ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Computer Scientists, Your Planet Needs You: An Introduction to Access-Optimized Data

July 24, 2024

Patrick Quinn
NASA ESDIS Architect
patrick.m.quinn@nasa.gov

NASA

Access-Optimized, a wishy-washy definition

Earth observation data structured so that retrieving relevant information from a data set is fast and inexpensive, on average, across all use cases and access patterns, at real-world scale

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Access-Optimized, a wishy-washy definition

Earth observation data structured so that retrieving relevant information from a data set is fast and inexpensive, on average, across all use cases and access patterns, at real-world scale

Trade-offs exist

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Why is this a new challenge in the cloud?

Why is this a new challenge in the cloud?

It’s not.

Why is this a new challenge in the cloud?

It's not.

What's happened…

- Canonical data formats optimize continual read/write, not read-only
- Things were *harder* historically. Getting full files was the challenge. Quickly getting partial files required services.
- Once data users had the right files they had the same problems.
- Cloud storage's high latency exaggerated some performance issues.

Why is this a new challenge in the cloud?

It's not.

What's happened…

- Canonical data formats optimize continual read/write, not read-only
- Things were *harder* historically. Getting full files was the challenge. Quickly getting partial files required services.
- Once data users had the right files they had the same problems.
- Cloud storage's high latency exaggerated some performance issues.

Speculation: The cloud made it easier to encounter data optimization challenges, allowing an influx of engineers to understand and care

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Why is this a new challenge in the cloud?

It’s not.

What’s happened…

- Canonical data formats optimize continual read/write, not read-only
- Things were *harder* historically. Getting full files was the challenge. Quickly getting partial files required services.
- Once data users had the right files they had the same problems.
- Cloud storage’s high latency exaggerated some performance issues.

Speculation: The cloud made it easier to encounter data optimization challenges, allowing an influx of engineers to understand and care

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

… so let’s understand and care!

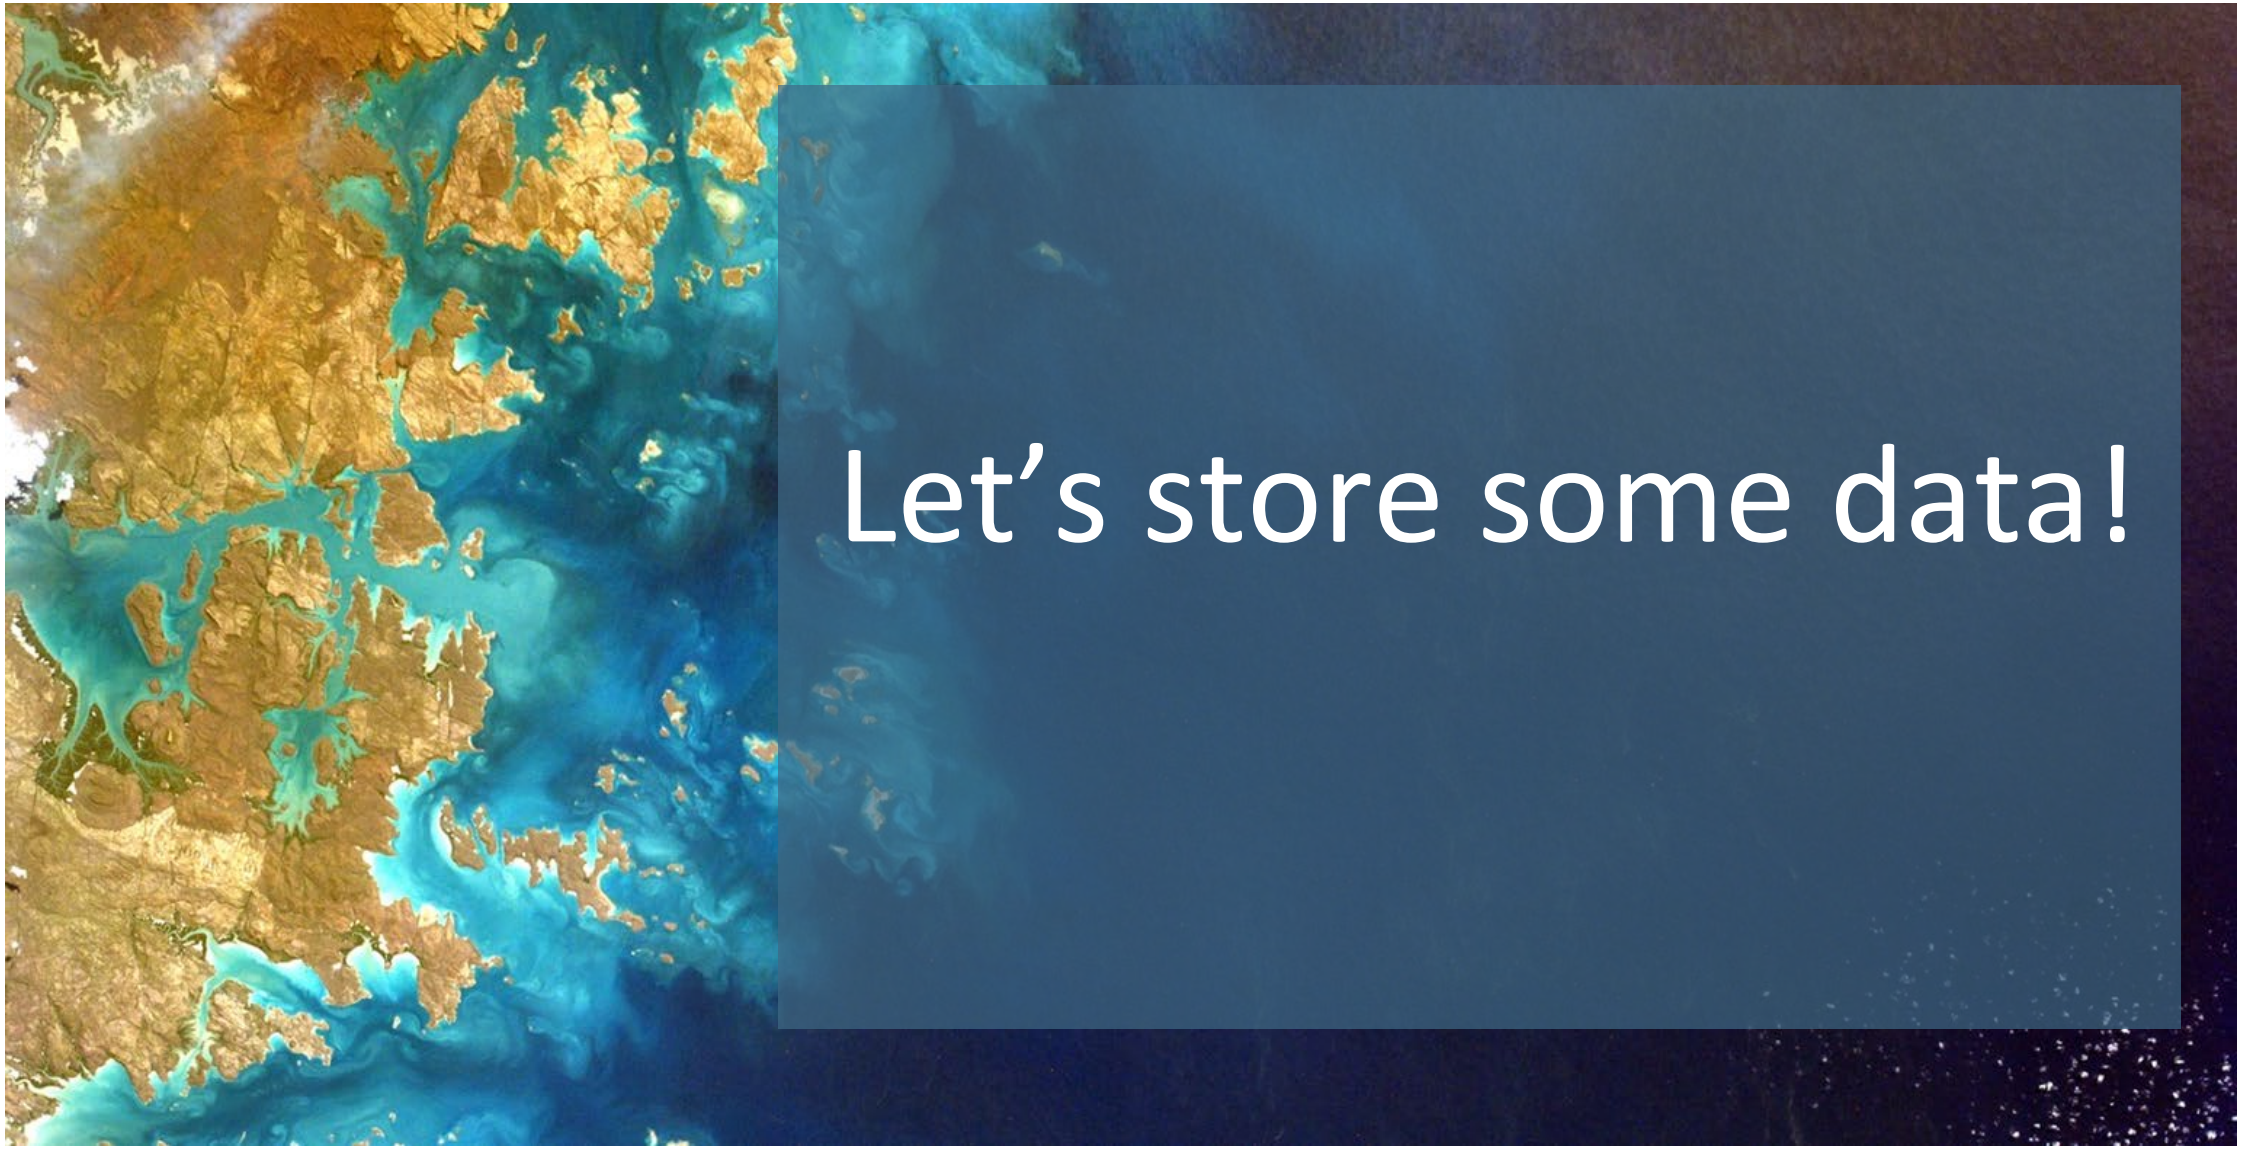
Let’s store some data!

Analysed Sea Surface Temperature

GHRSST/MUR[1] 2021-07-15

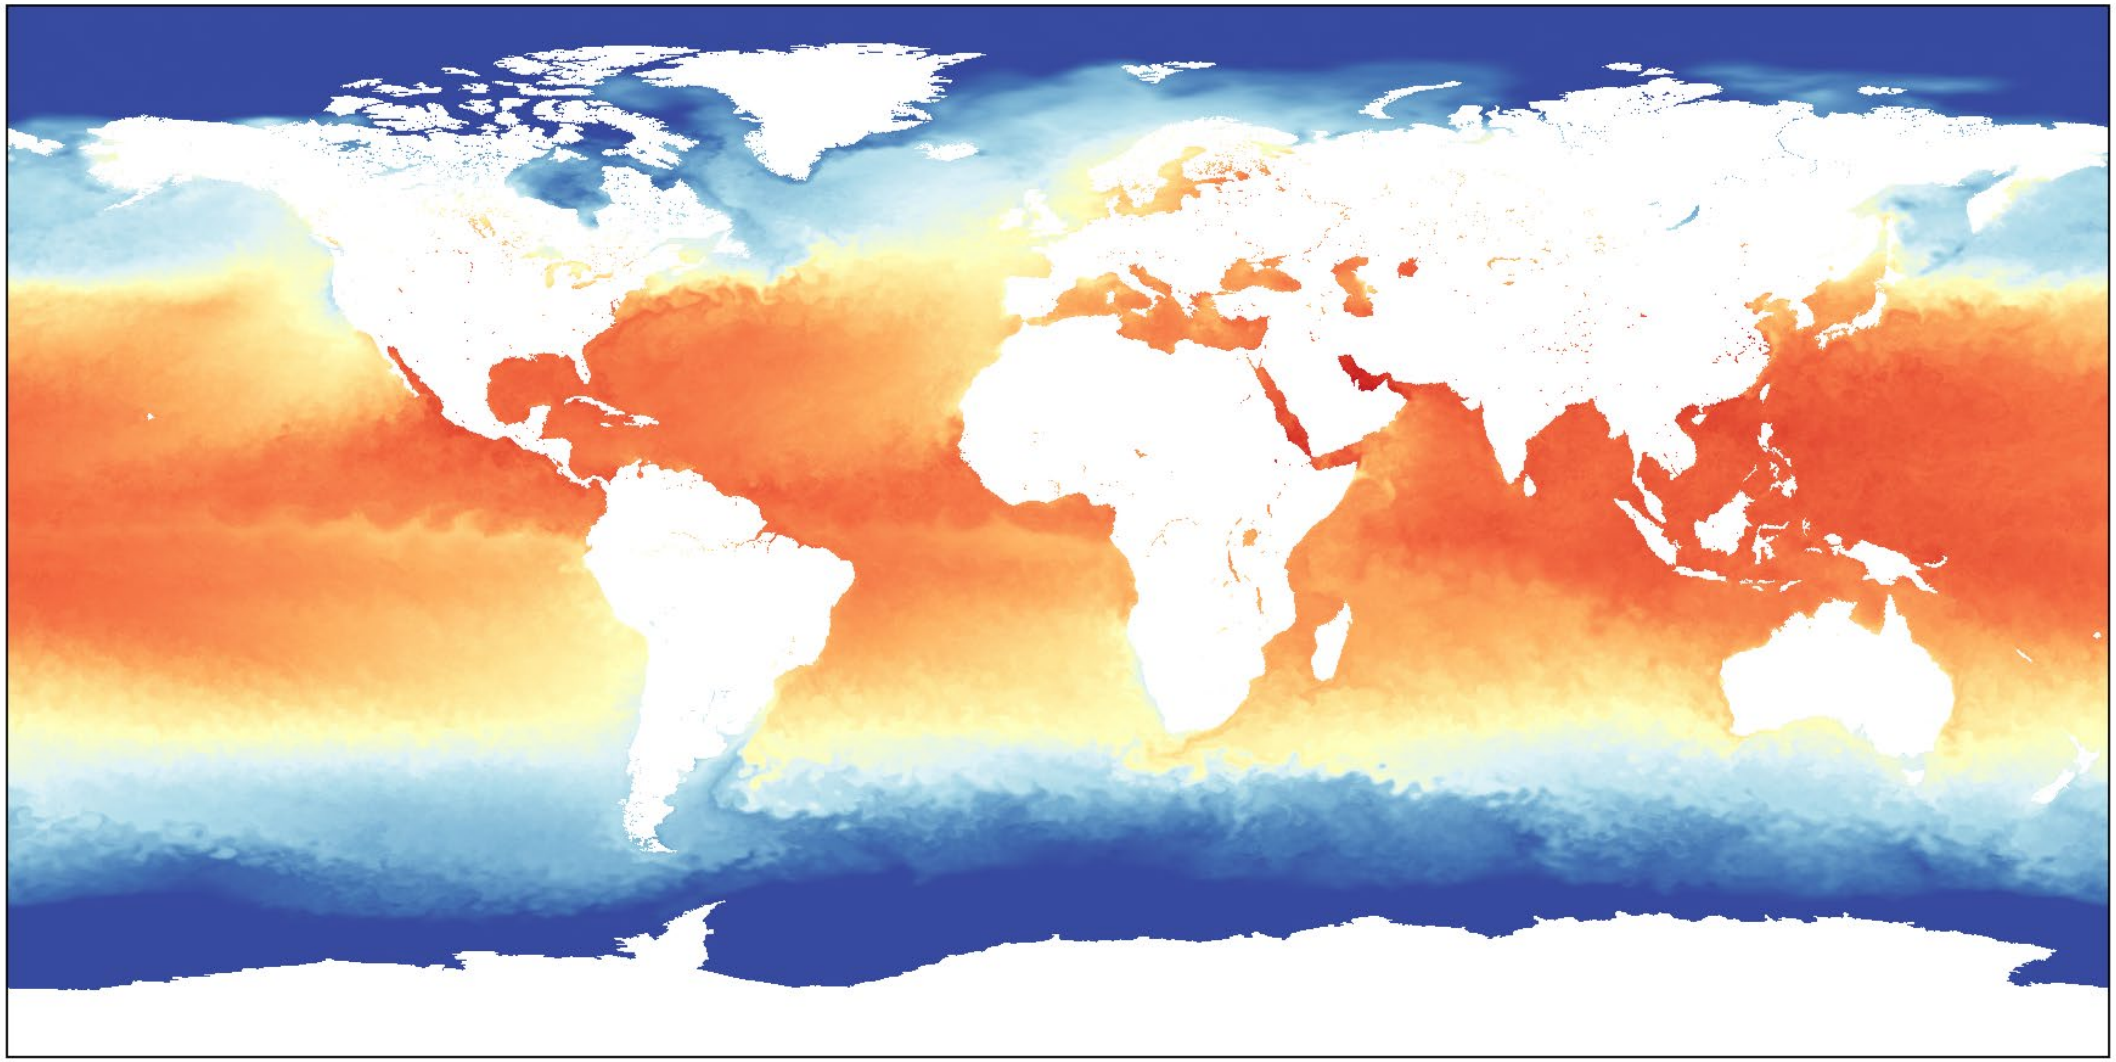

Analysed Sea Surface Temperature (K)

270 280 290 300 310

1. Group for High-Resolution Sea Surface Temperature / Ultrahigh Resolution
2. Digital Object Identifier (https://dx.doi.org)

ESDIS EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI[2] 10.5067/GHGMR-4FJ04

Dataset information

Key/value pairs

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

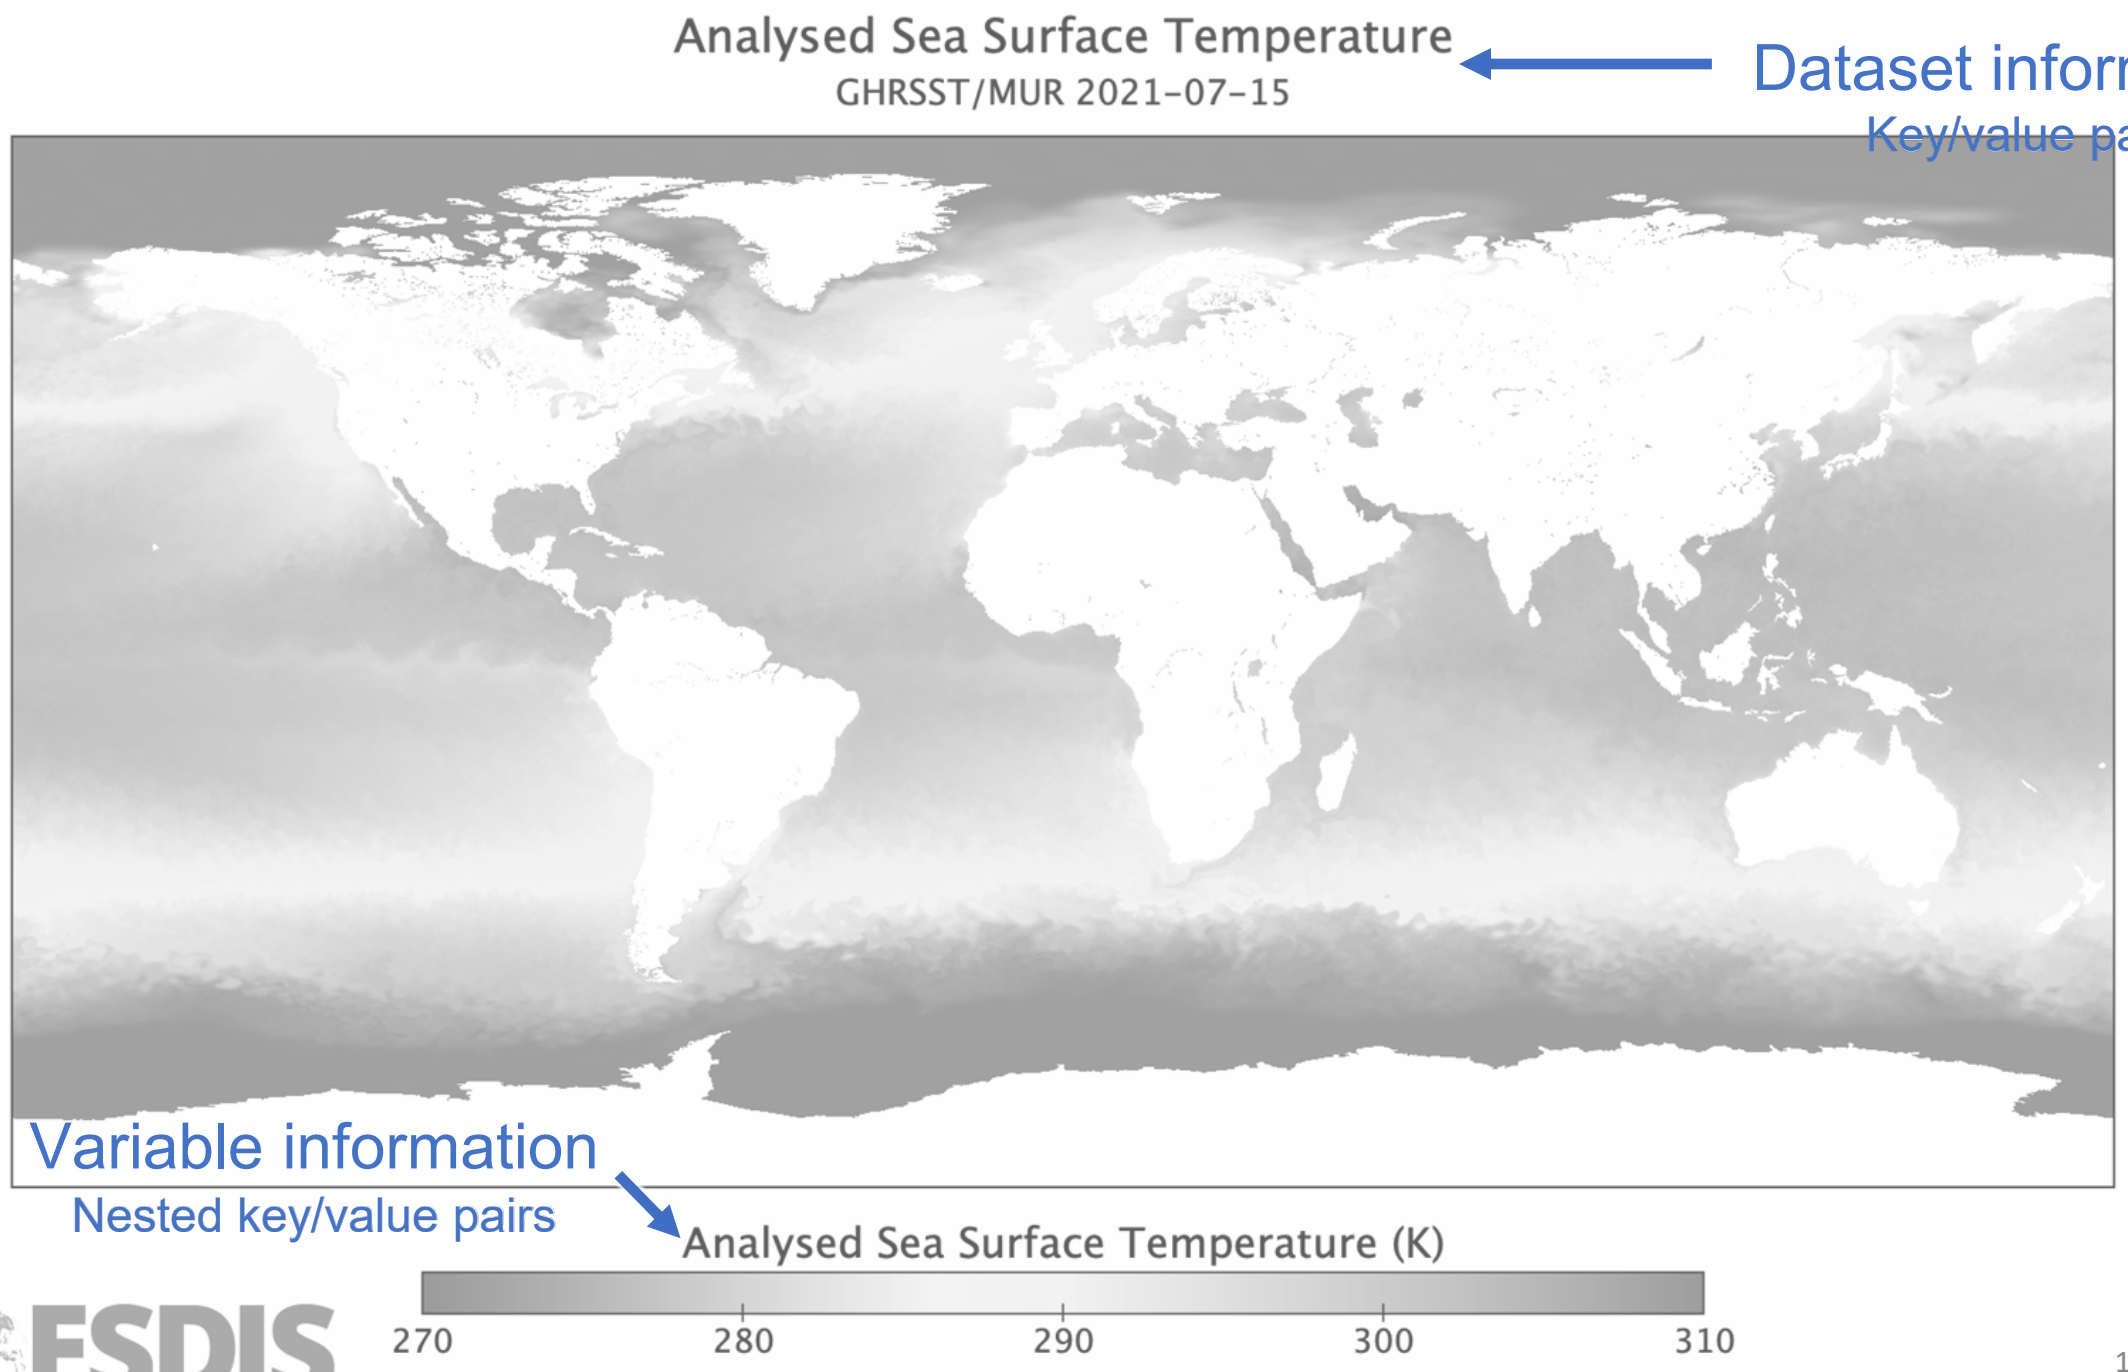

Source: DOI 10.5067/GHGMR-4FJ04

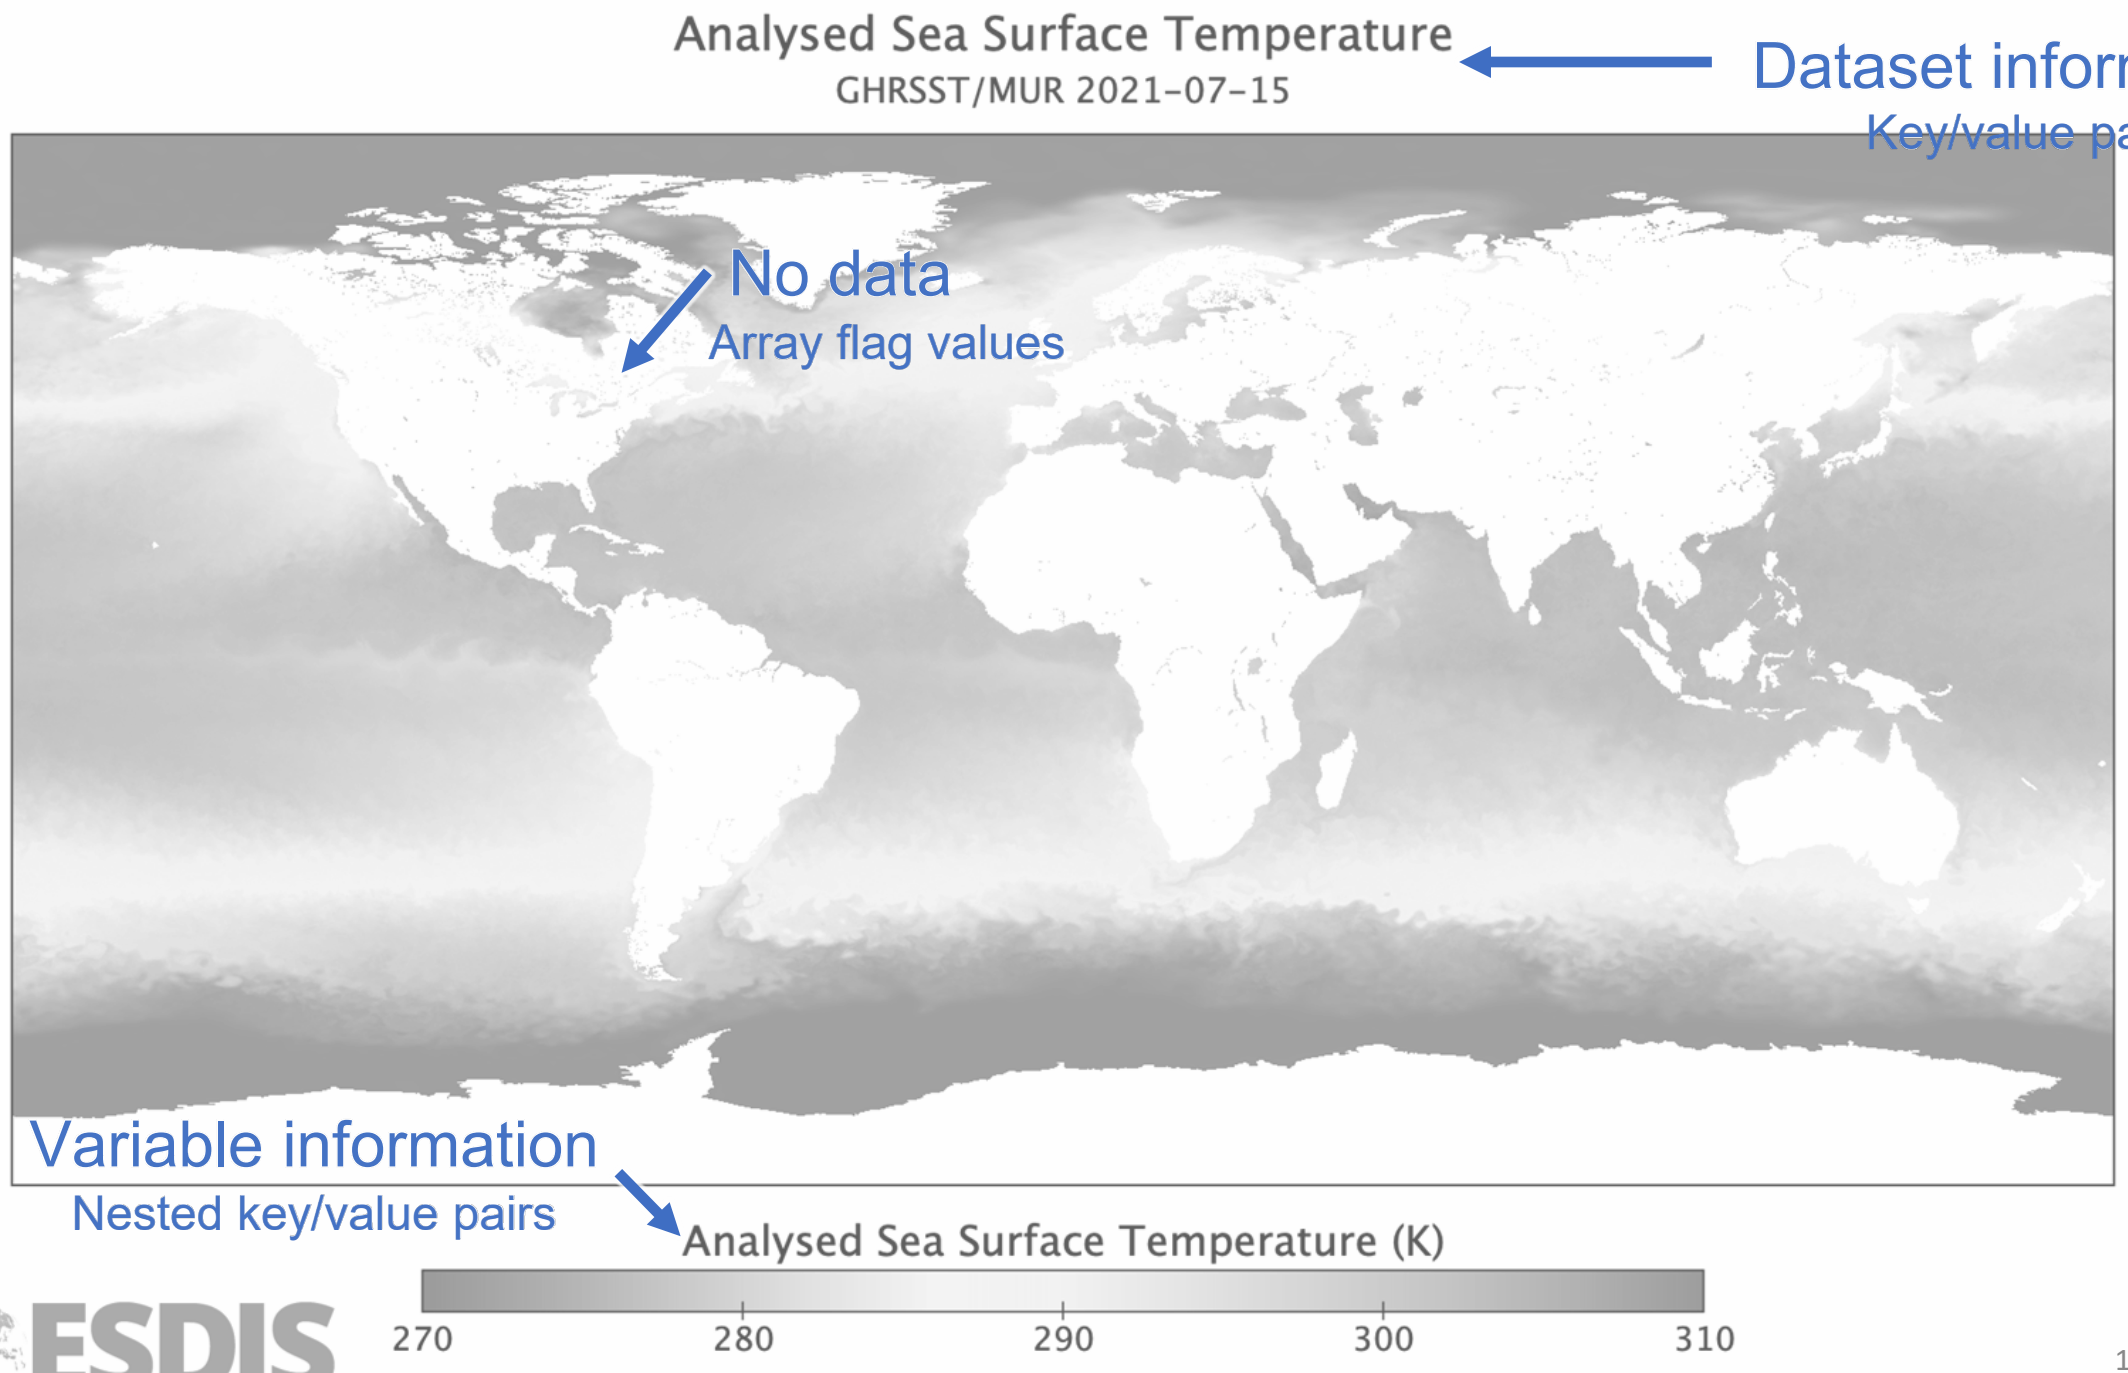

Source: DOI 10.5067/GHGMR-4FJ04

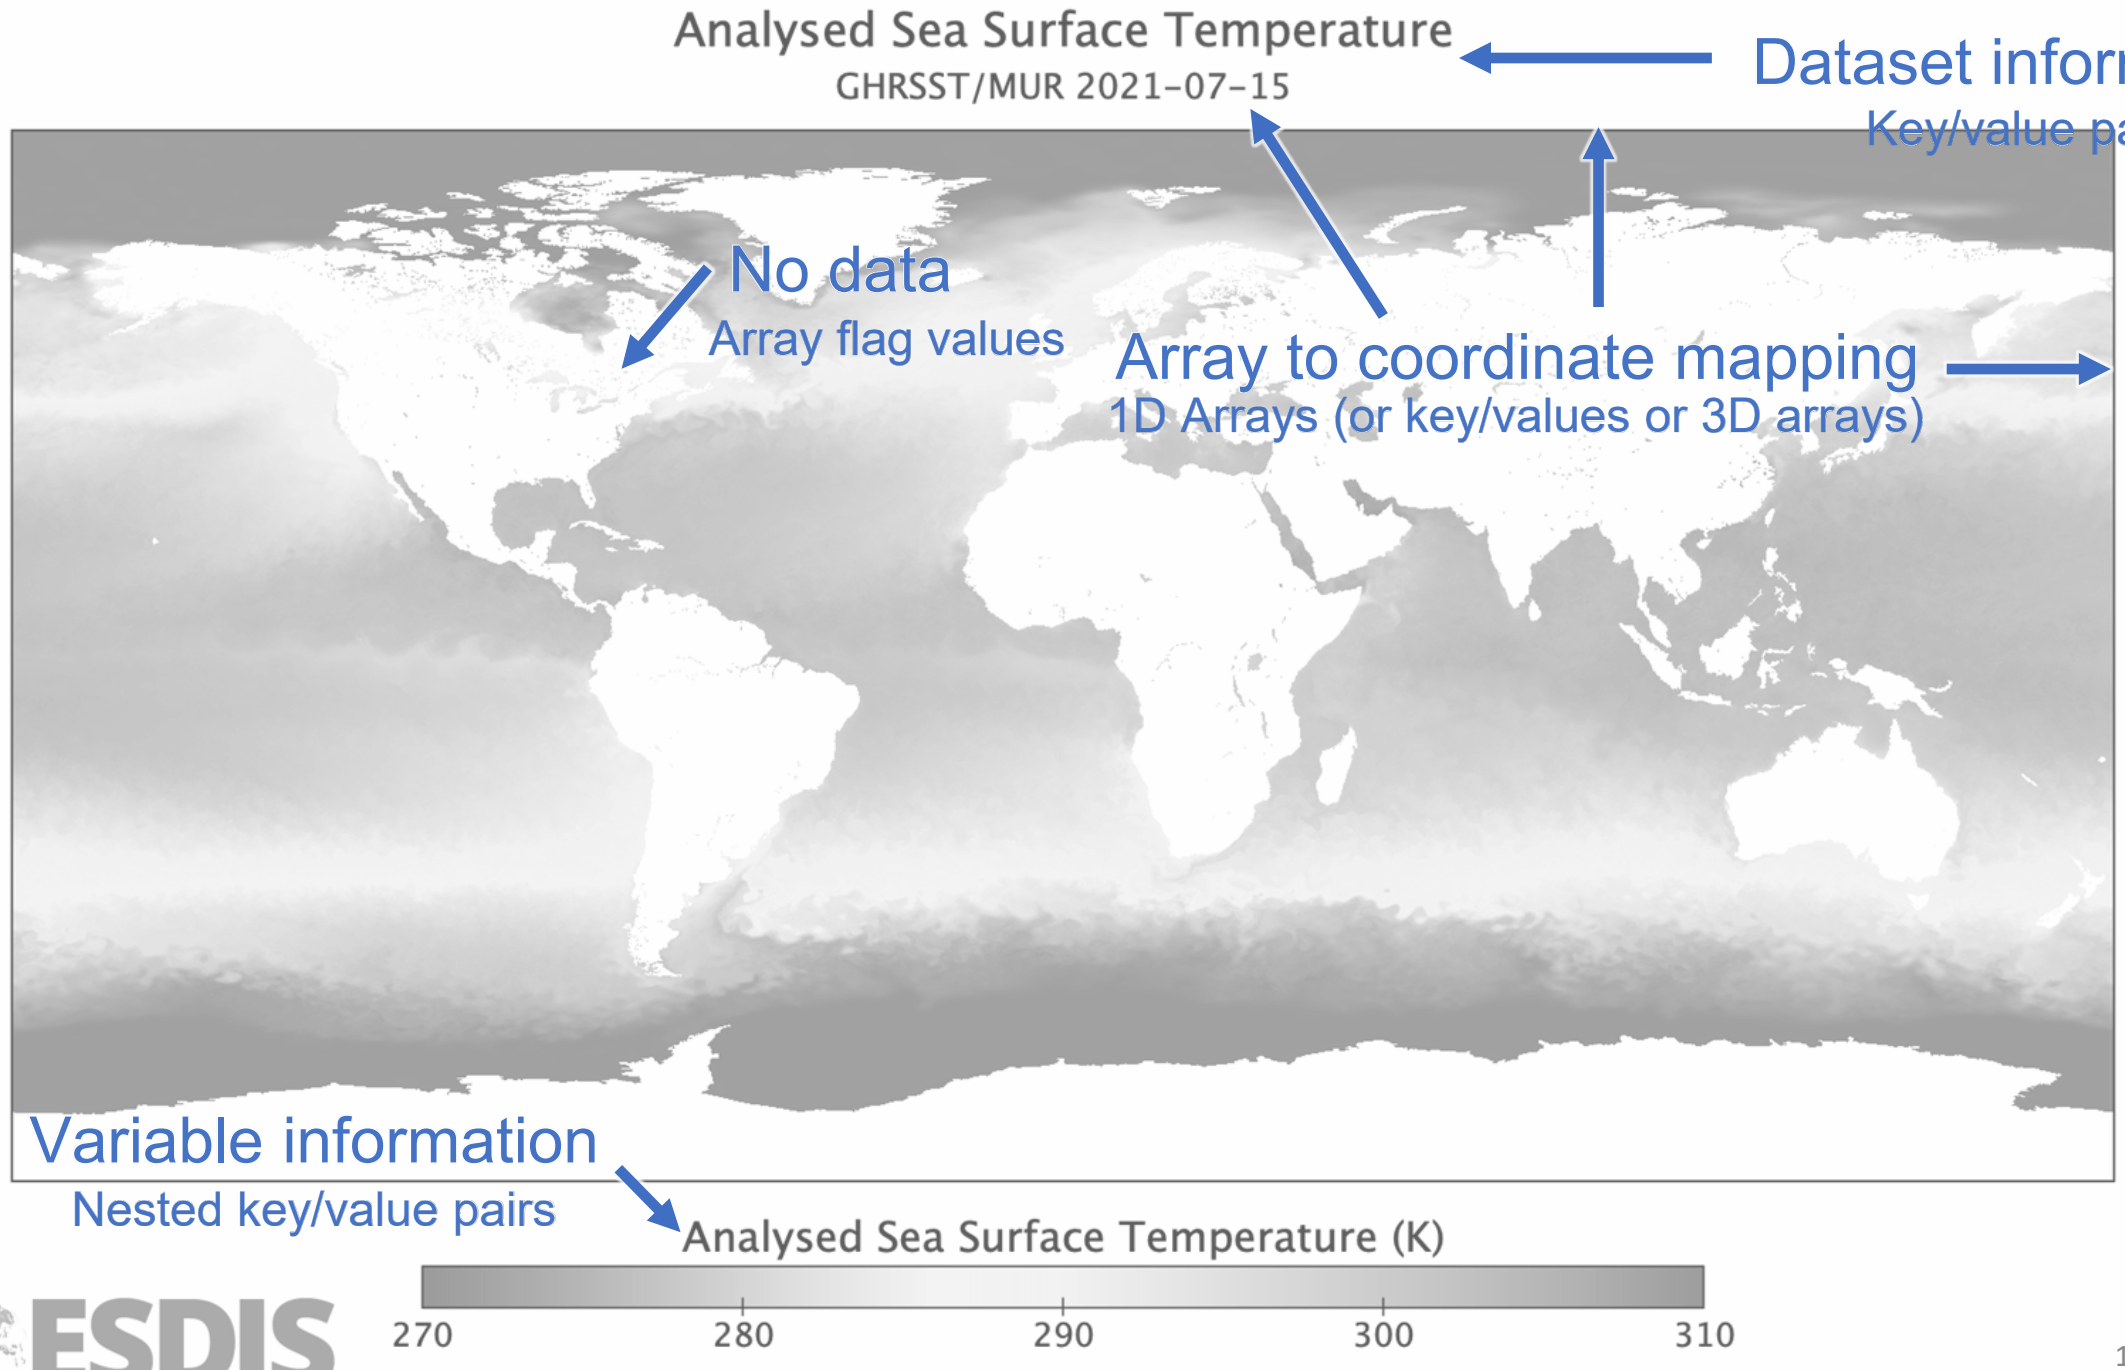

Source: DOI 10.5067/GHGMR-4FJ04

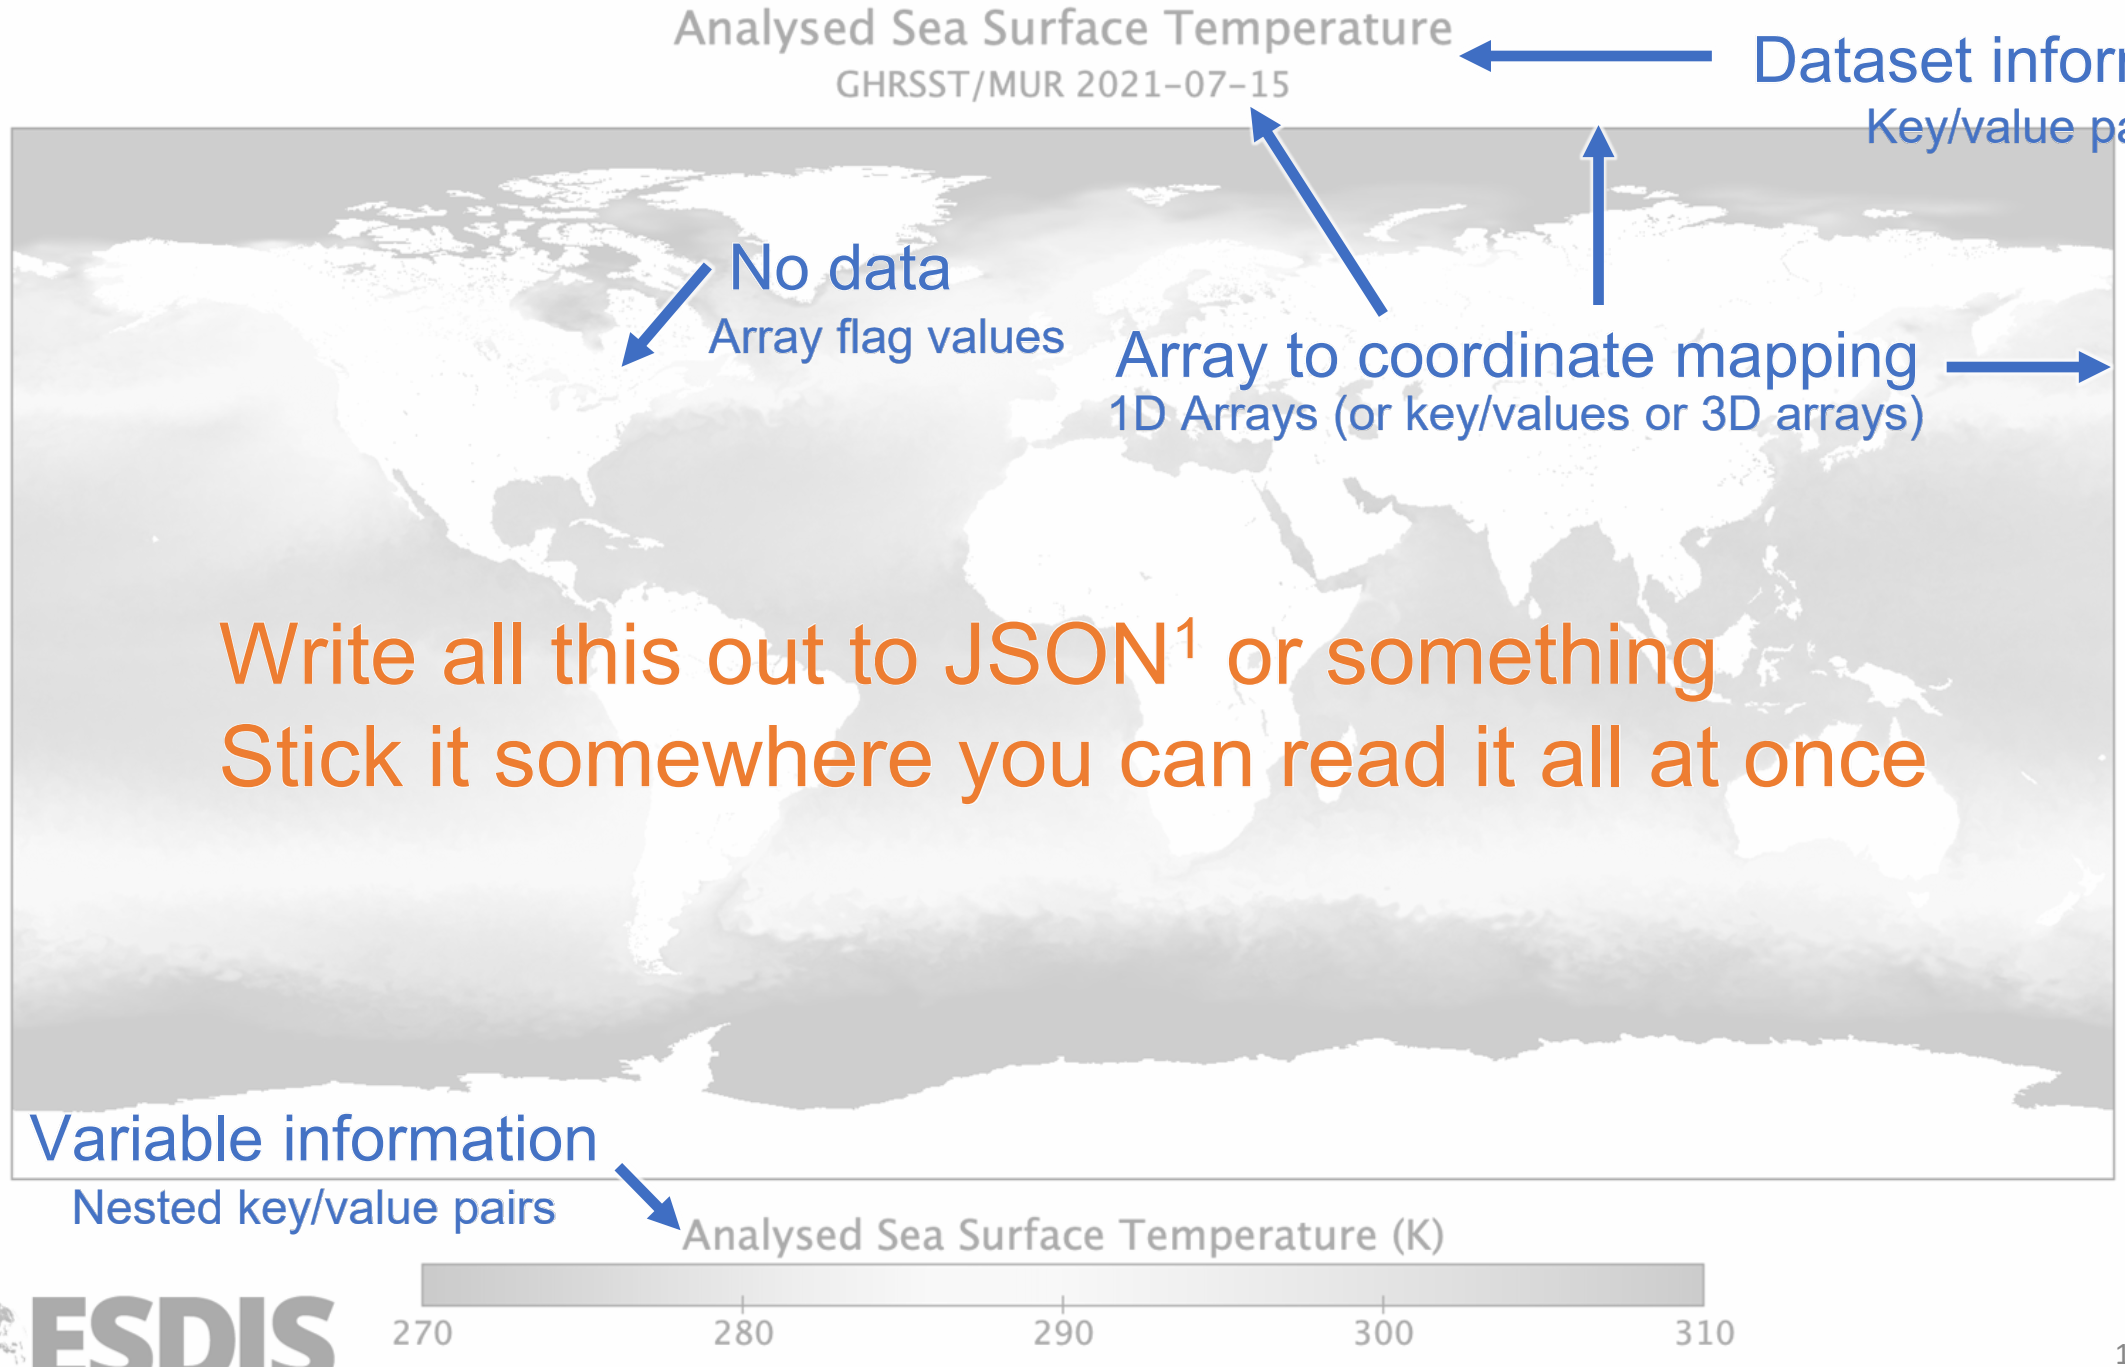

1. Javascript Object Notation

Source: DOI 10.5067/GHGMR-4FJ04

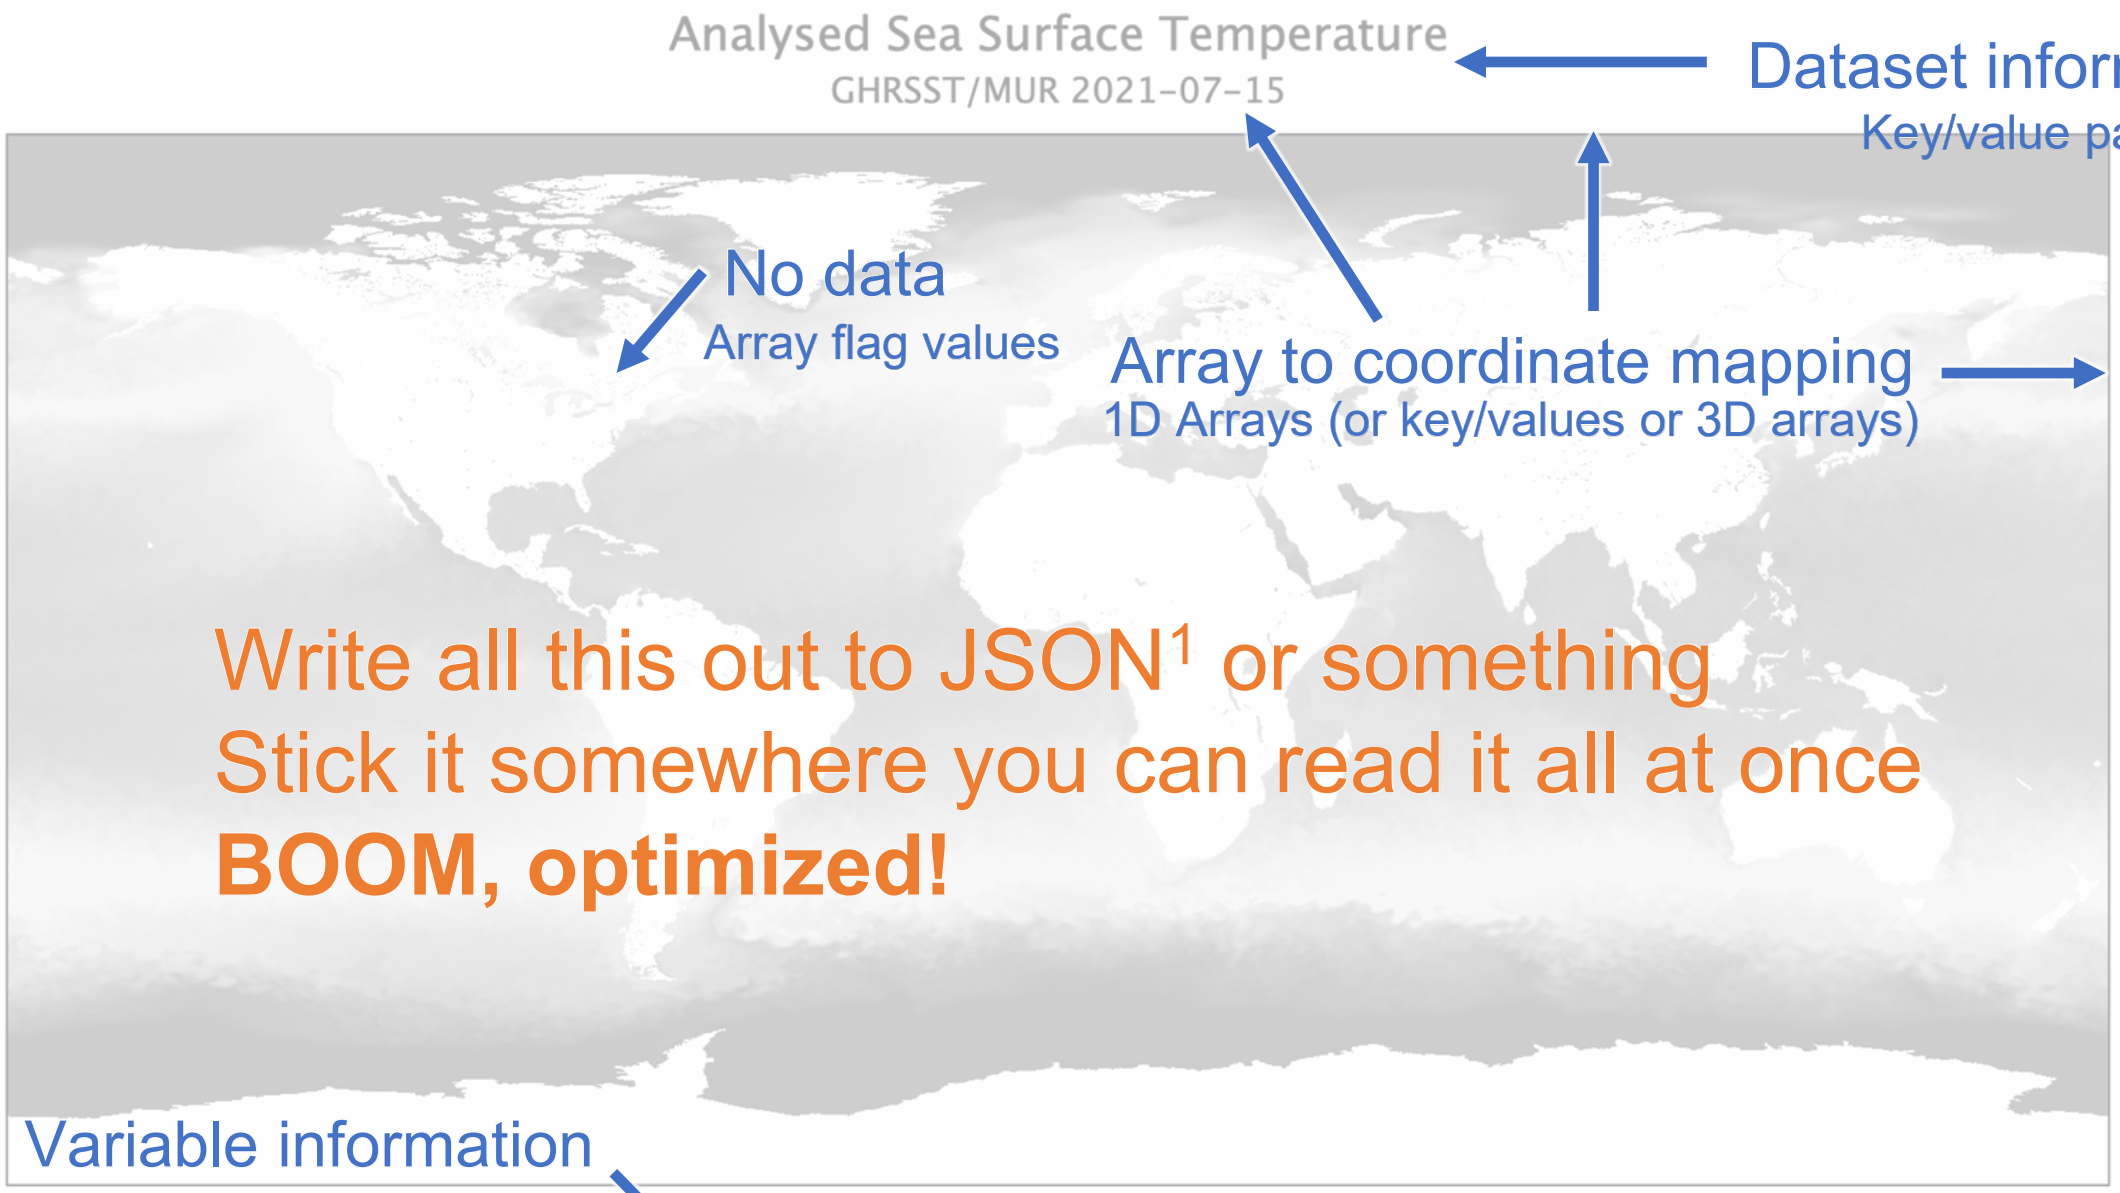

1. Javascript Object Notation

Source: DOI 10.5067/GHGMR-4FJ04

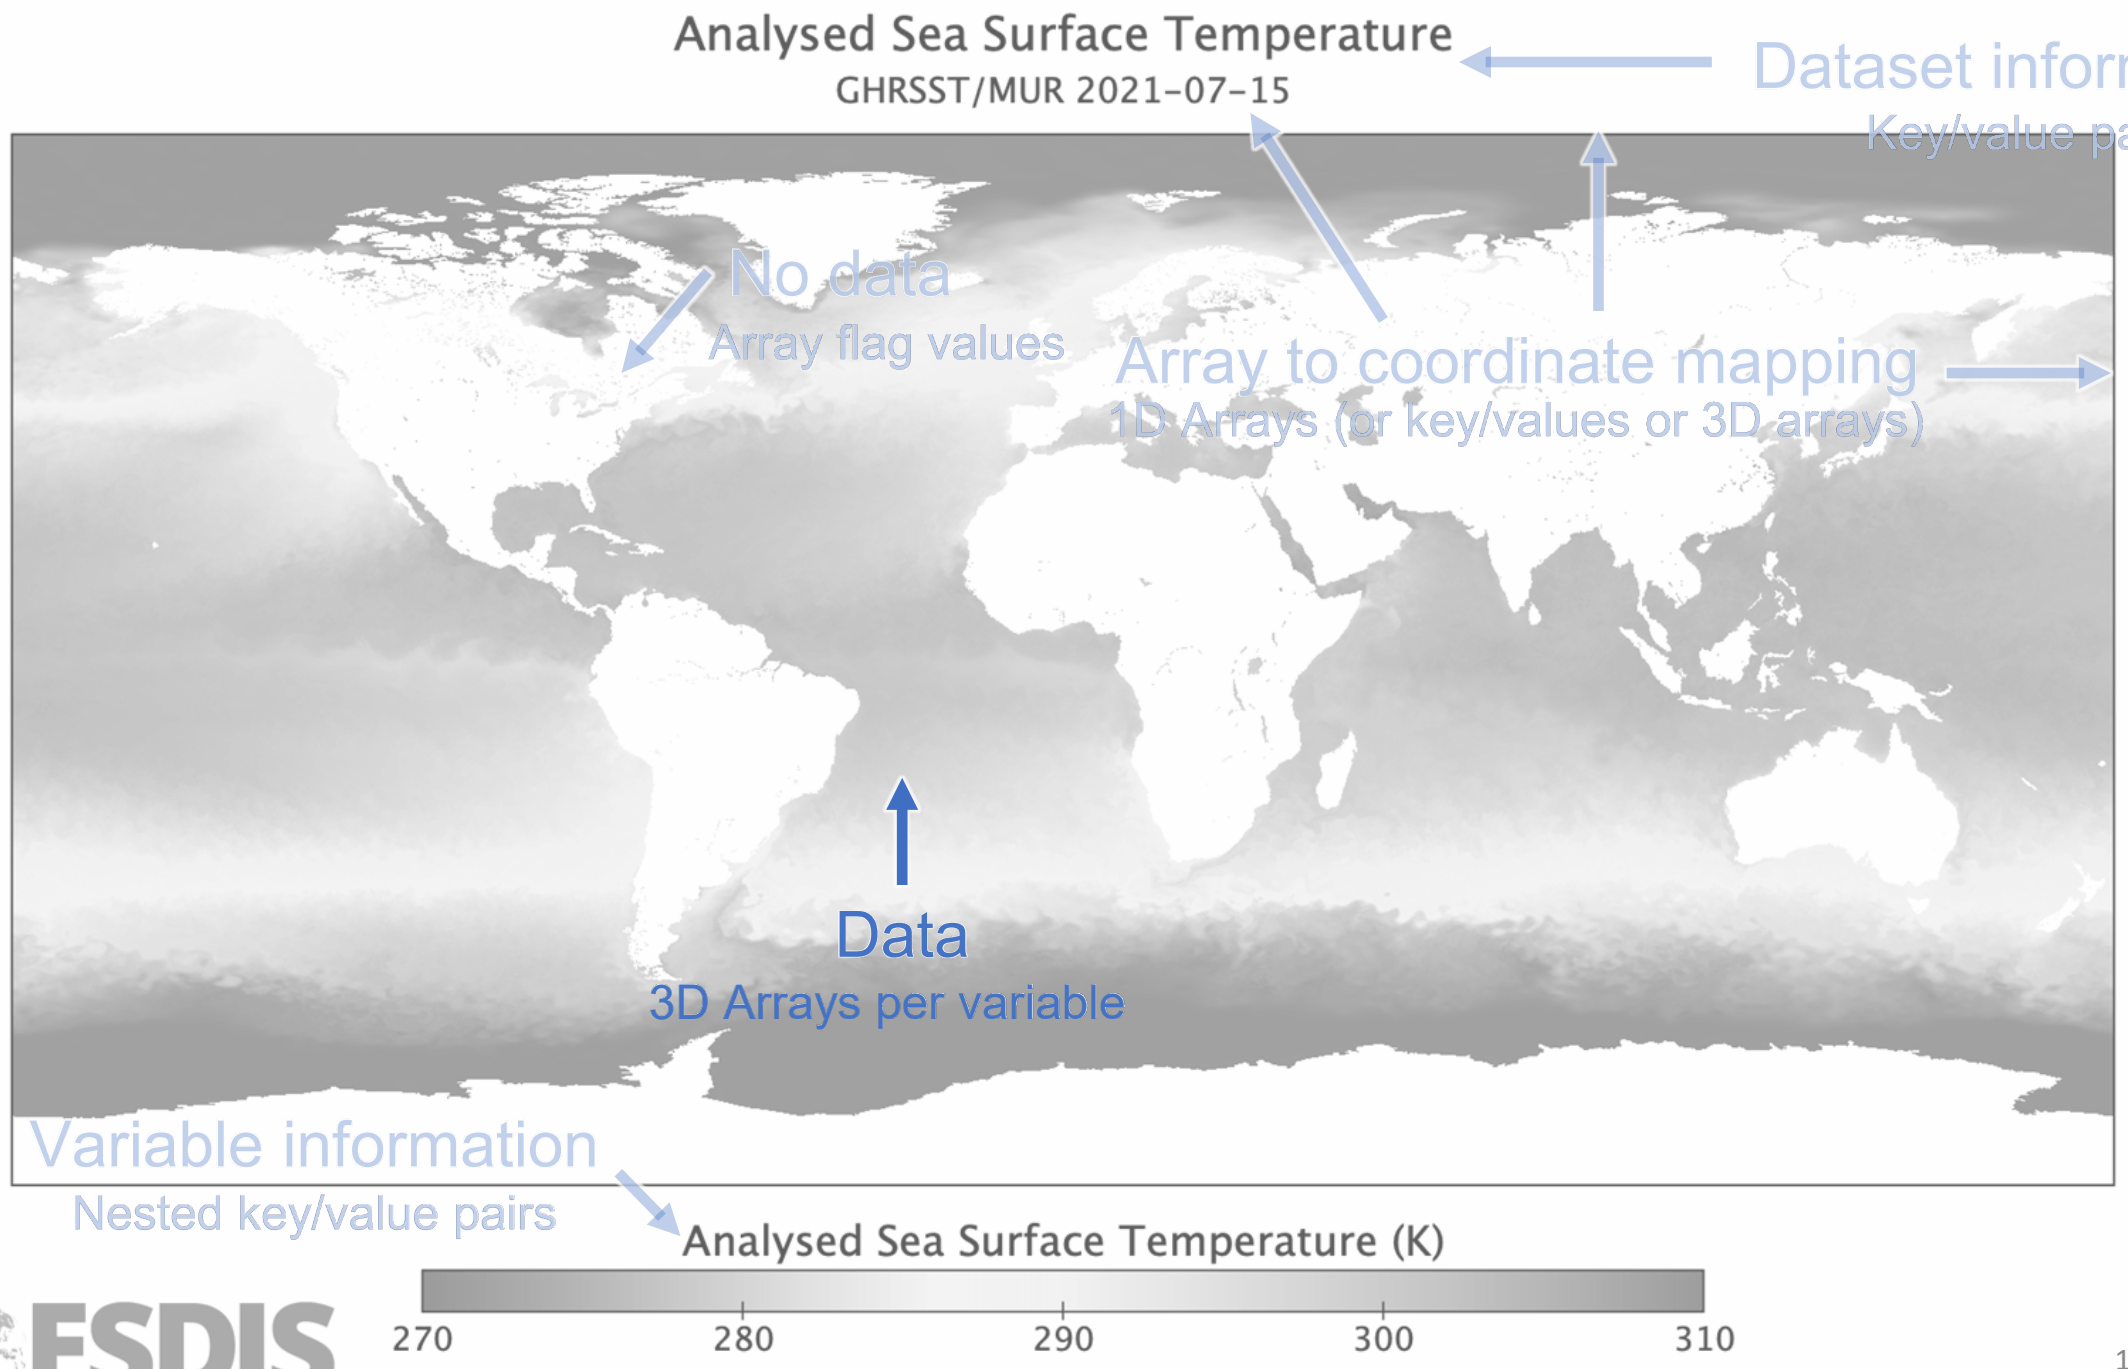

Source: DOI 10.5067/GHGMR-4FJ04

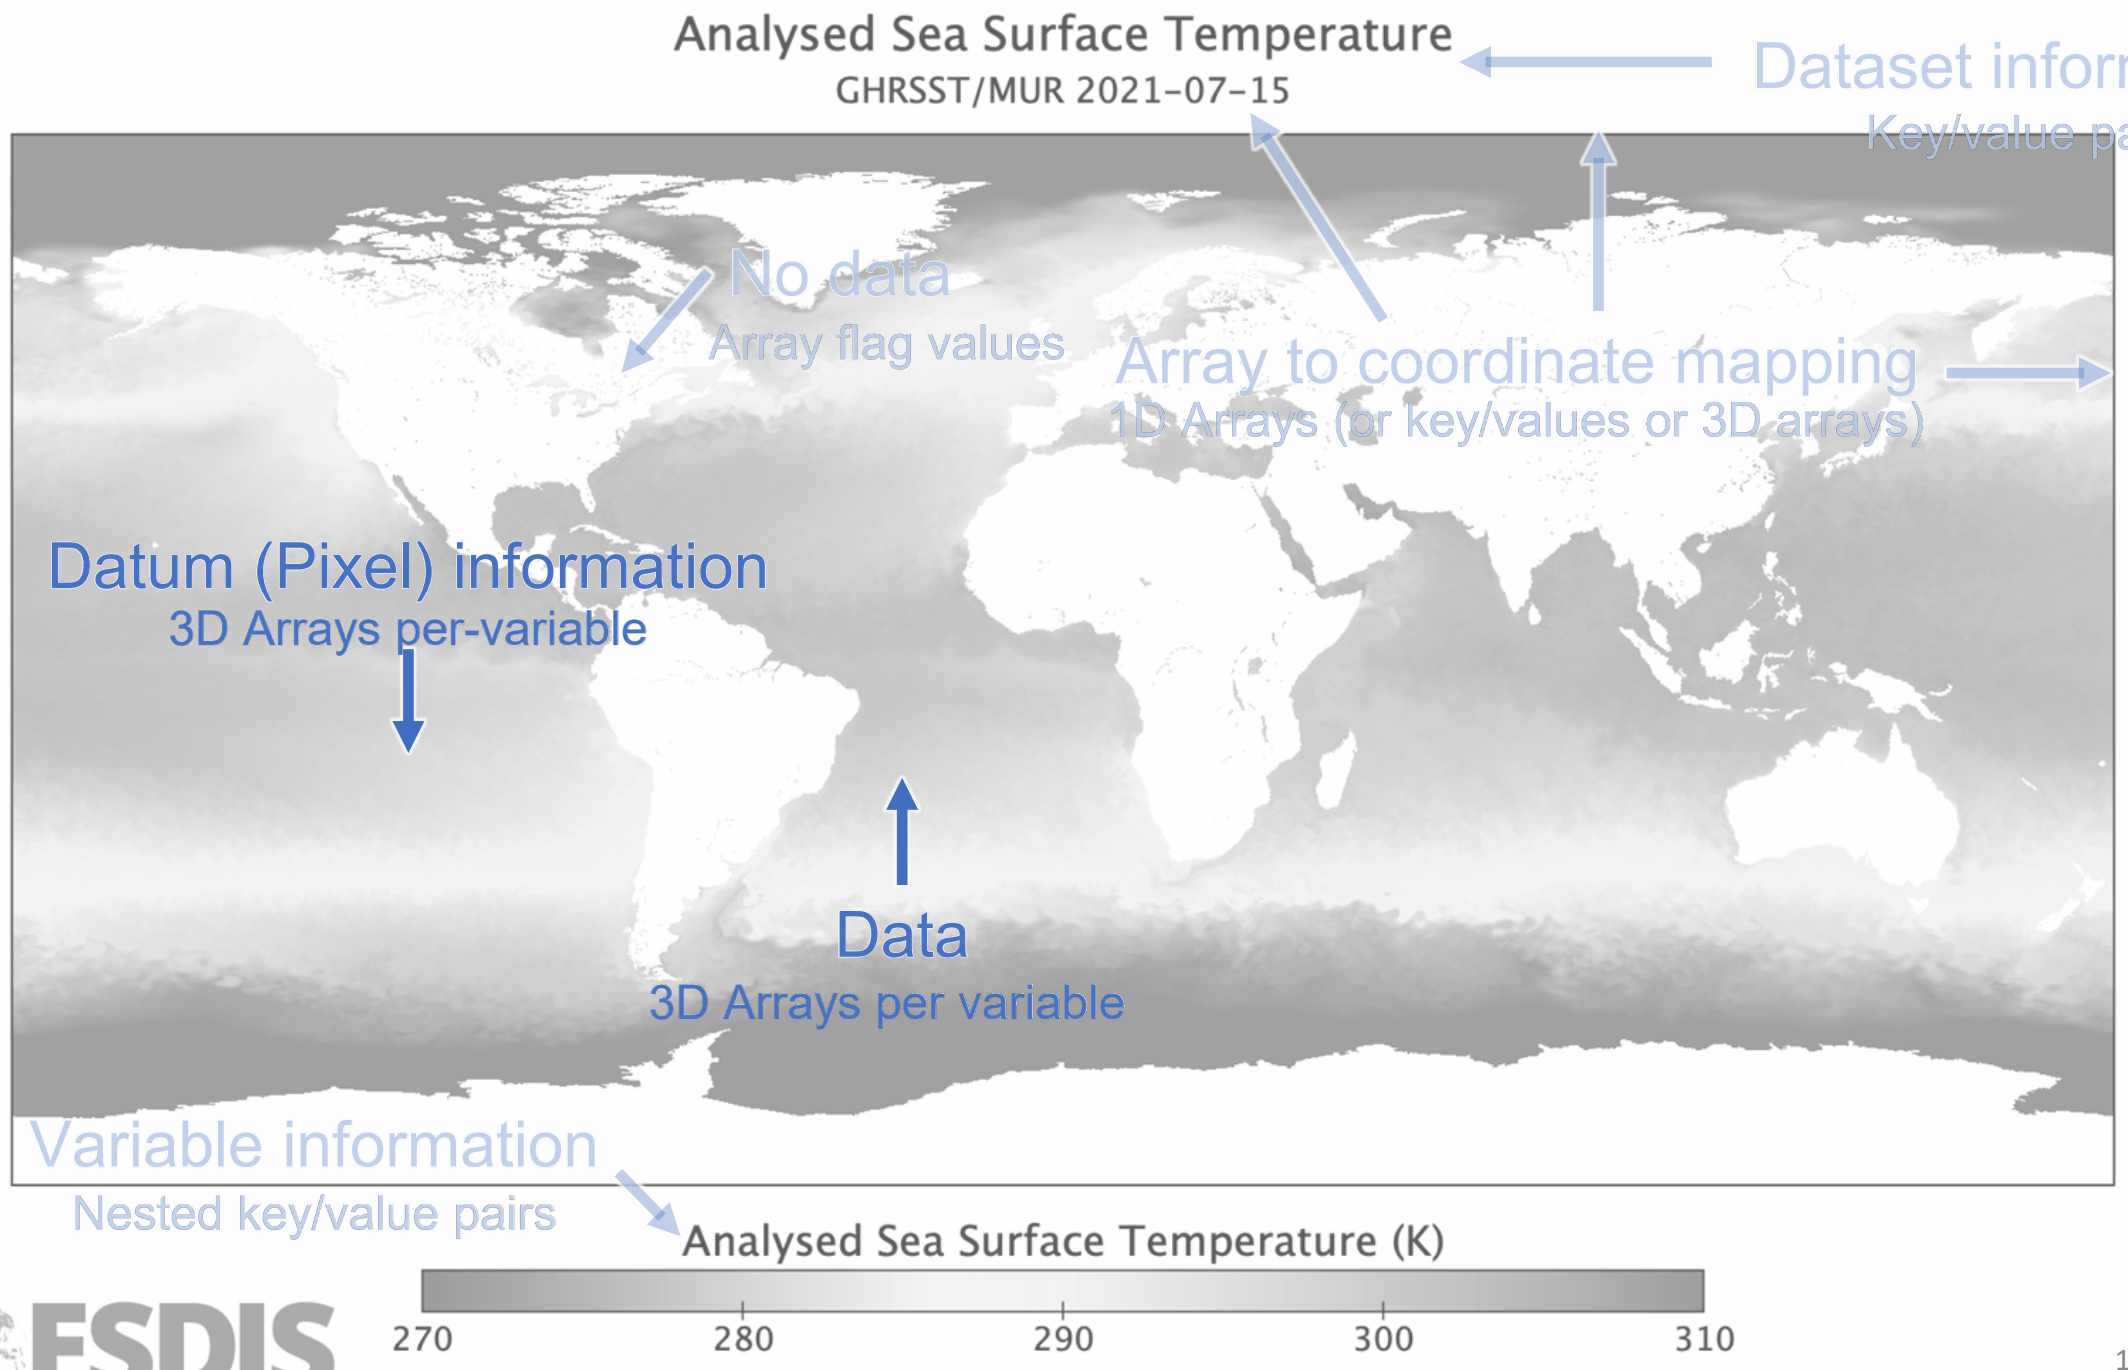

Source: DOI 10.5067/GHGMR-4FJ04

What's the data?

- It's an array. Literally an array as you know it. Not some special sciency overload of the word array:
  ```
  unsigned short data[TIME_SIZE][LAT_SIZE][LON_SIZE];
  ```

- It's big.
  ```
  LAT_SIZE == 17999
  LON_SIZE == 36000
  TIME_SIZE > 8000
  ```

  Naively, that's $17999 \times 36000 \times 8000 \times 2 \approx 10$ Terabytes per variable
  $17999 \times 36000 \times 2 \times 4.5 \approx 6$ Gigabytes (GB) per time step, all variables

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

What might we want to do with the data?

- Access and crunch on everything
  ```
  do_science(data[:, :, :])
  ```

- Investigate how single locations change over time
  ```
  do_science(data[:, y, x])
  ```

- Investigate a single-time snapshot over an area
  ```
  do_science(data[t, y0:y1, x0:x1])
  ```

- Look at seasonal changes over an area/time (ignoring leap years)
  ```
  do_science(data[t:(t+10):365, y0:y1, x0:x1])
  ```

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

What might we want to do with the data?

- Access and crunch on everything
  `do_science(data[:, :, :])` ← Fastest when `data` is serialized as row-major
- Investigate how single locations change over time
  `do_science(data[:, y, x])` ← Fastest when `data` is serialized as column-major
- Investigate a single-time snapshot over an area
  `do_science(data[t, y0:y1, x0:x1])` ← Fastest when serialization is more clever
- Look at seasonal changes over an area/time (ignoring leap years)
  `do_science(data[t:(t+10):365, y0:y1, x0:x1])`

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

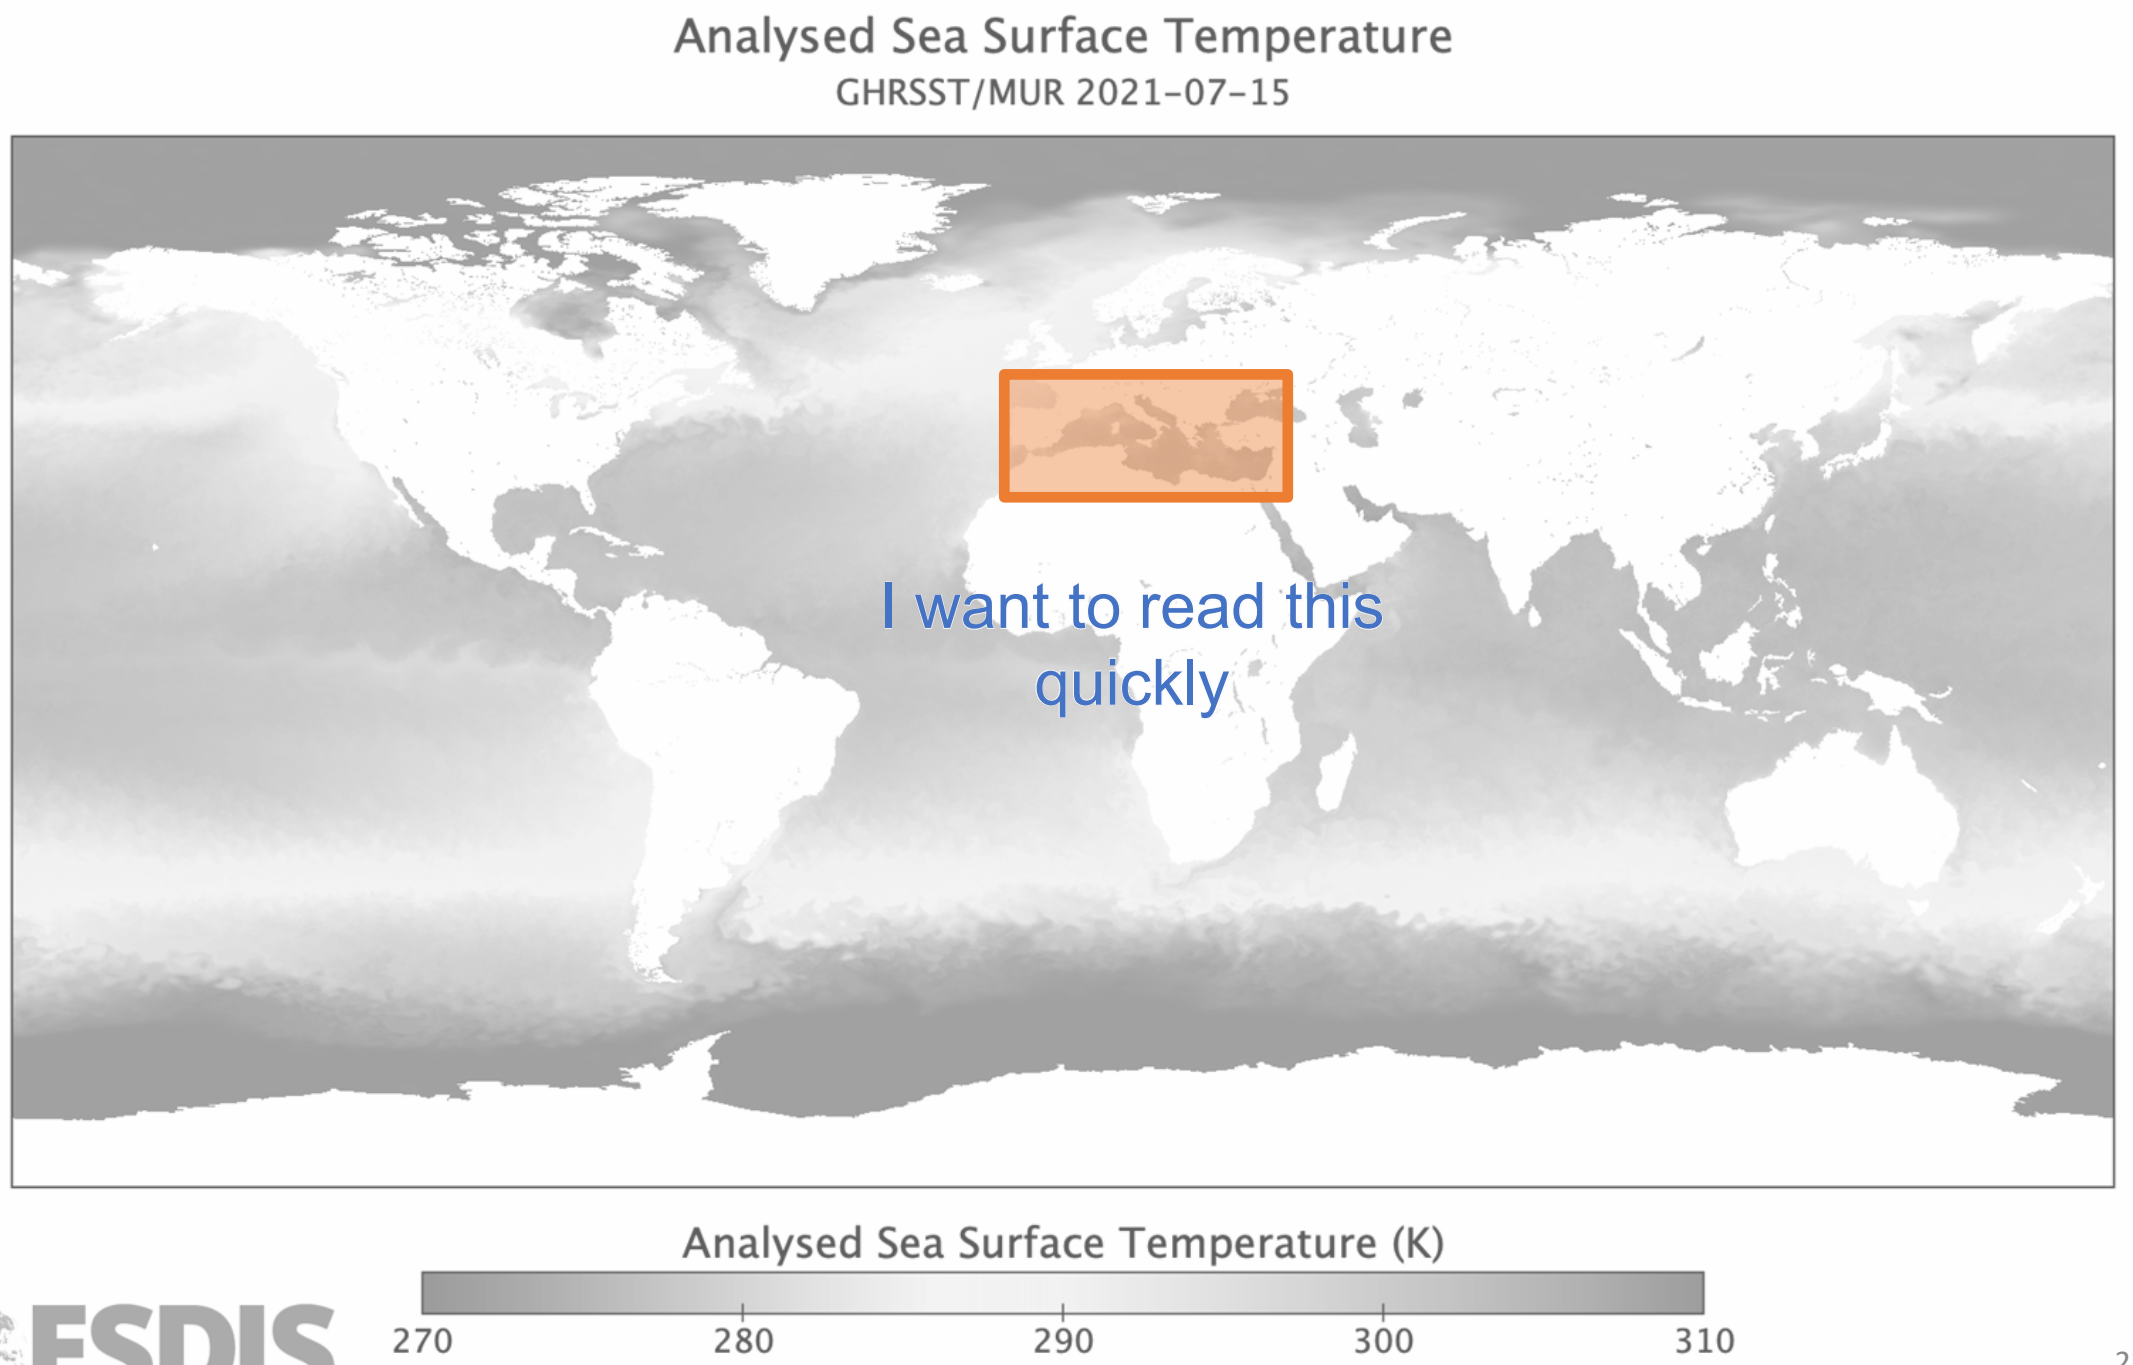

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

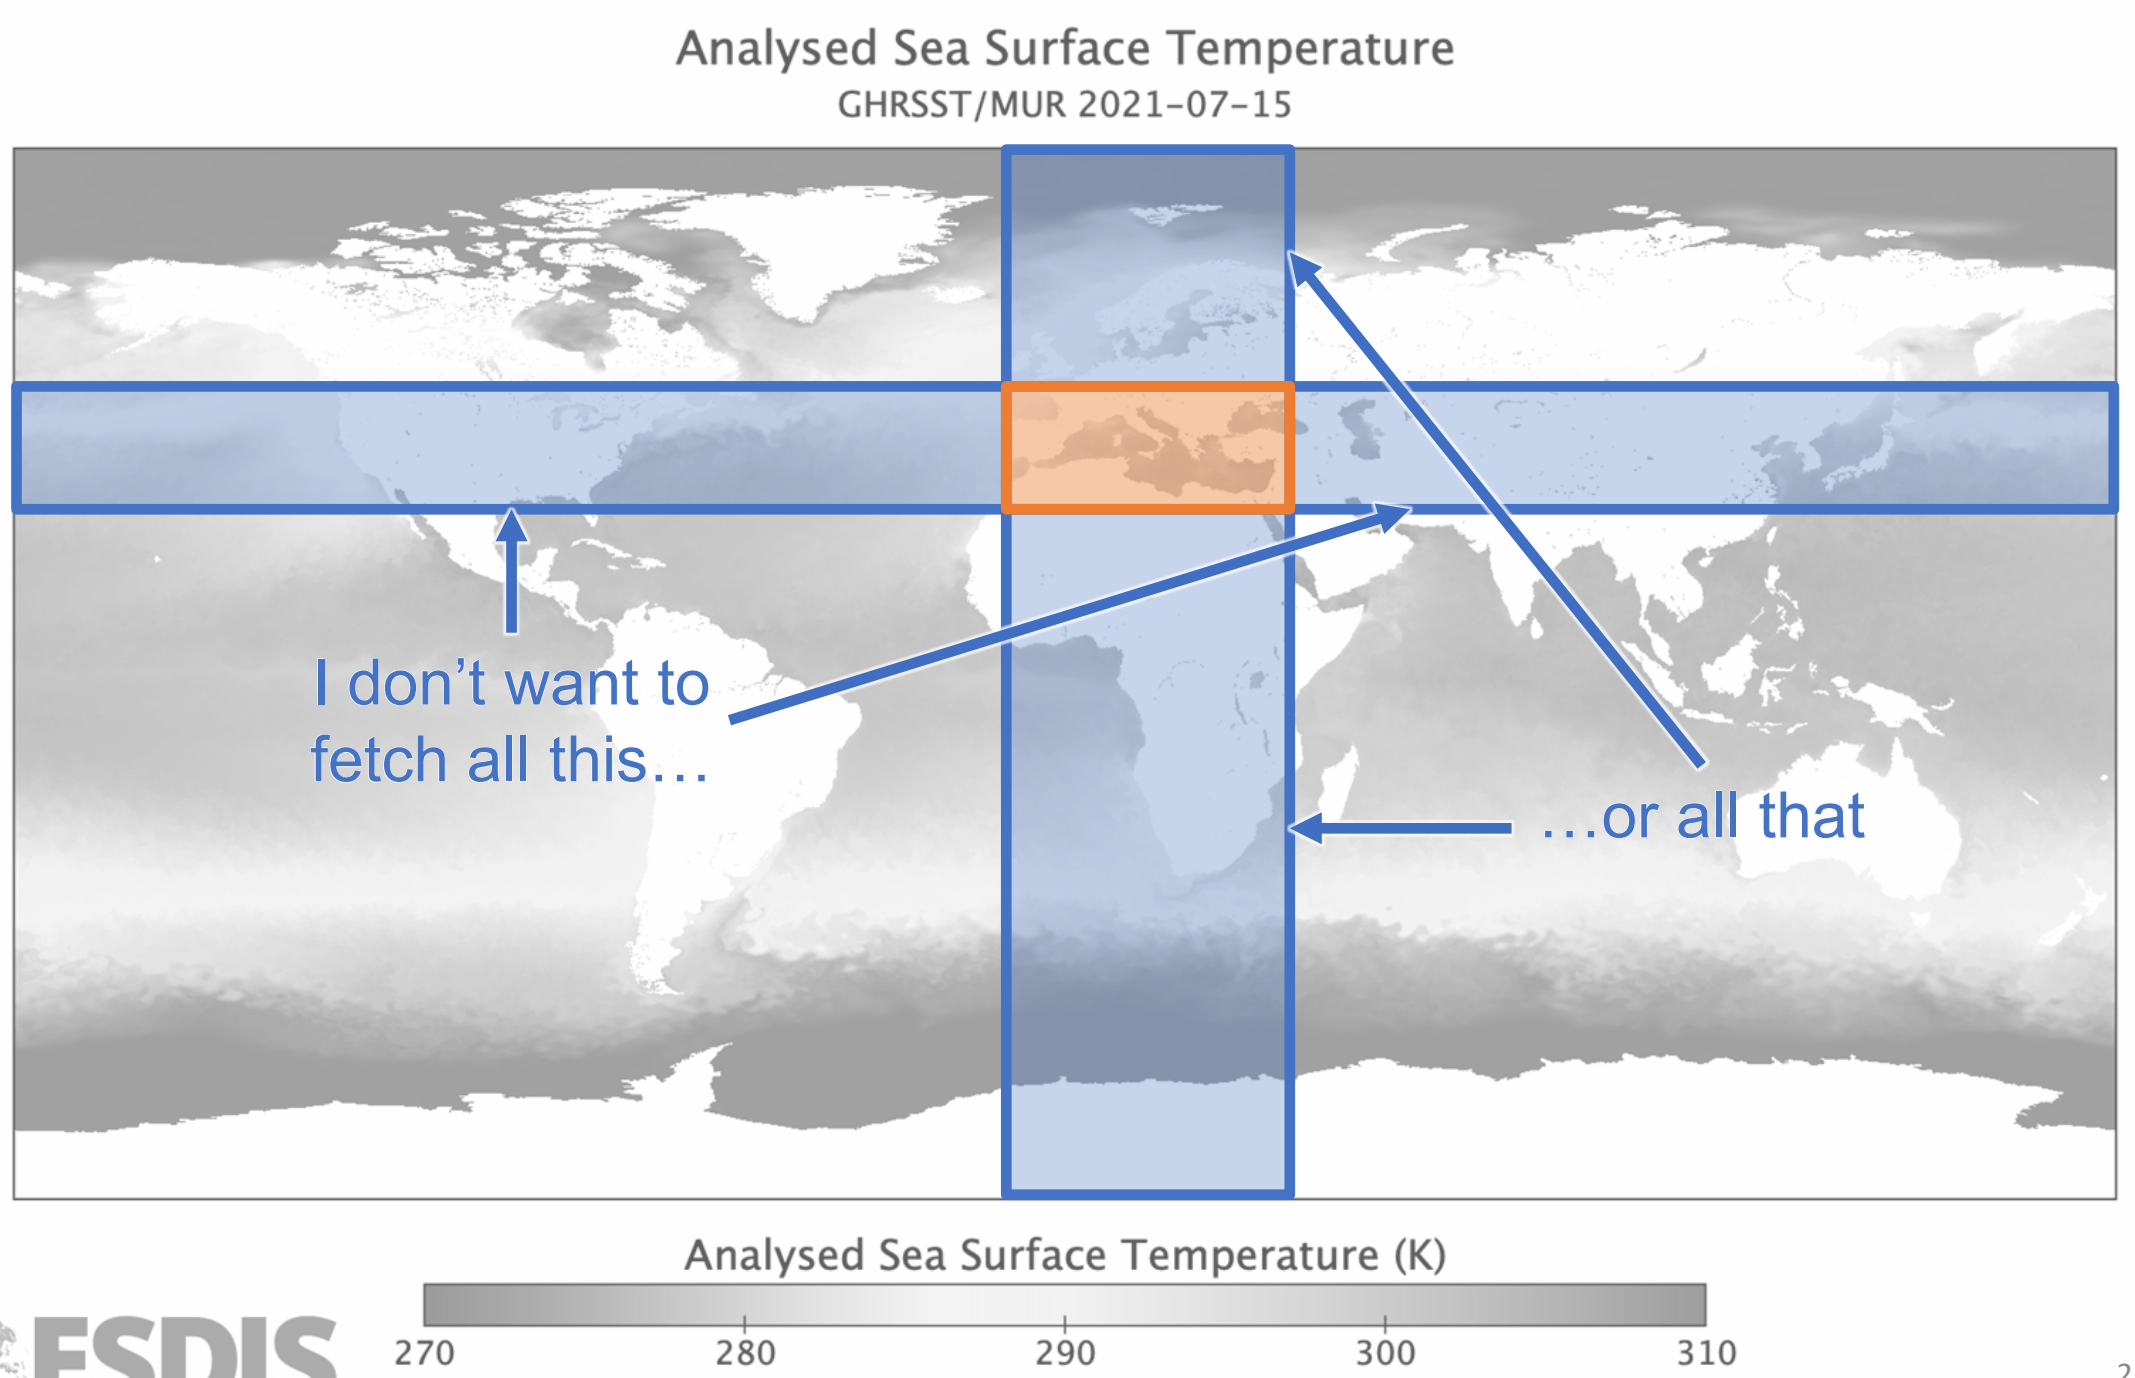

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

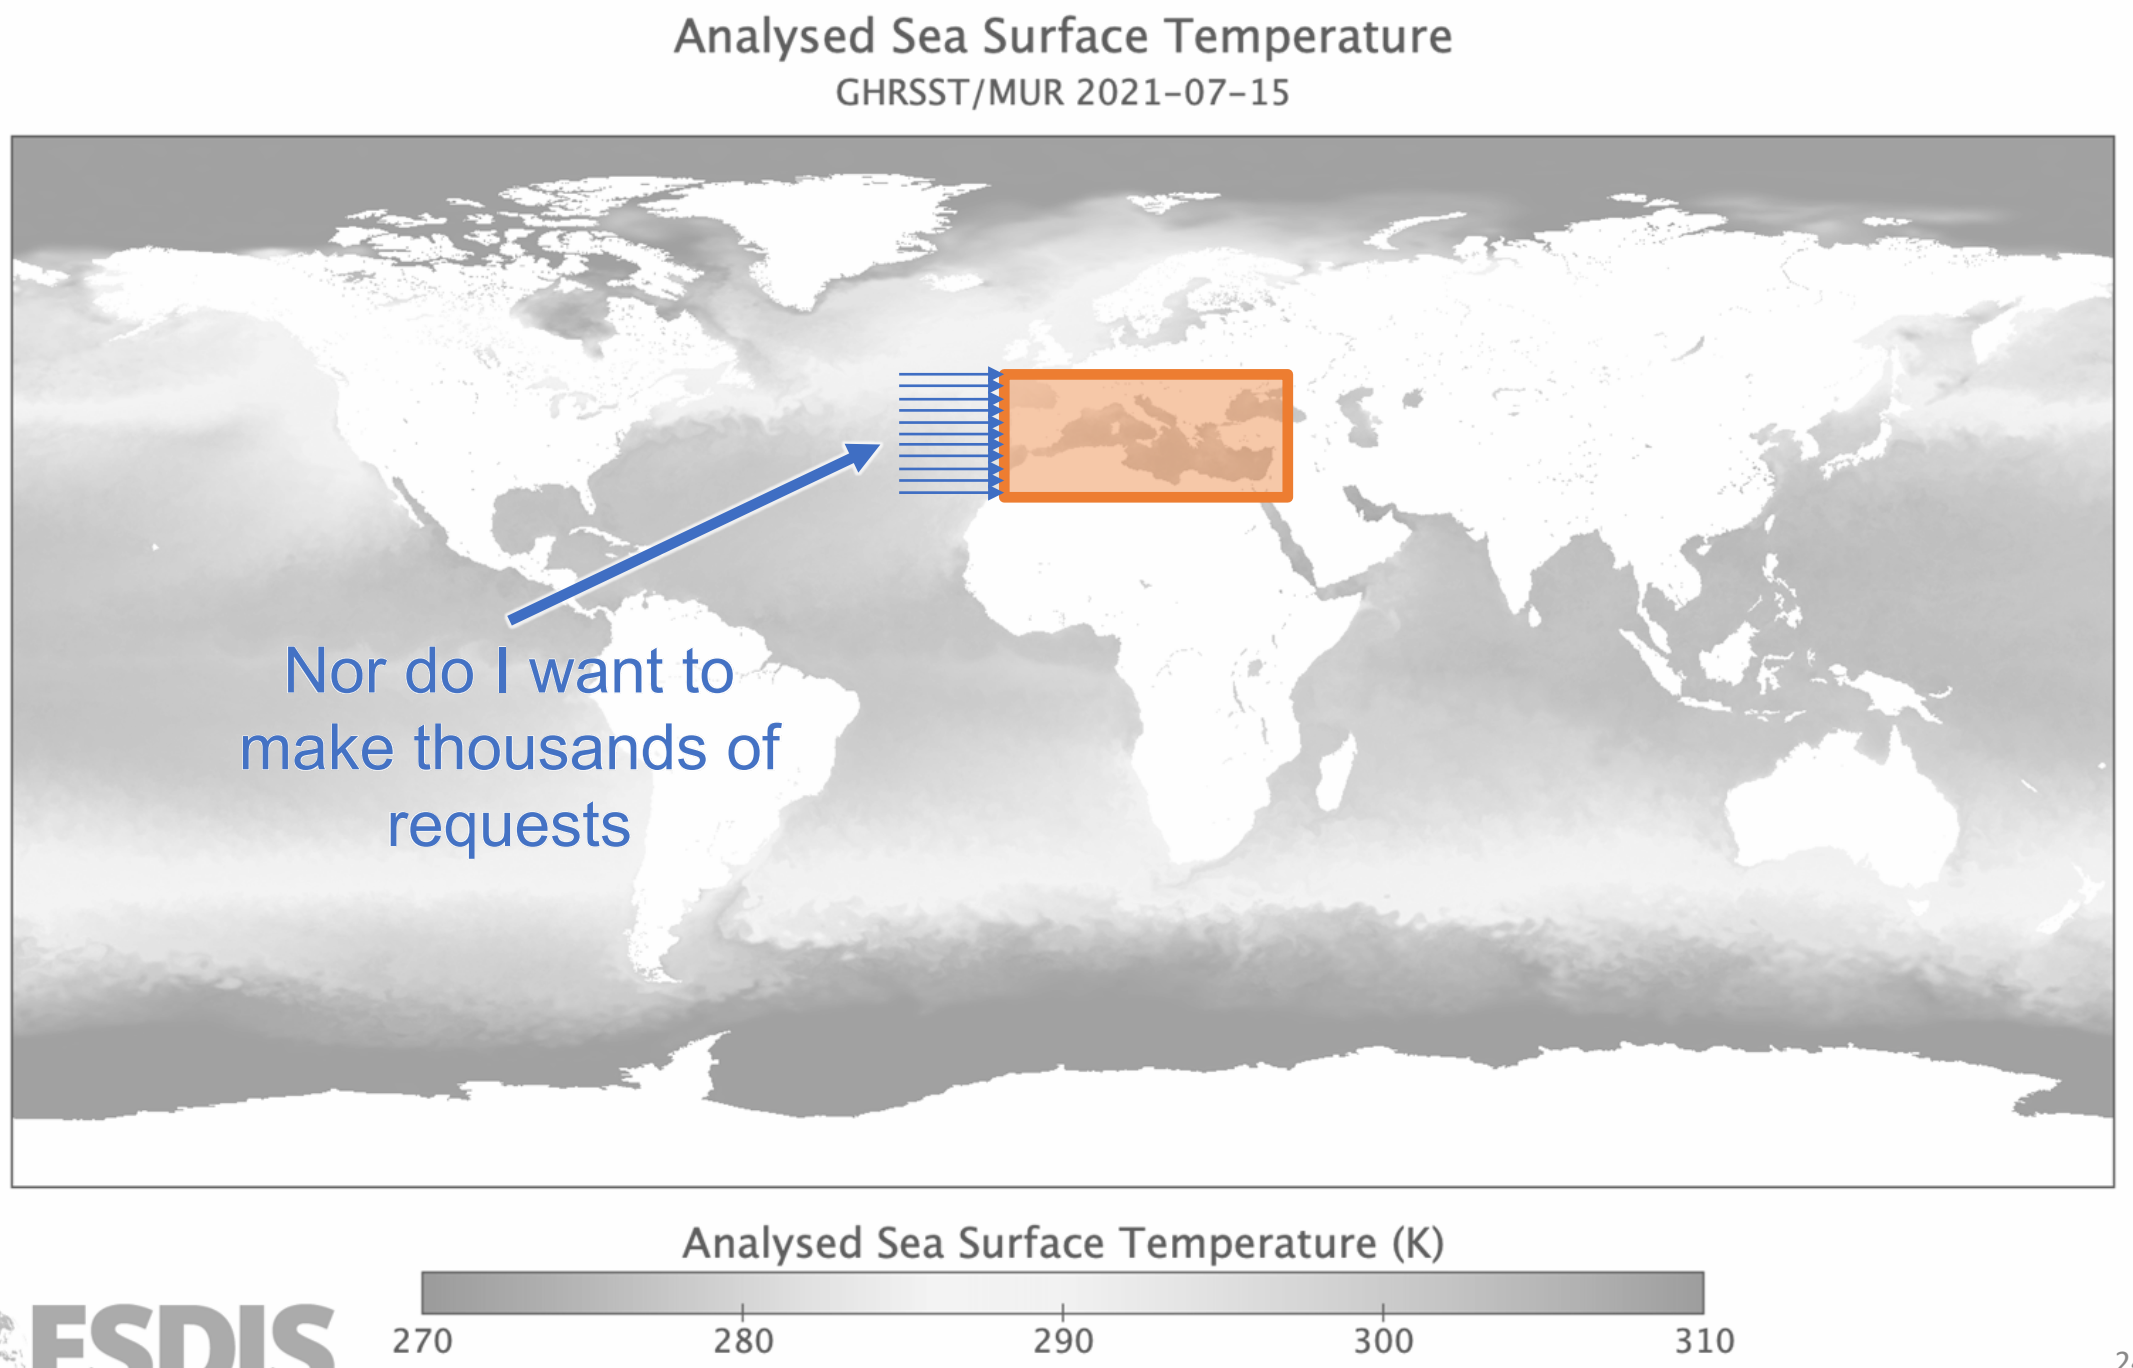

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

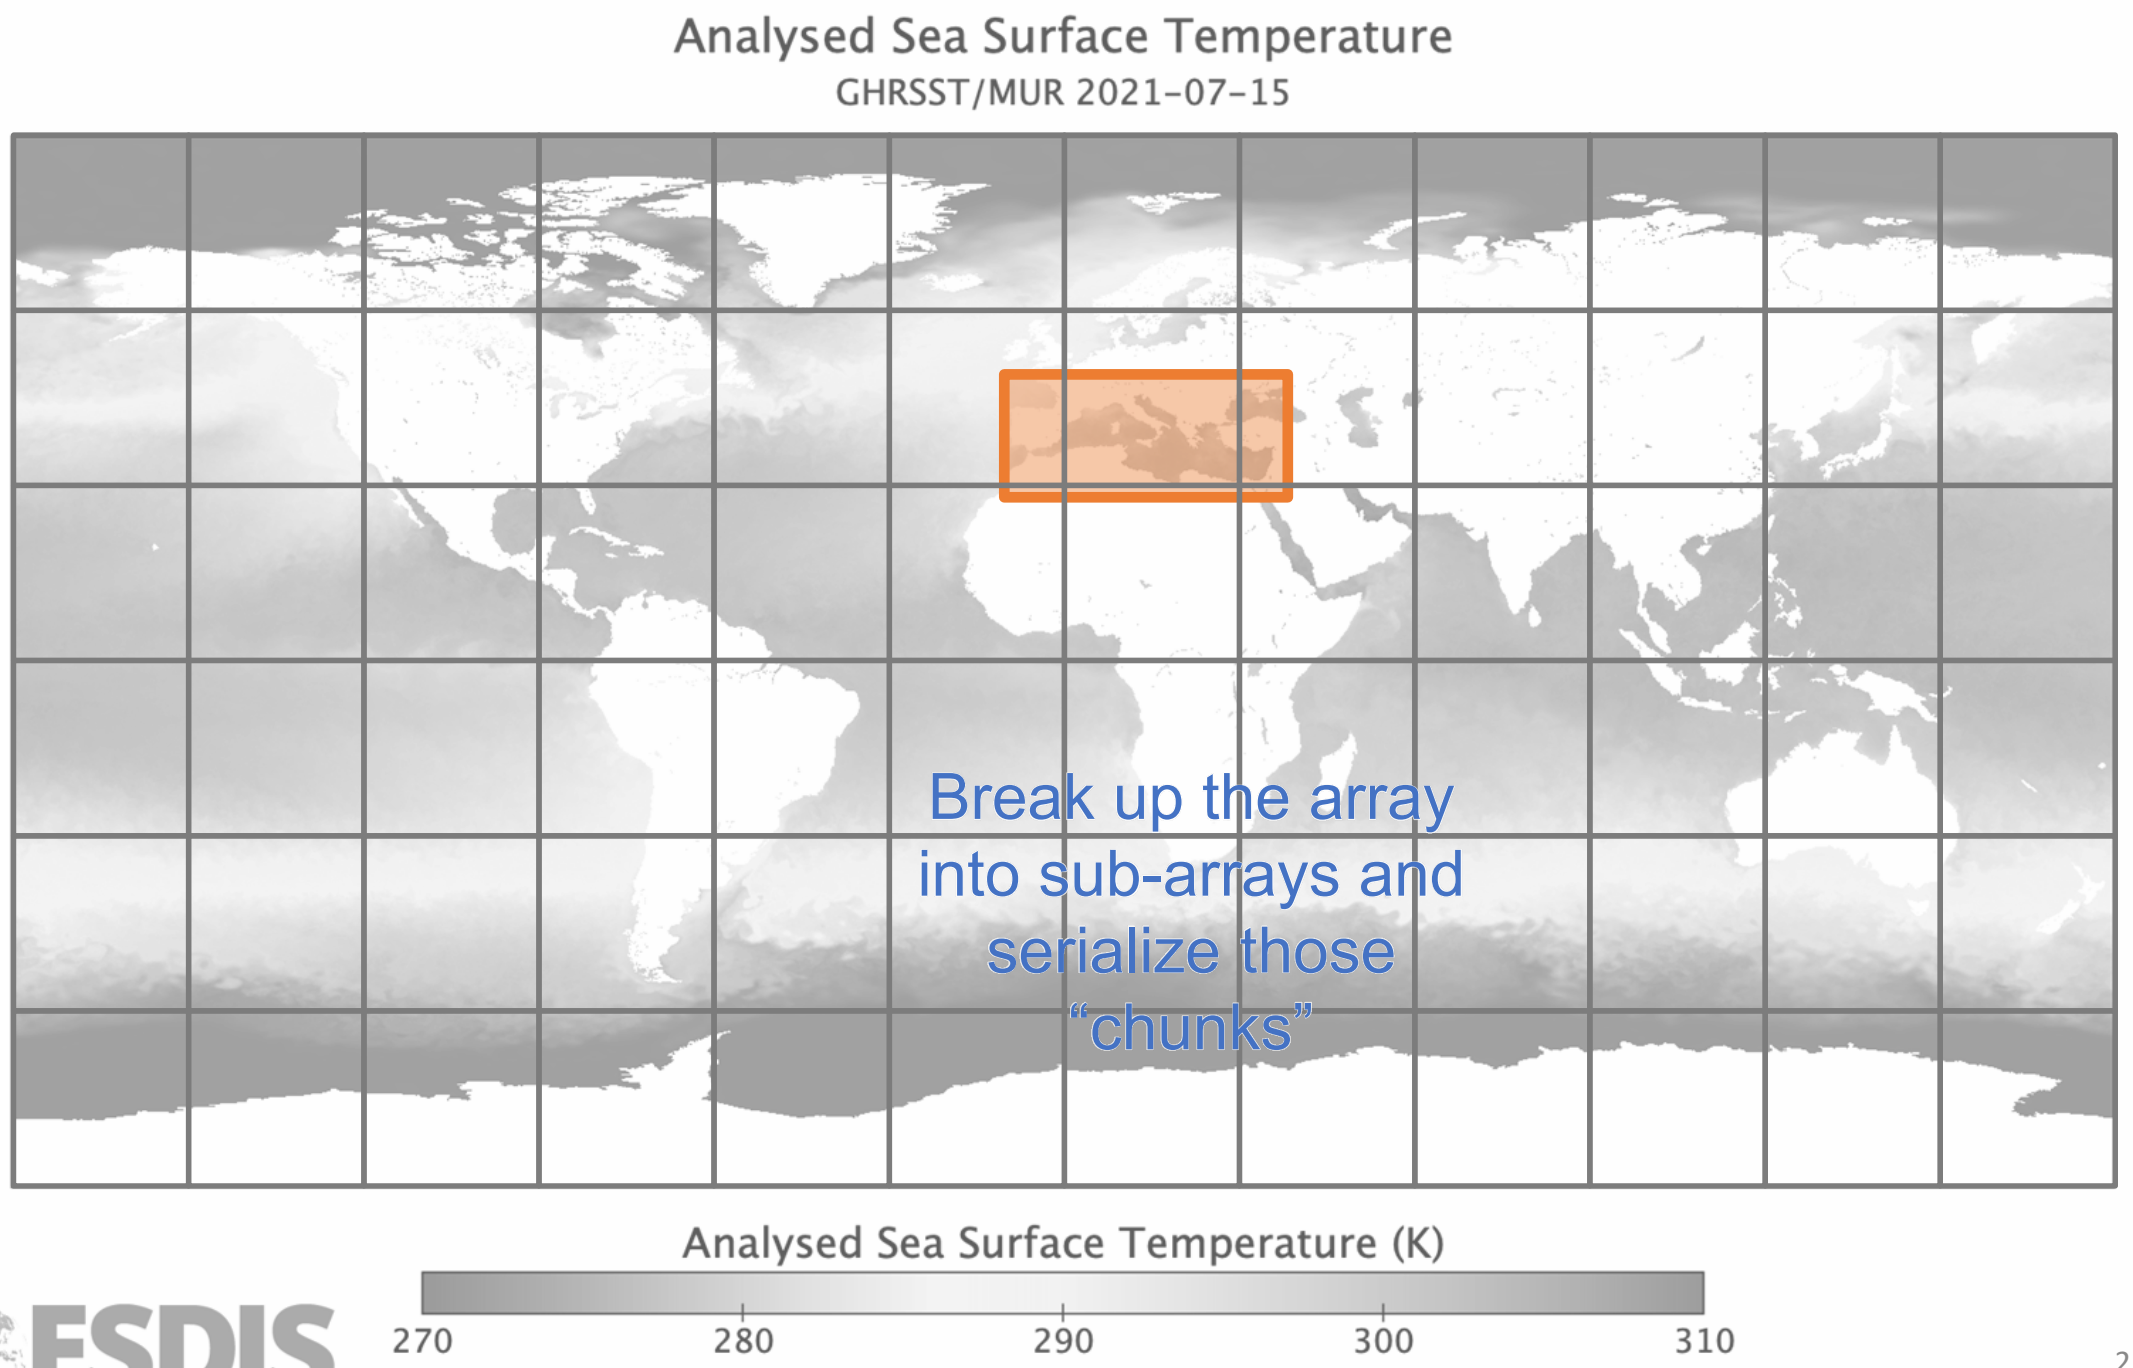

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

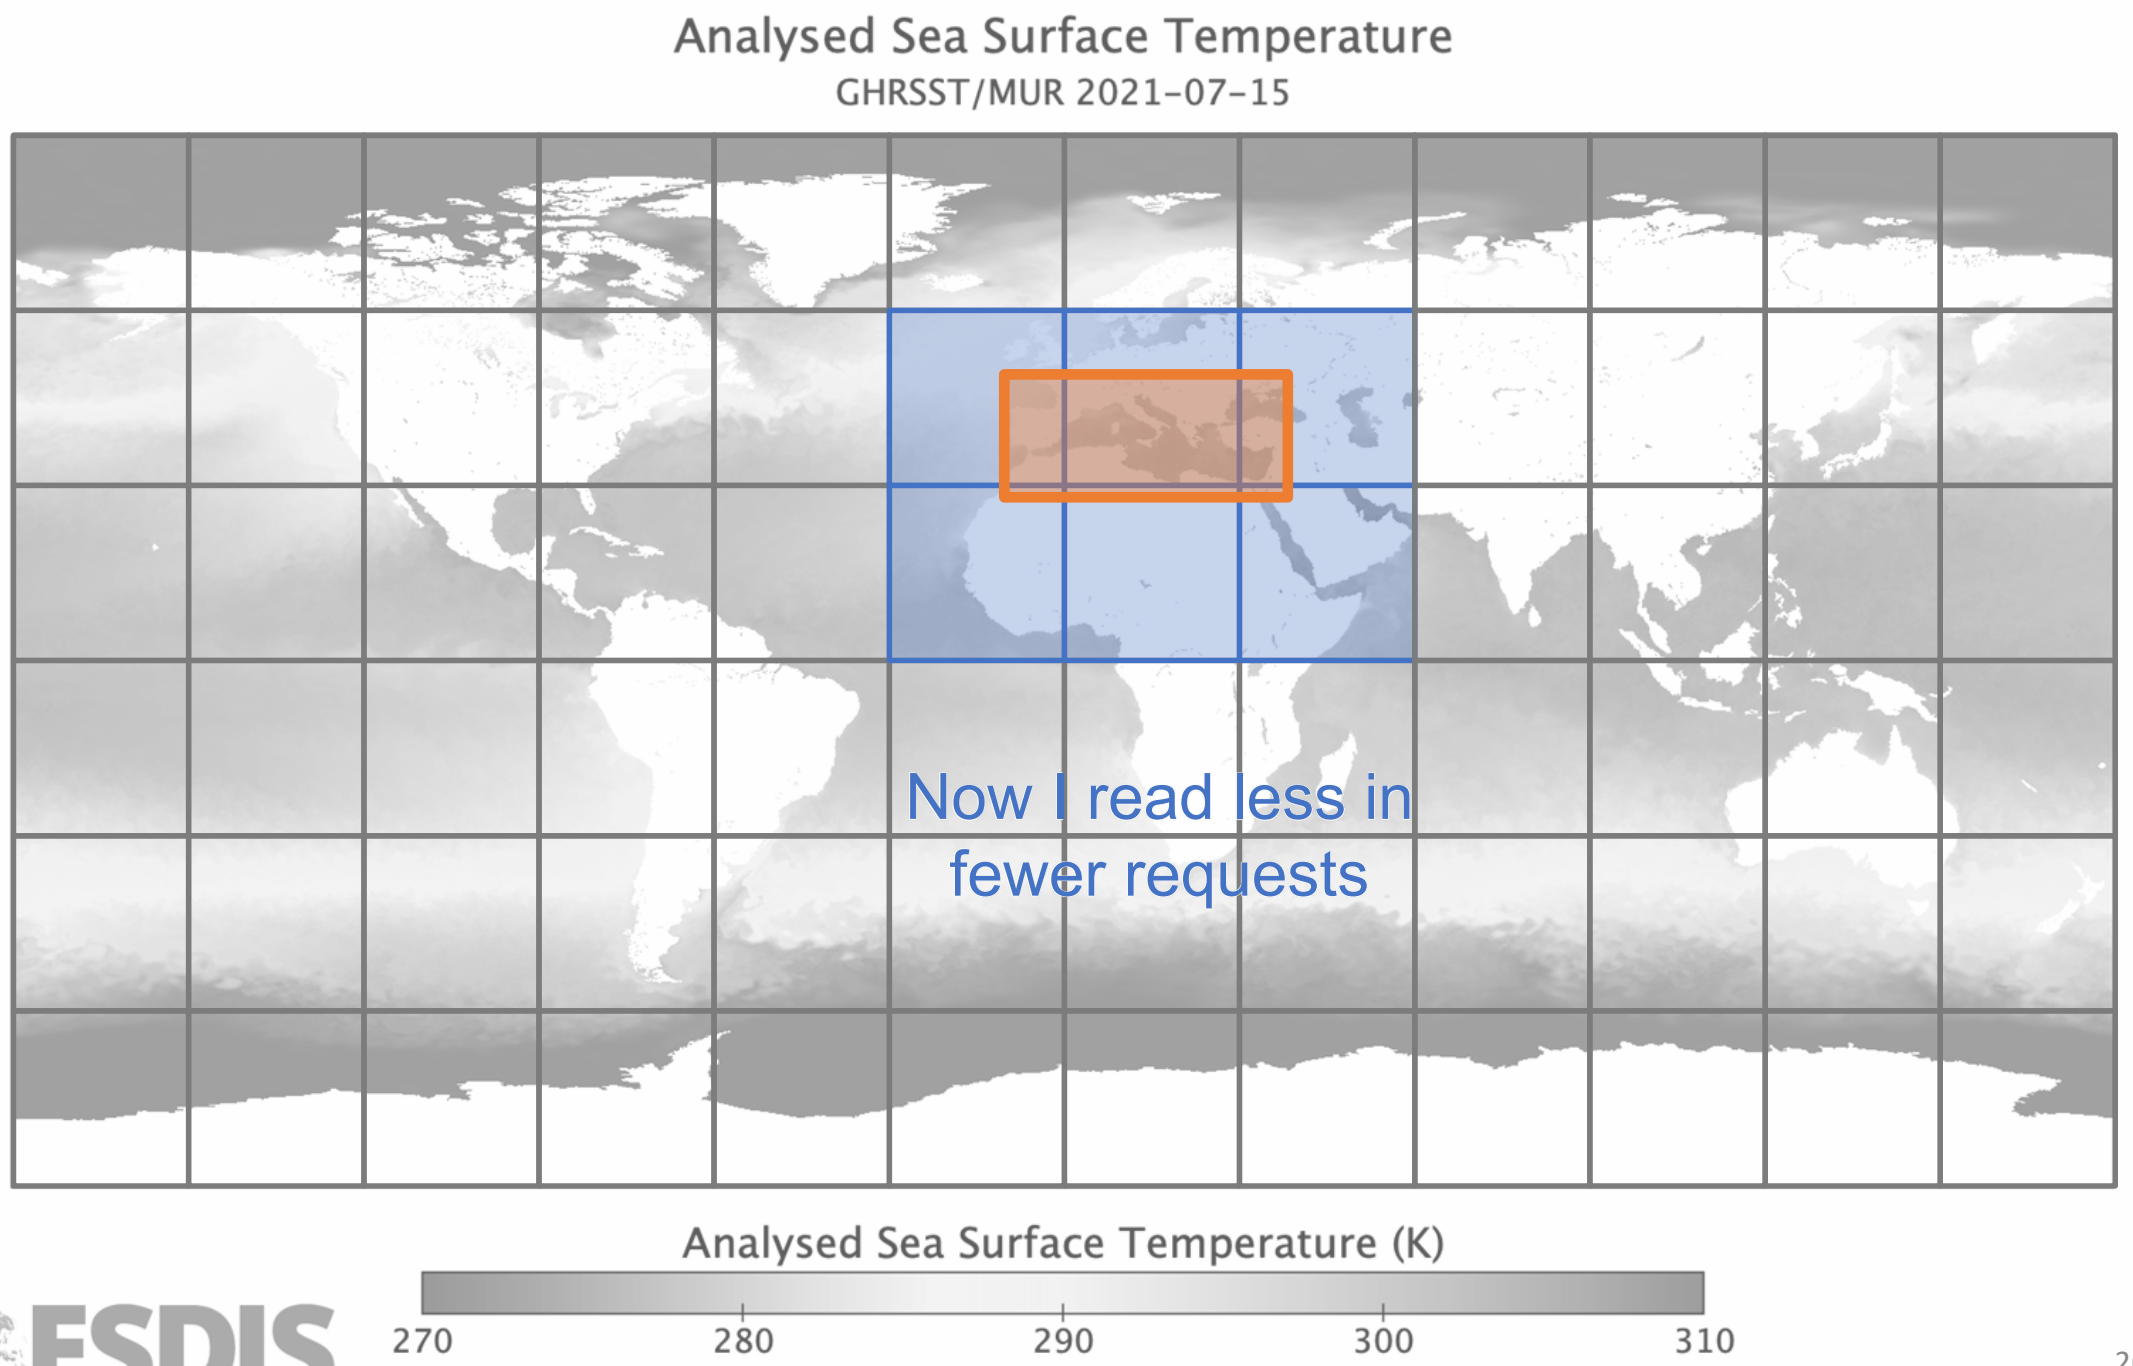

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

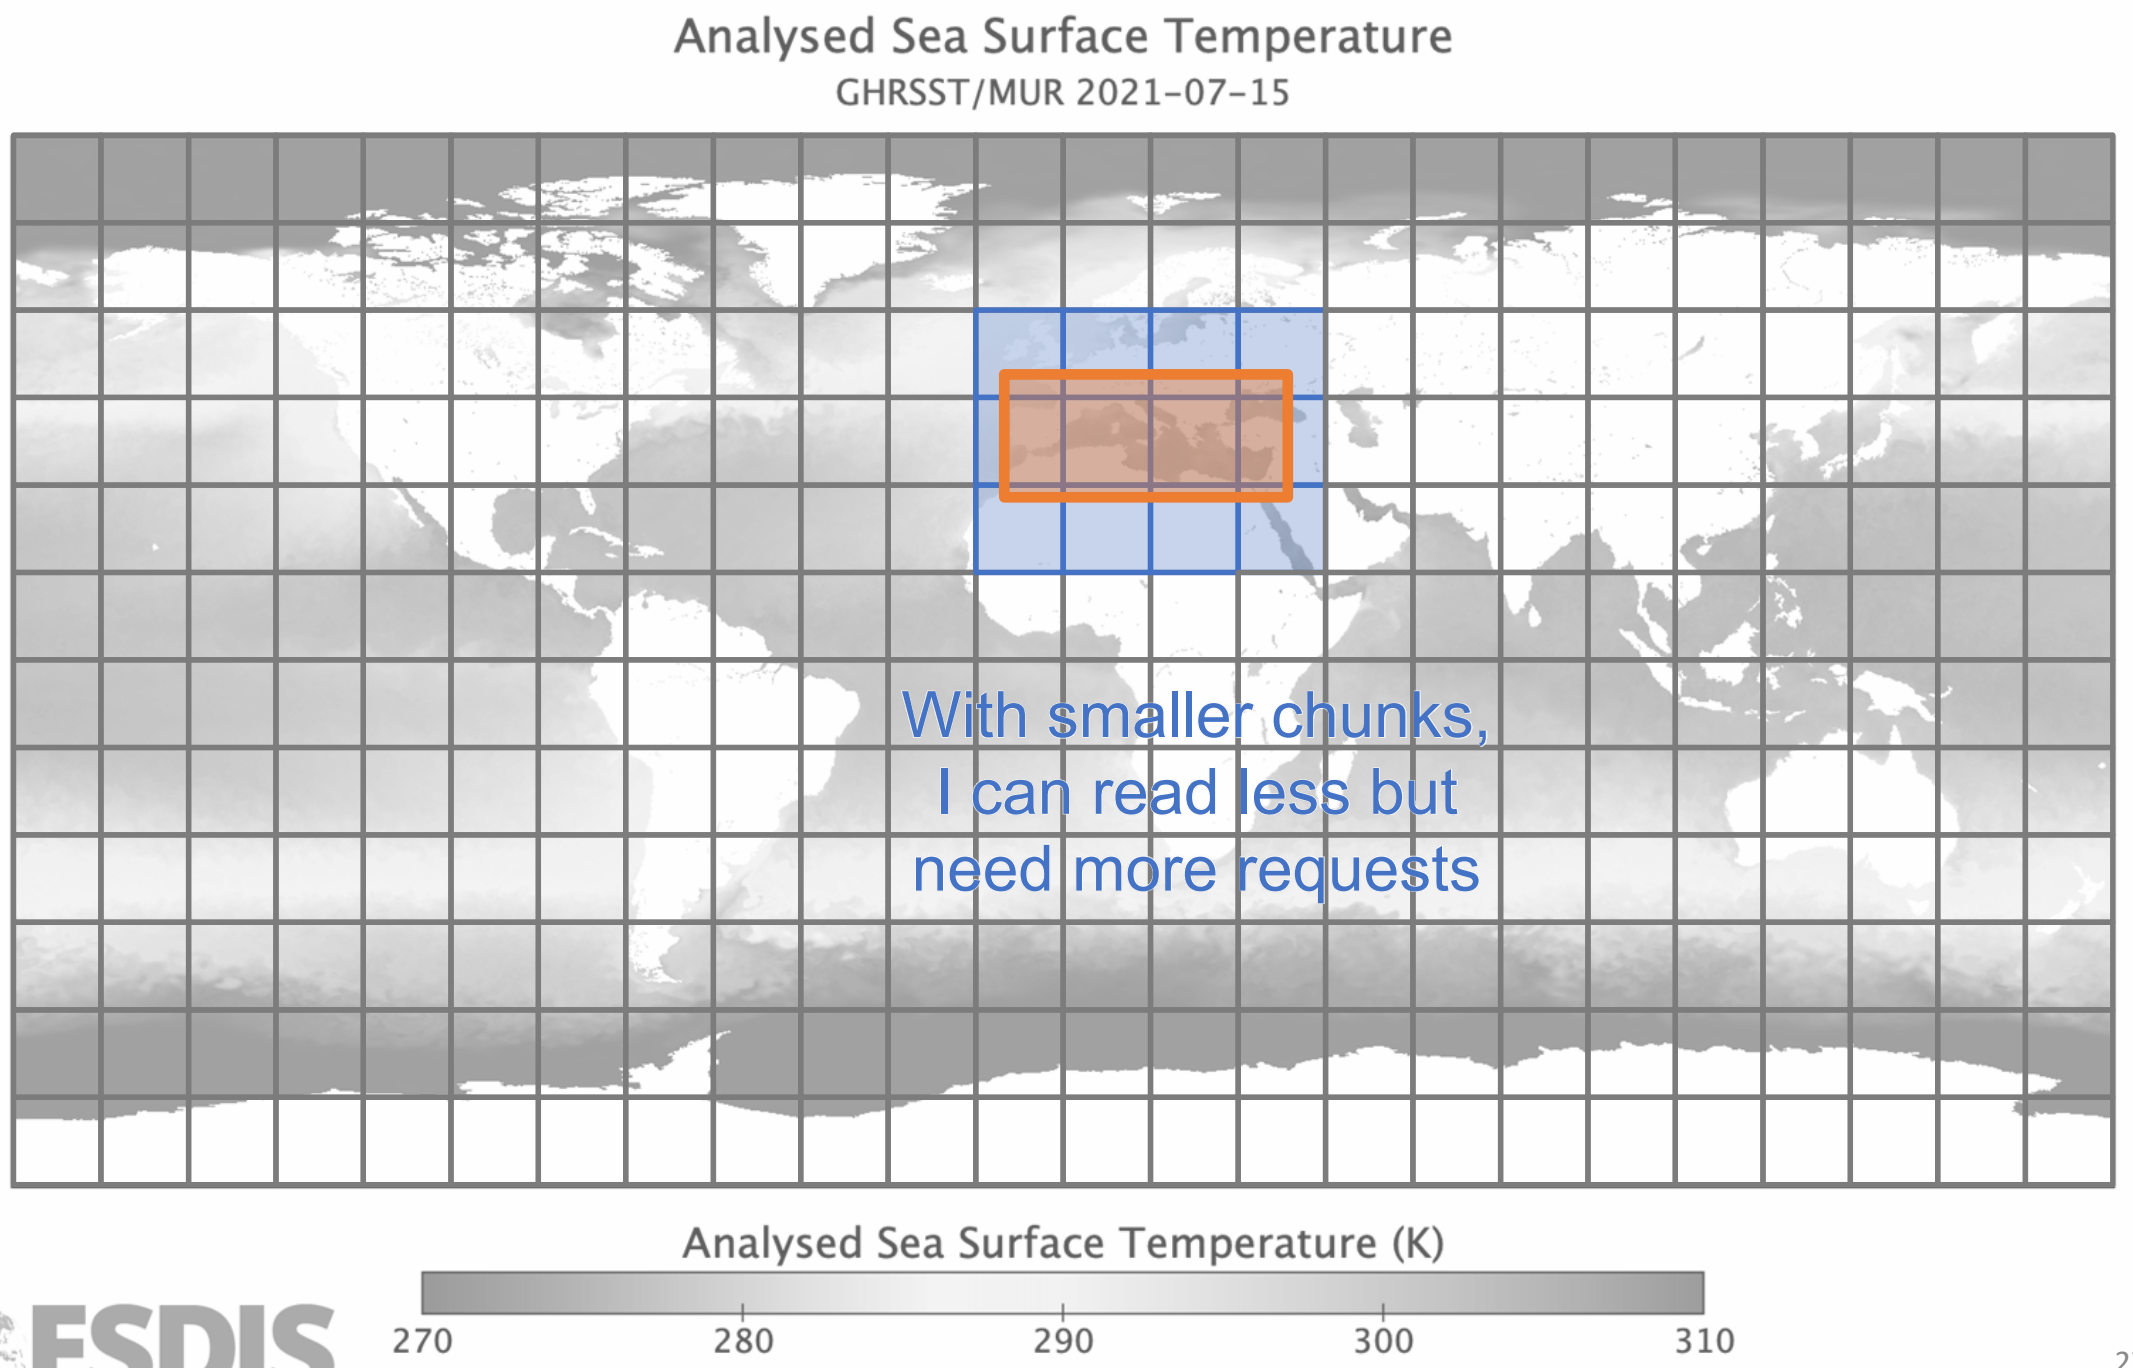

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

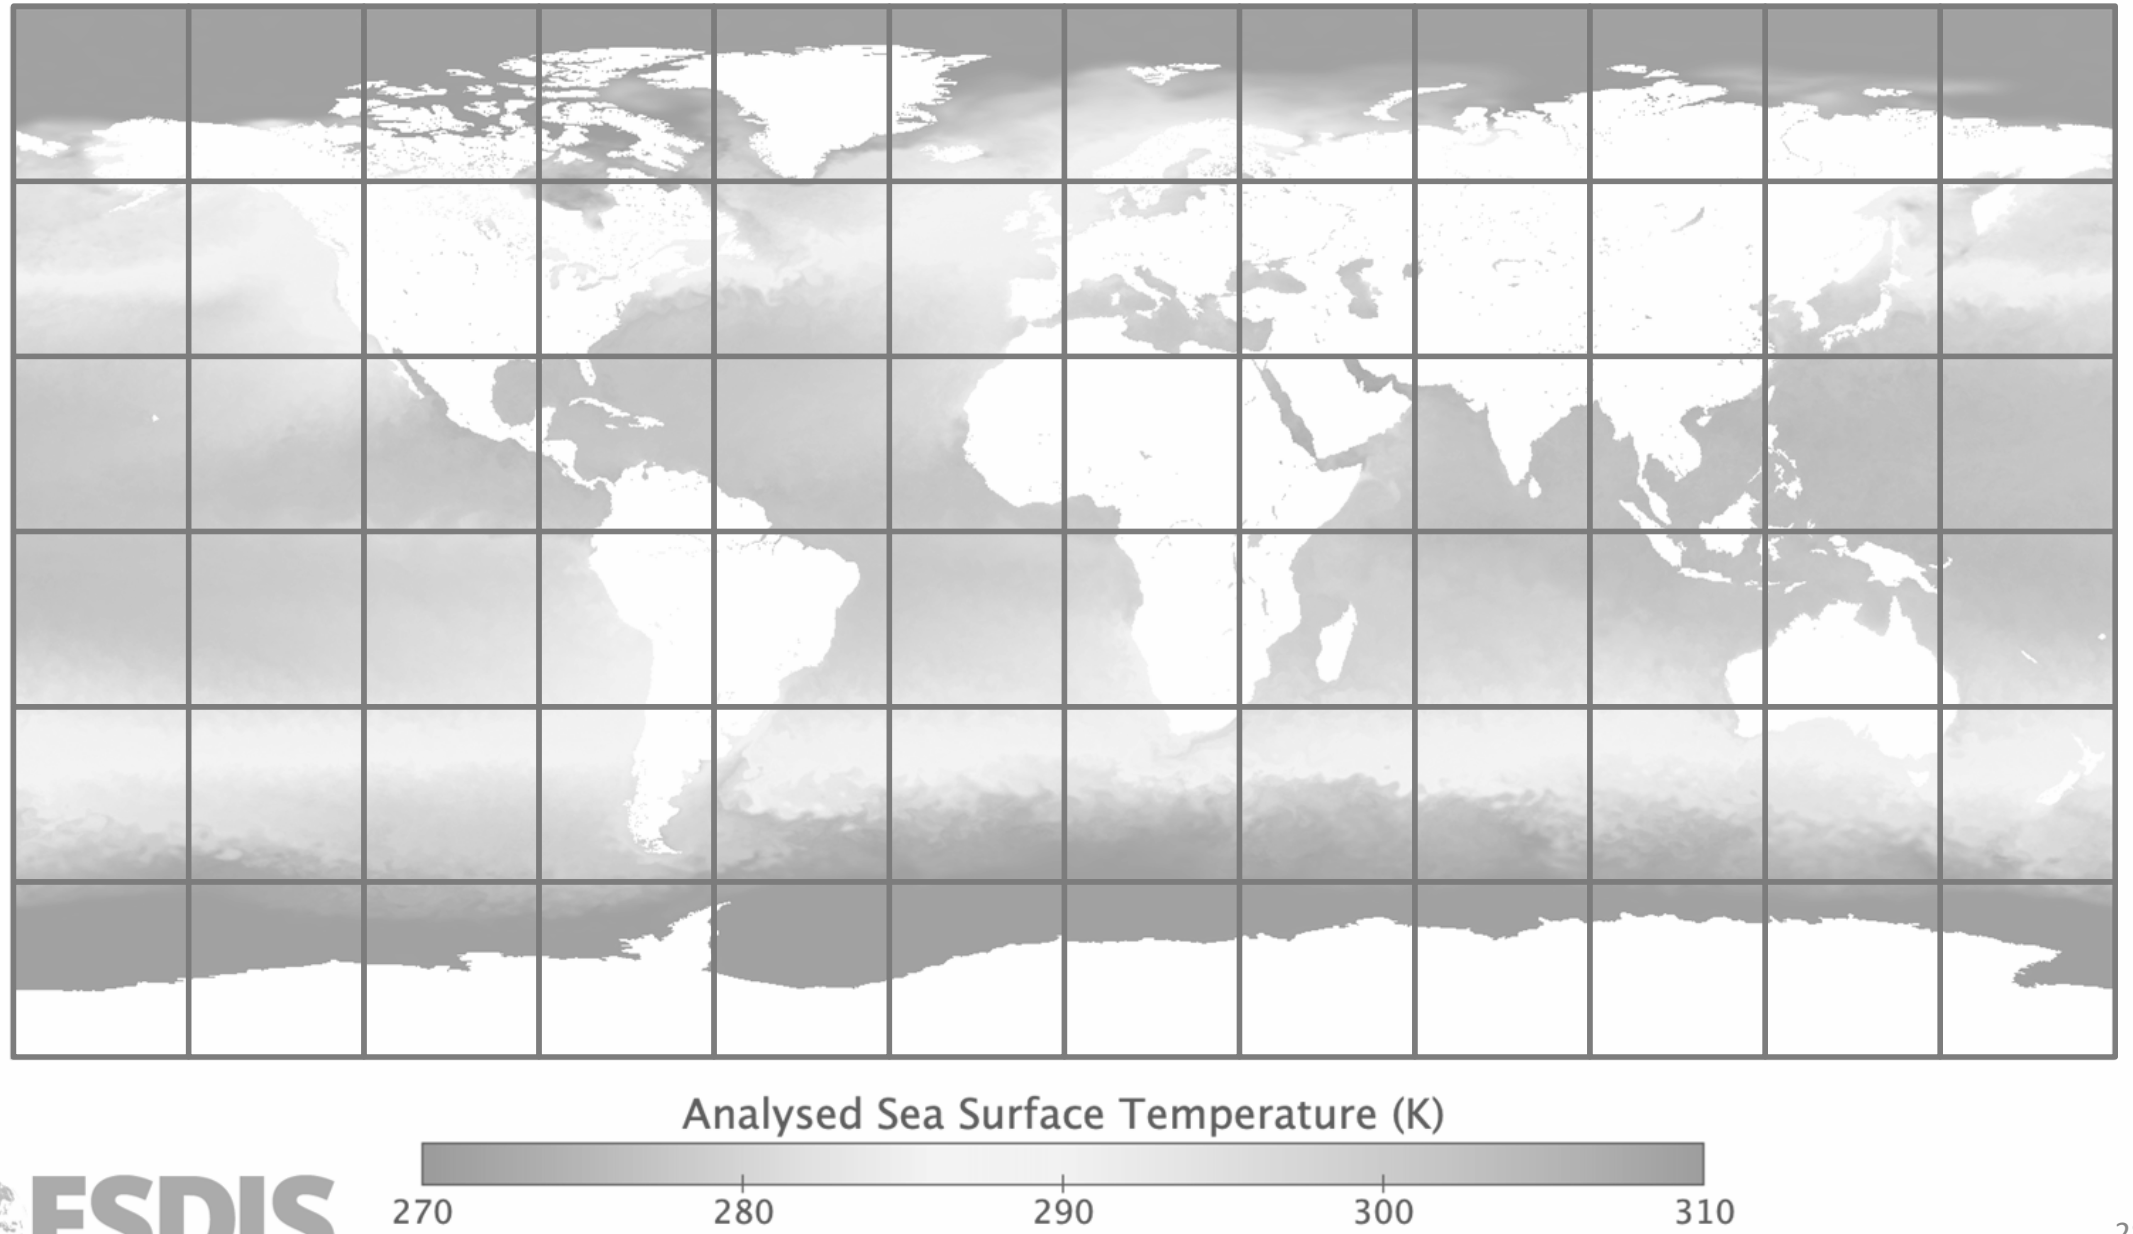

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

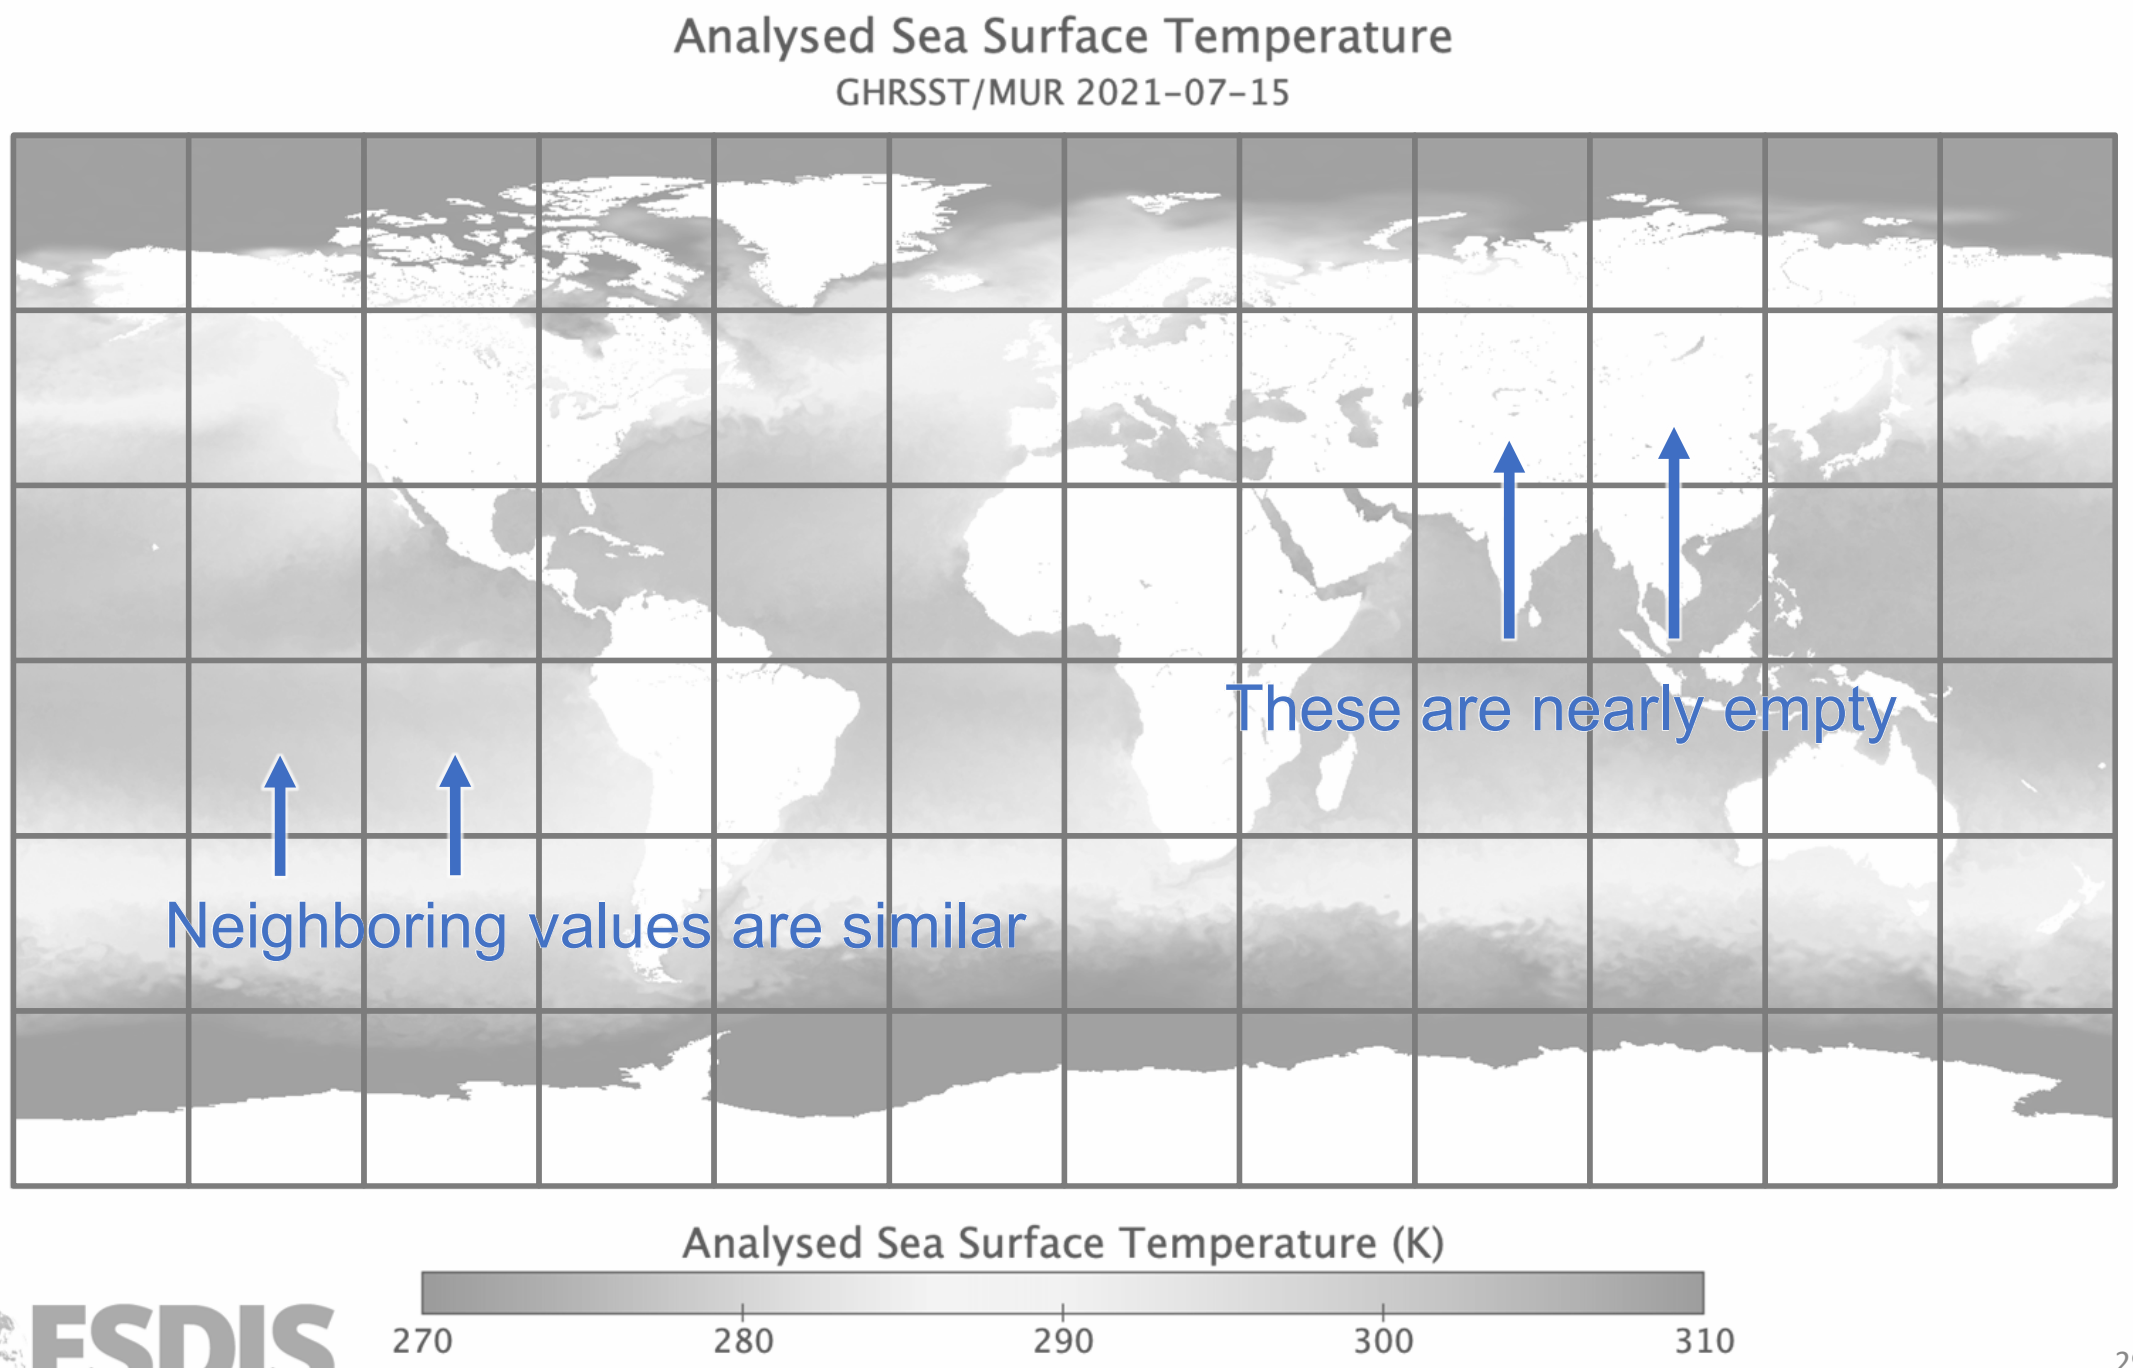

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

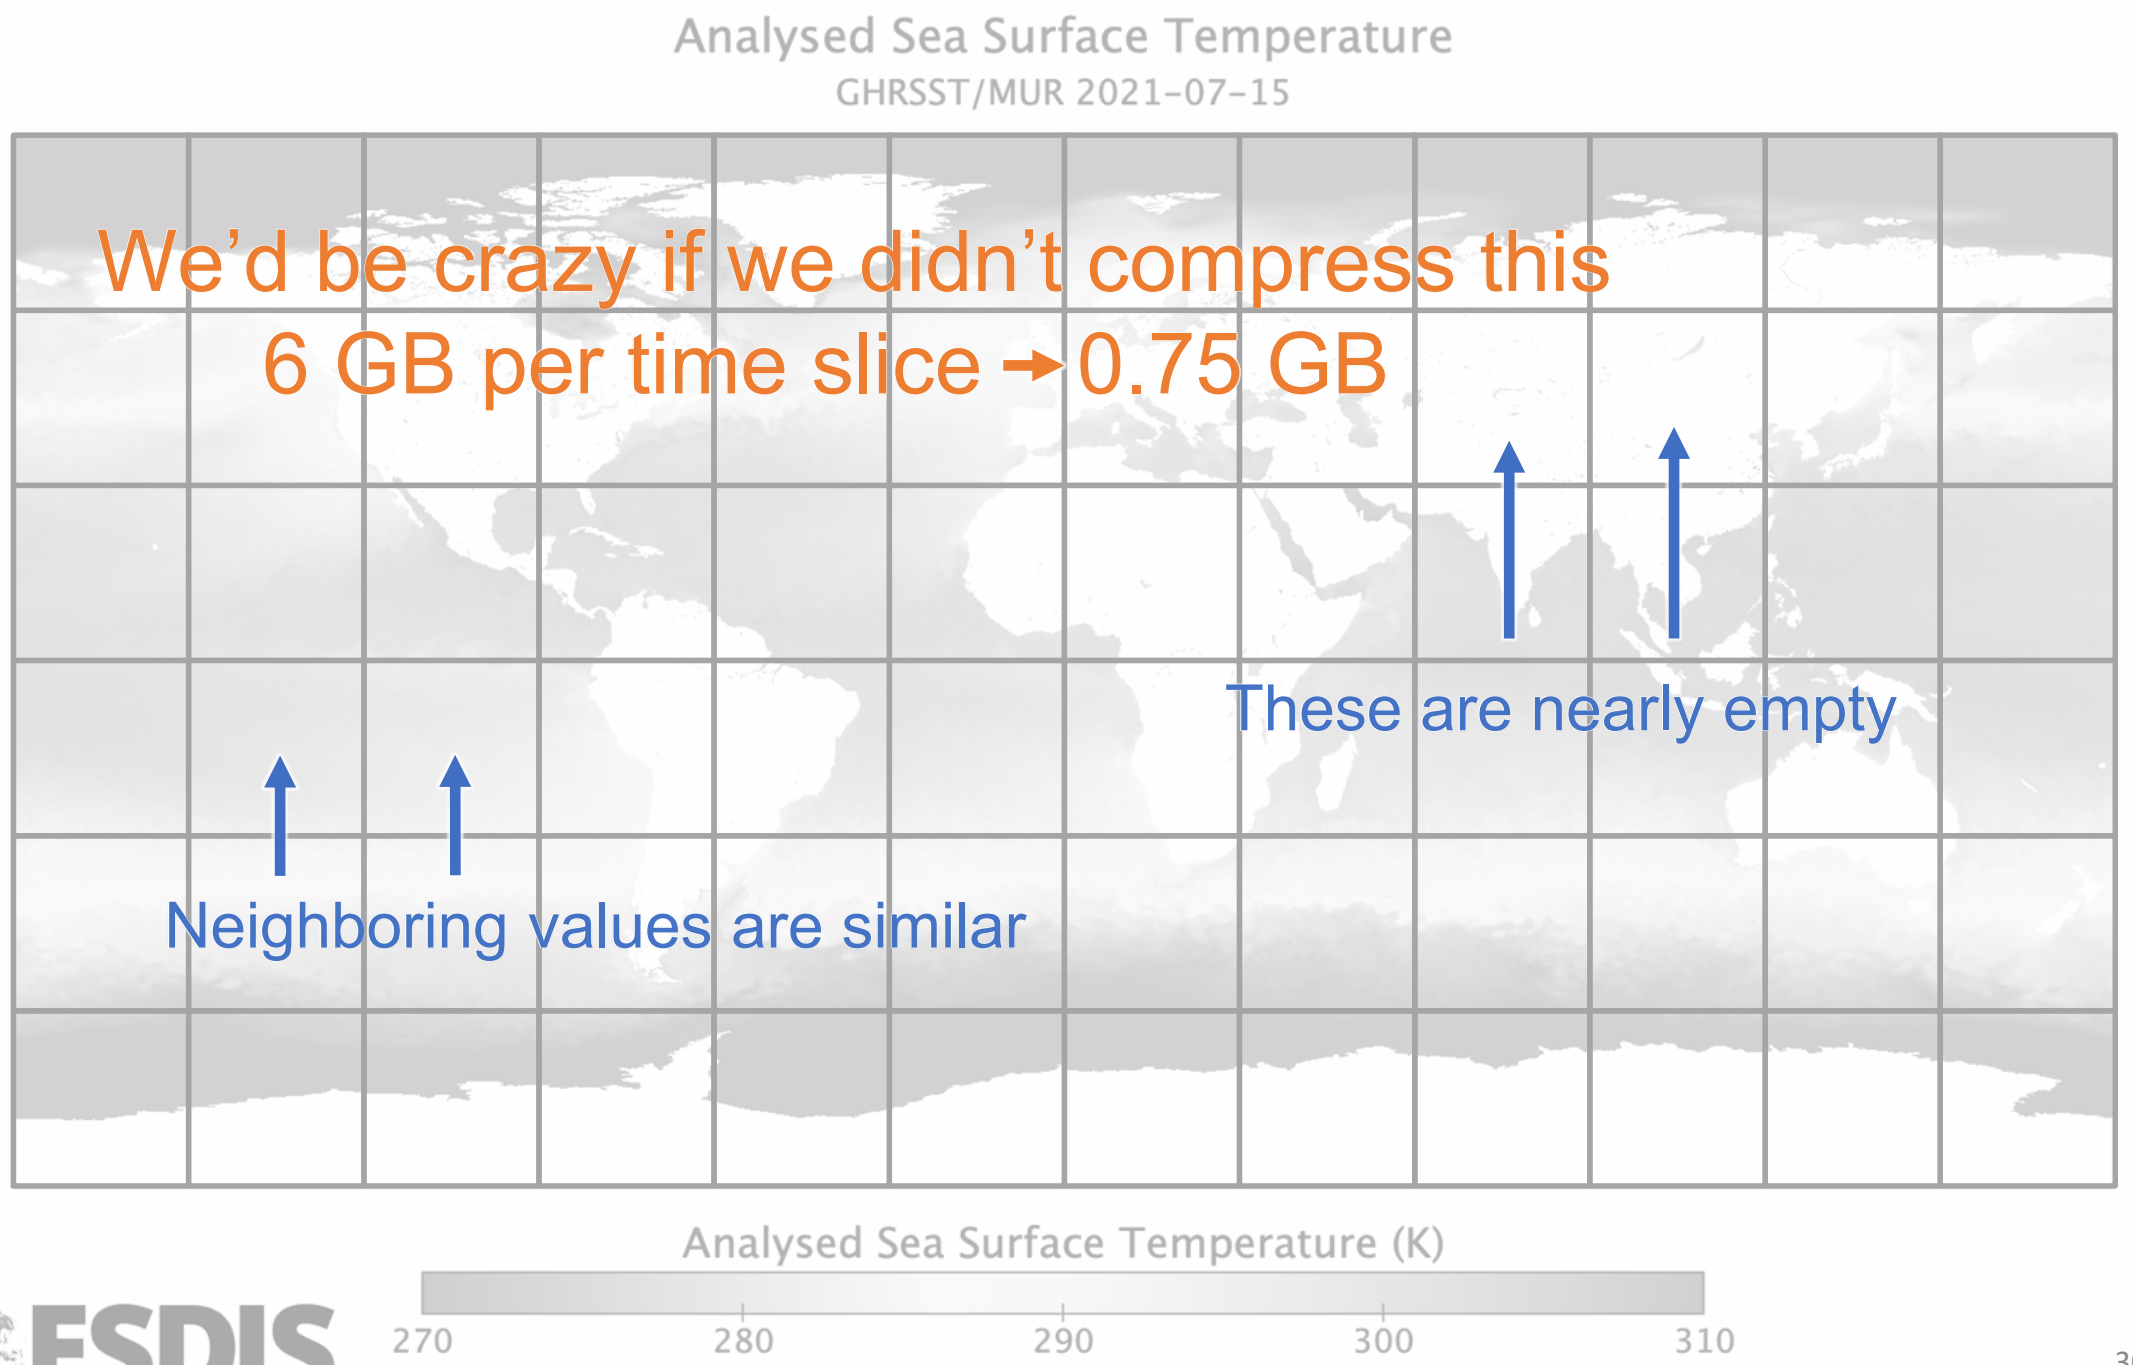

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Source: DOI 10.5067/GHGMR-4FJ04

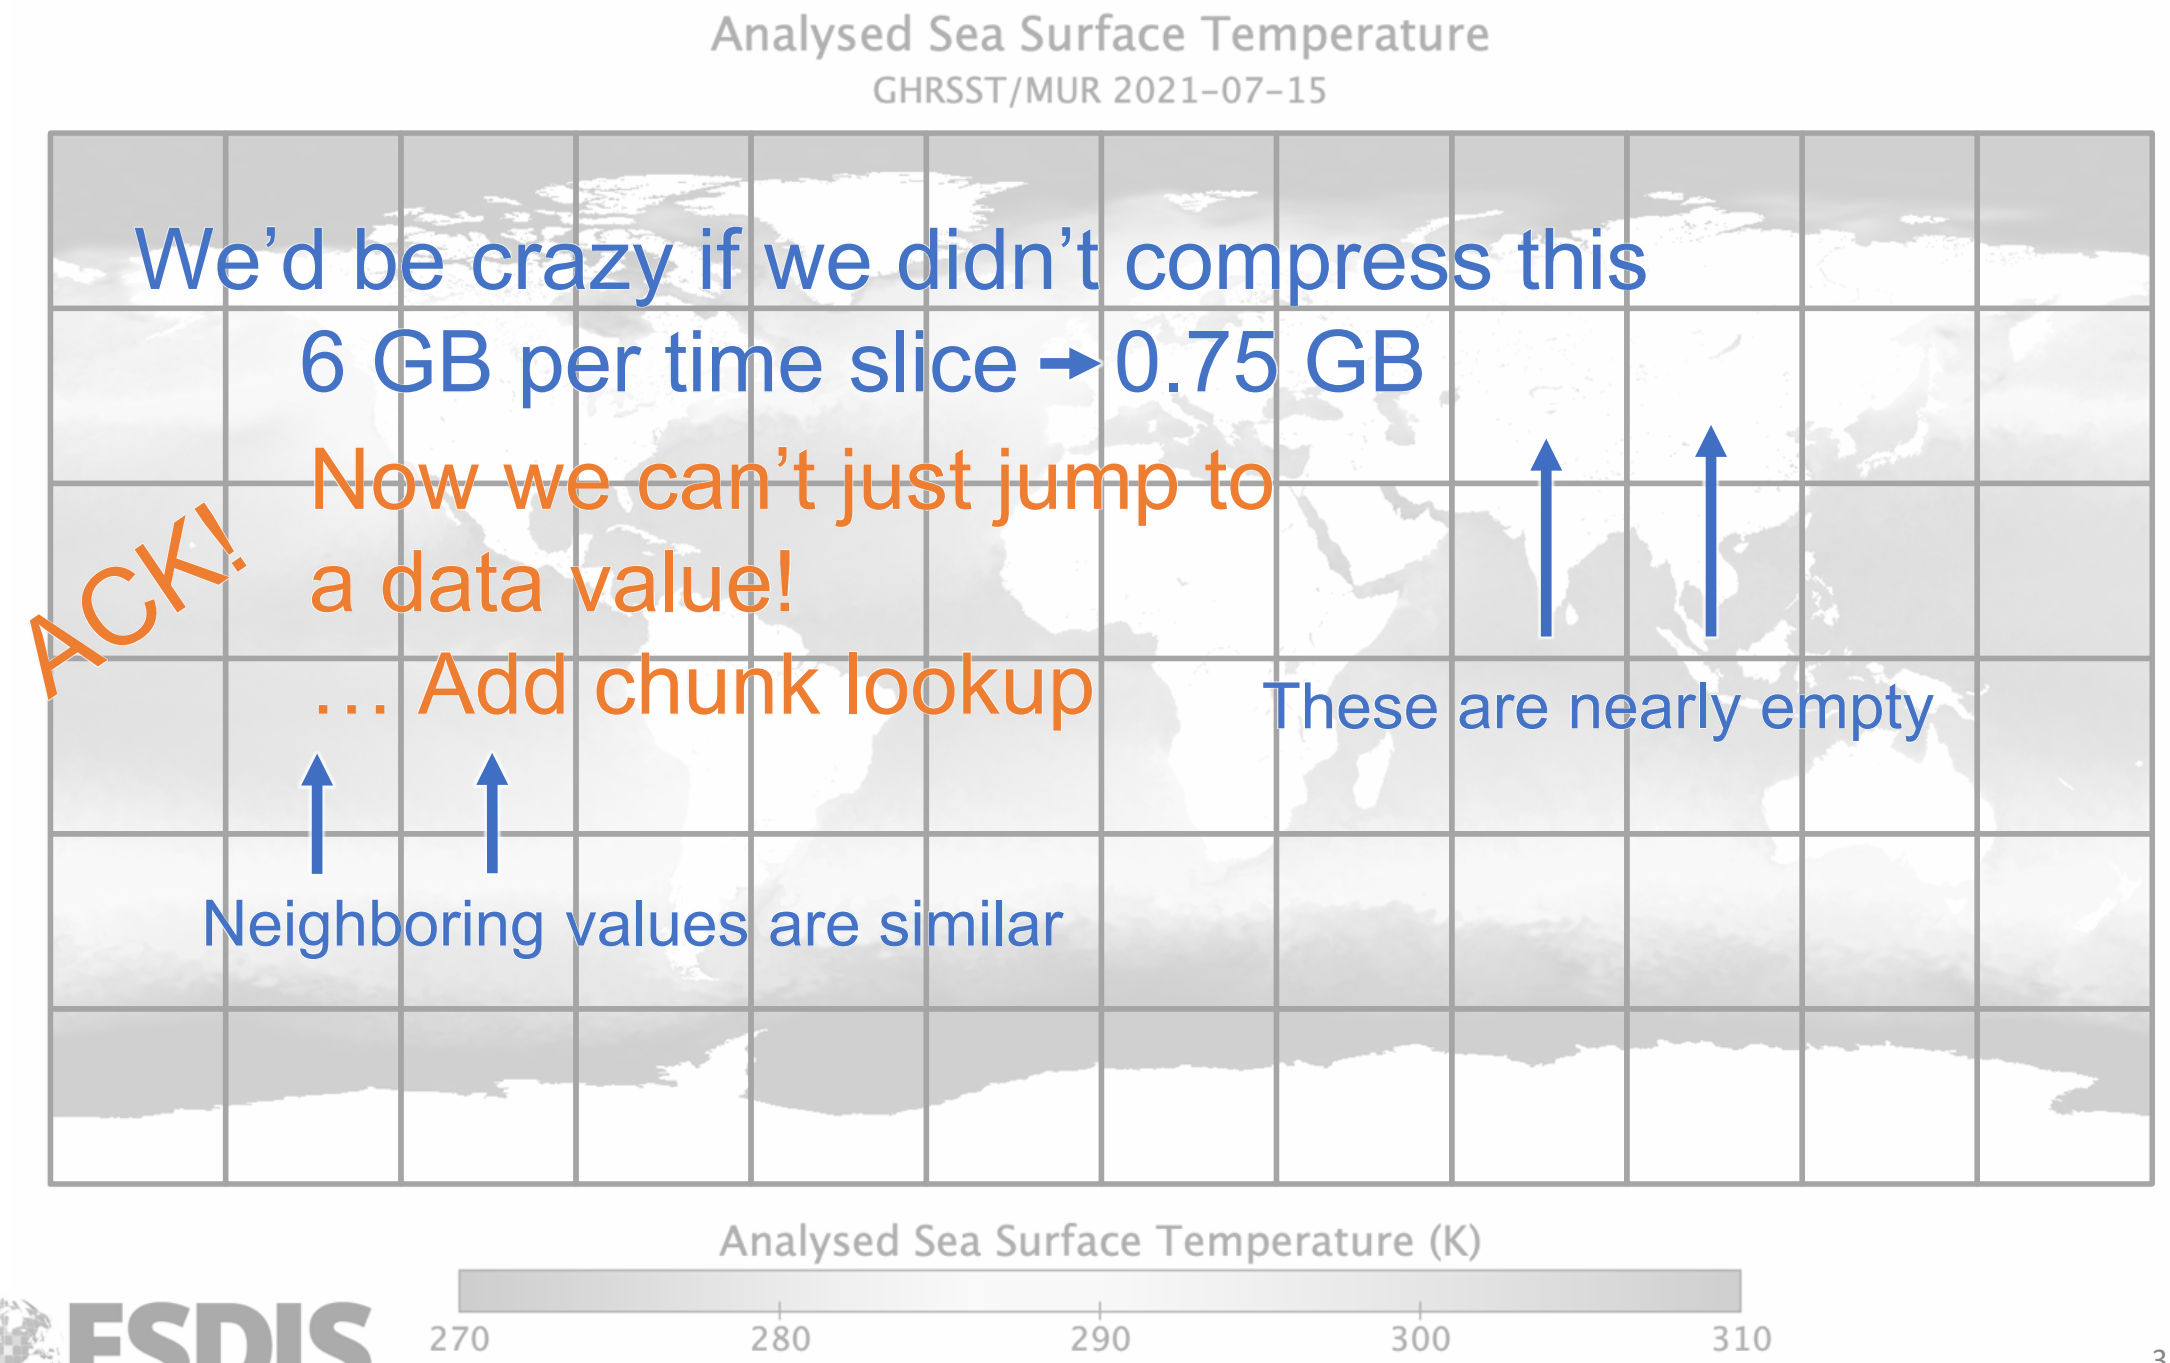

Source: DOI 10.5067/GHGMR-4FJ04

We're done!

Access-optimized multidimensional array data:

- Put all the small metadata stuff in a predictable place so it can all be read quickly, ideally in one read
- "Chunk" the data with a size and shape that balances number of reads and size of reads for most users (Hard!)
- Create an index of where to find the chunks that can be read quickly, ideally in one read

ESDIS
EARTH SCIENCE DATA AND INFORMATION SYSTEM

Solved problem?!

Not even close! We need more engagement on:

- What does an optimized archival copy look like?
- What is the role of the archive system in optimization?
- Can we improve current archive formats? (Seeking collab!)
- How do we add new data while users are reading older data?
- What about swath data? (Seeking collab!)
- How do we find the best chunk strategy?

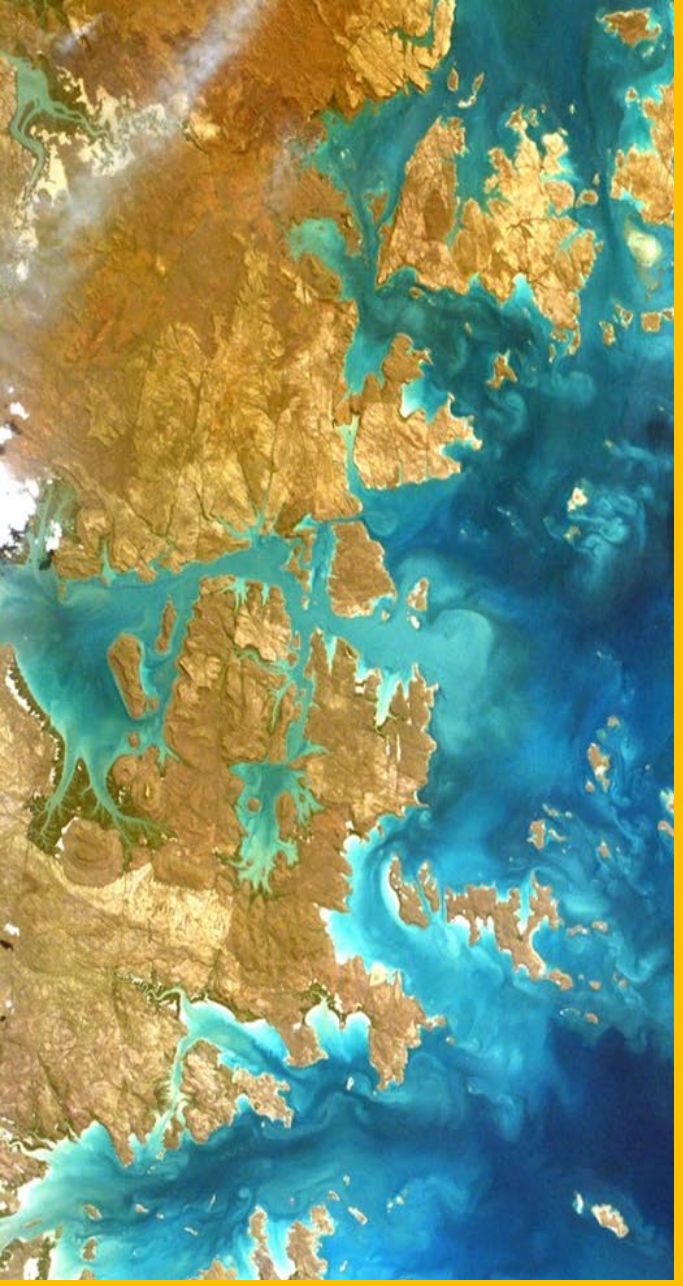

Thank you!

Patrick Quinn
ESDIS Architect
patrick.m.quinn@nasa.gov