



Part 2 Question & Answer Session

Please type your questions in the Question Box. We will try our best to get to all your questions. If we don't, feel free to email Amita Mehta (amita.v.mehta@nasa.gov) or Ryan O'Shea (ryan.oshea@ssaihq.com).

Question 1: How much does spatial resolution affect the processing time of the MDN?

Answer 1: I am trying to use MDN to estimate chl_a from Sentinel-2 images and it took 170 minutes to process one West Lake Erie image (10m res). Whereas for SF Bay (350m res) in the MDN tutorials, it took less than 5 mins. Most of the variability comes from the number of aquatic pixels in the given image. So for the same area, MSI would have ~900 times more pixels (300 m/10 m)², and take that much longer.

Question 2: What is the difference between PC & TSS?

Answer 2: As defined on Slide 20 in the presentation, PC stands for Phycocyanin, which is a pigment specific to cyanobacteria biomass, a potentially toxic algae. TSS stands for total suspended solids (such as sand which has been suspended in the water column).

Question 3: Is there any reason to select 50/50 Training and validation dataset? We usually consider 80/20 or 70/30.

Answer 3: While it is common practice in machine learning to use splits such as 70:30 or 80:20 (train:test), the characteristics of this dataset motivated a different choice. Sampling is uneven across aquatic conditions and geographic regimes, which can lead to instability in performance estimates when the test set is small. We have a large dataset so we are able to use 50:50. We use a regional split to test uncertainty.

To ensure statistically robust and representative validation across the range of observed conditions, we adopted a 50:50 split. Given the overall dataset size and the diminishing gains observed in model performance beyond this threshold, allocating half the data to independent evaluation improved the stability and reliability of generalization metrics.



In addition to the standard hold-out validation, we performed complementary experiments (e.g., regional leave-one-out analyses), which further confirmed the ability of the MDN to generalize to previously unseen regions.

Additional methodological details are provided in O'Shea et al. (2023) and Balasubramanian et al. (2025).

Question 4: Just to be sure, STREAM uses TOA radiances from satellites as input to the MDN model to derive water quality parameters?

Answer 4: STREAM has two different MDNs, one MDN retrieves Rrs (Remote sensing reflectance, or ocean color) from the TOA radiance, and a second MDN retrieves WQPs from the Rrs.

Question 5: The MSI scatter plots show a "classic" Heteroskedasticity" pattern of residuals. There are statistical corrections that can account for this and provide better predictions. Has anyone tried that? The residuals are CLEARLY a function of the X-variable.

Answer 5: The MSI plot should not be interpreted in the classical sense. In this case, the x-axis does not represent an independent predictor variable. Model predictions are generated from the input RS whereas the x-axis represents independent, in-situ measurements used solely for validation. We acknowledge the presence of systematic bias and range dependent error behaviour in certain areas of the dynamic range of the water quality variables. This is an active area of ongoing model refinement.

Question 6: As Modis temporal resolution is higher than Landsat or Sentinel, is there a way to use it as a support tool that works wherever there are cloud or glint effects and getting a higher amount of images?

Answer 6: Since we use MDNs trained on similar underlying datasets, we have seen that the retrievals of parameters are relatively harmonized retrievals between sensors, so yes you could use chlorophyll from MODIS to fill in for alternative sensors. Long term we are attempting to further improve harmonization between sensors (Landsat, Sentinel 2, and Sentinel-3) by leveraging more advanced machine learning models (e.g., domain adaptation), as a part of the Landsat Science Team. Yes, you can use the models to fill in and build time series.

Question 7: Can Planet data be used for MDN or is among the coming soon?

Answer 7: We do have a MDN available for Chla and Secchi Disk Depth from Planet Superdove: <https://doi.org/10.3389/frsen.2025.1624783>. It is not currently supported



as a part of the tutorials or Github. However we plan to add Planet, as well as other sensors we have models for, to the tutorials soon. Feel free to email me, and when we push the updates to our Github I can let you know.

Question 8: Have the MDN results been compared with the derived products from ACOLITE?

Answer 8: No, we have not done that direct comparison. We do compare against other operationally available algorithms (OCV4, CyAN) and show improved results across a wide dynamic range in Lake Erie and the Chesapeake Bay. We have seen the MDNs outperform.

Question 9: Can this code be used anywhere? Do we have to change the input which depends on the zone that we study?

Answer 9: Yes. The model relies only on Rrs (remote sensing reflectance, ocean color signal) as input, so it can be used in any region.

Question 10: Currently, I am working on estimating TSS using surface reflectance data. Since high TSS in situ data are not available for training, my model predicts low TSS values during extreme events, even though I know the actual TSS concentrations should be much higher than the estimates. Could you please suggest how to address this issue?

Answer 10: If you did have high TSS values with Rrs available, you could retrain the MDN.

Question 11: Is it possible to fine-tune MDN using in situ data?

Answer 11: Yes, It is possible, using transfer learning. We do not currently support that via our tutorials. Please send me an email if you want to discuss how to do this yourself. We may have some code for this.

Question 12: The uncertainty maps seem to look very similar to the water quality maps. High WQ estimate = high uncertainty. Does that mean that the model can be less reliable for higher values?

Answer 12: The uncertainty metrics shown on the slide represent the bulk uncertainty, i.e., the absolute predictive uncertainty associated with each pixel, expressed in the same units as the model output. As expected, absolute uncertainty increases at higher chlorophyll-a (chl_a) values due to the larger dynamic range of the dataset.



However, when uncertainty is expressed in relative terms (as a percentage of the predicted value), it remains broadly consistent across the full range of WQ parameters.

Extreme values in relative uncertainty are particularly informative, as they help identify pixels that are potentially problematic or predicted with low confidence.

Question 13: Can Uranium contamination in surface water be estimated?

Answer 13: I do not have experience with mapping Uranium contamination.

Question 14: Do you have any guidance or recommended resources for preprocessing the Rrs inputs? Specifically, I'm interested in MSI and OLI.

Answer 14: From our experience, ACOLITE works well in eutrophic waters, L2GEN/SeaDAS works well in more oligotrophic waters. We are producing an Aquaverse correction methods and will make it available.

Question 15: In the sample dataset histogram shown earlier, what is the temporal range and spatial extent of the data?

Answer 15: Similar to that of the GLORIA database. We have used it to extend with our in situ data. Spatially most of the data (80%) is from

Question 16: Is there no atmospheric correction embedded in the MDN? Do you suggest just starting with L2 data?

Answer 16: Yes, The MDNs available in the tutorial.

Question 17: Is it possible to apply these models to earlier versions of Landsat, such as Landsat 5 or 7? What would the limitations be?

Answer 17: I believe we have applied some of these models for Secchi disc depth retrieval, from Landsat 5 & 7.

Question 18: Are there any matchup studies for MDN-PACE predicted TSS and Chl-a?

Answer 18: Yes, we have but not yet published. We are funded as part of the PACE science and applications team and are working on doing this as part of that work. We have an in-house validation data set from the PACE data validation team. Short answer, we're working on it but they're not published yet.

Question 19: What is the accuracy of the MDN network?



Answer 19: There are multiple accuracies you can report depending on the sensor. It depends on the sensor's spectral capabilities. We have publications on accuracy that are cited at the bottom of the tutorials that you can look through. We try to publish accuracies for different scenarios for sensors that we support.

Question 20: Is there a reason why you did not compare with other available machine learning algorithms?

Answer 20: We used operationally available algorithms but are open to compare with others. While the primary focus of this work has been comparison against operational WQ algorithms, the MDN framework has also been evaluated against other state-of-the-art machine learning approaches in separate studies, considering both predictive accuracy and the validity of uncertainty estimates.

These comparative analyses indicate that MDNs achieve performance comparable to leading ML/NN models. In particular, MDNs tend to exhibit lower median error, while mean error can be comparable to or slightly higher than classical point-estimate neural network models. Importantly, MDNs provide statistically meaningful predictive uncertainty estimates.

Representative comparative studies include:

[1] Saranathan et al. (2024), *Frontiers in Remote Sensing* — Assessment of advanced neural networks for the dual estimation of water quality indicators and their uncertainties.

[2] Werther et al. (2025), *Remote Sensing of Environment* — On the generalization ability of probabilistic neural networks for hyperspectral remote sensing of absorption properties across optically complex waters.

While these prior studies establish the MDN's strong performance, additional comparisons to other ML/NN baselines could be included in future work as part of extended validation.

Question 21: Is the same network used for all sensors?

Answer 21: Yes, Multi-spec we try to use the same network.

Hyper-spectral (EMIT, PACE, PRISMA) we have to increase the number of neurons in the network. MDNs are used for both.



Question 22: I didn't understand that one MDN retrieves Rrs for IOP, ocean Color properties and the second MDN retrieves WQ parameters. Isn't it supposed to be the same?

Answer 22: For each sensor, there is diff spectra. Multi: Chl-*a*, TSS, and CDOM. IOPS are added to the hyperspectral sensors. In each case a single MDN is used per sensor.

Question 23: Which bands are the network using for Chl*a* and the other parameters?

Answer 23: It will depend on the sensor. We try to keep using all the bands between 415 - 750nm.

Question 24: Is there a document that lists the water bodies where the MDN is trained and validated?

Answer 24: We have a wide set of pubs available cited at the bottom of our tutorials. Spatial distribution is included. See: <https://doi.org/10.1016/j.rse.2023.113706> (especially Fig. 7).

Question 25: Can it distinguish between oil spills and HABs over oceans?

Answer 25: We have not performed this analysis.

Question 26: How do NASA ocean-color and inland water quality algorithms manage seasonal spectral shifts caused by monsoon-driven turbidity, CDOM variation, and changing solar geometry without introducing systematic bias in long-term time series?

Answer 26: The change in solar geometry is captured in the first MDN, which retrieves Rrs from TOA radiance. It was trained with hundreds of millions of samples with variations in the sensor and solar geometries. Turbidity and CDOM variations are dealt with by the MDN which retrieves water quality parameters (e.g., Chl-*a*, TSS, and Secchi disk depth) from Rrs. It was trained with ~5,000-10,000 different combinations of water quality parameters, with variations multiple orders of magnitude variations in TSS, CDOM, and Chl-*a* concentration. That said, we would expect adding additional data from the region into training would improve overall accuracy.

Question 27: Why was the in situ chlorophyll *a* data from USGS which covers the United States not used?

Answer 27: If Rrs are available with chl-*a*, we would happily include them for training future models. If they are just chlorophyll-*a* data, then we can use them for validation.



Question 28: Does STREAM recommend regional recalibration of retrieval algorithms, or does it rely on globally parameterized models with deviation correction layers?

Answer 28: We do not recommend regional algorithms. We would be happy to add data during model retrainings.

Question 29: Can the MDN models be used for volatile and non-volatile compounds as part of water quality predictions? In the case of nitrates, phosphates, ammoniacal nitrates and other ions, could MDN be used? If yes, in general practice how many in-situ observations are we looking here to train MDN models?

Answer 29: MDN with multi/hyper can only be used for products that change the ocean color signal. If you have an Rrs dataset, we could try to train using it. As little as 1000 samples have been used for in situ training.

Question 30: Can MDN model run on wetland water quality monitoring?

Answer 30: If the wetlands are optically deep, they should work for retrieving WQPs. But if the bottom can be seen, it will affect the retrievals.

Question 31: Should input L2 data be gridded before using the MDN? Or, does it not matter?

Answer 31: No, it should not matter. The MDN takes the Rrs as input.

Question 32: How was the Rrs obtained for Gloria and other datasets used. were they downscaled from any satellite data to the location of the sample site? If so can I do similarly for my in situ dataset which only have chlorophyll-a

Answer 32: Gloria dataset has hyper Rrs and a few diff in situ products. We take the hyper Rrs and apply the relative spectral response function for the sensor we are training the model for, and then retain the models. You will need ChlA WITH Rrs.

Question 33: I am currently working on a coastal water quality study area. This area covers a small region, and I want to use STREAM, but STREAM does not support this area with data. What can I do?

Answer 33: Please email me!

Question 34: Are these data accessible in Google Earth Engine (GEE)?



Answer 34: We have not included these data in Google Earth Engine, and do not currently have plans to.

Question 35: I'm from Mexico and I would like to use this tool in my country. But I saw it's only available in the U.S..... is there a possibility that we can use in Latin America?

Answer 34: For STREAM, if there is a specific site you are interested in, contact us. For MDNs, as long as you have Rrs, the model can be updated.

Question 36: In highly turbid Indonesian tropical river systems such as the Kapuas and Mahakam in Kalimantan, where CDOM and suspended sediments are both elevated, how does the MDN handle spectral ambiguity between chlorophyll-a and CDOM absorption in the blue bands?

Answer 36: The MDN is treated as a black box for our purposes, we do not analyze how the model differentiates. It is based on the input data, so if similar data are included during training, it should work well.

Question 37: I am wondering how these tools could be applied to sturgeon population reintroduction from an upland caviar farm.

Answer 37: I am not an expert on this area. One of our main products is chlorophyll-a because it serves as a proxy for aquatic plant biomass, which serves as the base of the food chain. Perhaps the location of optimal reintroduction could be selected using the chlorophyll-a maps. With our STREAM tool we will be including timeseries information, which may be used to identify climatology or seasonality, to predict the best timing for a given location (perhaps the typical timing that chlorophyll-a concentration increases).

Question 38: Is the training sample 80% Europe and 20% the rest of the world representative of the model? (Data or model use cases)

Answer 38: 80% of our in situ dataset is from North America and Europe and 20% is from the rest of the world. We are trying to increase representation from the rest of the world. If you have data you would like us to include in future retrainings, we would be happy to work with you to include that data in our model training, which should increase accuracy for your given regions.

Question 39: I had issues trying to install SEADAS a few years ago. I'm wondering what type of system/computer you recommend for easiest/smooth running of that software.



Answer 39: I find it works well on Linux systems, though I think Windows and Mac are supported as well. It should work on most laptops/desktops, of course speed will depend on system characteristics. For help installing this software on your system, you may want to reach out to the Earthdata Forum.

Question 40: I see you are using python for the coding, do STREAM and the MDN models function using R coding as well?

Answer 40: No, we only support python for the MDN.

Question 41: When you were pulling the Rrs, the image seems like the waves are also filtered out. I didn't notice what filtering is used to exclude these water movements/waves (if what I am seeing is waves).

Answer 41: We just use the atmospherically corrected aquatic pixels, we do not further filter the data. The pixels that get converted in this case were identified by ACOLITE as aquatic Rrs. Since this is OLCI imagery at 300 m, I don't expect these to be waves. The section in the bottom left is clouds, which are filtered based on the brightness in the NIR.

Question 42: Do you know of an open source data sensor that would cover the Fraser River with hyperspectral data?

Answer 42: PRISMA may be able to sense the Fraser river, since it has ~30 m resolution. The Fraser river is right at the edge of what is measurable by EMIT (90 m resolution).

Question 43: What method is used to produce STREAM Secchi depth data? Could you provide materials or references?

Answer 43: The SDD is recovered from MDNs with the other products. The exact model is not yet published. An example MDN retrieving SDD is shown in this publication: <https://doi.org/10.1016/j.rse.2023.113889>

Question 44: The global data set of in situ measurements could be better distributed. Are there any plans to extend MDN training with additional in situ datasets from parts of the world, which were previously not covered? If there are plans, how do you plan to access those in situ data?

Answer 44: We try to add all of the data that we have available to us. We plan to update the models as additional data is provided to us. We are happy to include datasets from other regions, please email me!



Question 45: Atmospheric correction is problematic over lakes due to high amounts of CDOM and turbidity and may introduce distortions. The STREAM app seemed to use a custom correction based on MDN. What do you suggest users do? Could just Level-1 data be used?

Answer 45: What we have seen in these regions, ACOLITE seems to work very well. STREAM uses a custom correction based on the MDN. We will provide further clarification.

Question 46: I noticed on the GitHub with MDN tutorials there are three python files, could you describe more on them? Should we focus only on the 3rd file?

Answer 46: There are 3 files and the third was covered today. The first covers individual estimates on in-situ Rrs. The second focuses on uncertainties. We will post an update for the first two tutorials.

Question 47: Does the input data format have to be netCDF?

Answer 47: For the functions we currently use, yes netCDF is required. If you pull in the data using the same function, then no you do not have to use netCDF.

Question 48: Can we use Landsat 5 for historical data?

Answer 48: We currently have Secchi Disk Depth models for Landsat 5, but we do not support chl_a or TSS products from Landsat 5.

<https://doi.org/10.1016/j.rse.2023.113889>

Question 49: 1. If STREAM were scaled to autonomously monitor every inland lake and coastal region on Earth in real time, what technological bottlenecks – satellite capacity, AI, or data infrastructure – would have to be broken first?

Answer 49: The largest current bottlenecks are compute resources and global in situ data availability.

Question 50: Can the model use more than one sensor data type input at once?

Answer 50: Currently it only uses Rrs of one sensor as input at one time.

Question 51: Is there a difference in model performance between biogeographic regions?

Answer 51: Yes, model performance is generally linked to data availability. We expect the highest performance in regions where we have lots of data (e.g., Lake Erie) and worse performance in optically complex regions where we do not have any representative data.



We have done a regional leave-one-out analysis, to see the impact of leaving a region out of training on the accuracy of the models. On average they achieve about 50% error for retrieving chlorophyll-a, but this varies substantially regionally: Fig. 7, <https://doi.org/10.1016/j.rse.2023.113706>

Question 52: I noticed the range of limits for the IOP/OAC were the same, 0.01 to 100. Shouldn't this be different for each parameter, Chl-a conc (ug/m3), TSS (mg/l), and others considering the way of representations are different? For determining uncertainties, did you use log scale , anti-log scale, gaussian distribution or any other?

Answer 52: For plotting the WQPs/IOPs users are able to change the range limit for the different products, via the 'pred_ticks' argument. In the example tutorials the colorbars have different ranges (.01-100 for TSS, .1-10 for CDOM, in section 4: https://github.com/ryan-edward-oshea/MDN_tutorials/blob/main/SC_3_MDNs_with_satellite_imagery.ipynb. The uncertainties are estimated from the output Gaussian layer, best covered in tutorial 2: https://github.com/ryan-edward-oshea/MDN_tutorials/blob/main/SC_2_MDN_product_uncertainty.ipynb or <https://doi.org/10.1109/TGRS.2023.3234465>

Your point regarding the differing dynamic ranges of the output variables is well taken. To address this, we apply scaling to the target variables prior to model training.

In our WQ models, this typically involves a combination of logarithmic transformation followed by Min–Max scaling. This ensures that all output variables are presented to the model on comparable numerical scales, which improves training stability and prevents variables with larger magnitudes from disproportionately influencing the optimization process.

Question 53: If we applied STREAM's methodology to extraterrestrial environments, say subsurface oceans on Europa, what adaptations would be necessary to make water quality monitoring possible off-planet?

Answer 53: In situ measurements would be ideal for obvious reasons, but all jokes aside you would have to make sure that the parameters would be relatively similar for extraterrestrial adaptations.



Prior to applying models trained on collocated terrestrial in situ measurements of Rrs and WQPs, it is essential to verify that the underlying relationships between reflectance and water quality parameters are comparable to those observed on Earth.

There are several reasons why this assumption may not hold. Differences in chemical composition, optical properties, and scattering/absorption regimes in extraterrestrial oceans could alter the Rrs–WQP relationship. In addition, atmospheric correction approaches and the associated error characteristics of extraterrestrial processing chains may differ substantially from those used for Earth observations.

Given these potential discrepancies, such models may not be directly applicable without at least some collocated Rrs/WQP pairs from the target environment to validate or recalibrate the relationships.

Even in cases where differences are moderate, more advanced machine learning strategies—such as domain adaptation or transfer learning—should be considered to leverage knowledge from terrestrial datasets while tailoring the model to the extraterrestrial application.

Question 54: You mentioned that the STREAM platform currently focuses on the United States and select regions. We are a Canadian-based company specializing in Hyperspectral data processing and might be interested in hosting the model to provide prediction to Canadian freshwater and coastal waters. Under what license are the MDN models and code released, particularly for commercial use?

Answer 54: It would be under the MIT license on the Github page and is available for commercial use.

Question 19: Do you find in terms of the overall water ecology, overall parameters follow a seasonal or daily pattern over time? Like can you pick up on salmon days of arrival etc?

Answer 19: If we look at regions such as Lake Erie, they do have seasonal parameters and have certain timings of various parameters.

Question 20: Inland and coastal water quality have different factors so different observations. Does the model apply for both?

Answer 20: Yes, the same model is used for both inland and coastal water bodies. The model has used a variety of water bodies for training.



Question 21: I am from Suriname and we are having a coastal problem, where the vegetation at the coast is being destroyed, is it possible to find out if it is the wave or sediments with remote sensing?

Answer 21: For waves, ocean surface height can be monitored using satellites and tools. We will look into this further.

Question 22: Is the model evaluated and calibrated with in-situ measurements?

Answer 22: We train the model with in-situ measurements only. We take one dataset for training and exclude the other in a method referred to as “leave one out”.

Question 23: Assuming the same, for example chl-a concentration, reflectances from coastal water are likely to be different from fresh water. Is this considered when MDN training/validation data sets are constructed?

Answer 23: We lump in in-situ data from different regions. We have over 10,000 chl-a measurements alone.

Question 24: STREAM and Aquaverse rely on satellite imagery and MDN models to estimate water quality parameters. How do you quantify and propagate uncertainty across the entire workflow from raw sensor calibration, atmospheric correction, and MDN parameter retrieval so that decision-makers can distinguish between algorithmic noise and true ecological signals in near-real time?

Answer 24: We were recently funded to include uncertainties in our measurements over time.

Question 25: Is the model more accurate with some sensors more than others?

Answer 25: Yes. It depends on how the atmospheric correction is applied. There are uncertainties when comparing satellite data (MSI and OLI data for example) to in-situ data.

Question 26: I was doing research on an integrated approach using In-Situ measurement, Drone measurement and Satellite imagery. Is there any opening or call for PhD?

Answer 26: There is not one in our group.