



Demonstrating multi-country calibration of a tuberculosis model using new history matching and emulation package - *hmer*

Danny Scarponi^{a,*}, Andrew Iskauskas^b, Rebecca A. Clark^a, Ian Vernon^b, Trevelyan J. McKinley^c, Michael Goldstein^b, Christinah Mukandavire^a, Armindeer Deol^a, Chathika Weerasuriya^a, Roel Bakker^a, Richard G. White^a, Nicky McCreesh^a

^a Centre for Mathematical Modelling of Infectious Diseases, Department of Infectious Disease Epidemiology, London School of Hygiene and Tropical Medicine, UK

^b Department of Mathematical Sciences, Durham University, UK

^c College of Medicine and Health, University of Exeter, UK

ARTICLE INFO

Keywords:

Tuberculosis
Vaccines
Mathematical modelling
Model calibration

ABSTRACT

Infectious disease models are widely used by epidemiologists to improve the understanding of transmission dynamics and disease natural history, and to predict the possible effects of interventions. As the complexity of such models increases, however, it becomes increasingly challenging to robustly calibrate them to empirical data. History matching with emulation is a calibration method that has been successfully applied to such models, but has not been widely used in epidemiology partly due to the lack of available software. To address this issue, we developed a new, user-friendly R package *hmer* to simply and efficiently perform history matching with emulation. In this paper, we demonstrate the first use of *hmer* for calibrating a complex deterministic model for the country-level implementation of tuberculosis vaccines to 115 low- and middle-income countries. The model was fit to 9–13 target measures, by varying 19–22 input parameters. Overall, 105 countries were successfully calibrated. Among the remaining countries, *hmer* visualisation tools, combined with derivative emulation methods, provided strong evidence that the models were misspecified and could not be calibrated to the target ranges. This work shows that *hmer* can be used to simply and rapidly calibrate a complex model to data from over 100 countries, making it a useful addition to the epidemiologist's calibration tool-kit.

1. Introduction

Complex mathematical models, also known as simulators, are widely used in many fields in science and technology, including infectious disease epidemiology. Their utility and reliability in predicting the behaviour of real-world systems strongly depends on our ability to calibrate them to the empirical data. Despite the wide variety of calibration methods available to date, the application of calibration methods to the analysis of complex models is often lacking. A key reason that comprehensive complex model calibration is uncommon is that most formal methods require a vast number of model evaluations, which is exacerbated for high-dimensional input spaces (Kennedy and O'Hagan, 2001). This is often infeasible when dealing with computationally expensive models. The lack of methodologies and tools to robustly calibrate and analyse complex models hinders the credibility and validity of subsequent inferences drawn from model predictions. This can

lead to overconfident or misleading recommendations being made to policy makers, potentially costing lives.

History matching with emulation is a robust calibration method that has been applied successfully to complex models in a variety of fields, e. g. cosmology (Vernon et al., 2010), climate science (Williamson et al., 2013), geology (Craig et al., 1996, 1997), and infectious disease epidemiology (Andrianakis et al., 2015, 2017; McKinley et al., 2018). Despite this, its routine use by infectious disease modellers has been hindered by the lack of accessible computer software to implement the various components of the method. Our new R package *hmer* addresses this problem, allowing the user to more easily and efficiently perform history matching with emulation.

We define calibration as identifying all parameter sets that match the observed data, given the included sources of uncertainties. It is important to appreciate that this is slightly different from the technical definition of calibration used in the Uncertainty Quantification literature,

* Corresponding author.

E-mail address: dannyscarp88@gmail.com (D. Scarponi).

<https://doi.org/10.1016/j.epidem.2023.100678>

Received 13 May 2022; Received in revised form 23 February 2023; Accepted 6 March 2023

Available online 7 March 2023

1755-4365/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

where the existence of a single best-fitting parameter set is assumed and calibration is the process of estimating the posterior distribution of the (unknown) best-fitting parameter set.

In this article, we show how we performed history matching with emulation on a complex deterministic model for the implementation of tuberculosis vaccines at country-level, calibrating the model to a total of 105 low- and middle-income countries. Such a large-scale task was made straightforward thanks to our package *hmer*, which allowed us to automate the various steps of the history matching and emulation process. This was possible thanks to inbuilt checks available in *hmer*, which monitored the performance of each stage of the process.

2. The model and the calibration task

Worldwide, tuberculosis was the leading cause of death from a single infectious agent in 2019, killing an estimated 1.4 million people (World Health Organization, 2020). The only licensed vaccine against tuberculosis, Bacille Calmette-Guerin, is effective at preventing disseminated disease in infants, but confers highly variable efficacy against pulmonary tuberculosis in adults. The development of a new tuberculosis vaccine that protects several categories of people, including adults, the elderly, and immunosuppressed patients, is critical for rapidly reducing disease incidence and mortality.

Contracted by the World Health Organization, the London School of Hygiene and Tropical Medicine Tuberculosis Vaccine Modelling Group developed a model to evaluate the potential impact of country-level implementation of novel tuberculosis vaccines, in order to estimate the broader health and economic impacts (Clark et al., 2022). Full details of the model are provided in Clark et al., 2022, and described briefly here.

A compartmental deterministic dynamic model of *Mycobacterium tuberculosis* (*Mtb*) transmission and progression was developed, structured by the following core dimensions: 82 age classes, 8 tuberculosis natural history classes, two socio-economic status (SES) classes defined by access-to-care, and, for countries with a moderate or larger contribution of HIV to the tuberculosis epidemic, three HIV and antiretroviral therapy (ART) status classes.

In countries where HIV was not simulated, we had 19 input parameters that were varied during calibration, and 9 calibration targets. Among the 19 parameters, 12 were related to tuberculosis natural history, one to tuberculosis mortality, two to background mortality, three to tuberculosis treatment and one to SES. The 9 calibration targets were all for 2019, and can be divided into the following categories: tuberculosis incidence (all ages, age 0–14, and age 15+), tuberculosis notifications (all ages, age 0–14, and age 15+), tuberculosis mortality (all ages), tuberculosis subclinical proportion (the proportion of prevalent tuberculosis cases that are asymptomatic) and tuberculosis prevalence ratio by socio economic status. Ranges for the proportion of active TB disease that is subclinical, and the TB prevalence ratio by access-to-care were inferred from previous relevant publications, while ranges for all other targets, came from the WHO Global TB database. References for each calibration target can be found in Supporting Materials, Section A.

In countries including the HIV structure, we had an additional 3 input parameters, related to HIV incidence and ART initiation, and an additional four targets: HIV prevalence, ART coverage, and tuberculosis incidence and mortality in people living with HIV.

The model running time was 10 s on average, with runs usually taking longer in countries including the HIV structure usually, compared to countries without the HIV structure.

For each target, a lower and an upper bound were specified. Parameter sets were considered to match the targets if all the model outcomes were between the bounds. For more details on model structure, parameters and targets, see Supporting Materials, section A, and (Clark et al., 2022).

Model calibration was attempted for 92 countries without HIV structure and 23 countries including HIV structure, where all the data

required to calibrate the model were available.

3. Calibration method

In this section, we briefly describe how history matching with emulation works. We then show in detail how such a method was implemented in the case of our model.

3.1. History matching with emulation

History matching concerns the problem of exploring the input space (i.e. all possible parameter sets) and identifying parameter sets that may give rise to acceptable matches between the model output and the empirical data. This part of the input space is referred to as *non-implausible*, while its complement is referred to as *implausible*. History matching proceeds as a series of iterations, called *waves*, where implausible areas of input space are identified and discarded. To identify those areas, we use emulators: statistical approximations of the model outputs which are built using a manageable number of model evaluations, and are typically several orders of magnitude faster than the model. If D is a set of model runs, we can use D to build an emulator for model output f . The emulator then provides us with a prediction $E_D[f(x)]$ for the value of f at parameter set x , and quantifies the uncertainty associated with that prediction, $\text{Var}_D[f(x)]$ (Andrianakis et al., 2017; Vernon et al., 2010). This information is then used to calculate the implausibility measure at the parameter set x for the model output f :

$$I_f(x) = \frac{|E_D[f(x)] - z|}{\sqrt{\text{Var}_D[f(x)] - z}}$$

where z is the target for the model output f , and the variance in the denominator includes the uncertainty $\text{Var}_D[f(x)]$ in the emulator prediction and any other relevant form of uncertainty, e.g. the observation uncertainty or any structural discrepancy (accounting for the fact that no model perfectly represents reality). A high value of $I_f(x)$ suggests that, even while factoring in all sources of uncertainty, the emulator prediction $E_D[f(x)]$ and the target z are too distant for $f(x)$ to plausibly match the empirical data (Vernon et al., 2010). For more details about emulators and the implausibility measure see Supporting Materials C. Fig. 1 shows the various steps of the process.

Step 1. We first run the model on the initial design points, a manageable number of parameter sets that are spread out uniformly across the input space, in accordance with established design principles (Vernon et al., 2010).

Step 2. : The obtained model outputs provide us with training data to build emulators. The *hmer* package follows a Bayes Linear approach (Goldstein, 2012) for the construction of emulators.

Step 3. : We test the emulators to check whether they are good approximations of the model outputs. Emulators which do not perform well enough are either modified and made more conservative, or omitted from the current wave.

Step 4. : We generate a set of non-implausible parameter sets. To do this, we evaluate the emulators on a large number of parameter sets and we discard the implausible ones, i.e., those with implausibility measure above a chosen threshold.

Step 5. : As this is an iterative procedure, each time we complete Step 4 we need to decide whether to perform a new wave or stop. In the former case, the model is run on the non-implausible parameter sets found in Step 4 and the process returns to Step 2. In the latter case, we progress to Step 6. One possible stopping criterion is if the uncertainty of the emulator is smaller than the uncertainty in the targets, indicating that further improving the performance of the emulators will make little difference to the rate of points generation. We may also choose to stop




Fig. 1. Workflow of history matching with emulation.

the iterations when we get emulators that provide us with points matching all targets at a sufficiently high rate. If our aim is to do forecasting, we would usually be happy to stop if we have found good matches and all of the areas of the non-implausible space produce roughly similar forecasts for key future outcomes of interest. Conversely, if there is substantial variation for key forecasts in the non-implausible space, then we would want to continue, even if we have found plenty of matches already. Finally, we might end up with all the input space deemed implausible at the end of a wave: this would raise doubts about the adequacy of the chosen model, or input and/or output ranges.

Step 6. : We use the emulators to generate non-implausible points, we run the model on them and we retain only those points that match all targets.

For a more detailed introduction to history matching with emulation we refer the reader to [Andrianakis et al. \(2015\)](#), [Vernon et al. \(2010, 2018\)](#). For a direct comparison with ABC see [McKinley et al. \(2018\)](#).

3.2. Performing history matching with emulation on the tuberculosis model

The *hmer* package, released on CRAN on April 14th 2022 ([Iskauskas, 2022](#); [Iskauskas et al., 2022](#)), contains a range of options at each step of the process, allowing the calibration process to be customised at each wave. The CRAN repository offers several vignettes on the history matching with emulation methodology and on the use of *hmer* for the calibration of deterministic and stochastic models. Further tutorials and workshops on *hmer* can be found on the *hmer* website (<https://hmer-package.github.io/website>). As we were calibrating the model to a large number of countries, we used an automated approach, which did not require any customisation choices across different countries or waves. The calibration of each country was set to run for at most a week, with the aim of finding as many full fitting points as possible. We describe below how each step of the process was implemented.

Step 1. : The model was first run at 400 input points, which were chosen through Latin Hypercube sampling. Of these, half were used to build the emulators (training points) and half to validate them (validation points). Note that it is not necessary to have the same number of validation and training points. The *hmer* package can in fact work with fewer validation points, which can be advantageous when the model is computationally very expensive. In our calibration procedure, we used the same number of points for training and validation for two reasons. First, our model was not particularly slow and we could afford to do the extra runs. Second, having automated the process, we wanted to make sure we had a large enough validation set that the process would pick up any problems without user intervention.

Step 2. : The model outputs at training points were then used to train the emulators.

Step 3. : The model outputs at validation points were then used by the *hmer* package to conduct two diagnostic tests on emulators. The first assessed how accurately emulators reflected model outputs.

In [Fig. 2](#), the emulator expectation ($E_D[f(x)]$) is plotted against the model output ($f(x)$) for each validation point, providing the dots in the




Fig. 2. An example of the first emulator diagnostic. The emulator expectation ($E[f(x)]$) is plotted against the model output ($f(x)$) for each validation point, providing the dots in the graph. A 95% credible interval for the emulator expectation at each validation point is shown in the form of a vertical interval. Points for which the model output is outside the interval provided by the emulator are highlighted in red. The lower the percentage of points in red, the more accurate the emulator’s approximation of the model output.

graph, with 95% credible intervals. Points for which the model output is outside the interval are in red. The exception to this is in places where the model output is far away from the target we want to match, as it is not important that the emulators are accurate in such places. An ‘ideal’ emulator would exactly reproduce the model results: this behaviour is represented by the green line. This diagnostic was used as an initial filter, to identify emulators that performed very poorly: if more than 25% of all validation points were in red, the emulator was discarded, and the corresponding output was not used to rule out implausible points for that wave. In the example shown above, the test was clearly successful (only one validation point in red), and the emulator proceeded to the second diagnostic.

[Fig. 3](#) compares the emulator implausibility to the equivalent simulator implausibility (i.e. the implausibility calculated by replacing the emulator output with the model output). There are three cases to consider:

- The emulator and model both classify a set as implausible (bottom-left quadrant) or non-implausible (top-right quadrant). This is fine. Both are giving the same classification for the parameter set.
- The emulator classifies a set as non-implausible, while the model rules it out (top-left quadrant). This is also fine. The emulator should not be expected to shrink the input space as much as the simulator does, at least not on a single wave. Parameter sets classified in this way will survive this wave, but may be removed on subsequent waves as the emulators grow more accurate on a reduced input space.
- The emulator rules out a set, but the model does not (bottom-right quadrant): these points are highlighted in red and are the problematic sets, suggesting that the emulator is ruling out parts of the input space that it should not be ruling out.

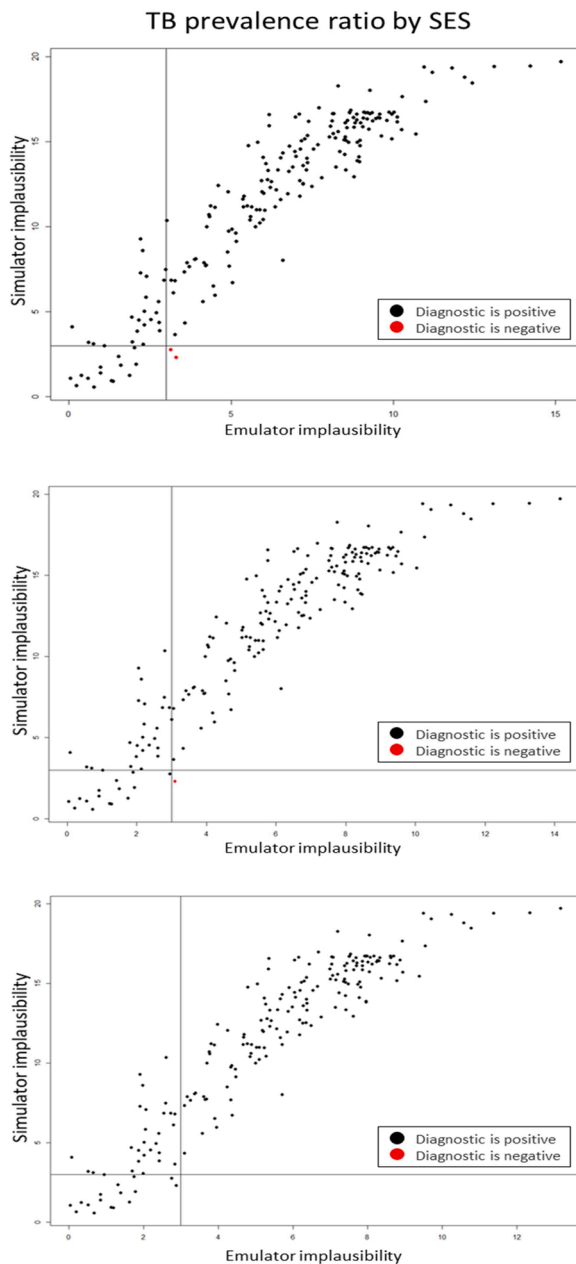


Fig. 3. An example of the second emulator diagnostic. The simulator implausibility is plotted against the emulator implausibility. The top plot refers to the emulator for the TB prevalence ratio by SES as first trained. The middle plot shows the diagnostic after the emulator uncertainty was increased by 10%. The bottom plot shows the diagnostic after a further 10% increase of the emulator uncertainty.

To ensure that no part of the input space was unduly ruled out, we required that no point was in red. We achieved this by iteratively increasing the emulator uncertainty by 10% (i.e. the length of the blue vertical bars in the first diagnostic), resulting in an emulator that was more cautious when ruling points out. In Fig. 3, increasing the uncertainty by 10% once removed one of the two red points (middle plot), and a further increase resulted in no red points (bottom plot).

Step 4. : A new set of 400 non-implausible points was generated. Using the *hmer* package, new points were generated according to the following strategy. An initial set of points was generated using a Latin Hypercube Design, rejecting implausible parameter sets (Fig. 4, top-left). Pairs of non-implausible parameter sets were then selected at random and more

sets were sampled from lines going through them, to identify those that are close to the boundary (Fig. 4, top-right). Finally, using all non-implausible points and boundary points found so far as seeding points, more parameter sets were generated using importance sampling to attempt to fully cover the non-implausible region (Fig. 4 bottom-left).

Step 5. : If the emulator uncertainty was larger than the target uncertainty, another wave was performed: the 400 points generated in STEP 4 were split in training and validation sets, the model was run on them and the process returned to STEP 2. Once the uncertainty of emulators was smaller than the uncertainty in the targets, the process continued to STEP 6.

Step 6. : Rejection sampling was performed at this stage: a set of points was uniformly sampled from the non-implausible region using all available emulators. The model was then run on such points, and those matching all targets were accepted. We repeated this process, stopping when seven full days had passed since the model calibration process started.

3.3. Derivative emulation

The calibration process described above could fail to find fully fitting points for one of two reasons: because no such points existed, or because history matching had not found them. Working with tools from the *hmer* package, we analysed countries where the fitting process failed, to identify the reason(s) why calibration could not be completed. We describe one of these tools, derivative emulation, here.

The term derivative emulation refers to the process of estimating the derivative of model outputs using the corresponding emulators. More precisely, once an emulator is trained for a model output f , we can use it to get a prediction $E[\frac{\partial f(x)}{\partial x_i}]$ of the partial derivative of f with respect to the i -th parameter at any given parameter set x , with associated uncertainty $\text{Var}[\frac{\partial f(x)}{\partial x_i}]$. This is a low cost and potentially powerful approach, since it does not require any additional model evaluations. The *hmer* package has a function that, given a set of emulators and a parameter set x not fitting all targets, estimates the partial derivatives of all model outputs at x . It then uses the gradient descent method to find the best direction to move along in order to improve (or at least not worsen) all emulator evaluations simultaneously, and then proposes a new parameter set. This function allows us to explore parameter sets near x , to determine if any of those are a better fit to the data.

4. Results

Of the 115 countries that we analysed, 105 were fully calibrated to all targets. In the time available, the number of parameter sets hitting all targets found by history matching with emulation varied by country, from a minimum of 232 to a maximum of 8857. For 101 of the 105 calibrated countries, at least 1000 full-fitting points were found. On average, for each country 13 waves were performed and 126 emulators were trained. Note that in general, the number of waves performed will depend on how conservative or restrictive emulators are. Previously published papers, where history matching was not automated and emulators were customised along the process, have used between 5 and 15 waves, depending on the number of inputs of the model.

We now show some visualisations produced using the *hmer* package for one of the calibrated countries. Fig. 5 shows (log-scaled) output values for non-implausible parameter sets at each wave. The plot clearly shows how the performance of non-implausible parameter sets improves each wave, with an increasing proportion of output values falling with the target ranges. In Fig. 6, output values for non-implausible parameter sets at each wave are shown for each combination of two outputs. For wave 7, we show only the fully fitting points. The main diagonal shows the distribution of each output at the end of each wave, with the vertical red lines indicating the lower and upper bounds of the target. Above and

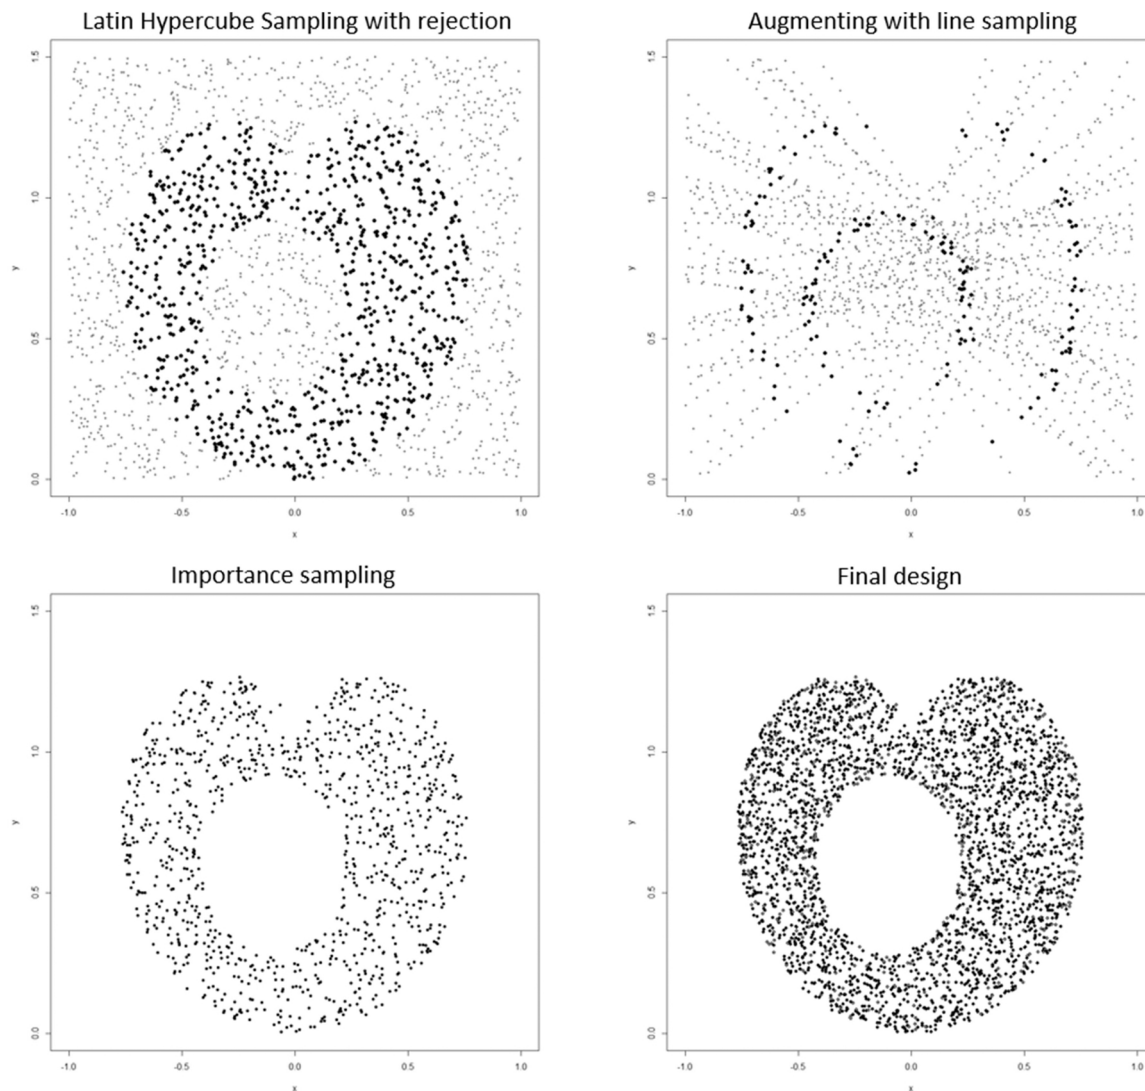


Fig. 4. Strategy used to generate non-implausible parameter sets. The top-left graph shows an example of Latin Hypercube sampling, with non-implausible points in black and implausible points in grey. Line sampling was then performed (top-right): non-implausible points found by Latin Hypercube (in grey) were connected by straight lines and points on those lines were explored. Parameter sets proposed by line sampling are in grey if implausible, or if in the interior of the non-implausible region and in black if close to the boundary. Finally, importance sampling, using mixture distributions, was carried out (bottom-left graph) using the non-implausible points found through Latin Hypercube and Line Sampling as seeding points. The final design is shown on the bottom-right graph.

below the main diagonal are plots for each pair of targets, with rectangles indicating the target area where full fitting points should lie (the ranges are normalised in the figures above the diagonals). These graphs can provide additional information on output distributions, such as correlations between them.

The third visualisation (Fig. 7) shows the distribution of the non-implausible input parameter space each wave. The plots in the main diagonal show the distribution of each parameter, which tends to narrow each wave. The off-diagonal boxes show plots for all possible pairs of parameters. These visualisations give us information about the extent to which each of the parameters has been constrained, and about correlations between parameters. For example in Fig. 7 the probability of transmission and the relapse rate have an extremely narrow posterior, while the background mortality rate has a much wider posterior. The plot also shows the presence of a negative correlation between the TB disease treatment rate and the *Mtb* infection self-clearance rate. Overall, for each country, history matching with emulation greatly reduced the input space to be searched for full-fitting points, with an estimated median reduction factor of 6×10^8 (see supporting materials, section D for more details). Fig. 8 shows a histogram of the log₁₀-transformed

reduction factors for all countries.

4.1. Computational cost of the calibration process

Each country was set to run on a single computer node (typically a 64 bit Intel Xeon processor at 2.6 GHz) for 7 full days, with the aim of finding as many full-fitting points as possible. For many countries, hundreds of full-fitting points were identified well before the end of the 7 allocated days. On average across all countries, model running time accounted for 25% of the total computational cost, emulator training/diagnostic time for 5%, and points proposal time for 70%. (Note that this study was conducted using a developmental version of the package and point generation is now around 5 times faster in Version 1.0.0 of the package. Future improvements are planned, to further optimise the generation of non-implausible points.) The points proposal time increased linearly with waves, as all previous emulators needed to be evaluated each wave. On average, 5000–10,000 points were investigated at each wave to find 400 non-implausible ones. Note that such a huge amount of evaluations were only possible thanks to the high speed of emulators: while a model run took 10 s on average, evaluating an




Fig. 5. Plot of simulator evaluations at non-implausible parameter sets for one of the calibrated countries. Each column refers to a model output, and parameter sets are coloured based on the wave they belong to, with darker colours for later waves. The model outputs (y-axis) were log-scaled to improve the readability of the graph. For each model output the lower and upper bounds of the correspondent target are shown as the white error bars.

emulator took just a few milliseconds.

4.2. Analysis of countries that could not be fitted to data

In 9 uncalibrated countries (7 countries without HIV structure and 2 countries with HIV structure), we used *hmer* package visualisation tools to understand why the calibration process could not be completed.

Let us consider country A, for which history matching found points matching 7 out of 9 targets at most. Fig. 9 (left plot) shows tuberculosis notifications in adults (ages 15–99) in 2019 and tuberculosis notifications in young individuals (ages 0–14) in 2019, with the red rectangle indicating the target area.

We see that the two outputs are highly correlated and that, in all waves, all points lie far from the red rectangle. This suggests that our model is not able to match the two targets simultaneously. Four more countries showed a similar clear incompatibility between two (or more) targets (see supporting material section B).

The remaining four countries (referred to as countries F, G, H, I) that could not be fitted to data also showed some evidence of an incompatibility between certain pairs of targets, although in a weaker form. For example, in country F, our model seemed incapable of matching overall tuberculosis notifications and tuberculosis disease incidence in young individuals (Fig. 9, right plot, where we see a “wall” in output space that prevents the model outputs from reaching the red target square. Such a wall is likely due to inherent and possibly non-obvious constraints in the model). The bounds for these two targets were almost identical. As our model did not include false positive notifications, this meant that nearly all TB in children would need to be diagnosed and notified for the model to fit the targets, which was incompatible with the plausible input ranges and other output target ranges used in the calibration.

Similar plots for countries G, H and I can be found in Supporting Material, Section B. Since the incompatibility between outputs was present but less pronounced in countries F, G, H and I, one could wonder if by moving around the last-wave points - which were often close to matching the potentially incompatible outputs - one could find full fitting points. We explored this using derivative emulation.

For countries G, H and I, the maximum number of targets met by a point in the last wave was 8 (out of 9). For these countries, even though derivative emulation increased the number of targets that some parameter sets matched, it could not find any parameter sets hitting all 9

targets. For country F, neither last wave points nor points found by derivative emulation matched more than 7 (out of 9) targets.

Finally, one country had inconsistencies in the data that made calibration unachievable.

5. Discussion

In this paper, we have demonstrated how the *hmer* R package allowed us to perform history matching with emulation on a complex deterministic tuberculosis model to a total of 115 low- and middle-income countries. For each country, the model was fit to 9–13 target measures, by varying 19–22 input parameters. Overall, 105 countries were successfully calibrated, while for ten countries history matching with emulation could not identify any full-fitting parameter sets. Such countries were analysed further, through derivative emulation and visualisations tools offered by the *hmer* package.

Calibrating 105 countries in just a few days was possible thanks to *hmer*, which allowed us to implement history matching with emulation in a straightforward and effective way. Using this method, we were able to fully explore the input space, and to identify a set of points that represent all possible conditions under which our model could match the data. In nine countries that could not be fitted to data, visualisation tools, paired with derivative emulation methods where necessary, provided us with evidence that calibration of the model was unlikely to be possible with the given input and output ranges.

Unlike most other full calibration methods, which attempt to make probabilistic statements about the posterior distribution over the input space, history matching can work with expectation and variances only, and aims to discard implausible areas of the input space. This, and the use of emulators, mean that the calculations involved in history matching tend to be more tractable and straightforward to implement. Furthermore, since the process is iterative, it is not necessary to work with all inputs and outputs simultaneously: for example, if an output cannot be emulated accurately in initial waves, it can be put aside and emulated in later waves, as emulation usually becomes easier as the waves progress (Vernon et al., 2010). All these characteristics make history matching a useful method for calibrating complex models. In addition to being a useful calibration method in its own right, when we are interested in making probabilistic statements about the posterior of the model’s parameters, history matching can be used as a pre-calibration procedure, in conjunction with probabilistic calibration

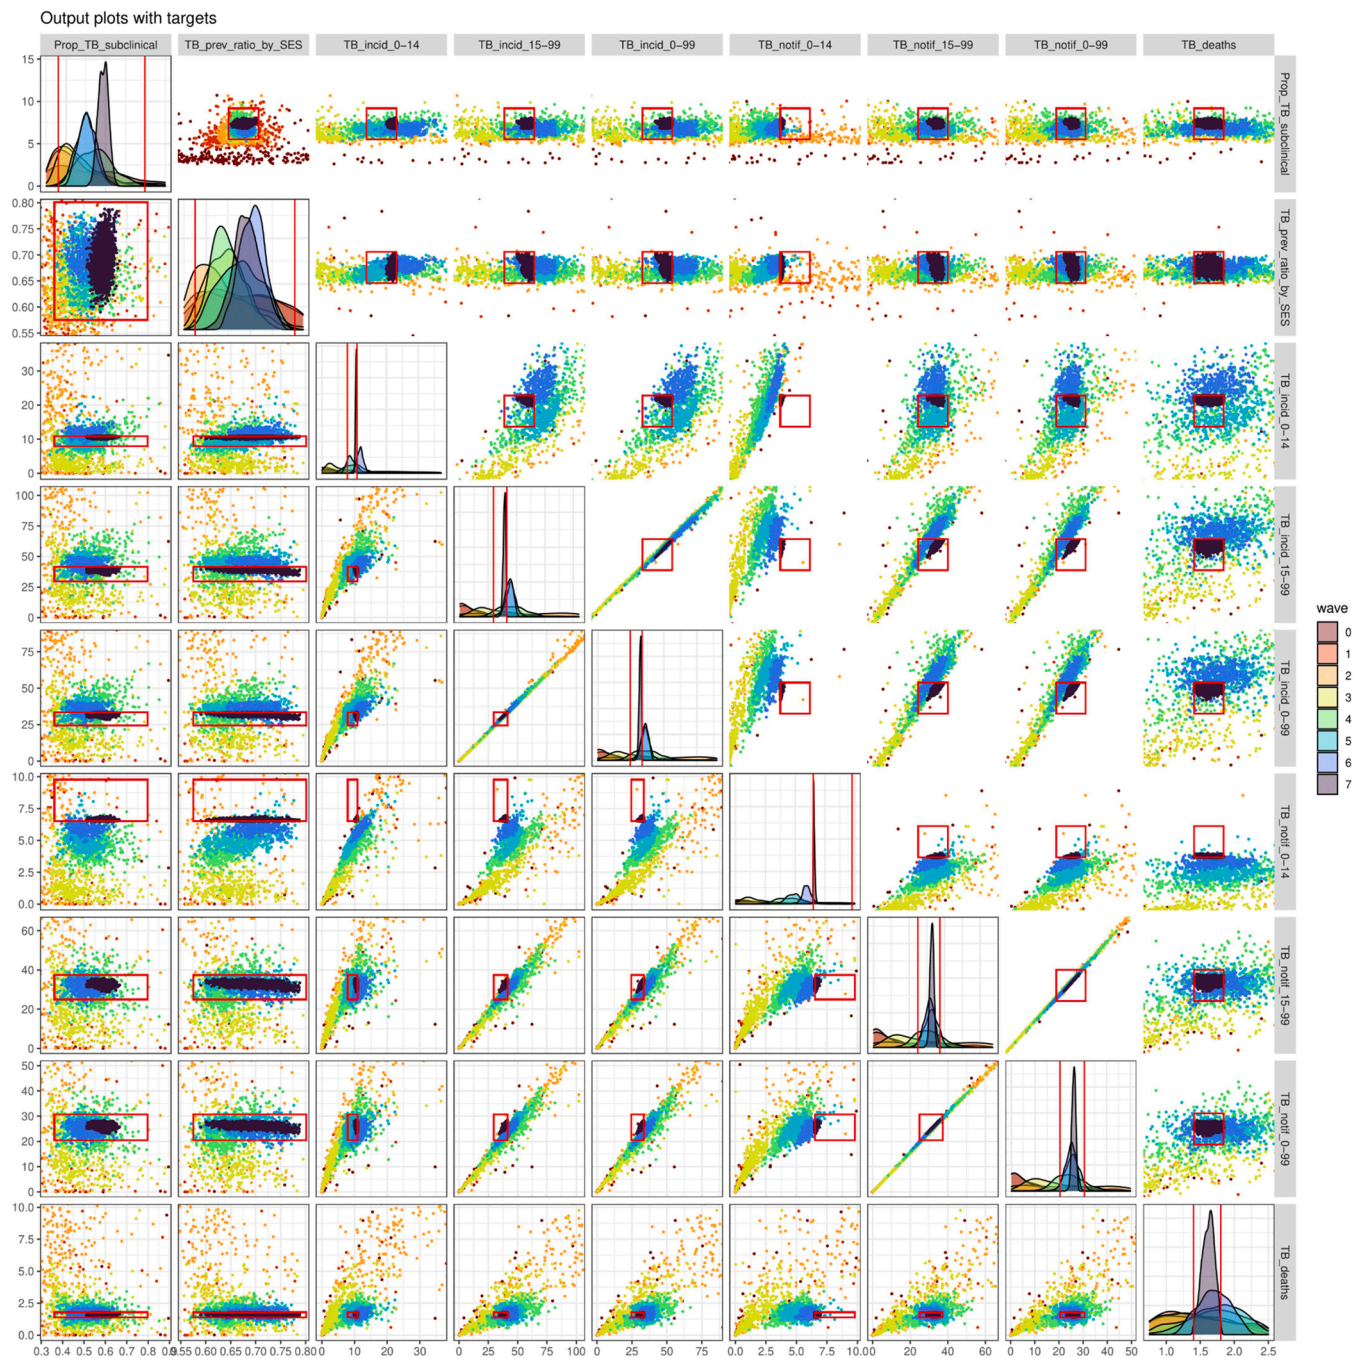


Fig. 6. Pair plots of the outputs for non-implausible parameter sets in different waves, for the same country shown in Fig. 5 and Fig. 6. The main diagonal shows the distribution of each output, with the vertical red lines indicating the bounds of the target. Below the main diagonal are plots for each pair of targets, with rectangles indicating the target area where full fitting points should lie. Above the diagonal we have similar plots, but the ranges have been normalised, so that the target area where full fitting points should lie is now a red square centred at the middle of the plot.

techniques (Vernon et al., 2018, 2010). While methods such as Approximate Bayesian Computation or Bayesian model calibration may struggle to adequately explore the large input space of a multi-output model, their application on the greatly reduced space that is the output of history matching can produce hybrid methods that are more tractable and combine the strengths of both approaches.

The closest approach to History Matching with Emulation would be Approximate Bayesian Computation (ABC). ABC relies only on the simulation model to explore the input space adequately, which for high-dimensional input spaces would likely require substantially more runs of the simulator model than we used here. We refer to a previous study

where we compared ABC-SMC (Sequential Monte Carlo) with history matching with emulation and showed that the latter was far more efficient when the computational load of the simulator model is substantial (McKinley et al., 2018). An in-depth comparison between history matching with emulation and Markov chain Monte Carlo methods can be found in (Vernon et al., 2018).

We would like to highlight the difference between history matching and emulation and calibration methods based on local optimisation. While history matching with emulation aims at finding all possible solutions that are consistent with the data, an optimisation procedure attempts to find a single point in the input space and treats that as a global




Fig. 7. Pair plots of the non-implausible parameter sets wave after wave, for the same country shown in Fig. 5. The plots in the main diagonal show the distribution of each parameter, while in the off-diagonal boxes we have plots for all possible pairs of parameters. Each colour corresponds to a specific wave, with darker colours representing later waves.

solution, ignoring the uncertainty in the input space. This is particularly pertinent when using the models to forecast, since different input points can produce similar fits to observed data, but show much more pronounced variation in forecasts. This is why it is fundamentally important to robustly explore the whole input space and not rely on optimisation approaches in our opinion.

Multi-output models, such as the one in this article, very often exhibit correlation between their outputs. This aspect was partially incorporated by building accurate 1-dimensional emulators that together captured the joint behaviour of the model outputs. However, full multivariate emulation would bring further benefits. Taking output correlation fully into account would improve the emulation process and subsequently the performance of history matching, but requires more sophisticated emulators, as well as a more detailed structure for model and observation uncertainty which takes into account these correlations.

These features, together with variance emulation, were not available at the time of this work, but are now part of the standard functionalities of version 1.0.0 of the *hmer* package. In particular, these features make it possible to use *hmer* to calibrate stochastic models, without having to average over multiple model runs for each parameter set. Also note that point generation is around 5 times faster in version 1.0.0 than it was in the developmental version of the package that was used in this work. Future improvements are planned, to further optimise the generation of non-implausible points. Features to be included in subsequent versions of the package will allow the user to address data quality problems, compare different models, and use emulators to make forecasts.

In conclusion, this work shows that the *hmer* package can play a key role in making history matching with emulation accessible to the community of epidemiologists, facilitating fast and efficient model calibration. By addressing the current lack of methodologies to robustly




Fig. 8. For each of the 115 countries we compared the initial input space and the final non-implausible space, estimating the correspondent reduction factor. The plot shows a histogram of the log10-transformed reduction factors for all countries. The estimated median reduction factor was 6×10^8 , with 95% of all reduction factors lying in the interval $[2 \times 10^5, 2.8 \times 10^{19}]$.




Fig. 9. On the left, pair plot of TB notifications in young individuals against TB notifications in adults, for parameter sets in waves 2, 3, 4, 5, 6 (country A). On the right, pair plot of TB incidence in young individuals and TB deaths for parameter sets in waves 2, 3,..., 19 (country F).

calibrate complex models and perform uncertainty analysis on them, *hmer* constitutes an invaluable addition to the epidemiologist’s calibration tool-kit.

Funding

This work was funded by Wellcome Trust (218261/Z/19/Z) and WHO (2020/985800-0). RGW is additionally funded by NIH (1R01AI147321-01), EDTCP (RIA208D-2505B), UK MRC (CCF17-7779 via SET Bloomsbury), ESRC (ES/P008011/1), BMGF (OPP1084276, OPP1135288 & INV-001754). TJM is supported by an Expanding Excellence in England (E3) award from Research England. RC is additionally funded by BMGF (INV-001754). IV is additionally funded by EPSRC funding (EP W011956).

CRediT authorship contribution statement

Danny Scarponi: Conceptualization, Data curation, Formal analysis, Methodology, Software, Visualization, Writing – original draft, Writing

– review & editing. **Andrew Iskauskas:** Conceptualization, Methodology, Software, Writing – review & editing. **Rebecca A. Clark:** Conceptualization, Data curation, Software, Writing – review & editing. **Ian Vernon:** Methodology, Funding acquisition, Writing – review & editing. **Trevelyan J. McKinley:** Methodology, Funding acquisition, Writing – review & editing. **Michael Goldstein:** Methodology, Funding acquisition, Writing – review & editing. **Christinah Mukandavire:** Data curation, Writing – review & editing. **Arminster Deol:** Data curation, Writing – review & editing. **Chathika Weerasuriya:** Software, Writing – review & editing. **Roel Bakker:** Software, Writing – review & editing. **Richard G. White:** Conceptualization, Funding acquisition, Supervision, Writing – review & editing. **Nicky McCreesh:** Conceptualization, Funding acquisition, Supervision, Writing – review & editing.

Declaration of interests

All authors declare no conflicts of interest.

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at [doi:10.1016/j.epidem.2023.100678](https://doi.org/10.1016/j.epidem.2023.100678).

References

- Andrianakis, I., Vernon, I.R., McCreesh, N., McKinley, T.J., Oakley, J.E., Nsubuga, R.N., Goldstein, M., White, R.G., 2015. Bayesian history matching of complex infectious disease models using emulation: a tutorial and a case study on HIV in Uganda. *PLoS Comput. Biol.* 11 (1), e1003968 <https://doi.org/10.1371/journal.pcbi.1003968>.
- Andrianakis, I., Vernon, I., McCreesh, N., McKinley, T.J., Oakley, J.E., Nsubuga, R.N., Goldstein, M., White, R.G., 2017. History matching of a complex epidemiological model of human immunodeficiency virus transmission by using variance emulation. *J. R. Stat. Soc. Ser. C (Appl. Stat.)* 66 (4), 717–740. <https://doi.org/10.1111/rssc.12198>.
- Clark, R.A., Mukandavire, C., Portnoy, A., Weerasuriya, C.K., Deol, A., Scarponi, D., Iskauskas, A., Bakker, R., Quaife, M., Malhotra, S., Gebreselassie, N., Zignol, M., Hutubessy, R.C.W., Giersing, B., Jit, M., Harris, R.C., Menzies, N.A., White, R.G., 2022. The impact of alternative delivery strategies for novel tuberculosis vaccines in low- and middle-income countries: a modelling study [Preprint]. *Public Glob. Health*. <https://doi.org/10.1101/2022.04.16.22273762>.
- Craig, P.S., Goldstein, M., Seheult, A.H., Smith, J.A., 1996. Bayes linear strategies for history matching of hydrocarbon reservoirs. *Bayesian Stat.* 5, 69–98.
- Craig, P.S., Goldstein, M., Seheult, A.H., Smith, J.A., 1997. Pressure matching for hydrocarbon reservoirs: a case study in the use of Bayes linear strategies for large computer experiments. In: Gatsonis, C., Hodges, J.S., Kass, R.E., McCulloch, R., Rossi, P., Singpurwalla, N.D. (Eds.), *Case Studies in Bayesian Statistics*, vol. 121. Springer, New York, pp. 37–93. https://doi.org/10.1007/978-1-4612-2290-3_2.
- Goldstein, M., 2012. Bayes linear analysis for complex physical systems modeled by computer simulators. In: Dienstfrey, A.M., Boisvert, R.F. (Eds.), *Uncertainty Quantification in Scientific Computing*, vol. 377. Springer Berlin Heidelberg, pp. 78–94. https://doi.org/10.1007/978-3-642-32677-6_6.
- Iskauskas, A. (2022). hmer: History Matching and Emulation Package. <https://CRAN.R-project.org/package=hmer>.
- Iskauskas, A., Vernon, I., Goldstein, M., Scarponi, D., McCreesh, N., McKinley, T.J., & White, R.G. (2022). Emulation and History Matching using the hmer Package. <https://doi.org/10.48550/ARXIV.2209.05265>.
- Kennedy, M.C., O'Hagan, A., 2001. Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* 63 (3), 425–464. <https://doi.org/10.1111/1467-9868.00294>.
- McKinley, T.J., Vernon, I., Andrianakis, I., McCreesh, N., Oakley, J.E., Nsubuga, R.N., Goldstein, M., White, R.G., 2018. Approximate Bayesian computation and simulation-based inference for complex stochastic epidemic models. *Stat. Sci.* 33 (1) <https://doi.org/10.1214/17-STS618>.
- Vernon, I., Goldstein, M., Bower, R.G., 2010. Galaxy formation: a Bayesian uncertainty analysis. *Bayesian Anal.* 5 (4) <https://doi.org/10.1214/10-BA524>.
- Vernon, I., Liu, J., Goldstein, M., Rowe, J., Topping, J., Lindsey, K., 2018. Bayesian uncertainty analysis for complex systems biology models: emulation, global parameter searches and evaluation of gene functions. *BMC Syst. Biol.* 12 (1), 1. <https://doi.org/10.1186/s12918-017-0484-3>.
- Williamson, D., Goldstein, M., Allison, L., Blaker, A., Challenor, P., Jackson, L., Yamazaki, K., 2013. History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. *Clim. Dyn.* 41 (7–8), 1703–1729. <https://doi.org/10.1007/s00382-013-1896-4>.
- World Health Organization. (2020). *Global tuberculosis report 2020*. World Health Organization. <https://apps.who.int/iris/handle/10665/336069>.