Climate Central Temperature Attribution System Methods

Andrew Pershing, Daniel M. Gilford, and Joseph Giguere

Version 2.0, March 2023

This documentation and software version supports the launch of the ERA5 update of the Climate Shift Index™.

The system builds on the peer-reviewed methodology published in Gilford et al. (2022, hereafter, G22). Most importantly, it uses their approach of computing scale factors (β, i.e. the regression coefficient between local temperatures, T and global mean temperature, GMT) and then uses them to transform an observed climatology into a set of modern and counterfactual climatologies used to compute an attribution metric. It also uses a similar strategy of directly attributing temperature distribution changes based on historical and counterfactual climate model simulations.

The major differences between this version and G22 are:

- The historical climatology used as the basis for our historical climate data and empirical Climate Shift Index calculations has been replaced with the fifth generation ECMWF reanalysis for the global climate and weather (ERA5). The switch to ERA5 allows us to quadruple the resolution of the CSI and produce more reliable, more robust, attribution estimates.
- The model-based method is computed with paired forced and pre-industrial CMIP6 models.
- The quantile-based empirical method (CDF values) for assigning the likelihood of the observed temperatures has been replaced with a method that fits a skew-normal distribution to the observations or models, and then uses that distribution to estimate likelihoods associated with modern and counterfactual climates.
- The occurrence ratio (OR) is used as the attribution metric instead of the probability ratio (PR).
- A screening procedure is implemented to mask out regions where weather model forecasting errors could confound daily attribution estimates.

Details and rationale are described below.

Historical Climatologies

We use the fifth generation ECMWF reanalysis for the global climate and weather (ERA5) as the basis for our historical climatologies. These climatologies drive both the climate model bias-adjustment scheme and the empirical attribution calculations (see below). ERA5 is a reconstruction of the state of the atmosphere that uses a sophisticated weather model to synthesize large volumes of meteorological data. It is one of the most widely-used datasets for studies of weather and climate. ERA5 is available starting in 1950 and running to present at 0.25° by 0.25° resolution (~16 miles or ~25 km).

The data are publicly available at 3-hourly resolution. To get daily maximum (T_{max}) and minimum (T_{min}), we use the xarray `groupby('time.day').max()` (or `min()`) function. We downloaded the entire record from 1979–2021 and processed to T_{min} and T_{max}.

We define a reference climatology for day = d as the ERA5 data for the 31 day period centered on day d over the years 1991–2020 (i.e. the NOAA climate normal period). To represent the seasonal variations in climatology, we compute attribution estimates for 24 day (d) periods across the year: the 1st and 15th calendar day of each month.

Improvement over CFS climatologies (previous version)

ERA5 represents a marked improvement over the Climate Forecast System (CFS) climatologies used in version 1.0 of the CC-TAS methods (June 2022). The horizontal resolution of the underlying data has increased fourfold, and now matches the resolution of the daily GFS forecasts (eliminating the need to regrid between the two with xesmf.Regridder). Unlike CFS, ERA5 has no model discontinuities that might cause spurious historical trend (and hence, attribution) analyses; its stability and wide usage across the climate science and attribution communities gives confidence to our approach and supports intercomparisons with other attribution estimates. ERA5's climatology also provides a more reliable basis (than CFS) for the climate model bias-adjustments underpinning the model-based approach.

Daily Observations

For our nowcast and forecasts operations, we get temperature data from NOAA's Global Forecast System (GFS). GFS is an operational model that is run several times a day to support weather forecasts. We download the forecast for the next 5 days and then compute each day's min and max temperatures. GFS is run at a 0.25° by 0.25° resolution, matching the ERA5 climatologies.

Note that we can also apply the attribution methods directly to the ERA5 record, e.g. to hindcast past conditions that may have been attributable.

Scale Factors

Scale factors (β values) are central to the empirical attribution methods. They are computed by regressing a temperature time series for a particular quantile against global mean temperature, which we take from the [Hadley Centre's HadCRUT5 temperature product].

For each quantile level, q_i, we computed $\beta_{q_i}(d, x)$ using the ERA5 climatologies at each location x and for each of the 24 d periods spread across the calendar year. By limiting the number of quantiles, we can gain computational efficiency without noticeable changes in the output. We found that using 21 evenly-spaced quantiles between 0.02 and 0.98—including the median ($q = 0.5$)—achieves this balance. The R^2 statistic and the p-value of scale factor regressions were retained to accomodate further statistical analyses.

Skew Normal Distributions

The most significant change between this version and the methods described in G22 is the use of parameterized distribution functions. We found that both T_{max} and T_{min} are well-approximated by the skew-normal distribution. This function is described by the location parameter (analogous to a median) = L, the scale parameter (analogous to the variance) = S, and the shape parameter = A that defines the level of skewness.

The SciPy stats package includes [a function] to fit each built-in distribution to data. We used this function to fit skew-normal distributions to the reference climatologies. The counterfactual distributions for median scaling are constructed from the reference distributions by adding $\beta * (GMT_{counter} - GMT_{reference})$ to the location parameter. The modern distributions for median scaling are created in the same way.

The quantile scaling method poses a challenge to using parameterized distributions like the skew normal. Because this method transforms the approximation of the CDF represented by the quantile distributions, it impossible to use the built-in routines. We use SciPy's `optimize.leastsq` method to find the A, L, S at each location that minimizes the residuals between the observed quantile distribution and the skew-normal CDF using A, L, S. We use the least-squares fit to the CDF for the modern and counterfactual distributions.

We also fit skew-normal distributions to the data from the climate models (see below).

Estimating the Modern Climate

The G22 methodology uses annual GMT to define the global climate state. Ideally, we would like to define $GMT(y)$ as the mean over some interval p of years (for example, 11) surrounding year y. However, this definition creates a problem for the last $\frac{p}{2}$ years in the record. This is especially troubling for our main application of estimating the attributable state of the climate for the current year.

One approach we considered was to define GMT as the mean of the prior p years. This is highly conservative in that it will always be cooler than the value we would get using the centered mean. For $p = 11$, it would be about 0.2°C below the "true" GMT. We decided to define $GMT(y)$ as,

- the centered 11 year mean, if $y < (2022 - 6) = 2016$
- $r_{30}y + c_{30}$, if $y \geq 2016$

where r_{30} and c_{30} are the coefficients from the linear regression of GMT against year over the 30 preceding years (e.g. 1992–2021 for $y = 2022$). Note that the values estimated using the 30 year regression are strongly related to those using the 11-year mean ($R^2 = 0.97$, $p < 0.01$).

Model-based Method

We accessed CMIP6 output (daily 'tasmax' and 'tasmin') from the Google Cloud archive. We found 24 models that had historical runs, forced projections (SSP3-7.0 if available, SSP5-8.5 in some cases), and pre-industrial control runs (Table 1). We concatenated the historical and projections for each model to create a single "forced" simulation for each model.

Table 1. List of climate models from CMIP6	
Organization	**Model Names**
Australian Community Climate and Earth System Simulator	ACCESS-CM2, ACCESS-ESM1-5
Alfred Wegener Institute	AWI-CM-1-1-MR
Euro-Mediterranean Center on Climate Change	CMCC-ESM2
Centre National de Recherches Météorologiques	CNRM-CM6-1-HR, CNRM-CM6-1, CNRM-ESM2-1
Canadian Centre for Climate Modelling and Analysis	CanESM5
EC Earth Consortium	EC-Earth3-AerChem, EC-Earth3-Veg-LR, EC-Earth3-Veg, EC-Earth3
Geophysical Fluid Dynamics Laboratory	GFDL-CM4, GFDL-ESM4
Institute for Numerical Mathematics, Russian Academy of Sciences	NM-CM4-8, INM-CM5-0
Institut Pierre Simon Laplace	IPSL-CM6A-LR
JAMSTEC, AORI, NIES, R-CCS	MIROC6
Max Planck Institute	MPI-ESM-1-2-HAM, MPI-ESM1-2-HR, MPI, ESM1-2-LR, MRI-ESM2-0
NorESM Climate Modelling Consortium	NorESM2-MM
Research Center for Environmental Changes, Academia Sinica	TaiESM1

We used xesmf's regridder to regrid each model to a common 0.25° by 0.25° grid. For models with coarser resolution, we usued xesmf's `method = bi-linear` , while for models with finer resolution we used `method = conservative` .

As in G22 we used the Lange (2019) methodology to bias-adjust the output of each individual climate model. We used the same 1991–2020 ERA5 climatology as in our empirical method as the reference data for the debiasing. The debiasing trained by the relationship between the reference climate and the forced simulations was likewise applied to debias each paired pre-industrial control simulation.

For each model, we identified the first year when its representation of GMT (smoothed using an 11-year running mean) was greater than or equal to the modern GMT as defined above. We then select the 31-year period centered around that year to assess the climate corresponding to that GMT. For each of the 24 periods in the year, as with the empirical method, we extract the 31-day period centered around the target day. We then fit skew normal distributions to the Tmin and Tmax data. We also extract the same range of days from the last 31 years of the pre-industrial control run and fit a counterfactual skew normal distribution.

Attribution Estimates

We now have several skew normal distributions representing the modern and counterfactual climates (Table 2). The pair of distributions created by each method allow us to independently estimate the change in the likelihood of the temperature in question.

Let $T(d, y)$ be the observed or forecasted temperature for day d in year y. We find d_j = the day from the 24 periods that is closest to d. For each method, we acquire $SN_{counter}(d_j)$ and $SN_{modern}(d_j)$, the files with the parameters for the skew normal distributions for day d_j.

Once we have the skew normal distributions, we use them to get $PDF_{modern}(T)$ and $PDF_{counter}(T)$, the PDF values from each climate and with each method. We then compute:

$$OR(T) = \frac{PDF_{modern}(T)}{PDF_{counter}(T)}$$

which we refer to as the "occurrence ratio."

The occurrence ratio is our main attribution metric. It approximates the ratio of the probability of encountering a temperature close to T in the two climates. Values of $OR > 1$ indicate that anthropogenic climate change has made those conditions more likely, and values < 1 indicate that climate change made those conditions less likely.

Table 2: Summary of the attribution methods. The first row shows the global mean temperature relative to the 1850–1859 mean for the three periods: reference, counterfactual, and modern. The remaining rows summarize the process for creating the skewed normal distributions that are used to estimate likelihoods.

	Reference (1991–2020)	Counterfactual (~1885–1915)	Modern (y=2022)
GMT (rel. 1850–	0.90	-0.06	1.27

	Reference (1991–2020)	Counterfactual (~1885–1915)	Modern (y=2022)
1899)			
Median-scaling (q_{50})	Direct fit to 1991–2020 ERA5 Data	Reference distribution with L shifted by $\beta(GMT_{counter} - GMT_{reference})$	Reference distribution with L shifted by $\beta(GMT_{modern} - GMT_{reference})$
Quantile-scaling (q_i)	Fit to 1991–2020 quantile distribution	Fit to reference quantile distribution shifted by $\beta_q(GMT_{counter} - GMT_{reference})$	Fit to reference quantile distribution shifted by $\beta_q(GMT_{modern} - GMT_{reference})$
Models	N/A	Fit to model M's pre-industrial control run	Fit to model M's forced run

Rationale for the Occurrence Ratio

Using the occurrence ratio as the primary attribution metric is a departure from established attribution methodologies. The probability ratio (sometimes called the hazard ratio) is used in most studies. It is defined as:

$$PR(T) = \frac{1 - CDF_{modern}(T)}{1 - CDF_{counter}(T)}$$

The PR makes a statement about the likelihood of temperatures greater than or equal to T. It is well-suited to the needs of most attribution studies which focus on extreme conditions.

A key difference between our application and most published studies is that we want to be able to make a statement for any T, not just extreme values. We also want to be able to discuss conditions like cold days that are becoming less likely. The OR provides a more intuitive way to do this as it refers to specific conditions and doesn't require the *a priori* choice of direction of change.

Combining Estimates and Accounting for Uncertainty

Our procedure provides two empirical estimates and 24 estimates from the paired climate models. G22 took the mean of the climate model estimates (24 models) and then took the mean of the model average and the two empirical methods. We use this same approach; however, additional testing may indicate a different weighting scheme.

G22 used a Monte Carlo procedure to estimate confidence intervals around the attribution metrics from the two empirical methods. A full implementation of their Monte Carlo approach would require us to fit multiple skew normal distributions for each time period. This is computationally challenging. We therefore focused on developing a conservative methodology that could be applied within our operational time limits. This includes the definition of our new attribution metric, the Climate Shift Index.

Our procedure is:

1. Convert each of the 26 estimates of OR into a climate factor (CF):

$$CF(T) = 2 * \log_2(OR(T))$$

This creates a linear scale centered on 0. Note that CF values above 8 or below -8 were replaced with 8 or -8, respectively.

2. Average the 22 model-based climate factors together to create a single model average. As in G22, we found that the model-based estimates are generally the most conservative estimates (CF values closest to 0).

3. On rare occasions when the median or quantile method could not make an estimate (something observed occasionally for very extreme temperatures), replace that method with the model-based method.

4. If there is any disagreement in the sign of the CF among the median, quantile, and model-average, then the average CF is set to 0.

5. If all three agree on the sign, then average the empirical methods together. Finally, take the average of the two averages.

This gives the value of the climate factor (which can be converted back to OR if desired) at the 0.25° ERA5/GFS grid points. For the initial Climate Central mapping tool, we averaged the climate factor into NOAA climate divisions across the U.S.

We then converted the average climate factor into the categorical Climate Shift Index. For warm temperatures, the Index is

CSI	Descriptor	OR range	CF range
0	No effect	1/5-1.5	-1.17 - 1.17
1	Moderate	1.5-2	1.17-2
2	Strong	2-3	2-3.17
3	Very Strong	3-4	3.17-4
4	Extreme	4-5	4-4.64
5	Exceptional	>5	4.64

The CSI values for cold conditions are defined analogously:

CSI	OR range	CF range
0	1/5-1.5	-1.17 - 1.17
-1	1/2 - 1/1.5	-2 - -1.17
-2	1/3-1/2	-3.17 - -2
-3	1/4-1/3	-4 - -3.17
-4	1/5-1/4	-4.64 - -4
-5	<1/5	-4.64

The CSI scale embeds several strategies meant to make the estimates conservative. First, it creates a very wide range of values around 0. This, plus the criterion that the three methods must agree on the sign, reduces the number of false positives. Second, the categorical nature of the scale means that we are effectively rounding toward zero. An event with an OR of 2.8 ($CF = 2.97$) becomes a 2 on the CSI scale. Finally, the CSI caps the maximum value at 5. This limits extrapolation into the tails of the distributions. The conservative approach and focus on lower bounds is consistent with other attribution studies that found a higher degree of certainty around lower bounds (e.g., Risser and Wehrner, 2017).

Forecast Uncertainty Screening

CSI levels are calculated by taking the latest temperature forecasts from GFS, and comparing them to known temperature distributions from ERA5 climatology. But discrepancies between daily forecasted temperatures and actual observed temperatures could lead to divergent CSI level calculations. We address these potential uncertainties by screening out regions and time periods where forecasting errors are expected to be large enough to confound attribution.

We can quantify the uncertainty in the CSI level due to an error between the daily forecast temperature ($T_{forecast}$) and the daily observed temperature (T_{obs}, which will later be recorded in ERA5). We begin by implementing a Taylor series expansion around $T_{forecast}$ to evaluate how the error could influence our climate factor estimates, CF, at T_{obs}:

$$CF(T_{obs}) = CF(T_{forecast}) \pm \frac{dCF}{dT} * (T_{error}) + [\ldots = 0]$$

Here $\frac{dCF}{dT}$ is the sensitivity of the climate factor to a local change in temperature, and $T_{error} = (T_{obs} - T_{forecast})$ is the magnitude of the temperature error we are considering influential on the daily CF estimates. Rearranging, we define the absolute value of the climate factor error arising from T_{error} as

$$\Gamma_{error} \equiv |CF(T_{obs}) - CF(T_{forecast})| = \frac{dCF}{dT} * (T_{error})$$

which, once evaluated, can be used to screen out error prone regions and time periods.

To solve for Γ_{error}, we assume that temperatures follow a normal distribution with a known mean ($\mu \sim L$) and standard deviation ($\sigma^2 \sim S$), i.e. $T \sim \mathcal{N}(\mu, \sigma^2)$. In the case of a uniform stationary shift in the temperature probability distribution attributable to climate change (i.e., the variance, S, remains fixed over time), the climate factor at a given location (x) and over a period (d) is a function of temperature in modern period (with mean L_{modern}) and counterfactual period (with mean $L_{counter} = L_{modern} - \beta * GMT$):

$$CF(T, x, d) = 2\log_2(\mathcal{N}(L_{modern}, S)) - 2\log_2(\mathcal{N}(L_{modern} - \beta * GMT, S))$$

From this expression, the sensitivity of the climate factor to a change in temperature is given by its derivative, which can be shown to be:

$$\frac{dCF}{dT}(x, d) = 2\log_2(e) * \beta * GMT * \frac{1}{S}$$

which is

$\approx 2.88*$(local sensitivity to mean global climate warming) / (local temperature variance).

We next consider the influence of the forecast error $T_{error} \equiv RMS$, where RMS is the root-mean square error between GFS forecasts and ERA5 historical reanalysis since the GFS update in June 2019; together with the CF sensitivity, we can now calculate Γ_{error}.

Finally, we screen out any locations and periods where $\Gamma_{error} \geq \Gamma_{mask}$, where Γ_{mask} is a free parameter that defines our degree of tolerance for forecast errors affecting CF (in units of CF levels). We tested a range of Γ_{mask} values, and found even modest degrees of screening sufficient to remove the regions with large forecast errors. We take $\Gamma_{mask} = 2$ to balance the retention of data with the curtailment of forecast uncertainties in live Climate Shift Index tool.

The conditions leading to $\Gamma_{error} \geq 2$ are most common at nighttime and in the tropics, where local temperature trends are high, temperature distributions are relatively narrow, and forecast RMS errors are relatively large.

Screened out locations are masked across our Climate Shift Index tools, assigned the value "Currently Unavailable." In the near future, daily real-time CSI levels based on GFS forecasts will be updated once ERA5 data is available (after approximately two weeks). This will allow us to calculate attribution estimates for most of the masked areas, and the updated CSI will then be visible in the CSI Global Map.

Example Results

Our initial public work focused on the contiguous U.S. Our main product is a set of daily maps of the climate factor averaged over the NOAA climate divisions (Figure 2). Here we show the results for June 29, 2021, during the Pacific Northwest heat dome, and December 16, during an unusual winter heat event.

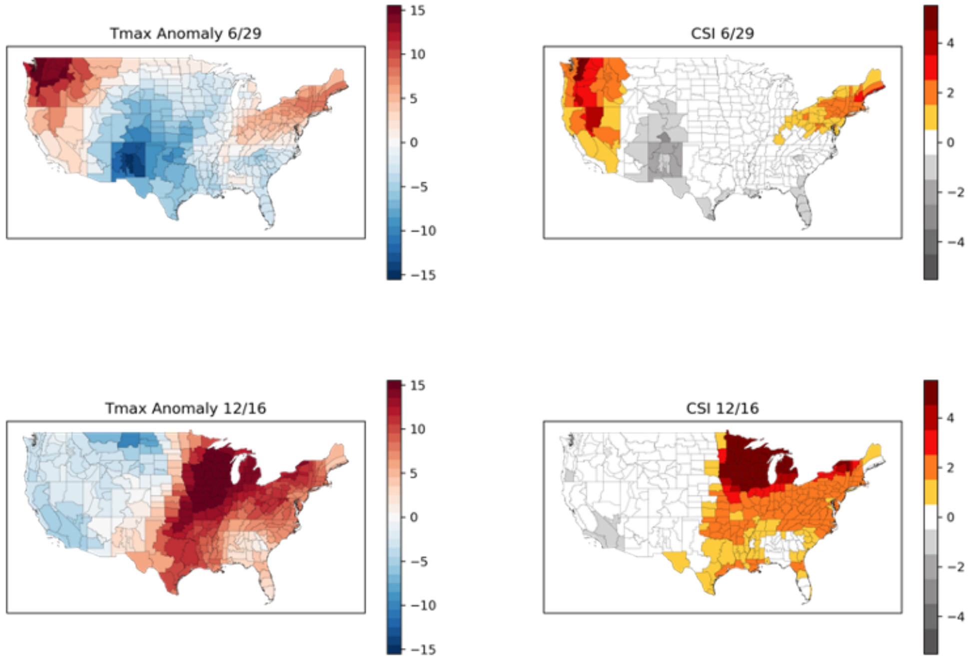

Figure 2. CC-TAS output for June 29, 2021 (top row) and December 16, 2021 (bottom row). Left. CFS T_{max} expressed as an anomaly from the 1991–2020 normal (anomaly is in °C). Right. The Climate Shift Index value for that day.

The Pacific Northwest heat dome is clearly identifiable and is quantified in our system with Climate Shift Index levels above 2 throughout Oregon and Washington and a band of level 4 and 5 in central Washington. Philip et al. (2021) estimated that the PR for this event was 175, a value much larger than what is indicated by the CSI. While our system is more conservative (by design), it clearly identified this as an important event and would have been able to do so days in advance. Our system also provides attribution estimates across the U.S.. For example, the CSI identifies the cool temperatures in New Mexico and Colorado as conditions that are less likely due to climate change. It also identifies an event in the Northeast with a significant climate fingerprint ($CSI > 1$).

In many ways, the December 16 event was even more striking than the Pacific Northwest event. A large region that encompassed all of Minnesota and Wisconsin and parts of Nebraska, South Dakota, North Dakota, Iowa, and Illinois reached level 5 on the CSI.

Acknowledgments

We are wish to express our gratitude to K. Haustein, F. Otto, C. Tebaldi, and M. Wehner who generously provided feedback on this project. We also want to acknowledge the technical team at Climate Central, especially Z. Bleki, D. Dodson, S. Katsnelson, and M. Vieuille who worked closely with us to translate our ideas into an elegant operational system.