

# Reviewer 1 – Kendra Murray

## General Comments

This manuscript presents a new approach for solving the crustal heat transfer problem in 1D and 3D using a recently developed deep-learning approach called physics-informed neural networks (PINNs). A PINN method can solve PDEs like the advection-diffusion equation with complex boundary conditions, especially as inverse problems, and it is therefore an attractive tool for calculating the thermal structure of the crust given changing surface elevation (1D) or topography (3D)—a common task in the interpretation of cooling ages. This manuscript presents forward and inverse solutions with a simple synthetic scenario in 1D, demonstrating that PINNs can predict crustal thermal profiles and time-temperature paths (forward) similarly to a more traditional numerical finite-element method and resolve the “true” exhumation histories, temperature histories, and synthetic cooling ages well (inverse). These 1D solutions are useful for demonstrating some of the model design choices that control optimization, search algorithms, etc. Then, this contribution moves into 3D, using an example of the Cretaceous-Recent landscape evolution of the Dabie Shan (the site of previous thermochronology studies and a published dataset). Forward and inverse models include changes to both topographic relief and rock uplift (and thus exhumation) rates over time. Inversion results are compared to PECUBE (an existing 3D finite-element method that is most similar to what PINNs can in principle do) and published thermochronologic data from the region. Overall, this manuscript demonstrates that PINNs are capable of being trained to solve relatively simple crustal heat transfer problems, and they have potential for tackling more complex thermokinematic scenarios with more statistical rigor than what is presented in this contribution.

This manuscript provides an unprecedented perspective on how deep-learning approaches might be applied in thermochronology, and this is important because I think the next generation of computational tools we use to interpret cooling ages could benefit from such advances. By applying PINNs to the crustal thermal field problem, even with simple scenarios that our current tools are very capable of solving, this manuscript represents an important step in introducing deep-learning approaches to those of us interested in interrogating the thermal history of rocks using thermochronology. As such, with some revision for clarity (most of my suggestions are primarily provided as comments in the attached PDF), I think this contribution is a good fit for this journal. I do note that my expertise lies in thermal history and thermokinematic modeling of thermochron data, and not in neural networks, so I cannot really rigorously review the details of the PINN methods, functions, optimization, and learning rates described here. Instead, I focus my comments on the packaging of the method and how to make it accessible to a thermochron audience who might be considering using such tools.

## Several specific comments that summarize my line-by-line feedback in the PDF

1. This manuscript would be stronger if it clearly articulated why deep-learning approaches will (eventually?) be an improvement over current numerical methods. On the face of them, PINNs seem like huge black boxes — with “multiple hidden layers with trainable parameters and nonlinear activation functions” (lines 96-7). Make a more clear case for why pursuing deep-learning approaches like PINNs is worthwhile. What will we get? Why should I spend time learning about this? It seems clear from this work that these methods are still in the early stages of development, and have not yet superseded currently available methods. Or, if the authors think the PINN approach presented here is already an improvement over current methods in some way (computational time? accuracy?), they do not make a clear case for this. Either way, a more explicit argument for the utility of these methods would improve the impact of this paper.

I will also add that the synthetic 1D scenario and model design is set up such that the inversion is a pretty simple problem to converge on the “true” answer for, which limits the impact of the example a bit. A suite of 6 perfect thermochronologic ages that span a 400°C temperature range produced by a simple two-stage monotonic cooling history in 1D is a very easy synthetic dataset to fit because it is distinctive — there is a narrow range of tT histories that can produce such a data set. This is especially true if, as is the case here, the inverse model design is narrow and dialed to only explore very similar/relevant histories (1D,

monotonic, only exploring for one change in exhumation rate, only exploring 50 Ma of time, etc). Essentially, it's not a hard problem to solve.

The synthetic 1D scenario presented here is an important starting demonstration, but many of the tools introduced in the beginning of this paper would perform similarly—no PINN black box required. I think a more interesting and compelling demonstration of the potential of this tool would be with a 1D (single-sample) synthetic dataset that is far less unique and more similar to a real dataset (fewer data) paired with a model design that explores whether these data could have been produced by a more complex or less constrained thermal history (i.e., starting earlier, non monotonic, multiple changes in exhumation rate possible). How well does this approach retrieve the “true” answer in this case? This is the realm of real thermochronologic datasets, and gets to the heart of the fundamental non-uniqueness of cooling ages, which is at the core of why thermochronologic studies rely on numerical tools for data interpretation.

The manuscript as it is already represents a complete contribution in my view, but I wanted to point out this additional opportunity.

REPLY: For inverse modeling, PINNs are coupled with a global search method to estimate the possible thermal history from thermochronological ages. As the reviewer pointed out, in a lot of cases for a given set of cooling ages, the possible time-temperature history is not unique. In this case, our search method won't be able to “retrieve the ‘true’ answer”, nor any other search methods can do. As the main purpose of this work is to introduce the PINN method, in the revision we compare the performance of inversions in which the thermal histories are solved using the PINNs vs. numerical methods. We modified the 1D forward model such that it features three periods with different rates.

Currently, the main advantages of PINNs over traditional numerical methods include its mesh-free setting, adaptability to configure complex boundary conditions, and relative efficiency in inverse analysis. In the revised paper, these advantages are discussed in more detail. However, as we still focus on the feasibility of the PINNs method and use numerical solutions to verify the result of PINNs, we did not design too complicate problems to demonstrate the full potential of PINNs.

2. Be more explicit about what is meant by “uplift” through out the paper (i.e., England and Molnar, 1990). The authors know that rock uplift only results in rock cooling if it is accompanied by exhumation, but that is not clear in the text as written. In the simple 1D scenarios, rock uplift = exhumation (both in rate and magnitude) because the surface elevation is not changing. But even then, using these terms interchangeably is not useful because there are many 1D scenarios where that would not be true. In the 3D example, the imposed rock uplift rate changes independently of the topographic evolution, so it is effectively decoupled from exhumation in space and time. Therefore, it is even more critical to be clear.

REPLY: Throughout the paper, “uplift” is “rock uplift” unless otherwise specified. In the revision, we will be careful and more specific by using “rock uplift” instead of “uplift”.

3. The model design and results are difficult to follow in places because it is hard to keep track of “model time” vs. “geologic time”, which are always opposites in such models. I suggest the authors can avoid some of this confusion by using “Myr” (millions of years) when referring to model time and timescales and “Ma” (millions of years ago) when referring to geological time and thermochronologic ages, in both the main text and figures.

REPLY: Good idea. The revised manuscript uses “Myr” and “Ma” for “model time” and “geological time”, respectively.

4. The Dabie Shan example would benefit from some additional text describing the previous work (data, what is known about the orogenic history and timing) in more detail, so the model design decisions made have more geologic context.

REPLY: We have added more background and introduction to the Dabie Shan study.

5. The paper would benefit from a short conclusions section that summarizes the main findings and looks ahead more broadly to future applications of deep-learning approaches in thermochronology.

REPLY: We have added a Conclusion section.

## RESPONSE to line-by-line feedback

L14 The text is changed to “how thermal structure of the crust has changed in time and space must be understood”.

L18 References of He and Tartakovsky (2021), Rasht-Behesht et al. (2022) are added.

L20 The method is easier to implement than numerical methods, and therefore has the potential to apply more complex boundary conditions. The introduction is be expanded to explain this point.

L21 “Rock uplift”. The text is modified.

L28 Agreed. “Quantifying” is a better term.

L30 Text is changed to “Collected on a vertical profile”.

L31 The term “adopts” is changed to “based on”. Next sentence is rewritten to explicitly say “the method does not require calculation of  $T_c$ ”. Limitation of AER outside “steady-state” is mentioned in the paragraph. The difficulty of collecting vertical profile will be added to the end of paragraph.

L40 Even when modeling data on a vertical profile, QTQt (at least the recent version) does not solve heat transfer; instead, it applies a linear interpolation to get Tt histories of samples and the calculate their ages. My focus here is not the T-t history modeling part of QTQt or HeFTy. Rather, some studies have used modelled T-t histories, which can be retrieved from the two apps or any other approaches, to estimate the geothermal gradient changes over time. This paragraph is rewritten to avoid confusion.

L42 QTQt has a forward modeling function to compute thermochron ages from T-t history. Hefty also has a similar function.

L64 QTQt and HeFTy are not considered apps or methods to calculate exhumation rates here, but the manipulation of the modeled T-t histories (as the reviewer stated, “it is up to the users to interpret these histories” ) is a method. Also, “the methods mentioned above” are only methods in this section. Text are modified to clarify.

L65 Text is changed.

L90 Numerical method can be difficult to implement. PECUBE has a function to calculate velocity fields associated with a single fault, or a set of faults with the same strike. A method which is convenient to implement more complex boundary conditions will be useful for solving exhumation in more complicate settings, although they have not been tested in this work yet.

L100 Text is modified.

L104 The value of the method is added.

L127 Model time will be changed to use the unit Myr.

L129 The text is modified to use “exhumation rate”.

L141 It is for demonstration purpose and a personal preference. Other machine learning libraries, e.g., PyTorch, will work too.

L150 The setup of FEM is explained.

L166 Here is a general statement. More detailed setup of the inverse problem will be described below.

L177 The sentence is modified.

L209 Yes. The result from a FEM + MCMC method is added as a comparison.

L212 Typo is corrected.

L214 A summary of the Dabie Shan evolution is added.

L219 It should be the last phase of orogeny and post-orogenic exhumation. Sentence is modified with more context added.

L221 Yes. It should be 'y'. The range of  $z$  includes potential maximum elevation of the Dabie Shan during mountain building, prior to post-orogenic decay.

L235 This should include the last phase of orogeny. Text in the model setup will be modified to reflect this.

L244 Yes. It is "rock uplift". Exhumation is not calculated here, but it should be the summation of rock uplift and topographic decay.

L247 PINN (physics-informed neural network) is a type of DNN (deep neural network) with the physical laws embedded (as partial differential equations).

L265 Here should be the model time, i.e., 150 Myr, or 0 Ma. Text is modified to clarify.

L269-L271 The  $l^2$  norm sums misfits from all data, so its value is dependent on the number of data points. Equations for calculating  $l^2$  and  $l^{inf}$  are included.

L275 Descriptions of the values is added.

L299 There is no need to put an equation here as I did not model it. The model setup is modified and this is needed.

L309 Text is rewritten..

L311 Text is rewritten.

L312 Text is rewritten.

L330 A conclusion section is added.

Figure 14 Good idea. The figures are changed to show the discrepancy.

## Reviewer 2 – David Whipp

### Summary

Jiao et al. presents an approach to determining the cooling and exhumation history of bedrock thermochronometer samples using physics-informed neural networks. The method presented here has not, as far as I am aware, been applied to thermochronology previously and is a flexible and modern alternative to the use of finite-element or finite-difference software for solving the heat transfer equation, calculating thermal histories, and predicting thermochronometer ages. The authors also demonstrate the potential use in forward models and for data inversion in one or multiple spatial dimensions in two example cases. Overall, the manuscript is a good fit for Geochronology, the method certainly seems appealing, and a nice case is made for its potential application. However, there are parts of the manuscript that would benefit from being expanded or otherwise edited to ensure readers have a good sense of the pros, cons, and potential for this method, and a handful of other considerations that could make this a stronger manuscript. In the comments below I note some of the general issues that should be considered followed by more detailed comments about the manuscript text and figures.

## Main comments

1. One of the main things I feel needs more attention in the manuscript is the case for why someone should consider using this method for interpreting their data. As detailed below, there are many interesting results that are presented, but the case for using this approach compared to existing methods is not really made/emphasized. As a new method being pitched, this seems like a key point to ensure readers are aware of how they might benefit.

REPLY: We have added more text in the Introduction and Discussion to explain the advantage of the method.

2. It would also be good, given that I would imagine many readers are not that familiar with neural networks, to provide a bit more background on topics like what neural networks are and how they differ in terms of solving equations from numerical methods like the finite-element or finite-difference methods. There are details for some things, but it also seems to be that familiarity with some neural network topics is assumed, which might not be true for many interested readers. A bit more background could significantly improve the ability of a broader group of readers to benefit from this work.

REPLY: We have added text in Section 3 for more explanation of neural networks and PINNs.

3. This is a bit tricky to handle with this type of paper, but topics typical for the methods section are mixed with the results in sections 4 and 5. This might make it hard for readers to find certain technical details, which would be unfortunate given this is a new method that readers might want to test. I would suggest considering whether it is possible to more cleanly split things between the results and methods topics in sections 4 and 5.

REPLY: We have restructure the manuscript. For each example, the Methods and Results are now in separate sections.

4. The discussion of differences between the FE and PINN solutions for the heat transfer equations is helpful in illustrating that the PINNs do reproduce the expected temperatures fairly well, but it might also be nice to have some basis for comparison. For example, it would be nice to see also some information about how much a typical finite-element solution differs from analytical solutions for heat transfer. In addition, I would suggest considering reporting percent differences between the two temperature solutions rather than absolute temperature differences, as the percent difference would not be linked to the temperature bounds the boundary conditions impose, for example. The same applies to the predicted ages. I have tried to note places where this could apply in the specific remarks below.

REPLY: We focus on the comparison between numerical method and PINNs, so an analytical solution is not provided. We added the percentage difference in the predicted ages, whose misfit function is based on age difference normalized to the age values. The loss function is constructed using the absolute temperature values, so a percentage difference will over emphasize the temperature difference in low values.

5. Finally, the manuscript is generally well written and clear, but there are some large contrasts in terms of how detailed different sections of the manuscript are. For example, the introduction section is shorter than the text describing the reason for using the differential evolution algorithm in the inverse problem description for the 1D case. As a result, some sections feel underdeveloped and lacking compared to sections that are rich with detail and information. I have tried to note sections where further development would help in the specific remarks below.

REPLY: The Introduction and Background (Sections 1, 2 and 3) have been rewritten to address this issue.

## Specific remarks (L = line number)

L1-11: Somewhere in the abstract it would be good to highlight the advantages of this approach over some of the alternatives that are currently used. It might also be helpful to note why the Dabie Shan case study is used (briefly).

REPLY: Highlight of PINNs is added. Dabie Shan has a well studied post-orogenic history, simple tectonic setting, and has been used as a site for applying or testing different modeling techniques. Abstract is modified.

L12-24: The introduction as a whole is quite short and doesn't really make a case for why this method would appeal to readers. For instance, it would be good to include something about some drawbacks of the current methods or things that motivate the use of PINNs. Essentially, similar to the abstract, it seems a case should be made for why this method should be used. The section that follows goes into more detail about existing approaches for finding thermal histories or linking ages to rates of exhumation, but before moving to that section it would be good to know what we might gain from the new approach.

REPLY: The Introduction has been rewritten to note the limitations of the conventional methods.

L31-32: I am not sure this statement is correct. Although the age-elevation method for calculating exhumation rates does make some assumptions about the geometry of the effective closure isotherm in the subsurface (see Figure 2 of Ehlers, 2005, Reviews in Mineralogy and Geochemistry, for example), there is no need to calculate an effective closure temperature to use an age-elevation profile to estimate an exhumation rate. You simply plot the sample age versus its elevation and the slope of a best-fit line is the exhumation rate estimate. I suggest not mentioning Dodson here, or at least clarifying that there is no need to actually consider the closure temperature to use this approach.

REPLY: The description of the AER method has been rewritten to avoid confusion.

L35-37: This is true, but it is also worth noting that there are other factors that can make this approach not work (such as those noted in the Ehlers paper mentioned above). Topographic and exhumational steady state can exist and still cause problems with the estimated exhumation rate due to topography affecting the effective closure temperature isotherm geometry.

REPLY: Other factors that limit the application of the AER have been added.

L40-41: I suppose here it might be good to introduce the thermal history modelling from HeFTy or QTQt first since many readers might know those models and then bring in the possible links to differences in sample depth or elevation. It seems quite handy to exploit the elevation to be able to calculate exhumation rates from codes like HeFTy, but for simplicity it might be good to present their general use first.

REPLY: We agree. The section has been rewritten to first introduce the T-t history modeling for clarification.

L52: This subsection might be more simply titled "Calculating crustal temperatures"

REPLY: The section head has been changed.

L64-72: Although we're in the process of writing a first article using the software, you could mention the Tc1D software (<https://zenodo.org/doi/10.5281/zenodo.7124271>) as another 1D code for calculating exhumation rates from thermochronometer data. The code solves the 1D heat transfer equation using finite differences and calculates (U-Th)/He and fission-track ages for apatite and zircon. Several different exhumation history models can be used to simulate time-varying exhumation rates and cooling scenarios.

REPLY: Reference to Tc1D has been added. We are looking forward to its application.

L91: As noted earlier in the introduction and abstract, it seems important that somewhere in this section you point out some of the deficiencies of the current software used for interpreting thermochronology data to make an appealing case for the use of PINNs. Without this, it feels somewhat like the method is simply

a neat computational development.

REPLY: The limitation of conventional analytical and numerical methods have been added to previous Introduction section, so it could be repetitive to state the again in this section.

L93-97: Given the range of potential readers (as noted in the general comments), it might be helpful to add a few sentences with general background about neural networks and how they differ from the numerical methods often used for similar heat transfer problems.

REPLY: The biggest advantage is that the method does not require discretization. One sentence explains this. Description of neural network and scientific machine learning have been added to indicate the difference.

L117: I would move the text from lines 126-130 here to first describe the general scenario this first forward model explores in words before describing the boundary conditions, exhumation rate model, etc.

REPLY: This has been done.

L131-135: It might be again beneficial here to explain a bit more about how the loss function components are calculated and why the weighting parameter has its assigned value. For instance, what is the basis for comparison in finding the residuals for  $L_f$ ?

REPLY: Assigning the weighting parameter to 1 in our experiments works for the optimization. For more complicate boundary conditions, it may require tuning by trial and error. The text has been modified.

L148-149: Figure 3 shows some "spikes" that appear in  $L_b$  after about 5000 iterations. While  $L$  appears to nicely decrease, is there any meaning to the "spikes" in  $L_b$  and their apparent growth with time? Should this be discussed somewhere?

REPLY: The spikes indicate that the learning rate was set too high. The learning rate has been changed to a exponential function, and such fluctuation has been reduced for forward models.

L156-157: The errors in degrees are helpful here, but perhaps it would also be good to note the percent error, at least for the infinity norm.

REPLY: We feel the percent error for temperatures can be misleading, as the loss function is based on temperature difference in degree. However, we have added percent error for predicted ages.

L163-164: Similar comment to above, that it might be nice to have a percent difference in age here.

REPLY: We have added percent error for predicted ages.

L177-189: This is a nice paragraph detailing and important choice in the design of the inverse model approach. However, the level of detail here is much greater than other sections of the text (I found this to be a good thing), so it makes those other sections appear to be lacking detail. In addition, this is an example of a topic that would seemingly belong in the methods section rather than in the results (although these are mixed in section 4).

REPLY: In the revised version, we use a different search algorithm for inversion, which requires less configuration. The text has been modified.

L208: Again, perhaps it is worth considering using percent difference here.

REPLY: We choose to show percent error for ages but not for temperatures.

L209-210: This sentence wording was somewhat awkward. Perhaps this could be rephrased.

REPLY: The paragraph has been rewritten. We have added a MCMC + FEM model for the inversion, and the accuracy between the two methods are comparable.

L212-213: It would be worthwhile to note that this a a considerably more complex and demanding example compared to the first case.

REPLY: This has been added.

L235-237: It would be helpful to have the text description of the topographic evolution moved to be prior to equation 17, for ease of reading. Also, I suggest using "Myr" to refer to model time (duration) and "Ma" to refer to millions of years ago (i.e., geochronological time). This is the first place I noted this issue in the text, but the "Myr" and "Ma" idea could be applied throughout the text and figures.

REPLY: The description has been relocated and "Ma" has been changed to "Myr" where appropriate.

L260: The "i.e." is not needed here. I suggest "...using the FEM code Pecube (Braun, 2003)."

REPLY: The text has been changed.

L264, 270-271: Again, please consider percent difference as an option for the infinity norm here.

REPLY: We choose to show percent error for ages but not for temperatures.

L273: The first sentence here is quite vague and should be made more specific to the case presented in this subsection.

REPLY: The sentence has been rewritten.

L276-277: Perhaps not here, but it would be worthwhile to know how much more demanding the inversion would be with all 5 free parameters. Could this be revisited in the discussion somewhere?

REPLY: The model has been modified to invert for all five parameters.

L286-288: As the total number of iterations is reported in the previous paragraph, could the total number of iterations be listed here as well?

REPLY: A more detailed description of the inverse modeling is provided.

L292-294: It seems that the convergence behavior would be another interesting point for further exploration in the discussion.

REPLY: The inverse algorithm has been changed, which allows less configuration. The convergence behavior is indeed interesting, but more discussion of the search algorithm is beyond the scope of this manuscript.

L296-301: This seems like a paragraph that belongs in the discussion section, perhaps revisiting the results of the examples that have been presented and how/why the results differ from the FEM solutions or estimates from other inversions. As the discussion section is currently fairly short, perhaps this could be moved and expanded?

REPLY: The discussion has been expanded substantially for the comparison between the PINN and FEM method.

L304-312: This section is probably the one I would seek out as a reader and it is currently quite brief. It would be nice to expand this and add some structure (i.e., multiple paragraphs) to discuss each of the potential benefits in more detail. For instance, I would be interested to know more about the computational efficiency of this approach compared to the alternatives, and this is only briefly mentioned in the current text. My sense is that the 3D PINNs can be quite fast, so this would be important to highlight if you want to interest readers in applying this method.

REPLY: This section has been expanded for more detailed comparison in the performance between the FEM and PINN>

L315: "CPU platform" seems a bit odd to me. Perhaps there is another term to use, such as "traditional



CPU-based systems"?

REPLY: The paragraph has been rewritten.

Figure 2: Perhaps the units for the lower x-axes could be "Myr"?

REPLY: This has been done.

Figure 4: My experience has been that it might be easier to see how the error varies if these plots were made as percent error, rather than temperature on the x axis. Perhaps this could be considered here.

REPLY: Plots have been added to show the temperature difference in degree.

Figure 5: Perhaps the log-log plot in panel b makes things look nicer, but it makes it a bit harder to visualize the error in predicted age. Perhaps a percent error could be considered here as well.

REPLY: Percent error in ages has been added to the figure.

Figure 6: It is quite hard to see the blue curves in general, and in panel a, it seems the value of the total loss is still decreasing for the green model. Is there any reason to think the loss would not decrease more for that case?

REPLY: The inverse model has been modified to use a different algorithm without different configuration. The loss may continue to decrease for more iterations, but it has been reduced to an acceptable level and the unknown parameters have been constrained to reasonable ranges.

Figure 7: The subplot labels are hard to see in some plots. Perhaps the background could be solid white?

REPLY: This has been done.

Figure 10: The lower x-axes could again use "Myr" for their unit.

REPLY: This has been done.

Figure 13: Could these be plots using percent difference?

REPLY: We added plots to show the temperature discrepancy but in degree.

Figure 14: It might be easier to compare the models if plotted as percent difference here (again). It would also require only two panels. And I would encourage using a colormap that is safe for readers with color blindness (e.g., <https://www.fabiocrameri.ch/colourmaps/>).

REPLY: All color maps have been changed to those in the suggested library.

Figure 17: Again, "Myr" could be used here.

REPLY: This has been done.

## Technical comments

L139: This line should not be indented.

L212: "Figure 9" misspelled in citation.

L296: "...with Braun and Robert..."

REPLY: The typos have been corrected.