Running eco-sharp-v1-1 in Google Earth Engine

ECOSTRESS LST Composite and 10m Sharpening

NASA ARSET Training | Step-by-Step GEE Tutorial

Link to Script

What this script does

eco-sharp-v1-1 produces two ECOSTRESS land surface temperature (LST) products for a user-selected area:

- LST 70m — a mean composite of all clear ECOSTRESS granules in the date range, in the selected temperature units
- LST 10m — the 70m composite sharpened to 10m spatial resolution using a Random Forest model trained on Sentinel-2 spectral bands and auxiliary data

Both products are exported as GeoTIFFs to Google Drive for visualization and analysis in QGIS or other GIS software.

i Prerequisites: a Google account with access to Google Earth Engine (code.earthengine.google.com) and enough Google Drive storage for your exports (typically 20–200 MB per tile).

1. Opening the Script in GEE

Google Earth Engine (GEE) is a cloud-based platform for large-scale geospatial analysis. All computation runs on Google's servers — you do not need to install any software locally beyond a web browser.

1. Open a web browser and go to **code.earthengine.google.com**
2. Sign in with your Google account. If this is your first time, you will need to register for GEE access at earthengine.google.com/signup.
3. In the GEE Code Editor, click **New** → **File** to create a new script, or paste the `eco-sharp-v1-1` code directly into the code editor panel.
4. Click **Save** (Ctrl+S / ⌘+S) to save the script to your GEE scripts library.

i The GEE Code Editor has three panels: • Left panel: Scripts, Assets, and Docs tabs — your saved scripts and uploaded data • Centre panel: Code editor — where you write and run JavaScript • Right panel: Console, Tasks, and Inspector tabs — output, exports, and pixel inspection • Map panel: Interactive map at the bottom — layers appear here after running

2. The Control Panel — UI Layout

When you click Run (►), the script loads two control panels: a left panel and a right panel. All settings and buttons are split between them to keep each panel compact.

Panel	Contents
Left panel	Title, date range, UTC hour window, timezone, cloud tolerance, additional years, selection mode, Activate/Clear/Run buttons, and status messages.
Right panel	Settings header, sharpening RF predictors (stacked vertically), output temperature units, and export options.

2.1 Left Panel — Date Range

UI Element	Description
Primary Date Range	Start and end date (YYYY-MM-DD). ECOSTRESS granules with acquisition dates within this range will be included in the composite. For a single overpass, set start and end to the same date.
Local Time Window (hour)	Filter granules by local acquisition hour in the selected timezone. Default 0–23 includes all overpasses. Narrow to e.g. 10–16 to restrict to afternoon acquisitions when urban heat is typically strongest.
Timezone	Dropdown covering 23 timezones worldwide. Defaults to America/Los_Angeles (Pacific Time). Select the timezone matching your study area so the hour window filters correctly. All major US, European, Asian, Australian, African, and South American timezones are included.
Cloud Tolerance (%)	Maximum percentage of invalid (clouded or missing) pixels allowed per granule. Granules exceeding this threshold are excluded. Default 30% is a good starting point — reduce to 10% for cleaner composites over longer date ranges.

2.2 Left Panel — Additional Summer Years

These checkboxes add Jun–Sep data from other years to build a richer multi-year composite. Useful when the primary date range has limited clear overpasses due to cloud cover or the ECOSTRESS observation schedule.

Option	What it adds
2019–2024	Ticking any year adds all ECOSTRESS acquisitions from June–September of that year over your selected area. The composite mean is computed across all selected years and the primary date range together.

i For a single-scene analysis (one specific day), leave all year checkboxes unticked. For a seasonal composite representing typical summer conditions, tick 2–3 years to increase sample size and reduce cloud gaps.

2.3 Right Panel — Sharpening RF Predictors

These checkboxes control which auxiliary datasets are used alongside Sentinel-2 imagery to train the 10m sharpening Random Forest model. Predictors are listed vertically — each has its own toggle row. Sentinel-2 reflectance bands are always included and cannot be deselected.

Predictor	Default	Description
Sentinel-2 reflectance bands (required)	Always on	Red (B4), Green (B3), Blue (B2) and Near-Infrared (B8) bands from the best available Sentinel-2 scene within ±2 months of the ECOSTRESS date. Primary drivers of spatial sharpening.
Land use / cover (ESRI LULC)	On	ESRI Global Land Use/Land Cover at 10m (2020). Helps the model distinguish heat signatures by surface type (buildings, vegetation, bare soil, water). Water pixels are excluded from RF training.
Built surface fraction (GHSL)	On	JRC Global Human Settlement Layer — fraction of built-up surface and non-residential built surface at 10m. Adds urban morphology information to help resolve heat island boundaries at fine scale.
Elevation (SRTM)	On	USGS SRTM 30m digital elevation model, resampled to 70m and 10m. Accounts for orographic temperature effects — particularly important in areas with significant terrain.
NDVI (from Sentinel-2)	Off	Normalised Difference Vegetation Index derived from Sentinel-2. Can improve sharpening accuracy in vegetated areas. Note: NDVI is already partially captured by B4 and B8 bands, so the marginal benefit is often small.

i The default predictor set (Sentinel-2 reflectance bands + LULC + GHSL + Elevation) is recommended for urban areas. The predictor combination used is encoded in all output file names so results are always traceable.

2.4 Right Panel — Output Temperature Units

A new dropdown lets you choose the temperature units for all LST outputs. This applies to both the 70m composite and the 10m sharpened product, and is encoded in the output filename.

Option	File suffix	Notes
Kelvin (K)	`_K`	Native ECOSTRESS units. Useful for scientific analysis and direct comparison with other thermal datasets.
Celsius (°C)	`_C`	Commonly used in international scientific publications and outside the US.
Fahrenheit (°F)	`_F`	Default. Useful for US-based applications, public communication, and city agency reporting.

i The sharpening accuracy metrics printed to the console (RMSE, bias, StdDev) are also reported in the selected units, so they are directly interpretable regardless of which unit you choose.

2.5 Right Panel — Export Options

Option	Default	Description
LST 70m	On	Exports the 70m mean LST composite as a GeoTIFF to Google Drive. Values are in the selected units, float32 precision, EPSG:4326.
LST 10m	On	Exports the 10m sharpened LST as a GeoTIFF to Google Drive. Same units and CRS as the 70m product. This file cannot be displayed in the GEE map due to its size — inspect it in QGIS after downloading.
RF model (asset)	Off	Saves the trained Random Forest sharpening model as a GEE asset. Useful for reusing the model on different dates without retraining. Requires the asset path field to be set correctly.
RF model asset path	Required if RF	

	model is on	

i Training samples are always exported as a GEE asset automatically — this is needed for eco-sharp-v1-1-Importance (feature importance extraction). You do not need to configure this separately.

2.6 Left Panel — Selection Mode

Choose how you define the geographic area for processing. The selection must be made before clicking Run.

Mode	How to use it
Draw custom rectangle	Click “Activate Selection” then drag a rectangle over any area on the map. Best for custom study regions or comparing specific neighbourhoods. The rectangle can cross ECOSTRESS tile boundaries.
Click to select ECOSTRESS tile	Click “Activate Selection” then click anywhere on the map. The script automatically finds the ECOSTRESS tile covering that point and outlines it in green. Best for standard single-tile analysis.
Click to select multiple ECOSTRESS tiles	Click “Activate Selection” then click multiple locations. Each click adds a tile — a running count is shown. The script combines the tiles into a single processing region. Hit Run when all tiles are selected.

i After selecting an area, the green outline shows the processing boundary. Use Clear Drawings to reset the selection and start over.

3. Running the Script — Step by Step

5. Click the ► **Run** button at the top of the code editor. The left and right control panels will appear on the map.
6. In the **left panel**: set your **date range**, **UTC hour window**, **timezone**, and **cloud tolerance**.
7. Select any additional years if building a multi-year composite.
8. In the **right panel**: choose your **predictor combination**. The default (Sentinel-2 reflectance bands + LULC + GHSL + Elevation) is recommended.

9. In the **right panel**: select your **output temperature units** (Kelvin, Celsius, or Fahrenheit). Default is Fahrenheit.
10. In the **right panel**: tick the **export options** you need. LST 70m and LST 10m are ticked by default.
11. In the **left panel**: choose a **selection mode** from the dropdown.
12. Click ■ **Activate Selection** and define your area on the map.
13. Click 🚀 **Run + Export**. The console (right panel) will show progress messages and sharpening accuracy metrics when complete.
14. When processing is complete, go to the **Tasks tab** (top-right panel).
15. Click **Run** next to each export task to confirm and start the export to Google Drive.

i Important: Clicking Run + Export registers the export tasks but does not start them automatically. You must go to the Tasks tab and click Run next to each task to begin the actual file transfer.

i The script prints progress messages to the Console tab on the right. Key messages include: granule counts after each filter step, the Sentinel-2 scene selected for sharpening, sharpening accuracy metrics (RMSE, bias, StdDev), and the list of exports registered.

4. Sharpening Accuracy Metrics

After the sharpening RF runs, eco-sharp-v1-1 automatically computes and prints accuracy metrics to the GEE console. These metrics describe how well the RF prediction reconstructs the original 70m LST composite when re-aggregated to 70m scale — the best available proxy for sharpening quality without an independent high-resolution thermal reference.

i Outlier pixels (water/shadow boundaries, RF extrapolation failures) are masked before computing stats using a unit-appropriate threshold. This gives robust metrics representative of well-observed land pixels.

Metric	Interpretation
RMSE	Root mean square error between original 70m LST and RF prediction re-aggregated to 70m. Lower is better. Typical good values are < 1°C / < 2°F.
Bias	Mean residual (original − RF prediction) at 70m. Should be close to 0 after the residual correction step. A large negative bias suggests the

	RF is overestimating LST, often caused by a temporally mismatched Sentinel-2 scene.
StdDev	Standard deviation of residuals at 70m. Captures spatial variability in RF fit quality across the tile. Large values indicate the RF fits some areas better than others.
IQR-based StdDev	Robust alternative to StdDev, computed as IQR / 1.349. If this is much smaller than the regular StdDev, there are still some outlier pixels influencing the stats.
5th–95th pct range	Range of residuals excluding the top and bottom 5% of pixels. Useful for understanding the spread of residuals across the tile.
Pixels used	Number of valid 70m pixels included in the stats after outlier masking. Very low counts (< 100) suggest most of the tile is masked — check cloud tolerance and water coverage.

5. Output Files and Naming Convention

All outputs follow a consistent naming convention encoding the date, predictors, and temperature units used, making results fully traceable.

5.1 Single-scene outputs (one granule, no extra years)

When only one granule is found for the date range, the exact UTC acquisition time is included in the filename:

```
eco-sharp-v1-1_LST_70m_2020-08-14_UTC0623_S2_LULC_GHSL_Elev_F.tif
eco-sharp-v1-1_LST_10m_2020-08-14_UTC0623_S2_LULC_GHSL_Elev_F.tif
```

5.2 Composite outputs (multiple granules or extra years)

When multiple granules or extra years are included, the date range is used instead:

```
eco-sharp-v1-1_LST_70m_2020-06-01_to_2020-08-31_S2_LULC_GHSL_Elev_F.tif
eco-sharp-v1-1_LST_10m_2020-06-01_to_2020-08-31_S2_LULC_GHSL_Elev_F.tif
```

The suffix encodes all key settings: `S2` = Sentinel-2, `LULC` = land cover, `GHSL` = built surface, `Elev` = elevation, `NDVI` = NDVI. Temperature units appended as `_K` (Kelvin), `_C` (Celsius), or `_F` (Fahrenheit).

5.3 GEE assets (always exported automatically)

Two GEE assets are always created regardless of the export checkboxes:

Asset	Purpose
`eco-sharp-v1-1_RF_model_{suffix}`	The trained RF sharpening model. Can be reloaded and applied to other dates without retraining.
`eco-sharp-v1-1_RF_samples_{suffix}`	The training sample table (~5000 pixels). Required as input for eco-sharp-v1-1-Importance to extract feature importance without reprocessing imagery.

i Finding your assets: In the GEE Code Editor, click the Assets tab in the left panel. Your assets appear under your project folder. They can take a few minutes to appear after the task completes.

6. Downloading Files from Google Drive

16. Go to **drive.google.com** and sign in with the same Google account used for GEE.
17. Navigate to **My Drive** → **GEE_ECOSTRESS**. The folder is created automatically on your first export.
18. Wait for the export tasks to complete. In GEE, the Tasks tab shows a spinning icon while the task runs and a blue checkmark when finished. Large tiles can take 5–20 minutes.
19. In Google Drive, right-click the GeoTIFF file and select **Download**. The file will download as a .tif or .zip depending on size.
20. Save the file to a local folder alongside any other files from the same run.
21. Open QGIS and follow the QGIS Visualization Guide to load and display the data.

i File sizes: A single ECOSTRESS tile (approx. 110x110 km) exports as: LST 70m: approximately 3–8 MB. LST 10m: approximately 50–200 MB. Multi-tile or multi-year composites will be proportionally larger.

7. What to Check in the Console

The GEE Console tab (right panel) prints diagnostic messages during the run. Key things to verify:

Console message	What to check
`Granules after date+hour filter: N`	Should be >= 1. If 0, widen the date range or hour window.

`Granules after cloud filter: N`	Should be >= 1. If 0, increase cloud tolerance or add extra years.
`Sentinel-2 scenes matching MGRS tile: N`	Should be > 0. If 0, a fallback scene is used — check the Sentinel-2 date printed next.
`Sentinel-2 date: YYYY-MM-DD`	The scene used for sharpening. Should be within a few months of your ECOSTRESS date. A very different date may indicate poor cloud coverage during the primary window.
`Sharpening RF predictor bands: [...]`	Lists all bands used in the RF. Verify these match your checkbox selections in the right panel.
`RMSE: X.XX [units]`	Sharpening accuracy at 70m scale. Lower is better. See Section 4 for interpretation.
`N export task(s) registered`	Confirms the number of exports queued. Go to the Tasks tab to start them.

i If you see "No granules found" the most common causes are: (1) the date range is too narrow, (2) cloud tolerance is too low, or (3) the selected area has no ECOSTRESS coverage on those dates. Try increasing cloud tolerance to 50% or adding an extra year.

8. Feature Importance (eco-sharp-v1-1-Importance)

Link to Script

After eco-sharp-v1-1 completes and the training samples asset has finished exporting, you can run eco-sharp-v1-1-Importance to extract and export the sharpening RF feature importance as a CSV. This tells you how much each predictor contributed to the 10m sharpening.

22. Wait for the **`eco-sharp-v1-1_RF_samples_...`** task to complete in the Tasks tab (green checkmark).
23. Open eco-sharp-v1-1-Importance in the GEE Code Editor.
24. Update the three configuration lines at the top of the script:

Variable	Set it to
`SAMPLES_ASSET_ID`	The full asset path printed by eco-sharp-v1-1, e.g. projects/your-project/assets/eco-sharp-v1-1_RF_samples_...
`PREDICTOR_BANDS`	Copy the list printed to the console as "Sharpening RF predictor bands" during the eco-sharp-v1-1 run.

`OUTPUT_NAME`	A descriptive name for the CSV, e.g. LST_RF_importance_2020-08-14_S2_LULC_GHSL_Elev

25. Click ► **Run** and then go to the **Tasks tab** to start the CSV export.
26. Download the CSV from Google Drive (same `GEE_ECOSTRESS` folder).

i The exported CSV has three columns: predictor (band name), importance (mean decrease impurity), and importance_pct (percentage of total). This can be plotted directly in MATLAB or Python to produce a ranked bar chart of feature contributions.