Q & A – Session C

DISCLAIMER:

The information, views and statements presented by speakers on this Q & A document are solely those of the speakers and do not reflect the views of PGO nor do they represent explicit or implied endorsement by PGO.

Q & A – Session C

What are some common sources of uncertainty and error when applying advanced machine learning techniques to mineral exploration, and how can geoscientists effectively mitigate these risks?

Data prep/collection:
Analytical Errors – With geochemical data from a single lab this is not too much of a problem as the detection limits and errors should be fairly well described (and can be easily verified by just plotting up the data). Need to be very careful comparing data between different labs or different techniques. It's best to mess with the data as little as possible, but if you really need to combine two relatively disparate datasets, you could try normalizing the data independently or apply a clip to get the ranges in line.

Random noise – Common for time series or spatial geophysical data. This noise can often be averaged away or dealt with through oversampling. You'll lose some resolution but it will bring the noise level down.

Representivity – We tend to oversample the ore body and not the host rock, so we're often dealt with an imbalanced dataset when trying to model/analyze data. Best way to mitigate this is to try and sample in a representative way from each lithological unit. Don't forget about the host rock! This is a particular problem for geophysics as density/conductivity measurements are often taken of the target rock but not of the host.

Machine Learning (ML):
Some techniques will provide a natural estimate of uncertainty. In the supervised case, as long as you hold back a test set (should be standard practice!), you can get an estimate of the accuracy of your prediction. If it's a classification problem, a confusion matrix is an easy way of displaying this. Some of the unsupervised routines also provide a way of estimating how good of a fit a single point is (e.g. distance to the nearest cluster center in a SOM). In both the supervised and unsupervised cases (but more so in the latter), you can run the data through a variety of algorithms. If the result is more-or-less similar using several techniques you can at least say that the answer you obtained is robust to perturbation. A similar exercise can be conducted by modifying hyperparameters and seeing how things change.

Once you're into implementing an algorithm, it's more about clearly communicating uncertainty and error than controlling for this as the data are what they are at that point.

Q & A – Session C

Do you know of educational opportunities within the industry to help ensure that geoscientists have the necessary training and expertise to effectively utilize advanced machine learning techniques in their work?

Aside from short courses that are put on at events like PDAC/Roundup, I'm not aware of any widely advertised training courses that are specifically tailored to geoscientists. I'm a proponent of self-study and some of the material available on Coursera as an easy onramp. You can't go wrong starting with "An Introduction to Statistical Learning" (**https://www.statlearning.com/**). This will get folks familiar with the background of a lot of the techniques that are already implemented in a lot of the commercial software packages. If folks would like to move on to practical implementation, there's a lot of great books available (e.g. "Hands-On Machine Leaning with Scikit-Learn, Keras and Tensorflow" **https://www.oreilly.com/library/view/hands-on-machine-learning/9781492032632**).

How important is a sensitivity analysis with ML and Large Language Models (LLM)?

I think, in general, sensitivity analysis is the analogue of accuracy in the unsupervised learning case (there's a great example with t-SNE **https://distill.pub/2016/misread-tsne/**). In that regard, it's an essential metric. Depending on temperature settings, LLMs are also very non-deterministic. At the very least, running iterations of similar prompts through the same LLM and comparing results is a useful way to understand the sensitivity of output. What metric to use to compare these is still a pretty healthy subject of debate, but nothing really beats having a human look at the outputs and rank them.

Q & A – Session C

How are geoscientists able to benchmark against industry standards for example CIM/PGO, etc? Are there guidelines available?

This is something I should be more aware of, but unfortunately, I don't think that there is a lot out there at this point that is specific to machine learning and computer science within geoscience. Geoscientists Canada is now working on some guidance for Staff and Volunteers on the use of Artificial Intelligence tools, though I suspect this will be more focused on appropriate use of generative language models.

PGO does address this in point 18 of the GENERAL PROFESSIONAL PRACTICE GUIDELINES FOR ECONOMIC GEOLOGISTS:
The P.Geo. is encouraged to integrate new technology or innovative techniques to advance the work and the project. The supervising P.Geo. should be able to understand and explain the effects on the new work procedures including any risk to the quality of the work, or the impact of the work on other aspects such as safety or the environment. In particular, the P.Geo. should consider areas where such innovation may be beneficial or could expose the project dataset to potential errors. If the new technology is to produce assay data to be used in project definition, then it is imperative that the samples are tested not only by the new technology but also by one or more industry-accepted methods to verify that the new technology is accurate and effective. This will assist the P.Geo. in understanding the error and the applicability of the technology to the work.

Implementation of any new technology should follow a Change Management Plan. When introducing new technology, the supervising P.Geo. should conduct:
1. An initial risk assessment to identify and address the changing conditions or new risks to the quality of the resulting geoscience work.
2. A risk assessment during commissioning to identify and address any new risks to the integrity of the geoscience work.
3. A final assessment following implementation to monitor and address any new hazards to the quality of the geoscience work.

Q & A – Session C

Any caveats or pitfalls to consider for AI implementation in mineral exploration?

In general, putting a geoscientist firmly in the loop is critical and will be able to mitigate many of the missteps I have seen during implementation. I think that, historically, subject matter expertise in disciplines has been downplayed as it was assumed that the difficult part of the problem (i.e. the creation of the algorithm) was already solved. This is not generally true, and particularly so in a discipline like the geosciences where data availability is comparatively quite sparse. You can't chuck a bunch of data at an algorithm and hope that it all works out as the parameter space is usually too large to be adequately constrained by the available data.

Great presentations! - the GSC for many years tried to do prospectivity maps to predict potential locations for undiscovered mineral deposits. One of the biggest problem they had was what volume/extent do you assign to a mineral deposit. Any comments on this?

That's a great question! I realize this answer is a vacillation, but I think a series of prospectivity maps at different scales would be useful. I've often seen them trying to jump to an outcrop/drill target scale, but I think something targeting a mineralized halo would be more useful (i.e. bigger is better). I think it's unlikely that many mid-large exploration companies would drill directly off of publicly available data and would instead do some internal mapping/geochemical vectoring/geophysics. I personally think that mineral prospectivity maps are more useful to remove 80-90% from a geoscientist's search area such that manual effort can be directed to the final 10-20%, versus targeting in on only a fraction of a percent of an area (i.e. directly letting the algorithm pick where to go).

.