2026 HLS Workshop Answered Q&A

NSITE highly recommends you review the recorded presentations and slides, particularly the Q&A and open ended discussion sections for more background and information not captured in the questions and answers here. Timestamps for the recorded video workshop session and some questions are included.

FAQ

On Reporting Issues and Concerns or Questions

Question: Who is the appropriate point of contact to report missing values or limitations identified in the SpatioTemporal Asset Catalog (STAC) item JSON for Landsat HLS datasets?
Answer: Please email the support team directly with the specific details, and they will route the inquiry to the appropriate technical team for investigation. If you have issues found via a resource through LP DAAC, you can report your issues to the Earthdata Forum.

On HLS

Question: Are the spectral bands of Landsat and Sentinel-2 fully compatible within the HLS dataset?
Answer: There are several overlapping bands between Landsat (versions 4 through 9) and Sentinel-2 satellites. The specific bands that are common between both sets of satellites are Blue, Green, Red, Near-Infrared (NIR), Shortwave Infrared 1 (SWIR1), and Shortwave Infrared 2 (SWIR2). While the band names align, the exact nanometer ranges are not perfectly identical. More information on the HLS spectral bands can be found here.

Question: If the datasets are harmonized, why are they available separately as Landsat 30 and Sentinel-2 30? Why are they not just one dataset? Do users need to access the Landsat (HLS L30) and Sentinel-2 (HLS S30) datasets separately, or are they combined into a single dataset?
Answer: This formatting is largely due to the bands being slightly different between the two sources. For example, Landsat 8 and 9 include thermal data, whereas Sentinel-2 does not. HLS L30 and HLS S30 are independent collections. While they are harmonized to work together, they are hosted and accessed separately.

Question: In the HLS pipeline, why re-grid the coarser-resolution Landsat data into the Sentinel-2 grid and not the other way around? Doesn't this degrade spatial clarity?
Answer: The Military Grid Reference System (MGRS) used for Sentinel-2 introduces a global fixed tiling system, which facilitates stacking and time series analysis compared to UTM. For more information on the resampling of input data, see Ju et al. (2025), The Harmonized Landsat and Sentinel-2 version 2.0 surface reflectance dataset. *Remote Sensing of Environment, 324*, 114723.

Question: How does the harmonization process handle the different spatial resolutions between Landsat and Sentinel-2 for high-precision mapping?

Answer: The Landsat 30 m surface reflectance and brightness temperature data at 30 m spatial resolution are resampled from the native WRS-2 Path/Row system to 30 m pixels in MGRS tiles (native to the Sentinel-2 data) using cubic convolution. For more information on the resampling of input data, see Ju et al. (2025), The Harmonized Landsat and Sentinel-2 version 2.0 surface reflectance dataset. *Remote Sensing of Environment, 324*, 114723.

Question: Are there any thoughts on harmonizing Sentinel-1 SAR with these multispectral datasets?
Answer: NSITE are currently unaware of attempts to harmonize SAR with multispectral datasets. That said, the OPERA products were designed to be based on both optical and SAR data. Rather than trying to create a SAR/optical data stream and use that as input, NASA JPL is producing an optical-based version, and a SAR-based version where users can attempt to use those two different data streams. There is some merit that this is proven to be more efficient than trying to harmonize optical and SAR data, and use two different outputs to best handle the pros and cons. OPERA DSWx, the global surface water extent product, is a great example of that. Certain aspects can be observed in the HLS-based version that are missing in the SAR-based version, and vice versa. It is possible you may lose those advantages if you were to try to harmonize the SAR and optical.

Question: Is there an avenue for the community to contribute to the algorithms used to generate HLS data, or is this exclusively handled by internal development teams?
Answer: The core algorithms are maintained by the internal HLS development teams; however, they welcome feedback and recommendations from the community. If you have specific suggestions for algorithm improvements, please contact the development team directly. Also answered verbally in Day 1 recording at 4:28:28.

Question: Are there any plans on developing an AOD product on HLS?
Answer: The HLS teams are not planning an AOD product but LaSRC provides quality flags associated with AOD (low, moderate & high). More information in the following pub, section 3.5 & fig 6.

HLS Temporal Aspects

Question: From a spectral perspective, how far back could an HLS-consistent Landsat archive realistically be extended? Would the Landsat 4/5 TM era be the most feasible backward extension, whereas the Landsat 1–3 MSS era would only allow for limited harmonization?
Answer: The L30 archive begins with the launch of Landsat 8 data in 2013. As of now, there are no plans to process backwards to the older satellites.The primary constraint for backward extension is the availability of Sentinel-2 data, as the satellite was not launched until 2015. Furthermore, the spectral bandpass adjustment is applied to Sentinel-2 data to match Landsat

data. Consequently, HLS processing on Landsat data prior to 2013 would not include this spectral adjustment. More information on the algorithms can be found here. Furthermore, there are no current plans to add legacy Landsat 4, 5, or 7 data to the HLS framework. Also answered verbally in Day 1 recording at 1:01:00.

Question: Since Sentinel data has been available since 2017, can the processing algorithms be applied retroactively to older data?
Answer: The HLS S30 (Sentinel-2) data archive is already available from 2015 onwards, which coincides with the launch of the Sentinel-2A satellite. There are HLS data available for the whole Sentinel-2 era.

Question: Are there sensor degradation effects or data gaps when conducting time-series analyses with HLS and Landsat data, and how are these overcome?
Answer: Sensor degradation is generally not a significant issue for newer sensors like Landsat 8 and 9. However, there were known issues and data gaps associated with older sensors like Landsat 5 and 7. For more information on historical sensor anomalies, please review here.

Question: Is the HLS dataset suitable for long-term land cover and land use change studies (e.g., going back 35 years)?
Answer: The HLS catalog starts from 2013, corresponding to the launch of Landsat 8, and includes Sentinel-2 data starting from its launch in 2015. Therefore, it does not cover decades-long historical data (such as imagery from 1995). For studies spanning 35 years, researchers would need to independently combine older Landsat archival data with the modern HLS dataset.

Question: When building a time-series analysis using HLS L30 data from 2016 to 2017, the temporal resolution was approximately 7 to 10 days, leading to some observation gaps. Has this improved in recent years?
Answer: Yes. Starting in 2021, the ingestion of Landsat 9 data into the production pipeline increased the overall temporal resolution of the HLS L30 data. A schematic of the HLS temporal frequency can be viewed here.

Question: Is anyone looking to develop an algorithm for past years, such as from 1985-2000?
Answer: Currently no plans to produce HLS L30 before 2013. But the Landsat data before 2013 is still compatible with the HLS data in the sense that it's the same spectral resolution and the same spatial resolution. The geometry might be a bit different but I'm not aware of any plans of producing L30 data back in time past 2013.

Question: Will pre Landsat-8 be added to HLS at any point?

Answer: There are no plans at this point to include Landsat data collected by platforms prior to Landsat-8.

Question: How will Landsat Next fit in with HLS?
Answer: Which HLS products that will include data collected by Landsat Next has not been decided.

BRDF Corrections

Question: How are the View Azimuth Angle (VAA), View Zenith Angle (VZA), Solar Zenith Angle (SZA), and Solar Azimuth Angle (SAA) angular bands used in HLS data?
Answer: These angular bands are provided because they are utilized in the Bidirectional Reflectance Distribution Function (BRDF) correction process. They are included in the HLS datasets to allow advanced users the flexibility to perform custom BRDF corrections if desired. Detailed information is available in the HLS User Guide.

Clouds and Cloud Masking

Question: Which cloud mask algorithm is used for the dataset?
Answer: The dataset currently implements the Fmask algorithm (version 4.7) for cloud masking. The open-source code is available here.

Question: How can users effectively remove clouds that are not detected by the standard HLS mask, especially when working on continental-scale analyses where manual checking is unfeasible?
Answer: Cloud masking on a large scale is inherently challenging and prone to missing clouds, shadows, or over-masking bright areas (like bare ground). Users often employ alternative methods, such as supervised classification, to detect and mask missed clouds. Other approaches include using the ATSA algorithm or a modified despiking approach (referencing Bolton et al., 2020), which benefit from the high temporal frequency of HLS. Additionally, the standard HLS product currently includes a cloud mask layer produced using the Fmask algorithm (version 4.7). Future production updates will implement Fmask v5, which is expected to offer significant improvements in cloud masking.

Question: I have created cloud free 10-day and monthly cloud free composites with HLS data and I'm still seeing gaps with clouds. Would you fill these areas before delivering?
Answer: The general recommendation is to not try to fill that. If there are no data, there's no data. If it's cloudy, and are unable to see the land surface, it is commonly preferred to keep this screened out. It is possible to artificially fill that, but you have to consider how much more informative that would offer than having those pixels screened out. Gap filling with AI

Foundation Models might look better, and NSITE is interested in what the community using monthly composites in biodiversity would prefer. If users do not prefer to have those screened out, even if it's artificial, they might still prefer something that looks better without a bunch of gaps.

HLS Vegetation Indices (HLS VI)

Question: What is the difference between the two HLS Vegetation Index (VI) products?
Answer: The two vegetation index products correspond to the different satellite sources used in HLS: one is derived from Sentinel-2 (HLS VI S30) and the other from Landsat (HLS VI L30). More information about these products and the specific spectral indices used can be found in the HLS VI User Guide.

HLS Low Latency

Question: How is low-latency access defined for HLS data?
Answer: Standard HLS data is typically available 2 to 3 days after the initial satellite observation. With the low-latency version, the data becomes available approximately 6 hours after observation.

Question: Is the low-latency HLS data currently accessible?
Answer: Low-latency HLS data is still in development and is tentatively expected to be released late next year. It has not yet been determined whether it will be delivered as a separate data stream or integrated into the primary HLS delivery.

HLS 10 m

Question: Is an HLS product at 10-meter resolution currently in development?
Answer: Yes. The plan is to begin the development of a 10-meter resolution product by the end of this year. Prototype HLS 10 m data (currently under research and development) is available for review.

Question: For the upcoming 10 m spatial resolution version planned for the United States, will the band offerings and band order remain identical to the current HLS L30 and HLS S30 products? Furthermore, should users expect any substantial modifications to the metadata beyond the increased spatial resolution?
Answer: The upcoming product will be highly comparable to current offerings, localized specifically to the United States at a 10 m resolution. The project team expects to maintain the same set of spectral bands and does not anticipate any significant additions or changes to the metadata structure.

Question: Will the upcoming 10-meter HLS data cover areas outside the Contiguous United States (CONUS), such as Canada?
Answer: The development team plans to release the open-source code for the HLS 10-meter solution. This will allow researchers interested in regions outside of CONUS to run the code and produce the 10-meter product for their specific study areas.

Question: Will the algorithm used to produce the 10-meter HLS data be open-source so that users can process data independently?
Answer: Yes, all the code is open-source, allowing users to access and utilize it as needed. Also answered verbally in Day 1 recording at 1:01:33.

Question: With Sentinel-2C now in orbit, will it be integrated with Sentinel-2A and 2B to generate the 10-meter HLS product?
Answer: The 10-meter HLS product will incorporate data from all operational Sentinel-2 satellites. This is expected to include Sentinel-2B and Sentinel-2C, as well as Sentinel-2A, provided it remains operational.

Question: Do you have imagery with a resolution of less than 10 meters for specific countries? For research projects, is it freely accessible, or is a fee required?
Answer: No, there is currently no HLS data available at less than 10 m resolution. You can review the intermediary 10 m prototype data (under research and development).

On Trainings and Resources

Question: Can folks suggest conferences that are relevant to this workshop? (especially for academic users, who are mostly using open-access data and platforms)"
Answer: To name a few: AGU, ForestSAT, IGARSS, Pecora

Question: Does ARSET offer additional free courses? I have already completed the introductory remote sensing and research modules and would like to know what upcoming courses are available.
Answer: Yes, ARSET offers a wide range of free courses. If you have already completed introductory modules, you can filter for intermediate and advanced training levels directly on the ARSET platform. You can also view upcoming ARSET trainings here. By creating an account, you will gain full access to all available training modules.

On Applications of HLS

Question: There is a strong case for combining both products but when would users choose L30 or S30 individually instead of sticking with conventional Landsat or Sentinel-2 imagery?
Answer: Because the L30 and S30 products are harmonized (geometrically harmonized, spectrally harmonized, and atmospherically corrected), it is likely easier for you to work with the HLS product rather than sticking with conventional Landsat and Sentinel-2 imagery to work with both Landsat and Sentinel-2 seamlessly. There is not necessarily any new information in the HLS L30 and S30 compared to standalone Landsat and Sentinel-2. It is the way the data is harmonized that makes it easier for you to work with.

Question: Is there a specific platform or repository where researchers can submit academic projects that specifically utilize Harmonized Landsat Sentinel-2 (HLS) data?
Answer: While NASA does not host a dedicated repository for all HLS-based projects, the agency does compile and share societally impactful stories demonstrating how HLS data improves operational efficiency and decision-making beyond academic research. The project team regularly monitors peer-reviewed publications, Zenodo, professional blogs, and LinkedIn for general updates on academic work involving HLS. You can explore existing impact stories on the NASA Earthdata Solutions Page.

Question: How can HLS data be integrated with AI/Machine Learning models to improve real-time monitoring of satellite communication infrastructure in remote areas?
Answer: A good example is the Prithvi Foundation Model which was trained on HLS because of its high temporal resolution; for more information and examples of downstream applications, see Szwarcman et al. (2025). Prithvi-eo-2.0: A versatile multi-temporal foundation model for earth observation applications. *IEEE Transactions on Geoscience and Remote Sensing, 64*, 4400120.

Question: How can VIIRS data be used as a thermal band for Sentinel?
Answer: The algorithms for resampling VIIRS thermal data to 30 m resolution is currently being developed by the solutions implementation team of the 2024 SNWG solutions High Resolution Land Surface Temperature. The team includes scientists from NASA Marshall, Ames, Goddard, and JPL.

Vegetation Applications

Question: How can HLS data be effectively used to monitor small-scale changes in vegetation in regions with frequent cloud cover?
Answer: In comparison to stand alone Landsat and Sentinel-2, HLS offers a near-daily temporal resolution which is valuable when studying the land surface in areas of frequent cloud cover. More specifically, different change detection algorithms handle time-series data

uniquely; for instance, CCDC adapts different regression curves to the varying frequency of clear-sky observations, whereas LandTrendr relies on temporal composites to detect changes. However, persistent cloud cover presents a fundamental limitation for optical remote sensing, where the integration of SAR data may help. Detecting small-scale changes also depends on the severity of the disturbance. But when a change footprint is far smaller than 30 meters, current HLS data may face spatial constraints due to mixed-pixel effects. The HLS teams are developing a HLS 10 m product, which will improve our ability to resolve highly localized, small-scale land changes.

Question: When producing a seasonal composite, what criteria are used to delineate the seasons?
Answer: Seasonal composites for existing NASA products (such as PACE) typically utilize standard three-month seasonal blocks.Quarterly composites would be based on the traditional quarter e.g. quarter one, two, three and four. The composites produced as part of the NASA biodiversity virtual center are not seasonal. They are 10 days and monthly. NSITE and BDEC may be able to provide some support e.g. tutorials on how to further create seasonal composites. HLS solutions will not produce seasonal composites, as is the plan as of right now.

Agricultural Applications

Question: The user community is asking for smaller-scale agricultural monitoring. How does this fit with the Commercial Smallsat Data Acquisition (CSDA) program and future agency needs?
Answer: The demand for high-resolution monitoring is well recognized. The CSDA program is actively growing to meet these demands, and the development of the new 10-meter HLS solution is a direct response to these community needs.

Water and Weather Related Applications

Question: Can HLS data be utilized for monitoring coastal vegetation, such as mangroves?
Answer: Yes, HLS data is actively being used for coastal vegetation research, including the study of mangroves.

Question: Are there plans to add surface water temperature and Chlorophyll-a products in addition to the vegetation indices? This would be highly beneficial for monitoring harmful algal blooms (HABs).
Answer: For water quality monitoring, there are other dedicated solutions available. Specifically for HABs, researchers may find the Global Algal Blooms Assessment Network (GABAN).

Question: At one point it was recommended against using HLS for water quality detection such as detecting HABs. If we were still interested in using 30-m data to do this, is it recommended that we use either native Sentinel-2 or Landsat data by itself?
Answer: While HLS is sometimes not recommended for precise water quality detection, native Sentinel-2 data is frequently used for HAB monitoring. Additionally, the USGS provides a Landsat 8/9-based provisional aquatic reflectance product.

Question: Is the atmospheric correction in HLS data accurate for both land and water surfaces? Is it appropriate to use HLS for aquatic reflectance?
Answer: No, because HLS uses the LaSRC algorithm which is most appropriate for land surface reflectance estimates. Through another project, investigating the applicability of aquatic reflectance algorithms with the HLS time series, but this is a separate project that is under development. If you are interested in water quality studies, it is not recommended to use HLS. Answered verbally in Day 1 recording at 1:02:52.

Question: What is the best source for monitoring global ephemeral surface waters?
Answer: Researchers are encouraged to explore the NASA OPERA project, which provides observational products for end-users. Additionally, OPERA will produce a 30-meter NISAR-based displacement time-series product for North America, and eventually a vertical land motion product (VLM-NISAR). More information is available here.

Question: Can HLS data be used for water surface temperature?
Answer: HLS surface reflectance data is ideal for land-based applications. However, HLS L30 products contain thermal data from Landsat 8/9 as brightness temperature (Bands 10 and 11) which could potentially be leveraged for applications involving water surface temperature.

Other Applications

Question: To what extent is AppEEARS useful for non-technical fields like cultural heritage or archaeological studies?
Answer: AppEEARS handles data pre-processing, but it does not perform data analysis or draw conclusions. Its primary objective is to simplify access to satellite data, allowing users to focus on scientific research rather than the technical complexities of managing and processing satellite datasets.

Question: Compared to fire fronts assessments, what measures (patterns/layers) would you recommend for regions in Japan that frequently experience earthquakes and tsunami warnings?
Answer: The atmospheric correction scheme used for HLS is geared toward land applications

and the land surface. Aquatic applications is not something that is well explored with HLS data. NSITE is building awareness on such applications. Answered verbally in the Day 2 recording at 2:01:11.

On HLS in Various Environments

Question: Do you recommend using QGIS or ArcGIS Pro for operations involving HLS data?
Answer: Users are encouraged to review the workshop's Day 1 presentations recordings to evaluate which of these software best aligns with their technical requirements. Consulting the user community regarding their practical workflows and individual experiences can also provide valuable guidance for your decision.

On NASA and NSITE

NASA

Question: Since NASA data is openly available, does it matter if geospatial analysis was not included in our initial project scope for a grant? We would like to incorporate this data, even though it was not in the original grant proposal.
Answer: This depends on the specific grant. Each NASA funding opportunity may have different requirements and restrictions regarding how to apply and utilize NASA data within the scope of the project.

NSITE

Question: Is there an RSS feed available to subscribe to the newsletter?
Answer: Not at this time, though this feature is being explored. Currently, you can sign up for occasional HLS updates using this form.

Question: Does the NSITE newsletter cover the launch of new NASA OPERA data products?
Answer: Yes, the NSITE newsletter features all SNWG Solution data product availability updates, including OPERA. You can also follow OPERA directly for updates. Additionally, the broad NASA Earthdata monthly newsletter features data updates that will include this information.

Question: As an individual outside of federal employment, is there a way to engage with or join working groups for researchers in academia or private industry?
Answer: While the SNWG survey is restricted to federal government employees, there are many opportunities open to the broader community. These include joining NASA science

teams, applying for funding opportunities, and participating in organizations like the Earth Science Information Partners (ESIP). Additionally, many NASA missions have Communities of Practice (CoP) that anyone can join.

NISAR Soil Moisture Product (NSITE)

Question: Is there a general timeline for when NISAR Soil Moisture products will be released?
Answer: NISAR soil moisture, another SNWG solution that's unrelated to HLS. At the end of February 2026, ASF released preliminary NISAR data (not soil moisture) that covers a lot of parts of the world but not completely, and their timeline is by July, the global collection should be out. That should include other areas that were collected but were not included in the initial release. The current plan is for NISAR soil moisture data to become available this summer. Sample data for NISAR soil moisture products are available through NASA's Alaska Satellite Facility (ASF).

OPERA and HLS Related OPERA Products (NSITE)

Question: How is vegetation disturbance determined in the OPERA product line, and can this dataset be effectively used to detect agricultural harvest events?
Answer: The OPERA vegetation disturbance product maps per pixel vegetation disturbance (specifically, vegetation cover loss). Vegetation disturbance is mapped when there is an indicated decrease in vegetation cover within an HLS pixel. The spatial resolution is 30 m and the displayed layer describes vegetation disturbance status based on confidence, magnitude of loss, and whether it is ongoing.The OPERA vegetation disturbance product maps per-pixel loss of vegetation cover. A disturbance event is flagged when an HLS pixel shows a distinct decrease in vegetation cover. This layer is delivered at a 30 m spatial resolution and characterizes the disturbance status based on detection confidence, the magnitude of the loss, and whether the event is ongoing.

- For technical specifications, visit the OPERA DIST Product Suite.
- To see a use case with DIST and agriculture relevant case study, review NASA's Earth Science article on OPERA.
- This dedicated use case demonstrates how standalone HLS data can be utilized to operationally assess cover crop termination, which may closely mirror your application: USDA Cover Crop Termination Assessment.

Question: Can the OPERA vegetation disturbance product be used to isolate cyclical events like agricultural harvesting?
Answer: Agricultural harvesting is often excluded from disturbance classifications. For instance, the OPERA DIST-ALERT uses the previous three years of observations within a matching seasonal window to identify disturbances. Regular agricultural harvesting is typically

excluded from these disturbance classifications. For instance, the OPERA DIST-ALERT framework evaluates observations from the previous three years within a matching seasonal timeframe to establish a baseline. Similarly, Continuous Change Detection and Classification (CCDC) relies on utilizing historical time-series patterns as a reference to detect disturbances/identify anomalies. Because routine harvesting is captured within these baseline references, it is not flagged as an active disturbance. For workflows specialized in mapping harvest events, please consult Feng Gao's Day 1 presentation.

Question: According to Global Forest Watch, OPERA detection is less effective in forests with less than 30% canopy cover. How can we identify the effectiveness of the earthdata.nasa data?
Answer: The effectiveness of OPERA to detect forest loss as presented by accuracy assessments can be found in the literature. It is not currently clear how to use Earthdata to investigate the performance of these different products. In general, the way to do that would be to draw a sample from the land surface and provide reference observations and compare those to the map and to the various map products. These tools may allow you to collect such data to perform that. Answered verbally in the Day 2 recording at 2:02:22.

On High Resolution Harmonized Land Surface Temperature (LST) Product (NSITE)

Question: What is the timeline for the SNWG High-Resolution Harmonized Land Surface Temperature (LST) Products within HLS?
Answer: The LST project is just getting started, but data should start being produced in the next couple years!

Day 1 Presentation Based Answered Q&A

Day 1 Session 2: HLS in Action [1:24:49]

Question: How can HLS data be integrated with Machine Learning (ML) or Deep Learning models for monitoring applications? Are these models based on ground truth data?
Answer: When integrating HLS data into monitoring models (such as for rangeland products), traditional approaches like Partial Least Squares (PLS) models are sometimes preferred. Research has shown that standard ML and Deep Learning models can perform poorly under novel environmental conditions not represented in the training data, such as extreme drought

or excessively wet years. A relevant publication on this topic can be found here. For other variables like Leaf Area Index (LAI) and Evapotranspiration (ET), different approaches are utilized, often involving the downscaling of MODIS products.

Question: How can the Leaf Area Index (LAI) be calculated using HLS data?
Answer: HLS Vegetation Index (VI) products are currently available and provide the most common vegetation indices, which can be leveraged to estimate LAI. More information on HLS VI products can be found here. Additionally, recent research has successfully downscaled the MODIS LAI product using HLS data, demonstrating a strong correlation with ground-measured LAI in rangeland and pasture environments.

Monitoring Cover Crops for Agroecosystems Assessment Using HLS Data – Feng Gao, USGS [1:31:13]

Question: For the remote sensing monitoring of cover crop termination - are there ever cases where farmers do not agree with the satellite estimate of termination date (and want to contest this?)
Answer: Yes, there are some cases. When farmers report, maybe they report after termination and they've started planting so they don't have time or the data are not really accurate. That's self-reported data. Usually the cost share program staff check those fields, they select some percentage of fields to check if that's valid, and that's one kind of probable error in the database. Of course remote sensing is not perfect, it depends on the observation. If there are no observations during the critical time period (e.g. harvest) and there's always clouds and then it regrows again, you'll never be able to capture this harvest event, even if you do the daily interpolation, you miss this kind of information. In some cases, like weak cover crop performance, there is not a lot of change seen from the VI time series, so that won't be able to be detected or detected accurately. The staff can, looking at those different fields, verify from the ground so you still need some ground work.

Question: Why is the detection rate of greenup when using Landsat alone so low. 44% compared to 100% for Planet and HLS.
Answer: Since using Landsat alone, in the early time the data had 16-day repeat cycle and especially during the early spring from winter to spring there is probably lacking clear observation for a long time and there's no way to do the daily interpolation based on this data set. One reason is that daily interpolation could fail. A second reason, if it doesn't fail and it's interpolated, it might miss the greenup/increasing trend. Greenup detection is based on the changing trend so if there's some change that happens in the time series, it can detect it. If there's no change e.g. one data point from last fall and one from later spring, you can do

interpolation but you won't be able to detect the emergence date, the changed date. That's related to the frequency of the observation. This also highlights why it's important for HLS data to have a very frequent revisit with more clear observations it is possible to detect more greenup events.

Question: The multi-peaks in cover crop is understandable. How about the multi-NDVI peaks in main crops, e.g. soybean/Maize?
Answer: Although summer crops typically exhibit a single major peak in NDVI, they may show an extended plateau or slight fluctuations. This can occur when vegetation indices become less sensitive at high canopy density, causing saturation effects that make the signal appear as multiple small peaks rather than one sharp maximum.

Question: Are you using the TIMESAT software developed by Eklundh and Jonsson?
Answer: Feng Gao's team developed their own processing program using the Savitzky–Golay filter and a flexible fitting strategy to fill temporal gaps and smooth the NDVI time series. For additional details, please refer to Section 2.4.1, "Daily NDVI time series generation," in our recent publication (Gao et al., Mapping winter cover crops using remote sensing of crop phenology, Remote Sensing of Environment).

Question: How often do you collect information (from monocropping farmers and multi-cropping farmers)?
Answer: At the local region, Feng Gao's team collects ground information 1-3 times every week during critical growth stages

Question: How did you calculate the score for the cover crop decision?
Answer: Once the growing up date and the termination date is decided, and the the growing period of the winter cover crop is known, and then the likelihood score will be built based on the phenology NDVI or other VI metrics those metrics include the maximum value (maximum during this period) and the amplitude (difference between maximum and the minimum value) and also use the small integral that's above soil line and large integral of NDVI during this green period. Please refer to our recent publication for details (Gao et al., Mapping winter cover crops using remote sensing of crop phenology, Remote Sensing of Environment. https://doi.org/10.1016/j.rse.2026.115480.)

Near-Real Time Processing of HLS for Rangeland Management – Sean Kearney, USGS [1:52:03]

Question: How do you analyze the long-term average? (Is it annual, seasonal, or monthly?) And does it vary by region? If so, what are the reasons for this?
Answer: Feng Gao's team currently calculates the long-term average pixel-wise at a weekly interval. I have done this two ways. One is using a Landsat-MODIS fusion product created using the STARFM algorithm (see Gao, F., Hilker, T., Zhu, X., Anderson, M., Masek, J., Wang, P., & Yang, Y. (2015). Fusing Landsat and MODIS Data for Vegetation Monitoring. IEEE Geoscience and Remote Sensing Magazine, 3(3), 47–60. This gives us daily reflectance at the Landsat scale (30m) for each year going back to 2000. I first calculate the daily prediction for each product (e.g., standing biomass). I then take the weekly average across all years for each week of the year for each product. This works well for small areas, but because applying the STARFM algorithm is computationally expensive, a quick approximation for larger areas is to just use the native Landast Level 2 product (16-day). I predict each product for the available imagery, then interpolate between predictions to get a 'daily' prediction, and then average to weekly the same way as with the Landsat-MODIS fusion product. This likely results in some flattening of the long-term average and missing peaks. It could also be biased against cloudy/smokey years, but is nice because it is fast and straightforward.

Question: What's the plan for expanding beyond experimental range?
Answer: There is a lot of ongoing work in other areas and focusing on the national grasslands with the Thunder Basin National Grasslands (TBNG) in Wyoming which covers about a million acres. There are other experimental sites and producer sites where data has been collected from all these sites. The research is trying to understand what the limitations of the models are, e.g. when is it necessary to fit a new model that's region specific and when can an existing model be used that may have been developed at an experimental site, but hopefully applied to a broader region. It depends on the vegetation product, but TBNG are working at the appropriate scale of models with ground data behind them and are trying to scale up this prototype to some of these other regions. Launches will be held for the Thunder Basin National Grasslands for folks to work with there as well as others at a couple other more experimental ranches in the meantime. Unfortunately, for most of our products it does not look like a CONUS-wide model or even large regional models can be developed, because it is an empirical model. It is very specific to individual regions and the dynamics there.

Question: Is the visualization with Python for the dashboard and its capabilities?
Answer: The visualization dashboard is built on Python and a library ecosystem called HoloViews and it leverages Xarray which is this library that works well with these spatial/temporal stacks of data. This is a more data scientist friendly way to produce these visualizations. It handles all of the web development in JavaScript behind the scenes and runs on Google Cloud. It allows you to work within the workflow that you're used to if you're just processing the data. This presenter's organization typically runs everything on SCINet, the USDA supercomputer, behind the scenes. On this platform, most of the recent HLS data is

pulled, and the data is processed there, and the products are saved to disk and then pushed up to Google Cloud to actually run the web application from.

Day 1, Session 3: Assessing HLS in Various Environments [2:34:48]

Getting Started with HLS Data Access – Masha Jami, USGS EROS [2:37:15]

On Programming and Other Resources with the LP DAAC

Question: Can HLS data be accessed in AWS S3 buckets without authentication? What is the duration of the credentials?
Answer: Authentication using NASA Earthdata credentials is required to download or access data via S3 links. These temporary AWS credentials can be retrieved from your Earthdata login and are typically valid for one hour. This 1-hour limitation applies regardless of the GIS software being used (such as QGIS or ArcGIS).

Question: For researchers who are new to Python, what is the learning curve for processing HLS data, and what resources are available?
Answer: For users new to Python, it is highly recommended to start with the AppEEARS platform to handle initial data preprocessing tasks. Once preprocessed, Python can be used for the core scientific analysis. To learn how to work with the data programmatically, the HLS production team provides notebook tutorials for data access and example applications. These resources can be found at the Learning Hub and the NASA-IMPACT repository. Answered verbally in Day 1 recording at 3:29:44.

Question: What is the best way to access the data programmatically using Python?
Answer: The earthaccess Python library is highly recommended for programmatic workflows. Furthermore, data access is free as long as your compute environment is in the same AWS region (us-west-2) as the data storage. Documentation and tutorials can be found in the HLS Data Resources GitHub repository.

Question: Are there tutorials available for using the AppEEARS API?
Answer: Yes, hands-on tutorials and resources for the AppEEARS API are available.

On Data in AppEEARS

Question: If a user makes a point selection in an area covered by clouds, how does the AppEEARS platform handle this in its chart visualizations?
Answer: AppEEARS features a built-in quality decoder that displays pixel quality directly within the visualizations. This quality information is also included in the CSV outputs generated for point requests, allowing users to easily identify cloud-affected data.

Question: What is the difference between accessing HLS data on the AppEEARS platform versus Google Earth Engine (GEE)? Are they essentially the same?
Answer: From a NASA archive perspective, the Earthdata Cloud (accessed via AppEEARS) is the authoritative and most up-to-date source. While the data is available in GEE, there can sometimes be a processing delay of about half a month. Additionally, if data granules in the NASA archive are replaced or removed after their initial publication, those subsequent updates may not be immediately reflected in GEE.

Question: How does the AppEEARS platform reduce data download volumes by up to 99%?
Answer: AppEEARS allows users to extract pixel values for specific point samples or subset areas rather than downloading entire data tiles. Because many requests are for areas much smaller than a single tile, this targeted extraction drastically reduces the overall data download volume. The 99% reduction primarily applies when users extract specific pixel values for point samples rather than downloading entire datasets. However, the data reduction is often still very high for area-based requests, as users frequently request subsets or regions that are significantly smaller than a full standard data tile. If a project requires all the pixels within a tile, a larger volume of data will naturally need to be downloaded.

Working with HLS Data QGIS: Access, Visualization, and Analysis – Qiusheng Wu, UTK [3:33:19]

On Working with HLS in QGIS Plugins

Question: Does the Earthdata plugin allow clipping HLS tiles to an AOI?
Answer: For the Earthdata plug-in, it currently does not, but for GEE it does. You can say “Show cloud free imagery for a particular location and time period”. Then just say, for the U.S., clip the imagery to San Francisco. You don’t have to provide the AOI, the agent will be able to figure that out for you or you can say clip to the current extent as well. But for the NASA Earthdata plugin itself because it doesn't really have cloud computing functionality, you're

using the GDAL to run the processing on the fly. But once you store the data in QGIS, you can draw a bounding box and then use the QGIS tools to extract the raster. You can certainly do that with just a few more steps. I recommend using the agent. You can click to any region without having to provide the extent yourself.

Question: Can you select both HLSL30 and HLSS30?
Answer: Not right now because you are searching the data by a particular dataset but once you're done with one, you can switch to another one. Answered during the demo [3:33:19].

Question: For the RGB composite, is there any processing done on the bands (e.g. stretching) to improve the appearance of the image?
Answer: When the plug-in pulls the data it will try to find the base visualization parameters. You can also customize it, but the visualization tries to find the optimal compute and the optimal parameter minimum maximum or the stress or the clip for you, but it doesn't modify the original data itself. It is like the symbology of visualization. For Google Earth Engine, there is a different issue. When rendering the data in QGIS it's already after the visualization parameters so in that one you need to control the parameters in Google sending the call, or you can speak into the agent, but once the data is being rendered it is a data layer that you cannot change the settings within QGIS. You can think of it as just a base map layer so it depends on which plugin that you're using. You can have fine control over the data that you want to look at though. Answered verbally in Day 1 recording at 4:29:12.

Question: Can I do a batch download? e.g. select all HLSL30 and HLSS30 data over a certain area and a certain time period?
Answer: Yes. As shown by the time-series in the demonstration, you can make a monthly selection. If you want all the data, for example, you want all the data for San Francisco, or wherever. You can select the source by Landsat data and then create the time series and select by date, and the frequency by date. If it is by date, you're going to find all the available images or whichever was your criteria, and then you can export the data directly. Another example, within a month you might have four images or five. You download all the available images then you can switch to Sentinel-2 and do the same thing, and export the data. You can download whatever data at whatever temporal frequency. You
can download it if by date it's going to download the original data. There's no composite, but we're doing bi-weekly by monthly because there might be multiple images. From there you can do something like a median composite, or depending on how you want to specify your pre-processing and data export. Answered verbally in Day 1 recording at 4:30:30.

Question: Can you plot the profile for the stacked images where the inspector was used to check the pixel values? Will it hallucinate?

Answer: The data is based on the pixel value. It is doing the image reducer from Google, this example is using the algorithm to write it and there is not much uncertainty in it. The agent might hallucinate, but if you just inspect the pixel value, there's no Python code under the hood to run the extract pixel values. If you're just asking the agent for example, say, extract a pixel value for certain coordinates, then it might not give you the exact coordinates, or when you pass the coordinate for example, you say, I want to extract the spectral reflectance values for a tree in Seattle. Here it might hallucinate because it doesn't know the exact location. It might ask you which location and then you can provide the coordinates, but if you have a generic question it might not give you the exact result. The workflow should still work. It's going to get a location interpolate and then figure out a time span, then you filter all the images and then extract the pixel values, so the algorithms should still be valid. You always need to double check if you're using any AI in your analysis to make sure that you are getting the expected result for your particular location for your time period. This can be used for any data, any time series, and any image collection. Answered verbally in Day 1 recording at 4:33:12.

Question: Is the data produced and downloaded by GEE under license for public or commercial use?
Answer: All the global data is mostly from the public domain and open access. If you generate the data, it can be for public use, but if you want to use the compute for commercial application then you need a commercial license. For academic researchers, or for personal projects, there are not many license issues and it is free. All those data for example from NASA from different agencies are public domain data. There are private commercial data but those are only accessible via for certain users for commercial applications. Still, if you're using the computing power for commercial applications you need to have a commercial license.Answered verbally in Day 1 recording at 4:31:47.

On the QGIS Application

Question: Does the temporary credential limitation apply when accessing LP DAAC data hosted in AWS S3 buckets through QGIS?
Answer: Yes, the temporary AWS credentials required to access the data are only valid for one hour, regardless of whether you are using QGIS or another platform. For more details on accessing data and managing credentials, please review the documentation. Additionally, resources for GIBS (Global Imagery Browse Services) GIS usage can be found here.

Question: What is the recommended troubleshooting step for a dependency error ("cannot import name 'PROTOCOL_TLS'") when using the Earthdata Search QGIS plugin?
Answer: For technical plugin errors, it is recommended to submit an issue directly to the developers on the plugin's GitHub repository. You can report this specific bug and seek assistance here.

Question: Is what is shown in the demonstration compatible with version 3.44 of QGIS?
Answer: Answered during the demo [3:33:19].

Question: Do you have a similar plug-in shown in this demo for ArcGIS Pro?
Answer: No, QGIS is much easier to develop with because it is essentially just a Python script. For ArcGIS Pro you may need to develop with C sharp which works best with Windows and is not as compatible with Linux. You can also just use a scripting environment. There are other plugins that work on the web like Jupyter notebook environments. For open Geo Agent, there is currently only a working version for Jupyter notebook and also web. You don't actually need to use a desktop. From what was shown in the demo, you will soon be able to use just the web browser without installing anything, but this hasn't been released yet. Keep an eye on social media for when Qiusheng Wu releases this**.** Answered verbally in the Day 1 recording at 4:27:20.

HLS Data and Services in ESRI's ArcGIS Online and ArcGIS Pro – Leah Schwizer & Nalin Ly, EGIST [4:37:01]

Question: Why are there different methods for accessing the same datasets in ArcGIS?
Answer: The demonstrations showcased different methods for discovery, access, and use to acknowledge a wide variety of user types and their preferences when utilizing NASA data in workflows. NASA consistently receives feedback from users on how they want and need to leverage data in GIS tools. The intent was to use this feedback to highlight multiple approaches that leverage downloaded data files, metadata records, cloud repositories, or geospatial services. The goal is to meet users where they are and expand data access to meet their needs.

Question: How can multiple HLS granules be processed simultaneously in ArcGIS (e.g., to calculate NDVI) instead of working with each granule separately?
Answer: One efficient method is to add the individual granules to a mosaic dataset (a container for rasters) and apply the NDVI calculations to the entire mosaic. Users can leverage built-in raster functions for quick calculations and symbology without needing to run separate geoprocessing tools. A guide on cloud storage connections in ArcGIS is available here.

Question: In QGIS, there are plugins available that connect directly to Earthdata and Google Earth Engine (GEE) data repositories. For the ArcGIS demonstration, however, the HLS data was pre-downloaded. Does ArcGIS Pro lack similar plugins for direct data access?

Answer: The initial demonstration utilized pre-downloaded files in ArcGIS Pro without requiring additional plugins. However, users can utilize the Cloud Storage Connection and STAC (SpatioTemporal Asset Catalog) Explorer in ArcGIS Pro to search and access data directly from cloud environments like AWS or the Microsoft Planetary Computer (MPC). Users can also utilize geospatial web services to connect and use HLS in ArcGIS (or ArcGIS Online).

Day 2 Presentation Based Answered Q&A

Day 2 Session 1: HLS in Action [0:00]

HLS in NASA's Fire Information for Resource Management System (FIRMS) – Diane Davies & Jenny Hewson, LANCE/FIRMS [6:25]

On Fire Fronts & Environmental Assessments

Question: To what extent is the Normalized Difference Moisture Index (NDMI) useful or beneficial in assessing active fire fronts?
Answer: NDMI is not typically used to discern and view active fire fronts. However, False Color Composite (FCC) imagery is very useful for direct fire front assessments, particularly the SWIR, NIR, and Red composite, along with active fire detections from Landsat 8/9. While NDMI can provide informative data for fire risk/danger by helping identify vegetation that is dry or stressed, a key consideration with NDMI is the temporal lag—specifically, how quickly a detectable change in the moisture signature (resulting from vegetation stress or immediate thermal damage) manifests in the satellite imagery.

Question: Why is NDVI used considering it has certain issues such as saturation? Is it driven by user requests?
Answer: NDVI is a very commonly used index that measures the greenness of vegetation which relates to its health and density (USGS). It has been used by the wildland fire community

for decades in assessing fire danger (USFS). While saturation is a consideration, recent studies on the use of NDVI for fire risk determined extremely high or low values of NDVI are typically not associated with ignitions and wildfire growth and instead occur in the mid-NDVI value range (Billal, 2025). Additionally, in FIRMS, the refined spatial resolution and detail provided by HLS NDVI is not intended to be used as an index in isolation, but with other contextual information that is available in FIRMS to aid in its interpretation. Answered verbally in the Day 2 recording at 1:01:05.
USGS reference and fire danger viewer. Additional sources (Burgan, Robert E.; Hartford, Roberta A. 1993. and Bilal 2025).

On FIRMS

Question: What features are available in FIRMS Global? Are all those features open for use and free?
Answer: All the features shown in the FIRMS presentation are free to use. FIRMS global is the original instance with all of the same near real-time and real-time imagery. The dynamic imagery includes VIIRS and MODIS, as well as geostationary imagery, active fire detections, and some realtime products with additional contextual information that are relevant to both the global and regional context. FIRMS US/Canada, developed in partnership with the USDA Forest Service, is designed to support the unique wildland fire context of the US/Canada and has the same layout / look and feel at FIRMS Global. It also provides additional contextual data and information relevant to wildland fire in the US and Canada. FIRMS US/Canada also includes ultra real-time active fire detections, which are available in under five minutes, as well as active fire detections from Landsat 8 and 9. Also answered verbally in the Day 2 recording at 1:03:39.

Prithvi-EO-2.0 Foundation Model Applications – Venkatesh Kolluru, ODSI [24:33]

The repository for the workshop demonstration can be accessed on GitHub: Prithvi Applications Repository.

On the Prithvi Foundation Model

Question: How can users access the base foundation model versus the fine-tuned variants?

Answer: For a specific fine-tuned application like burn scar mapping, you can access the model on Hugging Face at the Prithvi EO 2.0 Burn Scars Page. To access the entire suite of base foundation models and associated weights, please visit the main organizational repository: IBM-NASA Geospatial Hub on Hugging Face.

Question: For the flood modelling and crop classification models, what is the accuracy of the Prithvi model? Please suggest agriculture related applications.
Answer: For flood detection, on the Sen1floods11 benchmark, Prithvi-EO-2.0-600M achieves 90.3% mIoU and 83.1% IoU for the water class. When the team validated operationally across 19 flood events in 12 countries against Copernicus GFM and OPERA DSWx-HLS, accuracy ranged from 76% to 97%, with F1 scores between 0.55 and 0.80.
For crop classification, on GEO-Bench segmentation tasks the model achieves around 71% average mIoU. In operational validation across 18 test regions against USDA CDL, AAFC, and European reference datasets, median overall accuracy is 0.58. Performance varies by crop type—corn and soybean mapping in the US Corn Belt achieves F1 scores above 0.84 for corn and 0.78 for soybeans.
For agriculture applications, Prithvi supports crop type classification, agricultural flood damage assessment, burn intensity mapping for farmland, and GPP flux estimation for productivity monitoring.

Question: What is the output of the pre-training step?
Answer: The output of pre-training is the encoder weights. These weights contain learned representations of Earth's surface patterns and spectral signatures. The model learns these through a masked autoencoder approach where it predicts hidden portions of input images. When you fine-tune for a specific application like flood mapping, you take these encoder weights as your starting point and attach a task-specific decoder head.

Question: If the Foundation Model has an unsupervised pre-trained encoder, can it be shared with researchers to make it customized with their own supervised fine-tuning?
Answer: Yes, the pretrained encoder weights are openly available on Hugging Face at the IBM-NASA Geospatial repository. You can download these weights and fine-tune with your own labeled datasets for custom applications. ODSI provides example scripts and notebooks to help users get started. The only constraint is access to GPU resources, but even a single GPU is sufficient for fine-tuning with a few hundred labeled samples.

Question: Does the Foundation Model produce better results with less labels? Is there any evidence to support this?
Answer: When the team fine-tuned the model with Sen1floods11, it had only 11 flood events with around 446 labeled image chips, and still achieved strong accuracies. This was the case even when the team fine-tuned the model for fires, where data was taken from only three or

four years and the results are still promising. It all depends on the quality of the data, even if they are few. In another use case for landslides and trained the model, but the data was generated using different methods and standard deviations they computed on their own. The model worked well on that specific dataset but when attempting to operationalize it there were some issues. So even though you have fewer samples, the model will perform well if the data quality is good. Answered verbally in the Day 2 recording at 1:02:23

Question: Does finetune freeze encoder?
Answer: It depends on the approach. For our flood and burn scar applications, the entire model was fine-tuned including the encoder, which allows the weights to adapt to the specific task. You can also freeze the encoder and only train the decoder head, which is faster and requires less GPU memory. Freezing works well when you have limited labeled data and want to preserve the general representations learned during pretraining. Both approaches are supported in our codebase.

Question: Would it be easy to train the Prithvi model on other datasets, like the OPERA datasets (water, disturbance, displacement)?
Answer: You can use OPERA as a label and you can fine-tune. This is how the model was tested, even though it was pre-trained on HLS imagery, specifically when the burn scar or flood model was tested. These were again fine-tuned. If you have label data from OPERA, you can relatively easily fine-tune the model. The scripts to test this are provided online. The only constraint for most people is whether you have enough GPU resources, if you have GPU assets, it is relatively easy to test. Answered verbally in the Day 2 recording at 50:51.

Question: What is the main difference between your model and Google embedding model? What makes Prithvi different from other Foundation Models being developed today?
Answer: Prithvi is the first of its kind released back in 2023. There are around 60 to 70 new Geo-AI that are being released now. People might question why Prithvi hasn’t used SAR, but this was the first model that NASA and IBM built in collaboration and released using different embedding spaces or different architectures. When the collaboration developed version 2, the teams considered all possibilities of including location information, temporal information and utilizing all available imagery. The Google embeddings from the model representations were released as an embedding dataset, but the teams have not extracted and posted these as separate embeddings. Architecturally they might be a bit different, but the teams would need to dig deeper to understand how much they differ. Answered verbally in the Day 2 recording at 1:08:12.

Question: What are some future plans for Prithvi? Do you plan to pre-train on data like SAR?
Answer: Another team is currently working on fusing SAR. They are using some adapter networks and are not going to pre-train the entire model again with SAR, because that's not

optimized or cost efficient. They will freeze the encoder part and are using U-Net based architecture to induce SAR, this is currently being tested for cloud-based gap filling and flood mapping, but the results are underway. The current focus is on training the foundation models for other planetary bodies. Answered verbally in the Day 2 recording at 1:05:41.

On Training and Access to the Prithvi Foundation Model

Question: How do I know which model to import?
Answer: The Foundation Model can be found here. When you load the model in Terratorch, you specify the model name or the checkpoint path in the configuration. Terratorch has a model registry where Prithvi models are registered with specific identifiers. So when you call the inference function, the model name you pass in the config, like "prithvi_eo_v2_300" or the path to the fine-tuned weights, tells Terratorch exactly which architecture and weights to load. If you look at the config file in the notebook, you'll see the model name specified there. That's how it knows you're fetching Prithvi and not some other foundation model.

Question: What platform do you recommend for researchers or students who are entirely new to working with Python and machine learning models?
Answer: Google Colab is highly recommended for simplifying the initial setup environment. Both the free and pro tiers provide sufficient resources to execute notebooks and get started with geospatial AI.

Question: Are there opportunities for users to participate in online training sessions regarding Prithvi models and scripts?
Answer: Yes, the team frequently organizes technical workshops. Upcoming demonstrations include the joint ESA-NASA workshop in Huntsville, Alabama, as well as an upcoming session at the IEEE HDCRS workshop in Spain (scheduled for June 2026).

Question: All these tools are great but keeping them code-based makes "Research Grade" products. Is there any effort to make easy to use GUI tools for practitioners?
Answer: Yes, ODSI is developing a web-based platform accessible at explore.fm.odsi.io. This provides a GUI interface where you can draw an area of interest on a map, select a date, and run inference for flood mapping or burn scar detection without writing any code. The platform handles data download, preprocessing, inference, and visualization behind the scenes. This was built specifically to make the technology accessible to emergency managers and practitioners who may not have coding experience.

Session 2: Accessing HLS in Various Environments [1:09]

NASA Worldview – Minnie Wong & Ryan Boller, ESDIS [1:13:11]

Relevant Resources

- HLS Tour Story in Worldview
- GIBS API docs
- GIBS Web Examples

On Applications Using Worldview

Question: What are the known limitations or data anomalies to look out for when applying HLS data via NASA Worldview to cryosphere studies and glacial lake analysis in high-altitude, mountainous environments?
Answer: One recognized data quality constraint in HLS products is the occurrence of unrealistically high snow-surface reflectance values within the visible spectral bands. Please Users conducting cryospheric research are strongly advised to consult the official documentation detailing these anomalies: HLS v2.0 Known Issues Document. Also answered verbally in the Day 2 recording at 1:59:52.

On Data in Worldview

Question: Following along with the Worldview demo, I find that S2 and L8/9 have a 2-day latency for these beta-testing layers. Will that latency improve with time?
Answer: This is the expected latency until there's a near real-time product available. It also depends if you're using the dynamically generated imagery versus the true color and statically generated, or generated in advance imagery. The dynamic imagery that was demoed is available a little bit sooner since it works directly on the data. On the visualization side, Worldview relies on the data to be produced first and then from there the imagery is generated for GIBS which adds some latency to it. The dynamic data visualization is usually available faster because it works directly on that original data and doesn't need to wait to be generated. There is a mix of latencies available. There is a low latency version of HLS being investigated and that might replace the actual HLS core, or SNWG might deliver HLS with a six-hour latency. Within these platforms, you would be able to get HLS data at six hour latency, if/when this latency is possible. For the dynamically generated imagery, you'd be able to see it as soon

as the data is available or pretty close to it once that metadata is also published. There would be an additional probably couple of hours after that before the GIBS based pre-generated tiles are available, since it does take a little bit to create the image tiles there, but that latency would be reduced based on what's available from the HLS data production side. Answered verbally in the Day 2 recording at 1:56:09.

Question: The demo mentioned that Worldview is for visualization purposes only. When using that downloaded data as a GeoTIFF, is that downloaded GeoTIFF good for research analyses or still for visualization.
Answer: The short answer is it is for visualization if you're downloading it straight from worldview. The longer answer is wanting to work more towards working directly with the data and to be able to provide users extra flexibility. So stay tuned. Answered verbally in the Day 2 recording at 1:59:04. Note that while Worldview's "snapshot" feature lets you download imagery for visualization purposes, you can also get to the underlying source data for research/analysis purposes using Worldview's "data download" feature which is located in the "Data" tab of the layer panel on the left side of the app.

Question: From a user perspective, what are the major differences between NASA Worldview, NASA GIBS, and LP DAAC?
Answer: Worldview is meant to be an easy to use front-end mapping tool that showcases all the imagery available from NASA. GIBS is the engine that drives all of those visualizations you see in Worldview, as well as any for other clients like FIRMS and many other tools which use GIBS directly. If you're a GIS user or you want to develop your own tool or you have your own mapping tool, then you would talk directly to GIBS (or Earthdata GIS services). If you're an end user who just wants to use NASA imagery, yeah you would start with Worldview. If you're a scientist who wants to download the data, you could still start with Worldview, and then download the data from there using that data download feature showcased in the demo. You can also go to LP DAAC [demo in Day 1] and then download the data directly from there. If you already know the area that you want, the time and date, you should go straight to LP DAAC and get the data there. It really depends on what you want to do. Answered verbally in the Day 2 recording at 1:53:33.

AI-powered HLS Data Analysis with Microsoft Planetary Computer – Adam Loverro, Microsoft [1:55:04]

On Data in Microsoft Planetary Computer (MPC)

Question: Seeing the impressive capabilities of Earth Copilot and MPC Pro, how seamless is the process of resolving 'NoData' from cloud cover in HLS time-series? Does the platform natively leverage multimodal Geospatial Foundation Models (fusing SAR and optical) to serve ready-to-use, gap-filled data via API, or does the end-user still need to build a custom pipeline for this fusion
Answer: The platform does not leverage any AI or machine learning models to fuse multi-modal data. Planetary Computer is fundamentally a large-scale data catalog, and Earth Copilot offers a large-language model (LLM) based front-end to query and explore that data.

Question: I have noticed gaps in the historical HLS data in the Microsoft open cloud bucket compared with what's available on AWS and EarthData. Is that intentional?
Answer: Microsoft intends to deliver a complete HLS data record from 2020 onward. If you observe any gaps within this timeframe, or you have a need for data going back further than this date, please email Microsoft: planetarycomputer@microsoft.com.

Question: In the demo, it looks like you are filtering by cloud cover but not applying a cloud mask. Should you perhaps apply cloud masks to the collection first?
Answer: In the demo, the presenter applies the cloud-cover filter in our Spatio Temporal Access Catalog (STAC) query to help narrow down the list of images that should be used in our analysis. It is correct that applying a cloud-mask filter to computing NDVI would offer a more robust approach to computing scene-level statistics**.**

Question: What is the ingestion latency for HLS products on MPC? In other words, how soon do HLS data products become available on MPC after being published by LP DAAC?
Answer: There is no exact number. Microsoft's pipelines are typically processed within a few hours. There have been instances of delays sometimes of a day or two. Generally it can be anywhere between four to six hours and 24 at the outset which is what is expected on a regular basis. If there's someone who is relying for timely purposes, please reach out at planetarycomputer@microsoft.com to get better statistics. MPC may ask if you are using, or developing, or have an application where timeliness or availability is a key concern. There are additional options under development to access Planetary Computer data in general but also HLS data specifically with SLAs's, feedback from users about this would be helpful. General expectation is anywhere from four to six hours. If there's a delay in our processing, it could be a day delay, but typically not more than that. Answered verbally in the Day 2 recording at 2:56:19

Question: There appear to be gaps in the historical HLS data hosted in the Microsoft Planetary Computer's open cloud bucket compared to what is available on AWS and Earthdata. Is this intentional?

Answer: Any intrinsic gaps in the dataset are typically based on spatial coverage and are inherited directly from the source satellite data (Landsat 8/9 or Sentinel-2) during the harmonization process.

On MPC

Question: Do you have any free training for beginners?
Answer: There are no live or recording training sessions available for beginners, but it is recommended to start here for complete documentation: Documentation | Planetary Computer

Question: Please explain the difference between MPC open and MPC pro?
Answer: Microsoft Planetary Computer (MPC) and MPC Pro are the public and private offerings of Microsofts Spatio Temporal Access Catalog (STAC) product. MPC offers users access to 70+PB of publicly available data, while MPC Pro is an Azure cloud service, which enables users to create their own STAC catalog and to populate that catalog with their privately owned data.

Question: Why do I see Claude Opus 4.6 in your Visual Studio Code. I thought you were using CoPilot? Also, why does the Visual Studio chat icon not use the CoPilot icon instead or can we change the LLM chat model that we use in VS Code?
Answer: There are many chat LLMs available in the VS code environment. The presenter's version of the tool had access to the Claude Opus and Claude Sonnet, Chat GPT Codex, Claude Haiku, and Gemini models. There are various options at different token prices, and more options for the paid tiers of the co-pilot experience in VS Code. There are other different LLMs available, not just from different vendors like OpenAI, Anthropic, and Google, but also different versions of the LLMs from those vendors. The presenter notes that Microsoft tends to bounce around between models because models are changing so frequently, and often test which one's the best. Claude Opus 4.6 was used in the demo because the presenter found it has been mostly successful with this type of task. The demo could have been running the Gemini model just as easily. Answered verbally in the Day 2 recording at 2:58:01.

Question: Is this MCP different from the MCP for managing agents?
Answer: Model Context Protocol (MCP) technology is an agent to tool interaction & communication standard. The MCP server demonstrated in this session is deployed via VSCode, and provides a bridge between a VSCode agent and Microsoft Planetary Computer (MPC) and MPC Pro. This MCP tool for MPC Pro does not help users manage multiple agents.

Question: Could you run this notebook in MS Fabric instead of locally?

Answer: The notebook from the demo has not been tested to operate in MS Fabric, but this should be possible with some small modifications. If you would like further assistance using this demo notebook please reach out to our team at mpcpro@microsoft.com.

Question: What are the computation Azure costs?
Answer: The Azure environments are free to start and Planetary Computer is free to use at moderate and research scales. It is not designed for incredibly large scales today. If you have a project where you are needing to pull terabytes, hundreds of terabytes or petabytes of data from the service reach out to the MPC Pro team. It is free to use today for computational power. It is a cloud-based service. The demos from the workshop were all local compute. In a sense that is free to use on GitHub co-pilot. There is a free tier and an intermediate tier of $10 a month. The takeaway here is that getting started learning MPC is very low cost. If you are trying to do things at scale and not on your local machine that is where the Azure costs start to reveal themselves and the pricing is generally based upon your organization agreements and whether you are representing yourself or part of a larger organization. Also answered verbally on Day 2 recording at 2:29:09.

Question: What are the main differences between MPC and GEE?
Answer: MPC and GEE are separate and different in purpose. The planetary computer resource houses growing petabytes of HLS data and is free to use globally as an accessible service. The access point for MPC is with a set of APIs with documentation accessible via the planetarycomput.microsoft.com website. The best way to access those tools is through VS Code which is a free to use IDE environment. The best way to get started was shown in the workshop demonstration with an agent, which will help dive into some of the examples you can find on the planetary computer website. What planetary computer is not, is a computing environment. It is a part of the Azure environment where you stand up your processing tasks in more traditional Azure services like Azure machine learning or Azure data bricks or Azure fabric. Examples of this are on the documentation page and from the fabric website. The combination of MPC and VS code allow you to very easily have a local coding experience and a local investigation experience with HLS data in the planetary computer. MPC Pro is analogous to the open planetary computer resource except the data storage, database, indexing and API surface area that Planetary Computer provides to the world, can provide to your organization internally. This allows users to bring together different datasets, a combination of private and public datasets to do analysis. The fundamental computing engine is the various compute resources that are available in Azure, which is distinct from Google Earth Engine, which has its own development environment and compute backend for accessing these data. Answered verbally in the Day 2 recording at 2:53:27.

HLS within Planet Insights Platform – Grega Milcinski, Planet Lab [3:00]

Relevant Resources:

- Ongoing competition regarding select scenes released from Planet's Tanager
- CSDA assessed Planet data geometric quality report

On Data in Planetary Insights Platform

Question: During the Planet demo, it showed two options under Harmonized Landsat Sentinel: Landsat and Sentinel-2. Does that mean that users can only view one portion of HLS at a time in the platform, or is there a way to work with both Landsat and Sentinel portions of HLS at once?
Answer: You are able to work with both. That filter is optional, if you want to limit yourself to one of these two. For example you might want to work with bands that are only available in Sentinel-2 or just in Landsat. If you are working with bands that are common you might want to use both of them at the same time.

Question: What options exist for university researchers to access proprietary Planet data without an active commercial license?
Answer: Planet participates in the NASA Commercial Satellite Data Acquisition (CSDA) program, which allows federally funded researchers to access select data offerings at no cost. Information regarding eligibility can be found on Planet's NASA Partnership Page. Data access for authorized users is managed through the CSDAP Earthdata Portal.

Question: The data shown in the workshop demo, is it available for all the world?
Answer: The complete HLS archive is available both from time and the geography perspective. You can access all the data immediately.

Question: Are the collections shown in the demo private or when published they are open to other users?
Answer: Here are the collections with examples and guidance on accessibility**.**

On the Planet Insights Platform

Question: What should a user do if they receive an expiration error stating that their trial account has expired when attempting to view PlanetScope data?

Answer: If you are trying to access the data under the auspices of the federal program, you must first complete the formal CSDA application process. If you have already been officially approved and continue to experience account configuration barriers, please open a formal inquiry at the Planet Support Portal.

Question: Is the Planet Insights Platform free? If not, how does payment or licensing work?
Answer: Planet Insights Platform provides free 30-days trial (only e-mail required), which provides access to most of the platform capabilities and various public data collections, including HLS.
Planet also participates in the NASA Commercial Satellite Data Acquisition (CSDA) program, which allows federally funded researchers to access select data offerings at no cost. Information regarding eligibility can be found on Planet's NASA Partnership Page. Data access for authorized users is managed through the CSDAP Earthdata Portal.
Planet Insights Platform also provides payable access, more details here.

Question: How does the Planet analysis/development platform differ from Google Earth Engine, or is it like GEE for Planet data?
Answer: Planet Insights Platforms was designed to seamlessly integrate into the user workflows that can be running outside of the platform. This means you have your own computer or virtual machine, or your own infrastructure, and you just want to get this part of the data analysis process run on the platform but then you continue with your followup steps in your environment where you also have other data you might not want to share with Planet.

Question: For me, PlanetScope is not working. It says the trial account has expired. It requests to pay a certain amount per month.
Answer: You have to go through the CSDA application process if you are attempting to access via that program. If you have already been approved and are trying to access your Planet account, please enter a support ticket here: https://support.planet.com/hc/en-us

Demonstration of HLS data in Google Earth Engine – Emil Cherrington & Kelsey Herndon, NASA EarthRISE [3:45:25] & Retrieving HLS Data Locally from GEE with computerPixels – Paulo Arevalo, Boston University [4:13:52]

The script for the demonstration can be accessed via the following public GitHub Gist: GEE computePixels Demo Notebook.

Here is the Google Earth Engine repository from the demonstration.

If you are interested in analysis ready HLS composites, contribute to this effort by filling out this questionnaire.

Relevant Resources

- GEE Data Catalog
- HLS-Landsat collection in GEE
- HLS-Sentinel-2 collection in GEE
- GEE Code Editor
- GitHub repo w/ other resources
- GEE Repository
- Great GEE documentation
- An article on CloudScore+
- QGIS plugin for GEE Data Catalogs (Dr. Qiusheng Wu demo)
- Try out OpenGeoAgent if you don't know how to code. Simply use natural language and voice commands to interact with GEE
- GEE Developers Google Group for a community forum and troubleshooting with Google developers.

Question: Does GEE allow for logic to be built atop the Engine code? For example, can someone add QA or cloud masking logic to expand upon the default masks? As in, to tweak the QA and cloud masks. It is common to need to improve masks and QA flags for development of products via secondary logic, atop the main masks.
Answer: The GEE code that was demo'd in the workshop includes examples of these kinds of functions, but you can create your own. You can use your own logic and create your own functions and tools. You can create your own masks from scratch. You can build atop of the masking functions that are built in. The workshop demo didn't show a lot of the built-in tools and functions that Google Earth Engine has, but it does have its own masking processes, or you can create your own from scratch. You can customize in Python, writing your own functions, and creating your own processes that you can run in the code editor as well. There are a lot of different functions and tools that are already pre-made for you, however. Google has produced a Cloud Score for Sentinel-2 it's relatively easy to apply that to the Sentinel 2 collection itself. In principle, you could also use that with the Sentinel portion of the HLS data set. For a detailed community tutorial on how different masks (e.g. cloud, cloud shadow, etc.) can be applied for Sentinel-2 data in GEE, refer to the GEE Sentinel-2 Cloudless Tutorial. Answered verbally in the Day 2 recording at 4:57:42.

Question: To avoid memory timeouts in Google Earth Engine during complex calculations, it was mentioned that users can extract intermediate assets and resume operations. Is there a formal name for this method, and where can documentation be found?
Answer: If your complex analytical scripts are timing out due to compute limitations, the standard workaround is to export your intermediate computational layers directly into a static asset using functions such as `Export.image.toAsset`. Once this intermediate asset finishes rendering and is saved to your account repository, you can load it into a new script to complete your downstream processing steps. Alternatively, data can be exported to Google Drive for intermediate storage workflows. While there is no single tutorial dedicated exclusively to this approach, optimization techniques can be reviewed in the official documentation: Google Earth Engine Best Practices Guide. Here is an external resource giving a practical explanation of what the recent GEE terms of service changes mean.

Question: How can users circumvent the standard 250 GB storage limitation for assets in Google Earth Engine when working saving larger files with very large files datasets?
Answer: There are multiple ways of trying to reduce the amount of space that you take. The simplest one is thinking about what you can do with your data yourself to reduce how much space takes on disk. So let's say that you're exporting some quantity that is saved as a float. One of the easiest ways to save space is to scale that by any quantity 10,000 or so and then save your file as an integer. That will save you a lot of space and it will also make a lot of things faster. Another thing to keep in mind is that the 250 gigabytes are for a cloud project. Let's say you have different projects, maybe one project about fire, one project about land cover. The more separate you keep those things, you can also use different cloud projects for each that will also help you manage your space limitations a little better. In instances where your project exceeds the default quota, you will need to directly contact the Google Earth Engine support team to formally request an allocation uplift for your account resources. Answered verbally in the Day 2 recording at 4:56:11.

Question: Does a native MATLAB library exist to access the Google Earth Engine API?
Answer: No, Google Earth Engine does not currently feature native, dedicated libraries for MATLAB in the manner it does for Python or JavaScript.

Question: Is there a way to interact with Google Earth Engine for users who do not have extensive experience writing code?
Answer: Non-programmers can explore OpenGeoAgent, an open-source tool designed to let users execute commands and interact with Google Earth Engine using natural language and voice instructions: GeoAgent GitHub Repository.