



Technical Note: The Enhanced Controlled Random Search (CRS) Algorithm for Thermal History Analysis in HeFTy

Kendra E. Murray¹, Richard A. Ketcham², Andrea L. Stevens Goddard³

¹Department of Geosciences, Idaho State University, Pocatello, ID, 83209, USA

5 ²Jackson School of Geosciences, University of Texas at Austin, Austin, TX, 78712, USA

³Department of Earth and Atmospheric Sciences, Indiana University, Bloomington, IN, 47408, USA

Correspondence to: Kendra E. Murray (kendramurray@isu.edu)

Abstract. The thermal history modeling software HeFTy is one of several widely used tools for the numerical analysis of low-temperature thermochronologic data. HeFTy version 2 includes a new optional Controlled Random Search (CRS) strategy for posing the candidate time-temperature (tT) paths that are tested against data. Unlike the Monte Carlo (MC) strategy that is the default algorithm, the CRS procedure attempts to converge toward a solution. In order to overcome known limitations of CRS convergence (e.g., returning an artificially narrow range of tT solutions that exaggerates the capacity of the data to document a specific thermal history), this ‘enhanced’ CRS randomizes, expands, and reconverges on candidate tT-paths. Here, we use both synthetic and real apatite and zircon (U-Th)/He data, previously analysed using the MC algorithm in HeFTy, to explore the utility of this new tool. Overall, we demonstrate that the enhanced CRS can substantially improve computation time and is useful for finding families of tT paths that fit thermochronologic data. However, users should still expect the CRS to produce an uneven distribution of paths (i.e., ‘clustering’) in the permissible tT space. Therefore, we suggest that systematically performing multiple CRS and/or MC inversions is essential for building a robust assessment of how CRS inversion results are produced by the data and inversion design choices, or else it may not be clear what features of a CRS result are produced by the convergence algorithm rather than the data and a sample’s geologic context. Diagnosing such behaviors is essential for using thermal history inversion results for geological interpretations.

1 Introduction

One of the essential components of any software used for inverse thermal history analysis is the algorithm used to search for candidate time-temperature (tT) histories. Since its first version (Ketcham, 2005), the program HeFTy has used a simple Monte Carlo (MC) search that generates a large number (tens to hundreds of thousands) of independent tT paths in the permissible tT space defined by the user; these candidate thermal histories can be complex, with many degrees of freedom in tT space. Therefore, the MC approach embodies the principle that geological histories are likely to be complex, and the purpose of an inversion result is to map out the uncertainty that arises from this potential complexity.



30 This MC approach contrasts philosophically—and practically—with optimization algorithms, which iteratively converge on a
best-fitting solution. Such converging algorithms, which have been used in thermal history analysis for decades (e.g., Lutz and
Omar, 1991; Zeitler, 1993; Gallagher and Sambridge, 1994; Willett, 1997; Gallagher, 2012; McDannell and Issler, 2021),
learn from how well the already-attempted tT paths fit the data. Therefore, the search of tT space is optimized because it tends
to be more computationally efficient than a simple (i.e., non-learning) MC approach. The efficiency of optimization algorithms
35 also arises from the fact that they are commonly paired with a philosophy of parsimony (i.e., a preference for the simplest
solution) because when presented with many degrees of freedom, they can get trapped in local solutions. Indeed, the tendency
of convergence algorithms to produce artificially narrow ranges of tT solutions, and thereby exaggerate the capacity of
thermochronologic data to resolve thermal history information, prompted Ketcham (2005) to not include a converging
Controlled Random Search (CRS) in HeFTy, even though it was implemented in its precursor (AFTSolve; Ketcham et al.,
40 2000). Occasionally, converging and non-converging approaches have been presented as two philosophies in conflict (e.g.,
Vermeesch and Tian, 2014; Gallagher and Ketcham, 2018). However, from our perspective, it is advantageous to be fluent in
the abilities and limitations of both approaches.

Given the potential synergy between learning and non-learning algorithms, the expanding scope of the questions geoscientists
45 are posing in time-temperature space (e.g., Gautheron and Zeitler, 2020; McDannell and Flowers, 2020), and the increasing
prevalence and computational burden of multi-sample, multi-thermochronometer, and multi-aliquot studies (e.g., Hestnes et
al., 2024; Mackaman-Lofland et al., 2024), Ketcham (2024) updated and re-implemented an optional CRS algorithm in HeFTy
version 2. Here, we expand upon the description and demonstration of the HeFTy CRS in Ketcham (2024), and document
changes made since then.

50 Our first CRS examples are inversions of synthetic apatite (U-Th)/He (AHe) data, which we produced by forward modeling a
specific tT path. Such synthetic data are particularly instructive in assessing thermal history analysis tools (e.g., Murray et al.,
2022; Abbey et al., 2023) because the ‘true’ thermal history is known, various types of data can be predicted from a true
history, and it is simple to systematically explore how different data handling and model design choices (such as data inputs,
55 constraint box design, algorithm parameters) control the capacity of the CRS to find the ‘true’ history. Then, we present new
CRS inversions of published AHe and zircon (U-Th)/He (ZHe) data previously modeled using the MC algorithm, to consider
the utility of the CRS when exploring deep-time thermal histories using data from Paleoproterozoic crystalline rocks. Such
deep-time inversions commonly require long MC searches (millions of attempted paths) to find solutions, so the potential for
the CRS to make such analyses substantially more efficient is one motivation for using this tool. All examples utilized HeFTy
60 version 2.3.

Overall, these examples demonstrate the CRS returns solutions efficiently compared to the MC procedure. However, CRS
inversion results commonly include clusters of solutions in the permissible tT space, even though the CRS procedure includes



65 settings that limit such behavior (see Section 2). In other words, some solution clustering is inevitable; it should be expected and interpreted with caution. We demonstrate that this solution clustering arises from the fundamental trade-offs between time and temperature that are inherent in thermochronology, and how individual CRS solutions can magnify the consequences of these trade-offs by finding and oversampling specific families of solutions. We discuss not only how to identify CRS solution clustering but also how to leverage clustered solutions in order to find families of tT paths that can help one build a robust thermal history analysis.

70 **2 The Enhanced CRS Procedure**

The CRS procedure in HeFTy version 2 is “enhanced” with four Stages that are designed to overcome the tendency of optimization algorithms to over-converge, or leave parts of tT space underexplored. This procedure is described in detail by Ketcham (2024), but here we provide a brief summary and an overview of moderate changes that were more recently implemented in HeFTy version 2.3.

75 **2.1 Goodness-of-fit Criteria**

The combined goodness-of-fit (GOF) assessment used in the CRS procedure is the same as the default HeFTy MC algorithm, with one optional exception: the user may limit the minimum GOF that any single component of the thermochronometric data being fitted can have when combined probability is calculated (Ketcham, 2024, section 4.5). Enforcing a minimum GOF can help overcome the potential of one low-probability datapoint to dominate the combined GOF and prevent convergence. This 80 can happen if uncertainties are small, or if a remarkably specific resetting history is required to produce an age. Otherwise, the criteria and verbal descriptors are the same as the default MC algorithm: ‘good-fit’ ($GOF \geq 0.5$) and ‘acceptable-fit’ ($0.05 \leq GOF < 0.5$).

2.2 Stages of the CRS Procedure

Stage 1 uses a standard CRS algorithm to find a ‘set’ of 250 tT paths that have good or acceptable fits to the data. First a set 85 of 250 random paths are generated, and then random 5% subsets of these paths are reflected, averaged, and compared; better-fitting paths replace worse-fitting paths in the set until the ending condition is met: 20% of paths in the set have good fits, or all paths in the set have acceptable fits. A scaling factor (α , assigned a value of 1.7 by default) controls the reflection magnitude. As expected, at the end of Stage 1 the solution is over-converged compared to a full MC solution (e.g., Ketcham, 2024, Figure 4A). Stages 2-4 are designed to expand this result.

90

Stage 2 permits a radical re-exploration of independent parts of the solution space (i.e., the path segments between constraints) while preserving the top 1% of the best fitting paths from the CRS set produced in Stage 1. A random path from the CRS set is selected, and a portion of that path between two constraints is randomly expanded, while still conforming to any user-



imposed rules. If the resulting new path has an acceptable fit to the data, it is added to the CRS set, and then the path just *below*
95 the top 1% of best-fitting paths is removed from the set. Unlike in Stage 1, this “reverse-ordered” operation commonly results
in worse-fitting paths replacing better-fitting paths in the CRS set (though all are retained in the solution set). Thus, Stage 2
produces a revised CRS set that is expanded toward the margins of the solution space. The user specifies the number of tries
per path segment (default $n=1000$ in HeFTy v2.3).

100 Stage 3 continues to broaden the solution space, but it does so more gradually. It uses the standard CRS selection and reflection
procedure from Stage 1 with $\alpha = 4$ and the reverse-ordered path replacement operation from Stage 2. This incrementally
expands the edges of solution clusters in tT space. The user specifies the number of iterations ($n=2000$ in HeFTy v2.3).

Stage 4 reconverges using the standard CRS algorithm (with default $\alpha = 2$) and the CRS set that was expanded by Stages 2
105 and 3. Better-fitting paths replace worse-fitting paths as in Stage 1. This procedure fills in, while still incrementally expanding,
the solution space. The user specifies the number of iterations (default $n=2000$ in HeFTy v2.3). Finally, if Stage 1 did not find
any good-fit paths, but some were returned in Stages 2-4, a 5th stage is added in which the standard CRS convergence continues
until 20% of the CRS set has good fits to the data.

2.3 Updates

110 Two changes were made to the enhanced CRS procedure in HeFTy for version 2.3 used for the exercises/examples in this
study. These updates improve the procedure’s ability to find and converge on a solution set in situations when the algorithm
rarely returns good or acceptable fits after the maximum number of attempted paths.

First, limiting minimum combined probability (Ketcham, 2024; section 4.5) was made the default. No circumstances have yet
115 been identified in which it harms the search procedure, and it certainly helps in some cases.

Second, a ‘reboot’ condition was added to Stage 1 to aid in cases where no good or acceptable cases are found, or if no progress
is made after finding a few qualifying paths. If a large number of paths (corresponding to a given multiple, currently set to 21,
of the CRS set size) are generated without finding a new good or acceptable one, the reboot preserves the best-fitting 2% of
120 paths in the CRS set and re-randomizes the rest. This measure counteracts the occasional tendency of the CRS set to collapse
into a set of similar paths that surround local minima while bypassing large regions of the true solution space.

With this ‘reboot’ addition, the enhanced CRS routine does not report failure unless no good or acceptable paths are found
after one reboot. If even one such path is found, the procedure will continue searching and rebooting until a convergence
125 criterion is met. Whereas the previous CRS version could potentially run for an infinite amount of time if one or two acceptable



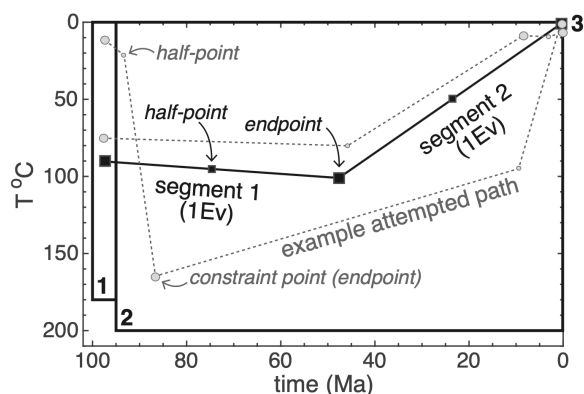
paths were found but the rest of the CRS set had an unfavorable arrangement, we have observed no such failure in the new version.

3 Examples

3.1 Synthetic Data

130 3.1.1 Single 40 Ma AHe Crystal

A single cooling age with little geologic context can be fit by a very wide range of thermal histories. Therefore, a single synthetic AHe age provides an instructive demonstration of the fundamental nonuniqueness of thermochronologic data (Wolf et al., 1998) and how this nonuniqueness manifests in the solutions generated by various thermal history analysis tools (e.g., Murray et al., 2022; Abbey et al., 2023). Here, the results of CRS inversions of a single 40 ± 4 Ma age of an apatite crystal
 135 (with 60 ppm U and a 60 μm radius, as in Murray et al., 2022 and Abbey et al., 2023) reveal how the enhanced CRS procedure performs when data can be fit with a wide range of tT solutions. Using identical inversion designs (Fig. 1), we compare several CRS solutions to a MC solution (Fig. 2). We also use this simple data to demonstrate how deviating from the default CRS parameters changes the inversion result.



140 **Figure 1. Our synthetic examples use identical inversion designs: two constraint boxes (labelled 1 and 2), plus a third constraint (labelled 3) of $T = 5 \pm 5^\circ\text{C}$ at 0 Ma, which together generate a wide range of monotonic (only cooling) and nonmonotonic (heating then cooling) thermal histories over 100 Myr (Murray et al., 2022). There are two independent segments, which each have one half-point and episodic, variable behaviour (“1Ev”). The black line that connects the segments also visually links the constraint boxes, as in the HeFTy interface in inversion mode. However, this is a simple geometric line that is not representative of a specific attempted tT history. The grey dotted lines, and the associated constraint- and half-points, are two example paths that would be posed by this design. Visualizing just the constraint points of an inversion result can reveal key features of a solution, including CRS clustering, but points alone lack the time-integrated context of the full tT paths.**

The enhanced CRS (v2.3) produces solutions that return a similar range of tT solutions to a MC inversion with identical model design, but more efficiently: a MC search finds 1 good path per ~ 23 attempts (Fig. 2a), whereas a CRS commonly returns a
 150 similar number of good paths in less than one third of the time (Fig. 2b). For this example, the CRS procedure returned 831 good-fit paths in 6730 attempted paths, whereas the MC procedure returned 831 good-fit paths after 20,845 attempts. The CRS



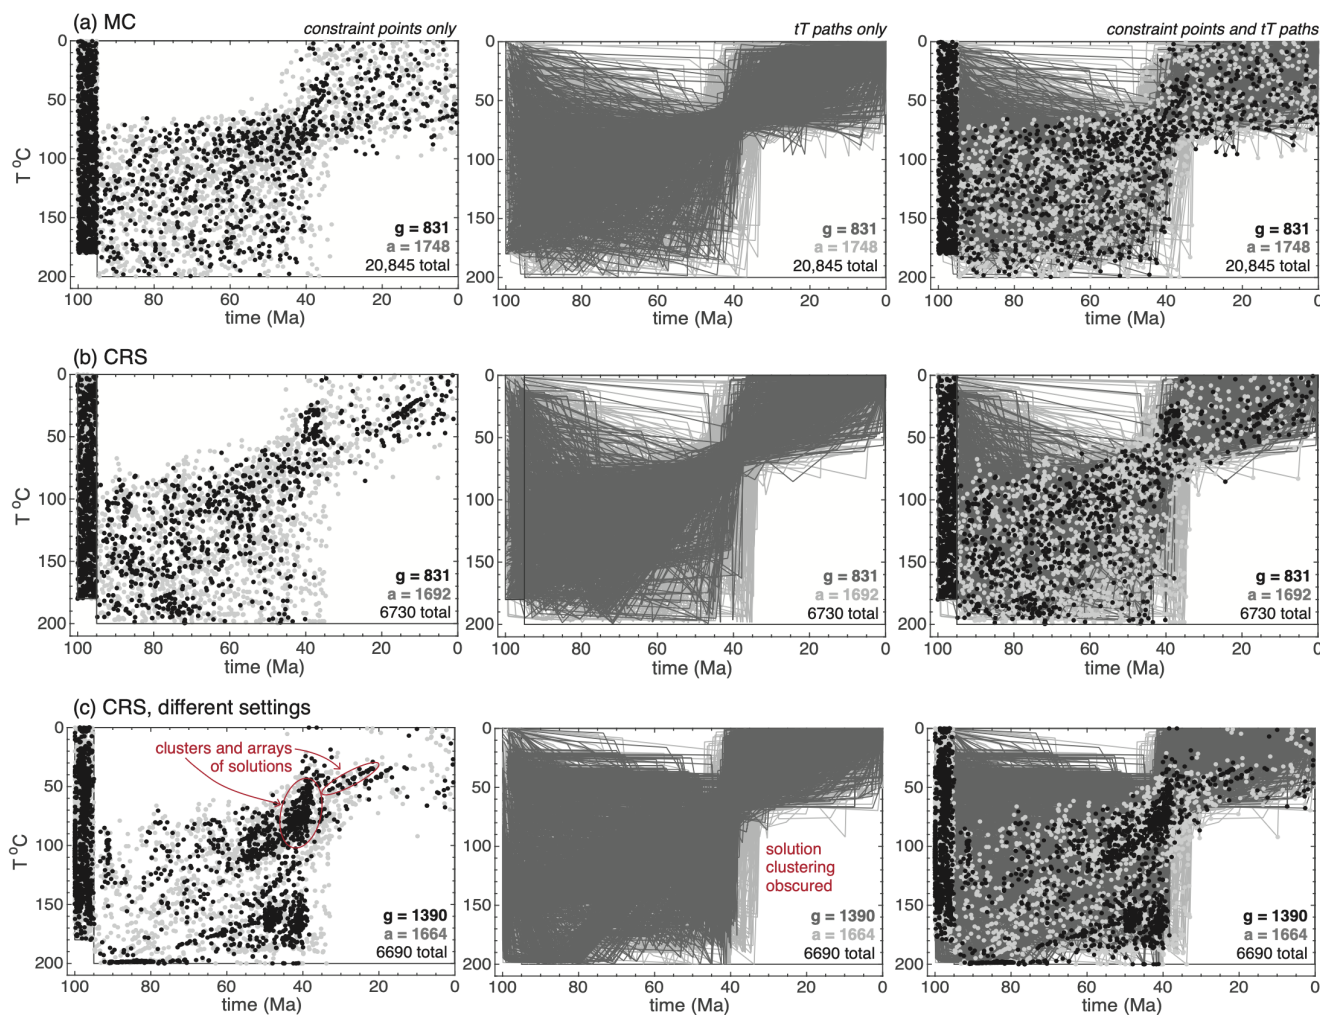
solution has some minor clustering in tT space, but overall the coverage of the tT solution space is similar to the MC solution. Therefore, the enhanced CRS procedure is performing well, i.e., finding a reasonable solution more efficiently than the MC procedure.

155

However, modest changes to the CRS settings can produce clustering in solution space that is unrelated to the resolving power of this single age. Some aspects of the CRS procedure have been updated since the HeFTy v2 was originally released (see section 2.3) and others (such as the α value and default number of iterations for each Stage) can be changed by selecting “Options” in the Inverse Modeling Window. Although our general recommendation is to not adjust these Options, software version control is also important. For example, an identically designed CRS inversion in HeFTy using the default path numbers from HeFTy v2.1, which had half as many tries per independent segment during Stage 2 and twice as many path interactions in Stages 3 and 4, produces a result that has clear clustering in tT space (Fig. 2c). Because this is a single cooling age, the fact that the clustering is a product of the CRS procedure is clear and confirms that recent adjustments to the CRS default settings have improved the inversion results. However, this result also highlights how sensitive some details of a CRS result can be to adjustments to the CRS Options. Thus, deviating from the current default Options should only be done deliberately.

160

165



170 **Figure 2. Inversions of a single 40 Ma AHe age ($[U] = 60$ ppm; radius = 60 μm) demonstrate the performance of the enhanced CRS procedure when data can be fit by a wide range of solutions. Left panels present just the constraint points (i.e., Fig. 1) for the good- and acceptable-fit solutions, which are useful for visualizing any clusters or arrays of solutions produced by the CRS procedure. Middle panels present the corresponding tT paths, which illustrate the time-integrated solution space but can obscure details such as solution clustering. Right panels overlay the constraint points and tT paths. (a) An inversion result produced by the standard MC procedure run to the same number of good-fit paths as the CRS solution. (b) An inversion result produced by the enhanced CRS procedure using the default CRS setting in HeFTy v2.3 (see section 2.2). (c) An inversion result produced by the enhanced CRS procedure using the CRS settings for Stages 2-4 in the originally released version of the enhanced CRS (HeFTy v2.2.1: in Stage 2, 1000 tries per independent segment instead of 2000; in Stages 3 and 4, 2000 iterations instead of 1000).**

3.1.2 A Synthetic Age-[eU] Pattern

In (U-Th)/He chronometers, radiation damage accumulation and annealing can produce distinctive data trends, which in some cases can only be produced by specific thermal histories. These trends are visualized as age-[eU] patterns, where effective U concentration ($[eU] = [U] + 0.238[Th] + 0.0012[Sm]$) is a proxy for radiation damage. HeFTy implements several kinetics



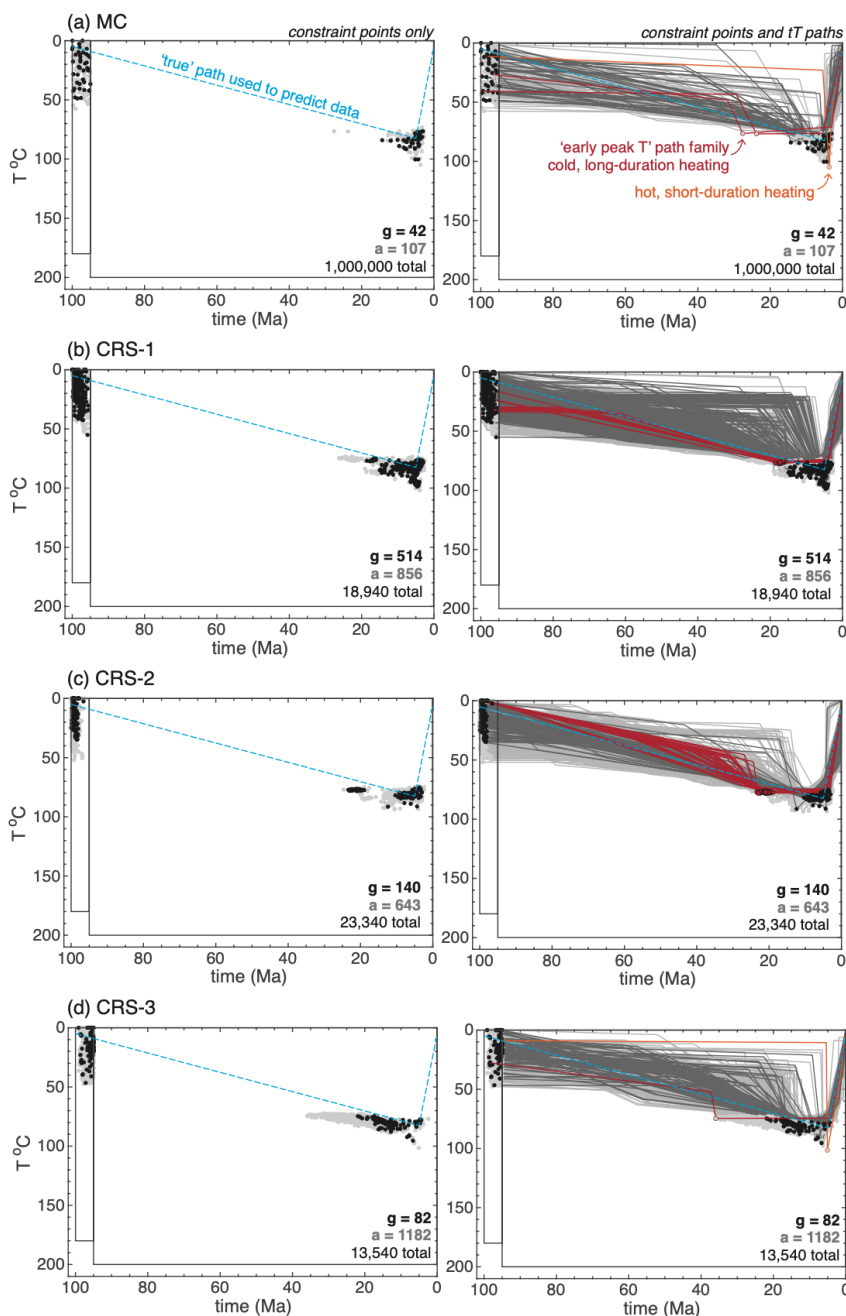
models for radiation damage accumulation and annealing in apatite (Flowers et al., 2009; Gautheron et al., 2009) and zircon (Guenther et al., 2013).

185 Burial histories in which samples reside at colder-than-closure temperatures for tens of millions of years prior to heating to partial-retention temperatures produce particularly distinctive age-[eU] patterns. A hypothetical 100-Myr long burial history—first used by Wolf et al. (1998, their “History 5”) to demonstrate the behavior of the AHe system and adjusted for the RDAAM kinetics (Flowers et al., 2009) by Murray et al. (2022, their “Path 5”)—starts at 0°C at 100 Ma, heats to 82.5°C at 5 Ma, and then cools to 0°C by 0 Ma. The RDAAM predicts an AHe age-[eU] trend with ages of 4-95 Ma from apatite crystals with 10-300 ppm [U].

190 A MC solution from inverting the age-eU trend predicted by Path 5 demonstrates the narrow range of solutions that fit this data trend (Fig. 3a). This is expected given the distinctive age-[eU] pattern. However, the MC procedure found this solution very inefficiently, due to the non-converging algorithm and our purposefully generic search design (Fig. 1); on average, one good-fit path was returned every ~24,000 attempts. Thus, the MC procedure is effective but slow.

195 In contrast, the enhanced CRS procedure returns 2-10 times as many good-fit paths in a fraction of the time (e.g., ~4 hours for MC vs. ~6 minutes for CRS); three repeated CRS inversions with identical model designs and CRS settings, made possible by the short run times, all returned solutions that include the ‘true’ path. However, differences among these inversion results were produced by the CRS algorithm, not the data. Visualizing the results as constraint points (Fig. 3, left panels) reveals the clustering of solutions in tT space. One CRS solution (CRS-1, Fig. 3b) returned 514 good-fit paths and had a weighted mean path very similar to the path used to produce the age-[eU] pattern. Indeed, the CRS-1 solution looks similar to the MC result, with a relatively uniform distribution of good-fit constraint points. In contrast, solutions CRS-2 and CRS-3 (Fig. 3c,d) returned fewer good-fit paths (140 and 82, respectively) but spanned a larger solution space.

200



205 **Figure 3.** Inversions of a synthetic AHe age-[eU] pattern (4-95 Ma, 10-300 ppm [U]) produced by the ‘true’ reheating path shown in
 blue dashed line on each panel. The inversion design is identical to the single-crystal inversions (i.e., Fig. 1) and we used a 10% error
 for the synthetic ages. Left panels present just the constraint points, whereas right panels also include each good- and acceptable-fit
 path. Red and orange paths exemplify the ‘cold and long’ and ‘hot and short’ path families that are discussed in the text. All good-
 fit paths are statistically indistinguishable from the ‘true’ path. (a) An inversion result produced using the MC procedure, which
 210 demonstrates the inefficiency—but effectiveness—of this non-converging approach. (b-d) Repeated identical CRS runs produce



inversion results that cluster differently in the solution space, but the CRS returns two-to-ten times the number of good-fit paths in a fraction of the time.

Whether such clustering in solution space—and the overall differences among CRS inversion results, e.g., Figure 3—matters to a thermal history analysis depends on the motivating questions and how the results are being used to address them. For example, if the motivating question was “Is reheating required?” then the answer is yes, no matter how the data are inverted (i.e., Murray et al., 2022). However, quantitative information about the reheating history supported by these data (e.g., the timing of the peak T, the peak T, when cooling started, etc.) was variably returned by the CRS inversions.

For example, when we analyzed the good-fit constraint points in Box 2 (which represent the time and temperature of peak reheating), we found that the *timing* varied substantially among the CRS results. The good-fit constraint points in both the MC result and similar-looking CRS-1 result (Fig. 3a, b) document peak T at 6 ± 2.4 Ma (range 3-14 Ma) and 6 ± 2.5 Ma (range 3-19 Ma). However, the CRS-2 and CRS-3 results (Fig. 3c, d) had peak T that skewed significantly older: 11 ± 6.9 Ma (range 3-23 Ma) and 13 ± 4.1 Ma (range 4-22 Ma), respectively. In the ‘true’ path used to produce the synthetic data, the peak T was reached at 5 Ma, so what is the significance of these variable CRS solutions that all fit the data equally well?

The variable timing of heating returned by the CRS inversions reflects the real trade-offs between time and temperature that are inherent in thermochronology. These trade-offs can be particularly clear in reheating scenarios like this, where the same data pattern can be fit by hot, short-duration heating or colder, commensurately longer heating (e.g., Reiners, 2009). Examining the constraint-point clusters as full tT paths—and thereby as time-integrated solutions (Fig. 3, right panels)—reveals specific time-temperature trade-offs. Thus, a path family approach (Murray et al., 2022; Stevens Goddard et al., 2023) can be used to demonstrate what is broadly required to fit the data, beyond the result of any single inversion, and thereby leverage the solution clustering that is common in CRS results.

For example, the reason *why* the timing of the peak T varies among these CRS results is because there is a family of paths where a slightly colder ($\sim 75^\circ\text{C}$) peak T was reached 10-30 Myr earlier than the ‘true’ path. In the CRS-2 solution (Fig. 3c) there is a cluster of good-fit paths with constraint points that indicate peak temperatures of $\sim 75^\circ\text{C}$ at ca. 20 Ma. All the tT paths in this family had cold initial temperatures, stayed at $\sim 75^\circ\text{C}$ for ~ 15 Myr, and then cooled to surface conditions starting at ca. 5 Ma. Indeed, an ‘early peak T’ path family was present to varying degrees in all the other inversion results presented in Figure 3 (red paths), including two acceptable-fit outliers in the MC solution (Fig. 3a) and an array of good- and acceptable-fit paths in the CRS-3 solution (Fig. 3d) that extended peak T conditions to as early as 35 Ma and pulled the average peak T timing to be significantly older than the other inversion results. Indeed, the CRS-3 result—despite returning an average of 13 ± 4.1 Ma for the timing of peak T instead of the ‘true’ timing of 5 Ma—was actually the most useful result for defining the permissible range of solutions that include early heating. In all cases, paths with early heating stayed at $\sim 75\text{-}80^\circ\text{C}$ until ca. 5 Ma and then had a similar timing, magnitude, and rate of recent cooling as the ‘true’ path.



245

In contrast, when we calculated the average *temperature* of each good-fit constraint point in Box 2 (i.e., the peak reheating T), we found this was not only consistent among all the inversion results but also within error of the peak T (82.5 °C) in the ‘true’ path used to produce the synthetic data. The average peak T of the good-fit paths was 85 ± 5.5 °C for the MC solution (range 77-100 °C), and 82 ± 5.2 °C (range 75-99 °C), 80 ± 2.9 °C (range 75-91 °C), and 81 ± 3.9 °C (range 75-95 °C) for the CRS solutions. A few acceptable-fit paths (Fig. 3a,d; orange paths) had peak T as hot as 105 °C, with commensurately shorter-duration heating and dT/dt as fast as 285 °C/Myr. However, paths from this hot, shorter-duration family of solutions were rarely found.

250

The contrast between the narrow range of average peak temperatures and the comparatively broad range of peak-T timing we extracted from these inversion results prompts additional questions. Why are the peak temperatures in these inversions much more consistent with each other and the ‘true’ path, whereas the timing of the peak T is more variably returned? In other words, why are the inversions returning the ‘cold, long duration’ family of solutions more extensively than the complementary ‘hot, short-duration’ family of solutions? The idea that the data pattern itself constrains the peak T to <100°C is appealing, but as stated previously, this is a thermochronologic story of trade-offs between time and temperature, modulated by the inversion design.

255

260

In this case, the permissible ‘hot, short-duration’ family of solutions was only marginally represented in inversions in Figure 3 for three reasons that are related to our inversion design (Fig. 1). First, the two-box design is limited in how frequently it attempts tT paths with segments in which heating and cooling is very rapid (>100 °C/Myr) because Box 2 is wide in both the time and temperature dimensions. Second, the ‘hot, short-duration’ solution space is extremely narrow compared to the tT dimensions of Box 2, because a rapid heating event has to be at just the right time to fit to the age-[eU] pattern, as was previously discussed by Murray et al. (2022). Third, the “average, reflect, and relax” method of generating new paths in the CRS algorithm naturally reduces the incidence of high-dT/dt path segments. Thus, our inversion design does not comprehensively explore the permissible ‘hot, short-duration’ family of solutions.

265

270

In one sense, the constraint-box design is biased toward an unstated geological assumption: the process responsible for reheating does not produce extremely rapid changes in temperature. This may be a reasonable assumption in many geological scenarios, such as the burial heating and erosional exhumation envisioned for the ‘true’ path used to generate the data. Indeed, although not done here, in HeFTy a maximum dT/dt can be set in order to enforce a specific maximum heating or cooling rate for any path segment. In the results shown in Figure 3, the agreement between this assumption and the ‘true’ path is why the average peak T for all inversion results consistently fit the ‘true’ peak T so well. However, if a process capable of ‘hot, short duration’ (e.g., hydrothermal or magmatic) heating and cooling was possible (i.e., our ‘true’ path was not Wolf’s Path 5 but

275



instead one with a very short-duration high-T thermal pulse that produces the same data trend), then these inversion results would not provide a useful solution.

280

We produced a simple test of this analysis by adjusting the inversion design to explicitly explore short-duration heating ca. 5 Ma, and indeed the peak T results are correspondingly much hotter (Figure 4). This alternative constraint-box design used three boxes between 6 Ma and 4 Ma to permit heating from $<30^{\circ}\text{C}$ to peak T and then cooling to $<30^{\circ}\text{C}$ during this 2 Myr time interval. The average peak T for a MC solution (Fig. 4a) were $114 \pm 7.4^{\circ}\text{C}$ (range $97\text{-}145^{\circ}\text{C}$), and three identical CRS runs (Fig. 4b-d) had average peak T of $112 \pm 4.2^{\circ}\text{C}$ (range $99\text{-}131^{\circ}\text{C}$), $114 \pm 5.4^{\circ}\text{C}$ (range $101\text{-}132^{\circ}\text{C}$), and $114 \pm 4.2^{\circ}\text{C}$ (range $101\text{-}129^{\circ}\text{C}$). In these solutions, the ‘cold, long duration’ family of solutions was precluded, so the minimum peak temperatures were around 100°C ; shortening the duration of heating even more would require even hotter peak T to fit the data. These inversion results also demonstrate how the CRS performs when the exploration space is very narrow compared to the overall solution space. The CRS solutions still clustered and thus did not fill the solution space as completely as a MC solution—despite returning far more good and acceptable fits to the data.

285
290

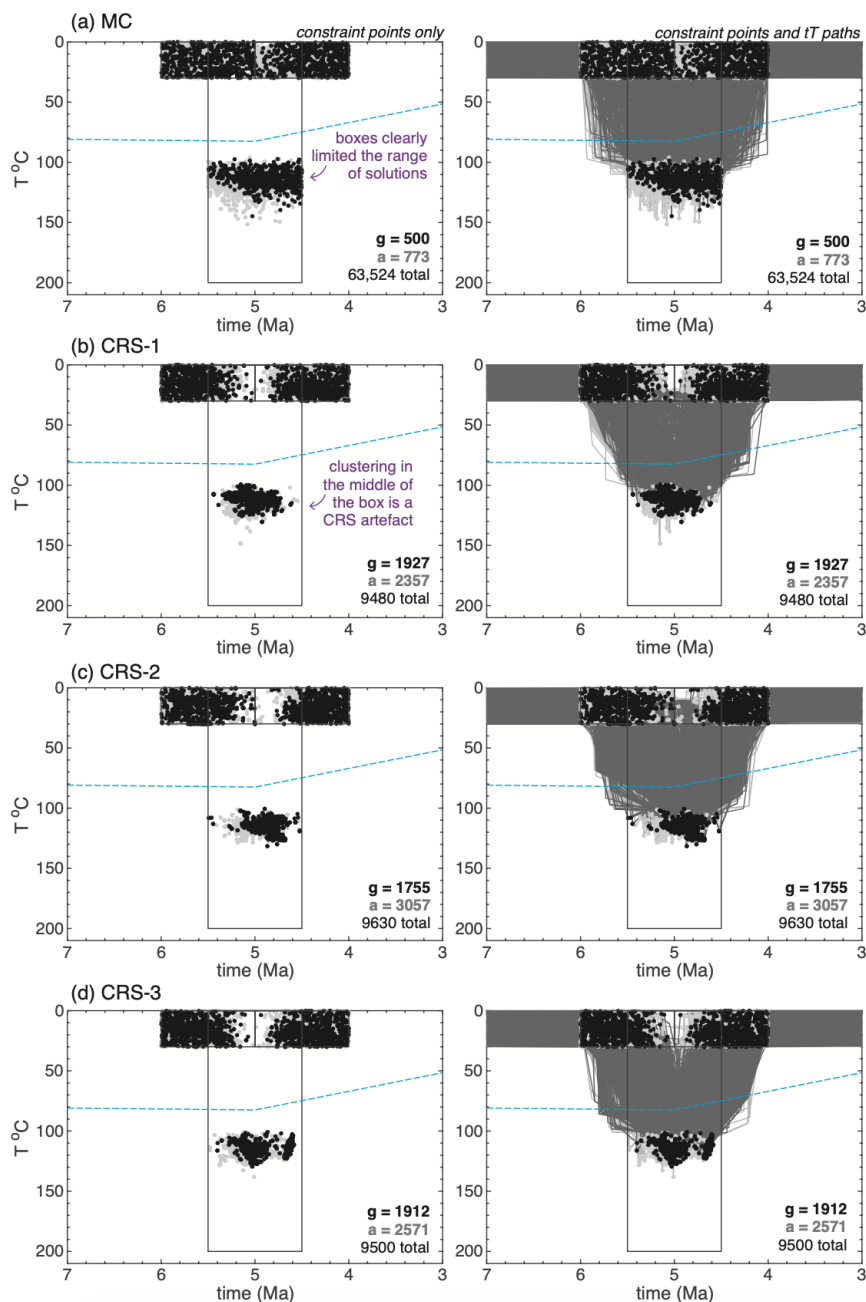


Figure 4. An alternative inversion designed specifically to explore hot, short-duration solutions by restricting heating to $T > 30\text{ }^{\circ}\text{C}$ to occur only between 6 Ma and 4 Ma. Data inputs are identical to the inversions in Figure 3. The blue dotted line is the tT history used to produce the data trend. (a) One MC inversion run to find 500 good-fit paths. This solution uniformly filled the available solution space; because the constraint points sit at the edges of the constraint boxes, it is clear where the boxes limited the range of solutions. (b-d) Repeated identical CRS inversion results. Despite the limited exploration space, the CRS results clustered and did not completely fill the available solution space, so it is less visually clear how the constraint boxes limited the CRS solution. The MC solution aids in diagnosing the source of the clustering in the CRS solutions.

295



If it is geologically reasonable to assume that a sample did not experience such extreme heating and cooling, and therefore the
300 'hot, short duration' family of solutions (Fig. 4) need not be considered, then it would be reasonable to argue that the results
in Figure 3 provided a comprehensive search of solution space. However, the presence of the 'cold, long duration' path family
revealed by solution clustering in some of the CRS results supports that *when* the peak T was reached and *how long* the sample
sat at the peak T is poorly constrained compared to the peak T itself. Thus, if the timing or duration of peak heating is the
motivating question, then additional independent constraints on the sample's history (geologic observations, other
305 thermochronologic data, etc.) are needed to further narrow the solution. Alternatively, the motivating question could be
thoughtfully shifted to a part of the thermal history that appears to be more consistently documented, such as the timing of
cooling below a reference temperature in the AHe partial-retention zone (note, this is not when cooling started, i.e., the
inflection point in each tT path at the constraint point, which is a different and less well-resolved piece of information here).
Overall, the general insights that arise from examining path families provide essential context for using any of the inversion
310 results to make geological interpretations; they are arguably more valuable information than any single quantitative metric
derived from one or an ensemble of inversions.

Finally, we note that the ability of a CRS inversion to find a particular path family depends upon the set of paths found in Stage
1 of the procedure, which is itself a function of how the search of tT space is designed. If a particular path family is found
315 during Stage 1, then it is likely that family of solutions will be iterated on in detail during subsequent CRS stages and will be
represented in the CRS solution. The broader the solution space, the more likely solutions are to be returned. In this example,
we inverted an extremely distinctive age-[eU] pattern using an extremely generic search of tT space, so the solution space was
very narrow. In the MC solution, this manifested in an inefficient but effective solution. In our repeated CRS inversions, this
manifested as occasional failure to find any fits to the data in Stage 1, although recent updates to the CRS procedure (see
320 Section 2.3) have reduced the instances of this outcome. As with the MC procedure in HeFTy, if no CRS solution is returned,
this does not necessarily mean there is a problem with the data inputs or kinetics models. Instead, it could be that the solution
space is very small relative to the search, and there may be a relatively distinctive—and thus geologically useful—solution.
The results of repeated CRS inversions, combined with complementary MC inversions that have 'optimization' constraint
boxes informed by CRS results (Ketcham, 2024), would reveal such a scenario.

325 **3.2 Real Data from a Deep-time Thermal History Analysis**

Computationally demanding inversions—those requiring complex kinetics models, integration over long time periods, or
multiple samples, chronometers, or aliquots—are an obvious target for the CRS algorithm. Deep-time thermochronology
(McDannell and Flowers, 2020) commonly involves all of these elements. In order to explore how the CRS performs in a
deep-time application, we compare a recently published MC inversion of AHe and ZHe data from Paleoproterozoic crystalline
330 rocks (Murray et al., 2025) that has been extensively sensitivity tested (Murray et al., 2022) to new CRS inversions of the same
data (Fig. 5).



All the deep-time inversions in Figure 5 parallel the published inversion. They use the same data: two ~60 Ma AHe ages and five ZHe ages that represent an age-[eU] trend spanning ~50-600 Ma and ~400-2300 ppm [eU]. The inversions start at 700-
335 800°C at 1.7 Ga and cool below 300°C by 1275 Ma. Four of the constraint boxes (Fig. 5a, A, B, C, D) explore time periods of interest using the variable ‘v’ segment parameter to permit, but not require, reheating. The remaining boxes are geological constraints from the rock record. Here, we simply compare the MC and CRS results in tT space.

The published MC inversion result (Fig. 5a) documented a late Neoproterozoic thermal event (constraint points in box C) and
340 peak temperatures in latest Cretaceous time (constraint points in box D). The good- and acceptable-fit paths sparsely but uniformly fill boxes A and B, indicating that there is no record of the thermal history prior to ca. 700 Ma. This is corroborated by all but one point in Box C indicating full resetting of all thermochronometers during late Neoproterozoic time.

As in the synthetic examples, the CRS procedure returned solutions in a fraction of the time of the MC procedure (e.g., ~24-
345 36 hours for MC vs. ~20-40 minutes for CRS). The CRS results feature strong clustering of good solutions in boxes A and B, more non-reset acceptable solutions in Box C, and a higher range of acceptable temperatures in box D. Which of these solution clusters are simply artefacts of the CRS procedure? It is essential to answer this question in order to use this result for geological interpretations.

In this inversion design, we used ‘variable’ segment behavior for the large Mesoproterozoic exploration boxes A and B, which
350 produced a wide range of paths that fall into two general families. In one, paths are cold in box A and reheat in box B. In the other, paths cool monotonically through both boxes. In Figure 5b, the solution clustering in boxes A and B reflects the tendency of the CRS procedure, when confronted with multiple path families that fit the data equally well, to pick one path family over the other rather than sampling them with equal probability. Although the ‘enhanced’ algorithm is capable of avoiding this
355 tendency in some circumstances (e.g., Fig. 2), we find it is prone to this behavior in this deep-time example.

If all thermochronometers were reset during the Neoproterozoic, then any distinction between paths prior to that time (boxes
A and B) is meaningless, and the clustering in those constraints is a model artifact. However, it is possible that the colder, non-
reset points in box C require a particular pre-700 Ma history. Moreover, the relationship between the more-abundant non-reset
360 points in C and the new higher-T points in D is not clear. These latter questions can be addressed by creating two new models that respectively enforce monotonic cooling and reheating in the Mesoproterozoic (Fig. 5c, d). These models both reproduce non-reset points in box C and higher-T points in D, reinforcing that no information can be extracted concerning the Mesoproterozoic. The many high-T points in box D and comparatively few low-T points in box C in Figure 5c indicates that these departures from the main body of solutions are not related to each other.

365



370 An additional insight into CRS model workings can be gleaned from comparing the various results. A cursory comparison of
Figures 5c and 5d might suggest that Mesoproterozoic cooling-only histories are easier to find, and thus somehow more likely
or favorable, than reheating ones. However, it is useful to notice that in Fig. 5d the CRS algorithm found “only” 50 good
paths. This indicates that no good paths were found in Stage 1 and at least one was found in Stages 2-4, leading to Stage 5
being added and running until 20% of the CRS set had good paths. The number of paths found is in this case a reflection of
when the algorithm found them during that particular run, and not a feature of the underlying data.

375 The CRS inversions all point to a possibility of higher Cretaceous temperatures that is allowed by the data but not discovered
in previous MC modeling. The MC approach would have found them eventually if allowed to run long enough, but there was
no indication of such a possibility in the initial 24-36-hour runs.

380 In contrast to the CRS solution artefacts in boxes A and B, the distinctive late Neoproterozoic (box C) and latest Cretaceous
(box D) good-fit solution clusters persist through the various CRS runs. This is because they are a product of the data, kinetics,
and geological constraints, as was previously demonstrated using a suite of MC inversions (Murray et al., 2022; 2025).

385 Clearly, the CRS has strong potential to accelerate parts of deep-time thermal history analysis and even reveal new possibilities.
However, this can come at the cost of adding nonuniform complexity to the solution patterns in each CRS result. Thus, the
tendencies of optimization make it even more essential to rigorously evaluate inversion results—whether they are produced
with learning or non-learning algorithms—with additional thermal history models informed by the fundamentals of
thermochronology.

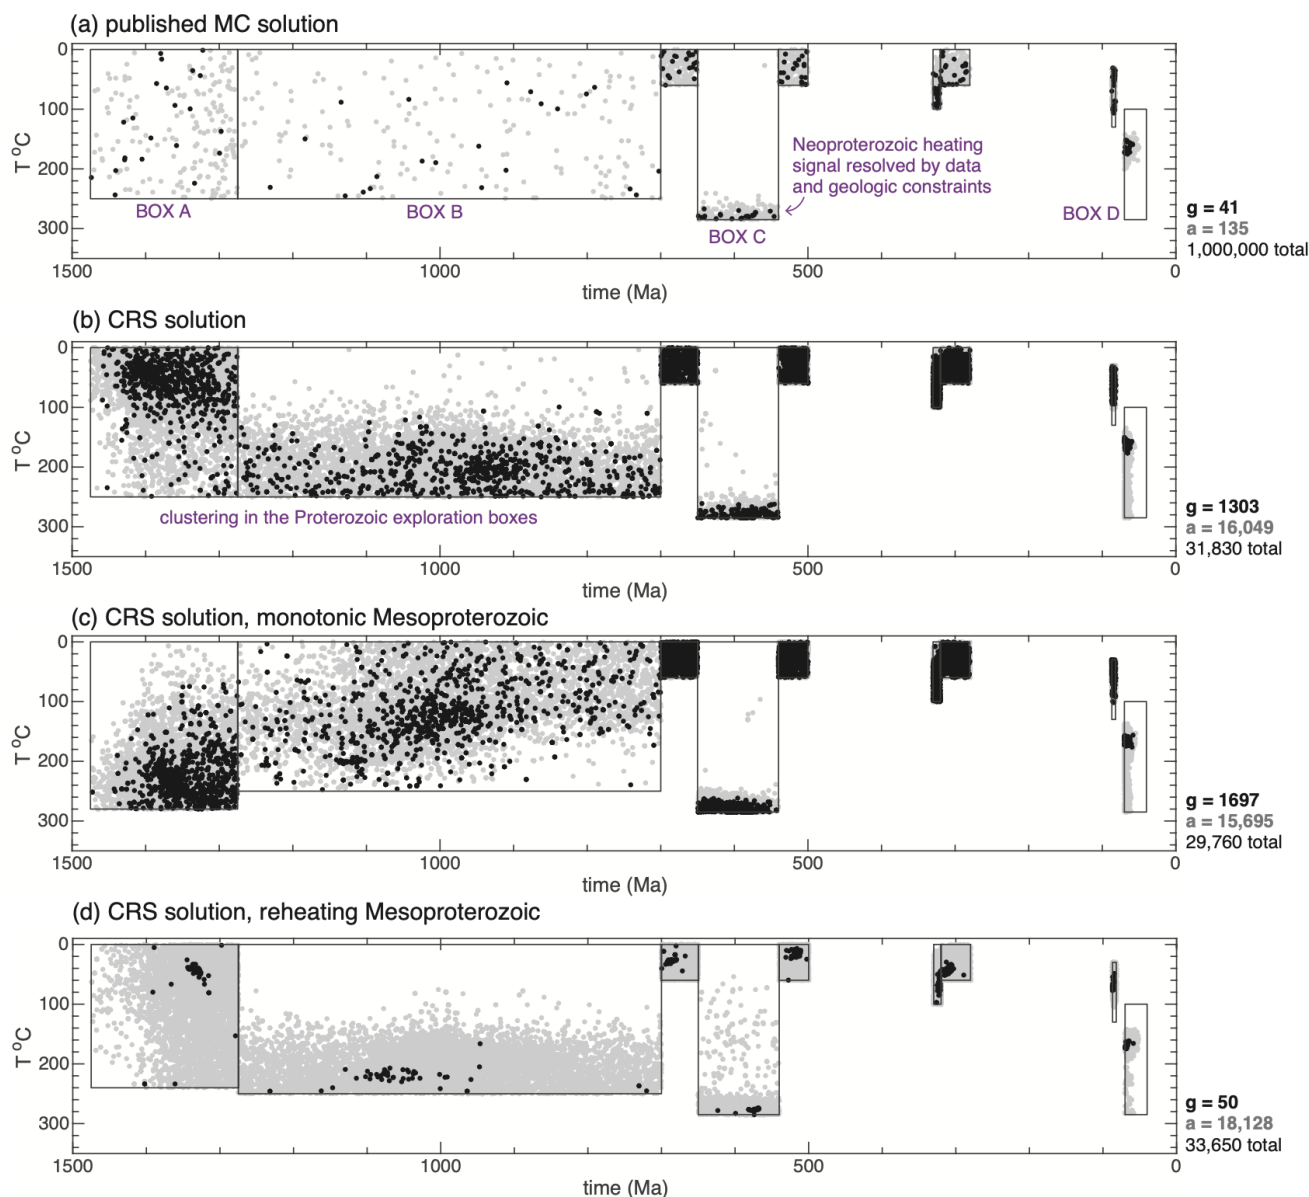


Figure 5. Inversions of AHe and ZHe data from Paleoproterozoic crystalline rocks demonstrate how the CRS performs in an inversion that spans 1.7 Gyr and is designed to investigate Proterozoic thermal histories. Four of the constraint boxes (A, B, C, D) explore time periods of interest. Starting conditions at 1.7 Ga not shown. Additional details about the inversion design, as well as the sensitivity testing and geological significance of the published MC result, can be found in Murray et al. (2022; 2025), but here we are simply interested in comparing the published MC result to new CRS inversions in tT space. Data inputs for all inversions are identical to those used in the published MC result. (a) The published MC result. (b) A CRS inversion result using identical model design as the MC result. Like the MC result, the CRS result includes isolated clusters of solutions in boxes C and D. However, solution clustering is evident in boxes A and B as well. (c) A CRS result with boxes A and B adjusted to only explore monotonic (i.e., always-cooling, no-reheating) Mesoproterozoic histories. (d) A CRS result with boxes A and B adjusted to only explore reheating Mesoproterozoic histories.

390

395



4 Summary and Recommendations

Overall, the enhanced CRS procedure in HeFTy substantially reduces computation time relative to the MC procedure but returns solutions that are less uniform and sometimes less complete, though sometimes more. CRS results should be expected to commonly exhibit clustered distributions of tT paths that are partly artefacts of the optimization algorithm itself. Distinguishing algorithm-driven clustering from clustering genuinely controlled by data, kinetics, and geological constraints is therefore essential. Solution clusters should not be interpreted as geologically meaningful unless they can be demonstrably shown to not be a CRS artefact. That said, CRS clustering can be useful for identifying families of tT paths that reflect the inherent trade-offs between time and temperature.

The advantages of a more efficient algorithm in terms of the human dimension should not be overlooked. The most valuable asset in any scientific endeavor is the time and effort of the scientist, and having to wait hours or days for a result necessarily limits how thoroughly the scientist can investigate and develop a relationship with their data. Conversely, faster execution allows a more interactive modeling approach, enabling the scientist to conceive, pose, and answer questions about their data and models iteratively, aiding their own convergence to a robust solution.

CRS shows strong potential to accelerate parts of deep-time thermal history analysis; however, this can come at the cost of adding nonuniform complexity to solution space that requires careful interpretation. In our deep-time example, CRS inversions ran more efficiently than MC and returned more good- and acceptable-fit paths, while supporting the same geological interpretations as previous work. The CRS results also exhibited substantial clustering of solutions in Proterozoic exploration boxes that would be perilous to interpret geologically without the context provided by additional inversions. Furthermore, the deep-time CRS solutions also returned acceptable-fit alternative path families, which the MC did not completely explore due to computational limitations, providing additional context for the geological interpretations of these results. This illustrates the complementary strengths of the two approaches. We recommend using CRS and MC together: for example, a CRS inversion could efficiently explore solution space and inform the design of a subsequent MC search (Ketcham, 2024), allowing the MC procedure to sample that space uniformly and without algorithmic bias.

As with a MC result, a single CRS result is not sufficient to fully evaluate the robustness of a thermal history interpretation. We recommend that users develop familiarity with CRS behavior on a familiar or synthetic dataset before applying it to new problems; this is the same best practice that applies to using MC for the first time. Users can also help CRS avoid artificially narrow solutions by posing clear, well-constrained questions: narrowing the search of tT space with geological constraints and qualitative data assessment, and testing end-member scenarios—such as reheating versus monotonic cooling—separately rather than simultaneously. Ultimately, understanding the fundamental time-temperature trade-offs inherent to thermochronologic data remains essential regardless of which inversion method is used.



430 **Code and data availability**

The HeFTy files used to produce the inversions in this contribution are provided in the supplementary materials. The data related to the deep-time inversion example are available at <https://doi.org/10.7302/g4eg-vr84>. HeFTy is written in a commercial programming environment. The compiled program is available at <https://doi.org/10.18738/T8/EZ5AA0> (Ketcham, 2025), and is free of charge to academic, not-for-profit users.

435 **Author contributions**

All authors conceptualized this Technical Note. KM carried out the thermal history modelling and prepared the manuscript with contributions from all co-authors.

Competing interests

440 KM and ALSG are co-coordinators (but not associate editors) for the special issue to which this paper belongs. Otherwise, the authors declare that they have no conflict of interest.

Acknowledgements

The authors thank the editors of *Geochronology* for supporting the Special Issue this Technical Notes is submitted to.

References

- 445 Abbey, A. L., Wildman, M., Stevens Goddard, A. L., and Murray, K. E.: Thermal history modeling techniques and interpretation strategies: Applications using QTQt, *Geosphere*, 19, 493–530, <https://doi.org/10.1130/GES02528.1>, 2023.
- Flowers, R. M., Ketcham, R. A., Shuster, D. L., and Farley, K. A.: Apatite (U-Th)/He thermochronometry using a radiation damage accumulation and annealing model, *Geochim. Cosmochim. Acta*, 73, 2347–2365, <https://doi.org/10.1016/j.gca.2009.01.015>, 2009.
- 450 Gallagher, K.: Transdimensional inverse thermal history modeling for quantitative thermochronology, *J. Geophys. Res. Solid Earth*, 117, 1–16, <https://doi.org/10.1029/2011JB008825>, 2012.
- Gallagher, K. and Ketcham, R. A.: Comment on “Thermal history modelling: HeFTy vs. QTQt” by Vermeesch and Tian, *Earth-Science Reviews* (2014), 139, 279–290, *Earth-sci Rev*, 176, 387–394, <https://doi.org/10.1016/j.earscirev.2017.11.001>, 2018.



- 455 Gallagher, K. and Sambridge, M.: Genetic algorithms: A powerful tool for large-scale non-linear optimization problems, *Comput. Geosci.*, 20, 1229–1236, 1994.
- Gautheron, C. and Zeitler, P. K.: Noble gases deliver cool dates from hot rocks, *Elements*, 16, 303–309, <https://doi.org/10.2138/GSELEMENTS.16.5.303>, 2020.
- Gautheron, C., Tassan-Got, L., Barbarand, J., and Pagel, M.: Effect of alpha-damage annealing on apatite (U-Th)/He
460 thermochronology, *Chem. Geol.*, 266, 157–170, <https://doi.org/10.1016/j.chemgeo.2009.06.001>, 2009.
- Guenther, W. R., Reiners, P. W., Ketcham, R. A., Nasdala, L., and Giester, G.: Helium diffusion in natural zircon: radiation damage, anisotropy, and the interpretation of zircon (U-Th)/He thermochronology, *Am. J. Sci.*, 313, 145–198, <https://doi.org/10.2475/03.2013.01>, 2013.
- Hestnes, Å., Gasser, D., Ketcham, R., Dunkl, I., Ksienzyk, A. K., Scheiber, T., Sirevaag, H., and Jacobs, J.: The Thermal
465 Evolution of Western Norway Based on Multi-Sample Models of an Elevation Transect: Implications for the Formation of High-Elevation Low-Relief Surfaces on an Elevated Rifted Continental Margin, *Geochem., Geophys., Geosystems*, 25, <https://doi.org/10.1029/2023gc010986>, 2024.
- Ketcham, R. A.: Replication Data for: “Technical note: Incorporating topographic deflection effects into thermal history modelling”, Texas Data Repository [data set], <https://doi.org/10.18738/T8/EZ5AA0>, 2025.
- 470 Ketcham, R.: Thermal history inversion from thermochronometric data and complementary information: New methods and recommended practices, *Chemical Geology*, 122042, <https://doi.org/10.1130/abs/2022AM-382455>, 2024.
- Ketcham, R. A.: Forward and Inverse Modeling of Low-Temperature Thermochronometry Data, *Rev. Mineral. Geochem.*, 58, 275–314, <https://doi.org/10.2138/rmg.2005.58.11>, 2005.
- Ketcham, R. A., Donelick, R. A., and Donelick, M. B.: AFTSolve: A program for multi-kinetic modeling of apatite fission-track data, *Geological Materials Research*, 2, 1–32, 2000.
- 475 Lutz, T. M. and Omar, G.: An inverse method of modeling thermal histories from apatite fission-track data, *Earth Planet. Sci. Lett.*, 104, 181–195, 1991.
- Mackaman-Lofland, C., Lossada, A. C., Fosdick, J. C., Litvak, V. D., Rodríguez, M. P., Llano, M. B. del, Ketcham, R. A., Stockli, D. F., Horton, B. K., Mescua, J., Suriano, J., and Giambiagi, L.: Unraveling the tectonic evolution of the
480 Andean hinterland (Argentina and Chile, 30°S) using multi-sample thermal history models, *Earth Planet. Sci. Lett.*, 643, 118888, <https://doi.org/10.1016/j.epsl.2024.118888>, 2024.
- McDannell, K. T. and Flowers, R. M.: Vestiges of the ancient: Deep-time noble gas thermochronology, *Elements*, 16, 325–330, <https://doi.org/10.2138/GSELEMENTS.16.5.325>, 2020.
- McDannell, K. T. and Issler, D. R.: Simulating sedimentary burial cycles - Part 1: Investigating the role of apatite fission track
485 annealing kinetics using synthetic data, *Geochronology*, 3, 321–335, <https://doi.org/10.5194/gchron-3-321-2021>, 2021.
- Murray, K. E., Goddard, A. L. S., Abbey, A. L., and Wildman, M.: Thermal history modeling techniques and interpretation strategies: Applications using HeFTy, *Geosphere*, 18, 1622–1642, <https://doi.org/10.1130/GES02500.1>, 2022.



- 490 Murray, K. E., Niemi, N. A., and Clark, M. K.: Evidence for the Neoproterozoic Rifting of Rodinia in the Rocky Mountain
Front Range, *Tectonics*, 44, <https://doi.org/10.1029/2023tc008216>, 2025.
- Reiners, P. W.: Nonmonotonic thermal histories and contrasting kinetics of multiple thermochronometers, *Geochim.
Cosmochim. Acta*, 73, 3612–3629, <https://doi.org/10.1016/j.gca.2009.03.038>, 2009.
- Stevens Goddard, A. L., Fosdick, J. C., Calderón, M., Ghiglione, M. C., VanderLeest, R. A., and Romans, B. W.:
Thermochronological Evidence for Eocene Deformation in the Southern Patagonian Andes: Linking Orogenesis
495 Along the Patagonian Orocline, *Tectonics*, 42, <https://doi.org/10.1029/2022TC007677>, 2023.
- Vermeesch, P. and Tian, Y.: Thermal history modelling: HeFTy vs. QTQt, *Earth Science Reviews*, 139, 279–290,
<https://doi.org/10.1016/j.earscirev.2014.09.010>, 2014.
- Willett, S. D.: Inverse modeling of annealing of fission tracks in apatite 1: A controlled random search method, *Am. J. Sci.*,
297, 939–969, <https://doi.org/10.2475/ajs.297.10.939>, 1997.
- 500 Wolf, R. A., Farley, K. A., and Kass, D. M.: Modeling of the temperature sensitivity of the apatite (U-Th)/He
thermochronometer, *Chem. Geol.*, 148, 105–114, [https://doi.org/10.1016/S0009-2541\(98\)00024-2](https://doi.org/10.1016/S0009-2541(98)00024-2), 1998.
- Zeitler, P. K.: Inversion of $^{40}\text{Ar}/^{39}\text{Ar}$ age spectra using the controlled-random-search method, *Eos Trans. AGU* 74 43, 1993.