HadSST.4.1.0.0 Product User Guide

C. G. Sandford

An update to the HadSST.4.0.1.0 Product User Guide
(J. J. Kennedy, N. A. Rayner, C. P. Atkinson, R. E. Killick)

February 2025

Quick start guide

What is HadSST?

HadSST is the Met Office Hadley Centre Sea-Surface Temperature anomaly data set. HadSST.4.1.0.0 is presented as a 200-member ensemble of bias-corrected SST anomalies, where anomalies are expressed relative to the 1961-1990 average. The data have been adjusted to minimize the effect of systematic errors associated with instrumentation changes as described in Kennedy et al. (2019), and are representative of SST measured at a depth of 20 cm. The dataset is presented on an equi-rectangular 5° latitude by 5° longitude monthly grid from January 1850 to the present day.

Uncertainty information is provided for the HadSST.4.1 anomalies in a combination of formats:

- Uncertainties arising from local measurement errors and local sampling errors are presented as gridded fields
- Uncertainties in the various bias corrections are represented by the spread of the 200-member anomaly ensemble
- A total uncertainty field combining local measurement, sampling and bias uncertainties is provided as a gridded field
- Uncertainties arising from correlated measurement errors are presented as error covariance matrices

Guidance on how to use this uncertainty information - for example, to calculate the uncertainty on global or regional averages - is provided in section 2.

HadSST.4.1 is based on in-situ observations from ICOADS release 3.0.0 (1850-2014) (Freeman et al., 2017), ICOADS release 3.0.1 (2015-2021), and ship data from the Met Office GTS feed and tropical moored buoy data from IQUAM (the NOAA In situ Quality Monitor, Xu and Ignatov (2014)) since January 2022. In addition it uses drifting buoy observations "Generated using E.U. Copernicus Marine Service Information" from CMEMS.

What products are available?

The following products are available from `https://www.metoffice.gov.uk/hadobs/hadsst4/` under the Open Government Licence v3:

- Global and regional monthly and annual SST anomaly timeseries with separate uncertainty components and total uncertainties (CSV)
- Gridded median SST anomaly and median actual SST (NetCDF)
- 200-member gridded SST anomaly ensemble and actual SST ensemble (zip archives containing one NetCDF file per member)
- Error covariance matrices representing correlated errors (zip archive containing yearly zip archives, each containing one NetCDF file per month)
- Gridded uncertainty due to uncorrelated measurement and sampling errors (NetCDF)
- Gridded total uncertainty due to uncorrelated and residual bias errors (NetCDF)
- Gridded observations counts (NetCDF)

All gridded products are monthly averages and contain entries for every month from January 1850 to the present day.

How do I read the data?

Timeseries data are stored in CSV (Comma Separated Value) files which can be read by a wide variety of tools, including Excel. The gridded data are stored in NetCDF format files. NetCDF is a platform-independent, self-describing binary format and there are a number of common tools that can be used to access the data, detailed in section 1.2.

How to acknowledge and cite the dataset

We recommend that users of the data specify which version of the data set they used, when the data were downloaded, provide a link to the website and a link to the licence. For example:

> HadSST.4.1.0.0 data were obtained from `http://www.metoffice.gov.uk/hadobs/hadsst4/data` on [date downloaded] and are ©British Crown Copyright, Met Office [year of download] provided under an Open Government Licence v3 `http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3/`.

In addition, users of HadSST4 are kindly requested to cite the following specific journal article:

> Kennedy, J.J., Rayner, N.A., Atkinson, C.P., and Killick, R.E. (2019). An ensemble data set of sea surface temperature change from 1850: the Met Office Hadley Centre HadSST.4.0.0.0 data set. Journal of Geophysical Research: Atmospheres, 124. https://doi.org/10.1029/2018JD029867.

Further information and contact

For further information please read the rest of the document. Kennedy et al. (2019) is the best place to find the technical details. For any questions not covered by this guide or references herein, please contact us at caroline.sandford@metoffice.gov.uk.

Contents

1 Getting started with HadSST

1.1 What is HadSST?

HadSST is the Met Office Hadley Centre Sea-Surface Temperature anomaly data set. HadSST.4.1 is presented as a 200-member ensemble of bias-corrected SST anomalies, where anomalies are expressed relative to the 1961-1990 average. The data have been adjusted to minimize the effect of systematic errors associated with instrumentation changes as described in Kennedy et al. (2019), and are representative of SST measured at a depth of 20 cm.

HadSST.4.1 is based on in-situ observations from ICOADS release 3.0.0 (1850-2014) (Freeman et al., 2017), ICOADS release 3.0.1 (2015-2021), and ship data from the Met Office GTS feed and tropical moored buoy data from IQUAM (the NOAA In situ Quality Monitor, Xu and Ignatov (2014)) since January 2022. In addition it uses drifting buoy observations "Generated using E.U. Copernicus Marine Service Information" from CMEMS.

HadSST is a non-infilled dataset, with gaps in both spatial and temporal coverage particularly in the early parts of the record. Data are presented on an equi-rectangular 5° latitude by 5° longitude monthly grid from January 1850 to the present day. The initial gridding is done via a "super-observations" 1° pentad (5-day) grid, to ensure a more equitable weighting of observations taken at different space-time points within the month. Gridded anomaly values are calculated from all of the measurements available within a particular grid cell and month, and cells with no available observations contain a "no data" value.

The products available from `https://www.metoffice.gov.uk/hadobs/hadsst4/` under the Open Government Licence v3 are listed below in tables 1-3.

File name	Variable
HadSST.4.1.0.0_[FREQ]_[REGION].csv	Regional average SST anomaly timeseries at monthly or annual frequency. Provided for the globe, Northern Hemisphere, Southern Hemisphere and tropics (20S-20N).
HadSST.4.1.0.0_median.nc	Gridded median SST anomalies compared to a 1961-1990 climatology
HadSST.4.1.0.0_ensemble.zip	200-member ensemble of SST anomalies compared to the 1961-1990 average, accounting for uncertainties in bias corrections
HadSST.4.1.0.0_unadjusted.nc	Gridded uncorrected SST anomalies. These should be used with caution - ideally only to quantify the bias corrections that have been applied during processing.
HadSST.4.1.0.0_actuals_median.nc	Gridded median actual SSTs (calculated from ensemble median by adding climatology)
HadSST.4.1.0.0_actuals_ensemble.zip	200-member ensemble of actual SSTs (calculated from each ensemble member by adding climatology)

Table 1: SST and anomaly products

File name	Variable
HadSST.4.1.0.0_number_of_observations.nc	Number of observations contributing to each monthly 5x5° grid cell
HadSST.4.1.0.0_number_of_superobservations.nc	Number of 1x1° pentad superobservations contributing to each monthly 5x5° grid cell

Table 2: Observation counts

File name	Variable
HadSST.4.1.0.0_total_uncertainty.nc	Total uncertainty in SST anomalies accounting for uncorrelated and bias components (see section 2)
HadSST.4.1.0.0_measurement_and_sampling_uncertainty.nc	Gridded uncorrelated measurement and sampling uncertainties (see section 2)
HadSST.4.1.0.0_error_covariance.zip	Gridded error covariances representing simply-correlated errors (see section 2)

Table 3: Uncertainty components

1.2   Accessing the data

The timeseries CSV files can be read by a wide variety of tools, including Excel. The pandas package provides a mechanism for manipulating these data in python.

The majority of HadSST data are provided as NetCDF files. A list of software tools that work with NetCDF files is maintained by UCAR at `https://www.unidata.ucar.edu/software/netcdf/software.html`. Some simple tools for viewing and manipulating NetCDF files in Linux include:

- **ncdump**: provided with the NetCDF library, produces a text rendering of a NetCDF file
- **Climate Data Operators (CDO)s**: a set of command line utilities for performing operations on NetCDF files including concatenation, editing and mathematics (`https://code.mpimet.mpg.de/projects/cdo`)
- **ncview**: a program to produce graphical displays of the contents of NetCDF files

In addition, packages are available in most commonly-used scientific programming languages for reading and working with NetCDF files. For example, in python:

- **netCDF4**: this basic package provides functionality to read NetCDF files and extract metadata (`https://pypi.org/project/netCDF4/`)
- **iris**: developed by the Met Office, iris provides functionality to read, write and process files in a variety of formats including NetCDF (`https://scitools-iris.readthedocs.io/en/stable/`)

And in R:

- **ncdf4**: `https://cran.r-project.org/web/packages/ncdf4/index.html`
- **raster**: `https://cran.r-project.org/web/packages/raster/index.html`
- **rcdo**: `https://github.com/r4ecology/rcdo`
- **RNetCDF**: `https://cran.r-project.org/web/packages/RNetCDF/index.html`
- **CM SAF R tools**: `https://www.mdpi.com/2220-9964/8/3/109`

Example: reading NetCDF data in python

The following code snippet shows an example of how to read a HadSST4 NetCDF file in python using the iris package:

```
>> import iris
>> cube = iris.load_cube('HadSST.4.1.0.0_median.nc')
>> print(cube)
```

which produces something like:

```
sea_water_temperature_anomaly / (K) (time: 2088; latitude: 36; longitude: 72)
    Dimension coordinates:
        time                            x               -               -
        latitude                        -               x               -
        longitude                       -               -               x
    Attributes:
        Conventions                 'CF-1.7'
        institution                 'Met Office'
        license                     'British Crown Copyright, distributed ...'
        reference                   'Kennedy, J. J., Rayner, N. A., ...'
        source                      'In-situ observations: ICOADS.3.0.0 ...'
        title                       'Ensemble median bias-corrected ...'
        version                     'HadSST.4.1'
```

The data are accessible as a numpy masked array in "cube.data". Further information on navigating iris cubes can be found here. Further details of the NetCDF metadata are included in section 3.2.

1.3   Things to think about when using the data

Do be aware that there are gaps in the data and that there can be considerable "noise" in less-well-observed grid cells. Auxiliary products like the uncertainty estimates and the number of observations files can provide useful additional information for identifying grid cells in which the uncertainty is likely to be large.

Do use the uncertainty information (see next section). It will help you to understand the relative reliability of the data as this changes markedly over time and in different places.

Don't compare HadSST.4.1.0.0 to globally complete SST analyses without taking into account the gaps in the data. Data coverage can affect the comparability of two data sets.

Do use the files of actual SSTs if you need actual SSTs. However, be careful when calculating area averages from the actual SSTs because the changing geographical coverage has a much greater effect when averaging actual SSTs than it does when averaging SST anomalies.

Do consult the paper associated with this dataset (Kennedy et al., 2019) for more information on the gridding and bias correction methods applied.

Do read the Open Government Licence v3 when using this data.

1.4 Contact us

If you have any questions about using the dataset that are not addressed by this guide, please contact us at caroline.sandford@metoffice.gov.uk.

2 Using the uncertainty information

The uncertainty analysis is one of the more complex parts of using HadSST4. The following provides a basic guide to using the uncertainty information provided with the dataset.

The uncertainty has been broken down into three separate components, each of which arises from errors with different degrees of correlation. The separate components should be propagated individually through any calculation (such as spatial averaging) and combined at the end to get an estimate of the overall uncertainty.

The three components provided as gridded uncertainty fields with HadSST are as follows:

- Uncertainties arising from uncorrelated measurement and sampling errors are provided as a single NetCDF file: **HadSST.4.1.0.0_measurement_and_sampling_uncertainty.nc**. In most cases these uncertainties can be propagated analytically.

- Uncertainties arising from simply-correlated measurement errors. These are provided as error covariance matrices representing systematic errors such as individual ship biases, which cannot be assumed to be independent across space and time. Covariances matrices are provided for each month, in yearly zip archives in **HadSST.4.1.0.0_error_covariance.zip**. In many cases the uncertainty for this kind of error can be propagated analytically.

- Uncertainties arising from bias correction uncertainty are represented by the spread of 200-member HadSST anomaly ensemble: **HadSST.4.1.0.0_ensemble.zip**. The ensemble represents non-Gaussian errors with complex correlation structures. In most cases, the uncertainty associated with this kind of error cannot be propagated analytically, which is why we use an ensemble.

2.1 Gridded uncertainties

The total uncertainty on the HadSST.4.1 ensemble median, covering local measurement and sampling uncertainties and the uncertainties associated with bias corrections, is provided as a gridded NetCDF file. This can be used to estimate confidence ranges at individual grid points. However, estimating the total uncertainty on global and regional averages requires treating the various uncertainty components separately according to their correlation structure.

2.2 Uncertainties on spatial and temporal averages

We can use the calculation of spatially-averaged timeseries as an example of how to propagate the different types of uncertainty described above.

Uncorrelated measurement and sampling errors

Uncertainties arising from uncorrelated errors are relatively easy to deal with using the standard propagation of errors formula. If the SST anomalies are being processed through a function $f(x_1, x_2, ..., x_n)$ where $x_1, x_2, ..., x_n$ are the gridded SST anomalies with uncertainties $\sigma_1, \sigma_2 ... \sigma_n$, then the uncertainty in f, σ_f is given by

$$\sigma_f^2 = \sum_{i=1}^{n} \left(\frac{\partial f}{\partial x_i} \right)^2 \sigma_i^2 \tag{1}$$

For a spatial average where each x_i is associated with an area weighting w_i this becomes simply:

$$\sigma_f^2 = \sum_{i=1}^{n} w_i^2 \sigma_i^2 \tag{2}$$

If there are multiple steps in the processing, the uncorrelated uncertainty can often be propagated through each step separately. For example, if we want to calculate an annual global average we can take the uncertainties in each of the twelve monthly global averages and combine them:

$$\sigma_{annual}^2 = \sum_{m=Jan}^{Dec} \left(\frac{1}{12} \right)^2 \sigma_m^2 \tag{3}$$

This only applies if the steps in the calculation are truly independent. If the same data values are used more than once in the calculation, then errors in different processing steps become correlated and can no longer be handled independently.

Simply-correlated measurement error

The propagation of uncertainties arising from simply-correlated errors is more complex than for uncorrelated errors. For an error covariance matrix $\mathbf{E}$ the uncertainty in f would be:

$$\sigma_f^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} \left(\frac{\partial f}{\partial x_i} \frac{\partial f}{\partial x_j} E_{ij} \right) \tag{4}$$

So for the example of a spatial average this would become:

$$\sigma_f^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} w_i w_j E_{ij} \tag{5}$$

For representation in code this is more cleanly expressed as $\sigma_f^2 = \underline{w} \mathbf{E} \underline{w}^T$, where the vector $\underline{w}$ contains the normalised area weightings for each grid cell. These kinds of matrix manipulations are straightforwardly handled using standard python libraries (e.g. numpy, scipy).

It is not straightforward to propagate simply-correlated errors through multiple processing steps. For annual timeseries in HadSST.4.1.0.0 it is assumed that the simply-correlated errors within a year are fully correlated, and completely uncorrelated with the adjacent year, in order to propagate the correlated component of timeseries error from monthly to annual.

Bias correction uncertainty

In many ways this is the simplest of the three types of uncertainty to deal with. If we have n ensemble members then we can calculate the uncertainty on f by calculating f for each ensemble member. For a regional average, therefore, we would simply calculate an average of each of the 200 HadSST.4.1 ensemble members. The standard deviation of the resulting 200-member ensemble of averages represents the uncertainty in the average.

Although the method for propagating uncertainty using the ensemble is simple, its usefulness is limited by the size of the ensemble and the distribution of the ensemble members. Although 200 members is a reasonable number for calculating a standard deviation, it may not be sufficient for calculating accurate quantiles or inferring statistical properties of extreme events. It should also be noted that the distribution of bias errors is non-Gaussian, so care should be taken with any statistical inferences. In all cases, it is sensible to examine the resulting distribution before performing calculations or inferences.

Coverage uncertainty

The process of spatial averaging generates an additional source of uncertainty relating to missing data. HadSST.4.1.0.0 is a non-infilled dataset, with gaps in both spatial and temporal coverage particularly in the early parts of the record. Coverage uncertainties on global and regional average timeseries are estimated by subsampling from a complete SST field (provided by HadISST.1.1.0.0) over a fixed historical time period, as described in Kennedy et al. (2011a). In HadSST.4.1.0.0 this coverage uncertainty is defined as the RMS difference of subsampled averages from the complete regional average, rather than the standard deviation (as used in HadSST3 and earlier versions of HadSST4). Estimates of coverage uncertainty are included in timeseries CSV products and form part of the total uncertainty estimates therein.

Combining uncertainty components

The four propagated uncertainty components of the spatially averaged SST anomaly are independent, and so can be combined into a total uncertainty as:

$$\sigma_{total} = \sqrt{\sigma_{uncorrelated}^2 + \sigma_{correlated}^2 + \sigma_{bias}^2 + \sigma_{coverage}^2} \tag{6}$$

2.3   Caveats and limitations

The uncertainty analysis for HadSST4 has some known deficiencies:

- The error covariance matrices can only be calculated where we have callsigns for ships. At some times (for example, the 1860s) there is very little call sign information and so the correlated error component will generally be underestimated.
- Currently, there is no mechanism for the correct combination of correlated errors from month to month. To do this analytically would require very large covariance matrices. There are various approximations that can be made - for example, assuming that errors are correlated within one year, but uncorrelated between years.
- The ensemble likely underestimates the uncertainty arising from residual errors in bias adjustments. In particular, it does not represent structural uncertainty arising from fundamental

choices made in the dataset construction process. We recommend that any analysis be repeated using a different SST data set. The ERSSTv5 and COBE-SST2 datasets are long-term climate data sets of a comparable maturity to HadSST4, and the new DCENT and COBE-SST3 datasets incorporate different approaches to bias estimation that may improve estimates of structural uncertainty.

- The uncertainty calculation does not include uncertainty associated with the errors in the climatology used to calculate the anomalies. These uncertainties are likely to be locally correlated in space and periodic in time.

As time goes on, new research also reveals previously unknown problems with the data. It is not always possible to incorporate this information in data sets during regular updates. Nonetheless it can be important for users to be aware of the wider literature. Some limitations which there are no current plans to address are detailed below.

Underestimation of correlated errors due to missing platform IDs

A number of studies have been published both during the HadSST4 development and since its release which have potential impacts on understanding of the data set. Carella et al. (2017) used a probabilistic algorithm to assign marine reports to ship tracks. In HadSST4 the error covariances are constructed using only those observations that have a non-missing and non-generic call sign or ID. This means the error covariances most likely underestimate the simply-correlated part of the total uncertainty, particulary during periods such as the 1860s when few reports have valid IDs in ICOADS. Assigning reports with erroneous IDs, missing IDs or generic IDs (such as SHIP or MASK) to ship tracks would increase the number of reports that can be included in the error covariances. This has not yet been done.

Underestimation of correlated errors due to deck-level correlations

Chan and Huybers (2019) and Chan et al. (2019) showed that there are errors in the data that correlate at the level of individual decks in ICOADS (decks refer to the decks of punched cards used to store and transmit the earliest digital data) or at a national level. However, there is no evidence that this affects trends in the global average.

Uncertainties in HadSST4 account for correlations at the level of individual ships and in large scale biases at a quasi-global level. Correlations at a deck or nation level would act to increase the uncertainty above what is estimated at present. However, in the period after 1950 the majority of ships use engine room measurements, and the spatially-temporally varying engine room adjustments in HadSST4 will account for much of this inter-deck or inter-nation variation. Prior to 1950, because fixed fields are used for the bucket adjustments, there is no component of the analysis that accounts for uncertainty correlating at the deck or nation level. Consequently, the simply-correlated measurement errors will be underestimated.

Deck-level biases

Chan et al. (2019) also identified a cold bias of around 0.4°C in Japanese data from specific decks (118), relative to nearby similar measurements (see reference for information on similarity groupings). This cold bias was due to truncation, rather than rounding, of decimal temperatures. Consequently, it is likely that the mean of HadSST4 is too cold by a varying amount in the affected period - roughly

from the 1930s to the 1960s - with the largest effect in the North Pacific. Further details of these impacts can be found in Chan and Huybers (2019) and Chan et al. (2019).

Truncation of measurements to the nearest integer °C should in theory introduce an absolute bias of close to 0.5°C, which is added to the deck 118 values in the COBE-SST3 dataset (Ishii et al., 2025). This is not inconsistent with the value from Chan et al. (2019), because that value is a relative offset from nearby measurements, and therefore includes some influence from the overall measurement bias common to that particular group of observations. Correcting for the 0.4°C relative offset in the DCENT dataset (Chan et al., 2024) is consistent with the method that first adjusts relative offsets, and then effectively calibrates SST data against coastal land surface air temperatures. However, the HadSST approach models absolute measurement bias, and we would therefore need to correct for the absolute error introduced by truncation, so that this is completely separable from the other components of bucket bias modelled downstream. The impact of adding a 0.5°C offset to deck 118 measurements will be explored for a future release.

3 File format details

3.1 Timeseries files

The timeseries file names follow a simple pattern: HadSST.4.1.0.0_annual_[R_ID].csv (and equivalent "monthly") where R_ID is the region identifier. Monthly and annual timeseries since January 1850 are provided for the following regions:

Region (R_ID)	Extended name	Lat-Lon extents (W, S, E, N)
GLOBE	Globe	-180, -90, 180, 90
SHEM	Southern Hemisphere	-180, -90, 180, 0
NHEM	Northern Hemisphere	-180, 0, 180, 90
TROP	Tropics	-180, -20, 180, 20

Timeseries are stored as CSV (Comma Separated Value) files which can be read by standard tools, including Excel and the python pandas package. Columns are as described in the following table:

3.2 Gridded files

Gridded products (listed in tables 1-3) are provided in NetCDF format with self-describing metadata, which can be read using tools such as those described in section 1.2. With the exception of the error covariances these files are 3-dimensional, containing data as time by latitude by longitude arrays with descriptive names, attributes (including "title") and units. CF standard names are used wherever possible. An example of typical NetCDF metadata for the median anomaly file is given below:

```
>> ncdump -h HadSST.4.1.0.0_median.nc
netcdf HadSST.4.1.0.0_median {
dimensions:
    time = 2088 ;
    latitude = 36 ;
    longitude = 72 ;
    bnds = 2 ;
```

Column name	Description	Units
year	Year	n/a
month	Month - in monthly file only	n/a
anomaly	SST anomaly relative to 1961-1990	K
total_uncertainty	1σ uncertainty combining all sources of uncertainty: uncorrelated measurement error, sampling error, correlated measurement error, bias correction and coverage uncertainties	K
uncorrelated_uncertainty	1σ uncertainty uncorrelated measurement and sampling errors	K
correlated_uncertainty	1σ uncertainty arising from correlated measurement errors	K
bias_uncertainty	1σ residual uncertainty in bias correction	K
coverage_uncertainty	1σ uncertainty arising from incomplete spatio-temporal coverage (see section 2)	K
lower_bound_95pct_bias_uncertainty_range	2.5% limit of the ensemble spread (minimum likely SST anomaly accounting for bias uncertainty only)	K
upper_bound_95pct_bias_uncertainty_range	97.5% limit of the ensemble spread (maximum likely SST anomaly accounting for bias uncertainty only)	K

```
variables:
    float tos(time, latitude, longitude) ;
        tos:standard_name = "sea_water_temperature_anomaly" ;
        tos:long_name = "Sea water temperature anomaly at a depth of 20cm" ;
        tos:units = "K" ;
    int time(time) ;
        time:axis = "T" ;
        time:bounds = "time_bnds" ;
        time:units = "days since 1850-01-01T00:00:00Z" ;
        time:standard_name = "time" ;
        time:long_name = "time" ;
        time:calendar = "standard" ;
    int time_bnds(time, bnds) ;
    double latitude(latitude) ;
        latitude:axis = "Y" ;
        latitude:bounds = "latitude_bnds" ;
        latitude:units = "degrees_north" ;
        latitude:standard_name = "latitude" ;
        latitude:long_name = "latitude" ;
    double latitude_bnds(latitude, bnds) ;
    double longitude(longitude) ;
        longitude:axis = "X" ;
        longitude:bounds = "longitude_bnds" ;
        longitude:units = "degrees_east" ;
        longitude:standard_name = "longitude" ;
        longitude:long_name = "longitude" ;
    double longitude_bnds(longitude, bnds) ;

// global attributes:
```

```
    ...
}
```

Data arrays are stored as 32-bit floats, and coordinate metadata are the same for all files. Ensemble zip files (both anomalies and actuals) contain one NetCDF file per member, each of which conforms to the description above.

Error covariance matrices

Due to their size, error covariance matrices are stored as monthly NetCDF files. These files are packaged first into yearly zip archives, which are then further compressed into the single error covariance zip file available to download.

Each monthly error covariance NetCDF file has the following structure:

```
>> ncdump -h HadSST.4.1_covariance_new_gridder_202201.nc
netcdf HadSST.4.1_covariance_new_gridder_202201 {
dimensions:
    location_index_1 = 2592 ;
    location_index_2 = 2592 ;
    bnds = 2 ;
variables:
    float tos_cov(location_index_1, location_index_2) ;
        tos_cov:standard_name = "sea_water_temperature_anomaly standard_error" ;
        tos_cov:long_name = "Error covariance for sea water temperature anomaly
            at a depth of 20cm" ;
        tos_cov:units = "K2" ;
        tos_cov:coordinates = "latitude_vector_1 latitude_vector_2
            longitude_vector_1 longitude_vector_2 time" ;
    int64 location_index_1(location_index_1) ;
        location_index_1:long_name = "location_index_1" ;
    int64 location_index_2(location_index_2) ;
        location_index_2:long_name = "location_index_2" ;
    double latitude_vector_1(location_index_1) ;
        latitude_vector_1:bounds = "latitude_vector_1_bnds" ;
        latitude_vector_1:units = "degrees_north" ;
        latitude_vector_1:standard_name = "latitude" ;
        latitude_vector_1:long_name = "latitude_vector_1" ;
    double latitude_vector_1_bnds(location_index_1, bnds) ;
    double latitude_vector_2(location_index_2) ;
        latitude_vector_2:bounds = "latitude_vector_2_bnds" ;
        latitude_vector_2:units = "degrees_north" ;
        latitude_vector_2:standard_name = "latitude" ;
        latitude_vector_2:long_name = "latitude_vector_2" ;
    double latitude_vector_2_bnds(location_index_2, bnds) ;
    double longitude_vector_1(location_index_1) ;
        longitude_vector_1:bounds = "longitude_vector_1_bnds" ;
        longitude_vector_1:units = "degrees_east" ;
```

```
        longitude_vector_1:standard_name = "longitude" ;
        longitude_vector_1:long_name = "longitude_vector_1" ;
    double longitude_vector_1_bnds(location_index_1, bnds) ;
    double longitude_vector_2(location_index_2) ;
        longitude_vector_2:bounds = "longitude_vector_2_bnds" ;
        longitude_vector_2:units = "degrees_east" ;
        longitude_vector_2:standard_name = "longitude" ;
        longitude_vector_2:long_name = "longitude_vector_2" ;
    double longitude_vector_2_bnds(location_index_2, bnds) ;
    int time ;
        time:bounds = "time_bnds" ;
        time:units = "days since 1850-01-01T00:00:00Z" ;
        time:standard_name = "time" ;
        time:long_name = "time" ;
        time:calendar = "standard" ;
    int time_bnds(bnds) ;

// global attributes:
    ...
```

The two top level dimensions "location_index_1" and "location_index_2" are numerical indices for a flattened latitude-longitude array. The coordinates "latitude_vector_1" and "longitude_vector_1" describe the latitude and longitude coordinates of each point along the "location_index_1" dimension, and equivalently for dimension 2. The "tos_cov" variable array describes the error covariances between each pair of locations for this particular month.

4 Frequently asked questions

4.1 Is HadSST4 the dataset for me?

There are a number of different SST data sets available covering different periods and intended for a variety of purposes. Table 4 compares properties of HadSST.4.1.0.0 with other gridded datasets that cover approximately the same period (mid-1800s to present). In general, we advise looking at more than one of these datasets to explore the structural uncertainty inherent in dataset construction.

In addition to these datasets, for applications where a long, homogeneous SST series is not required, there are a wide range of more recent satellite-based datasets available covering all or part of the period from the early 1980s to the present day.

4.2 What anomaly period have you used?

Anomalies are calculated as temperature differences from the 1961-1990 period. Note that in the calculation of the anomalies, an estimate of the climatological bias is also calculated and removed (see Kennedy et al. (2019) for details). Combining the anomalies with an SST climatology will typically lead to biased SSTs unless the climatology has been specifically bias-adjusted. A set of bias-adjusted actual SSTs, which accounts for biases in the climatology, is available as both an ensemble and an ensemble mean.

Name	Start	Resolution	Interpolated?	Bias-adjusted?	Updating?	Ensemble?	Uncertainty?	Reference
HadSST4	1850	5° monthly		Y	Y	Y	Y	Kennedy et al. (2019)
HadISST1	1870	1° monthly	Y	Y	Y			Rayner et al. (2003)
DCENT	1850	5° monthly	Y	Y		Y	Y	Chan et al. (2024)
ERSSTv5	1854	2° monthly	Y	Y	Y	Y	Y	Huang et al. (2017)
COBE-SST3	1880	0.25° daily	Y	Y			Y	Ishii et al. (2025)
COBE-SST2	1850	1° daily	Y	Y				Hirahara et al. (2014)
COBE-SST2	1850	1° monthly	Y	Y	Y			Hirahara et al. (2014)
ICOADS	1800	2° monthly			Y			Freeman et al. (2017)

Table 4: Gridded long term SST dataset characteristics

4.3 Why do you use anomalies and not actual SSTs?

The range of actual SSTs is quite large, typically ranging from around -2 to +33°C. In contrast the range of anomalies (at the spatial scales repesented by HadSST) is much smaller, typically in the range ± 5°C (although the upper limit is increasing as the climate warms). There are a number of reasons why it is beneficial to work with this more constrained data range:

1. Because HadSST is a non-infilled dataset, the uncertainty on regional averages has a "coverage" component relating to grid cells with no available data. This coverage uncertainty is smaller when there is a narrower range of possible values in the missing data grid boxes.

2. Systematic gaps in observational coverage such as those at high latitudes, where seasonal and permanent sea ice make measurement difficult, can lead to large biases in averages of actual SSTs. Although systematic gaps in coverage can also affect averages based on SST anomalies, the effects are typically much smaller.

3. Converting from actual SSTs to anomalies during quality control can help reveal problematic outliers. The largest difference between two nearby, good-quality observations is often due to the climatological average difference in temperature at those two points. By accounting for the climatological difference between the spatial and temporal locations of measurements when comparing nearby observations - i.e. by using anomalies - it is often possible to identify more subtle errors in the data.

In HadSST4, in-situ observations are converted to monthly gridded anomalies via a pentad (5-day) grid and climatology, allowing for changes in actual temperature over the course of the month to be accounted for.

For the reasons described above, when observations are averaged onto a regular grid, anomalies typically yield smaller uncertainties, unless methods are used that account for the gradients of SST across each grid cell. For this reason we recommend using anomalies for most purposes. However, we do provide an ensemble and median actual SST estimate alongside the primary SST anomaly products from HadSST4. These are constructed by adding the monthly grid-box average anomaly to the monthly climatology.

4.4 Why do the anomalies not average to zero over the climatology period?

The climatology used to calculate anomalies is spatially complete and represents the full annual cycle. This is achieved by combining observations with interpolation as described in Rayner et al. (2006). In some places coverage of observations during the 1961-1990 climatology period is very sparse, so the climatology is estimated from observations in a wider region. As a result, we do not expect the available observations expressed as anomalies to average to exactly zero over the period 1961-1990.

4.5 Why are there 200 different datasets?

The 200 datasets are known as the "ensemble". The HadSST dataset presents 200 different possible versions of the historical SST record, where differences between these versions ("ensemble members") represent a realistic range of uncertainties in the bias adjustment scheme.

The errors associated with residual biases have complex long-term correlations, and a multi-member ensemble is the only practical way to express them. To generate the ensemble, parameters in the bias adjustment scheme are varied within their plausible ranges to generate a selection of 200 different bias adjustments, which are each then used to produce a member of the bias-corrected ensemble. Further details of these parameters and their ranges are given in Kennedy et al. (2019).

4.6 What do the different parts of the version number mean?

The four components of the version number in HadSST.W.X.Y.Z are:

- W represents a **major** update which would usually be accompanied by a peer-reviewed publication. For example, the change from HadSST.3.1.1.0 to HadSST.4.0.0.0 involved fundamental changes in the way that the bias adjustments were calculated, as documented in Kennedy et al. (2019).

- X represents a **minor** update, which would not normally require a peer-reviewed publication, but does affect much of the record. For example, the change from HadSST.4.0.1.0 to HadSST.4.1.0.0 involved re-coding the system in Python, including some bug fixes (detailed in release notes), and updated the random seeds used to generate members of the bias ensemble.

- Y represents a **very minor** update which would not require a peer-reviewed publication. This would be the kind of change that affected a few months of data, or made a negligible change to the whole record. For example, the change from HadSST.3.1.0.0 to HadSST.3.1.1.0 reduced the size of the smallest component of the estimated uncertainties for seven years at the end of the record.

- Z is not currently used, but is kept for consistency with other Met Office Hadley Centre products.

Any update to the dataset would normally require users to redownload the complete dataset to get all of the changes. A major change would also mean that users might want to reconsider the details of how they use the data, in light of any changes to processing or expected uncertainties.

4.7 How does HadSST4 differ from HadSST3 and HadISST?

Key differences between different generations of HadSST and HadISST are summarised in table 5. Note that it is recommended always to use the most recent HadSST version available.

	HadSST4	HadSST3	HadSST2	HadISST1
Resolution	5° monthly	5° monthly	5° monthly	1° monthly
Time span	1850-present	1850-2018	1850-2014	1870-present
Source observations	ICOADS.3.0, CMEMS etc	ICOADS.2.5	ICOADS.2.1	Met Office Marine Database
Uses satellite data	For covariance estimation only	No	No	Yes
Spatially complete?	No	No	No	Yes
Bias adjustments applied	1850-present: space- and time-varying ERI adjustment, buoy-ship adjustment	1850-present: fixed ERI adjustment, buoy-ship adjustment	1850-1941, bucket corrections only	1870-1941, bucket corrections only
Uncertainty information	Ensemble, error covariance, uncorrelated and total uncertainty fields	Ensemble and fields	Fields	None
Reference(s)	Kennedy et al. (2019)	Kennedy et al. (2011a) and Kennedy et al. (2011b)	Rayner et al. (2006)	Rayner et al. (2003)

Table 5: Current and previous versions of HadSST and HadISST.

4.8 When can I use the total uncertainty vs the uncertainty components?

It will depend very much on your application, but in general a complete uncertainty calculation will require use of the individual components. The total uncertainty represents a combination of errors that correlate in very different ways, so there is no way to propagate the total uncertainty correctly. If you need to propagate the uncertainty, you will only be able to do this by using the individual components separately. An worked example of how to propagate uncertainty components through the calculation of a regional average is presented in section 2.

4.9 There are gaps in the data, what can I do?

You can still process data with gaps, but you should consider the impact of those gaps on the outcomes of your processing. For example:

- If you are taking a regional average, you should consider the uncertainty due to missing data ("coverage uncertainty") within the region you are interested in. Coverage uncertainty is usually estimated using a spatially complete dataset, either an infilled observations dataset or an analysis. One possible method is described in Kennedy et al. (2011a).

- If you are comparing datasets, you could consider reducing the datasets to their common coverage to allow direct comparisons to be made.

If you are using SST actuals rather than anomalies, you need to be extra careful how you deal with missing data.

As an alternative, you could choose to use a spatially infilled dataset such as one described in table 4. It is worth being aware that even with infilled products, there are still areas where the data are more uncertain, either due to lack of observations or because of limitations of the various infilling techniques. Some comparison of different SST datasets and their uncertainties is included in Kennedy (2014) and Kent et al. (2017).

5 HadSST version history

This section summarises the key differences between each successive version of HadSST published on the hadobs website. It has been taken with little modification from the HadSST.4.0.1.0 Product User Guide.

5.1 HadSST4

HadSST.4.1.0.0 (2025-present)

HadSST.4.1.0.0 is a minor release of the HadSST dataset that replaces HadSST.4.0.1.0. There have been no changes to the science in this release, which remains as documented by Kennedy et al. (2019). However, significant technical changes have been implemented including:

- Complete re-implementation of the HadSST system in Python (previously IDL)
- Addition of unit testing and stricter version control / QA procedures (GitHub)
- Move from manual to automated workflows (implemented in Cylc 8)

The product offering has also been slightly reduced, to remove additional uncertainty components that are not necessary for the propagation of uncertainties through calculations such as regional averaging.

The process of re-implementation modified some aspects of the HadSST ensemble through changes in numerical optimisation behaviour between Python and IDL, and from new random seeds used to generate perturbed parameters for the ensemble. It also highlighted some bugs in the original IDL code that have been fixed in this release.

The key differences between HadSST.4.1.0.0 and HadSST.4.0.1.0 are:

- Changes to the random seeds used to generate parameters for the ensemble at various stages of the processing
- A bug fix to the pre-1941 bucket bias corrections in which a mask had been incorrectly applied (affecting half of the ensemble members)
- A bug fix to the definition of coverage uncertainty for timeseries products (to base this on standard deviation rather than RMS difference)
- Bug fixes to estimation of the "fraction of correct metadata" affecting inferred measurement types (insulated bucket, non-insulated bucket or engine room intake) and overall bias corrections between 1945 and 1980

- Changes to spatial ancillaries affecting the spatial pattern of post-1930 bias corrections

Further details can be found in the HadSST.4.1.0.0 release notes, which are available alongside the data and this user guide. The magnitude of the resulting differences is small and falls well within the expected uncertainty range.

HadSST.4.0.1.0 (2021-2024)

Three changes were made between HadSST.4.0.0.0 and 4.0.1.0:

1. October, November and December 2020 were rerun using CMEMS data, having been originally updated using NOAA OMSC (for drifter data) in HadSST.4.0.0.0.
2. There was a bug in the code used to calculate the coverage uncertainty in the time series files in HadSST.4.0.0.0, which meant that the uncertainties were too large. A displaced mask meant that some land areas were misidentified as ice covered ocean areas. This has now been fixed.
3. A "fill_value" was added to the NetCDF files, the absence of which was causing problems for some users.

Fix to CMEMS input data (2021)

In November 2021, a problem with the drifting buoy data was identified which affected data from July 2021 to October 2021 (the latest data available at the time). Measurements from some drifting buoys, mostly in the Atlantic and Indian Oceans appeared twice in the files produced by CMEMS. Many of these observations were excluded by the HadSST4 quality control. Once the problem was identified we contacted CMEMS who rectified the problem. In early December, HadSST.4.0.1.0 was rerun from July to pick up the new files from CMEMS. **If you downloaded HadSST.4.0.1.0 between August 2021 and 10 December 2021, please redownload the files to obtain the corrected data.**

HadSST.4.0.0.0 (2019-2020)

This was a major update from HadSST.3.1.1.0, which involved changes in almost all aspects of the data set as described in Kennedy et al. (2019). Key changes included:

- The bias adjustment method was changed in almost all aspects, with particularly important changes to the way that engine room biases and measurement methods were estimated. These changes affected the whole record in some way.
- The ensemble was expanded from 100 to 200 members.
- A more complete breakdown of uncertainty components was provided to users, including error covariances that describe uncertainties associated with correlated errors.
- HadSST.4.0.0.0 was the first HadSST data set to use ICOADS release 3.0.0 and 3.0.1 for updates (HadSST3 used ICOADS 2.5).

October, November and December 2020 were updated using drifting buoy data from the NOAA OSMC service as important metadata identifying individual buoys was removed from the CMEMS drifter data files rendering them unusable. This metadata has since been reinstated and used to create HadSST.4.0.1.0 (see above).

5.2   HadSST3

HadSST.3.1.1.0

Bug fixes from HadSST.3.1.0.0 to 3.1.1.0:

- Reports from ships given the generic callsign MASKSTID, which denotes a ship whose call sign has been masked in the ICOADS Real Time feed, were not flagged as having generic callsigns in HadSST.3.1.0.0. This led to all reports with this callsign being dealt with as a single ship in the uncertainty calculation and consequently, the uncertainty was incorrect. MASKSTID was included in the list of generic call signs in HadSST.3.1.1.0. This led to a small reduction in the measurement and sampling uncertainty from 2007 to 2014 and in the total estimated uncertainty.

- A default metadata file was being read for May 2014. In this case, the default file was set to be WMO Publication 47 for May 2014, so this bug had no effect on the output. Nevertheless, it has been corrected.

HadSST.3.1.0.0

HadSST.3.1.0.0 was the first updating version of HadSST3. The random seeds used to create the ensemble were regenerated and consequently, ensemble members were different from those in HadSST.3.0.0.0.

HadSST.3.0.0.0

HadSST.3.0.0.0 (Kennedy et al., 2011a,b) was based on ICOADS release 2.5. Bias adjustments were applied to the whole data period from 1850 to 2006, and uncertainties in the bias adjustments were represented for the first time using a 100-member ensemble. The error model was rebuilt to include correlated error terms.

5.3   HadSST2

HadSST2 (Rayner et al., 2006) was based on ICOADS release 2.1. The gridding system was rebuilt to allow grids on flexible spatial and temporal resolutions. A new climatology for 1961-1990 was created. Uncertainties in the grid cell values were provided and uncertainties in the bias adjustments estimated using a Monte Carlo approach. The bias adjustments, which were applied to the period 1850-1942, were an adaptation of Folland and Parker (1995) that extended the adjustments back to 1850 and ramped them down over a three year period between 1939 and 1942. Between 1939 and 1942, new US data, available for the first time in ICOADS release 2.1, had smoothed out a previous step-change in bias at the start of the war.

5.4   HadSST

This was the first version named HadSST. HadSST was a development of the MOHSST dataset, with a variance correction applied to the gridded values to reduce spurious noise in poorly-sampled grid cells.

References

Carella, G., E. C. Kent, and D. I. Berry, 2017: A probabilistic approach to ship voyage reconstruction in ICOADS. *International Journal of Climatology*, **37**, 2233–2247, doi:https://doi.org/10.1002/joc.4492.

Chan, D., G. Gebbie, P. Huybers, and E. Kent, 2024: A Dynamically Consistent ENsemble of Temperature at the Earth surface since 1850 from the DCENT dataset . *Scientific Data*, **11 (953)**, doi:https://doi.org/10.1038/s41597-024-03742-x.

Chan, D. and P. Huybers, 2019: Systematic Differences in Bucket Sea Surface Temperature Measurements among Nations Identified Using a Linear-Mixed-Effect Method . *Journal of Climate*, **32**, 2569–2589, doi:https://doi.org/10.1175/JCLI-D-18-0562.1.

Chan, D., E. C. Kent, D. I. Berry, and P. Huybers, 2019: Correcting datasets leads to more homogeneous early twentieth-century sea surface warming. *Nature*, **571**, 393–397, doi:https://doi.org/10.1038/s41586-019-1349-2.

Folland, C. K. and D. E. Parker, 1995: Correction of instrumental biases in historical sea surface temperature data . *Quarterly Journal of the Royal Meteorological Society*, **121**, 319–367, doi:https://doi.org/10.1002/qj.49712152206.

Freeman, E., et al., 2017: ICOADS Release 3.0: a major update to the historical marine climate record. *International Journal of Climatology*, **37**, 2211–2232, doi:https://doi.org/10.1002/joc.4775.

Hirahara, S., M. Ishii, and Y. Fukuda, 2014: Centennial-scale sea surface temperature analysis and its uncertainty. *Journal of Climate*, **27 (1)**, 57–75, doi:https://doi.org/10.1175/JCLI-D-12-00837.1.

Huang, B., et al., 2017: Extended Reconstructed Sea Surface Temperature, Version 5 (ERSSTv5): Upgrades, Validations, and Intercomparisons. *Journal of Climate*, **30 (20)**, 8179–8205, doi:https://doi.org/10.1175/JCLI-D-16-0836.1.

Ishii, M., A. Nishimura, S. Yasui, and S. Hirahara, 2025: Historical High-Resolution Daily SST Analysis (COBE-SST3) with Consistency to Monthly Land Surface Air Temperature . *Meteorological Journal Vol2*, **advpub**, 2025–002, doi:https://doi.org/10.2151/jmsj.2025-002.

Kennedy, J. J., 2014: A review of uncertainty in in situ measurements and data sets of sea surface temperature. *Reviews of Geophysics*, **52**, 1–32, doi:https://doi.org/10.1002/2013RG000434.

Kennedy, J. J., N. A. Rayner, C. P. Atkinson, and R. E. Killick, 2019: An Ensemble Data Set of Sea Surface Temperature Change From 1850: The Met Office Hadley Centre HadSST.4.0.0.0 Data Set . *Journal of Geophysical Research: Atmospheres*, **124**, 7719–7763, doi:https://doi.org/10.1029/2018JD029867.

Kennedy, J. J., N. A. Rayner, R. O. Smith, D. E. Parker, and M. Saunby, 2011a: Reassessing biases and other uncertainties in sea surface temperature observations measured in situ since 1850: 1. Measurement and sampling uncertainties . *Journal of Geophysical Research: Atmospheres*, **116 (D14)**, doi:https://doi.org/10.1029/2010JD015218.

Kennedy, J. J., N. A. Rayner, R. O. Smith, D. E. Parker, and M. Saunby, 2011b: Reassessing biases and other uncertainties in sea surface temperature observations measured in situ since

1850: 2. Biases and homogenization . *Journal of Geophysical Research: Atmospheres*, **116**, doi: https://doi.org/10.1029/2010JD015220.

Kent, E. C., et al., 2017:   A Call for New Approaches to Quantifying Biases in Observations of Sea Surface Temperature . *Bulletin of the American Meteorological Society*, **98**, 1601–1616, doi:https://doi.org/10.1175/BAMS-D-15-00251.1.

Rayner, N. A., P. Brohan, D. E. Parker, C. K. Folland, J. J. Kennedy, M. Vanicek, T. J. Ansell, and S. F. B. Tett, 2006:   Improved Analyses of Changes and Uncertainties in Sea Surface Temperature Measured In Situ since the Mid-Nineteenth Century: The HadSST2 Dataset . *Journal of Climate*, **19**, 446–469, doi:https://doi.org/10.1175/JCLI3637.1.

Rayner, N. A., D. E. Parker, E. B. Horton, C. K. Folland, L. V. Alexander, D. P. Rowell, E. C. Kent, and A. Kaplan, 2003:   Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century . *Journal of Geophysical Research: Atmospheres*, **108 (D14)**, doi:https://doi.org/10.1029/2002JD002670.

Xu, F. and A. Ignatov, 2014: In situ SST Quality Monitor (iQuam). *Journal of Atmospheric and Oceanic Technology*, **31**, 164–180, doi:https://doi.org/10.1175/JTECH-D-13-00121.1.