



Using rare event algorithms to understand the statistics and dynamics of extreme heatwave seasons in South Asia

Clément Le Priol, Joy M. Monteiro, Freddy Bouchet

► To cite this version:

Clément Le Priol, Joy M. Monteiro, Freddy Bouchet. Using rare event algorithms to understand the statistics and dynamics of extreme heatwave seasons in South Asia. 2024. hal-04560861v1

HAL Id: hal-04560861

<https://hal.science/hal-04560861v1>

Preprint submitted on 26 Apr 2024 (v1), last revised 1 Dec 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike 4.0 International License

Using rare event algorithms to understand the statistics and dynamics of extreme heatwave seasons in South Asia

Clément Le Priol^{1,2}, Joy M. Monteiro^{3,4}, Freddy Bouchet²

¹ Laboratoire de physique à l'ENS de Lyon, UMR 5672, Lyon, France

² Laboratoire de Météorologie Dynamique, IPSL, ENS-PSL, CNRS, UMR 8539, Paris, France

³ Department of Earth and Climate Science, Indian Institute of Science Education and Research Pune, Pune, India

⁴ Department of Data Science, Indian Institute of Science Education and Research Pune, Pune, India

E-mail: `clement.le-priol@lmd.ipsl.fr`

Abstract.

Computing the return times of extreme events and assessing the impact of climate change on such return times is fundamental to extreme event attribution studies. However, the rarity of such events in the observational record makes this task a challenging one, even more so for "record-shattering" events that have not been previously observed at all. While climate models could be used to simulate such extremely rare events, such an approach entails a huge computational cost: gathering robust statistics for events with return time of centuries would require a few thousand years of simulation.

In this study, we use an innovative tool, rare event algorithm, that allows to sample numerous extremely rare events at a much lower cost than direct simulations. We employ the algorithm to sample extreme heatwave seasons, corresponding to large anomalies of the seasonal average temperature, in a heatwave hotspot of South Asia using the global climate model Plasim. We show that the algorithm estimates the return levels of extremely rare events with much greater precision than traditional statistical fits. It also enables the computation of various composite statistics, whose accuracy is demonstrated through comparison with a very long control run. In particular, our results reveal that extreme heatwave seasons are associated with an anticyclonic anomaly embedded within a large-scale hemispheric quasi-stationary wave-pattern. Additionally, the algorithm accurately represents the intensity-duration-frequency statistics of sub-seasonal heatwaves, offering insights into both seasonal and sub-seasonal aspects of extreme heatwave seasons. This innovative approach could be used in extreme event attribution studies to better constrain the changes in event's probability and intensity with global warming, particularly for events with return times spanning centuries or millennia.

1. Introduction

Due to global warming, most regions in the world are experiencing an increase in the frequency and intensity of heatwaves [Seneviratne et al., 2021] with dramatic consequences on human health [Romanello et al., 2021, Dimitrova et al., 2021], crop yields [Wegren, 2011] and ecosystems [Ciais et al., 2005]. Therefore, understanding changes in the probability of occurrence and characteristics of heatwaves is crucial to climate adaptation efforts globally. While the mechanisms leading to heatwaves are well understood, research remains to be done to assess at the regional scale the relative importance of each mechanism and how they will be affected by global warming [Perkins, 2015, Horton et al., 2016].

South Asia is particularly vulnerable to the impacts of extreme heat. This vulnerability is not just because of the magnitude of heatwaves experienced in this region but also because of its high population density which results in a many people being exposed to extreme heat [Romanello et al., 2021].

In South Asia, while there does not seem to be a significant increase in exposure to extreme heat (as measured by air temperature alone) over the historical record [Rogers et al., 2021], climate models project that exposure to extreme heat will increase under all emission scenarios [Mora et al., 2017, Im et al., 2017, Mishra et al., 2017]. However, few studies focused on the understanding of their mechanisms and of the relative importance of their main precursors [Ratnam et al., 2016, Rohini et al., 2016, Monteiro and Caballero, 2019].

Complementary to traditional research into the meteorology of extreme events, the field of extreme event attribution (EEA) focuses on assessing the influence of anthropogenic climate change on individual extreme weather events. While the way in which event attribution is framed is debated [Hulme, 2020], it is generally accepted that understanding the changes in probability and intensity of the events and associated physical processes at the regional scale are important both from a scientific and societal perspective.

Extreme event attribution research comes in different flavours. Most studies focus on estimating the changes in probability and/or intensity of a given event in the present world with respect to a "counterfactual" world where the global temperature would have remained to its preindustrial value [Stott et al., 2004, Philip et al., 2020, Zachariah et al., 2023]. This is known as the *probability-based* approach. Other studies follow the *storyline* approach which focus on the causal chain leading to the extreme event and how climate change has affected the drivers, like sea surface temperature (SST) or soil moisture deficit for heatwaves [Hoerling et al., 2013, Trenberth et al., 2015, Shepherd, 2016]. Finally, some works condition on the atmospheric circulation associated with the event and evaluate how climate change affected temperature, precipitation and wind speeds conditioned on this circulation [Jézéquel et al., 2018, Faranda et al., 2022]. They also provide information about how rare such atmospheric states are [Faranda et al., 2023].

In the probability-based approach, the estimation of unconditional probabilities, or return levels, usually relies on extreme value theory (EVT). The underlying assumption is that the ground truth distribution of the events lies in the basin of attraction of one of the generalised extreme value distributions [Coles, 2001]. EVT allows to perform statistical extrapolations to compute the probabilities of unseen events. However, it does not actually sample them. A second limitation of EVT is that the probabilities of unseen, record-shattering, events can be extremely difficult to estimate. For instance, the temperature of the record-shattering 2021 Pacific Northwest heatwave falls outside the range of possible values that could be estimated from the previous years, implying huge uncertainties on estimates of that event’s return time [Philip et al., 2022]. The possibility of such an event could be anticipated using climate models but without knowledge of its probability [Fischer et al., 2021]. [Noyelle et al., 2023] also showed that the maximal reachable temperature in western Europe are, in many locations, higher than the upper bounds provided by EVT fits. As record-shattering events are among the most detrimental [Robine et al., 2008], they should be the topic of active research. In this work, we introduce a new tool, namely *rare event algorithms*, that allows to actually sample extreme events in climate models and provides an unbiased estimator of their probabilities with narrow confidence intervals on the return levels up to return times of millennia or even more. We argue that this tool could complement EVT in extreme event attribution, especially regarding the most extreme events, the ones with return times of centuries or millennia.

1.1. Rare event algorithms

In numerous fields, *rare event algorithms* have been used to improve the sampling of extremely rare events and compute unbiased estimates of their probabilities. These algorithms concentrate the computational effort on the trajectories most likely to lead to the rare event of interest. Such algorithms have recently been applied to climate models to sample extreme hot seasons [Ragone et al., 2018, Ragone and Bouchet, 2020, Ragone and Bouchet, 2021], extreme winter precipitation [Wouters et al., 2023] and extreme tropical cyclones [Plotkin et al., 2019, Webber et al., 2019]. Since these algorithms provide samples of actual model trajectories, they allow computation of composite statistics on extreme events. Such statistics could, for instance, be used to diagnose how global warming affects different drivers of heatwaves, like the large-scale atmospheric circulation, depleted soil moisture, or SST.

An empirical rare event algorithm based on similar ideas has been proposed to simulate extreme hot seasons from observations and reanalysis data [Yiou and Jézéquel, 2020]. We also note that a different resampling protocol, called ensemble boosting, has been proposed to increase the sampling of extreme events in climate models [Gessner et al., 2021], but this protocol does not provide information about the probabilities of the simulated events.

1.2. Our contribution

We use the GKLT algorithm [Ragone et al., 2018, Ragone and Bouchet, 2021, Wouters et al., 2023] to sample extreme heatwave seasons in a subregion of South Asia in the global climate model (GCM) Plasim. We demonstrate the precision and quality of return time curves and composite maps for heatwaves with 100 year return times that were computed via the algorithm by comparing them with the outcome of a very long (8,000 year) control run. For the first time, we compare the return time curve estimate of the algorithm with the one of an EVT fit. We show that the algorithm estimate is both closer to the long control run than the EVT fit and has a much narrower confidence interval.

We show that the algorithm can help to gain insight into the atmospheric circulation associated with extreme heatwave seasons in South Asia. Composite maps created using samples obtained from the algorithm show that extreme heatwaves seasons are associated with an anticyclonic anomaly which is embedded in a large-scale hemispheric quasi-stationary wave-pattern. We also show that the algorithm correctly represents the intensity-duration-frequency statistics of individual subseasonal heatwaves. Thus, not only does the algorithm capture seasonal statistics accurately, but the results can also be used to estimate the distribution of intensity and duration of heatwaves that comprise an extreme heatwave season.

We define the region and event of interest, and describe the model and data used in section 2. Our results are presented in section 4. Finally we discuss the relevance of rare event algorithms for extreme event attribution in section 5 and conclude in section 6. A brief explanation of the algorithm is presented in Appendix A.

2. Event definition, model and data

2.1. Event definition

In South Asia, heatwaves occur during the pre-monsoon season, i.e. March to June, with May-June usually being the hottest months. [Ratnam et al., 2016] identified two distinct heatwave hotspots in India based on the leading EOFs of Tmax variability: one in the northwest and one on the east coast. These two regions are also the ones with the highest incidence of mortality attributed to heat waves (see fig. S9 in [Justine et al., 2023]). Heatwaves on the east coast of India are affected by the Indian and Pacific oceans and in particular by ENSO. These influences are not represented by our model's setting, which has no dynamical ocean. Due to this limitation, we restrict our study to the northwest region.

Some heatwave seasons experience multiple heatwaves, whose impact is observed in the seasonal mean as well. For instance, in the pre-monsoon season of 2022, our region of interest (defined below) experienced an unusually warm season with six extreme heat spells [Dash et al., 2022]. The March-April average temperature was estimated to be a 100 year event in present-day climate [Zachariah et al., 2023]. While this is an illustrative

example, it underscores the importance of looking not just at individual extreme events but also the extremes of long-term averages.

Our region of interest is defined by a box 70°-76°E, 25-34°N, which includes parts of the Indian states of Rajasthan and Punjab and the Pakistani province of Punjab. The region is marked in light blue in Fig. 1. Our target observable is the April-May-June (AMJ) average of the area-averaged 2-meter air temperature:

$$a = \frac{1}{T} \int_{t_0}^{t_0+T} A(t) dt \quad \text{where} \quad A(t) = \frac{1}{\mathcal{A}} \int_{\mathcal{A}} T_{2m}(\mathbf{r}, t) d\mathbf{r} \quad (1)$$

where \mathbf{r} is the space vector, t is time, T_{2m} is the 3-hourly 2-meter air temperature, \mathcal{A} is the area of the region considered (70°-76°E, 25-34°N), t_0 is the 1st of April and T is the AMJ duration, equal to 90 days in the model's setting.

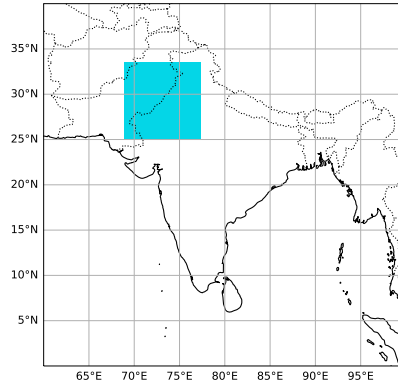


Figure 1. The region studied is marked in light blue.

2.2. Model: Plasim

All simulations are performed with Plasim [Fraedrich et al., 2005], an intermediate complexity climate model developed at the Theoretical Meteorology group of the University of Hamburg. Plasim solves the primitive equations for vorticity, divergence, temperature and logarithm of surface pressure. The subgrid-scale parameterization is simpler than the ones used by top-class GCM's, allowing it to run at dramatically lower computational cost while preserving the correct large-scale dynamics. Plasim also features land, ocean, and sea ice components. We run the model with prescribed seasonal cycles of sea surface temperature (SST) and sea ice, and a fixed CO₂ concentration which is representative of the late 20th century. We set the model at T42 horizontal resolution (about 2.8°) with 10 vertical sigma-levels. While coarse, this resolution is sufficient to capture large-scale atmospheric dynamics. The extremely low computational cost of this configuration (less than 1hCPU.yr⁻¹) makes it well suited for proof of concept and methodological studies. Plasim has been used in several studies of large scale

phenomena, e.g. to assess wind speed reductions by large-scale wind turbine deployment [Miller and Kleidon, 2016] or to study teleconnections and their changes in a changing climate [Herein et al., 2017]. It also proved to be relevant to study large-scale circulations associated with heatwaves: the mode 3 quasistationary Rossby wave pattern associated to European heatwaves observed with Plasim in [Ragone et al., 2018] was confirmed in subsequent studies using CESM [Ragone and Bouchet, 2021] and reanalysis [Miloshevich et al., 2023b].

We produced a control run of 2,000 years (10 independent simulations of 200 years) to serve as initial conditions for the rare event algorithm. A longer and independent control run of 8,000 years (80 independent simulations of 100 years, [Miloshevich et al., 2023a]) is used as a baseline against which the results of the algorithm are compared.

2.3. ERA5 data and detrending

We use ERA5 reanalysis dataset [Hersbach et al., 2020] to assess the ability of Plasim to represent teleconnection between the surface air temperature over the target region and the geopotential height. We use daily 2-meter air temperature (T2m) and geopotential at 500hPa. The temperature is selected at 9h UTC (14h30 Indian time) which corresponds to the approximate time of peak temperature over the target region. The geopotential, which does not have a daily cycle, is selected at 12h UTC. The horizontal resolution is $0.25^\circ \times 0.25^\circ$. We detrended the data by performing a cubic fit to the time series of yearly MAMJ averages of T2m averaged over the target region and of the global average of Zg500 (see fig. D1). The resulting trends were then subtracted uniformly to all days and grid points.

3. Rare event algorithm

In many fields, the return times for events of interest are orders of magnitude larger than the typical time scale of the dynamics. This is the case, for example, in chemical reactions or the conformational changes of polymers. In recent decades, rare event algorithms have been developed to simulate these rare events from the dynamics but at a much lower computational cost than required by direct simulations. The general principle is to bias the ensemble statistics towards the event of interest without biasing individual trajectories. For an out-of-equilibrium system such as the climate system, this must be done in the trajectory space[‡]. We use the GKLT algorithm [Ragone et al., 2018, Ragone and Bouchet, 2020, Ragone and Bouchet, 2021, Wouters et al., 2023]. It is a genealogical

[‡] For an equilibrium system, one knows the probability measure over the state of the system. It is the Gibbs measure: $\mathbb{P}(x) = \exp(\beta H(x))/Z$ where x denotes the state of the system, H is the Hamiltonian of the system (so $H(x)$ is the energy of the system in state x) and Z is a normalisation constant known as the partition function. Since one knows the probability measure, one can directly bias the probability measure to sample rare events. Out-of-equilibrium systems, on the other hand, don't have a well-defined energy and we don't know the probability measure. For the climate, the only thing we are able to do is to simulate the dynamics. So we have no choice but to work in the trajectory space.

selection algorithm, adapted from [Del Moral and Garnier, 2005] by [Giardina et al., 2011] to sample rare extrema of large-time averages of a target observable $A(t)$.

Consider first a classical ensemble simulation. Most trajectories will produce typical events and only a handful will yield extreme events. For instance, if we are interested in a centennial event, only one out of hundred trajectories will sample the event of interest and 99% of the trajectories will be wasted data and computation time. To save computation time, it is more efficient to discard, early in the simulation, trajectories that will not yield the event of interest and to focus instead on the ones expected to yield the desired event. This is precisely what a genealogical selection algorithm aims at doing.

Figure 2 illustrates the behaviour of the algorithm. $N = 200$ trajectories are simulated in parallel. Every 5 days, some trajectories are discarded while others are duplicated in such a way that the number of trajectories remains constant. The duration $\tau = 5$ days is called the *resampling time* of the algorithm. Most trajectories are discarded at some point before the end of the algorithm while a handful of initial members are duplicated at every step and survive until the end of the algorithm. A *biasing parameter* k controls the strength of the selection. The larger k , the more stringent is the selection of trajectories that are duplicated, leading to a larger probability for trajectories with very large time-average of $A(t)$ (see Appendix A for details).

Note that we know the probabilities with which we discard or duplicate trajectories at each resampling step. This knowledge allows us to compute a probability for each surviving trajectory, which in turn enables us to compute various statistics such as return times or any composite statistics. This is where the added value of rare event algorithms lies, in comparison to other approaches.

Our target observable is the surface air temperature averaged over the box $70^\circ\text{-}76^\circ\text{E}$, $25\text{-}34^\circ\text{N}$, (see figure 1). We performed $M = 6$ rare event experiments with $N = 200$ trajectories for a total time $T_a = 90$ days from April 1st to June 30th§. The resampling time is $\tau = 5$ days which correspond to the synoptic time scale. The biasing parameter is $k = 0.1 \text{ K}^{-1}.\text{day}^{-1}$. For each experiment $m = 1..M$, the initial condition of the n^{th} trajectory is taken as the 1st of April of the year n of control run m . Since there is no oceanic variability in our simulations, the longest time scale is the one associated to soil moisture variability which is seasonal. Therefore we can consider that the 6 experiments start from 6 independent sets of 200 independent initial conditions. The computational cost of each experiment corresponds to the simulation of 200 AMJ seasons, equivalent to 50 years of simulation, leading to a total cost of 300 simulated years.

4. Results

[Ragone et al., 2018, Ragone and Bouchet, 2021] have assessed the reliability of the rare event algorithm for European extreme warm summers by comparing its results against a long control run of 1,000 years. Here we demonstrate the relevance of the algorithm

§ The number of days per month in the simulation is constant, equal to 30.

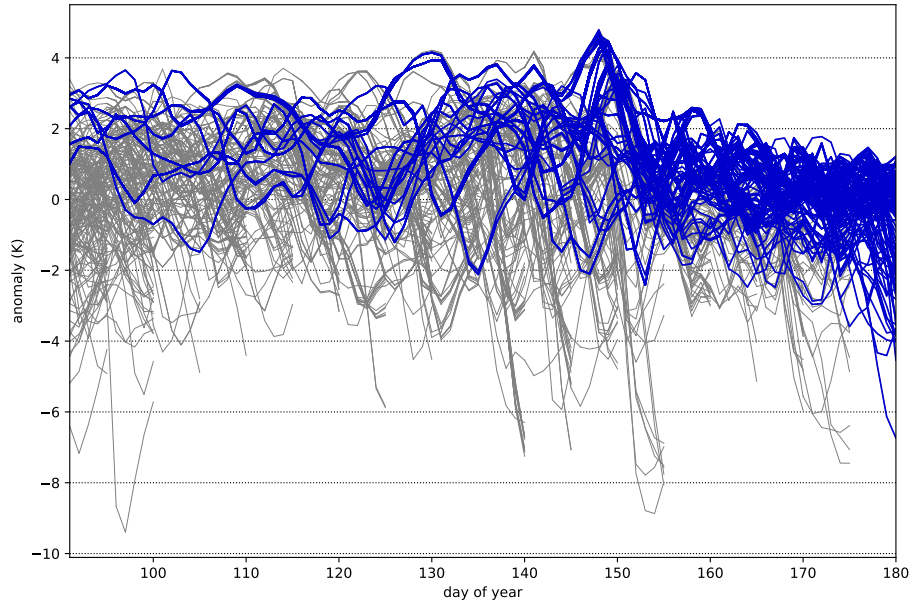


Figure 2. Time series of anomaly of daily averaged surface air temperature for one realisation of the rare event algorithm. $N = 200$ trajectories are initiated from N distinct years from the control run. Resampling is performed every 5 days: some trajectories are discarded while others are duplicated. The number of trajectories remains constant throughout the simulation. Trajectories that survive until the end of the algorithm are highlighted in blue. Grey lines correspond to trajectories that are eventually discarded.

for South Asian extreme heatwave seasons. We use a longer available control run of 8,000 years to extend the comparison to larger return times. We compare the return time curve estimated from the algorithm with an EVT fit and show that the algorithm estimate is more reliable with a much narrower confidence interval. The long control run also provides robust estimates of centennial heatwave season composite maps against which we can compare the composite maps derived from trajectories selected by the algorithm. Finally, we show that the algorithm also correctly estimates the intensity-duration-frequency statistics of subseasonal heatwaves.

We refer to appendix [Appendix A](#) or to the Supplementary Information of [[Ragone et al., 2018](#)] for the details of how to compute return time curves and composite statistics from the algorithm.

4.1. Importance sampling and return time curves

Figure 3a) shows the distributions of AMJ averaged temperature for the control run and the rare event algorithm (the 6 experiments have been pooled together). The algorithm effectively samples events above 1K anomaly and up to above 2K, a value never reached in the 8,000-year long control simulation. Most events lie in the far upper tail of the control run distribution ($\sigma_{\text{CTRL}} = 0.60$ K.). This represents a considerable gain in computation time for the simulation of these extreme events.

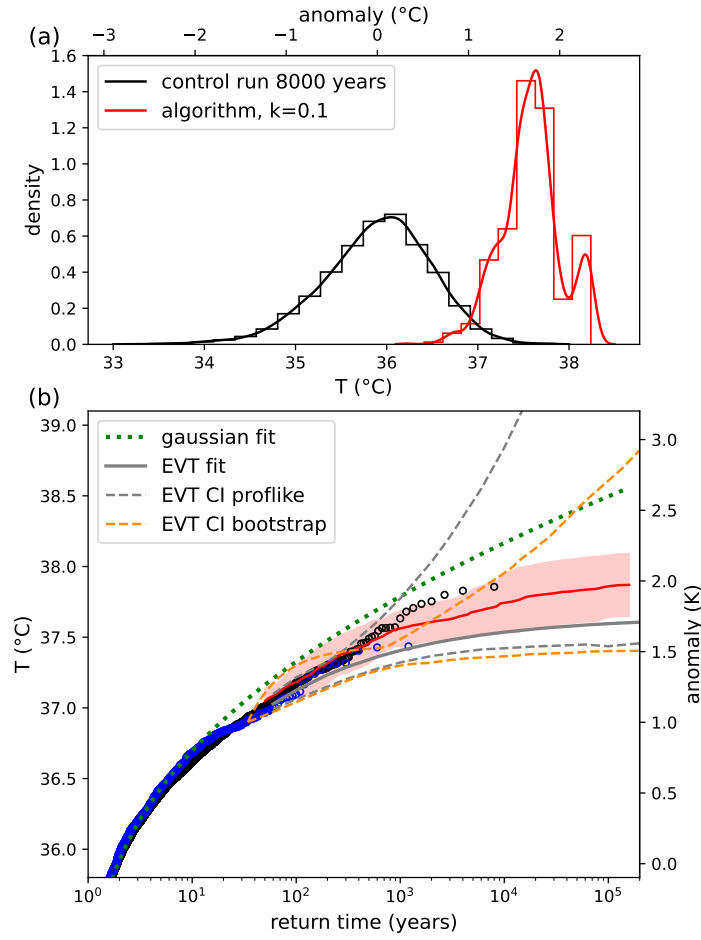


Figure 3. (a) Histograms and kernel density estimates of the AMJ averaged surface air temperature a in the 8,000-year control run (black) and the 6 rare event algorithm experiments with $k=0.1 \text{ K}^{-1} \cdot \text{day}^{-1}$ (red). (b) Return time curve of a built from the 8,000-year control run (black circles) and an independent 1,200-year control run (blue circles). The 1,200-year run was used to perform an EVT fit (grey) and served as initial conditions for the 6 algorithm experiments (red). Red shaded area corresponds to the 95% CI of the algorithm built using Student's t law. 95% CI for the EVT fit (dashed lines) is estimated via two different methods: (i) profile likelihood (grey) and (ii) a non-parametric bootstrap (orange).

On figure 3b), we compare different estimations of the return time curve of AMJ averaged temperature a . The black and blue circles correspond to the 8,000-year and an independent 1,200-year long control runs respectively. This 1,200-year run provided initial conditions for the algorithm whose return time curve (red line) is obtained by averaging together the 6 independent experiments. The spread of the experiments allows to build a 95% confidence interval (CI) on the return levels (red shaded area). Details about the estimation of the CI are provided in [Appendix B](#). The algorithm and control

run curves are in excellent agreement in the range 100-1,000 years. The two curves separate for return times larger than 1,000 years. However, the control run estimates for these extremely rare events are not very reliable since we are on the far edge of the empirical distribution.

How good is the rare event algorithm estimate compared to other approach? To answer this question, we added on the panel the return time curves of two statistical fits performed on the same 1,200-year run that provided initial conditions for the algorithm. The first one is a gaussian fit (green dotted line). While a gaussian fit may be a good approximation in the bulk of the distribution, it fails to capture the statistics of the extremes by overestimating their return levels. The second is an EVT fit (grey line) using the peak over threshold approach with a threshold of 36.9°C. The parameters are estimated by maximum likelihood estimation (MLE). The algorithm central estimate is much closer to the 8,000-year run return time curve than the EVT fit. The grey and orange dashed lines represent the bounds of the EVT fit 95% confidence interval computed with two different methods: profile likelihood (grey) and a non-parametric bootstrap (orange). The non-parametric bootstrap is widely used in the climate literature to compute confidence intervals on EVT fits (e.g. [Philip et al., 2020, Noyelle et al., 2023]). The profile-likelihood method consists in finding the range of parameters such that the likelihood exceeds a certain threshold [Coles, 2001] and seems more grounded theoretically. Details about the computation of the confidence interval and support for the choice of the threshold value are provided in [Appendix C](#). The bootstrap CI is too narrow for return times of order of millennia, failing to capture the empirical distribution of the 8,000-year run, before widening rapidly for larger return times. The profile-likelihood CI widens extremely fast for return time larger than a few centuries^{||} but has the merit of capturing the 8,000-year empirical return time curve. The algorithm confidence interval mostly captures the empirical distribution. Note that the empirical return time curve is not the ground truth: there may be fluctuations around the true return time curve, especially for the most extreme points. Moreover the algorithm CI is much narrower than the EVT ones for return times greater than 600 and 10^4 years, for the profile-likelihood and bootstrap methods respectively, and widens at a drastically lower pace. We conclude that the algorithm provides a much more precise estimate of the return time curve than the gaussian and EVT fits. Indeed, the central estimate is in better agreement with the empirical curve while the algorithm's CI is narrower than the one obtained from the EVT fit with the profile-likelihood. The fact that the non-parametric bootstrap confidence interval fails to capture the empirical distribution suggests that this method is not appropriate to estimate confidence intervals.

^{||} The upper bound of the CI corresponds to values of $\xi > 0$. As the shape parameter is almost always found to be negative in heatwave analyses [Van Oldenborgh et al., 2022], one could envision to constrain the result on $\xi \leq 0$. This possibility is not implemented in the package we used. However, an upper bound with $\xi = 0$ (exponential decay) would correspond to a straight line in the plot, still a rapidly growing CI.

4.2. Composite statistics of centennial heatwave seasons

As the algorithm produces heatwave seasons which are solutions of the model dynamics, we have access not only to the AMJ averaged temperature but to any observable at the model output frequency. This allows to compute a variety of composite statistics conditioned on extreme heatwave seasons. In this section, we compare composite statistics conditioned on centennial heatwave seasons computed from the rare event algorithm and from the control run. We define *centennial heatwave seasons* as the seasons for which the AMJ averaged temperature a is higher than a threshold a_{100} . We use a common threshold value which is the empirical 100-year return level of the 8,000 year long control run: $a_{100} = 37.18^\circ\text{C}$ corresponding to an anomaly of 1.28 K.

We study the composite statistics of the seasonal evolution of the temperature, of sub-seasonal heatwaves, and of the seasonal average of Zg500. We find that the statistics of trajectories generated by the algorithm are in very good agreement with the ones from the control run. In particular, the algorithm correctly reproduces the seasonal evolution of the temperature during 100-year heatwave seasons as well as the frequency-duration-intensity statistics of sub-seasonal heatwaves. Therefore the algorithm can be used, not only to compute the average temperature anomaly of a heatwave season, but also to anticipate what can be the intensity and duration of heatwaves within the season.

4.2.1. Seasonal evolution of the temperature We first compare the seasonal evolution of the temperature for the centennial heatwave seasons from the control run and for the ones from the algorithm in figure 4. Note that the months of April, May and June correspond respectively to day of year 91-120, 121-150 and 151-180. Panel (a) shows the daily-averaged absolute temperatures for the 80 centennial heatwave seasons from the control run. Panel (b) show the corresponding anomalies with respect to the climatology. The bulk of the trajectories concentrate between the climatology (solid black line) and the 99th percentile (upper dashed-dotted line). The temperature anomalies decrease towards the end of the season with more trajectories being below the climatology and almost no one being above the 99th percentile. Panels (c) and (d) show the same quantities for the centennial heatwave seasons produced by one realisation of the algorithm. Out of 200 trajectories, 161 correspond to centennial heatwave seasons. At the beginning of the season, the distribution of temperature is more restricted in the algorithm than in the control run because the number of independent trajectories is small (here all centennial heatwave seasons come from 4 independent initial conditions). The number of separated trajectories increases progressively during the simulation. All trajectories remain between the climatology and the 99th percentile during most of the time. The temperature anomaly also decline at the end of the season. However, note that visual comparison can be misleading to interpret the distribution of trajectories produced by the algorithm. Indeed, trajectories in the control run have equal probability while the ones produced by the algorithm have different probabilities which can differ by orders of magnitude but these differences are not visible on the figure. In panels (e-f)

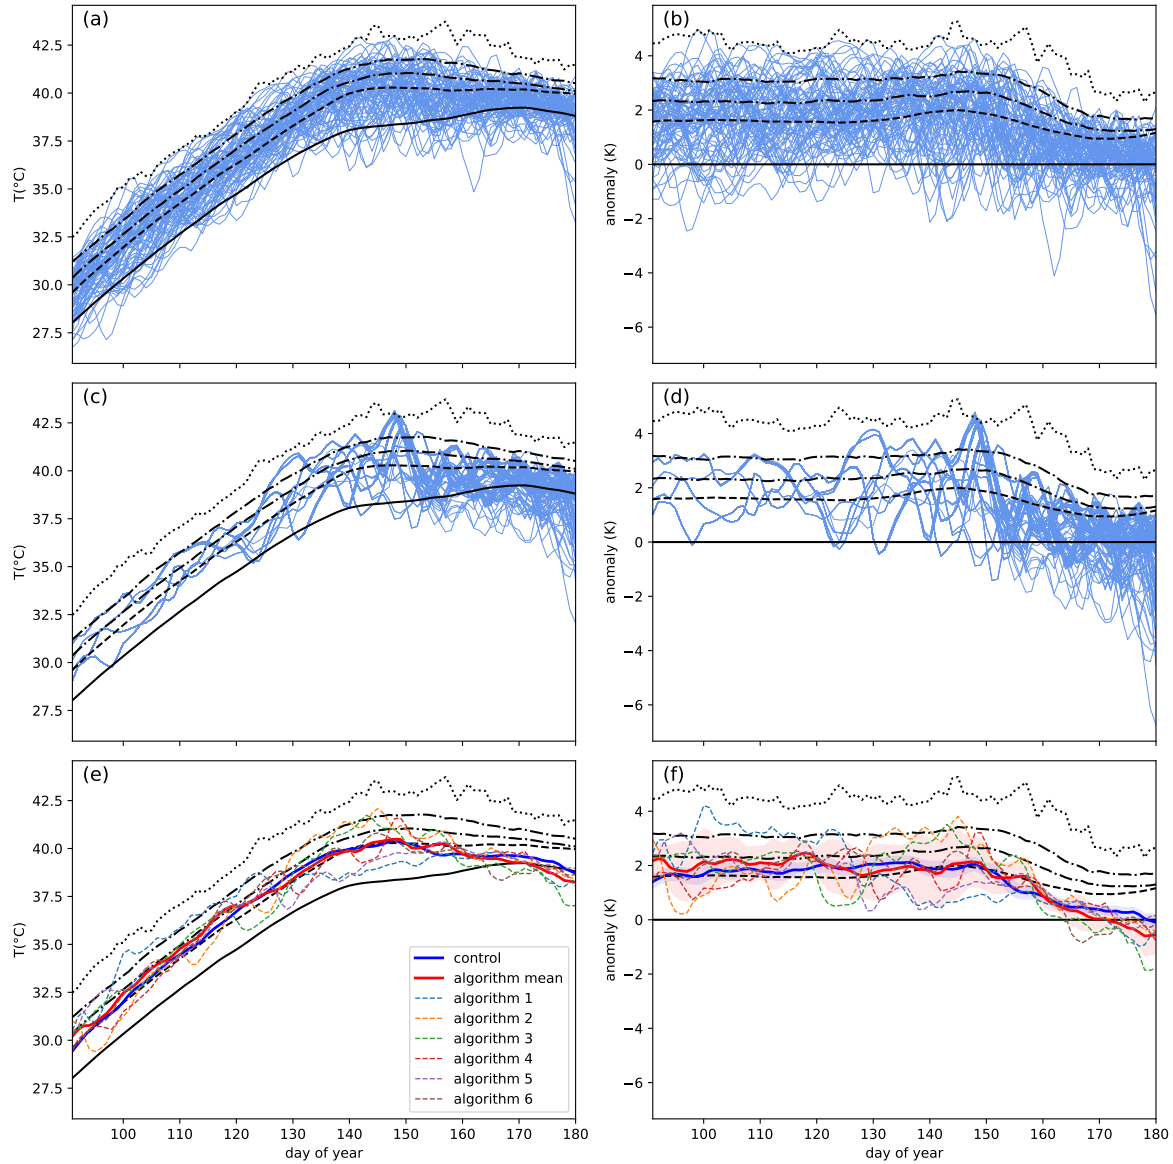


Figure 4. Daily-averaged temperature of centennial heatwave seasons in the control run (a-b) and one rare event algorithm experiment (c-d). The left column shows the absolute temperature while the right one shows the temperature anomaly with respect to the climatology. Black lines correspond to the climatology (solid), one standard deviation (dashed), the 95th and 99th percentiles (dash-dotted) and the maximum (dotted) of the long control run daily temperature distribution. Panels (e) and (f) show the composite mean of daily temperature conditioned on centennial heatwave seasons in the control run (blue) and the algorithm (red) as well as the mean in each individual experiment (dashed coloured lines). The blue and red shaded bands on panel (f) represent the 95% confidence interval for the control run and the algorithm respectively. Note that the months of April, May and June correspond respectively to day of year 91-120, 121-150 and 151-180.

we make a more quantitative comparison by plotting the composite mean temperature of centennial heatwave seasons, along with a 95% confidence interval in panel (f). The

composite means of the control run and the algorithm are close to one another and consistent within the uncertainty range from the algorithm.

We notice that, while the climatological mean of daily-averaged temperature keeps increasing along the season until the end of June, extreme heat (characterized by the 95th and 99th percentiles of daily temperature) peaks end of May and slowly decays afterwards. As a consequence, the variability of temperature is lower end of June. Regarding centennial heatwave seasons, the conditioned mean temperature is consistently at least one standard deviation higher than the climatology during the months of April and May (day of year 91 to 150). It peaks end of May and then decays, at a faster rate than extreme quantiles, to reach a value close to the climatology by the end of the season. This means that the conditions leading to anomalously high temperatures in April-May are different from the ones leading to high temperature anomalies in June

4.2.2. Sub-seasonal heatwave statistics During a heatwave season, periods of hot or extremely hot temperatures alternate with relatively cooler periods. Long heatwaves are the most detrimental to human health. Therefore, knowing that we have an extreme heatwave season, it is also important to correctly represent the sub-seasonal statistics of heatwaves. Hereafter, we call heatwave a period during which the daily-averaged temperature remains above a (time-dependent) threshold. We do not impose any restriction on the duration of the heatwave. We use two different thresholds: $A_{\text{th}}(t) = A_{\text{mean}}(t) + \sigma_A(t)$ and $A_{\text{th}}(t) = A_{q95}(t)$ where $A_{\text{mean}}(t)$, $\sigma_A(t)$ and $A_{q95}(t)$ are the climatological mean, standard deviation and 95th percentile of the long control run daily temperature distribution. We look at events where the temperature exceeds these thresholds. For the sake of convenience, we call these events *above 1σ events* and *above q^{95} events* respectively. When we don't need to be specific, we simply call them *heatwaves*.

For $1 \leq J \leq T_a = 90$ we define the observable $\#\text{HW}_J(\{x(t)\}_{0 \leq t \leq T_a})$ which counts the number of heatwaves lasting exactly J consecutive days along the season $\{x(t)\}_{0 \leq t \leq T_a}$. We first study the average number of heatwave lasting exactly J days during centennial heatwave seasons:

$$f_J := \mathbb{E} [\#\text{HW}_J(\{x(t)\}_{0 \leq t \leq T_a}) | a(\{x(t)\}_{0 \leq t \leq T_a}) > a_{100}]. \quad (2)$$

We compute the empirical estimation of f_J from the control run and from the algorithm using the following formula:

$$\hat{f}_{J,\text{ctrl}} = \frac{1}{K} \sum_{k=1}^K \#\text{HW}_J(\{x_k(t)\}_{0 \leq t \leq T_a}) \quad (3)$$

$$\hat{f}_{J,\text{algo}} = \sum_{n|a_n > a_{100}} w_n \#\text{HW}_J(\{x_n(t)\}_{0 \leq t \leq T_a}) \quad \text{with } w_n = \frac{p_n}{\sum_{n|a_n > a_{100}} p_n}, \quad (4)$$

$$p_n = \frac{\exp\left(k \int_0^{T_a} A(t) dt\right)}{N \tilde{Z}} = \frac{\exp(k a T_a)}{N \tilde{Z}}. \quad (5)$$

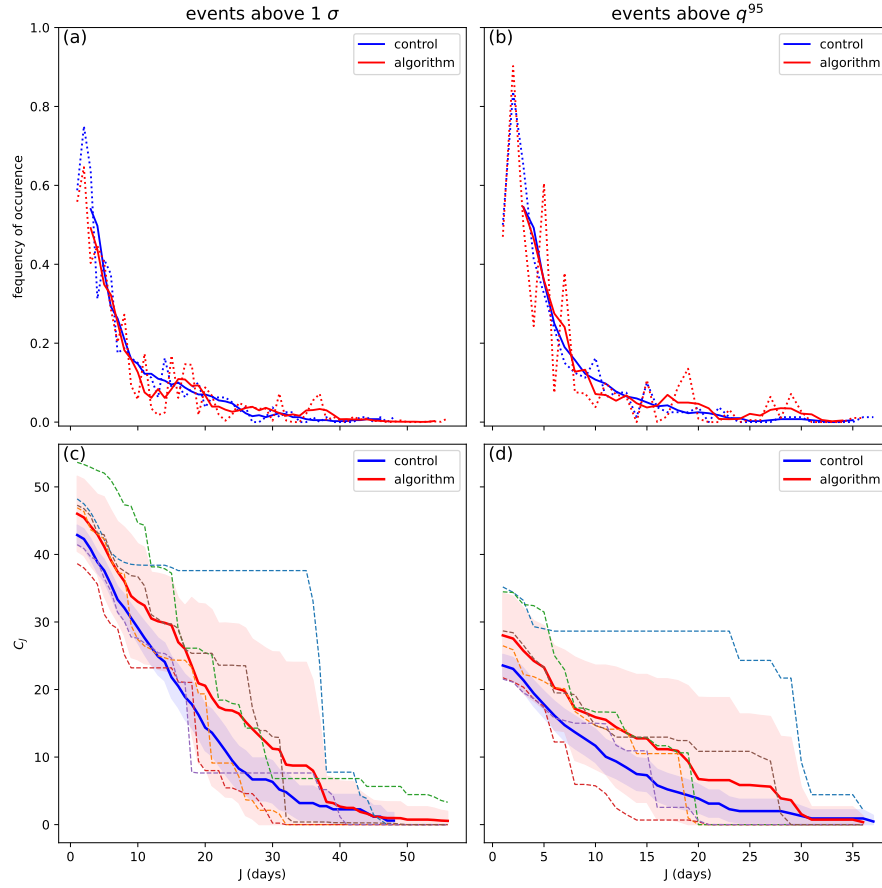


Figure 5. *Top:* Frequency of heatwave occurrence as a function of the heatwave duration J . The dotted lines represent the direct estimate computed with equations 3. The solid lines represent a smoothed version, using a 5-day centred mean. *Bottom:* Cumulative number of days C_J belonging to heatwaves lasting J days or more. Dashed lines correspond to the individual realisation of the algorithm. The blue and red shaded area correspond to the 95% confidence interval from the control run and the algorithm respectively.

In the first line, the sum runs over the $K=80$ years whose average temperature is above the 100 years return level a_{100} . The p_n 's are the probabilities attached to the algorithm trajectories and \tilde{Z} is a normalization factor that is different for each realisation of the algorithm. We refer to [Appendix A](#) for the details of the probabilities computation.

The top row of figure 5 shows the empirical frequencies $\hat{f}_{J,ctrl}$ and $\hat{f}_{J,algo}$ for the two heatwave thresholds. The algorithm and the control run are in excellent agreement: both the number of heatwave days and the decreasing trend are correctly captured. The second row of figure 5 shows the total number of days that belong to heatwaves that last J days or more:

$$C_J = \sum_{j \geq J} j \times f_j. \quad (6)$$

In particular, C_1 is the conditional expectation of the total number of heatwave days

during centennial heatwave seasons. The empirical estimate from the algorithm is 4 to 5 days higher than the control run for $J < 37$ (above 1σ events) or $J < 30$ (above q^{95} events). This can be traced to a larger occurrence of heatwaves of duration $J = 32, 37$ (above 1σ events) and $J = 27, 29$ (above q^{95} events). The control run curve is nevertheless within the uncertainty range of the algorithm, with 2 out of 6 algorithm realisations yielding estimates of C_1 that are lower than the control run value.

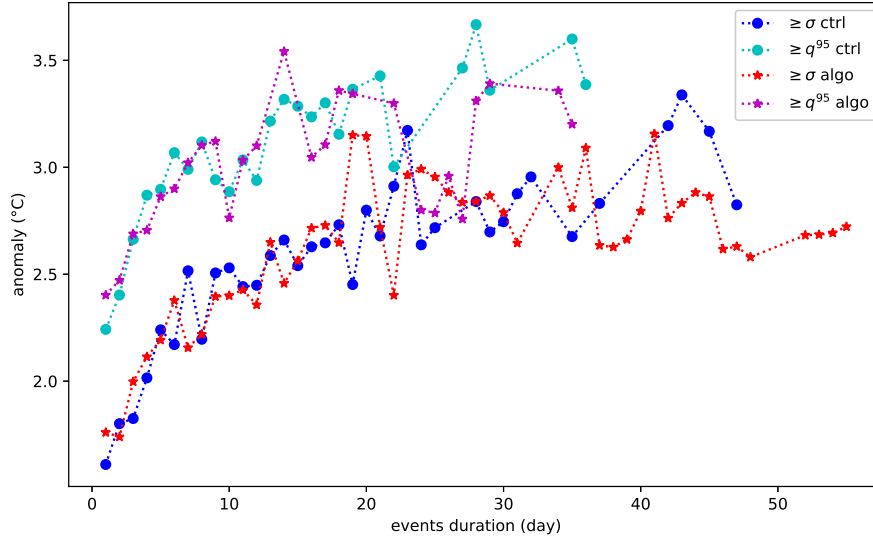


Figure 6. Heatwaves intensity versus heatwaves duration. Mean temperature anomaly of heatwaves versus their duration.

Finally, we assess the ability of the algorithm to correctly represent the intensity-duration relationship. Figure 6 shows the mean temperature anomaly of heatwaves as a function of their duration. Once again the algorithm is in excellent agreement with the control run despite the noisy curves (due to the low number of heatwaves for each duration). In particular we observe a first regime of rapid increase of the intensity with the events duration up to $J = 8$ days before a second regime of slower increase.

4.2.3. Composite maps and hemispheric quasi-stationary wave-pattern Rare event simulations can also be used to study dynamical quantities associated to the extreme events of interest. Here, we study the atmospheric conditions associated to centennial heatwave seasons. In figure 7 we compare Zg500 composite maps from the 8,000-year long control run (a) and the algorithm (b). The long control run composite map is an average over $M=80$ independent years, it is therefore an exceptionally good estimate of the ground truth for such a rare event, thanks to the unusually long dataset. The algorithm composite map is similar to the long control run one. In particular, the extreme heatwave seasons are associated to an anticyclonic anomaly which is centred on the northwest of the target region. This anticyclonic anomaly is part of a large-scale hemispheric quasi-stationary wave-pattern which features anticyclonic anomalies over

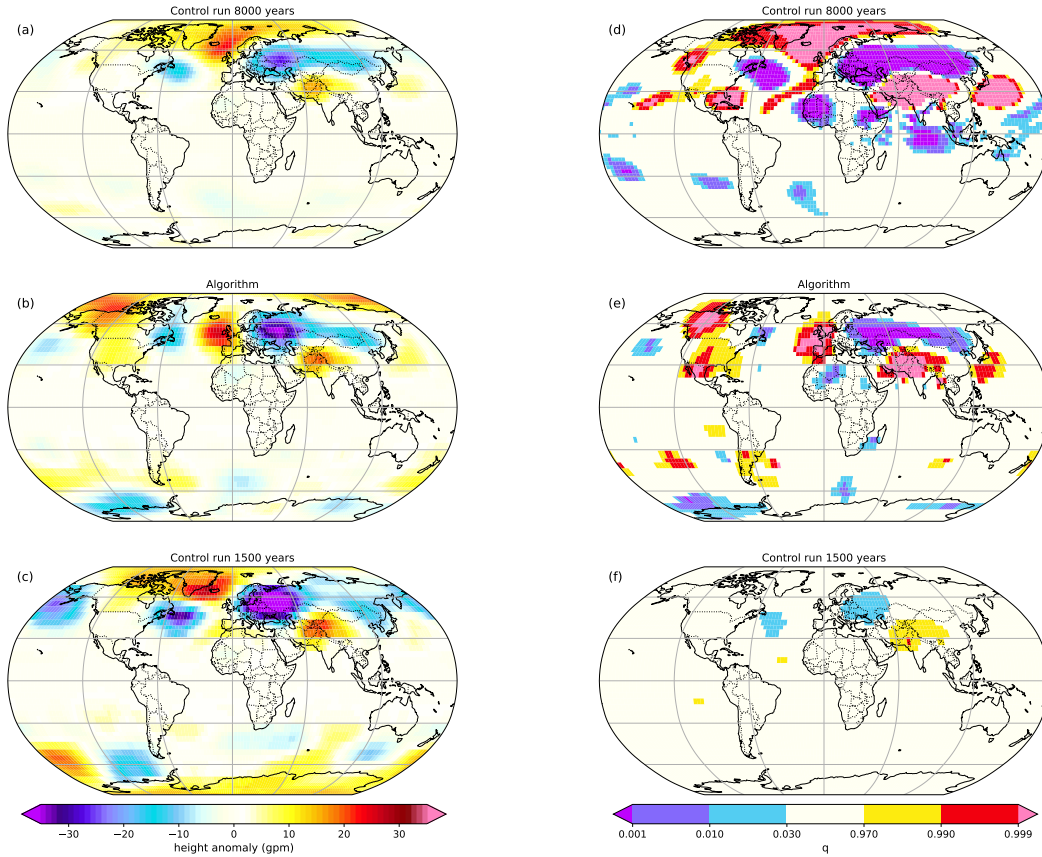


Figure 7. *Left:* Zg500 composite maps conditioned on $a > a_{100}$ for the 8,000 year long control run (a), the rare event algorithm (b) and an independent control run of 1,500 years (c). The maps (a) and (c) are averaged of $M=80$ and $M=10$ maps respectively while (b) is the mean of the 6 individual experiments. *Right:* Corresponding significance maps. q is a measure of the significance computed with a Student's t-test: the anomaly is positive (resp. negative) with probability q (resp. $1 - q$). See Appendix B for details.

the North-East Atlantic, western Canada and the western Pacific ocean (East of Japan) and cyclonic anomalies at the tip of Newfoundland and over Eastern Europe. The latter anomaly stretches to the East towards the Pacific ocean. However, the algorithm composite map presents larger anomalies in the Southern hemisphere. In the right column of figure 7, we plot the significance levels associated to the anomalies. Both the long control run and the algorithm have numerous significance regions (in particular all the anomalies mentioned above are significant). But the significance area for the algorithm are indeed less extended than for the 8,000-year long control run. However, the total computation cost for running the algorithm is only 1,500 years (a 1,200-year long control run to produce the 1,200 initial conditions and 300 years of computation cost for running the algorithm itself), i.e. more than 5 times lower. What would be the

quality of a composite map computed with a control run of similar cost? To answer this question, we add on panel (c) a composite map computed from a second independent control run of 1,500 years[†]. The general pattern mentioned above is still distinguishable but the map bears less resemblance with the long control run than the algorithm does (although the initial conditions for the algorithm come from this second control run.). There are also larger anomalies in the Southern ocean hinting at less significance. This is confirmed by the corresponding significance map (f). The 1,500-year composite map only has three regions of statistically significant anomaly. Therefore, it would not be possible, based on this sole map, to confidently assess the existence of the hemispheric wave-pattern mentioned above.

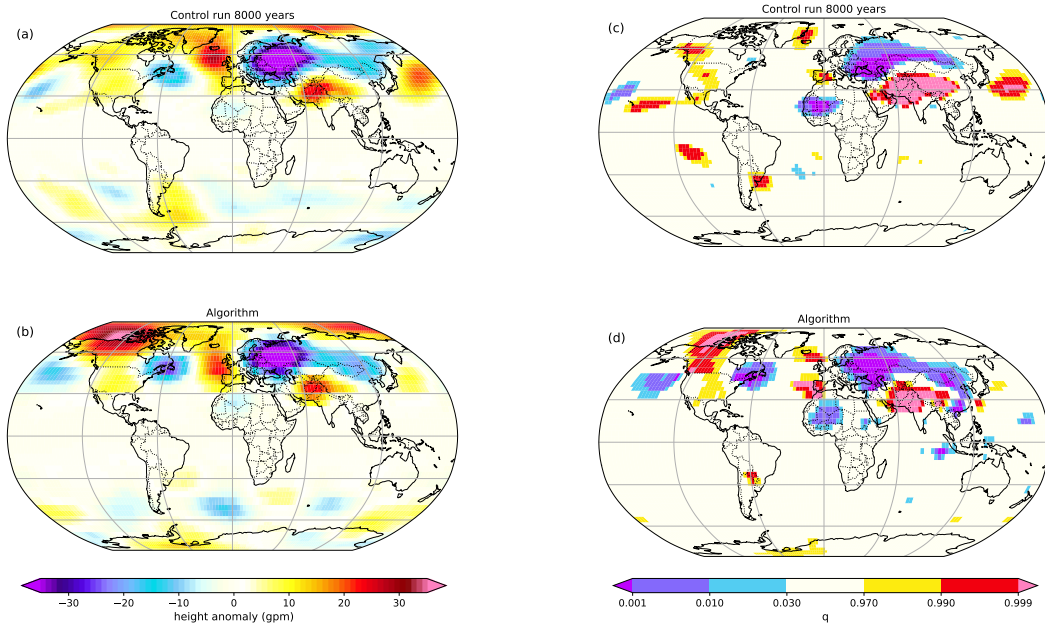


Figure 8. *Left:* Zg500 composite maps conditioned on $a > a_{1,000}$ for the 8,000-year long control run (average on $M=8$ maps) (a) and the rare event algorithm (b). *Right:* Corresponding significance maps. The interpretation is the same as in fig. 7.

The algorithm is not only more efficient than a control run of equivalent cost for studying 100-year events. It can be used to study even rarer events that are out of reach of a control run at equivalent cost. On figure 8 we show Zg500 composite maps for 1,000-year heatwave seasons. The same threshold $a_{1,000}$, computed from the 8,000-year long control run, is used for the control run and the algorithm[‡]. We distinguish the same quasi-stationary wave-pattern as we observed for the centennial heatwave seasons but with stronger cyclonic and anticyclonic anomalies. The pattern is confirmed by the

[†] In this second run there are only $M=10$ years in which the AMJ-averaged temperature exceeds the 100-year return level.

[‡] There is no event exceeding the threshold in the 1,500-year long run.

significance maps (left column of fig. 8). Note that the significance of the algorithm is almost the same as for 100-year heatwave. In contrast, the significance of the long control run has largely decreased.

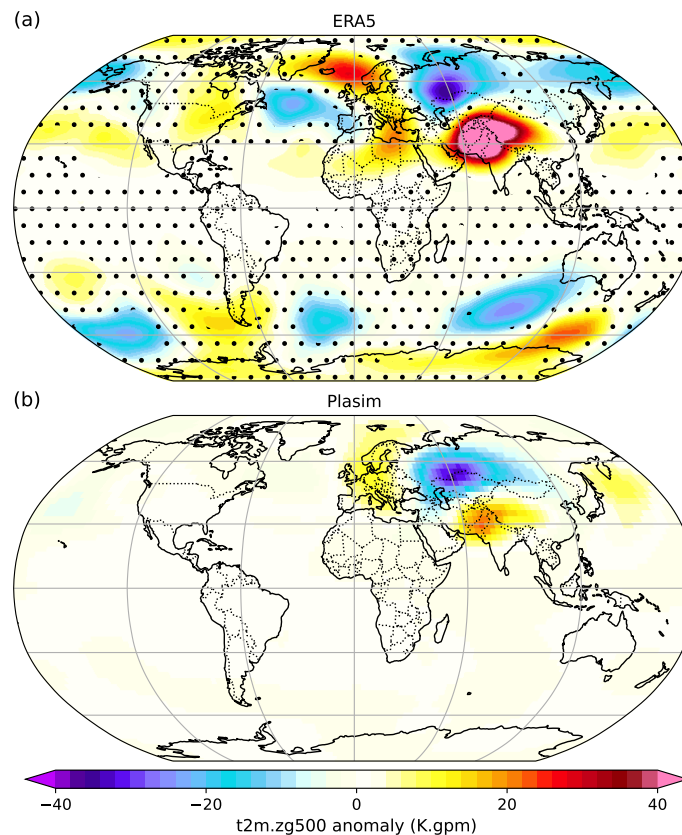


Figure 9. 3-day correlation maps between Zg500 and the temperature over the target region in ERA5 (a) and Plasim (b). In (a), we use a bootstrap test to assess significance at the 95% confidence level[§]. Stippling denote regions where anomalies are non-significant.

4.2.4. Fairness of the hemispheric wave-pattern The composite maps in figures 7 and 8 show that extreme heatwave seasons in Northwest South Asia are associated to an anticyclonic anomaly centred on the northwest flank of the target region embedded in a hemispheric quasi-stationary wave-train. Can we trust this pattern for the real world? Of course, we cannot use historical data to build composite maps of such extreme heatwave seasons. ERA5 data-set goes back to 1940 so event composite on 10-year heatwave seasons would have a poor significance. To go around this problem we assess the general ability of the model Plasim to correctly represent the teleconnections between the temperature over the target region and the 500hPa geopotential height anomaly. To

[§] The computation time to build a single correlation map in Plasim being 4 hours, it was not feasible to apply a bootstrap test in a reasonable computation time.

this end we compare maps of correlation between the temperature over the target region and Zg500 computed with Plasim and ERA5 reanalysis dataset. Before computing the correlation, we removed cubic trends from the ERA5 time series (see section 2.3 and fig. D1 for details) and coarse grain the time series by performing 3-day block average. Figure 9 show the resulting correlation maps for ERA5 (a) and Plasim^{||} (b). Plasim’s correlation map show a strong positive correlation region centred over Pakistan and two other ones over Central Europe and the Pacific West. A wide negative correlation region centred at the border between Kazakhstan and Russia separates the positive correlations over Europe and South Asia. ERA5 correlation map is much noisier due to a much lower number of years (83 versus 1,000). However, the features present in Plasim’s composite maps and mentioned above are all presents in ERA5 correlation map, including the negative anomaly at the tip of Newfoundland. This suggests that Plasim represents fairly well the teleconnection between temperature over the target region and Zg500 and that the hemispheric wave-train observed on the composite maps is something real.

4.3. Methodological study of the algorithm

Finally we address two methodological questions regarding the efficiency of the algorithm and the optimal way to run it. The first question pertains to the duration of extremes that the algorithm can effectively sample. The algorithm samples extreme hot events by selecting trajectories with persistent large temperatures. However, it does not produce new record of the daily temperature, as can be seen on the middle line of figure 4. So, what is the minimal event duration for which the algorithm effectively samples extremes? To answer this question, we show on figure 10 the return time curves for April-May (60 days) and April (30 days) averaged temperature, corresponding to running the algorithm for 11 and 5 resampling steps respectively[¶]. For 60-day events, panel (a) shows that the algorithm effectively samples extremes up to 1K hotter than in the control run used for initialisation. Moreover, the algorithm curve is in good agreement with the long control run. For 30-day events, panel (b) shows that the algorithm does not sample more extreme events than what is already present in the initialisation control run. Many algorithm trajectories stem from the same initial condition which corresponds to an outlier April temperature of the initialisation run. The remaining trajectories do not exceed the second largest April temperature of the initialisation run. Finally, the algorithm curve falls off the control run curve at low return time, with trajectory probabilities being overestimated.

While this result suggests that the algorithm does not improve the sampling of extreme 30-day events, we must mention that it is actually the case if one follows a slightly different approach than ours: when considering the largest 30-day averaged

^{||} We use 1,000 years of data to compute Plasim’s correlation map.

[¶] The number of resampling steps is the event duration divided by the resampling time τ (=5 days here) minus 1 because no resampling is performed after the last integration window.

anomaly during a four month long period, the algorithm does greatly improve the sampling of extreme 30-day events (see chapter 5 of the PhD thesis [Cozian, 2023]).

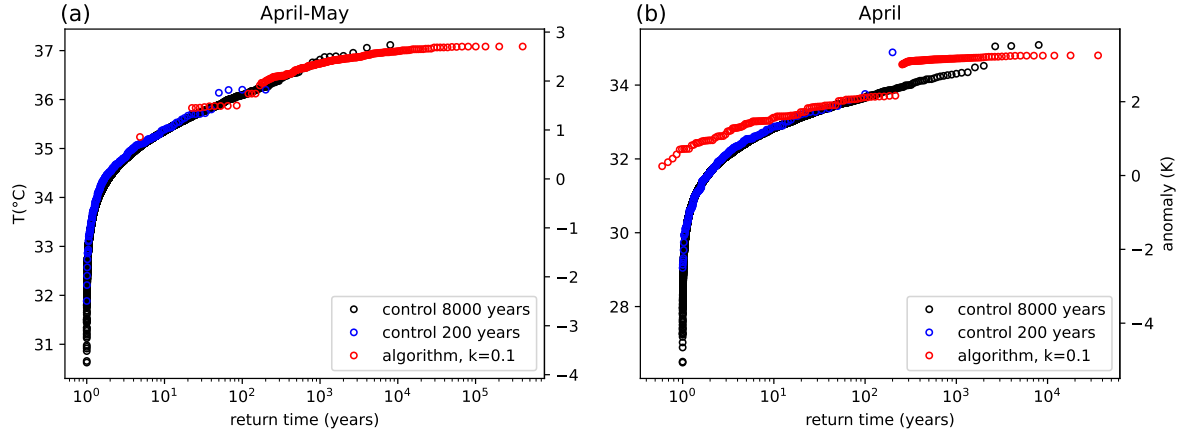


Figure 10. (a) Return time curves of April-May averaged temperature for the long control run (black), one algorithm realisation with $N = 200$ trajectories (red) and the 200-year long control run used to initialise the algorithm (blue). (b) Same as (a) for April averaged-temperature.

A possible alternative way to sample shortest events could be to decrease the resampling time τ while keeping the number of resampling steps constant. Although we did not try out this approach, we do not expect it to work. Indeed, an analysis of the trajectories reveals that clone trajectories barely separate after 5 days. If τ were smaller, clone trajectories would not separate during the integration window and the subsequent resampling would not discriminate among trajectories.

The second question is whether the algorithm is more effective when started a few resampling times ahead of the period studied. The underlying idea is that the algorithm could need a few resampling steps before reaching a statistical equilibrium among trajectories. On panel (a) of figure 11 we compare the set of algorithm experiments starting at the beginning of the period studied (April 1st for AMJ averaged temperature) with two other sets of experiments starting one and two resampling times earlier respectively (i.e. 5 and 10 days earlier). Each set contains $M = 6$ algorithm experiments with $N = 200$ trajectories. The return time curve of the set starting 5 days earlier is barely distinguishable from the standard experiment set. The uncertainty is smaller or larger than for the standard experiment depending on the return time. Thus, no gain is achieved globally on the precision of the estimated return levels. The 10-day earlier start curve is slightly under the two others but remains within the confidence intervals. Its inter-experiment standard deviation is consistently larger, thus yielding a larger uncertainty on the estimates. How to discriminate between the different curves? Is the lower curve of the 10-day earlier start a better estimate or does it suffer from a systematic bias? Panel (b) of figure 11 shows the effect of a start of the algorithm 30 days ahead of the period of interest. We show the return time curves for MJJ averaged temperature with the algorithm being started on April 1st. The algorithm falls

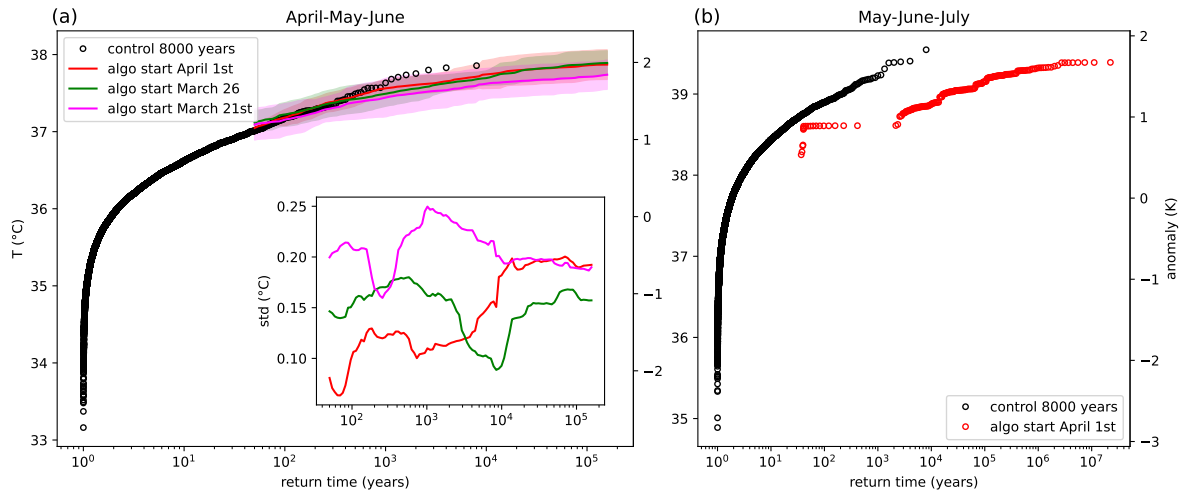


Figure 11. (a) Return time curves of April-May-June averaged temperature for the long control run (black) and three sets of algorithm experiments with starting dates on April 1st (red), March 26 (green) and March 21st (magenta). $M=6$ experiments were run for each starting date. Solid lines and shaded areas represent the inter-experiments mean and standard deviation respectively. **Inset:** Inter-experiments standard deviation for the three different starting dates. (b) Return time curves of May-June-July averaged temperature for the long control run (black) and one algorithm realisation that was started on April 1st.

completely off the control run curve by largely underestimating the events probabilities. The reason is that large temperatures in April give high weights, thus low probabilities, to the trajectories but do not enter the considered temperature average. It is probable that the 10-day earlier start suffers from the same bias, although in a smaller way. We conclude that an earlier start does not improve the efficiency of the algorithm and can actually be detrimental to the quality of the results by inducing a systematic bias where probabilities are underestimated.

5. Relevance for extreme event attribution

We showed that a rare event algorithm can provide better estimates of return levels than an EVT fit, with confidence intervals that are much narrower for large return times. Therefore, such algorithms could be used in attribution studies to better constrain the change in return time - or return level - with global warming. It would require running the algorithm at least two times, one for the factual world and one for the counter-factual world. Additional runs could be required if one wants to explore the change in a warmer (e.g. $+2^{\circ}\text{C}$) world. However, running the algorithm requires much more work (and computation hours) than performing a statistical fit on already existing data. Therefore one needs to consider if it is worth the effort. We give here a few considerations on this question.

We showed that the gain in precision on return levels increases drastically with the

rarity of the events. Therefore rare event algorithms would be particularly suited for the rarest events. It demands probably too much work for rapid attribution studies but could be used to refine the estimates in subsequent peer-reviewed studies. Robust attributions are performed when many models are combined with observations to obtain a *hazard synthesis* [Philip et al., 2020]. When there are many valid models, one might not want to run the algorithm for each model but it could be used to refine the estimate in the model with the highest uncertainties. In the alternative when there are only few valid models, one could envision to run it for all models.

We want to emphasize that an important added-value of rare event algorithms is that they actually sample the extreme events of interest, thereby allowing to compute composite statistics with good significance. Therefore they could be used for process-based studies and investigate how the process associated to very extreme events change with global warming.

6. Conclusion

In this study, we demonstrate how a rare event algorithm samples extremely rare events in a climate model and allows the computation of various related statistics. We compared the performance of the algorithm against a 8,000-year long control run which enabled us to demonstrate the accuracy of the statistics computed via the algorithm. First, we showed that the algorithm provides a good estimate of the return time curve with narrow confidence intervals. Secondly, since it actually samples extreme events, the algorithm can be used to compute statistics conditioned on the extreme event of interest. For instance, composite maps conditioned on 100-year heatwave seasons were found to be in good agreement with the long control run estimates. These maps suggest that extreme heatwave seasons in South Asia are associated with an anticyclonic anomaly which is embedded in a large-scale hemispheric quasi-stationary wave-pattern. Comparison with reanalysis give confidence in this result, since the model represents well teleconnections between the 2-meter temperature over the target region and Zg500 over the Northern hemisphere.

The algorithm also successfully estimates the intensity-duration-frequency statistics of individual subseasonal heatwaves. Thus, the algorithm can also be used to estimate the distribution of intensity and duration of heatwaves that comprise a heatwave season. Thus rare event algorithms could, for instance, be combined with seasonal forecasts to provide information regarding expected number of heatwave days and the distribution of the duration and intensity in an extreme heatwave season. This combination of seasonal mean and subseasonal variability could be useful for adaptation planning or risk management. In addition, the algorithm can be used to generate more samples of individual heatwaves and heatwave seasons. Such samples could be used in analysis of circulation features relevant to the predictability of seasonal weather such as atmospheric waveguides [White et al., 2022], for process-oriented studies and data-driven seasonal forecasting methods.

Most current literature on use of rare event algorithms to climate models, including our work, have analysed heatwaves as measured by air temperature, and in dynamical regimes where Rossby waves are the primary drivers of such heatwaves. Application to humid heatwaves (and other extreme events of interest such as droughts or precipitations) and to tropical regions remains an avenue for future research. Regardless, rare event algorithms are a promising methodological advance which allows for a systematic study of the statistics and dynamics of extreme events in climate models.

Acknowledgments

This project was provided with computing and storage resources by GENCI at TGCC on the partition ROME of supercomputer Joliot Curie thanks to the grant 20XX-A010110575. Data analysis was carried out using the resources of the Centre Blaise Pascal at ENS de Lyon. We are grateful to Emmanuel Quemener for his help with the platform. This work was supported by the ANR grant SAMPRACE, Project No. ANR-20-CE01-0008-01 (C.L.P). C.L.P. thanks F. Ragone for sharing its code of the algorithm, B. Cozian for help with running Plasim and D. Faranda for advices on EVT fits.

Appendix A. Detailed description of the algorithm

Let $x(t)$ be a vector that represent the states of the model at time t ($x(t)$ is the vector of values of all the variable used to integrate the model.). For any time-dependent observable A , we may make implicit the dependence on the system state by writing $A(t)$ instead of $A(x(t))$. We simulate N trajectories in parallel. Trajectories are integrated in pieces of duration τ for a total time $T_a = I\tau$. τ is called the *resampling time* of the algorithm. At each resampling time step $t_i = i\tau$, we assign to the n^{th} trajectory a weight

$$W_n^{(i)} = \frac{1}{Z_i} \exp \left(k \int_{t_{i-1}}^{t_i} A(x_n(t)) dt \right) \quad (\text{A.1})$$

where Z_i is a normalization factor such that $\sum_{n=1} W_n^{(i)} = N$. Each trajectory is then assigned a random number of clones $m_n^{(i)}$ such that $\mathbb{E}[m_n^{(i)}] = W_n^{(i)}$ and $\sum_{n=1} m_n^{(i)} = N$. Trajectories with small weight ($W_n^{(i)} \ll 1$) are likely to be discarded ($m_n^{(i)} = 0$) while trajectories with large weights are duplicated. The *biasing parameter* k controls the strength of the selection. Since the dynamics of the climate model is deterministic, we add a small perturbation after the resampling in order to allow clone trajectories to separate in the next integration step. The perturbation is introduced by adding to the coefficients of the spherical harmonics of the logarithm of the surface pressure a set of random numbers, sampled independently according to a uniform distribution in $[-\epsilon\sqrt{2}, \epsilon\sqrt{2}]$, with $\epsilon = 10^{-4}$ [Ragone et al., 2018]. For the separation to be effective, τ must be of the order of the autocorrelation time of the system. The algorithm samples extrema of $\frac{1}{T_a} \int_{t_0}^{t_0+T_a} A(t)dt$ by favouring the selection of trajectories with persistent large values of $A(t)$. The details of the algorithm, including the precise way to draw the number of clones, are given in algorithm 1.

Estimation of the biasing parameter k When running the algorithm, one must make a choice for the value of the biasing parameter k . This choice must be adapted to the goal of the study, more specifically to the target value of temperature anomaly a . To avoid a trial and error approach, we need to estimate a priori the order of magnitude k . We give here a heuristic derivation of this estimate by dimensional analysis. A more rigorous derivation, using large-deviation theory can be found in [Ragone and Bouchet, 2020]. First, note from (A.1) that $[k] = [A]^{-1}[t]^{-1}$ where $[..]$ denotes the dimension of variables. Note also that k changes sign under the transformation $a \rightarrow -a$. Clearly, $|k|$ must increase with $|a|$ and k changes sign under the transformation $a \rightarrow -a$. Moreover, it should decrease with the spread of the distribution $\frac{1}{T_a} \int_{t_0}^{t_0+T_a} A(t)dt$ (as it is easier to reach a given level if the variance is larger). Let Σ be the standard deviation of the latter distribution. Introducing T_a as a characteristic time of the algorithm, we can write $k = a^\alpha \Sigma^\beta T_a^\gamma$. The dimensional analysis imposes that $\alpha + \beta = -1$ and $\gamma = -1$. For the first relation, any solution with $\alpha > 0$ and $\beta < 0$ would meet the requirement. However, the anti-symmetry w.r.t. $a \rightarrow -a$ implies that α is an odd integer. We stick

to the simplest solution $\alpha = 1$, $\beta = -2$. Our estimate of the optimal k is thus:

$$k^* = \frac{a}{T_a \Sigma^2}. \quad (\text{A.2})$$

Note that the greater k , the more stringent the selection is. If k is too large compared to N , it can be the case that only a single initial trajectory survives after a few selection steps. We are thus loosing the diversity of the initial condition and the result can be largely biased. To avoid this issue, N must be large enough so that diversity of trajectories is preserved.

Algorithm 1 GKLT algorithm

Initiate N trajectories $x_n(t=0) = x_n^0$ for $1 \leq n \leq N$.

for $i = 1$ to I **do**

(0) Add a small random perturbation to each trajectory

(i) Integrate the trajectories between $t_{i-1} = (i-1)\tau$ and $t_i = i\tau$.

(ii) Assign to each trajectory a weight $\tilde{W}_n^{(i)} = \exp\left(k \int_{t_{i-1}}^{t_i} A(x_n(t)) dt\right)$

(iii) Normalize the weights: $W_n^{(i)} = \tilde{W}_n^{(i)} / Z_i$ where $Z_i = \frac{1}{N} \sum_{n=1}^N \tilde{W}_n^{(i)}$

(iv) For each trajectory draw a number of clones $m_n^{(i)} = \lfloor W_n^{(i)} + u \rfloor$ where $u \sim \mathcal{U}([0, 1])$.

(v) Duplicate or eliminate clones to keep a constant number of trajectories: Compute $\Delta N_i = \sum_{n=1}^N m_n^{(i)} - N$; if $\Delta N_i > 0$, select ΔN_i clones without repetition and delete them; if $\Delta N_i < 0$, select $-\Delta N_i$ clones with repetition and duplicate them.

end for

Computation of the probabilities Trajectories that survive until the end of the algorithm are reconstructed by attaching to trajectories at step $i+1$ their parent trajectories at step i , starting from the last resampling step of the algorithm. At each time step of the algorithm the probability of each trajectory to survive is biased by a factor $W_n^{(i)}$. We thus obtain a biased distribution $\mathbb{P}_k(\{x(t)\}_{0 \leq t \leq T_a})$ over the trajectories which is linked to the unbiased probability $\mathbb{P}_0(\{x(t)\}_{0 \leq t \leq T_a})$ to observe the trajectories under the model dynamics by the relation [Ragone et al., 2018]:

$$\mathbb{P}_k(\{x_n(t)\}_{0 \leq t \leq T_a}) = \left(\prod_{i=1}^{I-1} W_n^{(i)} \right) \mathbb{P}_0(\{x_n(t)\}_{0 \leq t \leq T_a}) \quad (\text{A.3})$$

$$= \frac{\exp\left(k \int_0^{T_a} A_n(t) dt\right)}{\tilde{Z}} \mathbb{P}_0(\{x_n(t)\}_{0 \leq t \leq T_a}). \quad (\text{A.4})$$

\tilde{Z} is a normalisation factor that can be computed from the algorithm as $\tilde{Z} = \prod_{i=1}^{I-1} Z_i$. By inverting equation (A.3) we can obtain an unbiased estimator for any observable $O(\{x(t)\}_{0 \leq t \leq T_a})$ of the dynamics:

$$\mathbb{E}_0 [O(\{x(t)\}_{0 \leq t \leq T_a})] \sim_{N \rightarrow +\infty} \sum_{n=1}^N p_n O(\{x_n(t)\}_{0 \leq t \leq T_a}) \quad (\text{A.5})$$

$$\text{with } p_n = \frac{1}{N} \left(\prod_{i=1}^{I-1} W_n^{(i)} \right)^{-1} = \frac{1}{N} \frac{\exp \left(k \int_0^{T_a} A(t) dt \right)}{\tilde{Z}} \quad (\text{A.6})$$

The probabilities p_n can be readily computed from the algorithm by keeping track of the weights $W_n^{(i)}$. Formulae (A.5, A.6) allow us to compute empirical expectations of any observable by performing an ensemble average weighted by the probabilities p_n . When we have several experiment, expectations are computed separately for each independent experiment. Then we perform a second average over the experiments.

Note that, in practice, the rate of convergence in (A.5) depends on the observable O . The rate of convergence is good (efficiency of the algorithm) to compute $\mathbb{P}_0 \left(\frac{1}{T_a} \int_0^{T_a} A(t) dt > a \right) = \mathbb{E}_0 \left[\mathbf{1} \left(\frac{1}{T_a} \int_0^{T_a} A(t) dt > a \right) \right]$ and composite statistics conditioned on $\frac{1}{T_a} \int_0^{T_a} A(t) dt > a$, provided that we choose a biasing parameter k adapted to the target value a . It would however be slower than a classical ensemble simulation to compute unconditioned statistics.

Appendix B. Estimation of confidence intervals and significance levels

Confidence intervals Let $\{X_i\}_{1 \leq i \leq n}$ be n i.i.d. samples from a random variable X . We want to build a confidence interval for the true mean μ_X of X . We compute its empirical mean $\bar{X} = (\sum_{i=1}^n X_i)/n$ and unbiased standard deviation estimator $S = \sqrt{\sum_{i=1}^n (X_i - \bar{X})^2 / (n-1)}$. We assume that the random variable

$$T := \frac{\bar{X} - \mu_X}{S/\sqrt{n}} \quad (\text{B.1})$$

follows the t -distribution with $(n-1)$ degrees of freedom t_{n-1} . We note $t_{\nu, \alpha}$ the quantile of order $1-\alpha$ from the distribution t_ν . Then T belongs to the interval $[-t_{n-1, \alpha/2}, t_{n-1, \alpha/2}]$ with probability $1-\alpha$. This is equivalent to:

$$\mathbb{P} \left(\mu_X \in \left[\bar{X} - t_{n-1, \alpha/2} \frac{S}{\sqrt{n}}; \bar{X} + t_{n-1, \alpha/2} \frac{S}{\sqrt{n}} \right] \right) = 1 - \alpha. \quad (\text{B.2})$$

We use (B.2) with $\alpha = 0.05$ to build 95% confidence intervals.

Significance levels We compute $\tilde{T} := \sqrt{n} \bar{X} / S$ and $q = F_{n-1}(\tilde{T})$ where F_{n-1} is the cumulative distribution function of t_{n-1} . \tilde{T} is a test of the null hypothesis that $\mu_X = 0$. If $q > 1/2$ then $\mu_X > 0$ with probability q . Conversely if $q < 1/2$ then $\mu_X < 0$ with probability $1-q$.

Appendix C. Extreme Value Theory fit

We performed an EVT fit, using the Peak Over Threshold approach with a threshold of 36.9°C. We use R package *extRemes 2.0* [Gilleland and Katz, 2016]. Figure C1 shows the estimates for different range of the threshold, supporting the choice of 36.9°C (see sections 4.3.4 and 4.4 of [Coles, 2001] for details). The resulting estimate for the shape

parameter is $\xi = -0.336$ ($[-0.470, 0.235]$ 95% C.I.) and the scale parameter is $\sigma = 0.251$ ($[0.179, 0.440]$ 95% C.I.). Figure C2 shows diagnostic plots of the estimated model. The confidence intervals on the fit parameters and return levels are estimated via the profile likelihood method.

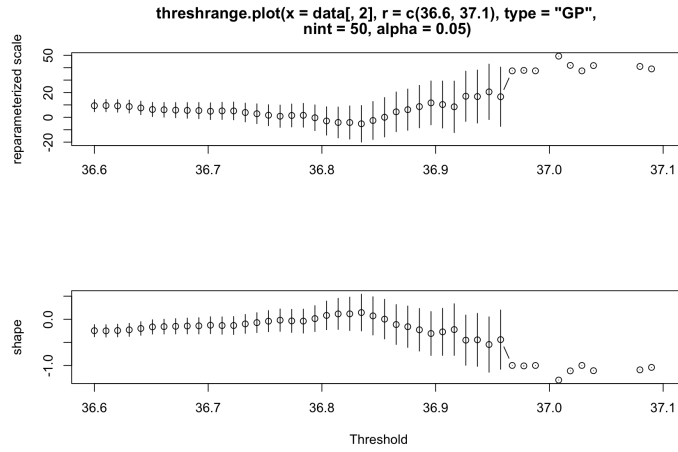


Figure C1. Reparametrized scale and shape parameters estimates for different choice of the threshold.

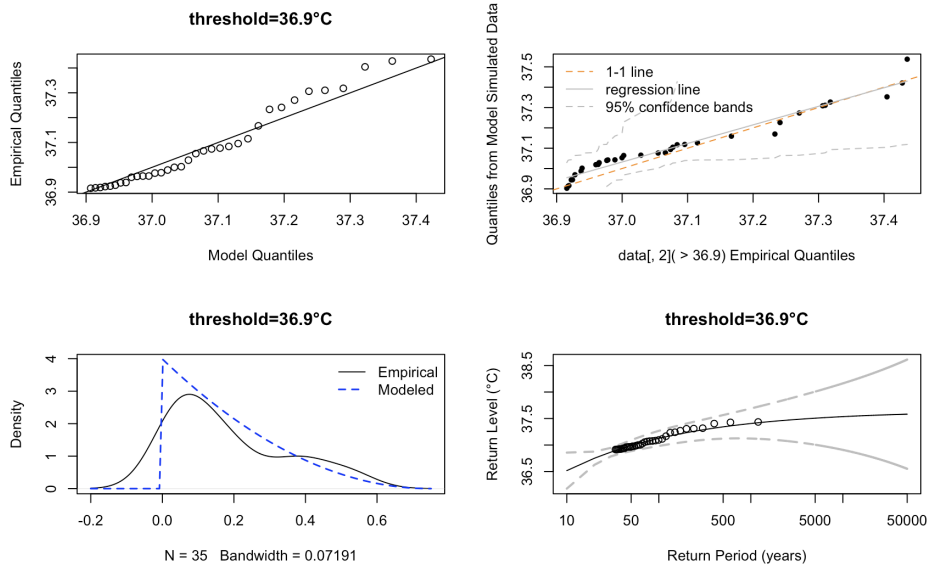


Figure C2. Diagnostic plots for the POT fit with threshold 36.9°C . Quantile-quantile plot (top left), quantiles from a sample drawn from the fitted GPD distribution against the empirical data quantiles with 95% confidence bands (top right), density plots of empirical data and fitted GPD distribution (bottom left), and return level plot with 95% pointwise normal approximation confidence intervals (bottom right).

In figure 3 we compare the 95% confidence interval of the algorithm with the one from the EVT fit estimated via two different methods: profile likelihood and a non-parametric bootstrap. We give below a brief description of each method.

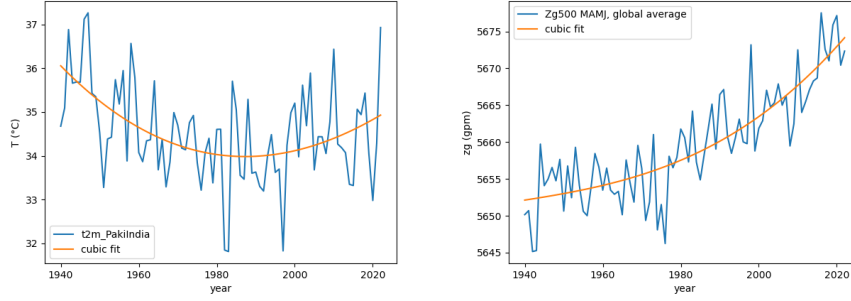


Figure D1. Time series of yearly MAMJ average of t2m averaged over the target region (a) and global average of Zg500 (b). A trend has been computed with a cubic fit (orange lines) and removed from the data before computing the correlation map shown in figure 9.

Profile likelihood Suppose the EVT model is determined by two parameters $(\theta^{(1)}, \theta^{(2)})$. Let $\ell(\theta^{(1)}, \theta^{(2)})$ be the log-likelihood of the model. The profile likelihood along $\theta^{(1)}$ is defined as $\ell_p(\theta^{(1)}) = \max_{\theta^{(2)}} \ell(\theta^{(1)}, \theta^{(2)})$. This can be generalized to an arbitrary number of parameters. Let $\hat{\theta}_0$ be the maximum likelihood estimate of the model $(\theta^{(1)}, \theta^{(2)})$. Under suitable regularity conditions:

$$D_p(\theta^{(1)}) = 2 \left(\ell(\hat{\theta}_0) - \ell_p(\theta^{(1)}) \right) \sim \chi_1^2. \quad (\text{C.1})$$

Let c_α be the $(1 - \alpha)$ quantile of the χ_1^2 distribution. Then $C_\alpha = \{\theta^{(1)} | D_p(\theta^{(1)}) \leq c_\alpha\}$ is a $(1 - \alpha)$ confidence interval for $\theta^{(1)}$. A more comprehensive explanation can be found in sections 2.6.6 and 2.7 of [Coles, 2001]. The method is implemented in R *extRemes 2.0*.

Non-parametric bootstrap The generalized Pareto distribution (GPD) was fitted on the exceedances over the 36.9°C threshold. We generated 2000 bootstrap samples by resampling with replacement from the observed exceedances and fitted a GPD for each sample. For each return time, we define the lower (resp. higher) bound of the 95% CI as the 2.5 (resp. 97.5) quantile of the bootstrap distribution.

Appendix D. Detrending of the reanalysis datasets

We detrended the reanalysis datasets by performing a cubic fit to the time series of yearly MAMJ averages of t2m averaged over the target region and global average of Zg500 (see fig. D1). The resulting trends were then subtracted uniformly to all days and grid points.

References

- [Ciais et al., 2005] Ciais, P., Reichstein, M., Viovy, N., Granier, A., Ogée, J., Allard, V., Aubinet, M., Buchmann, N., Bernhofer, C., Carrara, A., Chevallier, F., De Noblet, N., Friend, A. D., Friedlingstein, P., Grünwald, T., Heinesch, B., Keronen, P., Knohl, A., Krinner, G., Loustau, D., Manca, G., Matteucci, G., Miglietta, F., Ourcival, J. M., Papale, D., Pilegaard, K., Rambal, S., Seufert, G., Soussana, J. F., Sanz, M. J., Schulze, E. D., Vesala, T., and Valentini, R. (2005). Europe-wide reduction in primary productivity caused by the heat and drought in 2003. *Nature*, 437(7058):529–533.
- [Coles, 2001] Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. Springer, London Berlin Heidelberg.
- [Cozian, 2023] Cozian, B. (2023). *Computing Climate Extreme Events and Extremes of Production of Renewable Energy Using Rare Events Algorithms*. PhD thesis, Université Claude Bernard Lyon 1.
- [Dash et al., 2022] Dash, S. K., Jenamani, R. K., and Mohapatra, M. (2022). India’s prolonged heatwave linked to record poor summer rains. *natindia*.
- [Del Moral and Garnier, 2005] Del Moral, P. and Garnier, J. (2005). Genealogical particle analysis of rare events. *Ann. Appl. Probab.*, 15(4).
- [Dimitrova et al., 2021] Dimitrova, A., Ingole, V., Basagaña, X., Ranzani, O., Milà, C., Ballester, J., and Tonne, C. (2021). Association between ambient temperature and heat waves with mortality in South Asia: Systematic review and meta-analysis. *Environment International*, 146:106170.
- [Faranda et al., 2022] Faranda, D., Bourdin, S., Ginesta, M., Krouma, M., Noyelle, R., Pons, F., Yiou, P., and Messori, G. (2022). A climate-change attribution retrospective of some impactful weather extremes of 2021. *Weather and Climate Dynamics*, 3(4):1311–1340.
- [Faranda et al., 2023] Faranda, D., Messori, G., Coppola, E., Alberti, T., Vrac, M., Pons, F., Yiou, P., Saint Lu, M., Hisi, A. N. S., Brockmann, P., Dafis, S., and Vautard, R. (2023). *ClimaMeter: Contextualising Extreme Weather in a Changing Climate*. Preprint, Other aspects of weather and climate dynamics.
- [Fischer et al., 2021] Fischer, E. M., Sippel, S., and Knutti, R. (2021). Increasing probability of record-shattering climate extremes. *Nat. Clim. Chang.*, 11(8):689–695.
- [Fraedrich et al., 2005] Fraedrich, K., Jansen, H., Kirk, E., Luksch, U., and Lunkeit, F. (2005). The Planet Simulator: Towards a user friendly model. *metz*, 14(3):299–304.
- [Gessner et al., 2021] Gessner, C., Fischer, E. M., Beyerle, U., and Knutti, R. (2021). Very rare heat extremes: Quantifying and understanding using ensemble re-initialization. *Journal of Climate*, pages 1–46.
- [Giardina et al., 2011] Giardina, C., Kurchan, J., Lecomte, V., and Tailleur, J. (2011). Simulating Rare Events in Dynamical Processes. *J Stat Phys*, 145(4):787–811.
- [Gilleland and Katz, 2016] Gilleland, E. and Katz, R. W. (2016). **extRemes 2.0**: An Extreme Value Analysis Package in R. *J. Stat. Soft.*, 72(8).
- [Herein et al., 2017] Herein, M., Drótos, G., Haszpra, T., Márffy, J., and Tél, T. (2017). The theory of parallel climate realizations as a new framework for teleconnection analysis. *Sci Rep*, 7(1):44529.
- [Hersbach et al., 2020] Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.
- [Hoerling et al., 2013] Hoerling, M., Kumar, A., Dole, R., Nielsen-Gammon, J. W., Eischeid, J., Perlwitz, J., Quan, X.-W., Zhang, T., Pegion, P., and Chen, M. (2013). Anatomy of an Extreme

- Event. *Journal of Climate*, 26(9):2811–2832.
- [Horton et al., 2016] Horton, R. M., Mankin, J. S., Lesk, C., Coffel, E., and Raymond, C. (2016). A Review of Recent Advances in Research on Extreme Heat Events. *Curr Clim Change Rep*, 2(4):242–259.
- [Hulme, 2020] Hulme, M., editor (2020). *Contemporary Climate Change Debates: A Student Primer*. Routledge, Abingdon, Oxon ; New York, NY.
- [Im et al., 2017] Im, E.-S., Pal, J. S., and Eltahir, E. A. B. (2017). Deadly heat waves projected in the densely populated agricultural regions of South Asia. *Sci. Adv.*, 3(8):e1603322.
- [Jézéquel et al., 2018] Jézéquel, A., Yiou, P., and Radanovics, S. (2018). Role of circulation in European heatwaves using flow analogues. *Clim Dyn*, 50(3):1145–1159.
- [Justine et al., 2023] Justine, J., Monteiro, J. M., Shah, H., and Rao, N. (2023). The diurnal variation of wet bulb temperatures and exceedance of physiological thresholds relevant to human health in South Asia. *Commun Earth Environ*, 4(1):1–11.
- [Miller and Kleidon, 2016] Miller, L. M. and Kleidon, A. (2016). Wind speed reductions by large-scale wind turbine deployments lower turbine efficiencies and set low generation limits. *Proc. Natl. Acad. Sci. U.S.A.*, 113(48):13570–13575.
- [Miloshevich et al., 2023a] Miloshevich, G., Cozian, B., Abry, P., Borgnat, P., and Bouchet, F. (2023a). Probabilistic forecasts of extreme heatwaves using convolutional neural networks in a regime of lack of data. *Phys. Rev. Fluids*, 8(4):040501.
- [Miloshevich et al., 2023b] Miloshevich, G., Rouby-Poizat, P., Ragone, F., and Bouchet, F. (2023b). Robust intra-model teleconnection patterns for extreme heatwaves. *Front. Earth Sci.*, 11:1235579.
- [Mishra et al., 2017] Mishra, V., Mukherjee, S., Kumar, R., and Stone, D. A. (2017). Heat wave exposure in India in current, 1.5 °C, and 2.0 °C worlds. *Environ. Res. Lett.*, 12(12):124012.
- [Monteiro and Caballero, 2019] Monteiro, J. M. and Caballero, R. (2019). Characterization of Extreme Wet-Bulb Temperature Events in Southern Pakistan. *Geophys. Res. Lett.*, 46(17-18):10659–10668.
- [Mora et al., 2017] Mora, C., Dousset, B., Caldwell, I. R., Powell, F. E., Geronimo, R. C., Bielecki, C. R., Counsell, C. W. W., Dietrich, B. S., Johnston, E. T., Louis, L. V., Lucas, M. P., McKenzie, M. M., Shea, A. G., Tseng, H., Giambelluca, T. W., Leon, L. R., Hawkins, E., and Trauernicht, C. (2017). Global risk of deadly heat. *Nature Clim Change*, 7(7):501–506.
- [Noyelle et al., 2023] Noyelle, R., Zhang, Y., Yiou, P., and Faranda, D. (2023). Maximal reachable temperatures for Western Europe in current climate. *Environ. Res. Lett.*, 18(9):094061.
- [Perkins, 2015] Perkins, S. E. (2015). A review on the scientific understanding of heatwaves—Their measurement, driving mechanisms, and changes at the global scale. *Atmospheric Research*, 164–165:242–267.
- [Philip et al., 2020] Philip, S., Kew, S., Van Oldenborgh, G. J., Otto, F., Vautard, R., Van Der Wiel, K., King, A., Lott, F., Arrighi, J., Singh, R., and Van Aalst, M. (2020). A protocol for probabilistic extreme event attribution analyses. *Adv. Stat. Clim. Meteorol. Oceanogr.*, 6(2):177–203.
- [Philip et al., 2022] Philip, S. Y., Kew, S. F., Van Oldenborgh, G. J., Anslow, F. S., Seneviratne, S. I., Vautard, R., Coumou, D., Ebi, K. L., Arrighi, J., Singh, R., Van Aalst, M., Pereira Marghidan, C., Wehner, M., Yang, W., Li, S., Schumacher, D. L., Hauser, M., Bonnet, R., Luu, L. N., Lehner, F., Gillett, N., Tradowsky, J. S., Vecchi, G. A., Rodell, C., Stull, R. B., Howard, R., and Otto, F. E. L. (2022). Rapid attribution analysis of the extraordinary heat wave on the Pacific coast of the US and Canada in June 2021. *Earth Syst. Dynam.*, 13(4):1689–1713.
- [Plotkin et al., 2019] Plotkin, D. A., Webber, R. J., O’Neill, M. E., Weare, J., and Abbot, D. S. (2019). Maximizing Simulated Tropical Cyclone Intensity With Action Minimization. *J Adv Model Earth Syst*, 11(4):863–891.
- [Ragone and Bouchet, 2020] Ragone, F. and Bouchet, F. (2020). Computation of Extreme Values of Time Averaged Observables in Climate Models with Large Deviation Techniques. *J Stat Phys*, 179(5):1637–1665.
- [Ragone and Bouchet, 2021] Ragone, F. and Bouchet, F. (2021). Rare Event Algorithm Study of Extreme Warm Summers and Heatwaves Over Europe. *Geophys Res Lett*, 48(12).

- [Ragone et al., 2018] Ragone, F., Wouters, J., and Bouchet, F. (2018). Computation of extreme heat waves in climate models using a large deviation algorithm. *Proc Natl Acad Sci USA*, 115(1):24–29.
- [Ratnam et al., 2016] Ratnam, J. V., Behera, S. K., Ratna, S. B., Rajeevan, M., and Yamagata, T. (2016). Anatomy of Indian heatwaves. *Sci Rep*, 6(1):24395.
- [Robine et al., 2008] Robine, J.-M., Cheung, S. L. K., Le Roy, S., Van Oyen, H., Griffiths, C., Michel, J.-P., and Herrmann, F. R. (2008). Death toll exceeded 70,000 in Europe during the summer of 2003. *Comptes Rendus Biologies*, 331(2):171–178.
- [Rogers et al., 2021] Rogers, C. D. W., Ting, M., Li, C., Kornhuber, K., Coffel, E. D., Horton, R. M., Raymond, C., and Singh, D. (2021). Recent Increases in Exposure to Extreme Humid-Heat Events Disproportionately Affect Populated Regions. *Geophysical Research Letters*, 48(19):e2021GL094183.
- [Rohini et al., 2016] Rohini, P., Rajeevan, M., and Srivastava, A. K. (2016). On the Variability and Increasing Trends of Heat Waves over India. *Sci Rep*, 6(1):26153.
- [Romanello et al., 2021] Romanello, M., McGushin, A., Di Napoli, C., Drummond, P., Hughes, N., Jamart, L., Kennard, H., Lampard, P., Solano Rodriguez, B., Arnell, N., Ayeb-Karlsson, S., Belesova, K., Cai, W., Campbell-Lendrum, D., Capstick, S., Chambers, J., Chu, L., Ciampi, L., Dalin, C., Dasandi, N., Dasgupta, S., Davies, M., Dominguez-Salas, P., Dubrow, R., Ebi, K. L., Eckelman, M., Ekins, P., Escobar, L. E., Georgeson, L., Grace, D., Graham, H., Gunther, S. H., Hartinger, S., He, K., Heaviside, C., Hess, J., Hsu, S.-C., Jankin, S., Jimenez, M. P., Kelman, I., Kiesewetter, G., Kinney, P. L., Kjellstrom, T., Kniveton, D., Lee, J. K. W., Lemke, B., Liu, Y., Liu, Z., Lott, M., Lowe, R., Martinez-Urtaza, J., Maslin, M., McAllister, L., McMichael, C., Mi, Z., Milner, J., Minor, K., Mohajeri, N., Moradi-Lakeh, M., Morrissey, K., Munzert, S., Murray, K. A., Neville, T., Nilsson, M., Obradovich, N., Sewe, M. O., Oreszczyn, T., Otto, M., Owfi, F., Pearman, O., Pencheon, D., Rabbaniha, M., Robinson, E., Rocklöv, J., Salas, R. N., Semenza, J. C., Sherman, J., Shi, L., Springmann, M., Tabatabaei, M., Taylor, J., Trinanes, J., Shumake-Guillemot, J., Vu, B., Wagner, F., Wilkinson, P., Winning, M., Yglesias, M., Zhang, S., Gong, P., Montgomery, H., Costello, A., and Hamilton, I. (2021). The 2021 report of the Lancet Countdown on health and climate change: Code red for a healthy future. *The Lancet*, 398(10311):1619–1662.
- [Seneviratne et al., 2021] Seneviratne, S. I., Zhang, X., Adnan, M., Badi, W., Dereczynski, C., Di Luca, A., Ghosh, S., Iskandar, I., Kossin, J., Lewis, S., Otto, F., Pinto, I., Satoh, M., Vicente-Serrano, S., Wehner, M., and Zhou, B. (2021). Weather and climate extreme events in a changing climate. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 1513–1766. Cambridge university press edition.
- [Shepherd, 2016] Shepherd, T. G. (2016). A Common Framework for Approaches to Extreme Event Attribution. *Curr Clim Change Rep*, 2(1):28–38.
- [Stott et al., 2004] Stott, P. A., Stone, D. A., and Allen, M. R. (2004). Human contribution to the European heatwave of 2003. *Nature*, 432(7017):610–614.
- [Trenberth et al., 2015] Trenberth, K. E., Fasullo, J. T., and Shepherd, T. G. (2015). Attribution of climate extreme events. *Nature Clim Change*, 5(8):725–730.
- [Van Oldenborgh et al., 2022] Van Oldenborgh, G. J., Wehner, M. F., Vautard, R., Otto, F. E. L., Seneviratne, S. I., Stott, P. A., Hegerl, G. C., Philip, S. Y., and Kew, S. F. (2022). Attributing and Projecting Heatwaves Is Hard: We Can Do Better. *Earth’s Future*, 10(6):e2021EF002271.
- [Webber et al., 2019] Webber, R. J., Plotkin, D. A., O’Neill, M. E., Abbot, D. S., and Weare, J. (2019). Practical rare event sampling for extreme mesoscale weather. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(5):053109.
- [Wegren, 2011] Wegren, S. K. (2011). Food Security and Russia’s 2010 Drought. *Eurasian Geography and Economics*, 52(1):140–156.
- [White et al., 2022] White, R. H., Kornhuber, K., Martius, O., and Wirth, V. (2022). From Atmospheric Waves to Heatwaves: A Waveguide Perspective for Understanding and Predicting Concurrent, Persistent, and Extreme Extratropical Weather. *Bulletin of the American*

- Meteorological Society, 103(3):E923–E935.
- [Wouters et al., 2023] Wouters, J., Schiemann, R. K. H., and Shaffrey, L. C. (2023). Rare Event Simulation of Extreme European Winter Rainfall in an Intermediate Complexity Climate Model. J Adv Model Earth Syst, 15(4):e2022MS003537.
- [Yiou and Jézéquel, 2020] Yiou, P. and Jézéquel, A. (2020). Simulation of extreme heat waves with empirical importance sampling. Geoscientific Model Development, 13(2):763–781.
- [Zachariah et al., 2023] Zachariah, M., Arulalan, T., AchutaRao, K., Saeed, F., Jha, R., Dhasmana, M. K., Mondal, A., Bonnet, R., Vautard, R., Philip, S., Kew, S., Vahlberg, M., Singh, R., Arrighi, J., Heinrich, D., Thalheimer, L., Marghidan, C. P., Kapoor, A., Van Aalst, M., Raju, E., Li, S., Sun, J., Vecchi, G., Yang, W., Hauser, M., Schumacher, D. L., Seneviratne, S. I., Harrington, L. J., and Otto, F. E. L. (2023). Attribution of 2022 early-spring heatwave in India and Pakistan to climate change: Lessons in assessing vulnerability and preparedness in reducing impacts. Environ. Res.: Climate, 2(4):045005.