A Bayesian model of records

Jaime Sevilla[1] and Jonathan Lindbloom[1]

[1]Affiliation not available

May 28, 2022

Abstract

We derive a Bayesian model to forecast the continuation of a cumulative progression of records. We show that the model compares favourably to past least-square prediction models for records in running events. We also check the validity of the model for the case where we have access to the underlying attempts, using data from athletic events in the Olympics.

Highlights

- We explain how to perform **Bayesian inference** for time series where each data point is the **cumulative maximum (or minimum)** of an i.i.d. series.
- We **compare the results of this framework to a classic minimum mean square error** (MMSE) frequentist approach. We use world record data from six athletic events. We find a similar performance between the bayesian mean posterior estimate and the frequentist approach in terms of mean squared error.
- We explore the effect of the **choice of distribution of attempts.** We find that assuming a Weibull distribution marginally outperforms a Gaussian distribution and that both robustly outperform a Gumbel distribution of attempts.
- We **forecast world records for 11 categories of athletic events** for the 2022 to 2032 period.
- We introduce *fmax*, a Python open-source package to model and forecast time series of cumulative minima and maxima. The package can be found at `https://github.com/jlindbloom/fmax`.

How often and by how much are Olympic records beat? What score do we expect future machine learning systems to attain for classification tasks in the absence of new breakthroughs? With what probability will the fasted speed run for our favorite videogame be beaten within the next year?

In situations such as these, we are interested in characterizing how a historical record has evolved and will evolve in the future. And while properties of order statistics such as the maximum over a set of random variables are well-studied, the running maximum (or minimum) of a time series is markedly less so.

An example of such work is presented in (Tryfos & Blackmore, 1985), in which the authors present a model for the world record in six major running events using an i.i.d. distribution for attempts. In this article, we present the corresponding Bayesian approach to such models. We use this model to derive predictions for the men's and women's data for the same events considered in (Tryfos & Blackmore, 1985) and show comparable performance in terms of the squared error given the actual records that followed.

We also discuss the effects of the choice of attempt distribution. We find that a Weibull distribution gives the best fit in terms of loglikelihood, marginally outperforming a Gaussian distribution and robustly outperforming a Gumbel distribution.

Finally, we provide a forecast of records in the next decade for the 11 categories of athletic events we collected data on.

Previous work

(Tryfos & Blackmore, 1985) develops the minimum mean square error (MMSE) estimator for a series of cumulative minimums. They derive the estimator for both a normal distribution and extreme value distributions of attempts. In particular, they apply their approach to forecast the world records of six running events, assuming an underlying Gumbel distribution with density $f(x) = \exp\left(\frac{x-\mu}{\sigma} - e^{\frac{x-\mu}{\sigma}}\right)\frac{1}{\sigma}$. We use this article as a basis to compare our approach. In section we show how our approaches compare using the same data as the authors.

(Smith, 1988) expands the work of Tryfos and Blackmore to derive the maximum likelihood estimator for a series of cumulative minimums where the attempt distribution is the sum of an i.i.d. random variable X_n and a nonrandom drift trend.

$$Y_n = X_n + c_n$$

They consider random distributions including the Gaussian, Gumbel and Generalized Extreme Value (GEV) distributions. The drift trends considered include linear drift, quadratic drift, and exponential decay models. The author applies the method to model records in the mile and marathon races.

We could in theory adapt their approach assuming a zero-drift trend $c_n = 0$. However, in practice with the data, we considered and in short time scales this results in constant extrapolated forecasts $\hat{Y}_{n+k} = Y_n$. Future work may include extending the Bayesian framework presented in this paper to the case of non-zero drift and comparing it to this paper.

(Smith & Miller, 1986) follow Smith by considering a Gumbel model with linear drift, but work within a state-space approach to explicitly construct the forward-looking predictive distribution for the model. They also consider a Bayesian formulation, applying their method to the forecasting of athletic records similar to Tryfos and ourselves.

(Wergen et al., 2013) focuses on the related problem of modeling the probability that timestep n will be a new record given historical data. The authors assume a linear drift in the attempts.

(Kim & Seo, 2020) derive the Jeffrey prior for the Gumbel distribution, and derive the density function of the conditional forecasted distribution of records of i.i.d. Gumbel variables given previous observations. They compare their result to an ARIMA and DLM approach. As we will discuss in this article, the Gumbel distribution seems to underperform relative to the Weibull and Gaussian distribution, suggesting a natural extension to their work.

Beyond the aforementioned articles, we could not find much work on forecasting cumulative records, especially from a Bayesian perspective. This points to a gap in the literature that we aim to fill.

The General Model

In this section, we introduce a general framework for deriving the likelihood of the distribution.

First we establish the notation we will use through the paper. Then we derive the likelihood functions for a cumulative distribution of maxima and minima. These are, respectively:

$$\mathcal{L}_{Y_{1:n}}(y_1, \ldots, y_N) = \prod_{i \in R} f_X(y_i) \prod_{i \notin R} F_X(y_i)$$

$$\mathcal{L}_{Y_{1:n}}(y_1, \ldots, y_N) = \prod_{i \in R} f_X(y_i) \prod_{i \notin R} (1 - F_X(y_i))$$

Finally, we show how we can use these results to forecast future records.

Notation

Suppose we would like to model some time series for the running record for some task, where the record is the maximum or minimum over some sequence of attempts. We'll first consider the case where the record is of a minimum - the derivation for the maximum requires only a slight modification. We assume that the record is a continuous quantity, rather than discrete. Let $\{X_t\}_{t\ \in \mathbb{N}}$ denote a discrete-time stochastic process representing the results from some sequence of attempts at a task. We assume that the X_i are i.i.d. according to some random variable X with a common CDF given by F_X and PDF given by f_X. When we make inferences about X we will assume that X lies in some parametric family parameterized by $\theta \in \Theta$.

Our observed data of the record is some time-series $\{r_1, r_2, \ldots, r_n\}$ where n is the number of time periods for which the record has been observed. In the case the record is of a minimum, note that we must have $r_i \geq r_j$ whenever $i \leq j$. To match our observed data, we define a sequence $\{Y_t\}_{t\ \in \mathbb{N}}$ where

$$Y_i \ := \max\{X_1, \ldots, X_i\ \} .$$

We treat $\{r_1, r_2, \ldots, r_n\}$ as noiseless, truncated observations along some sample path $\omega = \{r_1, r_2, \ldots, r_n, \ldots\}$.

The Likelihood Function

To perform Bayesian inference on the model parameters θ, we need to be able to compute the likelihood function. The *marginal distribution* of each record is easy enough to derive:

Lemma 1: Marginal distribution of a historical record

Let $X_1, \ldots, X_n$ be a collection of i.i.d. i.id random continuous variables with PDF f_X and CDF F_X, and define $Y_n := \max_{i \leq n} X_i$. Then the marginal likelihood of Y_n is equal to:

$$\mathcal{L}_{Y_n}(o_n) = n[F_X(o_n)]^{n-1} f_X(o_n)$$

Proof: The CDF of Y_n is:

$$\begin{aligned}
F_{Y_n}(o_n) amp; &= P(Y_n \leq o_n) = P(\max\{X_1, \ldots, X_n\} \leq o_n) \\
amp; &= \prod_{i \leq n} P(X_i \leq o_n) \\
amp; &= \prod_{i \leq n} F_X(o_n) \\
amp; &= [F_X(o_n)]^n
\end{aligned}$$

Differentiation of the CDF gives us the desired result. □

However, to use all the data available to us in the inference we need to derive the *joint likelihood* of all the cumulative records:

$$f_{Y_{1:n}}(y_1, \ldots, y_n|\theta)$$

where $f_{Y_{1:n}}(\cdot|\theta)$ denotes the joint density of $Y_1, \ldots, Y_n$, given parameters θ. Since

$$\max\{X_1, \ldots, X_j, X_{j+1}\} = \max\{X_j, X_{j+1}\}$$

for any j, we conclude that each of the Y_i are independent of all of the previous except for Y_{i-1}. The likelihood function must then factorize like

$$f_{Y_{1:n}}(y_1, \ldots, y_n|\theta) = f_{Y_1}(y_1|\theta) \prod_{j=1}^{n-1} f_{Y_{j+1}|Y_j=y_j}(y_{j+1}|\theta).$$

Now we are ready to derive the likelihood function .

Proposition 1. Let $X_1, \ldots, X_n$ be a collection of *i.i.d.* continuous random variables with common PDF f_X and CDF F_X, and define $Y_j := \max_{i \leq j} X_i$. We assume that $f_X \in C^1$. Then the joint likelihood for the sequence $Y_1, \ldots, Y_N$ is given by

$$\mathcal{L}_{Y_{1:n}}(y_1, \ldots, y_N) = \prod_{i \in R} f_X(y_i) \prod_{i \notin R} F_X(y_i)$$

where $R \subseteq \{1, \ldots, n\}$ is the set of indices where the record was broken and a new maximum was established.

Proof: See appendix. $\square$

Proposition 2. Let $X_1, \ldots, X_n$ be a collection of *i.i.d.* continuous random variables with common PDF f_X and CDF F_X, and define $Y_j := \min_{i \leq j} X_i$. We assume that $f_X \in C^1$. Then the joint likelihood for the sequence $Y_1, \ldots, Y_N$ is given by

$$\mathcal{L}_{Y_{1:n}}(y_1, \ldots, y_N) = \prod_{i \in R} f_X(y_i) \prod_{i \notin R} (1 - F_X(y_i))$$

where $R \subseteq \{1, \ldots, n\}$ is the set of indices where the record was broken and a new minimum was established.

Proof: Analogous to proposition 1.

Forecasting future records

Given the likelihood, we can use a Bayesian posterior sampling method like the No-U-Turn Sampler (NUTS) (Homan & Gelman, 2014) to sample the posterior distribution of the model parameters θ. We draw samples from the conditional distribution of the records given the parameters (the posterior predictive distribution), and compare simulated data generated from these samples to the actual data for a sanity check of our model.

Given the distribution over the parameters, we can also generate n new samples of the attempt distribution $X_{N+1}, \ldots, X_{N+n}$ and take the cumulative maximum to generate a distribution of records in future timesteps.

To help us and others with this process, we have developed and released fmax, a Python library built on top of PyMC3 (Salvatier et al., 2016) to model and forecast future series. This expands on our previous article (Sevilla & Lindbloom., 2021).

Empirical results

With our theoretical framework established, we now study its applications with real-world data.

We will study the application of the framework to extrapolate the world record times for six athletic events (mile run, 1000 meters, 5000 meters, 10000 meters, 20000 meters, and marathon).

This is the same data discussed in (Tryfos & Blackmore, 1985). We extend the dataset they used with data up until the present day, and also include data from the corresponding women's events. We gathered the data from the World Athletics sports federation. A snapshot of the data is available in Figure 1.

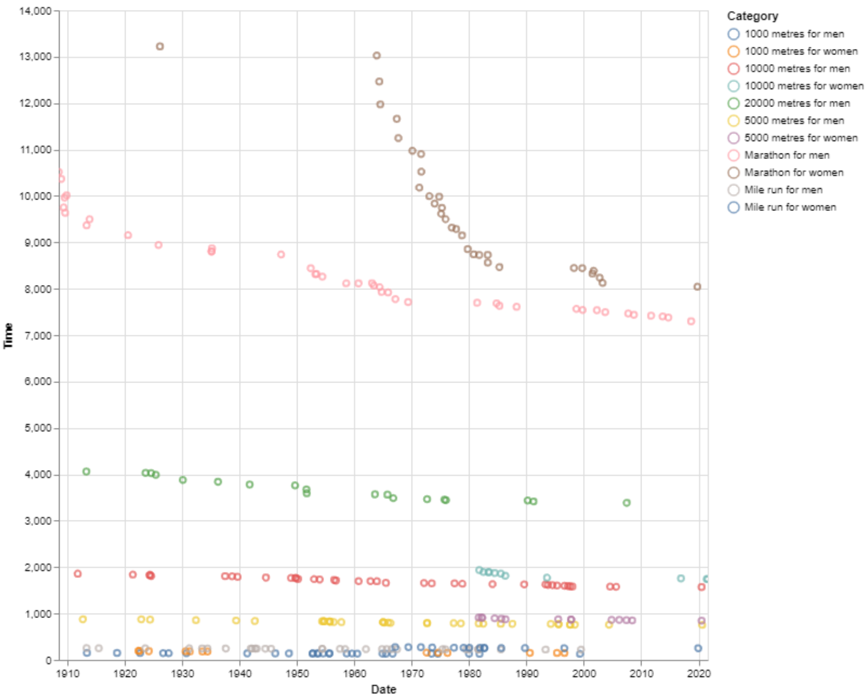

Figure 1: World records for 11 categories of athletic events. The vertical axis indicates the record times in seconds. Source: [X]

We use this data to:

1. Compare the results of the Bayesian posterior with a Gaussian attempt distribution to the reported results of the MMSE estimator.
2. Study the effects of modeling the problem using different attempt distributions and parameter priors.
3. Produce forecasts for the record progressions for each event over the course of the next decade.

Comparison with the Tryfos MMSE approach

(Tryfos & Blackmore, 1985) provided forecasts for future records of the six men categories between 1983 and 1997, for their model fit to the records in the previous years 1968 to 1982. Their approach is to derive the Minimum Mean Square Error (MMSE) estimator assuming an underlying Gumbel distribution of attempts. Here we compare the results from their approach to an analogous Bayesian model using our approach with a Gumbel attempt distribution. Note that the MMSE estimator of (Tryfos & Blackmore, 1985) is a maximum likelihood estimate that does not incorporate a prior for the solution. To make an accurate comparison, we opt to use highly uninformative priors for the parameters of the underlying Gumbel attempt distribution used in our approach.

Consider the mile run event. In Figures 2, 3, and 4, we present the results of the Bayesian model on the mile run in comparison to the MMSE estimator. Here we have used the No-U-Turn Sampler (NUTS) implementation of PyMC3 to compute 10 independent chains of 25,000 samples each, with posterior statistics calculated using samples from all chains.

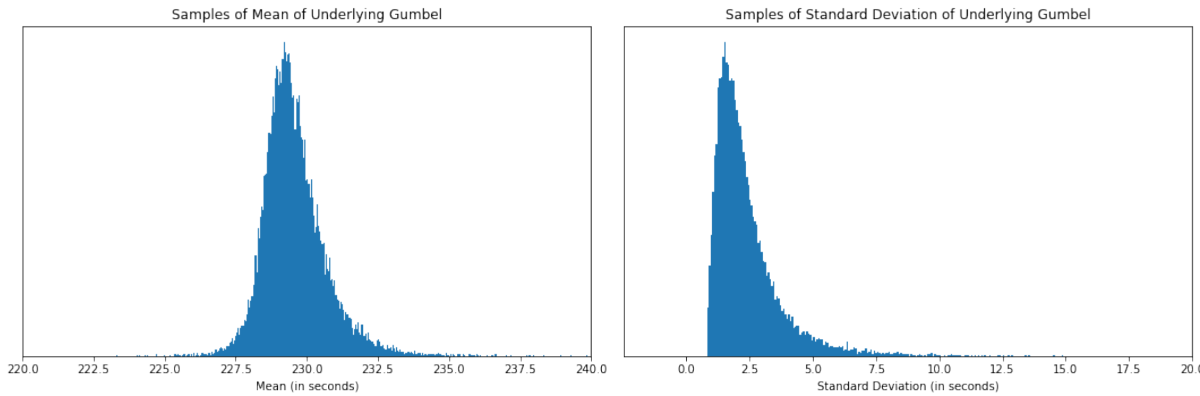

Figure 2: Histograms of samples of the marginal posterior distributions on the mean and standard deviation of the underlying Gumbel attempt distribution for the mile run.

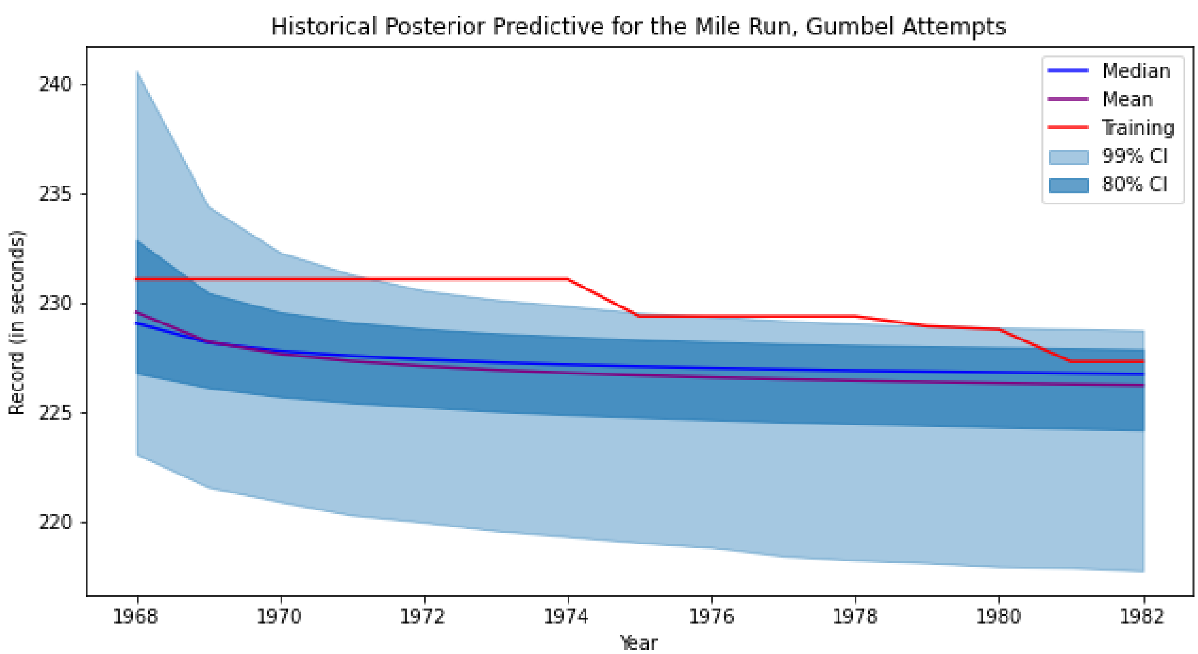

Figure 3: The posterior predictive distribution over the historical records for the mile run, using a Gumbel attempt distribution. The red curve tracks the historical record, while the blue curves summarize the posterior.

Event	Tryfos MMSE	Bayes Mean	Constant Baseline
Mile Run	1.923	**1.135**	3.425
1000 Meters	0.310	1.461	**0.000**
5000 Meters	43.371	**41.053**	65.494
10000 Meters	496.971	**427.221**	607.342
20000 Meters	192.800	**108.640**	383.957
Marathon	3500.800	**981.312**	5341.133

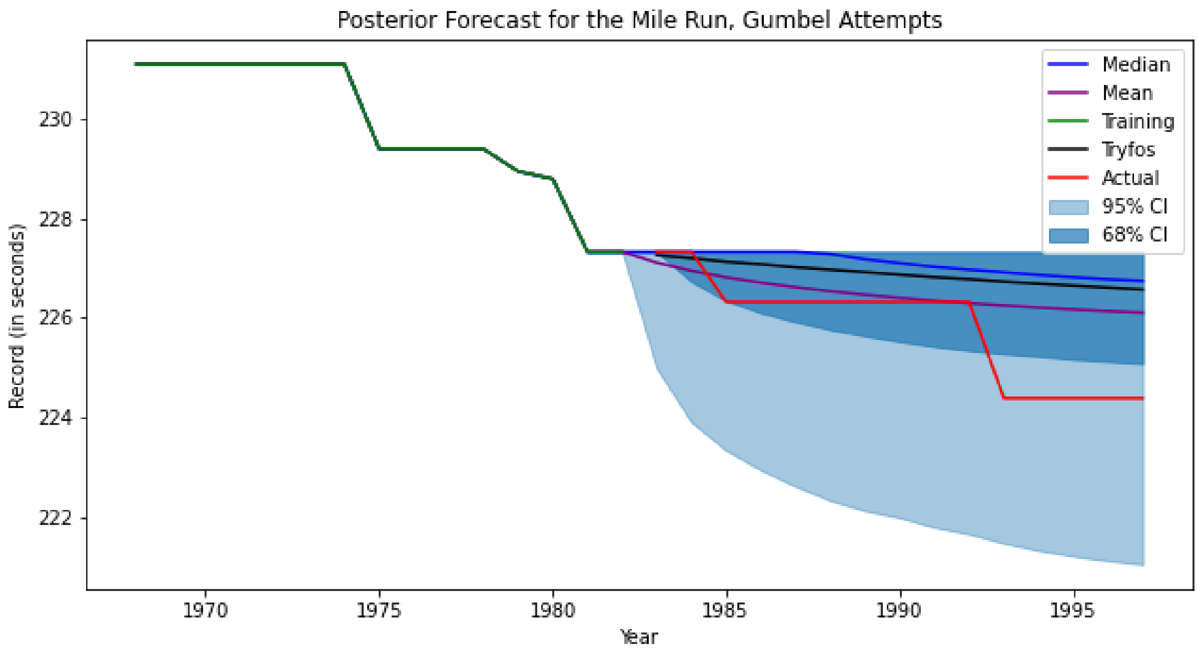

Figure 4: The posterior predictive distribution over the future records for the mile run, using a Gumbel attempt distribution. The red curve tracks the historical record, while the blue curves summarize the posterior. The actual observed record over the forecasted period (1983-1997) is shown in black.

We repeat this same calculation for each of the remaining events in the Tryfos dataset, but with only 5000 posterior samples per chain since we observe little difference in our results compared with the longer chains. To compare the performance of the Bayesian posterior with the MMSE, we compute the MSE of both the posterior mean and the MMSE estimator compared to the actual observed records over the forecasted period (1983-1997). We also include in our comparison a constant baseline estimator that always predicts the last observed record.

The performance of the Tryfos MMSE estimator and the mean of the Bayesian posterior are comparable, both strictly better than the constant baseline (except for the 1000 meters event, where the observed record saw no change during the hold-out period). The largest difference between the two predictors was observed for the Marathon event, where the posterior mean estimator performed significantly better than both Tryfos MMSE and the constant baseline. While the magnitude of the difference is partly driven by the scale of the marathon times, in Figure 5 we see that the posterior median estimator correctly tracks the actual observed records.

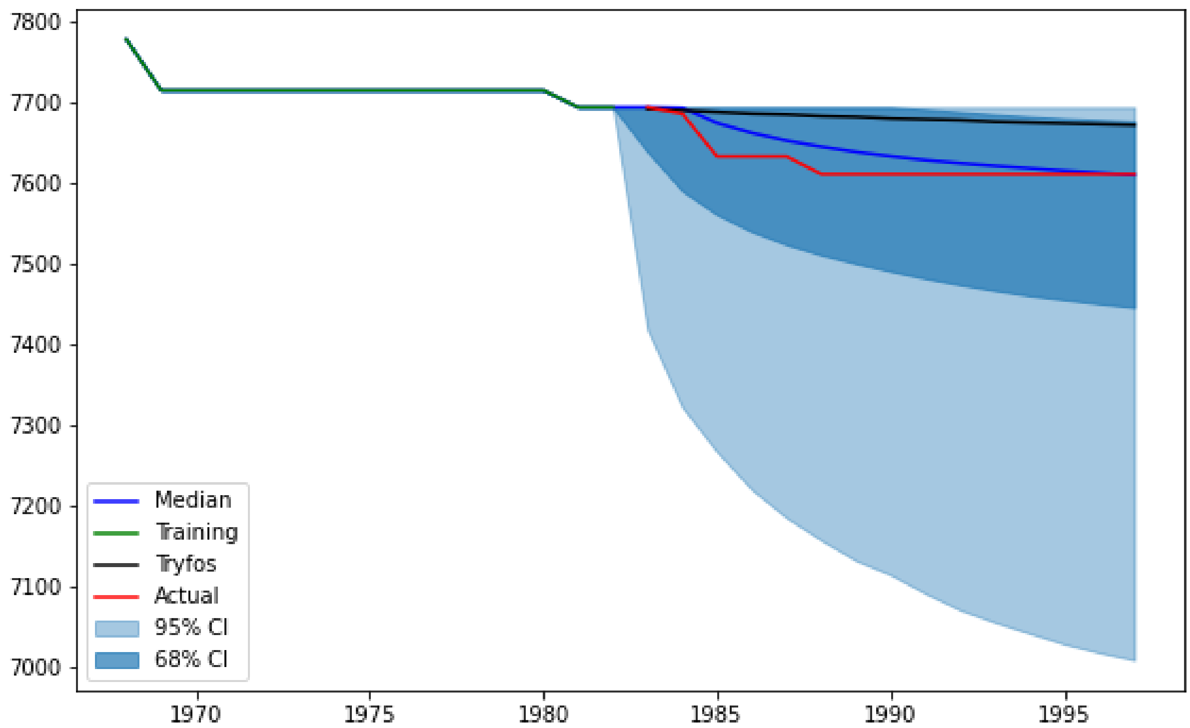

Figure 5: The posterior predictive distribution over the future records for the Marathon event, using a Gaussian attempt distribution.

Our approach allows us to produce credence intervals to quantify the uncertainty in our predictions. In this comparison, we cannot compare the credence intervals produced from the Trfyos MMSE estimator approach with our approach since they were not stated alongside their predictions.

We emphasize that even though the Tryfos MMSE estimator approach can produce credence intervals, they express uncertainty only in sampling variability and do not capture parametric uncertainty. Our approach also captures parametric uncertainty in the fit of our model, which provides a possible explanation for why the posterior mean estimator for the Marathon event performs better than the Tryfos MMSE.

Comparison of different attempt distributions

Our approach extends seamlessly to other attempt distributions. In this section, we repeat the same exercise on the Tryfos dataset but instead consider the effect of varying the form of the attempt distribution. Specifically, we compare the performance of a Gaussian attempt distribution with that of a Gumbel and of a Weibull attempt distribution. For this experiment, we again choose uninformative priors for each of the parameters of the underlying attempt distributions.

From our results, we cannot draw any strong conclusions about whether one of the underlying distributions tends to perform better than the others. However, the model using a Weibull attempt distribution appears to give a good balance between predictive accuracy and maximizing the average log probability.

These results match common sense - the Weibull distribution arises whenever we take the minimum out of a series of samples from an i.i.d. distribution with a finite lower bound (Fréchet, 1927). At each time step,

	Gaussian		Gumbel		Weibull
Event	Avg. Log Probability	MSE	Avg. Log Probability	MSE	Avg. Log Probability
Men's Mile Run	-11.794	0.462	-33.634	1.135	**-10.657**
Men's 1000 Meters	-2.486	4.443	**-2.333**	**1.461**	-2.808
Men's 5000 Meters	-42.163	33.407	-238.417	41.053	**-35.860**
Men's 10000 Meters	-133.376	**350.226**	-21846.803	427.221	**-77.867**
Men's 20000 Meters	**-16.289**	213.791	-21.031	**108.640**	-16.481
Men's Marathon	**-26.639**	**162.073**	-39.445	981.312	-26.733
Women's Mile Run	**-16.773**	7.754	-18.008	**0.865**	-17.554
Women's 1000 Meters	-16.975	436.421	**-16.120**	**93.398**	-17.477
Women's 5000 Meters	-54.615	565.344	-166.077	668.257	**-42.062**
Women's 10000 Meters	-36.218	**358.348**	-40.090	2722.824	**-36.177**
Women's Marathon	-17.989	**7129.153**	-18.243	24141.174	**-17.620**

the attempt to beat the record reflects the minimum time out of all runners that year. Since runs cannot obtain a time less than 0, the conditions are met for a Weibull distribution.

Predicting future records

In this section, we provide our best attempt at predicting the records for the next decade using our proposed framework. We fit Weibull models using all available records since 1968, presenting forecasts for the years 2022-2032. We hope this will help evaluate the performance of the method a decade from now, and other researchers who might want to benchmark their results against ours. We report our results for each year as a credence interval. We highlight the 5%, 15%, 50%, 85%, and 95% quantiles of the prediction. An example forecast is presented in Figure 6, and a panel of the forecasts for all of the men's events is presented in Figure 7. Tables containing the forecasts for all events, as well as individual plots for each event, can be found in Appendix B.

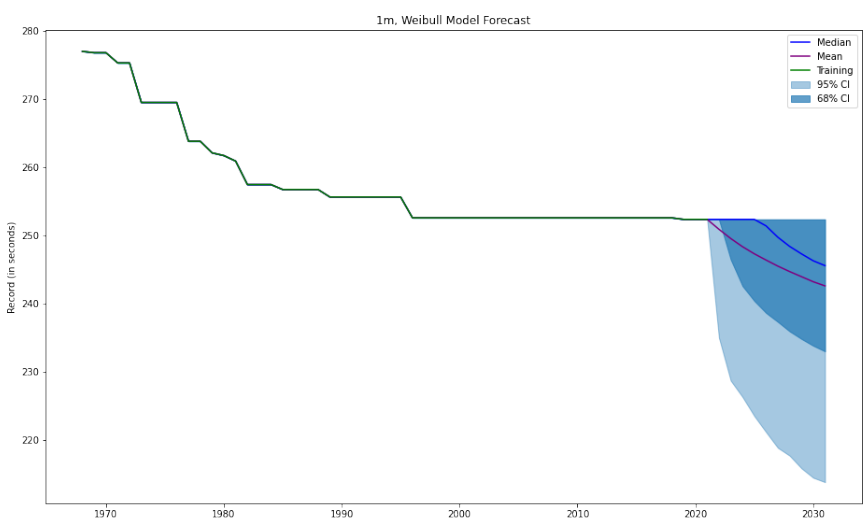

Figure 6: Forecast for the women's mile run event, using a Weibull attempt model.

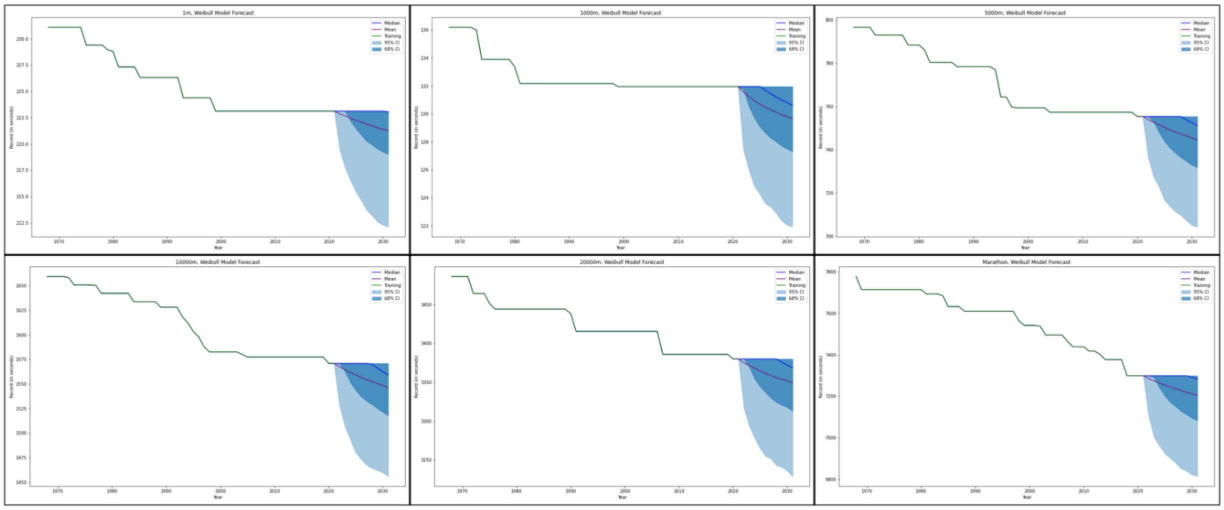

Figure 7: A collection of forecasts for the men's events. Full resolution plots available in Appendix B.

Conclusion

Predicting records is of paramount importance for science - with use cases ranging from predicting temperature records to forecasting the development of new technologies.

In this article, we have developed a Bayesian account for predicting records. Our work is quite general and extends to all situations in which the underlying attempts follow an i.i.d. distribution.

We have shown that our approach is competitive with a previous frequentist method by (Tryfos & Blackmore, 1985) in terms of MSE for predicting records in 6 athletic events.

We have also investigated the predictive accuracy of different attempt distributions on data for 11 athletic events. While the conclusions aren't clear-cut, the evidence suggests that the Weibull distribution results in a better fit for the data in terms of a log-likelihood loss.

Using our method we have forecasted the records for these 11 athletic events for the period of 2022 to 2032. We hope other researchers will be able to use this as a basis for comparison.

We have released an open-source PyMC3 package accompanying this paper, fmax. Researchers and practitioners can use this framework to model record distributions using their own data.

Acknowledgments

Jaime Sevilla works for Epoch and is funded by the Open Philanthropy Project. Jonathan Lindbloom is supported by the Dartmouth College mathematics department.

Appendix A: Proof of proposition 1

To reason about the conditional density $f_{Y_{i+1}|Y_i=y_i}$ it is useful to consider the procedure of drawing a sample of the corresponding random variable . Suppose at time $t = j$ the observed record for the minimum is y_j. Then, given that the CDF of X_{j+1} is $F_{X_{j+1}} = F_X$, the probability that the record Y_{j+1} exceeds y_j is

$$P(Y_{j+1} \geq Y_j | Y_j = y_j) = 1 - F_X(y_j)$$

and in this case we must have $y_{j+1} = y_j$ since the attempt doesn't change the record for the minimum. In the alternative case that Y_{j+1} is less than y_j, a sample is drawn from the distribution of Xconditioned on the event $X \leq y_j$. Let Z_j denote a random variable that admits a PDF

$$f_{Z_j}(z) = \frac{1}{F_X(y_{j-1})} f_X(z) \chi_{z \leq y_{j-1}}$$

where χ denotes the indicator function. Then our sampling procedure can be summarized as

$$u amp; \sim \text{Bernoulli}(F_X(y_j)),$$
$$y_{j+1} amp; \sim \begin{cases} Z_j, & amp; \text{if } u = 1, \\ \text{Constant}(y_j), & amp; \text{if } u = 0. \end{cases}$$

From this expression it is clear that the random variable $Y_{j+1}|Y_j = y_j$ is a *mixed* random variable, meaning that it consists of both a discrete and a continuous component. The PDF of such a random variable involves Dirac delta functions and the CDF has jump discontinuities. Omitting its derivation (see the Appendix), the likelihood function for a sequence of observed records $\{r_1, r_2, \ldots, r_n\}$ can be expressed as

$$f_{Y_{1:n}}(r_1, \ldots, r_n|\theta) amp; = \left(\prod_{j \in C} F_X(r_{j-1}) f_{Z_j}(r_j) \right) \left(\prod_{j \in D} (1 - F_X(r_{j-1})) \right)$$
$$amp; = \left(\prod_{j \in C} F_X(r_{j-1}) \cdot \frac{1}{F_X(r_{j-1})} f_X(r_j) \chi_{r_j \leq y_{j-1}} \right) \left(\prod_{j \in D} (1 - F_X(r_j)) \right)$$
$$amp; = \left(\prod_{j \in C} f_X(r_j) \right) \left(\prod_{j \in D} (1 - F_X(r_j)) \right)$$

where C denotes the set of time indices for the records which changed the running minimum and D denotes the set of time indices for the records which didn't change the running minimum. Note that we can determine the sets C and D by checking successively checking whether or not arecord changed since the previous record. Note also that we drop the factors $\chi_{z \leq y_{j-1}}$ from the likelihood since they are redundant ; $j \in C$ if and only if $r_j \leq y_{j-1}$. Presuming that we can evaluate the CDF and PDF of X for any θ, the evaluation of this form of the likelihood is straightforward. If we seek a model for a running maximum rather than a minimum, we can make a similar argument to find that the likelihood in this case becomes

$$f_{Y_{1:n}}(r_1, \ldots, r_n|\theta) amp; = \left(\prod_{j \in C} (1 - F_X(r_{j-1})) \cdot \frac{1}{(1 - F_X(r_{j-1}))} f_X(r_j) \chi_{r_j \geq y_{j-1}} \right) \left(\prod_{j \in D} F_X(r_j) \right)$$
$$amp; = \left(\prod_{j \in C} f_X(r_j) \right) \left(\prod_{j \in D} F_X(r_j) \right)$$

with the appropriate modifcations in notation for switching to the maximum.

Appendix B: Forecasts for all events using a Weibull attempt distribution.

The plots show the historical record in green, and the forecasted distribution in blue. In purple we show the mean.

In the tables we include the 5%, 15%, 50%, 85%, 95% percentiles and the mean of the distribution for each year between 2023 and 2032.

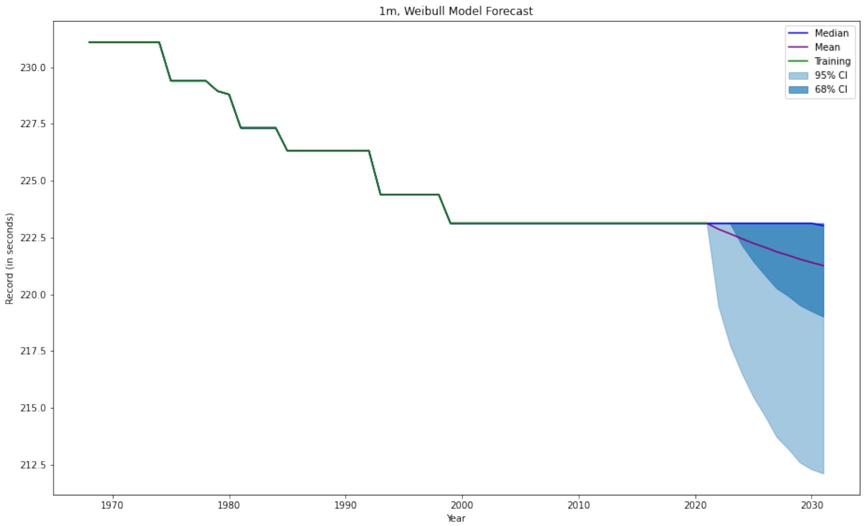

Figure 8: Forecast for the men's mile run event, using a Weibull attempt model.

Men's Mile Run										
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032
5%	221.700	219.869	218.541	217.771	217.238	216.563	216.140	215.528	215.176	214.816
15%	223.130	223.032	222.037	221.249	220.630	220.089	219.710	219.338	219.067	218.806
50%	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.023
85%	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130
95%	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130	223.130
Mean	222.872	222.667	222.452	222.247	222.070	221.875	221.717	221.549	221.402	221.270

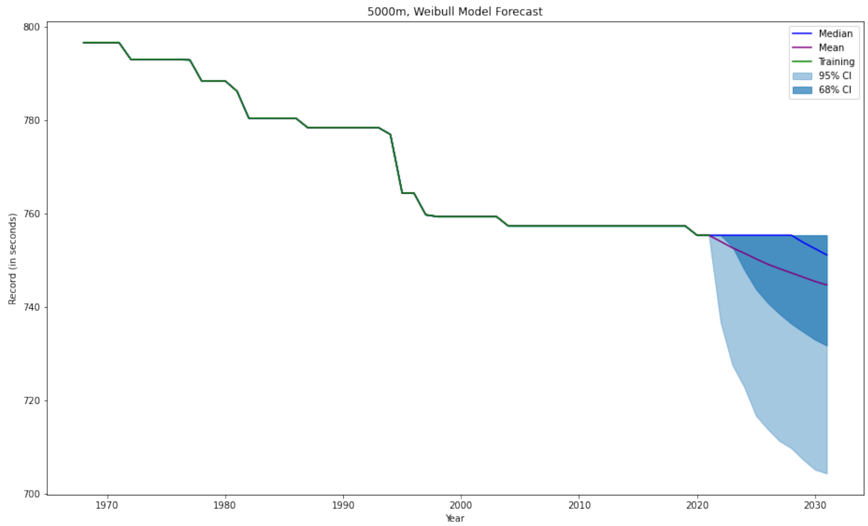

Figure 9: Forecast for the men's 5000 meter run event, using a Weibull attempt model.

Men's 5000 M										
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032
5%	746.244	735.782	731.595	727.079	723.798	721.009	719.030	717.754	715.459	714.333
15%	755.360	751.862	746.984	742.987	739.913	737.276	735.411	733.734	732.101	730.642
50%	755.360	755.360	755.360	755.360	755.360	755.360	755.360	753.827	752.492	751.146
85%	755.360	755.360	755.360	755.360	755.360	755.360	755.360	755.360	755.360	755.360
95%	755.360	755.360	755.360	755.360	755.360	755.360	755.360	755.360	755.360	755.360
Mean	754.027	752.664	751.529	750.305	749.156	748.189	747.294	746.391	745.469	744.723

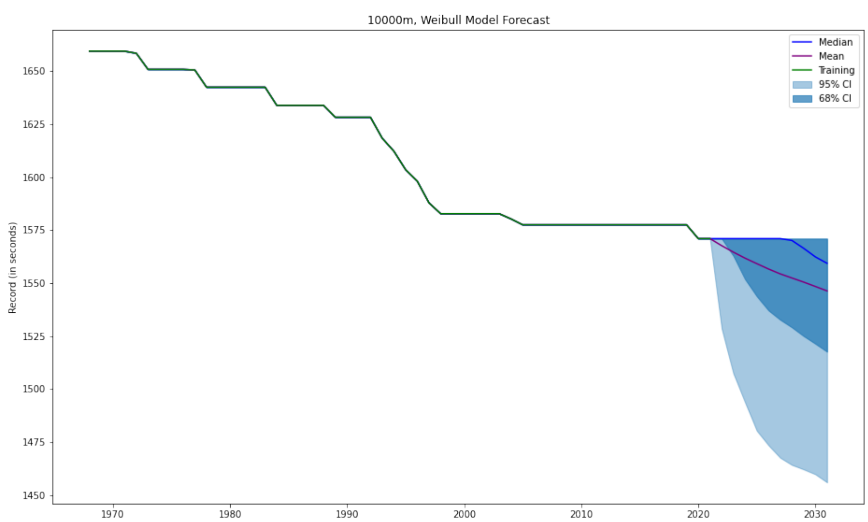

Figure 10: Forecast for the men's 10000 meter run event, using a Weibull attempt model.

Men's 10000 M									
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031
5%	1549.133	1529.001	1515.671	1507.667	1499.729	1493.763	1488.453	1483.908	1480.733
15%	1571.000	1560.467	1549.666	1541.676	1535.418	1530.620	1526.658	1522.914	1519.520
50%	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1570.128	1566.447	1562.402
85%	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000
95%	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000	1571.000
Mean	1567.684	1564.612	1561.785	1559.197	1556.670	1554.383	1552.431	1550.498	1548.441

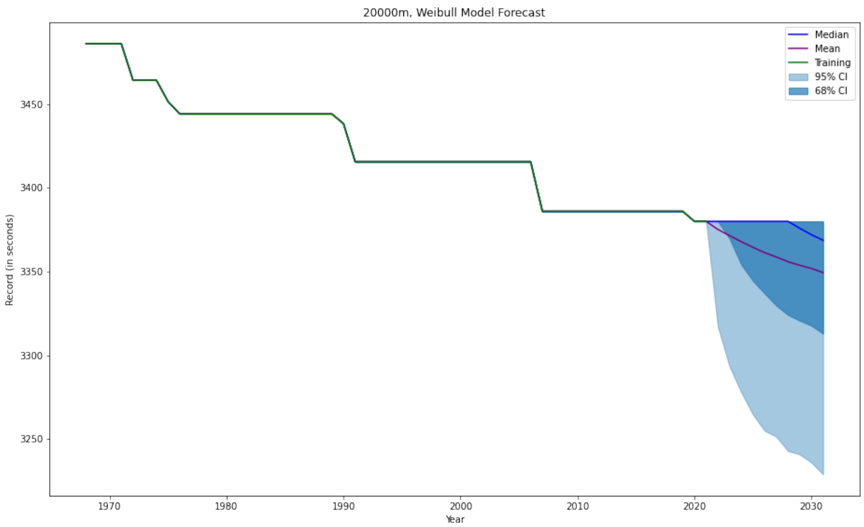

Figure 11: Forecast for the men's 20000 meter run event, using a Weibull attempt model.

Men's 20000 M									
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031
5%	3346.052	3322.007	3306.363	3294.180	3286.358	3279.879	3271.824	3270.657	3265.552
15%	3380.020	3366.939	3351.314	3341.700	3332.459	3326.468	3321.829	3317.942	3314.704
50%	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3375.903	3372.039
85%	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020
95%	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020	3380.020
Mean	3375.214	3371.503	3367.896	3364.506	3361.333	3358.737	3355.898	3353.788	3351.885

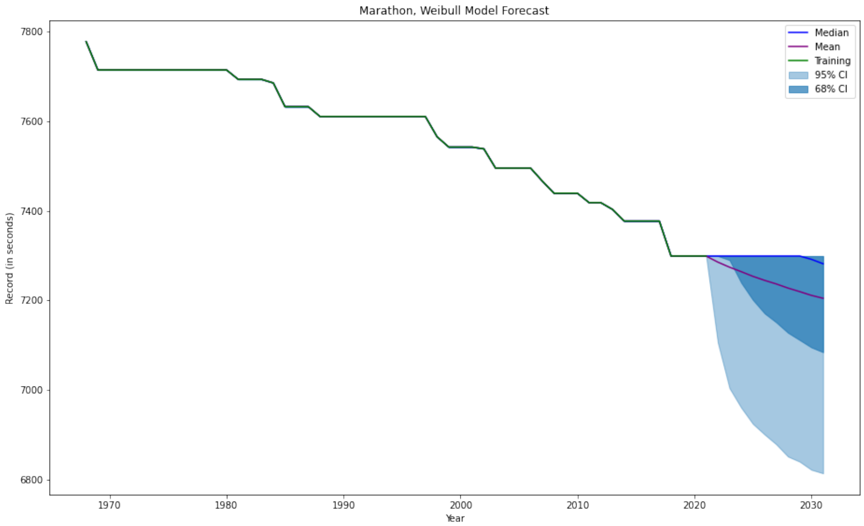

Figure 12: Forecast for the men's Marathon event, using a Weibull attempt model.

Men's Marathon									
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031
5%	7210.013	7120.834	7072.999	7025.397	6995.733	6977.074	6949.472	6934.098	6913.576
15%	7299.000	7282.595	7230.468	7191.317	7161.068	7141.663	7119.901	7100.728	7085.708
50%	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7291.665
85%	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000
95%	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000	7299.000
Mean	7285.462	7274.152	7264.177	7253.715	7244.807	7236.743	7227.567	7219.616	7211.340

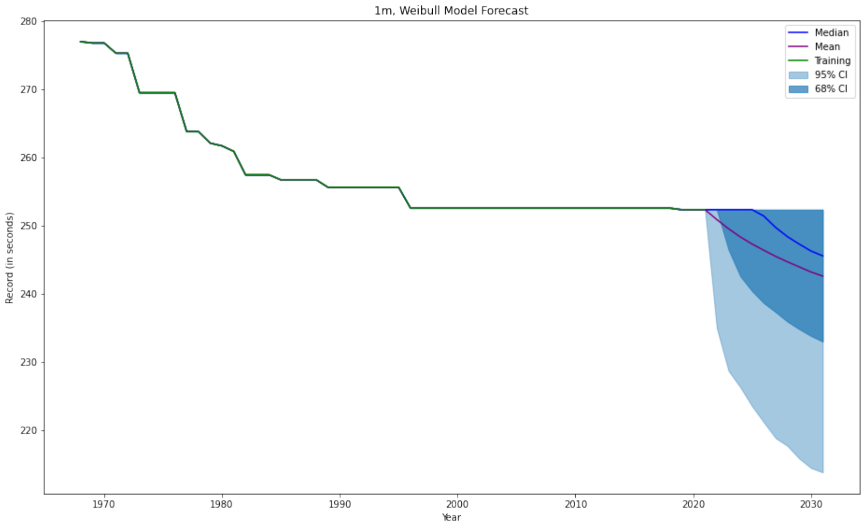

Figure 13: Forecast for the women's mile run event, using a Weibull attempt model.

Women's Mile Run										
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032
5%	241.485	234.801	231.596	229.318	227.781	226.558	224.863	223.427	221.937	220.9
15%	252.283	245.752	241.983	239.613	238.127	236.486	235.269	234.183	233.128	232.2
50%	252.330	252.330	252.330	252.330	251.367	249.670	248.360	247.266	246.249	245.5
85%	252.330	252.330	252.330	252.330	252.330	252.330	252.330	252.330	252.330	252.3
95%	252.330	252.330	252.330	252.330	252.330	252.330	252.330	252.330	252.330	252.3
Mean	250.877	249.520	248.320	247.274	246.355	245.462	244.665	243.919	243.170	242.5

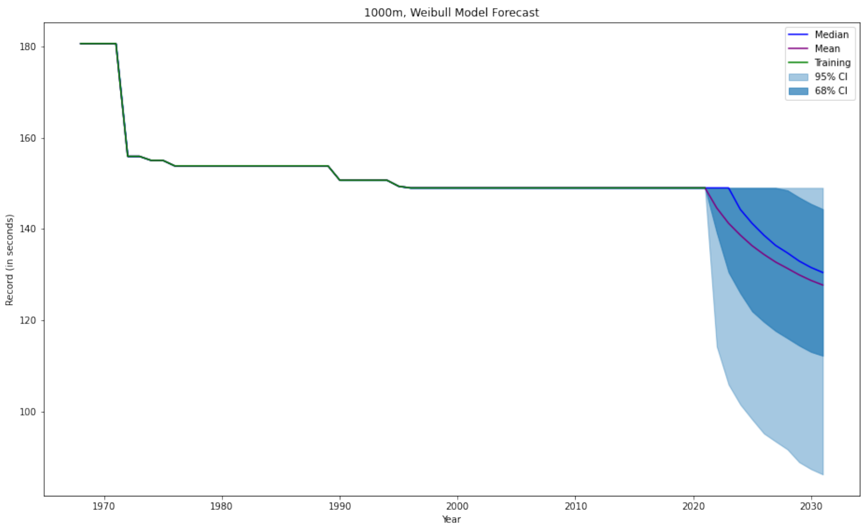

Figure 14: Forecast for the women's 1000 meter run event, using a Weibull attempt model.

Women's 1000 M										
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032
5%	123.219	115.159	110.267	106.359	103.893	101.900	100.132	98.667	96.914	95.522
15%	138.311	129.567	124.976	121.006	118.971	116.652	115.022	113.370	112.215	111.413
50%	148.980	148.980	144.257	141.208	138.617	136.360	134.727	132.935	131.534	130.445
85%	148.980	148.980	148.980	148.980	148.980	148.980	148.908	147.303	145.996	144.783
95%	148.980	148.980	148.980	148.980	148.980	148.980	148.980	148.980	148.980	148.980
Mean	144.572	141.231	138.629	136.306	134.407	132.723	131.338	129.911	128.699	127.696

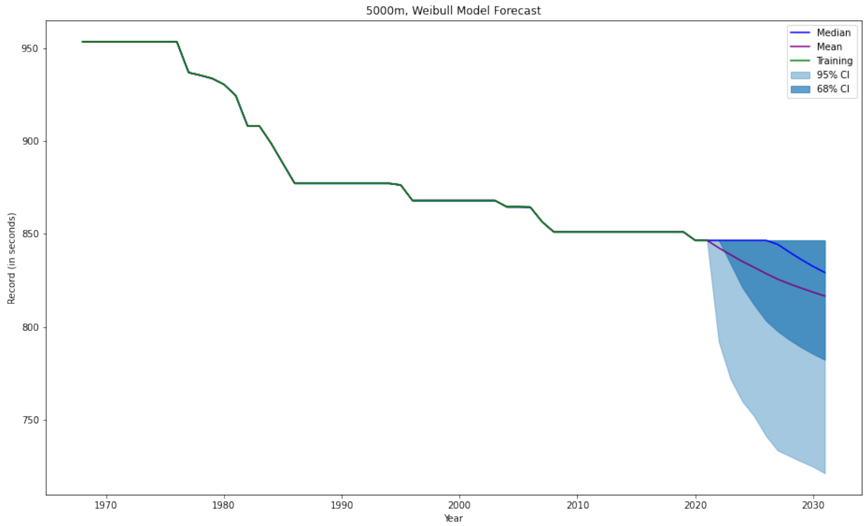

Figure 15: Forecast for the women's 5000 meter run event, using a Weibull attempt model.

Women's 5000 M										
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032
5%	815.144	792.258	779.824	772.488	765.834	758.765	754.099	750.127	745.510	741.99
15%	846.620	831.699	818.445	809.310	801.140	794.899	790.372	786.619	782.870	780.14
50%	846.620	846.620	846.620	846.620	846.620	844.468	840.244	836.266	832.596	829.30
85%	846.620	846.620	846.620	846.620	846.620	846.620	846.620	846.620	846.620	846.62
95%	846.620	846.620	846.620	846.620	846.620	846.620	846.620	846.620	846.620	846.62
Mean	842.557	838.839	835.247	832.066	828.752	825.707	823.162	820.880	818.680	816.64

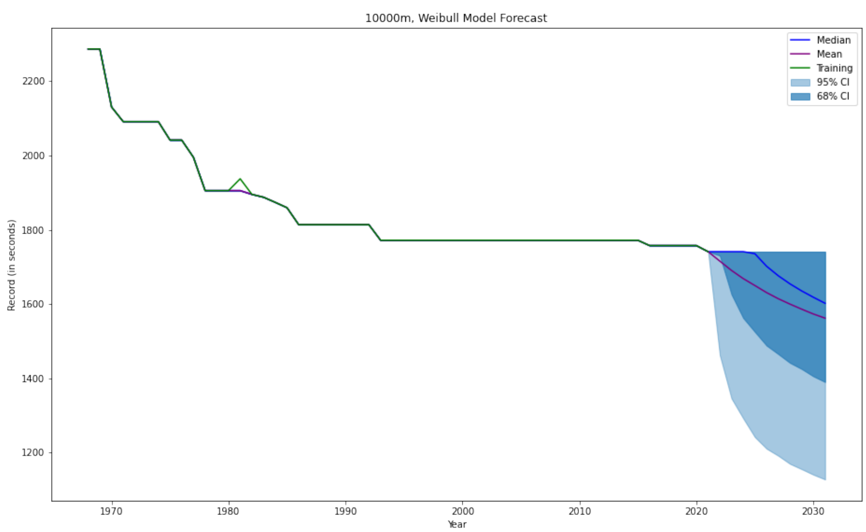

Figure 16: Forecast for the women's 10000 meter run event, using a Weibull attempt model.

Women's 10000 M									
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031
5%	1550.304	1448.169	1384.777	1345.584	1313.976	1292.429	1269.354	1248.860	1231.7
15%	1720.029	1614.853	1549.713	1515.500	1479.115	1455.277	1432.242	1414.355	1393.7
50%	1741.030	1741.030	1741.030	1735.545	1701.818	1676.206	1654.447	1635.369	1618.3
85%	1741.030	1741.030	1741.030	1741.030	1741.030	1741.030	1741.030	1741.030	1741.0
95%	1741.030	1741.030	1741.030	1741.030	1741.030	1741.030	1741.030	1741.030	1741.0
Mean	1715.697	1690.610	1668.621	1649.936	1630.922	1614.456	1599.926	1586.273	1573.3

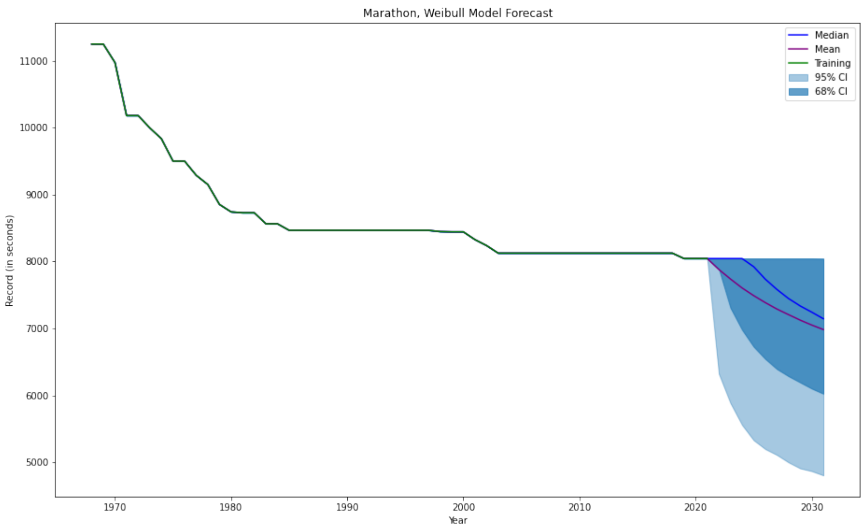

Figure 17: Forecast for the women's marathon event, using a Weibull attempt model.

Women's Marathon									
Year	2023	2024	2025	2026	2027	2028	2029	2030	2031
5%	6861.300	6359.315	6031.158	5803.483	5642.323	5550.896	5435.129	5310.331	5254.
15%	7826.202	7255.647	6913.038	6664.624	6488.961	6335.872	6245.672	6128.398	6048.
50%	8044.000	8044.000	8044.000	7922.313	7736.367	7582.565	7446.883	7336.415	7243.
85%	8044.000	8044.000	8044.000	8044.000	8044.000	8044.000	8044.000	8044.000	8044.
95%	8044.000	8044.000	8044.000	8044.000	8044.000	8044.000	8044.000	8044.000	8044.
Mean	7880.817	7738.886	7604.671	7490.326	7385.662	7289.743	7206.064	7124.282	7050.

References

Forecasting Records. (1985). *Journal of the American Statistical Association*, *80*(389), 46–50. https://doi.org/10.2307/2288039

Forecasting Records by Maximum Likelihood. (1988). *Journal of the American Statistical Association*, *83*.

A Non-Gaussian State Space Model and Application to Prediction of Records. (1986). *Journal of the Royal Statistical Society. Series B (Methodological)*, *48*.

Record occurrence and record values in daily and monthly temperatures. (2013). *Climate Dynamics*, *42*(5-6), 1275–1289. https://doi.org/10.1007/s00382-013-1693-0

Objective Bayesian Prediction of Future Record Statistics Based on the Exponentiated Gumbel Distribution: Comparison with Time-Series Prediction. (2020). *Symmetry*, *12*(9). https://doi.org/10.3390/sym12091443

The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. (2014). *J. Mach. Learn. Res.*, *15*.

Probabilistic programming in Python using PyMC3. (2016). *PeerJ Comput. Sci.*, *2*, e55.

Modelling a Time Series of Records with PyMC3. (2021). *Authorea*.

Sur la loi de probabilité de l'écart maximum. (1927). *Annales De La Société Polonaise De Mathématique*, *6*.