RAMAN/LIBS DATA FUSION VIA TWO-WAY VARIATIONAL AUTOENCODERS Mario Parente[1], I. Gemp[2], [1]ECE Department, University of Massachusetts, 01003 Amherst, MA; mparente@ecs.umass.edu; [2] CICS, University of Massachusetts, 01003 Amherst, MA

Introduction: Several upcoming missions will feature Raman and LIBS spectroscopy for planetary exploration on Mars (e.g. Mars 2020 SuperCam and SHERLOC). Supporting the scientific effort will require a deeper understanding how to interpret Raman spectra of mineral assemblages in soils and rocks. Unfortunately, there is no underlying theoretical understanding of mixtures for Raman spectroscopy, and developing the necessary theoretical basis for understanding Raman mixing phenomena will take decades.

As an alternative, we are investigating a completely different, original solution to extracting mineral abundances from Raman data by combining Raman results with LIBS using deep learning and data fusion. Raman alone constrains what phase is present and the relative cation proportions in each phase. LIBS provides the total numbers of those cations, so the combination of the two allows the proportion of each phase (the modal abundance) to be obtained. This requires only knowledge of the mineral formulas of each phase identified by Raman, the associated algorithm for interpreting its individual composition from peak position and the LIBS total amounts of those cations.

***Two-way* Variational Autoencoder [1, 2, 3]:** We consider the marriage of two probabilistic models to describe the data. The first is a probabilistic model (M2 in [4]) that describes the spectra, $\mathbf{x}$, as being generated by a composition vector[1] $\mathbf{y}$ in addition to a latent, nuisance vector $\mathbf{z}$. The joint distribution is assumed to factorize as $p(\mathbf{x},\mathbf{y},\mathbf{z}) = p(\mathbf{y})p(\mathbf{z})p(\mathbf{x}|\mathbf{y},\mathbf{z})$, so the data are explained by the *generative process*:

$$p(\mathbf{y}) = Dir(\mathbf{1}) \tag{1}$$
$$p(\mathbf{z}) = U(-\mathbf{1.5},\mathbf{1.5}) \tag{2}$$
$$p_\theta(\mathbf{x}|\mathbf{y},\mathbf{z}) = f(\mathbf{x};\mathbf{y},\mathbf{z},\theta) \tag{3}$$

Here, $p(\mathbf{y})$ and $p(\mathbf{z})$ are prior distributions and $f(\mathbf{x};\mathbf{y},\mathbf{z},\theta)$ is a distribution whose parameters are non-linear functions of $\mathbf{y}$ and $\mathbf{z}$ (e.g., diagonal Gaussian $\mathcal{N}(\mu_\theta(\mathbf{y},\mathbf{z});\mathbf{\Sigma}_\theta(\mathbf{y},\mathbf{z}))$). We choose a uniform prior over the simplex for compositions, $Dir(\mathbf{1})$, and deep neural networks with weights θ for $\mu_\theta(\mathbf{y},\mathbf{z})$ and $\mathbf{\Sigma}_\theta(\mathbf{y},\mathbf{z})$. The second is a probabilistic model that describes the *reverse process*: nuisances and compositions are generated by spectra,

$$q(\mathbf{x}) = U(-\gamma,\gamma) \text{ e.g., } \gamma \gg 0 \tag{4}$$
$$q_\phi(\mathbf{y}|\mathbf{x}) = g(\mathbf{y};\mathbf{x},\phi) \tag{5}$$

[1]By composition we will indicate mineral abundances for Raman data or elemental compositions for LIBS data.

$$q_\phi(\mathbf{z}|\mathbf{x},\mathbf{y}) = h(\mathbf{z};\mathbf{x},\mathbf{y},\phi) \tag{6}$$

where $q(\mathbf{x})$ is an uninformative, uniform prior, $q_\phi(\mathbf{z}|\mathbf{x},\mathbf{y})$ is a diagonal Gaussian parametrized by a deep neural network as before, and $\gamma \gg 0$. To define $q_\phi(\mathbf{y}|\mathbf{x})$ with support limited to the simplex, we first draw an intermediate random variable, $\tilde{\mathbf{y}}$, from a diagonal Gaussian. We then pass $\tilde{\mathbf{y}}$ through a non-linear transformation called a *normalizing flow* [5] which conforms the Gaussian distribution to the simplex.

To learn the parameters, θ and ϕ, we optimize variational lower bounds on the marginal likelihoods of our data samples [1, 2, 3]. The marginal likelihood for the entire dataset is

$$\mathcal{J}_f = \sum_{(\mathbf{x},\mathbf{y})\sim\tilde{p}_l} \mathbb{L}_f(x,y) + \sum_{\mathbf{x}\sim\tilde{p}_{u_x}} \mathbb{L}_f(x) \tag{7}$$
$$\mathcal{J}_r = \sum_{(\mathbf{x},\mathbf{y})\sim\tilde{p}_l} \mathbb{L}_r(x,y) + \sum_{\mathbf{y}\sim\tilde{p}_{u_y}} \mathbb{L}_r(x) \tag{8}$$

for the forward and reverse models, respectively [1, 2, 3]. As in [4], we introduce an additional discriminative objective to each model that can be learned from the labeled data:

$$\mathcal{J}_f^d = \mathbb{E}_{(\mathbf{x},\mathbf{y})\sim\tilde{p}_l} L(\bar{\mathbf{y}},\mathbf{y}), \quad \mathcal{J}_r^d = \mathbb{E}_{(\mathbf{x},\mathbf{y})\sim\tilde{p}_l} L(\bar{\mathbf{x}},\mathbf{x})$$

where $\bar{\mathbf{y}}$ can, for example, either be a sample from $q_\phi(\mathbf{y}|\mathbf{x})$ or the mean of the distribution and L can, for example, be $\mathrm{KL}(\mathbf{y} \parallel \bar{\mathbf{y}})$; $\mathbf{x}$ is treated similarly using $p_\theta(\mathbf{x}|\mathbf{y},\mathbf{z})$ and $L = \parallel \bar{\mathbf{x}} - \mathbf{x} \parallel^2$. We jointly optimize both models as a weighted sum:

$$\mathcal{J} = \alpha_f \mathcal{J}_f - \alpha_f^d \mathcal{J}_\mathrm{f}^\mathrm{d} + \alpha_r \mathcal{J}_r - \alpha_r^d \mathcal{J}_\mathrm{r}^\mathrm{d}. \tag{9}$$

We learn the parameters θ and ϕ by maximizing (9) using Monte Carlo samples for the latent variables —a technique known as stochastic gradient variational Bayes [6] or stochastic backpropagation [7].

Data Fusion: The data fusion model incorporates both LIBS and Raman data simultaneously. The approach requires that LIBS and Raman spectra be collected from the same samples, thereby creating a dataset consisting of N trials, each measuring spectral intensities at $S = L + R$ channels ($L = \#$ of LIBS channels, $R = \#$ of Raman channels) with corresponding mineral abundances for M minerals. Raman spectra are captured for 567 binary mineral mixtures created in the lab. For our training dataset, $N = 189$ samples, $M = 23$ mineral types, $L = 5485$ LIBS channels, and $R = 1715$ Raman channels. We hold out 378 samples to validate our data fusion approach. We use the

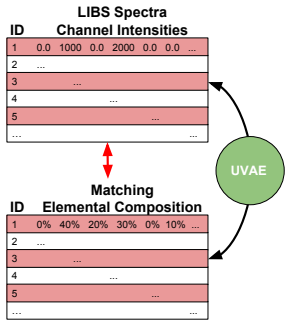

Figure 1: *We have* 130 *LIBS lab spectra obtained under Mars conditions with matching known elemental compositions. We also have* 500 *additional samples acquired in the same conditions with unknown elemental compositions. Each spectrum contains* 5485 *channels. By training the UVAE model to capture the relationship between these two representations, we acquire the ability to generate* synthetic *LIBS spectra. Specifically, training the UVAE results in learning the distribution over LIBS spectra given elemental composition,* $p(\mathbf{x}|\mathbf{y})$*. We can synthesize a new LIBS spectrum,* $\mathbf{x}$*, from this distribution conditioned on a specific elemental composition,* $\mathbf{y}$*.*

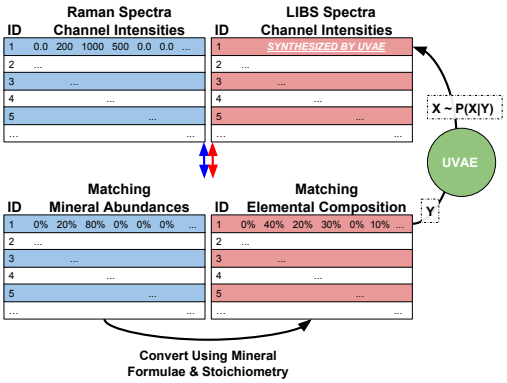

Figure 2: *We synthesize LIBS spectra for our mineral abundance prediction challenge by 1) converting mineral abundances to elemental composition using stoichiometry, 2) training the UVAE to model the relationship between LIBS spectra and elemental composition, and then 3) generating LIBS spectra conditioned on these elemental compositions using the UVAE model trained in step 2.*

deep generative model described above, namely the Untapped Variational Autoencoder (UVAE), to predict the percent mineral abundances for 23 minerals given the S-dimensional LIBS+Raman spectral data.

Synthesizing an Appropriate Dataset: The *data fusion* scenario we consider assumes both Raman spectra and LIBS spectra have been acquired in conjunction with matching mineral abundance for each rock sample. In practical in situ analyses one might only have access to LIBS spectra and corresponding percent elemental compositions for some samples and Raman spectra corresponding mineral abundances on other samples. In order to simulate this scenario, we devise a process for synthesizing matching LIBS spectra. First, we train a UVAE to model the relationship between LIBS spectra and elemental composition (see Figure 1). Next, we use stoichiometry to convert atomic weights for each mineral to percent elemental composition. For each measurement trial in our Raman dataset, we calculate the percent elemental composition of a mineral mixture by combining the elemental compositions of each mineral in corresponding proportions. Then, using the trained UVAE model, we generate LIBS spectra conditioned on the derived elemental compositions from the Raman dataset. This results in a LIBS+Raman spectral dataset where the LIBS portion of the spectrum is synthetic, however, deliberately crafted to mimic the real matching LIBS spectra although with sometimes significant loss in fidelity. The process is shown in Figure 2.

Predicting Mineral Abundances: We compare three models: 1) a UVAE trained on Raman+LIBS spectra, 2) a UVAE trained on Raman only, and 3) PLS trained on Raman only. In order to make predictions of mineral abundances using the UVAE model, we utilize $q(\mathbf{y}|\mathbf{x})$. We can either sample from this distribution or take the mean of this distribution; in our results we use the mean. Preliminary results show that including the LIBS spectral data (even though synthetic!) is helpful in driving down the error of predicting mineral abundances. Here,

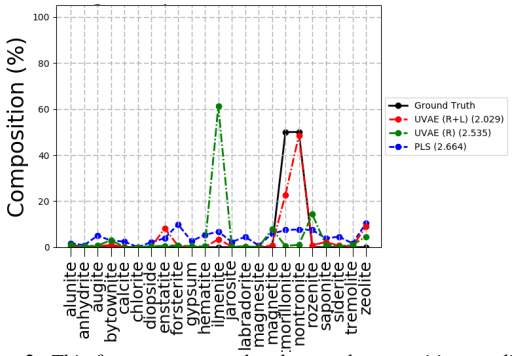

Figure 3: *This figure compares the elemental composition predicted by each model for a single rock sample. The average KL divergence for each model is given beside the model name in the legend. Mineral names are on the* $\mathbf{x}$*-axis while % abundance is on the* $\mathbf{y}$*-axis.*

we measure error as the average KL divergence between the true mineral abundance of a rock sample and the predicted mineral abundance given by the model (lower is better) on the held out validation set. Figure 3 shows the percent mineral abundance for a selected sample as predicted by UVAE (red), PLS (blue), and actual (black). The average KL divergence over all samples in the validation set is quantified for each model in the legend.

References: [1] I. Gemp, et al. (2016) *ArXiv e-prints.* `arXiv:1608.05983`. [2] I. Gemp, et al. (2017) in *NIPS AABI,* Long Beach, CA, USA. [3] M. Parente, et al. (2017) in *IEEE IGARSS* Forth Worth, TX, USA. [4] D. Kingma, et al. (2014) in *NIPS* 3581–3589. [5] D. J. Rezende, et al. (2015) *arXiv preprint arXiv:150505770.* [6] D. Kingma, et al. (2013) *arXiv preprint arXiv:13126114.* [7] D. J. Rezende, et al. (2014) *arXiv preprint arXiv:14014082.*