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

Kendra E. Murray[1], Richard A. Ketcham[2], Andrea L. Stevens Goddard[3]

[1]Department of Geosciences, Idaho State University, Pocatello, ID, 83209, USA
[2]Jackson School of Geosciences, University of Texas at Austin, Austin, TX, 78712, USA
[3]Department 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 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 one purpose of an inversion result is to map out the uncertainty that arises from this potential complexity.

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 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., 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 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.

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, 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 version 2.3, which was released in June 2025. The HeFTy files we used to produce the example inversions are available in the supplementary materials.

Overall, these examples demonstrate that 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 enhanced CRS procedure includes settings that limit such behavior (see Section 2). In other words, some CRS 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, highlighting 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 to find families of tT paths that can help one build a robust thermal history analysis.

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.

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 can happen if uncertainties are small, or if a remarkably specific resetting history is required to fit the data. 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 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.

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* 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).

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 (default 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 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

Two changes were made to the enhanced CRS procedure in HeFTy for version 2.3. 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 been identified in which it harms the search procedure, and it certainly helps in some cases.

Second, a 'reboot' condition was built into 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 internally 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 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 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 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

3.1.1 Single 40 Ma AHe Crystal

A single cooling age with little geologic context can be fitted 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 He age of an apatite crystal (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.

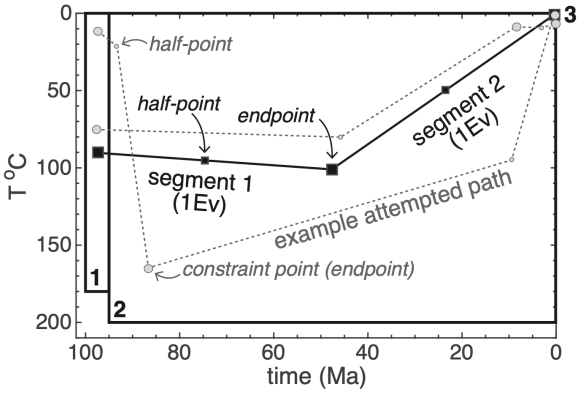


Figure 1. Our synthetic examples use identical inversion designs: two constraint boxes (labelled 1 and 2), plus a third constraint (labelled 3) of T= 5±5˚C at 0 Ma, which together generate a wide range of monotonic (only cooling) and nonmonotonic (heating then cooling) thermal histories over 100 Myr (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 constraint points alone lack the time-integrated context of the full tT paths.

The enhanced CRS (v2.3) returns 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 similar number of good paths in less than one third of the computation 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 (visualized as constraint points and tT paths) 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.

However, modest changes to the CRS settings can produce clustering in solution space that is unrelated to the resolving power of this single AHe age. Some settings for 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 a single AHe age is being inverted, 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. Additionally, we note that when opening and running a HeFTy file in v2.3 that was originally built in an older version of HeFTy v2, the CRS Options should be double checked.

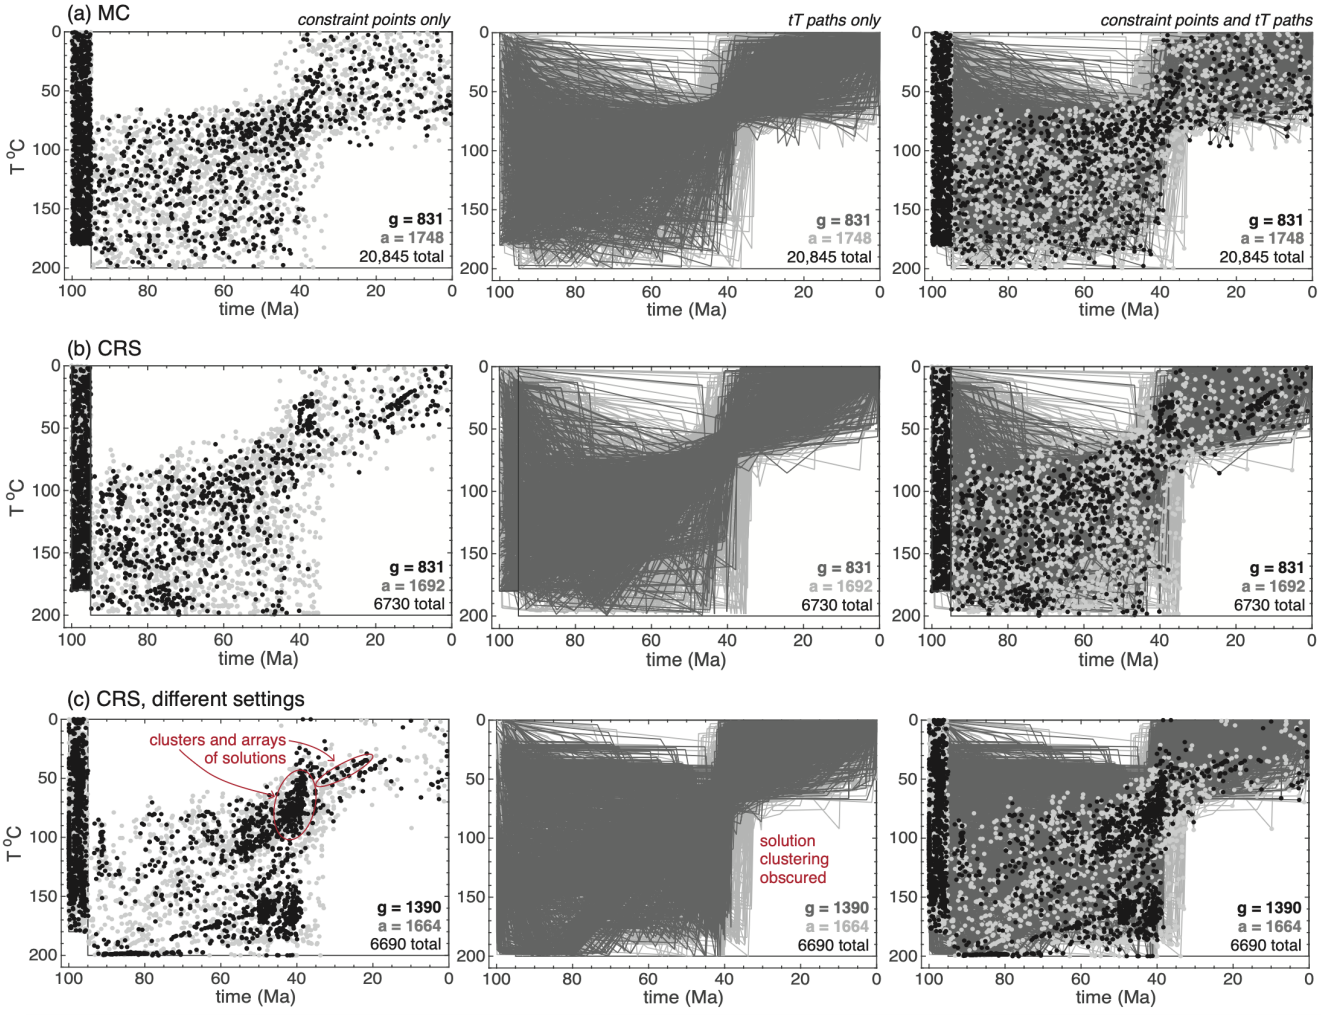


Figure 2. Inversions of a single 40 Ma AHe age ([U] = 60 ppm; radius = 60 μm) demonstrate the performance of the enhanced CRS procedure when data can be fit by a wide range of solutions. Left panels present just the constraint points (i.e., the endpoints of the independent segments of each tT path, Fig. 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. We present the same results in three different ways because constraint points and tT paths visually—and quantitatively—encompass the "solution space" (i.e., the tT conditions necessary to fit the data) in complementary but fundamentally different ways. How the solution space is described and evaluated depends on the questions motivating a thermal history analysis. (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). We only present one CRS result here, unlike the examples described later in this study. This is because when the enhanced CRS procedure in HeFTy v 2.3 is used repeatedly to invert this simplest example, it produces CRS solutions with no significant differences. (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, changes in thediffusivity of He caused by 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 (Guenthner et al., 2013).

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] pattern with ages of 4-95 Ma from apatite crystals with 10-300 ppm [U].

A MC solution from inverting the age-[eU] pattern predicted by Path 5 demonstrates the narrow range of solutions that fit these data (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.

In contrast, the enhanced CRS procedure returns 2-10 times as many good-fit paths in a fraction of the computation 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 stem from 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 Box 2 around the 'true' peak thermal conditions. In contrast, solutions CRS-2 and CRS-3 (Fig. 3c,d) returned fewer good-fit solutions (140 and 82, respectively) but a larger tT solution space, in which the good and acceptable constraint points in Box 2 are unevenly clustered and as old as 23 Ma and 35 Ma, respectively.

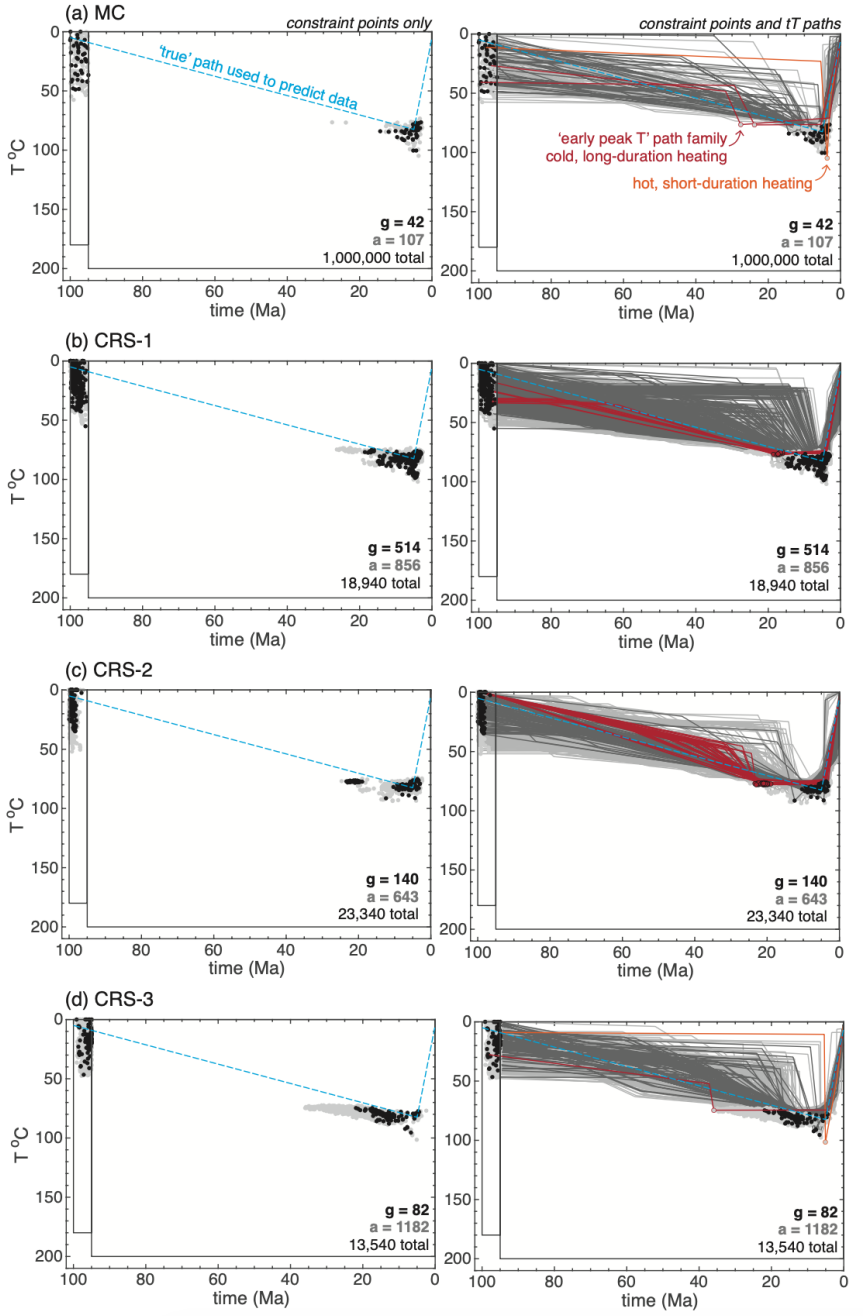


Figure 3. Inversions of a synthetic AHe age-[eU] pattern (4-95 Ma, 10-300 ppm [U]) produced by the 'true' reheating path shown in blue dashed line on each panel. The inversion design is identical to the single-crystal inversions (i.e., Fig. 1) and we used a 10% error for the synthetic ages. Left panels 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 fits to the data are statistically indistinguishable from the 'true' path. (a) An inversion result produced using the MC procedure, which 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 for 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 itself, when subsequent cooling started, etc.) was variably returned by the CRS inversions.

For example, when we analyze the constraint points in Box 2—which represent the time and temperature of peak reheating—we find that the *timing* varies 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) have peak T that skew 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 is 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 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 (~75˚C) peak T is reached 10-30 Myr earlier than the 'true' path. In the CRS-2 solution (Fig. 3c) there is a cluster of good-fit solutions with constraint points that indicate peak temperatures of ~75˚C at ca. 20 Ma. All the tT paths in this family have cold initial temperatures, stay at ~75 ˚C for ~15 Myr, and then cool to surface conditions starting at ca. 5 Ma. Indeed, an 'early peak T' path family is 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 extends peak T conditions to as early as 35 Ma and pulls the average peak T timing to be significantly older than the other inversion results. 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—is actually the most useful result for defining the permissible range of solutions that include early heating. In all cases, examining the full tT paths reveals that the solutions with early heating stay at ~75-80˚C until ca. 5 Ma and then have a similar timing, magnitude, and rate of recent cooling as the 'true' path.

In contrast, when we calculate the average *temperature* of each good-fit constraint point in Box 2 (i.e., the peak reheating T), we find this is 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 is 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) have 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.

The contrast between the narrow range of average peak temperatures and the comparatively broad range of peak-temperature 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.

In this case, the permissible 'hot, short-duration' family of solutions is 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 replicate 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.

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 to enforce a specific maximum heating or cooling rate for any path segment. In the results shown in Figure 3, the agreement between our unstated 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 instead one with a very short-duration high-T thermal pulse that produces the same data pattern), then these inversion results would not provide a useful solution.

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 from both the MC and CRS procedures are correspondingly much hotter (Fig. 4). This alternative constraint-box design used three boxes between 6 Ma and 4 Ma to permit heating from <30 ˚C to peak T and then cooling to <30˚C during this 2 Myr interval. The average peak T for a MC solution (Fig. 4a) are 114 ± 7.4 ˚C (range 97-145˚C), and three identical CRS runs (Fig. 4b-d) have average peak T of 112 ± 4.2 ˚C (range 99-131 ˚C), 114 ± 5.4 ˚C (range 101-132 ˚C), and 114 ± 4.2 ˚C (range 101-129 ˚C). In these solutions, the 'cold, long duration' family of solutions were precluded, so the minimum peak temperatures are 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.

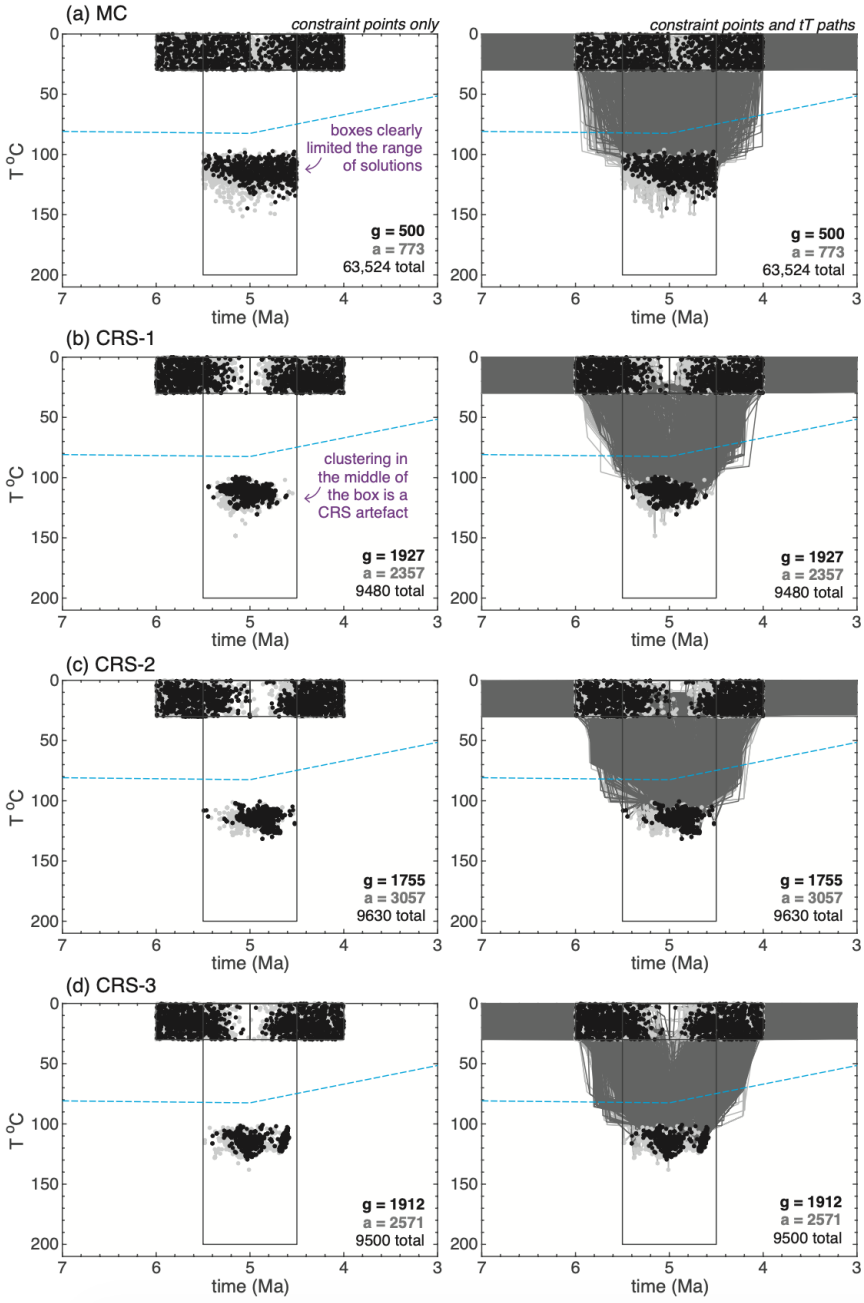


Figure 4. An alternative inversion designed specifically to explore hot, short-duration solutions by restricting heating to T >30 °C to occur only between 6 Ma and 4 Ma. Data inputs are identical to the inversions in Figure 3. The blue dotted line is the tT history used to produce the data trend. (a) One MC inversion run to find 500 good-fit paths. This solution uniformly fills 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 cluster and do 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.

If it is geologically reasonable to assume that a sample did not experience such extreme rates of heating and cooling, and therefore the '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 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 (we note however that 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 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 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 larger the tT solution space is within the tT space being explored, 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 small. In the MC solution, this manifests in an inefficient but effective solution. In our repeated CRS inversions, this manifests as occasional failure to find any fits to the data in Stage 1, although recent updates to the CRS procedure (see 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 perhaps 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), could reveal such a scenario.

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 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 inputs: two ~60 Ma AHe ages and five ZHe ages that represent an age-[eU] pattern spanning ~50-600 Ma and ~400-2300 ppm [eU]. The inversions start at 700-800˚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 compare the MC and CRS results in tT space and examine the fits to the ZHe data visualized on age-[eU] plots.

The published MC inversion result (Fig. 5a, produced using HeFTy v1.9.3) documented a tectonically significant late Neoproterozoic thermal event (constraint points in box C) and narrow range of 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 the He content and radiation damage in the zircon crystals during late Neoproterozoic time.

In order to compare this deep-time MC inversion result to those produced with the CRS algorithm, we re-ran the MC inversion in HeFTy v2.3 (Fig. 5b) because the statistical schema used to calculate the composite GOFs in HeFTy version 2 can potentially return more solutions regardless of the search algorithm being used (Ketcham, 2024). Indeed, this new MC result includes not only that same good-fit solution as the published inversion but also a new family of acceptable-fit solutions (green points and paths in Fig. 5b, n=214) that do not have Neoproterozoic heating to T > 240˚C in box C. This new path family sparsely but uniformly fills boxes A and B. However, all but three of the paths have peak temperatures in box D of >200˚C and correspondingly predict ZHe ages that are ~60 Ma across the full range of [eU] (green field on the age-[eU] plot in Fig. 5b). Therefore, this path family appears inconsistent with the observed old, low-[eU] ZHe ages, despite being returned as an "acceptable" fit—an instructive reminder to examine the fit to the data during thermal history analysis (Gallagher, 2016; Flowers et al., 2016). The trade-offs between the peak temperatures in boxes C and D that this path family highlights—as well as why fitting the observed ZHe age-[eU] pattern requires hot conditions in late Neoproterozoic time and relatively cold conditions in late Cretaceous time—were previously reported in the thermal history analysis of Murray et al. (2025), which also includes discussion of other features of the new MC result, such as why the permissible thermal conditions in boxes B and C can produce rare outliers and the key role of the young, highest-[eU] ZHe ages in resolving the heating in box C. Although not specifically related to the main point of our technical note, this is a useful demonstration of how data handling choices (e.g., assigning an uncertainty to data inputs; Flowers et al., 2022) that were calibrated to version 1 of HeFTy may not be directly applicable to version 2 of HeFTy. The MC solution generated in v2.3 (Fig. 5b) provides the most instructive comparison between the MC and CRS algorithms regarding how they handle time-temperature trade-offs on billion-year timescales.

As in the synthetic examples, the CRS procedure returns deep-time solutions in a fraction of the computation time of the MC procedure (e.g., ~24-36 hours for MC vs. ~20-40 minutes for CRS). The CRS results feature strong clustering of good solutions in boxes A, B, C, and D; a relatively small number of acceptable solutions in Box C with T < 240˚C; and a large 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 before using this result for geological interpretations.

In this inversion design, we used 'variable' segment behavior for the large Mesoproterozoic exploration boxes A and B, which 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 5c, 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 adequately, to pick one path family over the other rather than sampling them with equal probability. Although the 'enhanced' algorithm is capable of avoiding this tendency in some circumstances (e.g., Fig. 2), we find it is prone to this behavior in this deep-time example.

If the ZHe system was completely reset during Neoproterozoic time, 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, without the existing MC sensitivity testing (Fig. 5b and Murray et al., 2025), the relationship between the low-T points in box C and the high-T points in box D is not clear. These latter questions can be addressed by creating two new inversion designs that respectively enforce only monotonic cooling and only reheating during Mesoproterozoic time (Fig. 5d, e). These modified inversions both reproduce non-reset points in box C and higher-T points in D, reinforcing that no information can be extracted concerning Mesoproterozoic time.

An additional insight into the CRS procedure can be gleaned from comparing the various results. A cursory comparison of Figures 5d and 5e 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 Figure 5e 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 consequence of when the algorithm found them during that particular run, not a feature of the underlying data.

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).

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

410 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.

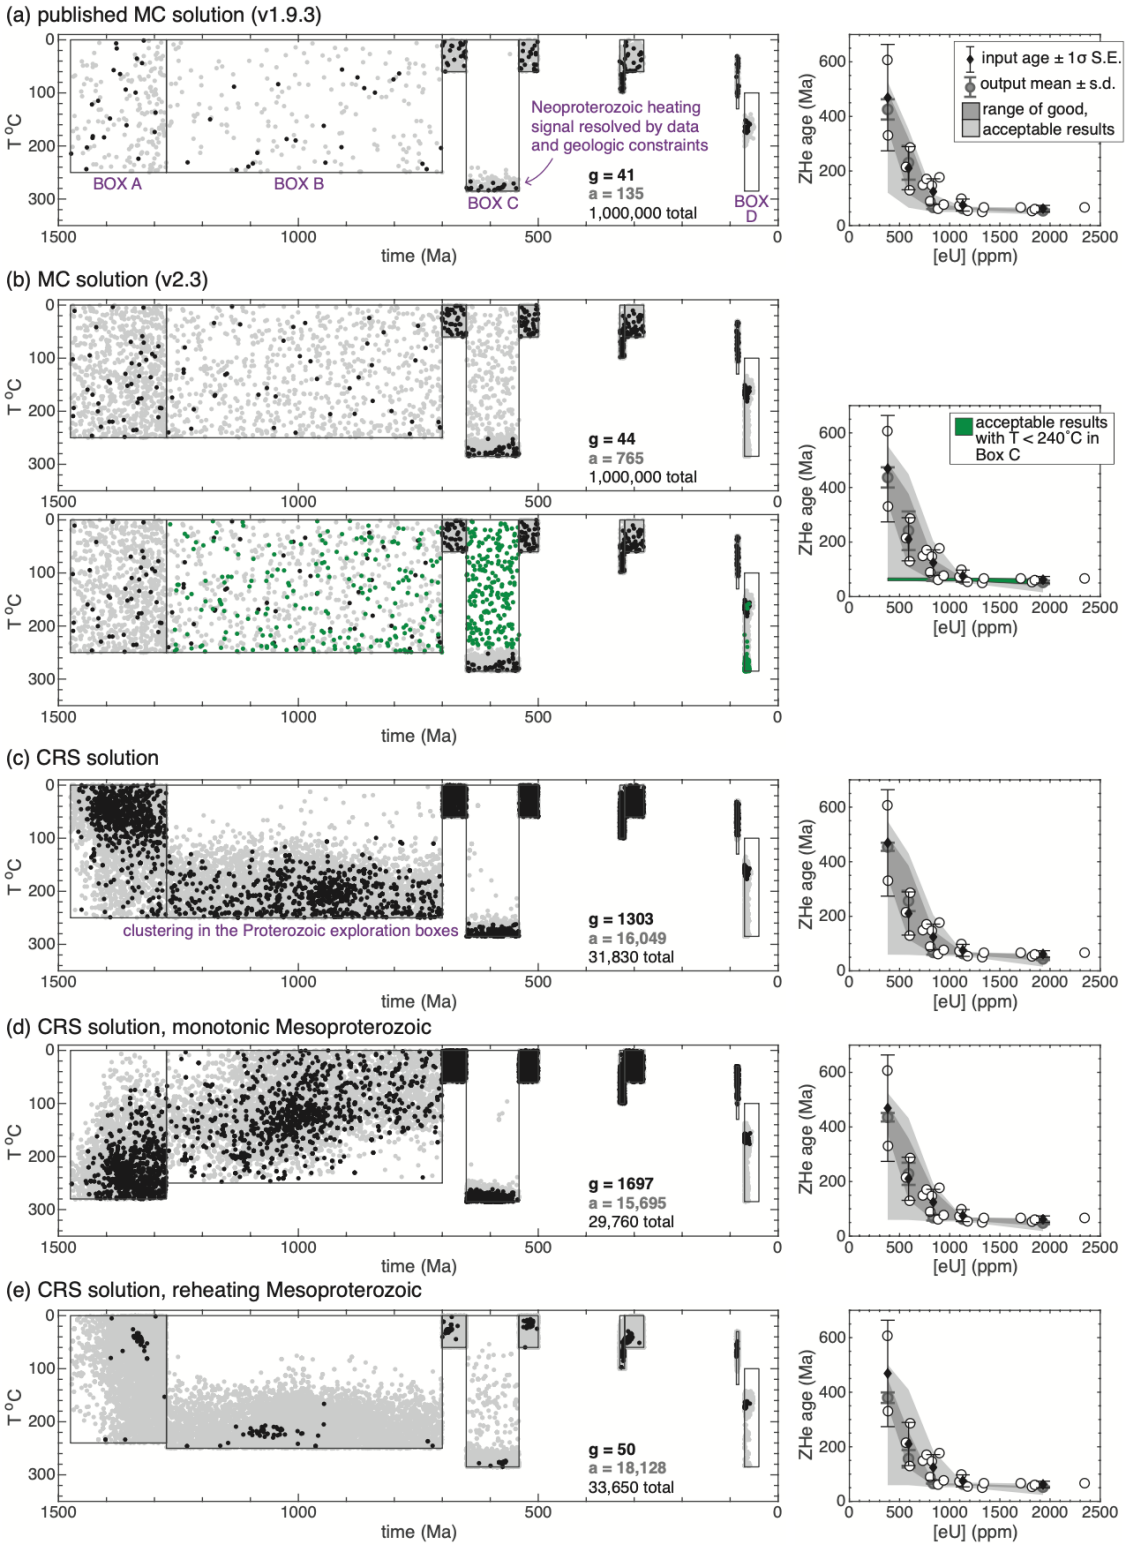

(a) published MC solution (v1.9.3)
T °C
time (Ma)
BOX A
BOX B
BOX C
BOX D
Neoproterozoic heating signal resolved by data and geologic constraints
g = 41
a = 135
1,000,000 total
ZHe age (Ma)
[eU] (ppm)
input age ± 1σ S.E.
output mean ± s.d.
range of good, acceptable results
(b) MC solution (v2.3)
g = 44
a = 765
1,000,000 total
acceptable results with T < 240˚C in Box C
(c) CRS solution
clustering in the Proterozoic exploration boxes
g = 1303
a = 16,049
31,830 total
(d) CRS solution, monotonic Mesoproterozoic
g = 1697
a = 15,695
29,760 total
(e) CRS solution, reheating Mesoproterozoic
g = 50
a = 18,128
33,650 total

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 MC solution to new CRS inversions. Data inputs for all inversions are identical to those used in the published MC result. Left panels present the solutions in tT space. Right panels present the good and acceptable fits to the ZHe age-[eU] pattern. (a) The published MC result, which was produced in HeFTy v1.9.3. (b) A new MC result produced using the identical inversion design in HeFTy v2.3; two identical tT panels present the same result with and without annotations. This solution includes more acceptable-fit paths and a new family of acceptable solutions (green points in bottom panel) that permit colder temperatures in Box C and hotter temperatures in Box D, because the statistical schema used for acceptable- and good-fit criteria was updated in HeFTy version 2 (Ketcham, 2024). This new MC solution is most relevant for comparison to the CRS solutions. (c) 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. (d) A CRS result with boxes A and B adjusted to only explore monotonic (i.e., always-cooling, no-reheating) Mesoproterozoic histories. (e) A CRS result with boxes A and B adjusted to only explore reheating Mesoproterozoic histories.

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. As demonstrated here, the CRS can even find broader solution spaces than MC, if the MC is constrained by the practicalities of computational time.

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. 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. We also recommend choosing one version of HeFTy to work with during an analysis, and/or listing all CRS algorithm parameters when reporting thermal history results, in order to avoid unnecessary version-driven complexities during the assessment of inversion results.

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 an understanding of the CRS behavior on a familiar or synthetic dataset before applying it to new problems; this is the same best practice that applies to using a MC procedure for the first time. Users can also help avoid producing artificially narrow CRS solutions by posing clear, well-constrained questions: for example, 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.

Code and data availability

The HeFTy files used to produce the inversions in this contribution are provided in the supplementary materials. Inversion results were exported from HeFTy for further analysis and plotting. 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.

Author contributions

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

Competing interests

KEM 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 Note is submitted to.

References

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.

Flowers, R. M., Farley, K. A., and Ketcham, R. A.: Response to comment on "A reporting protocol for thermochronologic modeling illustrated with data from the Grand Canyon," Earth Planet. Sci. Lett., 441, 213, https://doi.org/10.1016/j.epsl.2016.02.024, 2016.

Flowers, R.M., Ketcham, R.A., Enkelmann, E., Gautheron, C., Reiners, P.W., Metcalf, J.R., Danišík, M., Stockli, D.F., and Brown, R.W., 2022, (U-Th)/He chronology: Part 2. Considerations for evaluating, integrating, and interpreting conventional individual aliquot data: GSA Bulletin, doi:10.1130/b36268.1.

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.

Gallagher, K. and Sambridge, M.: Genetic algorithms: A powerful tool for large-sclae non-linear optimization problems, Comput. Geosci., 20, 1229–1236, 1994.

Gallagher, K.: Comment on 'A reporting protocol for thermochronologic modeling illustrated with data from the Grand Canyon' by Flowers, Farley and Ketcham, 441, 211–212, https://doi.org/10.1016/j.epsl.2016.02.021, 2016.

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 thermochronology, Chem. Geol., 266, 157–170, https://doi.org/10.1016/j.chemgeo.2009.06.001, 2009.

Guenthner, 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 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.

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.

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 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 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.

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 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.

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, 1998.

Zeitler, P. K.: Inversion of $^{40}Ar/^{39}Ar$ age spectra using the controlled-random-search method, Eos Trans. AGU 74 43, 1993.

545