

Increasing risk of mass human heat mortality if historical weather patterns recur

In the format provided by the
authors and unedited

Contents

This Supplementary Information file contains: Supplementary Figures 1–10, Supplementary Tables 1–2.




Figure S1: Temperature anomalies in each European region. Each distribution shows the range of June-July-August daily temperature anomalies (relative to day-of-year mean) in ERA5 (black) and the CMIP6 data used for machine learning training (red), in the three regions used in training.




Figure S2: Machine learning predictions of unseen GCM data. Relationship between predicted values and true values simulated by GCMs, when the machine learning algorithm has been trained on a subset of GCM data and evaluated on the held-out sample.




Figure S3: GCM-specific warming trends and mortality projections. Top row: Change in JJA mean temperature averaged across the three European regions from the raw GCM training data when moving from 0 to 3 °C (x-axis), plotted against the change in event intensity from 0 to 3 °C predicted by the machine learning approach (y-axis). Colors denote different GCMs, with their equilibrium climate sensitivity (ECS) noted in parentheses in the upper left panel. Different dots for a given GCM correspond to predictions using different random seeds in the machine learning training. Bottom row: Peak mortality projected for each event for annual GMT 3 °C above the pre-industrial baseline, separated according to the GCM used to train the neural network. Bar height shows average prediction and line spans the 95% confidence interval across random seeds and regression bootstraps ($N = 2,500$ for each GCM). Lower text denotes the ECS for each GCM from Meehl et al.⁵².




Figure S4: Temperature for each event in IPCC AR6 regions. As in the lower points in Fig. 2, but for each of the three IPCC regions for which we train the CNNs.




Figure S5: Observed and simulated rates of warming in Europe. Each line plots the annual global mean temperature change relative to 1850-1900 on the x-axis against 20-year running mean of Europe-wide temperature relative to 2000-2019 on the y-axis. “Europe-wide” temperature means the average across NUTS regions. 20-year running means are used because we find that the 20-year average 2000-2019 temperature of each NUTS region shapes the exposure response function in Fig. 2. Gray lines show each individual CMIP6 model, black line shows the ensemble mean, dark red line shows the E-OBS observations, and red dashed line shows a linear fit to the observations. X-axis is truncated to focus on lower warming levels to enable visual comparison between observations and models. 1999 is used as the first year of the slope calculation because using a 20-year running mean drops every previous year between 1980 and 1999.

Regional scaling with global warming




Figure S6: Pattern scaling coefficients across European regions. Linear coefficient between annual global temperature and running 20-year mean annual temperature for each region. Coefficients are averaged across 100 random samples of pooled model-year populations. Map was created using cartopy v0.24.0 with borders from the Eurostat database.




Figure S7: Alternative exposure-response curves. Panel (a) shows our main exposure-response function, which uses all-age mortality over 2015-2019 (same as Fig. 2). Panel (b) shows the same regression specification using over-65 mortality. Panel (c) shows the all-age response over 2000-2019 instead of 2015-2019. Panel (d) again shows our main exposure-response function, which uses daily mean temperature. Panels (e) and (f) show the same specification using daily maximum (e) and minimum (f) temperature. Note that the x-axes are scaled differently in (e) and (f) to account for the different observed ranges of the temperature metrics. All y-axes are standardized. Solid line shows mean response and shading shows 95% confidence interval across 500 regression bootstraps in each case.




Figure S8: Heterogeneity in exposure-response function by climate and income. Both panels show results from a single regression model where the polynomials of daily mean temperature are interacted with continuous values for both a region's mean climate and its mean log income. The left panel shows responses for the coolest, middle, and warmest tercile (analogous to Fig. 2), when evaluated at the middle tercile of income. Conversely, the right panel shows the responses for the poorest, middle, and richest tercile of income, when evaluated at the middle tercile of climate.




Figure S9: Effect of a hot day across lags. Both panels show the mortality effect of a 30 °C day relative to a 20 °C day, at a series of lags relative to the week of mortality. Lag 0 means contemporaneous temperature, lag 1 means temperature the week before, and so on. In our main analysis, we use 3 lags (left panel), but we also test a model with 6 lags (right panel). Dots and lines show mean response and shading shows 95% confidence interval across 500 regression bootstraps.




Figure S10: Alternative approaches to estimating adaptation. As in Fig. 4, but with two additional versions of adaptation. Orange bars show our main result without adaptation. Green bars show the effect of adaptation using CMIP6 pattern scaling, as in main text Fig. 4. Light blue bars show adaptation when regional 20-year-mean temperatures are assumed to scale 1-to-1 with annual GMT. Dark blue bars show adaptation when the rate of regional warming (relative to GMT change) is extrapolated linearly from observations, with the slope calculated over 1999-2023 (Fig. S5). Bar heights show average projections and error bars show 95% range across each combination of GCM, random seed, and regression bootstrap ($N = 12,500$).

Event	GMT	Peak mortality	Cumulative mortality
July 1994	1.5	14300	32100
July 1994	2.0	17800	40700
July 1994	3.0	26500	62000
July 1994	4.0	37500	89100
August 2003	1.5	17800	26500
August 2003	2.0	21900	33000
August 2003	3.0	32000	48900
August 2003	4.0	45100	69500
July 2006	1.5	13000	21700
July 2006	2.0	16600	27900
July 2006	3.0	25600	43500
July 2006	4.0	37500	63800
June 2019	1.5	10800	11000
June 2019	2.0	13100	13800
June 2019	3.0	18800	20900
June 2019	4.0	26000	29800
August 2023	1.5	11000	18300
August 2023	2.0	13800	23500
August 2023	3.0	20900	36700
August 2023	4.0	30200	53800

Table S1: Europe-wide mortality for each event. Each row shows the maximum weekly excess deaths (“peak”) and cumulative excess deaths for each event at each global mean temperature (“GMT”). We note that because the events differ slightly in their durations, peak single-week mortality is more directly comparable across events than cumulative mortality. Values are rounded to three significant figures to avoid excessive precision given multiple dimensions of uncertainty in these calculations.

Event	GMT	Peak mortality from warming	Percent from warming
July 1994	1.5	7700	54%
July 1994	2.0	11200	63%
July 1994	3.0	20100	76%
July 1994	4.0	31300	84%
August 2003	1.5	8800	49%
August 2003	2.0	12900	59%
August 2003	3.0	23000	72%
August 2003	4.0	36100	80%
July 2006	1.5	7600	59%
July 2006	2.0	11200	67%
July 2006	3.0	20300	79%
July 2006	4.0	32100	86%
June 2019	1.5	5300	49%
June 2019	2.0	7600	58%
June 2019	3.0	13300	71%
June 2019	4.0	20500	79%
August 2023	1.5	5900	53%
August 2023	2.0	8700	63%
August 2023	3.0	15800	76%
August 2023	4.0	25100	83%

Table S2: Climate change-driven mortality for each event. The “peak mortality from climate change” row shows the peak weekly excess deaths for each event at each GMT relative to the peak of the event at 0 °C, meaning only the component of mortality due to anthropogenic intensification of the event. The “percent from warming” column shows the percent of overall peak mortality (Table S1) due to climate change. Values are rounded to three significant figures to avoid excessive precision given multiple dimensions of uncertainty in these calculations.