

Dynamic Uncertainty Estimation of LIBS Chemistry on Mars with Gaussian Process Regression. T.S.J. Gabriel¹, R.B. Anderson¹, R. Wiens², and the Mars 2020 SuperCam instrument team. ¹U.S. Geological Survey Astrogeology Science Center, Flagstaff, AZ, USA (tgabriel@usgs.gov), ²Purdue University, West Lafayette, IN, USA.

Introduction: Obtaining accurate chemical compositions of Mars rocks along the *Perseverance* rover traverse is critical to Mars 2020 mission goals, including the assessment of samples planned for return to Earth [1]. SuperCam aboard *Perseverance* rover is a multi-modal instrument that can be used to produce estimates of elemental abundances of rocks using its laser-induced breakdown spectroscopy (LIBS) capability. After a short laser burst at a remote target, SuperCam's spectrometers collect light from the resulting plasma (emission spectra), which is indicative of the target's elemental composition (Fig. 1). This function is in addition to SuperCam's time-resolved Raman spectroscopy, visible and infrared reflectance spectroscopy (VIS/IR), remote microimager, and microphone capabilities [2,3].

For roughly a decade of ongoing chemometric activities across the SuperCam and ChemCam (a similar instrument on *Curiosity* rover) team members, machine learning (ML) algorithms have been explored and benchmarked in trade studies to improve elemental quantifications [e.g., 4, 5]. *Regressors* are a particular type of ML algorithm that can make predictions based on a learned functional response. They are useful in producing models that account for non-linear responses of LIBS spectra to the composition of target materials.

Traditionally, uncertainties on regressors have been treated in a static sense by the ChemCam and SuperCam team, *i.e.* the uncertainty of the estimation is predetermined and is only a function of the estimated chemical abundance. As part of SuperCam Trace and Minor Element Working Group activities [6], *we have benchmarked a unique regressor method, the Gaussian Process Regressor (GPR), which can generate chemical uncertainties that can respond in a dynamic manner to LIBS measurements. The method shows promise in producing similar and sometimes better performance to other regression algorithms, while providing more robust, per-prediction uncertainty quantification.*

Background: LIBS quantification is aimed at determining chemical quantities and their uncertainties from

photon emission spectra. Learning this functional response demands a robust set of existing LIBS spectra where the composition of the target is known (here referred to as *library* spectra). We leverage a large library of spectra collected at Los Alamos National Laboratory with a replica of the SuperCam instrument on Mars.

Developing a robust Machine Learning regressor to learn the functional response of a target's composition and the resulting spectra involves training, validation, and test steps. When a regressor is properly trained to LIBS spectra and compositions, elemental concentration(s) of a rock can be inferred from its emission spectrum. Training is performed by an optimization scheme on a subset of the library data. Models are then tested against data that was not included in the previous steps in order to assess model performance.

The performance of predictions of the test set are often used to characterize the uncertainties of model predictions. As shown in Fig. 2 (top), the prediction of a library target's elemental abundance can deviate from the true (known) abundance. To estimate this deviation, the root mean square error of the prediction (RMSEP) is computed locally (Fig. 2, bottom). This local-RMSEP method is a *static* approach, where uncertainty is held constant for a given abundance [5]. However, this method may over or underestimate uncertainty. For example, if a library set contains many samples with both high SiO₂ and Al₂O₃, but very few samples with both high SiO₂ and *low* Al₂O₃, then predictions of rocks with like the former should be more certain than those of the latter. However, the local-RMSEP method would assume an identical uncertainty for both. GPR, however, provides an estimate of uncertainty that is computed on a *per-prediction* basis, making it robust to these factors.

Methodology: We assess the performance of GPR using the open-source Python package scikit-learn [7]. Data handling, manipulation, and plotting were performed using the Pandas, NumPy, SciPy, and Matplotlib packages. Spectra were pre-processed according




Figure 1 – SuperCam LIBS spectrum (adapted from [5]) with coverage between wavelengths 244 – 341, 382 – 467, and 535 – 853 nm provided by three optical spectrometers [2,3]. Photon emission from the laser-generated plasma provides spectral peaks diagnostic of elemental concentration.

to instrument response characterization, calibration, and preparation procedures [5]. For each element, several hundreds of models were developed and the best model for each element was chosen based on validation performance. A multivariate approach (using multiple spectral features/wavelengths) was implemented instead of univariate (e.g. methods based on a single parameter, like the strength of an emission line) as we found generally better performance from multivariate methods.

Models were developed using different permutations of regressor parameters and preprocessing steps that were developed as part of SuperCam Trace and Minor Element Working Group activities [6]. These parameters included different GPR kernel designs, as well as three different elemental normalization methods, spectral masks, and selected features within the spectral window. Spectral masks are domains (wavelengths) of the spectra that are removed from the data, resulting in a smaller dataset. Domains are chosen based on the location of the dominant emission features for each element. For example, the full spectrum, the 375 – 465 nm, and the 620 – 710 nm range were explored for barium, which has strong lines at 454 and 650 nm, among others. Within each window, different numbers of features (specific wavelengths) of the spectra were provided to the models, using a feature selection ML algorithm, that itself is cross validated. In total, some elements reached nearly 1000 permutations of GPR models.

We also applied a ± 1 pixel shift in wavelength to the library spectra and independently tested the behavior on shifted spectra. This was done to characterize the effect of variations in wavelength calibrations and to ensure that the models were robust against this

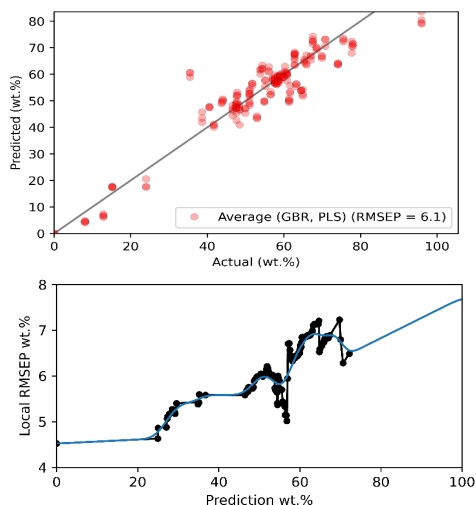


Figure 2 – (top) Actual (known) SiO_2 abundances of library targets and the predicted abundances from an ML regressor model that has a global RMSEP of 6.1 wt% in this case [adapted from 5]. The black (1:1) line defines where predictions match the actual (known) abundance. The local RMSEP of this model varies as a function of the predicted abundance.

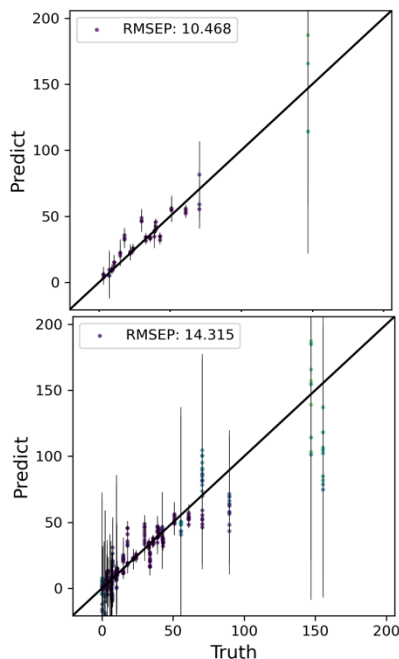


Figure 3 – (top) Actual (known) Li abundance (ppm) and the predicted Li abundance (ppm) from GPR for unshifted library data. (bottom) The same figure, with the addition of shifted and calibration targets.

effect. We also tested model performance on specially-designed calibration targets that are both in the library dataset and on *Perseverance* rover,

to understand the quality of calibration transfer. All calibration targets were in the test set for this reason and for certain elements we removed outlier samples with very high quantities from the library.

Results: *We find that GPR can provide accurate chemical predictions of library targets, while producing per-prediction uncertainties for targets where the prediction is expected to be less certain* (see Fig. 3). Higher uncertainties are generally associated with target types that are poorly populated in the library. For Li, high abundances are less populated in the library set, producing large error bars near ~ 150 ppm in Fig. 3. We also report larger uncertainties for calibration targets and targets at longer distances. We are currently exploring the source of this behavior and distance measurements are being considered in the development of an updated spectral library. Assessment of the correlations between GPR-reported uncertainties and library set, target factors, and other aspects is ongoing.

As part of SuperCam Trace and Minor Element Working Group activities, we found that GPR often gives comparable results to other regressors in terms of global RMSEP, but with the advantage of providing more information about the reliability of individual predictions. We are also expanding the spectral library and are exploring GPR for new major element models.

References: [1] Farley, *et al.* (2022), *Science*, 377(6614), [2] Wiens, *et al.* (2021), *SSR*, 217:4, [3] Maurice *et al.* (2021), *SSR*, 214:77, [4] Wiens, *et al.* (2013), *Spect. Acta. B*, [5] Anderson, *et al.* (2021), [6] Anderson, *et al.* 53rd LPSC, [7] Pedregosa, *et al.* (2011), *JMLR*.