Response to the comments of Reviewer 1 (Maarten Blaauw)

We thank the reviewer for his time and helpful feedback. Our responses to the reviewer’s comments (black) are recorded below in blue.

Reviewer 1

This manuscript proposes a new method to align multiple proxy records based on assumed synchroneity (e.g. appearance of key trilobite fossils); additional data such as radiometric dates or known ages of fossils can also be added. The model draws a Bayesian cubic spline (Heaton et al., 2020) per to-be-aligned proxy, using evenly-spaced knots and smoothness parameters. The model is applied to some synthetic and real-world examples.

I like the fact that not just one alignment is chosen, displayed and discussed, but a range of alignments (e.g., Fig. 6 and section 5.1). This clearly shows the probabilistic and uncertain nature of aligning multiple records, and thus the need and potential for a Bayesian framework. Could the age-depth relationships of the three solutions from Fig. 6 also be shown in a Figure akin to Fig. 7, to see how variable the reconstructed rates and hiatuses are?

We thank the reviewer for this positive assessment and have included an additional figure in Appendix B to show the age-depth relationships of the three different solutions from Fig. 6. (i.e. Fig. B2).

Sometimes stratigraphical correlation is the only way to obtain a chronology for a proxy record, e.g. where no absolute/radiometric age estimates are available. However, it would be good to also highlight potential problems with aligning records based on their assumed synchroneity, e.g. problems with circular reasoning, possible erroneous choice of tie-points, and the introduction of a dependence between records. These problems are reviewed by Blaauw 2012 (doi:10.1016/j.quascirev.2010.11.012).

We expanded the discussion in section 5.2.3 to point out those challenges:

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. δ^{13}C) are conditional on the assumption of synchronicity.

Line 76, would it be useful to mention Trayler et al. 2024's Astrobayes age-model, which includes hiatuses (doi:10.5194/gchron-6-107-2024)?

Reference added in lines 68-70: *A Bayesian age-depth modelling approach by Trayler et al. (2024) considers hiatuses and uses astrochronological interpretations to inform sedimentation rate priors.*

Lines 228-32 and 646-52 list an important limitation of the proposed model; assumed linear sedimentation rates will not cause chronological uncertainties to widen further away from age constraints. Some of the reconstructed age-model uncertainties seem very narrow indeed, e.g. 7d. Does setting spline knots at regular intervals not help?

The age model uncertainties at 7d (Siberia) are low because of the high-amplitude variations in δ^{13}C which are matched with similar high-amplitude signals at Oued Sdas, leading to comparatively low uncertainties. Spline knots are placed at regular intervals, but for age uncertainties to widen away from age constraints, additional sedimentation rate changes would have to be included (see discussion in section 5.2.3, lines 672 – 690).

For a frequently-used Bayesian age-depth model that includes priors on sedimentation rates and variability, please cite Bacon (Blaauw & Christen 2011, https://projecteuclid.org/euclid.ba/1339616472). Bacon is a piece-wise linear model much like what is proposed here; it also includes time hiatuses, slumps (depth 'hiatuses') and changes in sedimentation rates. It uses the t-walk, a flexible MCMC (Christen & Fox 2010, http://projecteuclid.org/euclid.ba/1340218339). Although Bacon is most often used on radiocarbon-dating timescales, it has also been applied to much longer time-scales. That said, the usage of dozens of parameters per site (owing to long cores with thin sections) would probably cause the MCMC to run much, much slower than the 5 days reported here.

We thank the reviewer for pointing out this oversight and have expanded the introduction to refer to Bacon and other Bayesian age-depth modelling tools (lines 65 – 68): *Bayesian approaches are commonly employed in age-depth models that interpolate between absolute age constraints or tie points, e.g. Bchron (Haslett and Parnell, 2008) and Oxcal (Ramsey 1995); Bacon (Blaauw and Christen 2011) also includes priors on sedimentation rates and variability.*

I ran a quick toy age-model in R using the vignette provided and all ran fine. This is important, because other recently proposed methods I've seen rely on many additional packages and on software external to R such as JAGS to run (often resulting in failure). Pity though that only binary versions are provided - could the source c++ code also be provided? That would enable users on other operating platforms to also run the code, would enable users to get a better idea of what exactly is done, and would be much more future-proof.

We are glad that the reviewer found that the package ran smoothly. We agree that releasing the source code would be beneficial, and plan to do so once we have more clarity on the long-term direction of this software project.

The MCMC runs multiple chains but only retains the samples from one chain (both a burn-in and thinning are applied afterward). Is this a standard approach?

The discarded chains are tempered chains; we are using a parallel tempering framework for easier sampling of multimodal posteriors. We have expanded the explanation in lines 136 – 139 to clarify this:

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, and frequently propose swaps between chains. For the posterior estimates, we retain samples only from the primary (cold) chain.

Could you clarify μ in section 2.1: is this a hypothetical target to which all sites are tuned, or is this akin to target/reference Site 1 as in Fig. 1?

μ in section 2.1 is a hypothetical target, to which all sites except the reference site are tuned (the reference site remains fixed in this scenario, but is informing μ alongside the tuned sites). We have added a clarification in section 2.1 (avoiding the term "target" and referring to a hypothetical composite curve instead): *Here, μ can be interpreted as an underlying common signal of which the observations from each site, including the reference site, are noisy realisations.* (lines 154 – 155)

Fig. 1 of the hypothetical sample: can the α and γ values of the placement in c) be depicted as vertical lines overlying the prior distributions of panel b)? This because in this example, site 2 is compressed a lot (2.8 times faster than site 1), and it would be nice to see where it falls on the log-normal prior (as well as of course the placement on the uniform prior, 12.5 m). In this example, site 2 accumulates linearly over time.

This is a helpful improvement of Fig. 1, we've added the reviewer's suggestion to the updated version of Fig. 1.

Eq. 8, shouldn't the hiatuses δ be expressed as gaps/jumps in time, not depth/height?

In Eq. 8, correlation is done on a reference height (or reference depth) scale, rather than a time scale, so hiatuses are expressed as an interval of heights on the reference scale that is not represented in the correlated section. A clarification has been added to section 2.2.3 (lines 216 – 217): *In a correlation on an absolute age scale (Sect. 2.2.5), hiatuses would instead be expressed as durations, not heights.*

Citation: https://doi.org/10.5194/egusphere-2025-1355-RC1