Temporal aggregations

- Statistical mosaics:
  - Filter images and calculate a stack statistic
    - Filter collection (DOY, Cloud Cover, …)
    - Mask clouds (FMASK, NDCI, …)
    - Calculate a statistic (mean, median, etc.)
    - https://code.earthengine.google.com/47f2685d752c1614fe6b842942c7e3dc
- Best available pixel (BAP) mosaics:
  - Sort each pixel so the best is at the top
    - Distance to clouds and shadows masks
    - Atmospheric Opacity
    - Day of the year, …
    - https://code.earthengine.google.com/e27240a92ecf64bbadf8a082b91c711c?hideCode=true

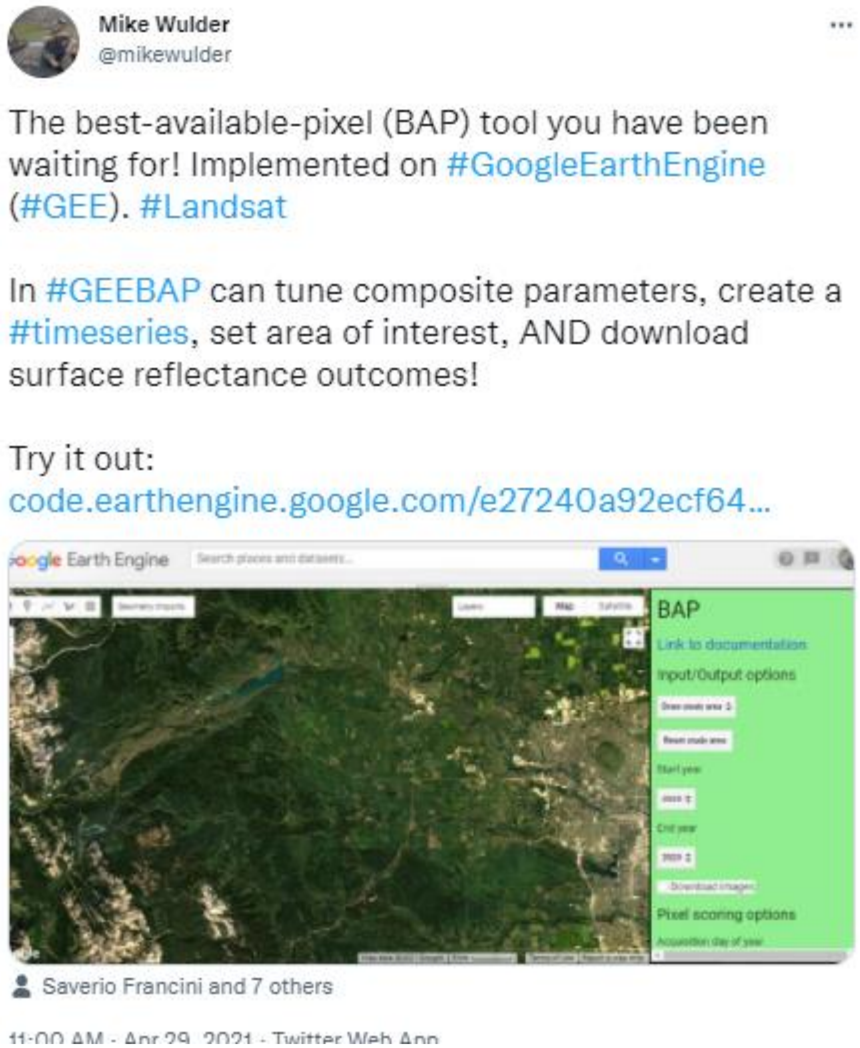

Gap filling:
Impute NoData

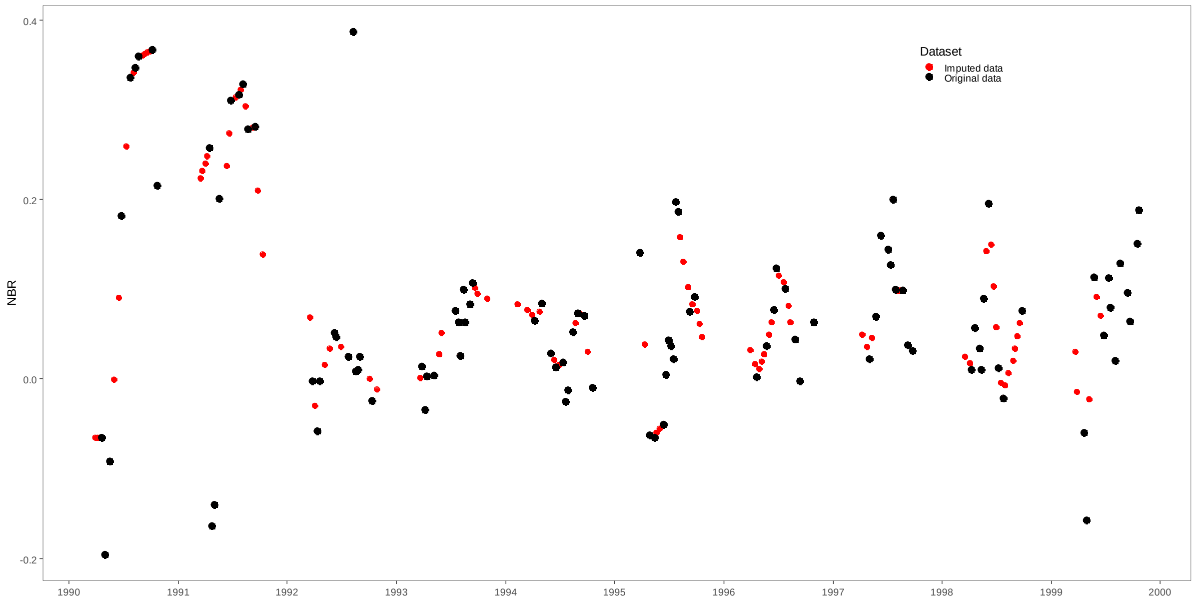

Detrend a timeseries

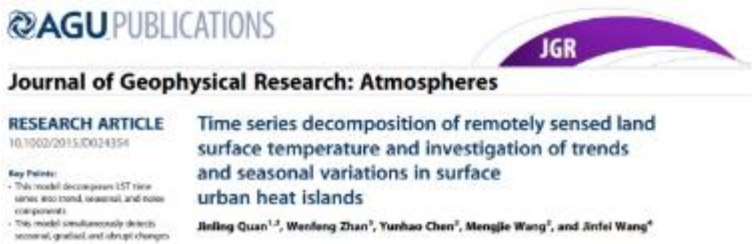
AGU PUBLICATIONS
JGR
Journal of Geophysical Research: Atmospheres
RESEARCH ARTICLE
Time series decomposition of remotely sensed land surface temperature and investigation of trends and seasonal variations in surface urban heat islands

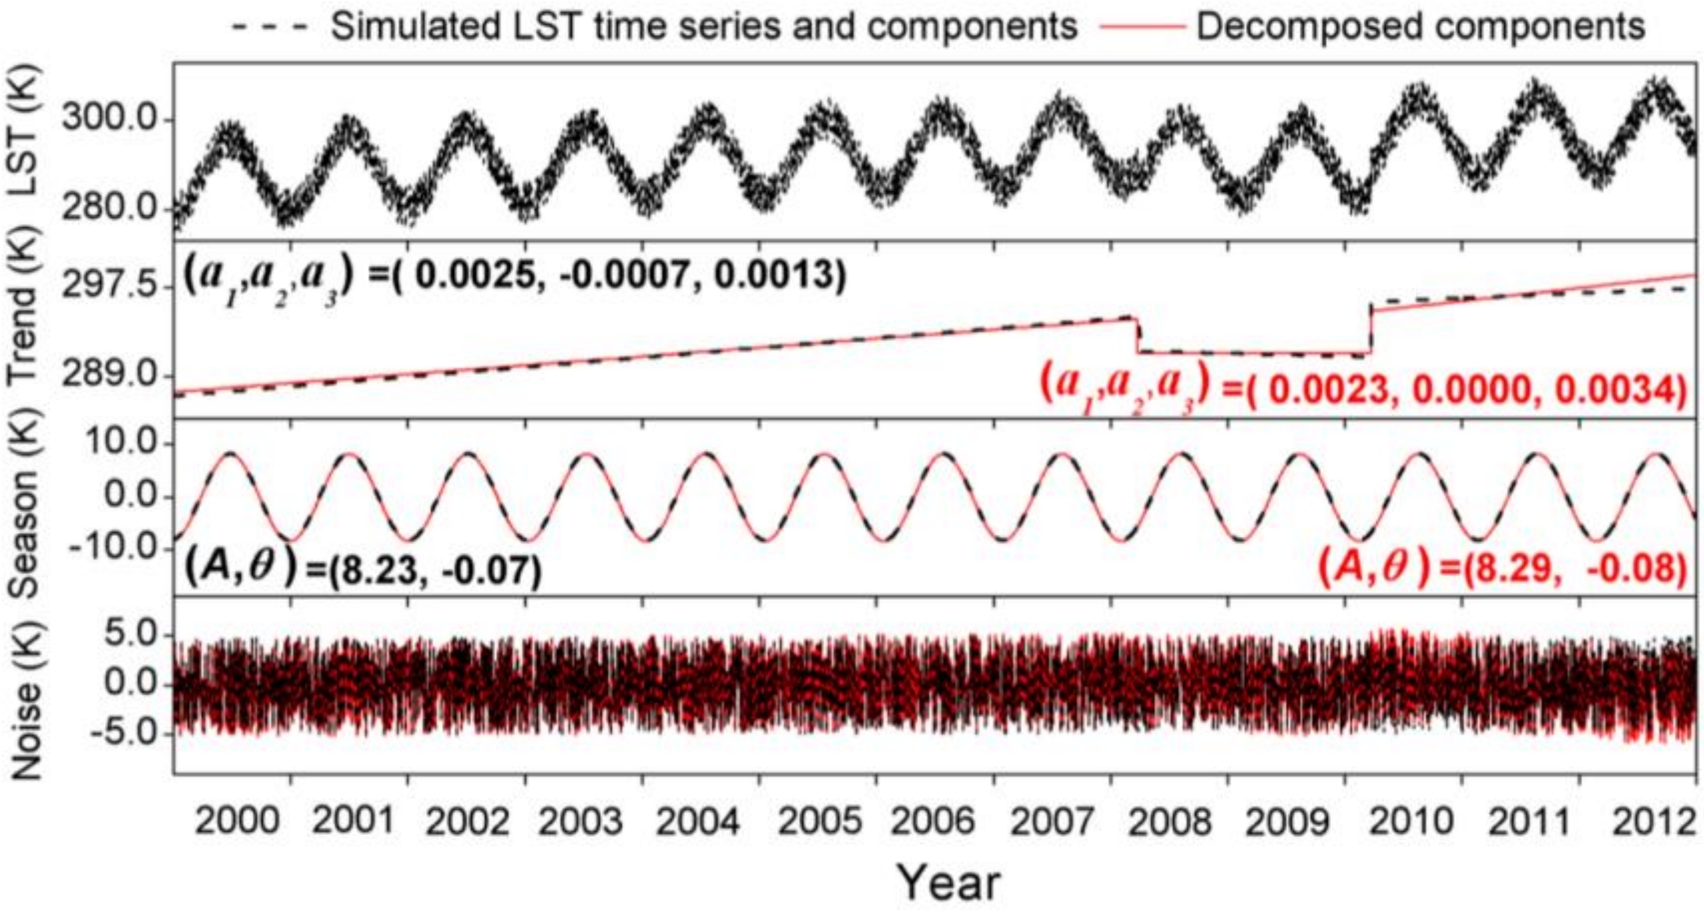

https://machinelearningmastery.com/time-series-trends-in-python/

Time series smoothing and interpolation

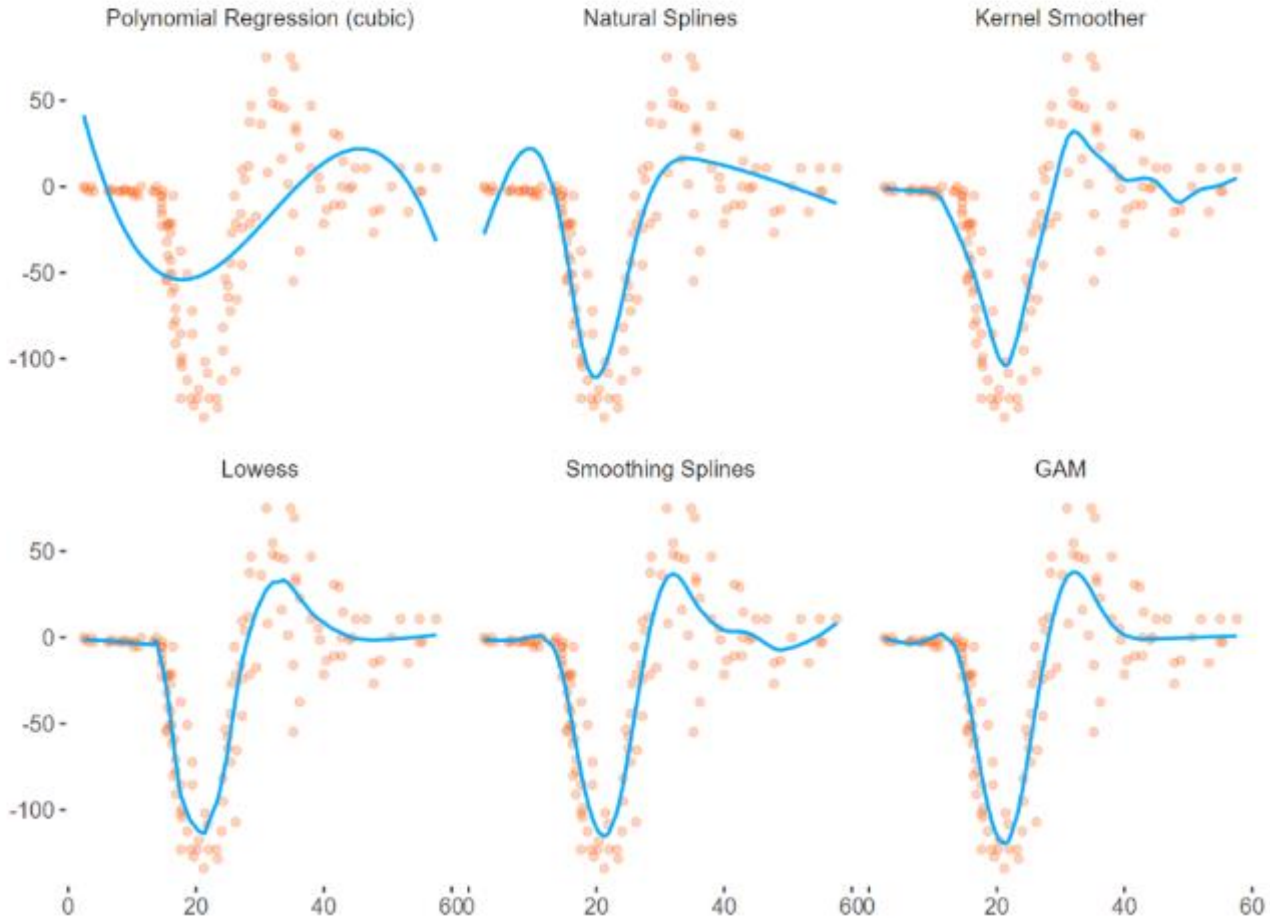

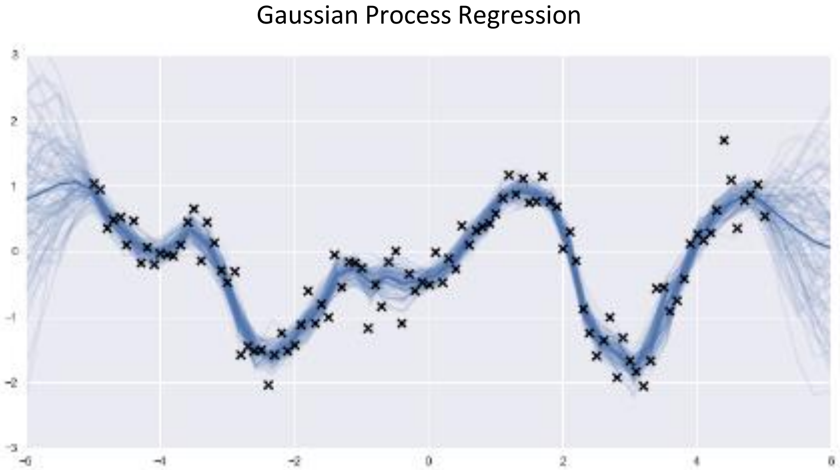

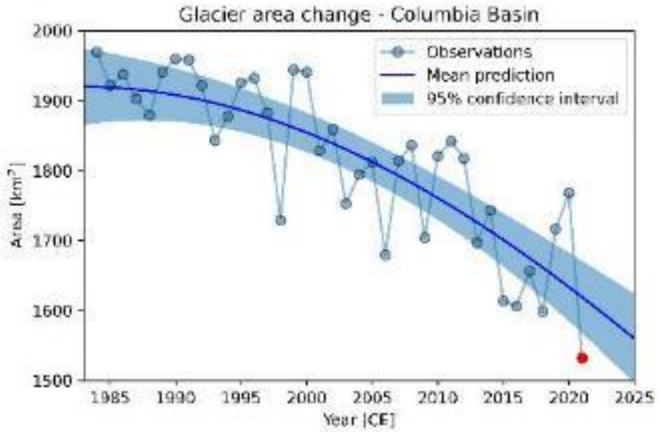

Harmonic and Seasonal Interpolation

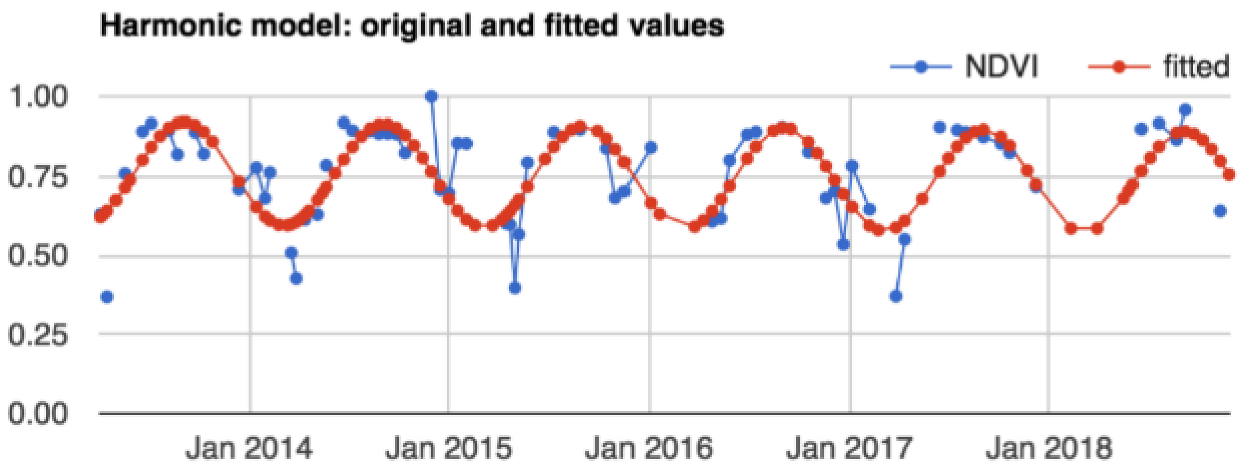

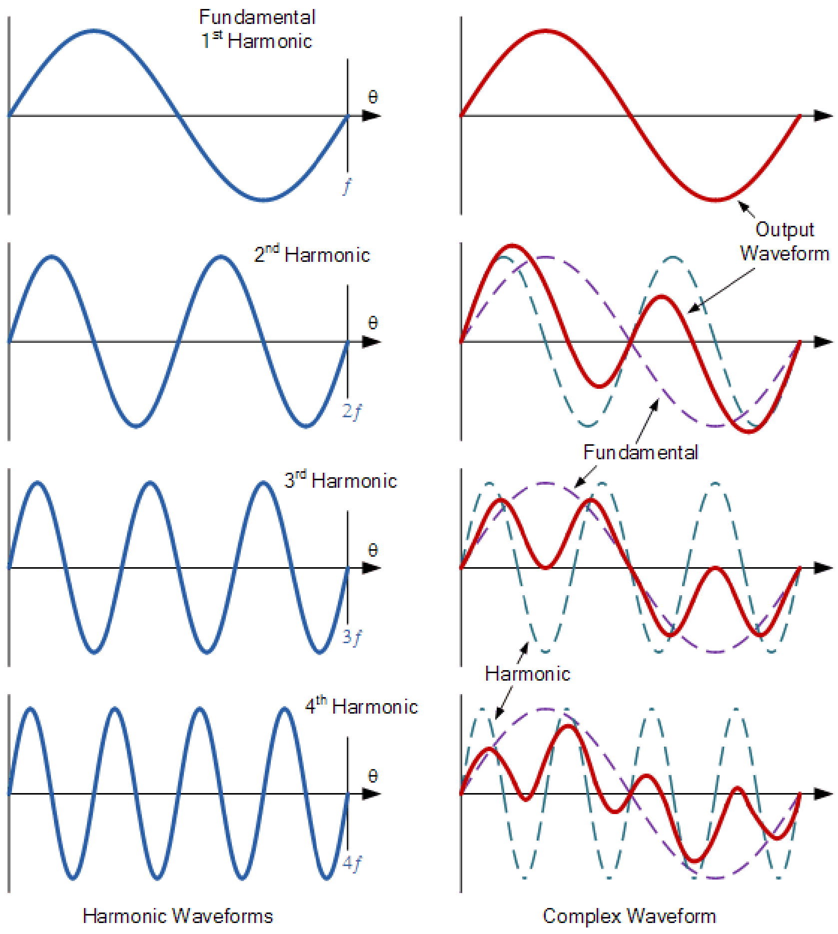

https://docs.google.com/document/d/1mNIRB90jwLuASO1JYas1kuOXCLbOoy1Z4NlV1qIXM10/edit

LandTrendR Algorithm

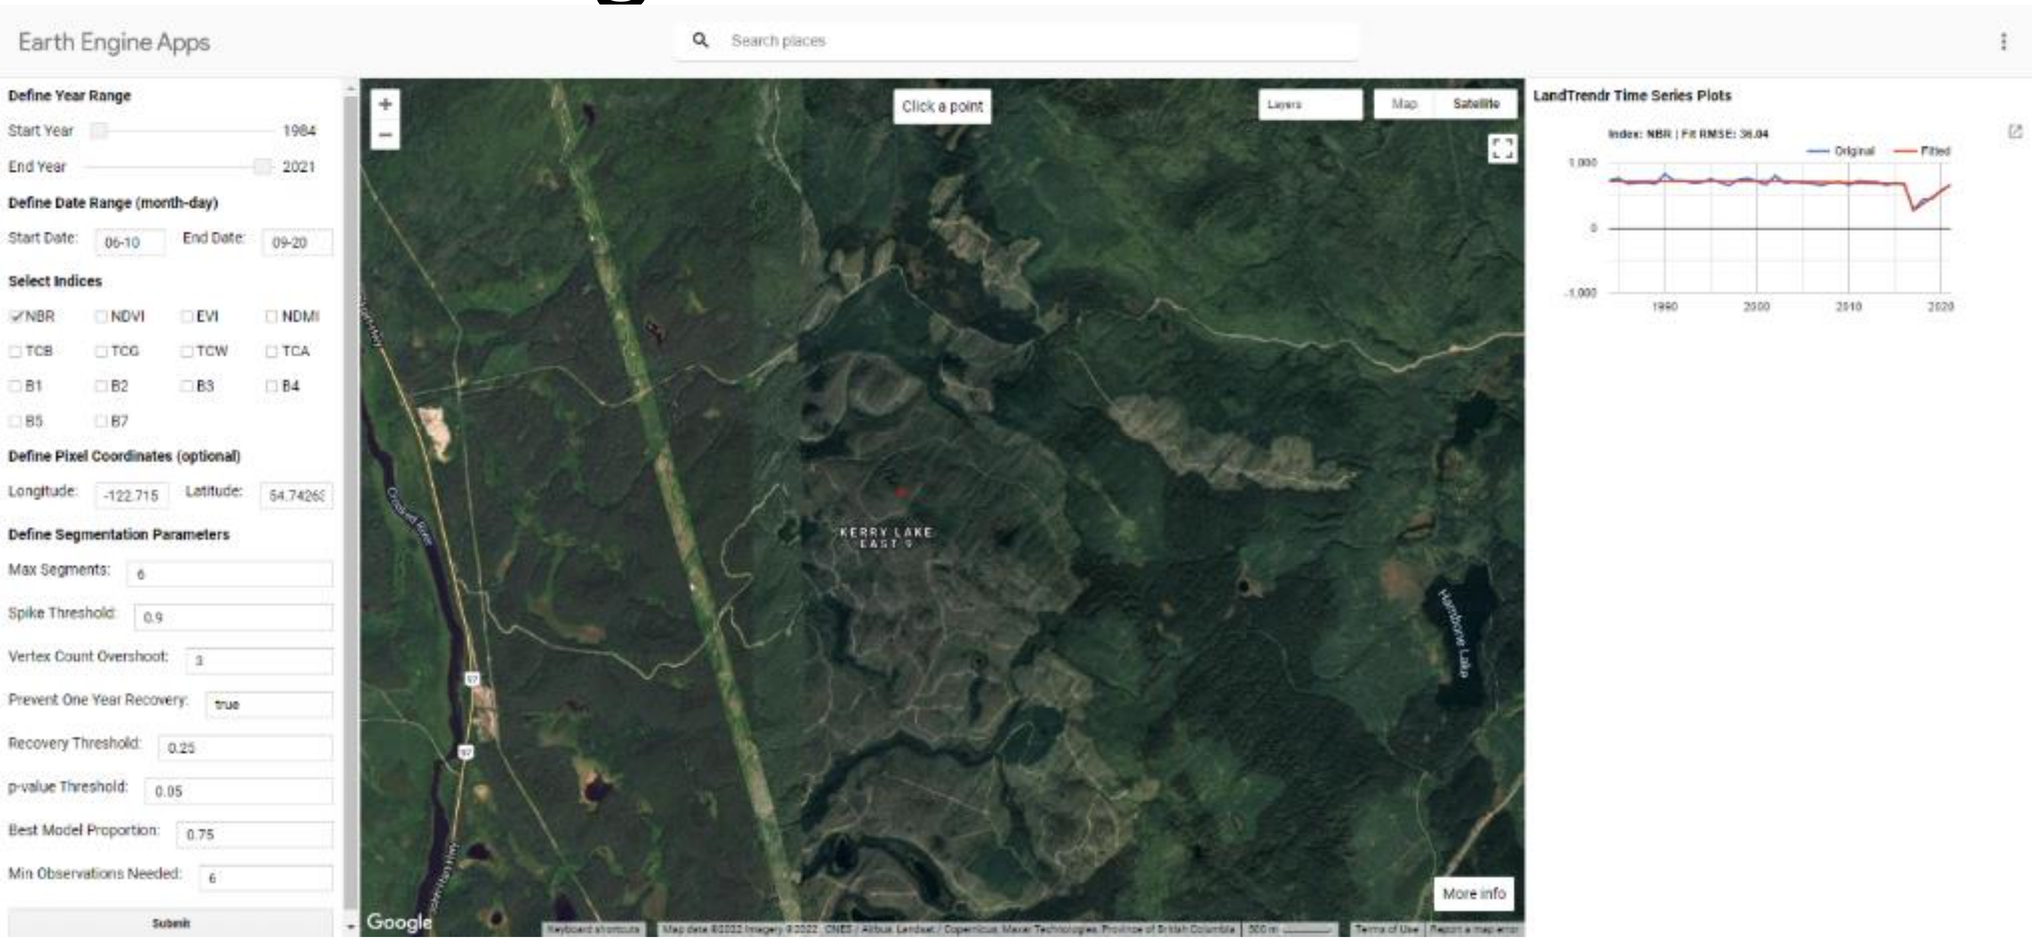

https://emapr.github.io/LT-GEE/ui-applications.html#ui-landtrendr-pixel-time-series-plotter

https://emaprlab.users.earthengine.app/view/lt-gee-pixel-time-series

Continuous Change Detection and Classification (CCDC) Algorithm

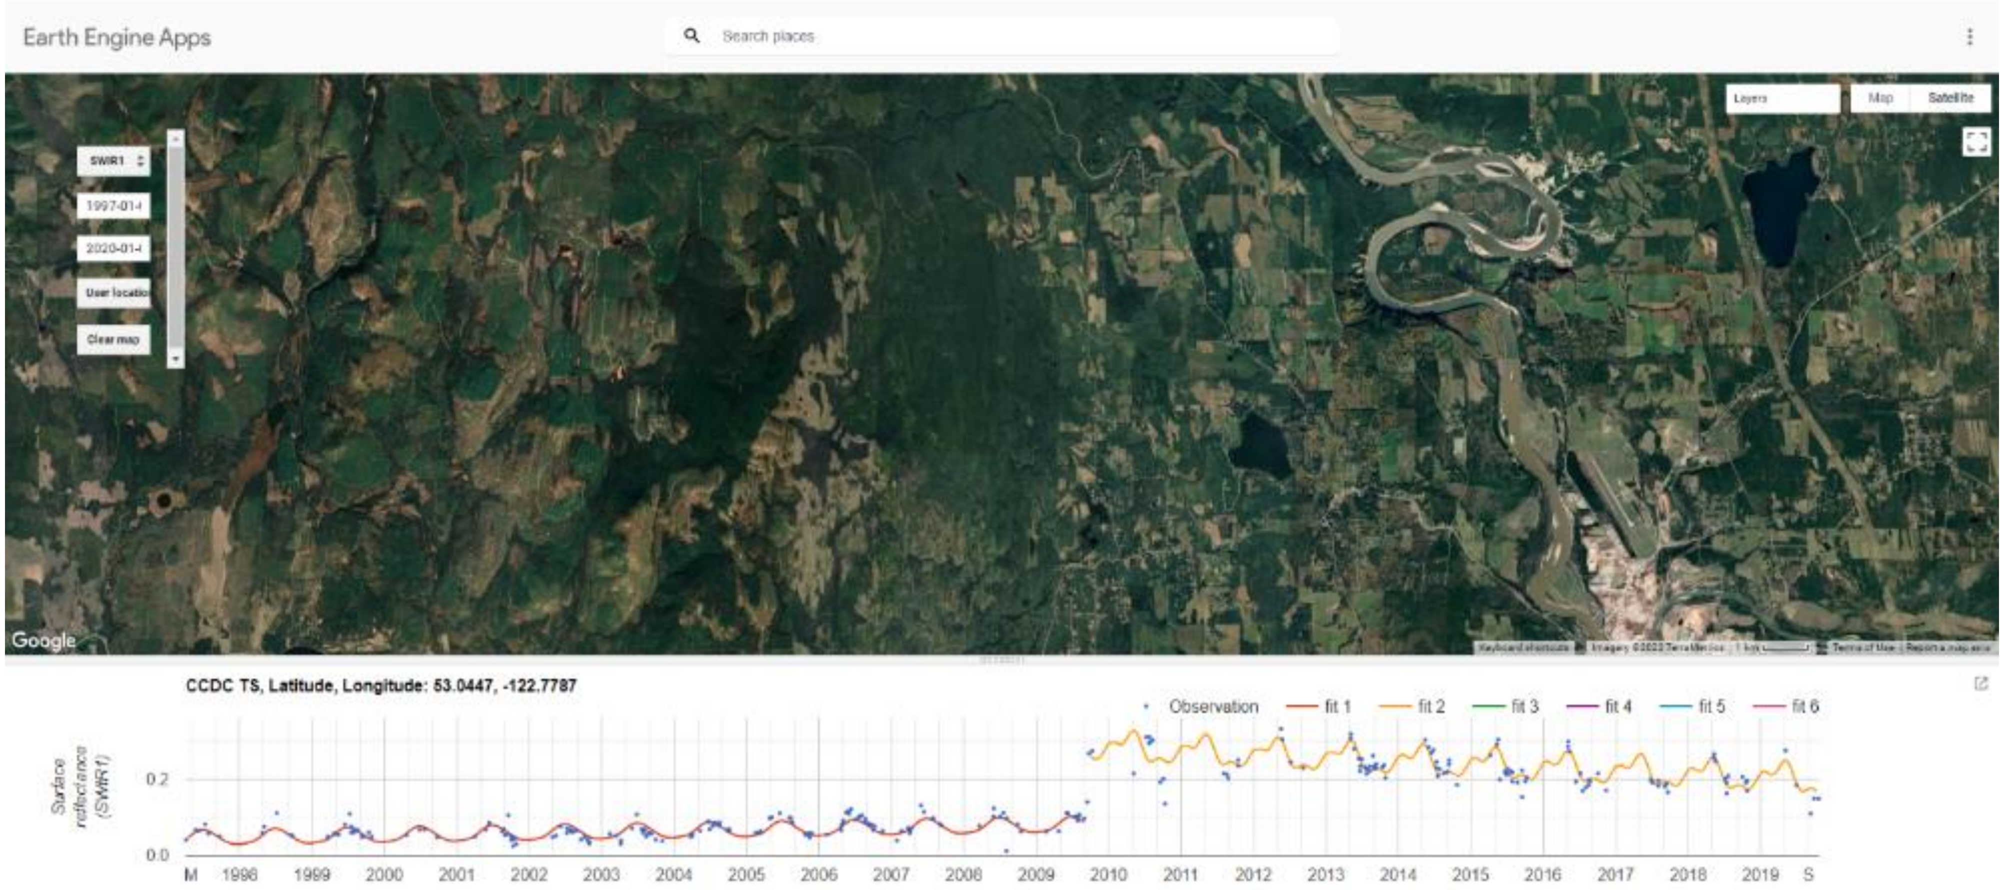

https://parevalo-bu.users.earthengine.app/view/quick-tstools

Dynamic time warping

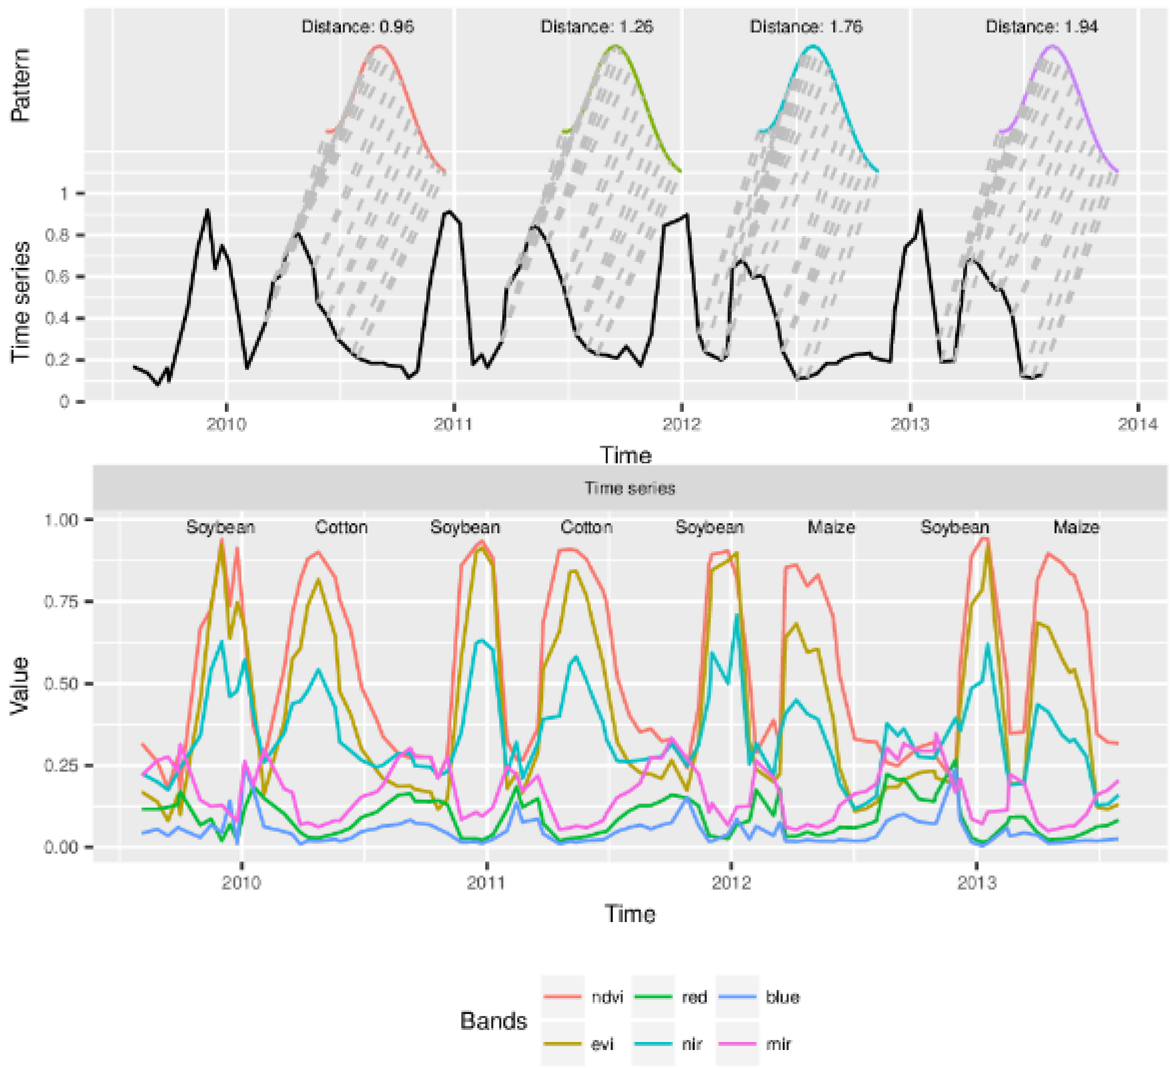

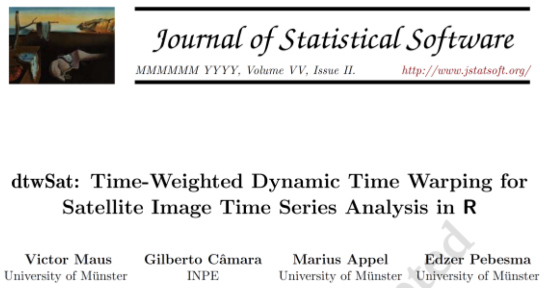
Journal of Statistical Software
MMMMMM YYYY, Volume VV, Issue II. http://www.jstatsoft.org/

dtwSat: Time-Weighted Dynamic Time Warping for Satellite Image Time Series Analysis in R

Victor Maus, University of Münster; Gilberto Câmara, INPE; Marius Appel, University of Münster; Edzer Pebesma, University of Münster

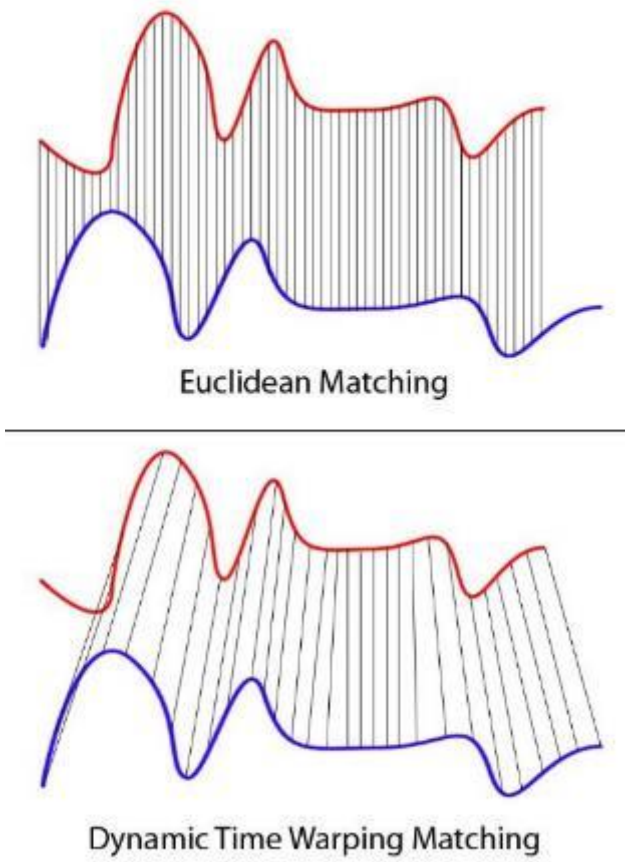

https://towardsdatascience.com/dynamic-time-warping-3933f25fcdd

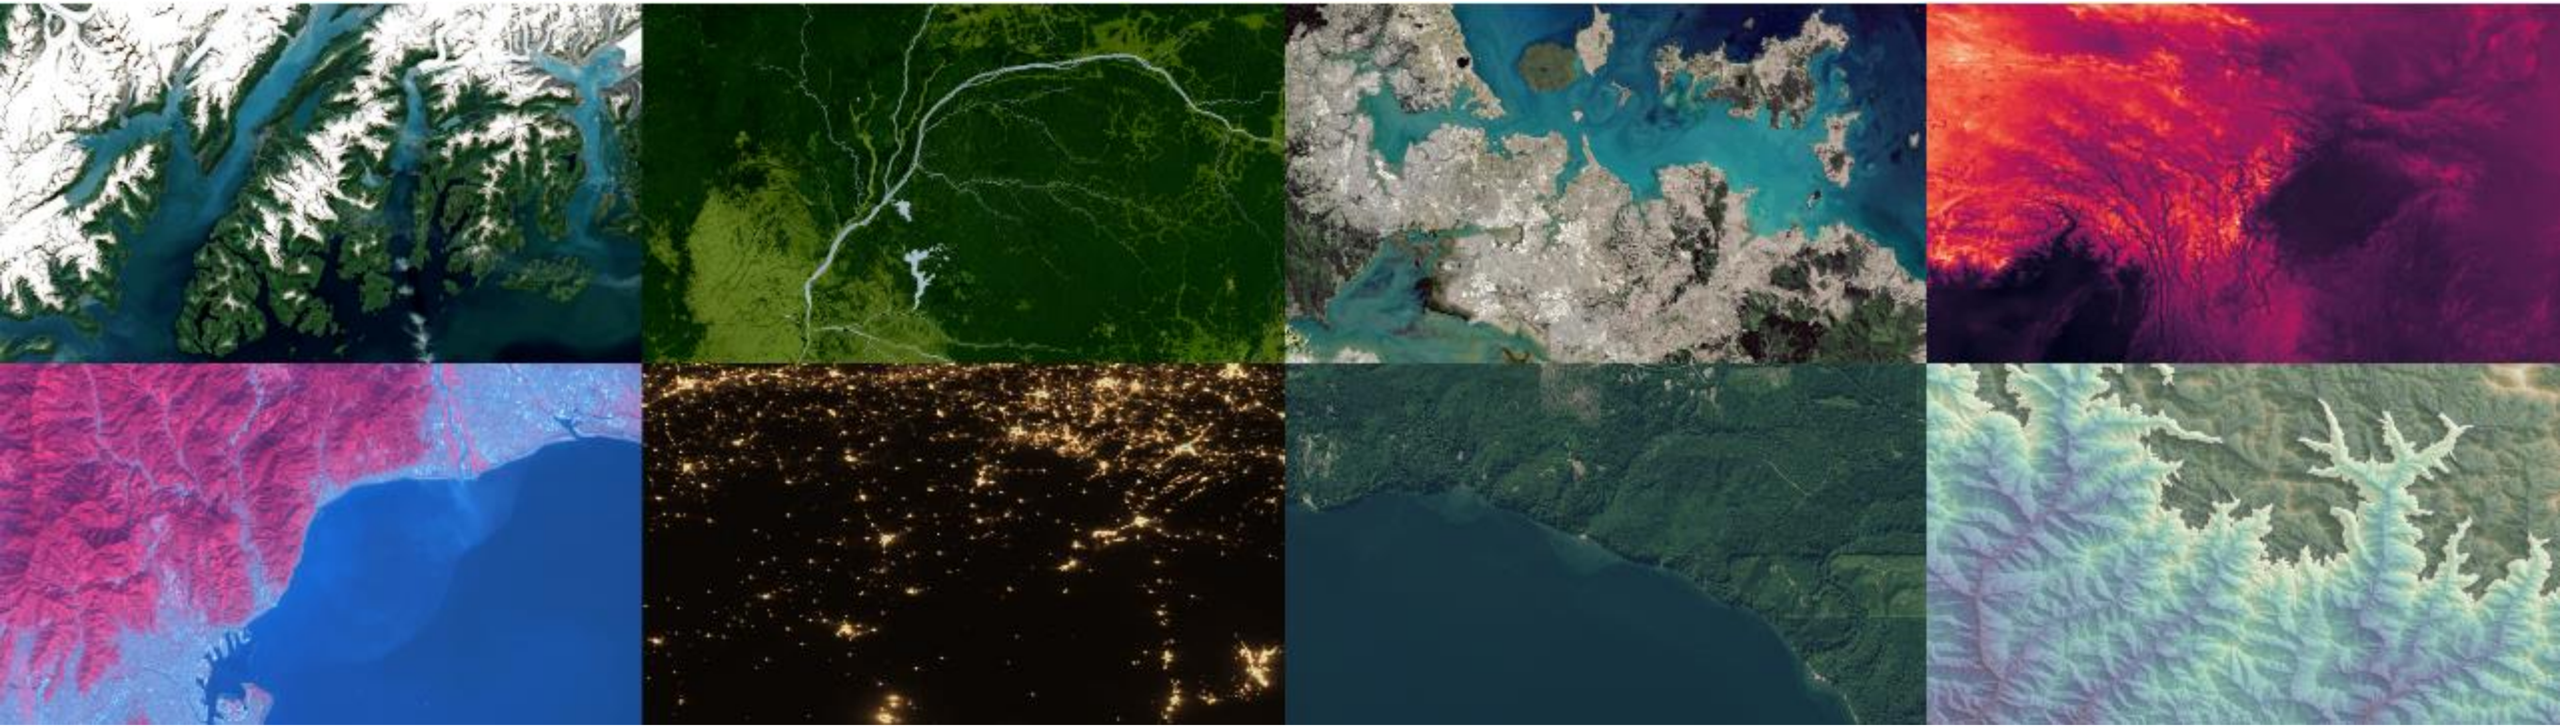

Working with `large` datasets

March 7, 2024

Alexandre.Bevington@gov.bc.ca

What are large datasets?

- More capabilities = Bigger questions

```
from sentinelsat import SentinelAPI
api = SentinelAPI('username','pw')

S1 = api.query(date=('2021-01-01T00:00:00Z','2021-12-31T23:59:59Z'),platformname='Sentinel-1',producttype='SLC',sensoroperationalmode='IW')

api.get_products_size(S1)

~ 2.07 Petabyte of Sentinel-1 IW SLC data (NOT INCLUDING GRD)

S2 = api.query(date=('2021-01-01T00:00:00Z','2021-12-31T23:59:59Z'),platformname='Sentinel-2',producttype='S2MSI1C')

api.get_products_size(S2)

~ 0.94 Petabyte of Sentinel-2 L1C dat
```

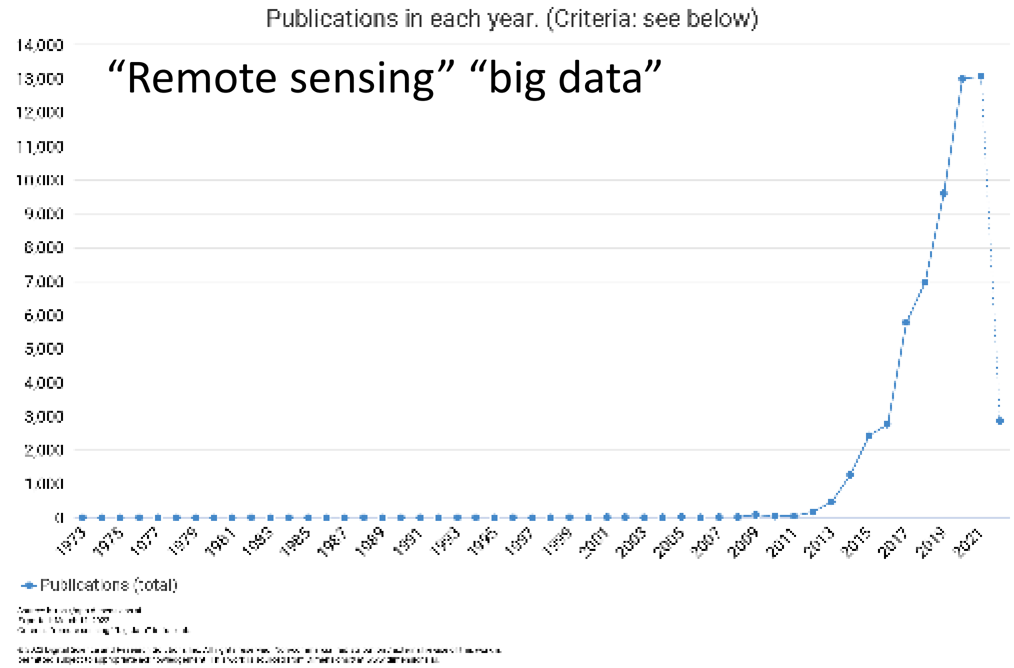

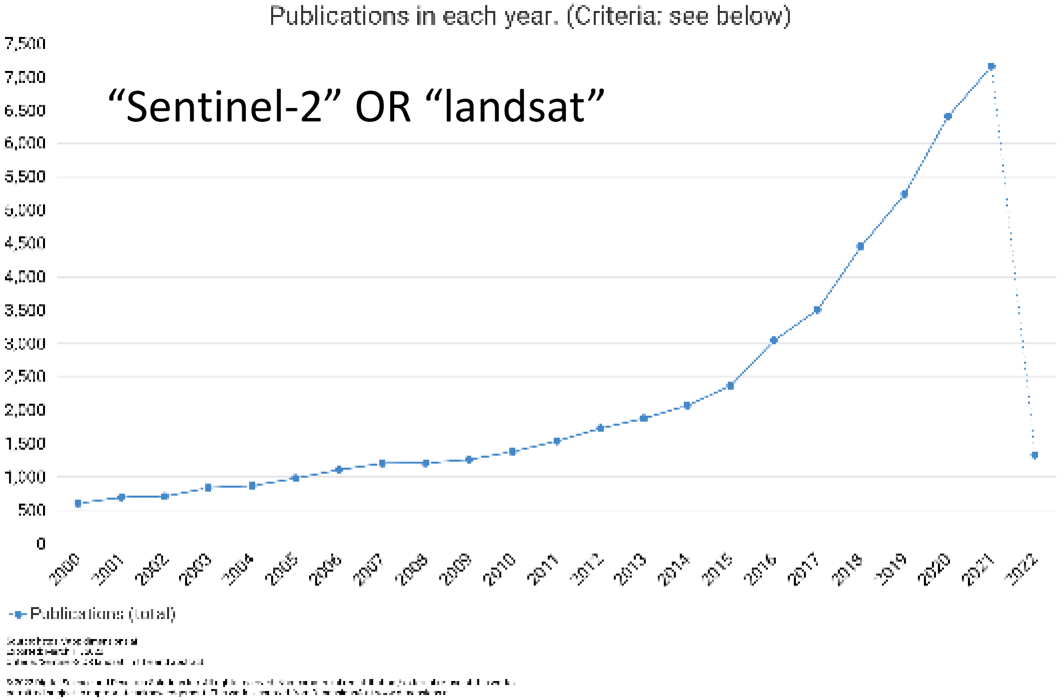

Local workflow

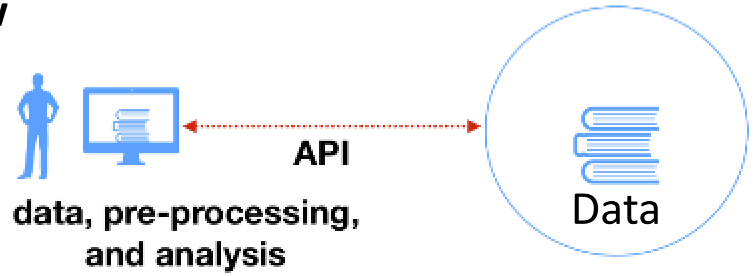

Cloud workflow

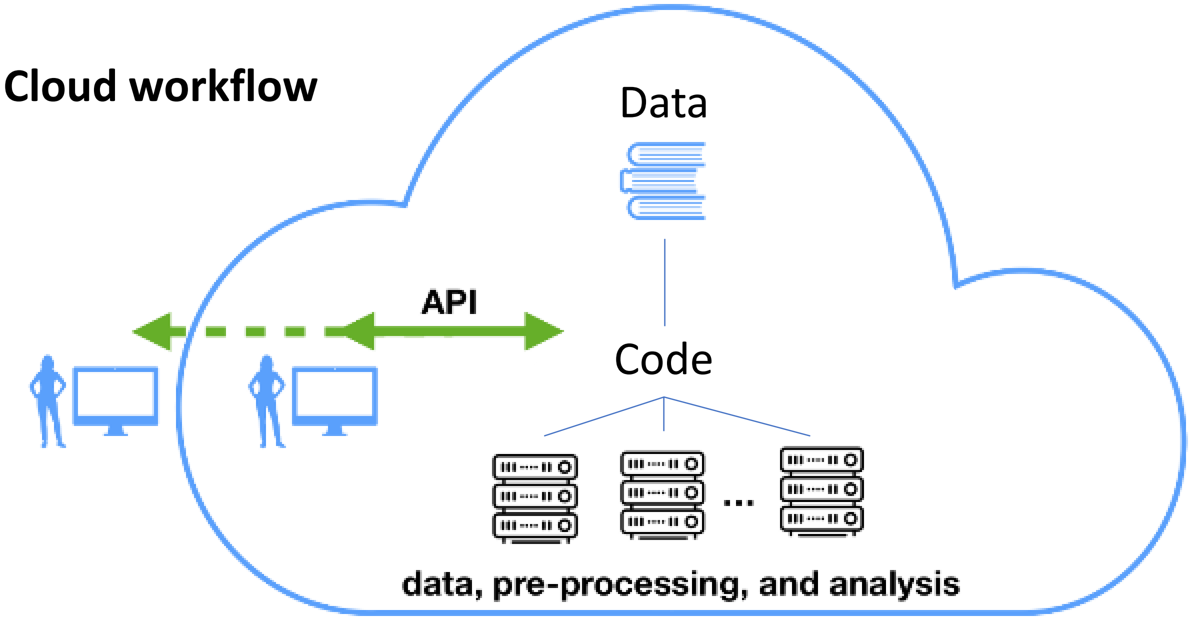

https://escience.washington.edu/uw-researchers-are-changing-the-way-we-use-satellite-observations-of-earth/

Computer speed

- **CPU – Central Processing Unit**
  Speed of processor, number of cores (How many workers)
- **RAM – Random Access Memory**
  Store working data and machine code (Multitasking)
- **SSD vs HDD – Solid State vs Hard Disk**
  Read/write speed and total storage
- **GPU - Graphics Processing Unit**
  Speeds up image visualization and processing, optimal for some tasks
- etc

- 8 bit = 256
- 10 bit = 1024
- 12 bit = 4096
- 16 bit = 65,536

Res (m)	Pixels	8 bit	16 bit
1	1 trillion	1 TB	2 TB
10	10 billion	10 GB	20 GB
100	100 million	100 MB	200 MB
1000	1 million	1 MB	1 MB

- **Serial computing**
  - A problem is broken into instructions
  - Executed sequentially on a single processor
  - One instruction executed at a time
- **Parallel Computing**
  - A problem is broken into parts
  - Each part is broken into instructions
  - Execute simultaneously on different processors
  - Requires orchestration, sometimes not worth it
- **Hyper-threading**
  - Better task scheduling
  - Minimizes processor downtime
  - Works for both serial and parallel computing
  - Not equivalent to more cores

Serial computing

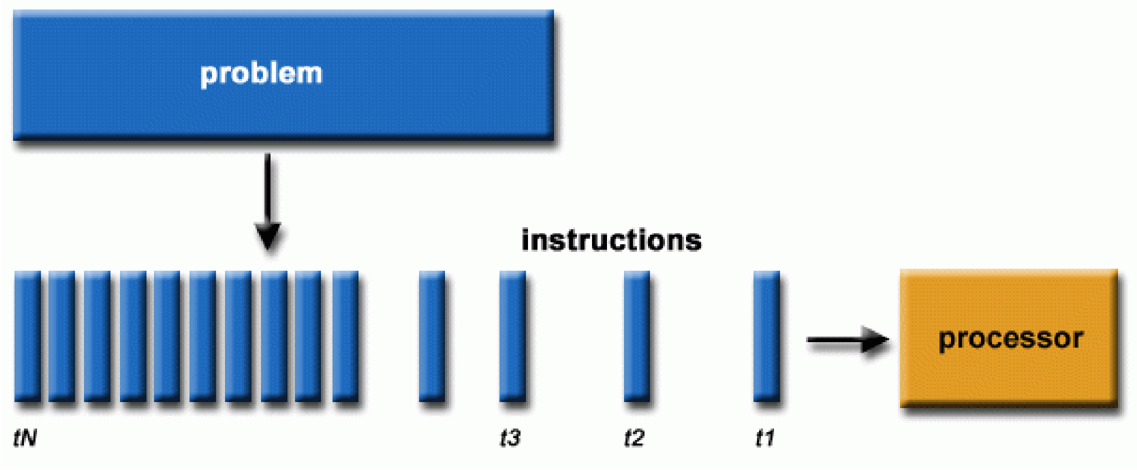

Parallel Computing

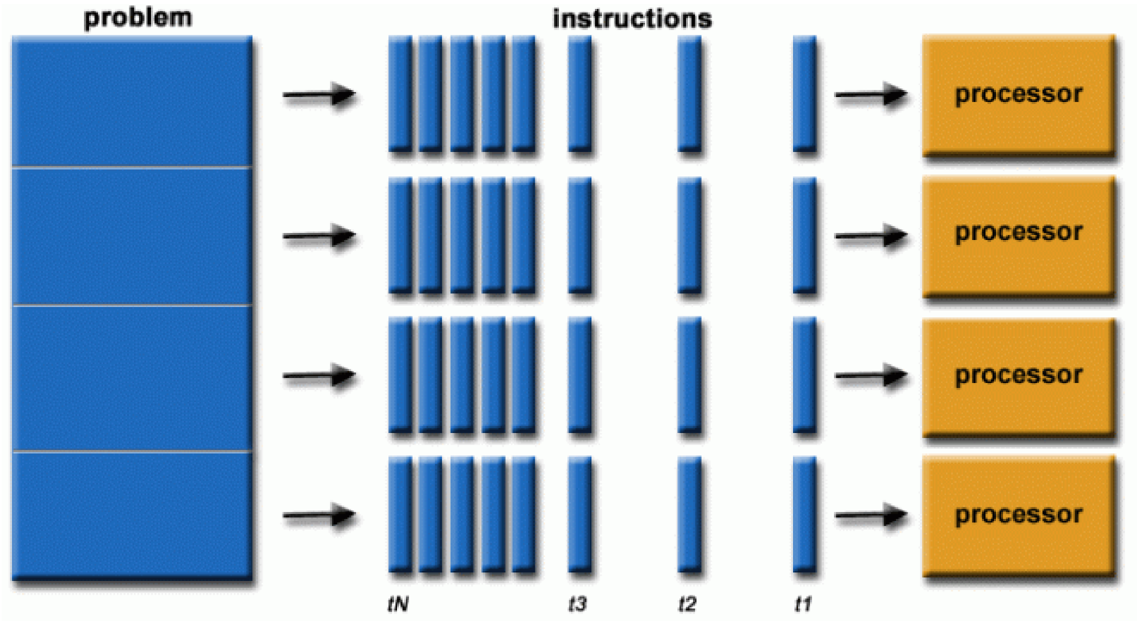

https://hpc.llnl.gov/documentation/tutorials/introduction-parallel-computing-tutorial

How to work in Parallel?

- Not all functions can be run in parallel
- SAGA GIS runs in parallel by default
- GRASS can be parallel with OpenMP
- Python: use `Dask`
- R: use `Future`
- ArcGIS: Available for some functions, not all

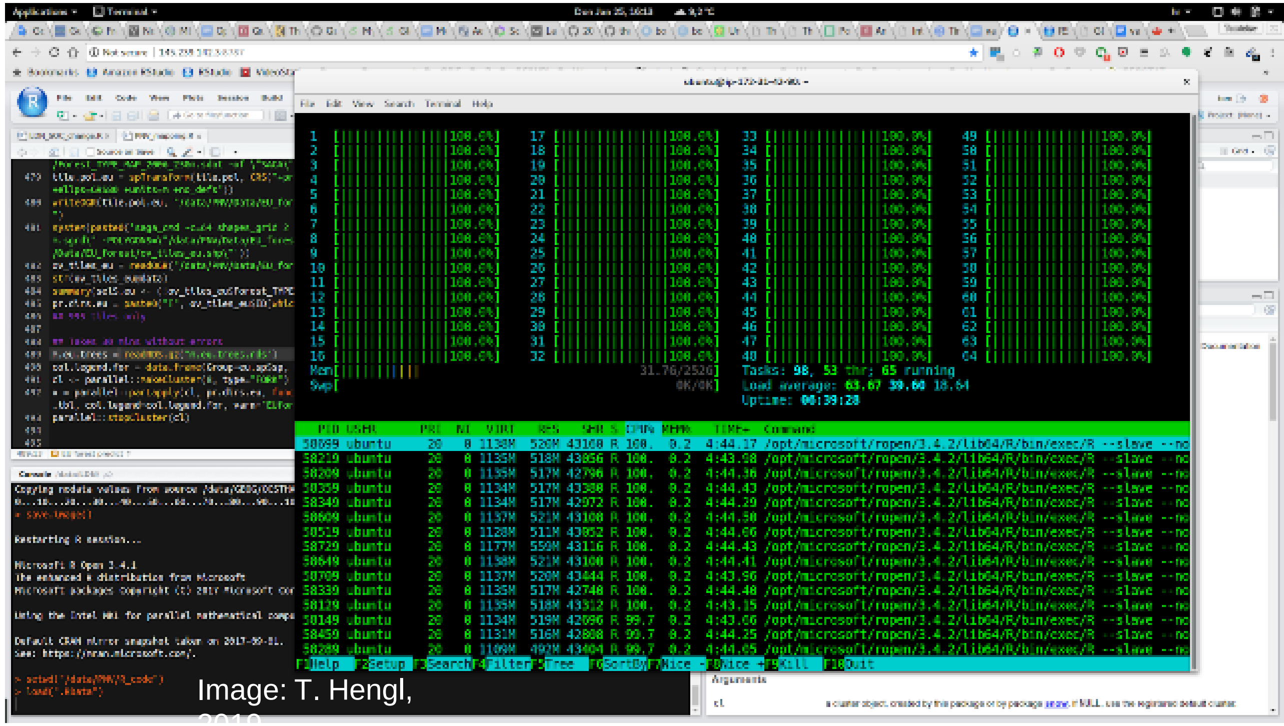

Image: T. Hengl,

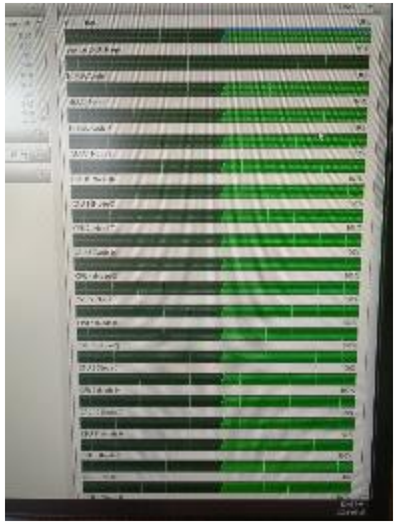

Managing large images locally

- Tiles
  - Easier to manage as small tiles
- Visualize all tiles at once
  - **Virtual raster (gdalbuildvrt)**
    Creates index of all tiles (small file)
  - **Mosaic (gdalmerge)**
    Combines all images into a large file
- Speed up visualization
  - **Overviews (gdaladdo)**
    Creates multiple reduced resolution layers that are used at different zoom levels, speeds up visualization. Layers stored in *.ovr file

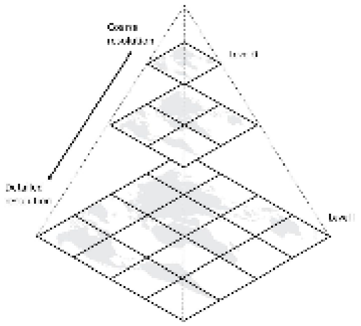

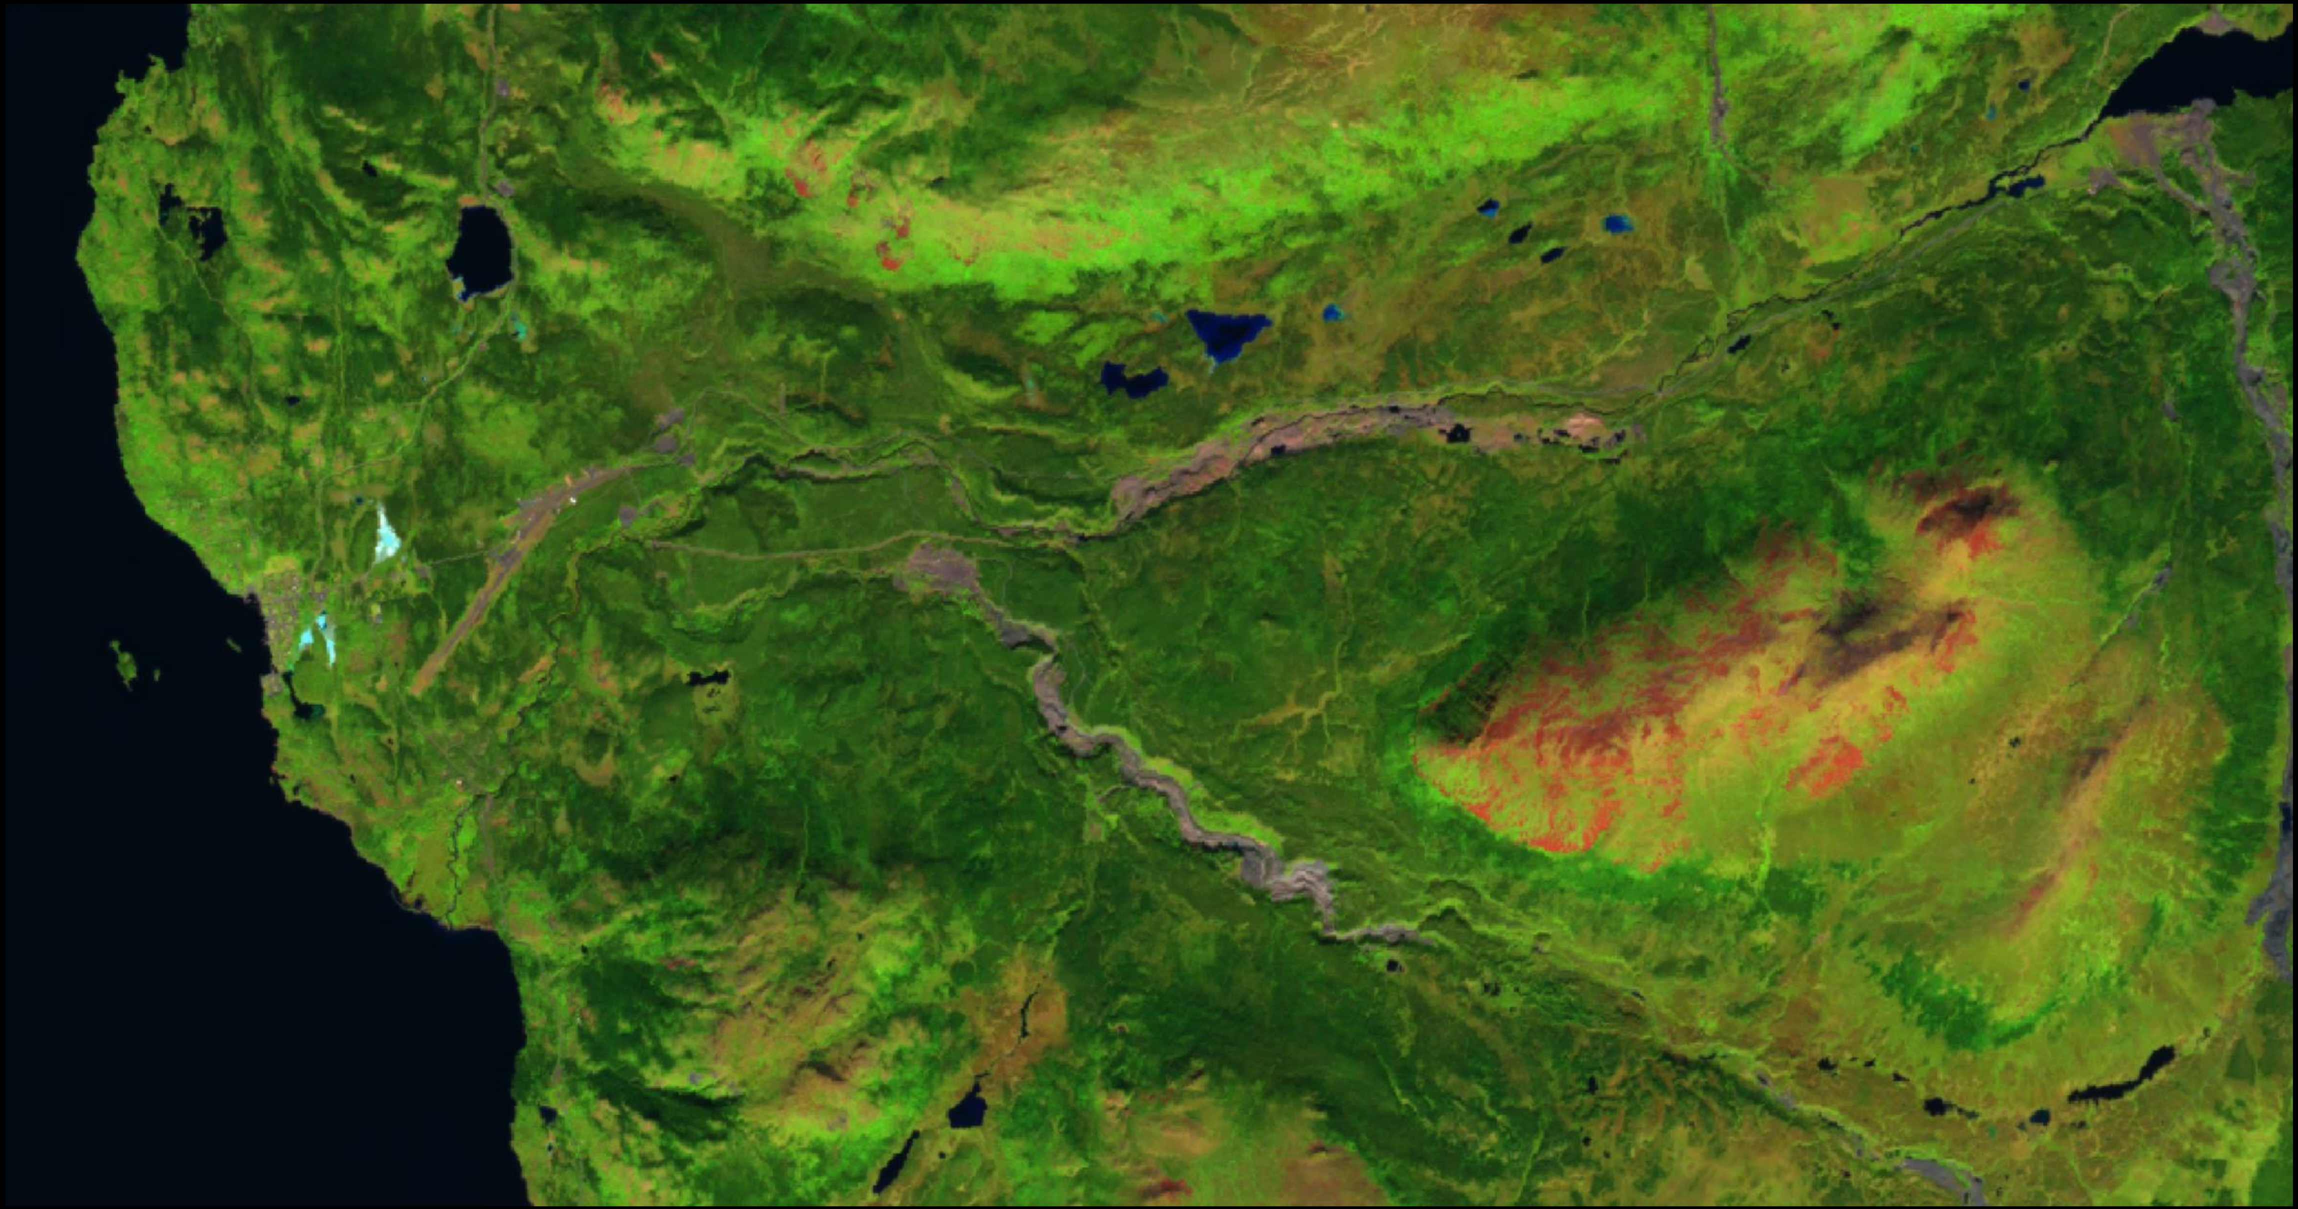

Cloud optimized geotiffs

- A COG is a regular GeoTIFF
- COGs have an internal organization that supports efficient access via HTTP GET range requests
- Supports overviews
- Clip rasters BEFORE processing downloading
- Used by:
  - STAC Index
  - Google Earth Engine
  - DigitalGlobe/Mazar
  - USGS
  - etc.

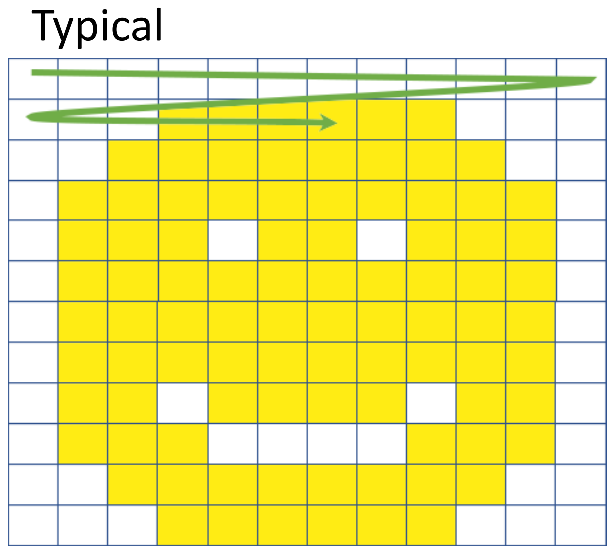

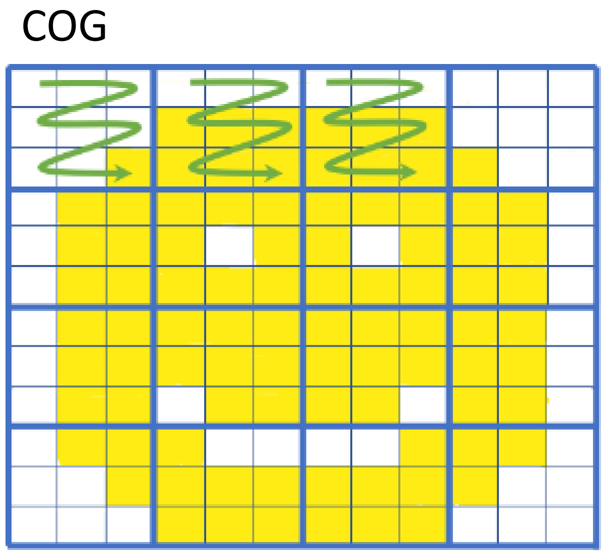

https://www.element84.com/blog/cloud-optimized-geotiff-vs-the-meta-raster-format

44_cog_stac.ipynb

File Edit View Insert Runtime Tools Help

+ Code + Text Copy to Drive

For this demo, we will use data from https://www.maxar.com/open-data/california-colorado-fires for mapping California and C

List of COGs can be found here.

```
[ ] import ee
    import geemap
```

```
[ ] Map = geemap.Map()
    Map
```

```
Map(center=[40, -100], controls=(WidgetControl(options=['position'], widget=HBox(children=(ToggleButton(value=.
```

```
[ ] url = 'https://opendata.digitalglobe.com/events/california-fire-2020/pre-event/2018-02-16/pine-gulch-fire20/10
```

```
[ ] geemap.cog_bounds(url)
```

```
[-108.63447456563128,
 38.963980238226654,
 -108.38008268561431,
 40.025815049929754]
```

```
[ ] geemap.cog_center(url)
```

```
(-108.5072786256228, 39.49489764407821)
```

```
[ ] geemap.cog_bands(url)
```

```
['band1', 'band2', 'band3']
```

```
[ ] geemap.cog_tile(url)
```

```
'https://titiler.xyzcog/tiles/WebMercatorQuad/{z}/{x}/{y}@1x?url=https%3A%2F%2Fopendata.digitalglobe.com%2Fever
```

```
[ ] Map.add_cog_layer(url, name="Fire (pre-event)")
```

```
[ ] url2 = 'https://opendata.digitalglobe.com/events/california-fire-2020/post-event/2020-08-14/pine-gulch-fire20/
```

```
[ ] Map.add_cog_layer(url2, name="Fire (post-event)")
```

DEMO

https://colab.research.google.com/github/giswqs/geemap/blob/master/examples/notebooks/44_cog_stac.ipynb

Spatiotemporal Asset Catalog (STAC)

- Manage geospatial data with a single language
- Ideal for searching and managing datasets (not for processing or visualization)
- Built for the cloud using an open standard format for simple geographical features, along with their non-spatial attributes (GeoJSON)
- STAC consists of:
  - Catalogue with collections (e.g. Sentinel-2),
  - Collection has Items (e.g. Multiband image)
  - Items have assets (e.g. Single band image)

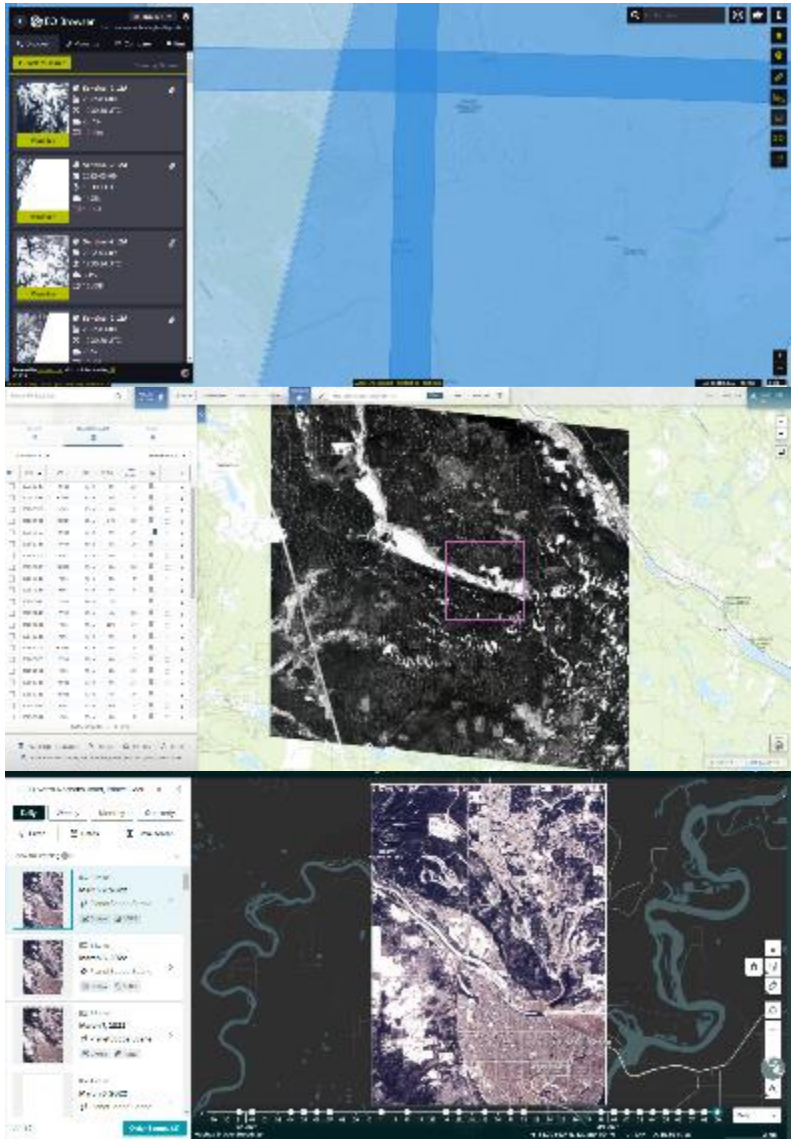

https://stacindex.org/catalogs/planet-labs-stac-catalog#/?t=1

Amazon Web Services (AWS)

- Common home for large geospatial datasets (e.g. Climate, Imagery, OpenStreetMap, Terrain)
- Typically data are hosted as “COGs” using “STAC”
- Access bands individually in seconds
- Can be accessed from a PC or from another cloud service
- There is a COST
- E.g.: https://aws.amazon.com/earth/

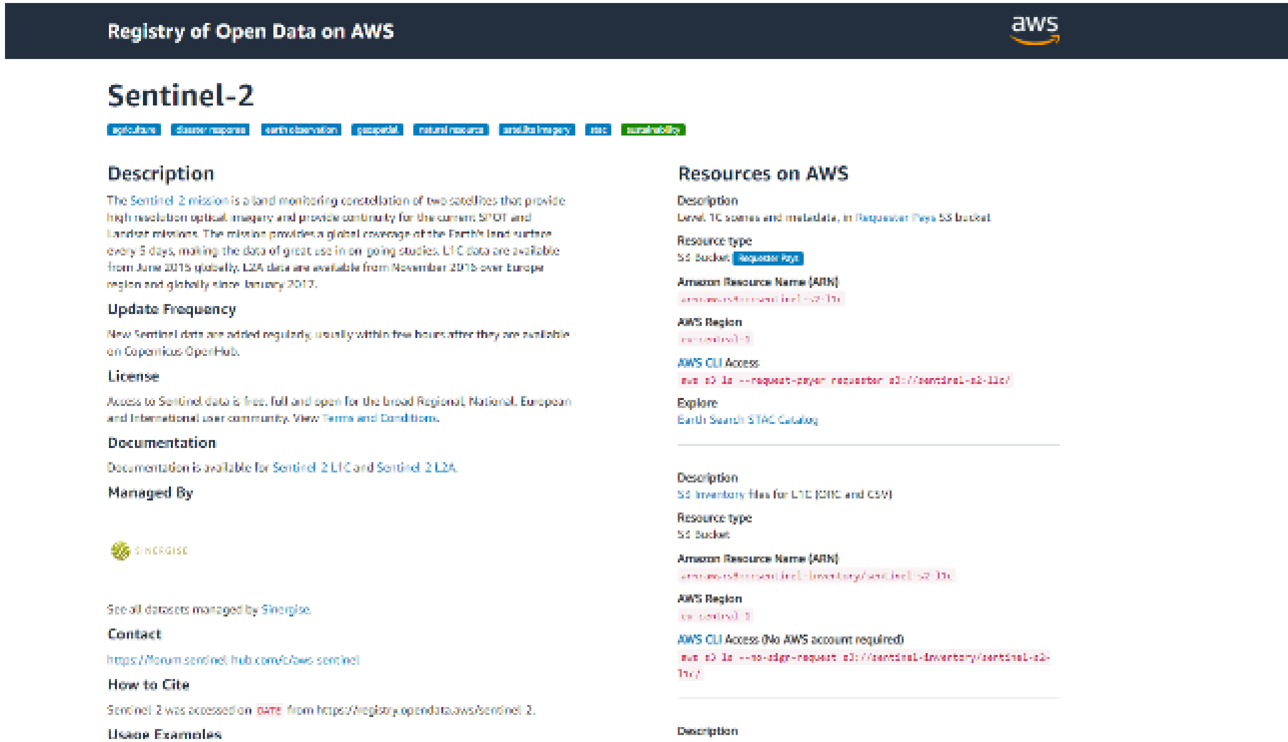

Google Earth Engine

- Available via: Python (geemap), R (rgee) and JavaScript
- Will use in next lab
- Very common in research and not-for-profit remote sensing
- Free for research education and not-for-profit use, must apply for commercial license

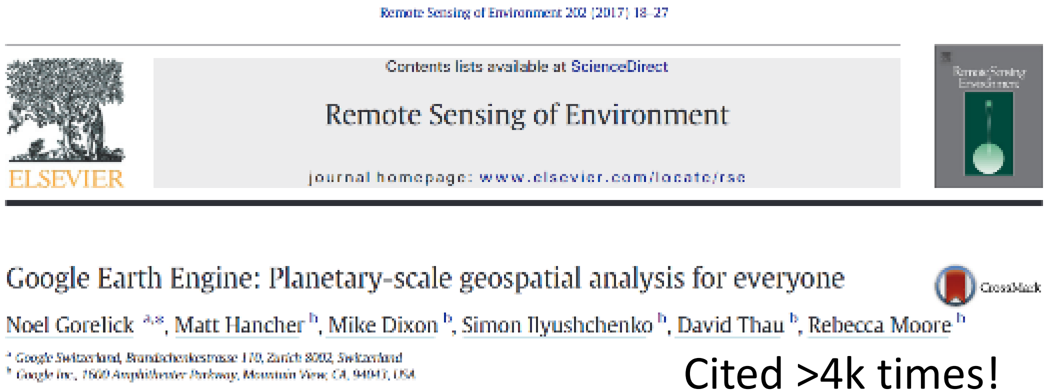
Remote Sensing of Environment 202 (2017) 18–27

Contents lists available at ScienceDirect

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

Google Earth Engine: Planetary-scale geospatial analysis for everyone

Noel Gorelick [a,*], Matt Hancher [b], Mike Dixon [b], Simon Ilyushchenko [b], David Thau [b], Rebecca Moore [b]

[a] Google Switzerland, Brandschenkestrasse 110, Zurich 8002, Switzerland
[b] Google Inc., 1600 Amphitheater Parkway, Mountain View, CA, 94043, USA

Cited >4k times!

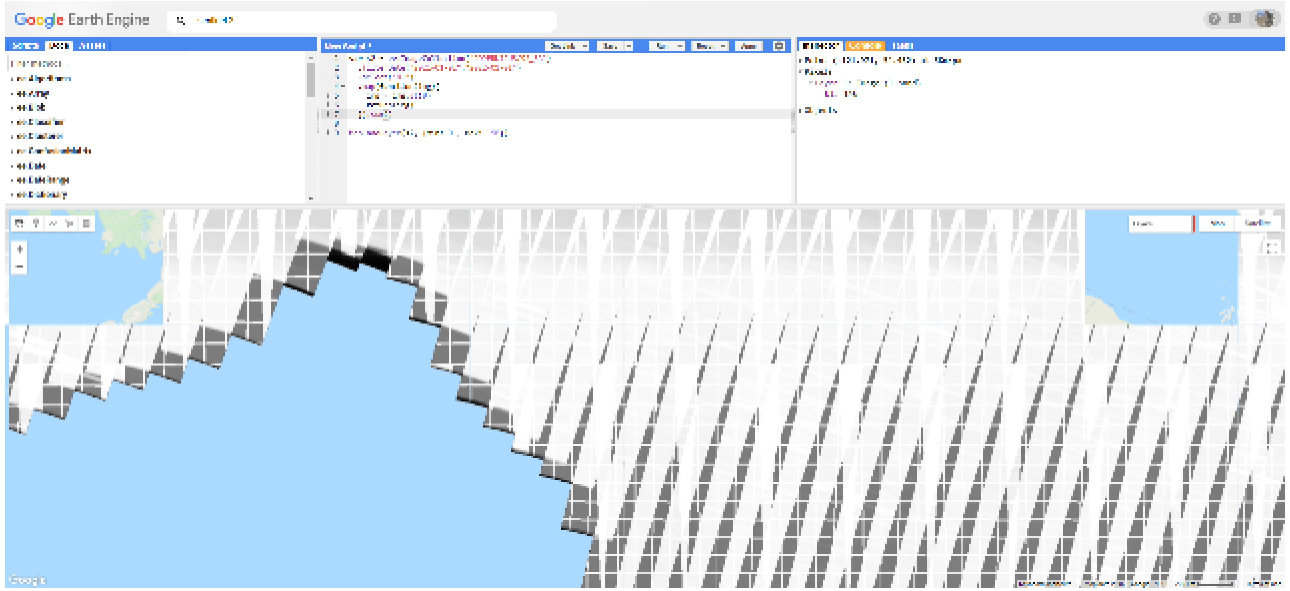

High-Resolution Global Maps of 21st-Century Forest Cover Change

M. C. Hansen,[1*] P. V. Potapov,[1] R. Moore,[2] M. Hancher,[2] S. A. Turubanova,[1] A. Tyukavina,[1] D. Thau,[2] S. V. Stehman,[3] S. J. Goetz,[4] T. R. Loveland,[5] A. Kommareddy,[6] A. Egorov,[6] L. Chini,[1] C. O. Justice,[1] J. R. G. Townshend[1]

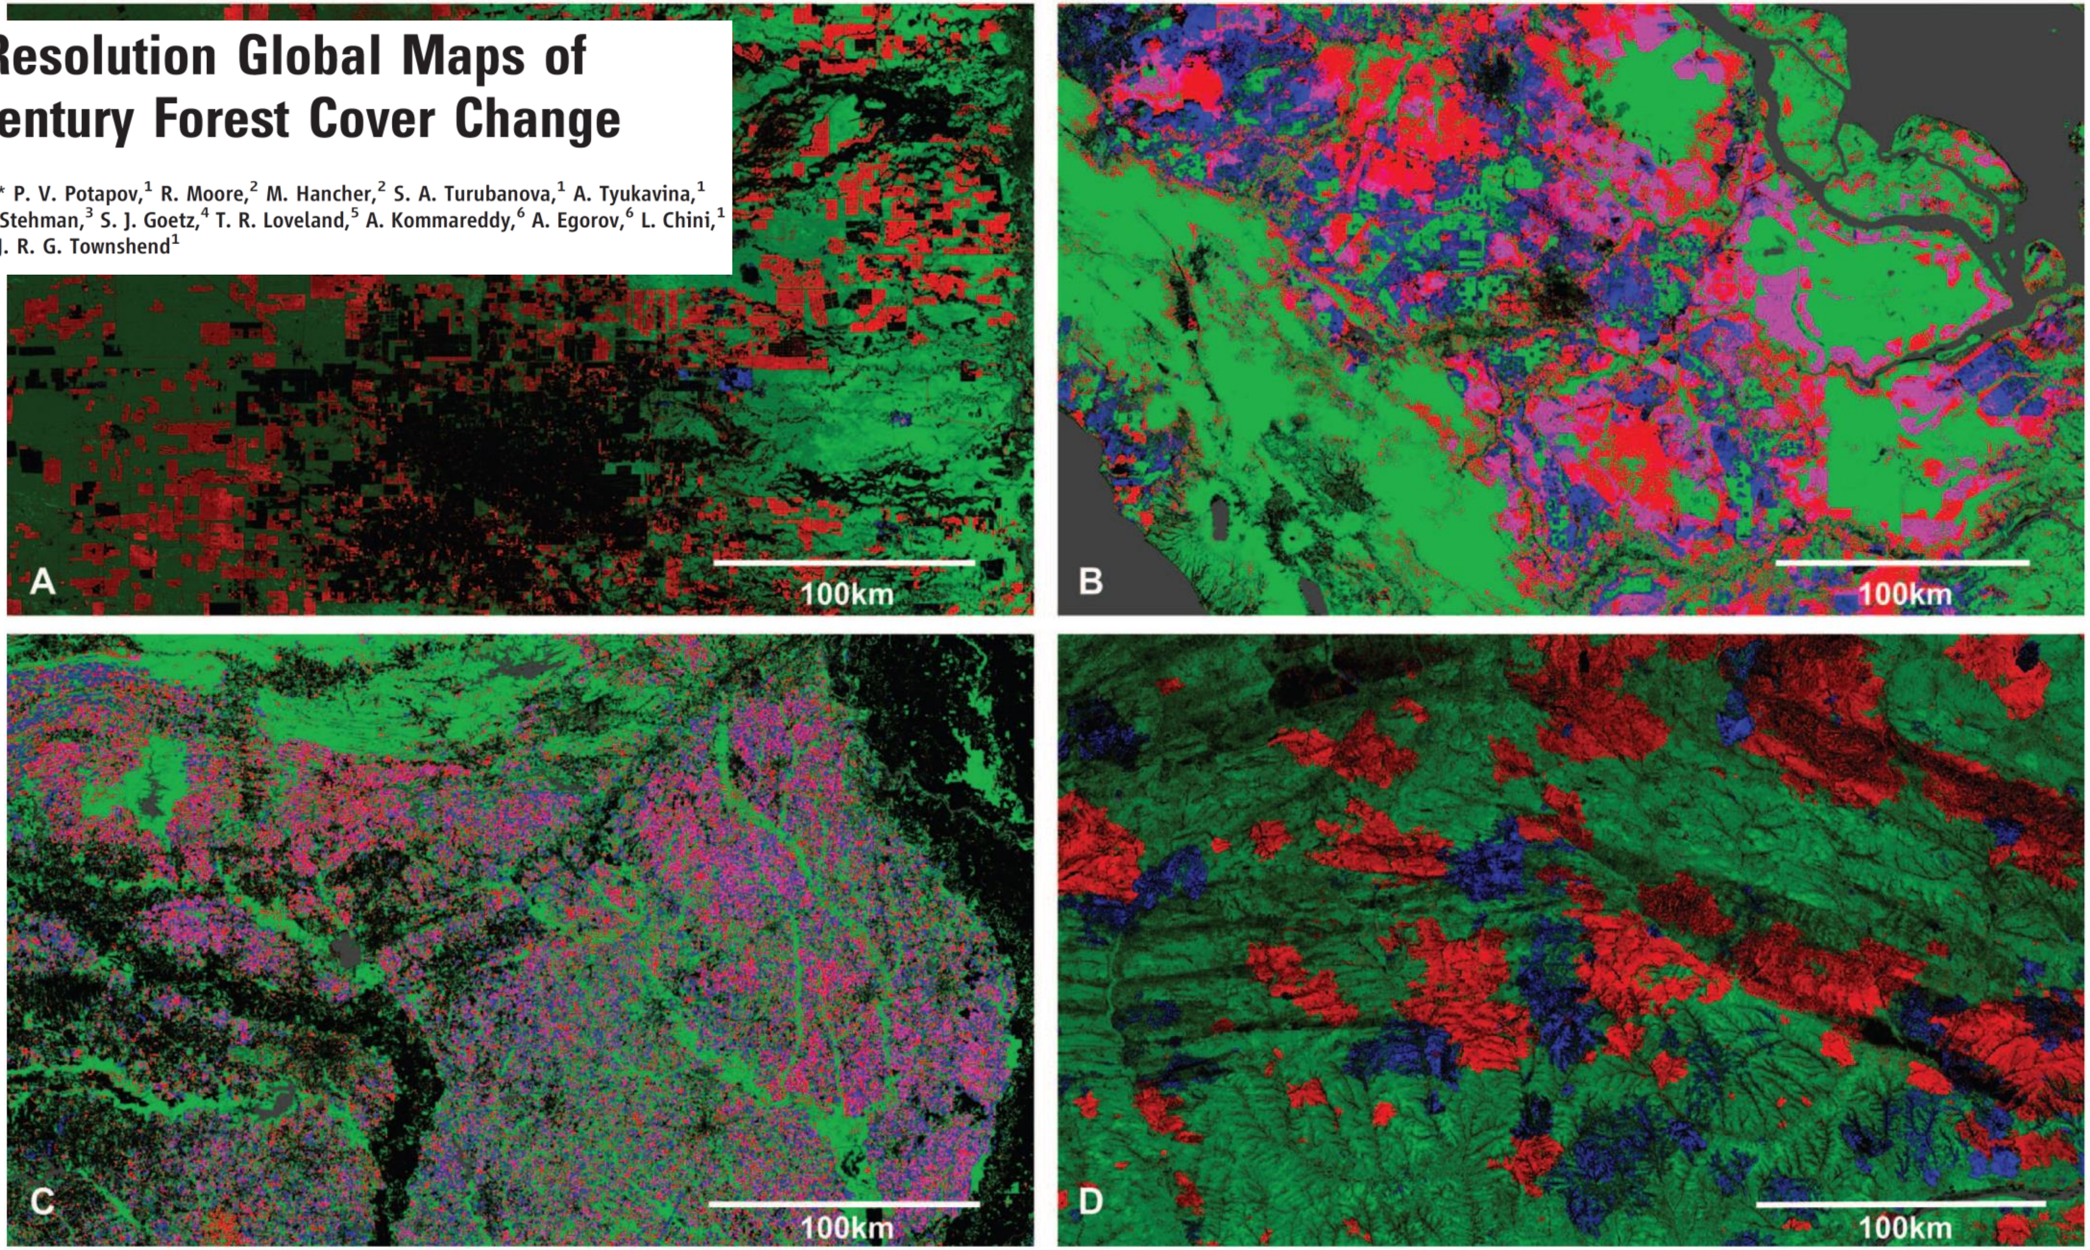

Fig. 2. Regional subsets of 2000 tree cover and 2000 to 2012 forest loss and gain. (**A**) Paraguay, centered at 21.9°S, 59.8°W; (**B**) Indonesia, centered at 0.4°S, 101.5°E; (**C**) the United States, centered at 33.8°N, 93.3°W; and (**D**) Russia, centered at 62.1°N, 123.4°E.

Cited >8k times!

Microsoft Planetary Computer

The Planetary Computer combines a multi-petabyte catalog of global environmental data with intuitive APIs, a flexible scientific environment that allows users to answer global questions about that data, and applications that put those answers in the hands of conservation stakeholders.

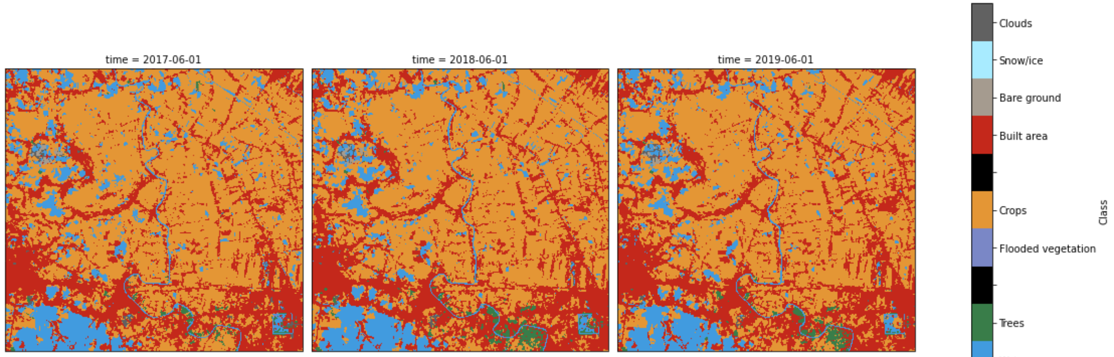

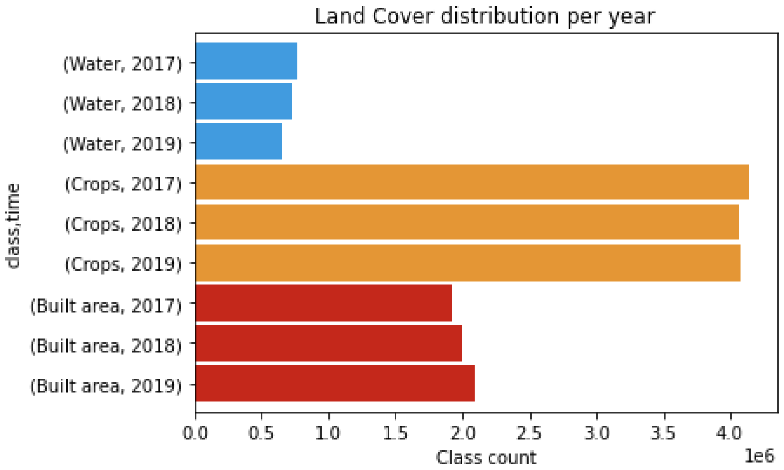

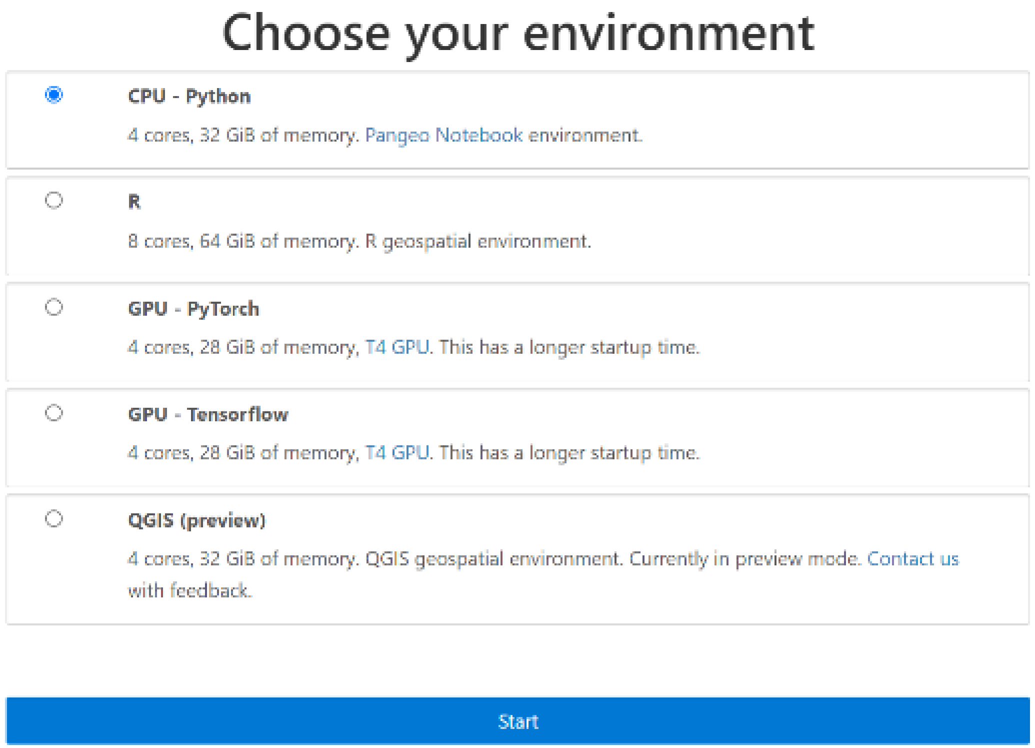

https://planetarycomputer.microsoft.com/dataset/io-lulc-9-class#Example-Notebook

https://planetarycomputer.microsoft.com/

Conclusion

- Tricks exists to make local processing more efficient
  - More cores, more RAM, better GPU
  - Tiling large datasets
  - For visualization: 8-bit compression, VRTs, Overviews
  - Parallelization processing
- Cloud infrastructure keeps data close to processing
  - COGS/STAC to query catalogues and grab only the data you need
  - Some free/cheap options for cloud computing
  - Commercial solutions can be very efficient but can also be costly