary non-linearly along a vertically sampled sedimentary section, with substantial incompleteness in shallow-water records (Curtis et al., 2025). These discontinuities can lead to drastically altered shapes of proxy curves from different depositional settings, and cycles from periodic proxy fluctuations may be missed due to insufficient

preservation or sampling (Curtis et al., 2025). This issue is evident in the Sukharikha section, where it is somewhat ambiguous whether the hiatus represents a fraction of a $\delta^{13}\text{C}$ excursion (alignment 1 and 2) or extends over more than one full cycle (alignment 3, Fig. 6). For correlations within sedimentary basins, the method of Bloem and Curtis (2024) could help resolve ambiguous alignments by reconstructing depositional histories through geological process modelling, but this method requires exceptionally high-resolution sampling and its utility has yet to be demonstrated with real-world data sets.

Besides the completeness, the sampling density of proxy records may influence correlations. In StratoBayes, densely sampled sections or parts of sections exert more influence on the shape of the spline than those that are thinly sampled, which can be seen in the spline curve primarily following the densely sampled Oued Sdas and Talat n'Yissi records in Fig. 6. Despite this, our Cambrian case study demonstrates that sections with differing sampling densities – both between and within sites – can still be effectively aligned. Varying sampling density would, however, pose a challenge for reconstructing a global average proxy curve from local records, as the global curve would primarily reflect the more densely sampled sites.

StratoBayes introduces a simplification in modelling sedimentary histories by forcing uniform sedimentation rates within pre-defined segments of a stratigraphic section. An effect of this simplification can be seen in the age-depth plots in Fig. 7: Due to sedimentation rates being modelled as uniform within stratigraphic partitions, the uncertainty of age estimates does not necessarily increase away from the radiometric dates. We acknowledge that this may underestimate the uncertainty associated with potential sedimentation rate variability (De Vleeschouwer and Parnell, 2014), especially when allowing for few sedimentation rate changes. Similarly, our method currently only allows for specifying potential hiatuses with an unknown duration at fixed, predetermined heights.

In principle, our method could be used to divide stratigraphic sections into an arbitrary number of segments with differing sedimentation rates, and with an arbitrary number of potential hiatuses. In practice, estimating the parameters of a model with more than a low double-digit number of alignment parameters (shift parameters, sedimentation rates, hiatuses) represents a challenge for the current implementation of the MCMC algorithm within StratoBayes, as finding and exploring the posterior becomes increasingly difficult as more parameters are added. This limitation could be alleviated by incorporating MCMC methods suited for higher dimensional problems and difficult posterior geometries. Alternatively, a continuous process model such as the compound Poisson-gamma process of BChron (Haslett and Parnell, 2008) might be integrated with our model for the proxy signal, but again the complexity of the MCMC would increase. Another approach would be to divide the alignment problem into sub-problems, e.g. by multiple pairwise correlation of sites (e.g. Hagen et al., 2024; Sylvester, 2023), or by correlating shorter subsections.

5.3 Towards quantitative stratigraphy

Quantitative stratigraphic correlation and age modelling of diverse geological data represent a long-standing challenge in stratigraphic research. Although many algorithms exist for correlating geochemical and geophysical stratigraphic data (e.g. Baviile et al., 2022; Bloem and Curtis, 2024; Hay et al., 2019; Sylvester, 2023); few can readily provide uncertainty estimates or incorporate different types of data simultaneously (e.g. Al Ibrahim, 2022; Edmonson and Dyer, 2024; Lee et al., 2022). Consequently, integrated statistical approaches have only rarely been applied to complex real-world stratigraphic problems (Hagen et al., 2024; Lee et al., 2022).

Our new method has the potential to be applied to diverse datasets; examples range from shallow borehole data from the Holocene (Finlay et al., 2022) to Proterozoic carbonates (Halverson et al., 2010). The ability of our model to incorporate multiple proxy records simultaneously opens new possibilities for refining stratigraphic correlations. For instance, correlations involving both $\delta^{13}\text{C}$ and $\delta^{87}\text{Sr}$ records could benefit from a probabilistic framework that accounts for their respective uncertainties (Bowyer et al., 2022). The integration of multiple proxies, e.g. multiple element ratios, in the StratoBayes framework could allow correlations based on the entire record of all proxies, rather than a few visually distinct transitions (Craigie, 2015).

Beyond geochemical records, our approach could also be applied e.g. to geophysical well-logs such as gamma ray or density logs, and magnetostratigraphic records could be correlated directly rather than relying on visually interpreted polarity reversals (Langereis et al., 2010). While index fossils can currently be integrated as tie points, the modelling framework could be expanded to explicitly model first and last occurrences to better incorporate biostratigraphic uncertainty. Similarly, astrochronological constraints can be expressed as priors on sedimentation rates, but an additional model component would be needed to incorporate all astrochronological information from a given site (Sinnesael et al., 2024).

Conclusions

StratoBayes is a Bayesian modelling framework for the probabilistic alignment of stratigraphic proxy records and age modelling. It correlates quantitative proxy signals such as isotope ratios, and integrates additional stratigraphic information such radiometric dates, to construct probabilistic age models. Applying our model to both simulated data and real-world stratigraphic records from the lower Cambrian of Morocco and Siberia, we have demonstrated its ability to account for uncertainty from all model components and to identify multiple plausible alignment solutions. Our lower Cambrian case study provides a fully

probabilistic estimate for the first appearance of trilobites in Siberia, and quantifies the temporal gap between their first occurrence and the oldest Moroccan trilobites. While our results remain dependent on model assumptions, they represent a step towards a more objective and reproducible approach to early Palaeozoic stratigraphy; they also highlight sources of uncertainty and identify targets for future research. Beyond this case study, StratoBayes has broad applicability to stratigraphic problems across all time intervals that involve the correlation of quantitative proxy records.

Appendix A: Markov chain Monte Carlo sampling scheme

Appendix A details the Metropolis-within-Gibbs sampling scheme and the parallel tempering framework that are used within the StratoBayes software to sample from the posterior of the unknown model parameters.

Sampling strategy

The MCMC sampling scheme used in this study includes an adaptive phase. During this phase, proposal distributions and the probabilities with which different proposal types are selected for the Metropolis-Hastings updates are adjusted based on the history of the MCMC chains to improve acceptance rates and mixing. Additionally, the temperature ladder of the parallel tempering framework is updated to improve the swap rates of chains. After the adaptive phase, the proposal distributions and probabilities, as well as chain temperatures, remain fixed for the remainder of the run to ensure proper sampling from the posterior.

In the current implementation, the length of the adaptive phase is pre-determined by the user, specified as a fixed number of iterations. However, the user has the option to extend the adaptation period by continuing the run if needed. More generally, adaptation could also be stopped automatically based on criteria such as mixing within chains (Yang and Rosenthal, 2017) or convergence criteria.

Adaptive MCMC algorithms do not always preserve the stationarity of the target distribution during the adaptive phase (Roberts and Rosenthal, 2009). Therefore, all samples from the adaptive phase are discarded as burn-in. Additionally, if diagnostic checks suggest that the MCMC has not converged by the end of the adaptive phase, further samples may need to be discarded.

Gibbs sampling scheme for the cubic B-splines

The following sampling scheme was adapted from Heaton et al. (2020). The spline coefficients are sampled from a multivariate normal distribution of the form:

$$\beta \sim MVN(b\mathbf{Q}, \mathbf{Q}), \quad (16)$$

where b is given by .

$$b = (\mathbf{B}(h))^T \frac{y}{\sigma^2}, \quad (17)$$

$\mathbf{B}(h)$ are cubic B-splines (Eilers and Marx, 1996) at a set of k knots evaluated at heights h at which y , the composite stratigraphic signal of all sites, was observed. Here, σ is the residual standard deviation.

The other element needed for sampling from the posterior of b is $\mathbf{Q}$, given by

$$\mathbf{Q} = (\mathbf{H} + \lambda \mathbf{D})^{-1}, \quad (18)$$

where λ is a smoothing parameter, $\mathbf{D}$ is a penalty matrix to prevent the spline from overfitting the data, and

$$\mathbf{H} = \left(\frac{\mathbf{B}(h)}{\sigma} \right)^T \frac{\mathbf{B}(h)}{\sigma} \quad (19)$$

The standard deviation σ can be fixed as

$$\sigma = \frac{1}{S} \sum_{s=1}^S \sigma_s, \quad (20)$$

where S is the number of sites, and σ_s is the standard deviation of individual splines fitted to the data of site s . This often provides a good approximation of σ , while removing an unknown model parameter, potentially facilitating quicker convergence of the model run.

Alternatively, σ can be estimated within the Gibbs sampling scheme from the data, by placing a conjugate gamma prior on the inverse of the variance (precision, $\tau = 1/\sigma^2$):

$$\sigma^{-2} \sim \text{Gamma} \left(a_\sigma + \frac{n_y}{2}, b_\sigma + \frac{1}{2} \sum_{y=1}^{n_y} (y - \beta \mathbf{B}(h))^2 \right) \quad (21)$$

The smoothing parameter λ is estimated by placing a gamma prior on λ :

$$\lambda \sim \text{Gamma} \left(a_\lambda + \frac{k}{2}, \frac{1}{\frac{1}{b_\lambda} + \frac{1}{2} \sum^k \beta \mathbf{D} \times \beta} \right) \quad (22)$$

Metropolis-Hastings step

The starting heights or ages α , sedimentation rates ν , site multipliers ζ and gaps δ are updated in a Metropolis-Hastings step. For each unknown parameter, a new value is randomly sampled from a proposal distribution. Initially, proposals are sampled independently for each parameter from its respective prior, or alternatively from a custom proposal distribution.

In the following, the current set of parameter values is labelled θ , and the proposed set is labelled θ' . To decide whether to accept or reject the new set of parameters, an acceptance probability A is calculated, and the proposal is randomly accepted or rejected with a probability of A . This probability is calculated as

$$A = \min\left(1, \frac{\pi(\theta')}{\pi(\theta)}\right), \quad (23)$$

where $\pi(\theta)$ is the unnormalised posterior probability of the current values, and $\pi(\theta')$ is the unnormalised posterior probability of the proposed values. These can be calculated as

$$\pi(\theta) = p(\theta) \times L(\text{data}|\theta), \quad (24)$$

where $p(\theta)$ is the prior probability of θ , and $L(\text{data}|\theta)$ the likelihood of the data given θ .

We calculate the likelihood of the data given θ as a product of the probability densities of each data point of the signal y (recorded at two or more sites) and of all absolute age information. For the signal, we assume that the observed values y are normally distributed and centred around the values predicted by the splines, μ , at height h , with a standard deviation σ which has been introduced earlier. The likelihood of a data point i from the signal y is thus

$$L(y_i|\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \times e^{\left(-\frac{(y_i-\mu_i)^2}{2\sigma^2}\right)} \quad (25)$$

and the log-likelihood for all data points of the signal is calculated as

$$\ln L(y|\theta) = \sum_i \ln L(y_i|\theta) \quad (26)$$

If more than one type of signal is used, the log-likelihood of additional signals can be calculated analogously and added in Equation 29.

Age constraints are incorporated by using an age estimate from radiometric dates d with, for example, mean ages a_{mean} and uncertainties given by standard deviations a_{sd} . The probability density of a date d_i is then calculated as

$$L(d_i|\theta) = \frac{1}{\sqrt{2\pi a_{sd,i}^2}} \times e^{\left(\frac{a_{mean,i} - a_{predicted,i}}{2a_{sd,i}^2}\right)} \quad (27)$$

where $a_{predicted,i}$ is the age predicted by the age-height transform at the height $h_{s,i}$, the height at the site at which date d_i was obtained.

The log-likelihood for all age constraints is calculated as

$$\ln L(d|\theta) = \sum_i \ln L(d_i|\theta) \quad (28)$$

and the overall likelihood, if absolute age constraints are included, is

$$\ln L(y, d|\theta) = \ln L(y|\theta) + \ln L(d|\theta) \quad (29)$$

Proposal types

In order to allow for a broad search of the parameter space, proposals are initially selected independently for each parameter, and are selected independently of the current parameter values. These proposals lead to a decreasing acceptance rate over time, and the chain tends to arrive at a single set of values with high posterior probability, $\pi(\theta)$, remaining there for many iterations due to frequent rejections. Therefore, different types of proposals are used after an initial period:

- 1) Proposing from the prior or a custom distribution: This proposal is used exclusively for a small number of initial iterations and is alternated with other proposals later on.
- 2) Adaptive independent (univariate) proposals: Proposals for each parameter are selected independently from other parameter values. Proposals are dependent on the current state of the parameter θ_i , and sampled from a normal distribution $N(\theta_i, \sigma_i)$, where σ_i is a standard deviation that is estimated based on the history of the MCMC chain, i.e. based on the sampled θ_i from previous iterations.
- 3) Adaptive dependent (multivariate) proposals (Roberts and Rosenthal, 2009): Proposals for the parameters are selected jointly and are dependent on the current state of the parameters θ . Proposals are sampled from a multivariate normal distribution $MVN(\theta, \Sigma)$, where Σ is a

covariance matrix that is estimated based on the history of the MCMC chain, i.e. based on the sampled θ_i from previous iterations.

- 4) Shifting some or all α and or δ parameters while keeping the other parameters constant. This can accelerate the convergence of the MCMC in cases where some sites are aligned with each other, but offset relative to other sites.

Proposal types are chosen with a probability that broadly corresponds to the relative acceptance probability of the respective proposal type, i.e. proposal types that are rejected often are chosen less frequently. Adaptation for types 2) and 3), and the adjustment of proposal type probabilities ends after the adaptive phase. Posterior samples from the adaptive phase have to be discarded as burn-in, to ensure the correct convergence of the chain.

Parallel tempering

To avoid the MCMC chain becoming trapped at isolated peaks of the posterior probability distribution, we implement a parallel tempering framework, following Sambridge (2014). This involves running multiple chains in parallel. The target chain, the chain from which the posterior samples will be taken, is left unaltered (“cold chain”). The other chains are tempered, i.e. their unnormalised log posterior probabilities are raised to the power of $1/T$, with T being the temperature. The higher T , the more “flattened” the posterior probability landscape becomes, and the easier it is for the chain to explore the landscape. Frequently, chain swaps are proposed, during which the model parameter values of different chains are exchanged with a Metropolis-Hastings acceptance probability based on the ratios of posterior probabilities of the states of the two chains, evaluated at both temperatures as in Appendix A2 of Sambridge (2014).

The initial temperatures for a number of chains n_c are selected using a geometric spacing, with $T_1 = 1$ (cold chain) and $T_{n_c} = \infty$ (hottest chain). The infinite temperature of the hottest chain implies that all proposals during the MCMC will be accepted, and we let that chain sample from the prior probability distributions of the parameters. If $n_c > 2$, intermediate chain temperatures are selected as

$$T_c = 10^{\sum_{j=2}^c d_j}, \quad (30)$$

where

$$d_c = \frac{(n_c - 1)^{(2/3)}}{n_c - 2} + \frac{c - 1 - (n_c - 1)/2}{1.5 * n_c}, \quad c = 2 \dots n_c - 1 \quad (31)$$

This leads to the spacing of temperatures decreasing with increasing number of chains, and temperature spacing is narrower for lower temperatures on the log scale. A small amount of white noise from a normal distribution with zero mean and a standard deviation of $(5 \times n_c)^{-1}$ is added to each d_c to vary the initial temperature ladders between independent model runs. Temperatures are updated in the adaptive phase of the MCMC to increase the swap rates of chains (Vousden et al., 2016).

Appendix B: Inspecting the posterior of the lower Cambrian case study

Appendix B provides additional details on the posterior of the inference with lower Cambrian $\delta^{13}\text{C}$ data and radiometric dates.

Trace plots

Trace plots visualise the evolution of chains from an MCMC and, together with tools such as the potential scale reduction factor (Gelman and Rubin, 1992; Vats and Knudson, 2021), allow for assessing convergence of model runs. The trace plot indicative of a well-behaved model run should be stationary after the burn-in phase, with different chains mixing well (Gelman et al., 1995). An example of a well-behaved trace plot is the first panel of Fig. B1. Inspecting the trace plots of the 18 model parameters of the lower Cambrian case study reveals that all parameters seem to have reached stationarity, this said; some chains occasionally visit distinctly different values (e.g. Fig. B1, column 1, row 2). The chains are not mixing well in those regions of the parameter space. Running the model for considerably more iterations is likely to overcome this problem. However, this affects only the less likely alignments; the most likely alignment (alignment cluster 1) is well explored across all parameters.




Figure B1: Trace plots of the 18 alignment parameters. Each colour corresponds to a distinct run. For visual clarity, only 250 samples are displayed per run. The burn-in phase (the first 150,000 iterations) is omitted.

Age-depth models for different alignments

The age-depth models for each of the four sites are shown for each alignment cluster separately in Fig. B2 (instead of for all samples combined as in Fig. 7).




Figure B2: Age-depth model for each of the four sites. The solid lines indicate the median posterior ages corresponding to the respective heights; the shaded interval denotes the 95% credible interval of posterior ages. Colours correspond to the three different alignment clusters and outlier samples. Circles indicate the mean age estimates of radiometric dates, with vertical lines spanning two standard deviations around the mean of these age estimates. Crosses denote the first appearances of trilobites in Morocco and Siberia.

Variation within alignment clusters

Summarising the posterior by grouping samples into clusters of similar alignments facilitates discussion of the results but risks oversimplifying the variation within each cluster. Each cluster represents a set of posterior samples that share similar inferred ages for the partition boundaries, but differences still exist between individual samples within the same cluster. As an example, three distinct alignments from cluster 1 are visualised in Fig. B3. An alignment from a sample not assigned to any cluster is shown in Fig. B3d.




Figure B3: Alternative alignments, each corresponding to a single sample from the posterior. (a) A sample from the most likely cluster 1, corresponding to that shown in Fig. 6a. (b, c) Alignments corresponding to other samples from cluster 1. (d) Alignment corresponding to an outlier sample that was not assigned to any cluster. The curved dark lines show the cubic B-splines corresponding to each alignment.

864 Posterior of alignment parameters

865 The posterior distributions of the alignment parameters are summarised in histograms in Fig. B4.




Figure B4: Comparison of prior and posterior probability densities. Histograms in colour denote the posterior probability densities of the 18 alignment parameters; the grey, smooth shadings represent prior probability densities. The four colours correspond to the four independent model runs.

Code and data availability

The StratoBayes R package is available for download as a binary at stratobayes.github.io. The data and R scripts used to generate the results are available at <https://zenodo.org/records/15065336>.

Author contribution

MRS and ARM conceived the idea. MRS, ARM and MS designed the initial model version, and all authors contributed to the design of the final model. MS, ARM and MRS coded the software prototype. KE, MRS and ARM modified and expanded the software. MS curated the lower Cambrian data. KE conducted the analyses and visualised the results. KE wrote the initial manuscript draft, and all authors contributed to the final version. MRS coordinated the project. MRS and ARM acquired the funding.

Competing interests

The authors declare that they have no conflict of interest.

Acknowledgements

We thank the participants of the StratoBayes workshop (Durham, 2024) for feedback on an early version of the StratoBayes software. We thank Tim Heaton for sharing his code for spline fitting. We thank Andrew Valentine for discussions on the model. We thank Prof. Nasrddine Youbi and Kamal Mghazli for hosting an exploratory field visit to the lower Cambrian sections in Morocco. We thank Maarten Blaauw, Andrew Curtis and Norman MacLeod for their helpful feedback on an earlier version of this manuscript.

Financial support

This research was funded by Leverhulme Research Project Grant 2019-223 to MRS.

References

- Agterberg, F. P.: Automated stratigraphic correlation, Elsevier, 1990.
- Al Ibrahim, M. A.: Uncertainty in automated well-log correlation using stochastic dynamic time warping, *Petrophysics*, 63, 748–761, <https://doi.org/10.30632/PJV63N6-2022a11>, 2022.
- Blaauw, M. Out of tune: the dangers of aligning proxy archives. *Quaternary Science Reviews*, 36, 38–49, <https://doi.org/10.1016/j.quascirev.2010.11.012>, 2012.
- Blaauw, M. and Christen, J. A. Flexible paleoclimate age-depth models using an autoregressive gamma process. *Bayesian Analysis* 6, 457–474, <https://doi.org/10.1214/11-BA618>, 2011.
- Baville, P., Apel, M., Hoth, S., Knaust, D., Antoine, C., Carpentier, C., and Caumon, G.: Computer-assisted stochastic multi-well correlation: Sedimentary facies versus well distality, *Marine and Petroleum Geology*, 135, 105371, <https://doi.org/10.1016/j.marpetgeo.2021.105371>, 2022.
- Bloem, H. and Curtis, A.: Bayesian geochemical correlation and tomography, *Scientific Reports*, 14, 9266, <https://doi.org/10.1038/s41598-024-59701-4>, 2024.

898 Bowyer, F. T., Zhuravlev, A. Y., Wood, R., Shields, G. A., Zhou, Y., Curtis, A., Poulton, S. W., Condon,
899 D. J., Yang, C., and Zhu, M.: Calibrating the temporal and spatial dynamics of the Ediacaran-Cambrian
900 radiation of animals, *Earth-Science Reviews*, 225, 103913,
901 <https://doi.org/10.1016/j.earscirev.2021.103913>, 2022.

902 Bowyer, F. T., Zhuravlev, A. Y., Wood, R., Zhao, F., Sukhov, S. S., Alexander, R. D., Poulton, S. W.,
903 and Zhu, M.: Implications of an integrated late Ediacaran to early Cambrian stratigraphy of the Siberian
904 platform, Russia, *Bulletin*, 135, 2428–2450, <https://doi.org/10.1130/B36534.1>, 2023.

905 Campello, R. J., Moulavi, D., Zimek, A., and Sander, J.: Hierarchical density estimates for data
906 clustering, visualization, and outlier detection, *ACM Transactions on Knowledge Discovery from Data*
907 (TKDD), 10, 1–51, <https://doi.org/10.1145/2733381>, 2015.

908 Condon, D., Schoene, B., Schmitz, M., Schaltegger, U., Ickert, R. B., Amelin, Y., Augland, L. E.,
909 Chamberlain, K. R., Coleman, D. S., Connelly, J. N., et al.: Recommendations for the reporting and
910 interpretation of isotope dilution U-Pb geochronological information, *Geological Society of America*
911 *Bulletin*, 136 (9-10): 4233–4251, <https://doi.org/10.1130/B37321.1> 2024.

912 Craigie, N. W.: Applications of chemostratigraphy in Middle Jurassic unconventional reservoirs in
913 eastern Saudi Arabia, *GeoArabia*, 20, 79–110, <https://doi.org/10.1016/j.jafrearsci.2015.01.003>, 2015.

914 Cramer, B. and Jarvis, I.: Carbon isotope stratigraphy, in: *Geologic time scale 2020*, Elsevier, 309–343,
915 <https://doi.org/10.1016/B978-0-12-824360-2.00011-5>, 2020.

916 Curtis, A., Bloem, H., Wood, R., Bowyer, F., Shields, G. A., Zhou, Y., Yilales, M., and Tetzlaff, D.:
917 Natural sampling and aliasing of marine geochemical signals, *Scientific Reports*, 15, 760,
918 <https://doi.org/10.1038/s41598-024-84871-6>, 2025.

919 De Vleeschouwer, D. and Parnell, A. C.: Reducing time-scale uncertainty for the Devonian by integrating
920 astrochronology and Bayesian statistics, *Geology*, 42, 491–494, <https://doi.org/10.1130/G35618.1>, 2014.

921 Edmonson, S. and Dyer, B.: A Bayesian framework for inferring regional and global change from
922 stratigraphic proxy records (StratMC v1. 0), *EGUsphere*, 2024, 1–45, [https://doi.org/10.5194/egusphere-](https://doi.org/10.5194/egusphere-2024-2579)
923 [2024-2579](https://doi.org/10.5194/egusphere-2024-2579), 2024.

924 Eilers, P. H. and Marx, B. D.: Flexible smoothing with B-splines and penalties, *Statistical science*, 11,
925 89–121, <https://doi.org/10.1214/ss/1038425655>, 1996.

926 Finlay, A., Bates, R., Bensharada, M., and Davies, S.: Applying chemostratigraphic techniques to shallow
 927 bore holes: Lessons and case studies from Europe's lost frontiers, in: Europe's lost frontiers: Volume 1
 928 context and methodology, Archaeopress, 137–153, <https://doi.org/10.1214/ss/1177011136>, 2022.

929 Gelman, A. and Rubin, D. B.: Inference from iterative simulation using multiple sequences, *Statistical*
 930 *science*, 7, 457–472, <https://doi.org/10.1214/ss/1177011136>, 1992.

931 Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B.: Bayesian data analysis, Chapman; Hall/CRC,
 932 <https://doi.org/10.1201/9780429258411>, 1995.

933 Geyman, E. C. and Maloof, A. C.: Facies control on carbonate $\delta^{13}\text{C}$ on the Great Bahama Bank,
 934 *Geology*, 49, 1049–1054, <https://doi.org/10.1130/G48862.1>, 2021.

935 Hagen, C., Creveling, J., and Huybers, P.: Align: A user-friendly app for numerical stratigraphic
 936 correlation, *GSA Today*, 34, 4–9, <https://doi.org/10.1130/GSATG575A.1>, 2024.

937 Hagen, C. J. and Creveling, J. R.: An algorithm-guided ediacaran global composite $\delta^{13}\text{C}_{\text{carb}}$ –Bayesian
 938 age model, *Palaeogeography, Palaeoclimatology, Palaeoecology*, 112321,
 939 <https://doi.org/10.1016/j.palaeo.2024.112321>, 2024.

940 Halverson, G. P., Wade, B. P., Hurtgen, M. T., and Barovich, K. M.: Neoproterozoic chemostratigraphy,
 941 *Precambrian Research*, 182, 337–350, <https://doi.org/10.1016/j.precamres.2010.04.007>, 2010.

942 Haslett, J. and Parnell, A.: A simple monotone process with application to radiocarbon-dated depth
 943 chronologies, *Journal of the Royal Statistical Society Series C: Applied Statistics*, 57, 399–418,
 944 <https://doi.org/10.1111/j.1467-9876.2008.00623.x>, 2008.

945 Hay, C. C., Creveling, J. R., Hagen, C. J., Maloof, A. C., and Huybers, P.: A library of Early Cambrian
 946 chemostratigraphic correlations from a reproducible algorithm, *Geology*, 47, 457–460,
 947 <https://doi.org/10.1130/G46019.1>, 2019.

948 Heaton, T. J., Blaauw, M., Blackwell, P. G., Ramsey, C. B., Reimer, P. J., and Scott, E. M.: The IntCal20
 949 approach to radiocarbon calibration curve construction: A new methodology using Bayesian splines and
 950 errors-in-variables, *Radiocarbon*, 62, 821–863, <https://doi.org/10.1017/RDC.2020.46>, 2020.

951 Holmes, J. D. and Budd, G. E.: Reassessing a cryptic history of early trilobite evolution, *Communications*
 952 *Biology*, 5, 1177, <https://doi.org/10.1038/s42003-022-04146-6>, 2022.

953 Keller, C. B., Schoene, B., and Samperton, K. M.: A stochastic sampling approach to zircon eruption age
 954 interpretation, *Geochemical Perspectives Letters* (Online), 8, 10.7185/geochemlet.1826, 2018.

955 Kouchinsky, A., Bengtson, S., Pavlov, V., Runnegar, B., Torssander, P., Young, E., and Ziegler, K.:
 956 Carbon isotope stratigraphy of the Precambrian–Cambrian Sukharikha river section, Northwestern
 957 Siberian platform, *Geological Magazine*, 144, 609–618, <https://doi.org/10.1017/S0016756807003354>,
 958 2007.

959 Landing, E. and Kruse, P. D.: Integrated stratigraphic, geochemical, and paleontological late Ediacaran to
 960 early Cambrian records from southwestern Mongolia: comment, *Bulletin*, 129, 1012–1015,
 961 <https://doi.org/10.1130/B31640.1>, 2017.

962 Landing, E., Bowring, S. A., Davidek, K. L., Westrop, S. R., Geyer, G., and Heldmaier, W.: Duration of
 963 the early Cambrian: U–pb ages of volcanic ashes from Avalon and Gondwana, *Canadian Journal of Earth*
 964 *Sciences*, 35, 329–338, 1998.

965 Landing, E., Schmitz, M. D., Geyer, G., Trayler, R. B., and Bowring, S. A.: Precise early cambrian u–pb
 966 zircon dates bracket the oldest trilobites and Archaeocyaths in Moroccan West Gondwana, *Geological*
 967 *Magazine*, 158, 10.1139/e97-107, 219–238, 2021.

968 Langereis, C. G., Krijgsman, W., Muttoni, G., Menning, M., et al.: Magnetostratigraphy-concepts,
 969 definitions, and applications, *Newsletters in Stratigraphy*, 43, 207, 10.1127/0078-0421/2010/0043-0207,
 970 2010.

971 Lantink, M. L., Davies, J. H., Ovtcharova, M., and Hilgen, F. J.: Milankovitch cycles in banded iron
 972 formations constrain the Earth–Moon system 2.46 billion years ago, *Proceedings of the National*
 973 *Academy of Sciences*, 119, e2117146119, <https://doi.org/10.1073/pnas.2117146119>, 2022.

974 Lee, T., Rand, D., Lisiecki, L. E., Gebbie, G., and Lawrence, C. E.: Bayesian age models and stacks:
 975 Combining age inferences from radiocarbon and benthic $\delta^{18}\text{O}$ stratigraphic alignment, *EGUsphere*,
 976 2022, 1–29, <https://doi.org/10.5194/cp-19-1993-2023>, 2022.

977 Lisiecki, L. E. and Lisiecki, P. A.: Application of dynamic programming to the correlation of
 978 paleoclimate records, *Paleoceanography*, 17, 1–1, 10.1029/2001PA000733, 2002.

979 Magaritz, M., Kirschvink, J. L., Latham, A. J., Zhuravlev, A. Y., and Rozanov, A. Y.:
 980 Precambrian/cambrian boundary problem: Carbon isotope correlations for Vendian and Tommotian time

981 between Siberia and Morocco, *Geology*, 19, 847–850, [https://doi.org/10.1130/0091-7613\(1991\)019%3C0847:PCBPCI%3E2.3.CO;2](https://doi.org/10.1130/0091-7613(1991)019%3C0847:PCBPCI%3E2.3.CO;2), 1991.

983 Maloof, A. C., Schrag, D. P., Crowley, J. L., and Bowring, S. A.: An expanded record of early Cambrian
984 carbon cycling from the Anti-Atlas margin, Morocco, *Canadian Journal of Earth Sciences*, 42, 2195–
985 2216, <https://doi.org/10.1139/e05-062>, 2005.

986 Maloof, A. C., Ramezani, J., Bowring, S. A., Fike, D. A., Porter, S. M., and Mazouad, M.: Constraints on
987 Early Cambrian carbon cycling from the duration of the Nemakit-Daldynian–Tommotian boundary $\delta^{13}\text{C}$
988 shift, morocco, *Geology*, 38, 623–626, <https://doi.org/10.1130/G30726.1>, 2010.

989 Miall, A. D.: Updating uniformitarianism: Stratigraphy as just a set of “frozen accidents,” *Geological*
990 *Society, London, Special Publications*, 404, 11–36, <https://doi.org/10.1144/SP404.4>, 2015.

991 Middleton, J. L., Gottschalk, J., Winckler, G., Hanley, J., Knudson, C., Farmer, J. R., Lamy, F., and
992 Lisiecki, L. E.: Evaluating manual versus automated benthic foraminiferal $\delta^{18}\text{O}$ alignment techniques for
993 developing chronostratigraphies in marine sediment records, *Geochronology*, 6, 125–145,
994 <https://doi.org/10.5194/gchron-6-125-2024>, 2024.

995 Paterson, J. R., Edgecombe, G. D., and Lee, M. S.: Trilobite evolutionary rates constrain the duration of
996 the Cambrian explosion, *Proceedings of the National Academy of Sciences*, 116, 4394–4399,
997 <https://doi.org/10.1073/pnas.1819366116>, 2019.

998 Ramsey, C. B. Radiocarbon calibration and analysis of stratigraphy: the OxCal
999 program. *Radiocarbon*, 37(2), 425–430, <https://doi.org/10.1017/S0033822200030903>, 1995.

1000 Roberts, G. O. and Rosenthal, J. S.: Examples of adaptive MCMC, *Journal of computational and*
1001 *graphical statistics*, 18, 349–367, <https://doi.org/10.1198/jcgs.2009.06134>, 2009.

1002 Rudman, A. J. and Lankston, R. W.: Stratigraphic correlation of well logs by computer techniques,
1003 *AAPG Bulletin*, 57, 577–588, 1973.

1004 Sadler, P. M.: Sediment accumulation rates and the completeness of stratigraphic sections, *The Journal of*
1005 *Geology*, 89, 569–584, <https://www.jstor.org/stable/30062397>, 1981.

1006 Sakoe, H. and Chiba, S.: Dynamic programming algorithm optimization for spoken word recognition,
1007 *IEEE transactions on acoustics, speech, and signal processing*, 26, 43–49, 10.1109/TASSP.1978.1163055,
1008 1978.

1009 Saltzman, M., Thomas, E., and Gradstein, F.: Carbon isotope stratigraphy, *The geologic time scale*, 1,
1010 207–232, <https://doi.org/10.1016/B978-0-444-59425-9.00011-1>, 2012.

1011 Sambridge, M.: A parallel tempering algorithm for probabilistic sampling and multimodal optimization,
1012 *Geophysical Journal International*, 196, 357–374, <https://doi.org/10.1093/gji/ggt342>, 2014.

1013 Schmitz, M., Gradstein, F., Ogg, J., et al.: Radiogenic isotope geochronology, *The geologic time scale*, 2,
1014 115–126, <https://doi.org/10.1016/B978-0-444-59425-9.00006-8>, 2012.

1015 Sinnesael, M., Millard, A. R., and Smith, M. R.: A Bayesian astrochronology for the Cambrian first
1016 occurrence of trilobites in West Gondwana (Morocco), *Geology*, 52, 205–209,
1017 <https://doi.org/10.1130/G51718.1>, 2024.

1018 Smith, E. F., Macdonald, F. A., Petach, T. A., Bold, U., and Schrag, D. P.: Integrated stratigraphic,
1019 geochemical, and paleontological late Ediacaran to early Cambrian records from southwestern Mongolia,
1020 *Geological Society of America Bulletin*, 128, 442–468, <https://doi.org/10.1130/B31248.1>, 2016.

1021 Sylvester, Z.: Automated multi-well stratigraphic correlation and model building using relative geologic
1022 time, *Basin Research*, 35, 1961–1984, <https://doi.org/10.1111/bre.12787>, 2023.

1023 Trayler, R. B., Meyers, S. R., Sageman, B. B., and Schmitz, M. D.: Bayesian integration of
1024 astrochronology and radioisotope geochronology, *Geochronology*, 6, 107–123,
1025 <https://doi.org/10.5194/gchron-6-107-2024>, 2024. Tucker, M. E.: Carbon isotope excursions in
1026 Precambrian/Cambrian boundary beds, Morocco, *Nature*, 319, 48–50, <https://doi.org/10.1038/319048a0>,
1027 1986.

1028 Uhlein, G. J., Uhlein, A., Pereira, E., Caxito, F. A., Okubo, J., Warren, L. V., and Sial, A. N.: Ediacaran
1029 paleoenvironmental changes recorded in the mixed carbonate-siliciclastic Bambuí basin, Brazil,
1030 *Palaeogeography, palaeoclimatology, palaeoecology*, 517, 39–51,
1031 <https://doi.org/10.1016/j.palaeo.2018.12.022>, 2019.

1032 Varlamov, A., Rozanov, A. Y., Khomentovsky, V., Shabanov, Y. Y., Abaimova, G., Demidenko, Y. E.,
1033 Karlova, G., Korovnikov, I., Luchinina, V., Malakhovskaya, Y. E., et al.: The Cambrian system of the
1034 Siberian platform. Part 1: The Aldan-Lena region, in: XIII field conference of the Cambrian stage
1035 subcommission working group, 2008.

1036 Vats, D. and Knudson, C.: Revisiting the Gelman–Rubin diagnostic, *Statistical Science*, 36, 518–529,
1037 <https://doi.org/10.1214/20-STS812>, 2021.

- 1038 Vats, D., Flegal, J. M., and Jones, G. L.: Multivariate output analysis for Markov chain Monte Carlo,
1039 *Biometrika*, 106, 321–337, <https://doi.org/10.1093/biomet/asz002>, 2019.
- 1040 Vousden, W., Farr, W. M., and Mandel, I.: Dynamic temperature selection for parallel tempering in
1041 Markov chain Monte Carlo simulations, *Monthly Notices of the Royal Astronomical Society*, 455, 1919–
1042 1937, <https://doi.org/10.1093/mnras/stv2422>, 2016.
- 1043 Wheeler, L. and Hale, D.: Simultaneous correlation of multiple well logs, in: SEG international
1044 exposition and annual meeting, SEG–2014, <https://doi.org/10.1190/segam2014-0227.1>, 2014.
- 1045 Yang, J. and Rosenthal, J. S.: Automatically tuned general-purpose MCMC via new adaptive diagnostics,
1046 *Computational Statistics*, 32, 315–348, <https://doi.org/10.1007/s00180-016-0682-2>, 2017.
- 1047 Zhang, X., Ahlberg, P., Babcock, L. E., Choi, D. K., Geyer, G., Gozalo, R., Hollingsworth, J. S., Li, G.,
1048 Naimark, E. B., Pegel, T., et al.: Challenges in defining the base of Cambrian series 2 and stage 3, *Earth-*
1049 *Science Reviews*, 172, 124–139, <https://doi.org/10.1016/j.earscirev.2017.07.017>, 2017.