



## Applied Remote Sensing Training (ARSET) Program

### Introduction to Thermal Remote Sensing and Applications in Urban Heat Island Mapping

## Homework Questions

### Question 1

A scientist is designing a new spaceborne instrument to map land surface temperature over urban areas at night. She needs to select a wavelength range that captures purely emitted surface radiation with no interference from reflected solar energy.

Which part of the electromagnetic spectrum should the sensor target, and which law best explains why Earth-surface objects emit most of their energy in that region?

Answers: (bold correct)

- 0.4–0.7  $\mu\text{m}$  (visible), because Wien's Displacement Law predicts peak emission in the visible for objects at 300 K
- 0.7–2.5  $\mu\text{m}$  (near- and shortwave infrared), because this range avoids all atmospheric absorption bands
- 8–14  $\mu\text{m}$  (thermal infrared), because Wien's Displacement Law predicts that objects near 300 K emit peak radiation around 9–10  $\mu\text{m}$ , and this window contains purely emitted energy with no solar contamination**
- 3–5  $\mu\text{m}$  (mid-wave infrared), because this range captures both emitted and reflected energy equally, enabling day/night observations

Feedback (correct): Wien's Displacement Law ( $\lambda_{\text{max}} = b/T$ , where  $b = 2898 \mu\text{m}\cdot\text{K}$ ) tells us that for a surface at  $\sim 300$  K (a typical Earth temperature), peak emission occurs around 9.7  $\mu\text{m}$ , squarely within

the 8–14  $\mu\text{m}$  thermal infrared window. Crucially, this window contains purely emitted thermal radiation with no solar contamination, enabling reliable nighttime LST retrieval.

Feedback (incorrect): All other ranges rely on reflected solar energy or contain mixed signals. The visible range (0.4–0.7  $\mu\text{m}$ ) captures reflected light only and no signal at night. The 0.7–2.5  $\mu\text{m}$  range also detects reflected solar energy. The 3–5  $\mu\text{m}$  MWIR range contains a mixture of reflected solar and emitted thermal energy during daytime. Apply Wien's Law:  $\lambda_{\text{max}} = 2898 / 300 \text{ K} \approx 9.7 \mu\text{m} \rightarrow$  thermal infrared.

## Question 2

Scenario: You are conducting an Urban Heat Island (UHI) study in a city undergoing rapid redevelopment, where old industrial warehouses are being replaced by green "cool roofs" and asphalt parking lots. You have access to two datasets: MODIS MOD11 (which uses a Split-window algorithm) and ECOSTRESS (which uses the TES algorithm).

In this rapidly changing urban landscape, why does the TES algorithm provide a more robust Land Surface Temperature (LST) retrieval than the MOD11 Split-window algorithm?

Answers: (bold correct)

- TES uses a higher spatial resolution, allowing it to see individual buildings that the MOD11 product cannot resolve
- Split-window algorithms rely on static land-cover maps to "assign" emissivity values; if these maps are outdated, the resulting LST can have substantial systematic biases.**
- The TES algorithm is the only method capable of correcting for atmospheric water vapor, which is higher in humid urban corridors.
- Urban surfaces like asphalt are "Blackbodies," meaning they have an emissivity of 1.0, which only the TES algorithm can measure.

Feedback (correct): The key difference is how emissivity is handled. The Split-window (MOD11) method "cheats" by looking at a land-cover map (e.g., "This pixel is Urban") and assigning a pre-defined emissivity. As noted in the Mauna Loa example, if that assignment is wrong (e.g., assuming 0.98 when it is actually 0.88), you can see errors as large as 6–12 K. TES is more robust because it calculates the emissivity directly from the spectral data for every pixel, making it far more accurate for heterogeneous or rapidly changing cities where land-cover maps might be years out of date.

Feedback (incorrect): While spatial resolution (Option A) varies between sensors, it is a hardware difference, not an algorithmic one. The fundamental conceptual advantage of TES is that it doesn't "guess" emissivity based on a map. Remember the "rule of thumb": a mere 1.5% error in emissivity leads to a  $\sim 1$  K error in LST. In a city where asphalt, grass, and rooftops are mixed together, a static land-cover map often fails to capture the true, current emissivity of the surface.

### Question 3

A city planner wants to identify which neighborhoods remain dangerously hot overnight during summer heat waves, to prioritize placement of cooling centers and tree-planting programs. She is choosing between Sentinel-2 multispectral imagery and ECOSTRESS thermal infrared data.

Which statement best explains why ECOSTRESS data is more appropriate than Sentinel-2 for this application?

Answers: (bold correct)

- a. Sentinel-2 has coarser spatial resolution than ECOSTRESS, making it less useful for neighborhood-level analysis
- b. ECOSTRESS has more spectral bands than Sentinel-2, allowing finer discrimination of urban vegetation cover
- c. Sentinel-2 lacks an atmospheric correction algorithm, while ECOSTRESS automatically corrects for water vapor
- d. ECOSTRESS detects emitted thermal radiation from the surface and can observe at night, while Sentinel-2 detects reflected solar energy and produces no usable imagery in the dark**

Feedback (correct): The fundamental distinction is the source of the measured energy. Optical sensors like Sentinel-2 detect reflected solar radiation - when the sun is down, there is no signal. Thermal sensors like ECOSTRESS detect energy emitted by surfaces themselves based on their temperature, which they do 24 hours a day. Since the application requires nighttime observations (when heat-health risk peaks), ECOSTRESS's thermal emission capability is essential. Note also that ECOSTRESS's ISS orbit means it samples different times of day, including evening hours.

Feedback (incorrect): The key distinction is physical: optical sensors (Sentinel-2) need reflected sunlight and cannot produce useful imagery at night. Thermal sensors (ECOSTRESS) detect emitted surface radiation, operating equally day and night. For monitoring overnight heat retention, only a

thermal sensor can provide the required data. Spatial resolution is not the primary reason: Sentinel-2 (10–20 m) actually has finer resolution than ECOSTRESS (70 m).

## Question 4

Scenario: An irrigation water manager in California's Central Valley wants to identify which fields are being insufficiently irrigated during a summer drought. She knows that water-stressed crops can't cool through transpiration and therefore show elevated surface temperatures relative to well-watered neighbors.

Which thermal remote sensing product would most directly support her goal?

Answers: (bold correct)

- a. MODIS NDVI (MOD13), to track vegetation greenness as an indirect proxy for water availability
- b. Landsat multispectral imagery, to classify irrigated vs. non-irrigated land cover from reflectance data
- c. ECOSTRESS evapotranspiration (ET) or land surface temperature (LST) products, to directly quantify crop water use and identify thermally stressed fields**
- d. ASTER emissivity data, to distinguish soil mineralogy between irrigated and rain-fed parcels

Feedback (correct): LST is a key input for estimating evapotranspiration (ET) — the combined flux of water evaporation and plant transpiration. When crops are water-stressed, stomata close and transpiration drops, causing LST to rise. ECOSTRESS ET and LST products directly capture this signal and are widely used for agricultural water management and drought monitoring. ECOSTRESS's diurnal sampling capability (including peak afternoon heat) is especially valuable here.

Feedback (incorrect): NDVI and land cover maps provide useful context but do not directly measure water use or thermal stress. Emissivity from ASTER is related to surface composition, not water status. The thermal response to drought — suppressed ET and elevated LST in stressed fields relative to irrigated neighbors — is what makes thermal RS uniquely suited. As discussed in Part 1: "LST is a key input to land surface models for drought monitoring, soil moisture and ET estimation."

## Question 5

A research team wants to study how the urban heat island intensity of a large city changes hour by hour over a 24-hour period — including the peak afternoon hot spot around 3:00 PM and the residual overnight warmth after midnight. They need multi-temporal thermal observations of the same location within a single day.

Which sensor is best suited for capturing the diurnal temperature cycle needed for this study, and what unique characteristic enables this capability?

Answers: (bold correct)

- a. MODIS (Terra/Aqua), because its daily revisit provides two overpasses per day at consistent morning and afternoon times
- b. ASTER, because its 5 TIR bands and 90 m resolution enable accurate TES emissivity retrieval throughout the day
- c. Landsat 8/9, because its 16-day repeat cycle and consistent ~10:30 AM overpass establish a reproducible thermal baseline
- d. ECOSTRESS on the International Space Station, because the ISS non-sun-synchronous orbit allows observations at varying local times across different days, enabling diurnal cycle sampling including peak afternoon and nighttime temperatures**

Feedback (correct): ECOSTRESS is uniquely suited for diurnal temperature studies because the ISS does not follow a sun-synchronous orbit. Unlike polar-orbiting platforms (MODIS, ASTER, Landsat), which observe a given location at predictable, fixed local times, the ISS orbit drifts across local times — so ECOSTRESS accumulates observations at many different times of day over repeated passes. This "4-day diurnal repeat cycle" is explicitly highlighted in Part 1 as the key advantage of ECOSTRESS for urban heat island work.

Feedback (incorrect): MODIS (Terra + Aqua) provides only two fixed overpass times (~10:30 AM and ~1:30 PM Terra; ~10:30 PM and ~1:30 AM Aqua), not the flexible multi-time sampling needed here. ASTER and Landsat follow sun-synchronous orbits with a fixed ~10:30 AM overpass — always before peak afternoon heat. The defining feature of ECOSTRESS is the non-sun-synchronous ISS orbit, which is what enables diurnal cycle sampling.

## Demo R Script Exercise

Question 6 is based on Demo 1: Nighttime Cooling Patterns During the LA September 2024 Heatwave. The two nighttime acquisitions remain the same as in the original demo: Night 1 = Sep 3, 11:14 PM PDT; Night 2 = Sep 4, 10:26 PM PDT.

Before answering this question, complete the following steps:

1. **Download** the Demo 2 R script from the [Github page](#).
2. **Modify the AOI coordinates.** Replace the Chatsworth AOI objects (lines ~80–88) with the Hansen Dam study areas below (or simply uncomment that section of code):  

```
# Hansen Dam Wildlife Preserve  
aoi_preserve <- ext(-118.38833, -118.37166, 34.26197, 34.26767) %>%  
  vect(crs = wgs84_wkt)  
  
# Adjacent Commercial Area  
aoi_commercial <- ext(-118.40428, -118.39430, 34.25404, 34.25793) %>%  
  vect(crs = wgs84_wkt)  
  
# Updated Study Region (covers both AOIs)  
study_region <- ext(-118.42, -118.36, 34.25, 34.28) %>%  
  vect(crs = wgs84_wkt)
```
3. **Run the full script.** Allow all sections to complete, including QC masking and AOI extraction
4. **Locate your results.** The printed df\_summary table in your R console contains the mean nighttime LST values you need. Use the mean\_lst column, rounded to one decimal place

**Complete the Demo 2 R script setup described in the instructions above before answering.**

## Question 6

Use the printed df\_summary table to answer.

After modifying the R script to use the Hansen Dam Wildlife Preserve and adjacent commercial AOIs, what was the thermal gap (mean LST of the commercial area minus mean LST of the Hansen Dam Wildlife Preserve) on Night 2, rounded to the nearest 0.1°C?

Answers: (bold correct)

- a. 4.8°C**
- b. 4.2 °C
- c. 3.7 °C
- d. 4.3 °C

Feedback (correct): The thermal gap is calculated as:  $\text{gap} = \text{mean\_lst}(\text{Commercial Area, Night 2}) - \text{mean\_lst}(\text{Nature Preserve, Night 2})$ . A positive gap confirms that impervious commercial surfaces remain warmer overnight than the vegetated wildlife preserve - the urban heat island effect. Compare the gap on Night 1 vs. Night 2: if it persists or grows, the preserve is maintaining its cooling function even after a full day of accumulated extreme heat. In this case, we saw an \*increase\* in the temperature gap, growing from 4.5°C on Night 1 to 4.8°C on Night 2, indicating that the preserve is maintaining its cooling function and supporting the argument for urban green space as heat resilience infrastructure.

Feedback (incorrect): Check that you (1) updated all three AOI extents in the script, (2) set the correct CRS (`wgs84_wkt`) on each vector, (3) ran the QC and cloud masking steps without error messages, and (4) subtracted in the correct direction (Commercial – Preserve, not the reverse). Also confirm you are using Night 2 values, not Night 1. If your gap is negative or implausibly large, re-check that the AOI extents did not accidentally overlap or fall outside the ECOSTRESS coverage footprint for these dates.

## Question 7

Scenario: A public health analyst wants to map heat exposure at the scale of individual city blocks and parcels to identify which households lack access to cooling. Native ECOSTRESS LST pixels are 70 m across — larger than many urban lots. She uses the Random Forest downscaling workflow in Google Earth Engine (GEE) to sharpen ECOSTRESS to 10 m resolution using Sentinel-2 spectral predictor bands (NDVI, NDBI, MNDWI, and shortwave infrared bands).

Which of the following best describes why high-resolution optical indices (such as NDVI and NDBI) from Sentinel-2 are effective predictor variables for downscaling ECOSTRESS LST?

Answers: (bold correct)

- a. Sentinel-2 shortwave infrared bands measure emitted thermal radiation at 10 m, which can directly replace coarser ECOSTRESS thermal pixels
- b. NDVI and NDBI encode information about surface cover type (vegetation density, impervious surface fraction) that strongly determines how surfaces absorb and retain heat, creating a statistical relationship between spectral indices and LST that the Random Forest model can learn and apply at 10 m**

- c. The GEE Random Forest model generates new thermal observations by interpolating ECOSTRESS pixel values spatially across urban areas
- d. Sentinel-2 has a shorter revisit time than ECOSTRESS, so its bands capture more representative daily average temperatures when fused with ECOSTRESS data

Feedback (correct): Downscaling works by exploiting the strong physical relationship between surface cover (captured by optical indices at 10 m) and land surface temperature. Vegetation (high NDVI) cools via evapotranspiration; impervious surfaces (high NDBI) retain heat. A Random Forest model trained on paired ECOSTRESS LST and Sentinel-2 predictor values learns this relationship at the 70 m scale, then applies it at 10 m — producing spatially detailed LST estimates that respect the thermal physics of the landscape.

Feedback (incorrect): Sentinel-2 does not measure thermal (emitted) radiation — its bands capture reflected solar energy. The sharpened LST is a model-predicted estimate, not a direct thermal measurement. It is not spatial interpolation between ECOSTRESS pixels, which would simply smooth existing values. The revisit time of Sentinel-2 is irrelevant here. The key insight is that optical spectral indices are physical proxies for surface properties (vegetation, imperviousness, moisture) that causally control LST — enabling the statistical downscaling model to transfer thermal patterns to higher resolution.

## GEE Exercise

Questions 8 and 9 refer to this [GEE script](#).

After running `eco-sharp-v1-2` over downtown Los Angeles on a summer afternoon, you get the following console output:

Sharpening RMSE (70m, clipped): 1.8 °F

Bias: 0.1 °F

StdDev: 1.6 °F

LST display stretch — vMin: 89 vMax: 118 °F (auto)

The 10m sharpened layer shows a large commercial rooftop registering 117°F, while the surrounding street canyon reads 108°F. The 70m composite for that same pixel reads 111°F.

## Question 8

What does the RMSE of 1.8°F represent in this context?

Answers: (bold correct)

- a. The error between the 10m sharpened LST and ground-based thermometer measurements
- b. The difference between the RF-predicted LST and the actual 70m composite LST used for training, evaluated at training pixel locations**
- c. The temperature uncertainty in the original ECOSTRESS L2T product
- d. The difference between the 70m and 10m spatial averages over the tile

Feedback (correct): The reported RMSE of 1.8°F represents the model's prediction error evaluated at the native 70m scale used for training. In the Earth Engine downscaling workflow, the Random Forest model learns the statistical relationship between the predictors (the 10m Sentinel-2 bands aggregated up to 70m) and the target variable (the actual 70m ECOSTRESS LST). The RMSE quantifies how closely the model's fitted predictions match the true ECOSTRESS observations across the training pixels. It serves as our best proxy for model confidence before projecting those learned rules down to the 10m scale.

Feedback (incorrect):

- Ground-based measurements: The workflow is entirely remote-sensing based; no ground-based weather station or thermometer data is fed into the GEE script, so the RMSE cannot represent ground truth error
- ECOSTRESS L2T uncertainty: The inherent instrument or atmospheric correction uncertainty of the ECOSTRESS sensor is a separate systematic value (typically around 1–2 K globally) and is not what this machine learning validation metric is measuring.
- Spatial averages over the tile: the difference between the overall 70m and 10m global spatial averages is captured by the Bias (0.1°F), not the RMSE. RMSE is a pixel-by-pixel metric tracking the magnitude of localized error, not a global average shift.

## Question 9

The 70m pixel reads 111°F but the 10m rooftop pixel reads 117°F. Which of the following best explains this difference?

Answers: (bold correct)

- The sharpening algorithm introduced an error — the 10m value should equal the 70m value
- The 70m pixel is a spatial average that mixes the hot rooftop with cooler adjacent surfaces; the 10m pixel isolates the rooftop signature that was blurred in the coarse composite**
- The Sentinel-2 NIR band incorrectly classified the rooftop as vegetation, inflating the temperature
- The residual correction step overcorrected by subtracting the mean LST bias

Feedback (correct): This scenario illustrates the core purpose of thermal sharpening. In a dense urban environment, a native 70m ECOSTRESS pixel acts as a spatial average (thermal mixture), blending together the intense heat of a dark commercial rooftop with the cooler, shaded microclimates of adjacent street canyons and vegetation.

By leveraging 10m Sentinel-2 optical indices (like NDBI for built-up infrastructure and NDVI for vegetation), the Random Forest model successfully "unmixes" this footprint. It recognizes the specific 10m sub-pixels that belong to the highly absorbent roof structure, correctly adjusting its predicted temperature upward to 117°F, while identifying the cooler street canyon at 108°F.

Feedback (incorrect):

- Algorithm error: The algorithm did not fail; expecting the 10m pixel to exactly equal the 70m pixel defeats the entire purpose of downscaling, which is meant to resolve sub-pixel heterogeneity.
- NIR and vegetation inflation: high vegetation density (high NDVI) typically lowers predicted LST because plants cool their environment via evapotranspiration. If the model mistook the roof for vegetation, it would have erroneously decreased the predicted temperature, not inflated it to 117°F. Furthermore, rooftops display high Normalized Difference Built-Up Index (NDBI) signatures
- Residual correction step: The residual correction step distributes the broad-scale optimization bias evenly across the scene to preserve energy conservation; it does not introduce dramatic, highly localized 9°F thermal adjustments to individual buildings.