

Machine learning with spatial data seminar.

Questions

CSC general and AI

- Is it always one GPU per one node? What is a node?
 - Node=computer, it's just a fancy HPC term for one computer in the cluster. CSC's Puhti and Mahti supercomputers have 4 GPUs per node (computer), and LUMI technically has 4 as well, but since these are dual-chip cards, it's essentially 8 GPUs per node on LUMI.
- I can also recommend using the "pen" tool instead of pure mouse cursor while presenting. (In my experience, this significantly increases the ability to convey information). It is PowerPoint instrument. The icon is visible in the lower left corner. Thank you.
 - Thank you for this suggestion, we will check for next power point presentation after the break. And I forgot about it. :((noted to self for next time)
- A general question: Could there be a potential in gathering training data from different projects? Or is there already an initiative targeting this?
 - This could be discussed eg in the framework of the Geoportti project. Thank you for bringing this up.

Leon-Friedrich Thomas, UHEL, Deep Learning based Semantic Segmentation Workflow for the QGIS Plugin EnMAP-Box

- How is the training/validation dataset 'real' land use class defined? It doesn't require human eye? Have you tested results against results where data is validated by humans?
 - The data in the Screen shots were just place holder - sorry if that caused confusion. I did not talk about any data in specific and just presented tools to process data or to train a model on any raster image data. The plugin I presented does a fully supervised semantic segmentation. So for training a model your training and validation data has to have annotated rasters, for the model to learn the features. This is the case in all fully supervised learning approaches for deep learning models. to test a model you would need to measure the model performance on a test dataset. Which can be done with the deep learning inference modul of my app. How well the model will perform depends on the quality of the data and annotation you used for training.
- Is the plug-in already available for others to use?
 - Not yet. It is open source and as soon as its ready it will be available with an update of the enmapbox - this update will be probably between mid-end of 2024.

- Could the plugin be used with Sentinel-2 time-series? For example, a simple two-date stack of images.
 - The plugin is designed for mono-timestemps. So if you want to use pretrained weights, this is not possible. If you use No pretrained weights you c
- How many bands can input raster have? Is there limitation to number of classes to be classified at once? Is unet available?
 - There is no limit of band number and no limit to class numbers. And yes the whole Deep learning trainer module is a 2d unet. So 2d unet is the obly model you can use. and you can chose a backbone of your choice.

Janne Mäyrä, Syke, geo2ml: python package for converting geospatial data to ML formats

- What is the size of the tiles and in what format are they saved?
 - Size can be specified by the user and they are saved as geotiffs. Default values are 256x256px for semantic segmentation and 320x320 for COCO and YOLO. Those functions also save the chipped labels as intermediate files, and that can be specified to be either geojson (one folder with multiple geojsons) or a geopackage (one geopackage with each chip as a layer). I recommend geojson as it is at the moment a lot faster.
- Question to audience. Should this be added to geoconda module in Puhti?
 - Why not, if it is not too much effort.

Imangholiloo (online), UHEL, A New Approach for Feeding Multispectral Imagery into Convolutional Neural Networks Improved Classification of Seedlings

- Do you use lidar for canopy height treshold?
 - We used UAV-RGB photogrammetric point clouds to create CHMs, but indeed LiDAR could be another (better) option.

Maria Yli-Heikkilä, LUKE, Agricultural monitoring with satellite remote sensing

- In what form did you use Sentinel-1 data?
 - Sentinel-1 was a composition of 11-days S1A images. The mean per pixel.

Parvez Rana, LUKE, 3-D modeling of biodiversity and ecosystem services

- Not a question, but a comment. This is very interesting and timely topic There is a hackathon organized about similar biodiversity challenge next week in Helsinki, Beyond Dead Wood by Ultrahack.

Pyry Kettunen, FGI, Vector ML - case Gaze-Aware Interactive Map System

- how do you transform polygons into tensor?
 - Coordinates are entered into a tensor (a multi-dimensional matrix). My 4 th slide shows the basic principle.

Jorma Laaksonen, Aalto, Hyperspectral superresolution and multi-image fusion for remote sensing forest studies

- What algorithm did you use to detect/segment trees in 3D?
 - That is a novel algorithm developed in our research group. The work has been submitted to IEEE IGARSS 2024 symposium and will hopefully appear there.
- How do you do super resolution?
 - We studied a number of methods. The work has been submitted to IEEE IGARSS 2024 symposium and will hopefully appear there. In any case, it will be available within a month in <https://aaltodoc.aalto.fi> as Yuvrajsinh Chudasama's M.Sc. Thesis. TL;DR <https://ieeexplore.ieee.org/document/10044141> was the method that gave the best results.
- Do you have any results for tree species classification? What are the accuracy metrics and what is the legend (classes) of the classification?
 - No published results yet with the current data and code. Our earlier results with other data and codebase can be found in <https://ieeexplore.ieee.org/document/9673792> .

Emma Turkulainen, FGI, Monitoring Bark Beetle Damage: Technical Insights into UAS Image Classification with Deep Neural Networks

- Where do you run Agisoft Metashape? What kind of resources does it consume for your process?
 - I have not used Agisoft Metashape personally, as this step of data processing has been done by a different team member. The software is run on a local computer. I can't answer the question about resource consumption at this moment but will come back to this once I have the answer.
 - Kylli: Agisoft Metashape licensing makes it easier to run it in cPouta cloud than supercomputers. Set up instructions by LUKE: https://github.com/csc-training/geocomputing/tree/master/pouta/metashape_with_VNC
- have you tried feeding in spectral indices instead (or in addition to) the images themselves?
 - In this study using convolutional neural networks we have not been using spectral indices in addition to the spectral values of the images. However we have also studies about bark beetle monitoring using traditional machine learning methods such as random forest, where spectral indices have been used.

Andras Balazs, LUKE, Predicting tree species specific growing stock using ALS point clouds and a 3D-Convolutional Neural Network

- Have you tried also drone lidar data (or other more dense data)?
 - I would love to use drone, or more dense data, but the CNN requires around a thousand samples to converge. If there is such a dataset I'd very much like to know about it!
- Nice interactive(?) point cloud viz! How did you get powerpoint to do that? :)
 - Thx :) I've prepared the plot using plotly package in R, published the plot in my plotly profile and used the Web viewer add-in in PowerPoint to display the plot. I'm happy to help if it's needed.

Ilja Vuorinne, UHEL, Mapping invasive plant species using airborne hyperspectral data and weakly supervised deep learning

- do you think that extending the training data to cover the crown edges (and not just the crown middle) might help improve detection?
 - Thanks, that's a good question. For the shallow models it is likely to reduce the accuracy, but about the deep models I'm not so sure. The best current model doesn't seem to have major problems to detect mixed pixels at the edges of canopies. However, including them in training could still help and should at least be investigated.

Jere Raninen, FGI, Building detection from Lidar data using 3DUNet

- could you include missing points in the training? they seem quite recognizable
 - We have not tried any method for this, but I have been thinking about that as well.

Emilia Hattula, FGI, How the diversity of training data improves the 2D UNet building detection

- What do you mean with 'hard features'?
 - For example bridges, large buildings, containers etc. have been hard for the model to learn.

Discussion

- What are your usual bottlenecks in Geospatial ML? What part of the process do you spent most time with?
 - Data labelling/annotation. New approaches or workflows to increase the efficiency of data labelling.
 - Geospatial data is often/always also temporal, so there can be severe mismatches between the raw data and its human annotations that need to be accounted for. This is time and human resource consuming.

- Anything to improve with Puhti/Mahti/LUMI installations or documentation?
 - In LUMI, it would be better if projects could be continued instead of closing old ones and initiating news ones.
- How would you summarize the status of automated data annotation/labeling, what is the current 'leading tech' for this and how good are the results compared to traditional manual labeling? In general, CSC builds infra that has wider benefit for researchers, and this could be one topic to bring forth/improve
 - Segment-anything and it's variants are somewhat used, at least I've been surprised how well they work. Main advantage is that instead of drawing masks, you can just prompt it with points or bounding boxes and have it draw the mask for you. There's also segment-geospatial -library. Downside is that the models are somewhat heavy, but in CSC environment thats not a problem