mineralML: Leveraging Machine Learning for Probabilistic Mineral Classification

Sarah C. Shi[1,2], Penny E. Wieser[1], Charlotte Gordon[1], Norbert Toth[3], Paula Antoshechkina[4], Matthew L.M. Gleeson[1], Kerstin Lehnert[2],

[1]University of California, Berkeley. Berkeley, CA USA
[2]Lamont-Doherty Earth Observatory, Columbia University. New York, NY USA
[3]University of Cambridge. Cambridge, Cambridgeshire UK
[4]California Institute of Technology. Pasadena, CA USA
[*]Corresponding Author (sarahshi@berkeley.edu)

The following document is a non-peer reviewed preprint submitted to EarthArXiv. The paper has been submitted to *Geochemistry, Geophysics, and Geosystems* for peer review. Please note that the manuscript has not undergone peer review and that subsequent versions of this paper may have slightly different content. If accepted, the final version will be available through a DOI link on this webpage.

The Python package can be accessed on GitHub at github.com/sarahshi/mineralML, and installed with PyPi. The documentation on ReadtheDocs details multiple usages of the mineralML package at mineralml.readthedocs.io/en/latest, working with several different data formats:

· Tabular Data
mineralml.readthedocs.io/en/latest/examples/mineralML_neuralnetwork.html
· Microanalytical Tabular Data, straight from an CAMECA, Probe for EPMA, or AZtec instrument
mineralml.readthedocs.io/en/latest/examples/mineralML_microanalysisoutput.html
· Mapped EDS Data
mineralml.readthedocs.io/en/latest/examples/mineralML_mapping.html

An implementation working with tabular data can be found on Google Colab at:
colab.research.google.com/github/SarahShi/mineralML/blob/main/mineralML_colab.ipynb

If you encounter any issues, please reach out or raise an issue on GitHub. We welcome constructive feedback.

mineralML: Leveraging Machine Learning for Probabilistic Mineral Classification

Sarah C. Shi [1,2], Penny Wieser [1], Charlotte Gordon [1], Norbert Toth [3], Paula Antoshechkina [4], Matthew L.M. Gleeson [1], Kerstin Lehnert [2]

[1]Department of Earth and Planetary Sciences, University of California, Berkeley. Berkeley, CA USA
[2]Lamont-Doherty Earth Observatory, Columbia University. New York, NY USA
[3]University of Cambridge. Cambridge, Cambridgeshire UK
[4]California Institute of Technology. Pasadena, CA USA

Key Points:

- Neural networks trained on high-quality datasets can be used to accurately and probabilistically classify mineralogy in igneous rocks
- `mineralML` streamlines e-beam analysis workflows with automated mineral classification, stoichiometric calculations, classification diagrams
- `mineralML` helps identify errors in large geochemical databases (e.g., GEOROC, PetDB), improving the reliability of big-data studies

Corresponding author: Sarah C. Shi, `sarahshi@berkeley.edu`

Abstract

Characterizing phase assemblages in igneous rocks and the chemical variability within these phases is the fundamental basis of many petrological investigations. We present `mineralML` (mineral classification using Machine Learning), an open-source `Python` package that classifies common igneous minerals based on oxide chemical data, with prediction scores. `mineralML` employs a two-stage neural network: a variational Bayesian classifier providing probabilistic mineral classifications with uncertainty estimates, along with an autoencoder that projects compositions into a 2-D latent space for visualization. Trained on ∼128,000 curated electron probe microanalyzer analyses spanning 23 mineral classes, `mineralML` achieves $> 99\%$ accuracy on validation data. Applied to > 1.1 million analyses from the GEOROC database, `mineralML` achieves $> 95\%$ classification accuracy, with many of the $\sim 5\%$ misclassifications demonstrating the utility of this package for identifying database misclassifications and data entry errors. When applied to co-collected EBSD-EDS maps, `mineralML` determines phase proportions within 2–3% of the EBSD-derived values, with added advantages over EBSD in identifying non-crystalline and difficult-to-index phases.

Following mineral classification, `mineralML` functions streamline common data processing workflows such as mineral stoichiometry and crystallographic site calculations, assignment of mineral subtypes (e.g., andesine versus labradorite), and production of mineral classification diagrams (e.g., feldspar, pyroxene, amphibole and Fe-Ti oxides ternary and quadrilateral diagrams). `mineralML` functions can also rapidly process EDS maps and generate publication-ready figures of mineral phases, their abundances, and their zoning patterns. As an open-source tool with high classification accuracy, `mineralML` presents immense potential for big data approaches when working with large geochemical databases and EDS maps.

Plain Language Summary

Identifying the minerals present in igneous rocks and understanding their chemical variability is central to understanding how magmas form and evolve prior to volcanic eruptions. Minerals are identified by probing polished rock samples with an electron beam and measuring the characteristic X-rays released, revealing their chemical composition. Classifying minerals from these chemical data, especially across large datasets or detailed chemical maps, is time-consuming and often performed manually, rendering these classifications difficult to reproduce. We have developed `mineralML`, an open-source `Python` tool using machine learning to automate mineral classification from chemical compositions. `mineralML` assigns mineral names, provides inbuilt prediction scores, and also streamlines downstream tasks such as calculating stoichiometries and crystallographic sites, performing sub-classifications, and converting chemical maps into publication-ready phase and zoning maps. When tested against over 1.1 million analyses from the GEOROC database along with co-collected EBSD-EDS maps, `mineralML` achieves greater than 96% accuracy and reproduces phase proportions within 2-3%. By automating and standardizing mineral classification, `mineralML` opens new possibilities for big data approaches to understanding igneous systems and curating large geochemical databases.

1 Introduction

Crystals in igneous rocks record the temporal, compositional, and thermal evolution of magmatic systems. Different mineral phases, and even distinct crystals of a single phase, encode various aspects of this history, including growth from compositionally heterogeneous melts, mixing of magma batches, and changing magma chamber conditions (Putirka, 2008; Blundy & Cashman, 2008; Cashman & Blundy, 2013; Neave et al., 2013; Mutch et al., 2019; Cheng et al., 2020; Wieser et al., 2025; Ubide et al., 2014). Petrologic studies therefore begin by identifying the mineral phases present, determining their

abundances, and documenting textures and microstructures. The presence and absence of different mineral phases within a given rock dictates which quantitative tools can be used, and thus what scientific questions may be addressed. For example, zircon allows for U-Pb dating, while K-bearing phases such as feldspar and glass allow for K-Ar and Ar-Ar dating (Schoene et al., 2013). Certain minerals exhibit variations in chemistry as a function of the temperature, pressure, and melt composition from which they grow, allowing for the application of thermobarometric and hygrometric methods to estimate these parameters within the magmatic system (e.g., amphibole (Ridolfi, 2021), pyroxene (Putirka, 2008), and plagioclase (Waters & Lange, 2015)). Furthermore, mineral assemblages, including modal phase abundances, are themselves key observables that provide important information about the rock. The presence or absence of specific phases defines equilibrium assemblage fields in pressure–temperature-composition (P-T-X) space, underpinning the construction of a P-T-(time) path revealing the history of the rock (Blatter et al., 2013; Weller et al., 2013; Spear et al., 2017; Lanari & Engi, 2017; Voigt et al., 2022).

The push towards automated, data-driven igneous petrology has driven the rapid growth of Findable, Accessible, Interoperable, and Reusable (FAIR)-aligned data infrastructures that aggregate results from individual petrological studies (Stall et al., 2019; Klöcking et al., 2023). Databases such as EarthChem and GEOROC (Lehnert et al., 2000), and LEPR/traceDs (Hirschmann et al., 2008) enable synthesis across studies and have supported analyses of global petrologic trends (Reubi & Blundy, 2009; C. B. Keller & Schoene, 2012; Pitcher & Kent, 2019; Barber et al., 2021; Stracke et al., 2022; Rasmussen et al., 2022) and mineral-melt equilibria (Mutch et al., 2022). However, these compilations inherit heterogeneous analytical practices and labeling conventions, and contain a number of errors, which are challenging issues to resolve given the size of the databases—GEOROC alone contains more than 1 million individual mineral analyses. Analyses may have been incorrectly assigned a mineral group (e.g., pyroxene labeled as an amphibole), columns may have been mismatched during data compilation (e.g., FeO_t values entered into the MnO column), and entire columns of data may not have been compiled during curation. There are also issues with analysis quality: because of the relatively large interaction volumes associated with electron beam techniques, a reported analysis may be a mixture of two neighboring or intergrown phases, termed a 'mixel' or mixed pixel in an EDS map. Literature analyses may also suffer from low totals, low cation sums, or poor correspondence to theoretical mineral stoichiometry. There is currently no community consensus on the best practices to filter out low-quality data in these large datasets, or how to derive mineralogy from oxide data. These limitations directly impact studies utilizing these datasets. Machine learning models using these large datasets to make petrological inferences rely on data reliability, and remain inherently vulnerable to 'garbage in, garbage out'.

Another major advance in petrological research has been the increase in the speed of Energy Dispersive Spectroscopy (EDS) detectors, and improvements in software allowing for fully quantitative analyses to be extracted with comparable quality to Wavelength Dispersive Spectroscopy (WDS) measurements on electron probe microanalyzers (Allaz et al., 2026). These advances make it feasible to collect a quantitative chemical map of an entire thin section or epoxy stub, comprised of around a million individual analyses (pixels), in one overnight session. There has also been an increase in the availability of SEM systems that allow EDS maps to be collected at the same time as electron backscatter diffraction (EBSD) crystallographic maps. High-throughput, spatially resolved chemical datasets at the thin-section scale are therefore becoming increasingly common in the petrological literature (Gottlieb et al., 2000; Pirrie et al., 2004; Neave et al., 2014, 2017; Neave & Maclennan, 2020; Gleeson et al., 2025; Toth et al., 2025). These hyperspectral SEM-EDS maps are powerful as they enable the concurrent interpretation of chemistry and microstructure at the grain scale, with minerals in their original textural context.

Extracting mineralogy from EDS maps can replace tedious tasks such as point counting to determine bulk mineralogy. Existing closed-source commercial softwares, including QEMSCAN (Pirrie et al., 2004), ZEISS Mineralogic (Graham et al., 2015; Taylor et al., 2024), and AZtecMineral (Jones et al., 2020), are designed to provide mineral classifications of each pixel using EDS data, but typically rely on *a priori* information on the mineral assemblage present for peak intensity matching along with iterative user input and adjustments to tune compositional windows. Each adjustment can require reclassification of large maps, making this workflow time-consuming and difficult to reproduce. More recently, open-source tools have been developed to perform unsupervised classification of peak intensities in EDS maps, through dimensionality reduction and clustering. These unsupervised methods group pixels by compositional similarity in their EDS spectra without requiring predefined mineral labels or compositional windows. For example, SIGMA (Tung et al., 2022) and GPyEDS (Toth et al., 2025) use neural network autoencoders to learn efficient, low-dimensionality encodings of EDS peak intensities and then cluster pixels in this reduced space. These autoencoder models can be re-used for the examination of similar datasets of samples with similar phases and compositions, but necessitate re-training for the examination of more variable datasets of samples with different phases and compositions. As these methods are unsupervised, compositional clusters are returned rather than mineral identifications, so the user must then manually assign a mineral label to each cluster after the fact. Additionally, XMapTools (Lanari & Tedeschi, 2025) has expanded to provide common clustering tools for working with mapped intensity data, but also requires *a priori* information with representative pixel X-ray intensities from each phase, or manual thresholding to initially identify clusters.

We present the first open-source tool, `mineralML` (**mineral** classification using **M**achine **L**earning), for classifying mineralogy from quantitative elemental or oxide concentration data, requiring no prior knowledge of the mineralogy or chemical composition of different phases in the sample. `mineralML` is built in `Python` and trained on common igneous minerals. `mineralML` allocates a mineral name (e.g., pyroxene, olivine, feldspar) based on elemental or oxide concentration data, and can be applied to spot analyses or map pixels acquired with a range of microanalytical techniques, including EPMA, WDS, and EDS. `mineralML` also calculates stoichiometries, crystallographic sites, and common mineral components (e.g. anorthite and/or jadeite content), allowing quality filters to be easily applied (e.g, cation sums), and helps streamline common igneous workflows using these parameters (e.g., thermobarometry, plotting ternary diagrams). `mineralML` also uses these additional calculations to provide further sub-classifications where relevant, such as classifying an amphibole as a tremolite, a feldspar as labradorite, a pyroxene as aegirine, an oxide as a rhombohedral oxide, among others. Thus, we envision petrologists incorporating `mineralML` as a key part of their workflow after extracting data from an electron beam instrument, to confirm the mineralogy, filter for poor data, perform component calculations, and make common petrological classification diagrams (Figure 1). These capabilities extend to entire EDS maps, shifting the bottleneck from manual phase segmentation to data acquisition, allowing users to focus on optimizing the quality of the input data, knowing that they can process the resulting data in minutes.

`mineralML` is also vital for the identification of misclassified mineral phases and low-quality data in compiled datasets prior to their use in 'big data' petrological studies. We evaluate the `mineralML` workflow on independent test datasets from GEOROC, demonstrating that `mineralML` holds the potential to rapidly classify millions of analyses within minutes, identifying their mineral group and sub-classifications, while flagging incorrectly classified data. We anticipate that this approach will be integrated into existing geochemical data processing pipelines, to improve the quality of large geochemical datasets used to train petrological models.

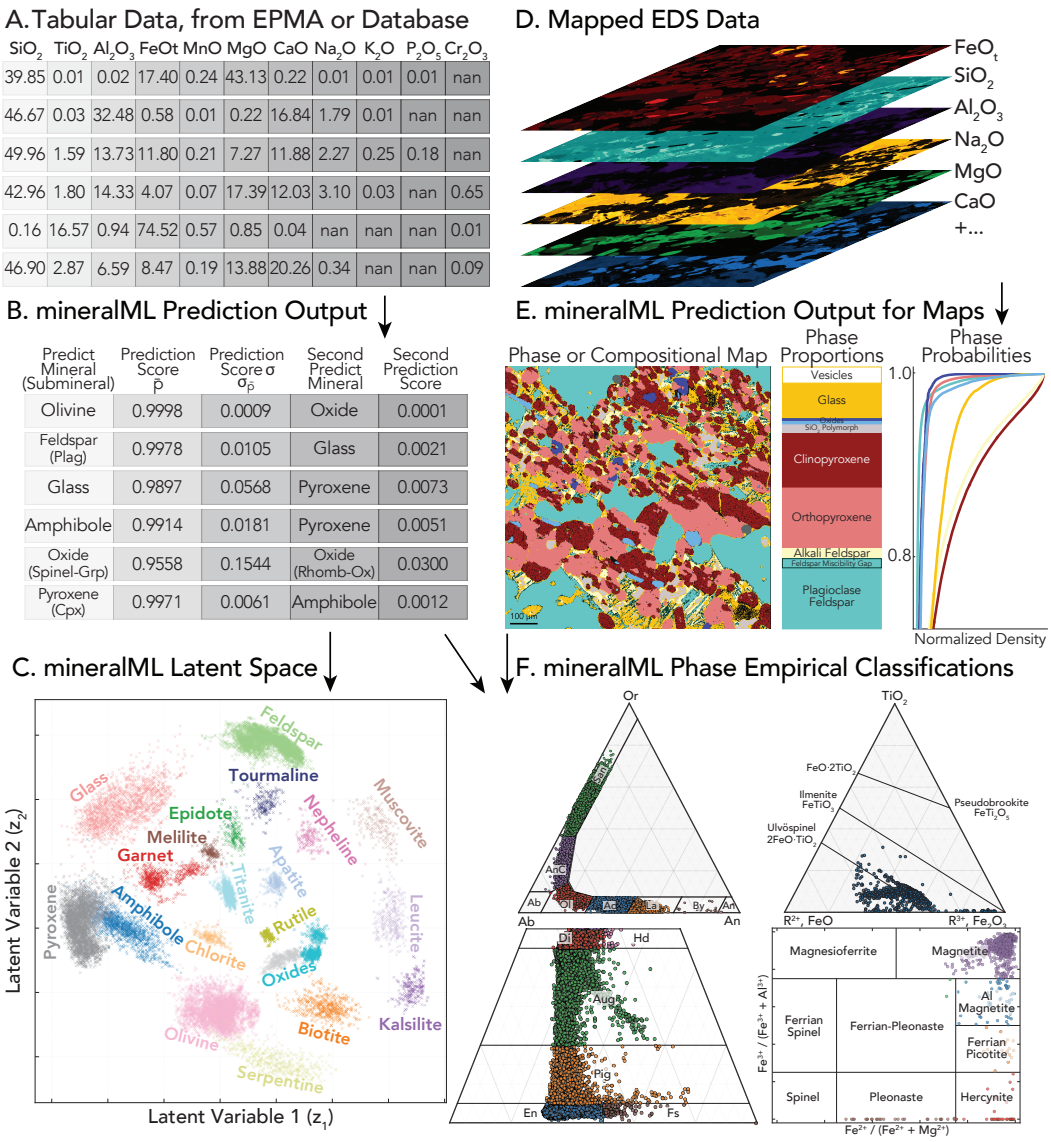

Figure 1. Overview of the `mineralML` workflow presented highlighting applications to both: A) Tabular data, straight from EPMA/EDS spot analyses or a geochemical database and D) Mapped EDS data. Functions are modular and enable probabilistic mineral predictions for either form of data in one code function call. B) `mineralML.predict_class_prob_nnwr` returns a dataframe containing the predicted mineral, its prediction score and associated uncertainty, and the second most likely mineral classification and score. C) The compositions can be projected into the two-dimensional latent space of the training and visualization dataset, to assess whether the predicted mineral is likely and if the data are well represented by the training dataset with `mineralML.plot_z2_overlay`. E) `mineralML.run_map` returns publication-ready figures of phase or compositional maps, phase proportions, and a cumulative density function showing phase probabilities. F) Tabular compositions can then be plotted in petrological classification diagrams including the feldspar, pyroxene, amphibole and Fe-Ti oxides ternary and quadrilateral diagrams.

2 Methods

`mineralML` is a `Python` package containing a neural network that jointly performs probabilistic mineral classification and compositional reconstruction from chemical anal-

yses. The neural network uses a joint probabilistic classifier and autoencoder, allowing for robust generalization and easy visualization of high dimensional data; the training overview workflow is shown in Figure 2 and the application workflow is shown in Figure 1. After initial empirical filters for compositionally simple phases (carbonates, SiO_2-polymorphs, zircon), the network operates on concentrations for 11 oxides (SiO_2, TiO_2, Al_2O_3, FeO_t, MnO, MgO, CaO, Na_2O, K_2O, P_2O_5, Cr_2O_3). The model outputs the predicted, most-likely mineral class with a prediction score, with values closer to 1 more likely belonging to that predicted mineral class. These prediction scores are the output of this probabilistic model, providing a metric for assessing the likelihood of the model providing an accurate classification. `mineralML` classifies minerals based on the following labels provided in Table 1: amphibole, apatite, biotite, calcite, chlorite, epidote, feldspar, feldspathoids, garnet, melilite, mica, olivine, oxides, pyroxene, quartz (inclusive of all SiO_2 polymorphs), rutile, serpentine, titanite, tourmaline, and zircon. The model also provides a 2-D latent representation which captures the discriminative geochemical features, enabling visualization of the geochemical differences between and within mineral groups, through a compositional reconstruction (trained representation in 1C). These latent representations serve as a lower-dimensionality representation of the input data in which samples with greater chemical similarity are positioned closer together. Unlike linear dimensionality reduction methods such as principal component analysis (PCA), the neural network learns non-linear mappings from the original oxide space to capture more complex compositional relationships between and within mineral groups (Supporting Information S4). By requiring the network to reconstruct input oxide compositions, the network learns to encode geochemically meaningful information into the latent space. This latent space is also vital for assessing the degree of overlap between user-supplied data and the range of compositions represented by the training dataset.

The following sections describe how each component of `mineralML` is constructed, from data compilation to model training and application. Section 2.1 presents the training, validation, and testing datasets compiled for model development and evaluation. Section 2.2 describes the stoichiometric filtering framework used to curate and filter the compiled training and validation data. This stoichiometric framework is also re-used for later examination of mineral compositions. Section 2.3 presents the synthetic mineral generator developed to address class imbalances between mineral groups. Section 2.4 outlines the imputation and pre-processing steps applied to oxide data prior to model training. Section 2.5 details the neural network architecture and training procedure. Finally, Section 2.6 demonstrates the application of the pre-trained model to new data for mineral classification and latent space visualization.

2.1 Training, Validation and Testing Data

We compiled two distinct datasets: one which is split between training and validation, and one for testing. The training dataset was used to learn the core patterns and relationships between oxide compositions and their mineral labels, forming the foundation of the model. The validation dataset was used to tune hyperparameters, which are the parameters configuring a machine learning model, and monitor for overfitting by providing checkpoints and continuous feedback on the model during training. The testing dataset was held as the final auditor of the machine learning model, to evaluate the performance of the model on entirely unseen data, as an unbiased assessment of model accuracy and generalization capabilities. We ensured there was no overlap between the combined training and validation dataset with the test dataset of GEOROC (Lehnert et al., 2000) to ensure robust model evaluation and model extensibility without overfitting.

The training and validation dataset, totaling ∼128,000 EPMA analyses of 23 classes (mineral and melt), was manually compiled from literature and will be published on EarthChem as an expert dataset upon article acceptance (Table 1). Training a model to perform mineral classification based on oxide data alone is non-trivial, as real datasets may

contain published phase labels that are incorrect or compositions with swapped entries (see Section 3.2 for examples). The data inputs are concentrations for 11 oxides in oxide %. Point analyses with at least three measured oxides were retained for further use in training the model. Analyses were carefully culled and filtered with `mineralML` stoichiometric calculators to determine stoichiometric site occupancies (see Section 2.2). Analyses falling outside 10% of expected site occupancies were removed, to ensure the generation of a dataset with accurate mineral phase labels.

To achieve class balance in the training dataset, 1,000 analyses of most minerals were sampled. The 1,000 analyses for each label is a relatively small dataset, due to the differential availability of data, as analyses from labels such as tourmaline and melilite are less abundant than those from olivine or feldspar. For minerals showing highly variable chemistry (e.g., amphibole, glass), or continuous coupled substitutions (e.g., Na–Ca and Si–Al in plagioclase and alkali feldspar), we sampled 2,000 analyses.

The distinctions between related minerals—for example, clinopyroxene, pigeonite, and orthopyroxene—may be defined by differences in crystallographic criteria, compositional criteria, or both. This poses a challenge for `mineralML`, which classifies minerals based only on chemical criteria and does not have access to crystallographic information. In some cases, the compositional and crystallographic boundaries align. Mineral species such as monoclinic clinopyroxene and orthorhombic orthopyroxene are defined by their different crystallographic structures, but this difference correlates strongly with composition. The minerals can therefore be distinguished empirically based on chemistry. SiO_2-polymorphs differ in their crystallographic structure but are chemically identical, so cannot be distinguished by `mineralML`. Other mineral species subdivisions are empirically defined by chemical criteria (e.g., anorthite content boundaries to distinguish types of plagioclase). `mineralML` therefore has access to all of the required information to make these subdivisions, but it is challenging for the model alone to learn these subdivisions given the continuous nature of the substitution, and somewhat arbitrary boundaries.

We grouped some minerals together: feldspar (merging plagioclase and alkali feldspar), pyroxene (clinopyroxene, pigeonite and orthopyroxene), oxides (rhombohedral oxides including hematite and ilmenite, and spinel group minerals including magnetite and ulvöspinel). Consolidating these minerals ensures that the model identifies the underlying geochemical relationships within and between minerals. To ensure representation across these minerals with a broader range of compositions, we used a two-step sampling procedure. First, analyses for each mineral were clustered in composition space using k-means clustering (Lloyd, 1982; Pedregosa et al., 2011), producing k clusters per mineral. Second, we performed stratified random sampling across clusters by selecting an equal number of analyses from each cluster. The sampling procedure ensures that the training dataset spans the observed compositional space within each mineral, rather than being dominated by the densest compositional modes. The glass phase was sampled with stratification within total-alkali silica space (Le Bas et al., 1992); TAS field classifications were assigned with `pyrolite` (Williams et al., 2020). In practice, we were unable to compile 1,000 natural analyses of certain mineral phases showing chemical diversity (e.g., zircon, kalsilite-nepheline, SiO_2-polymorphs, hematite). To ensure class balance in the training dataset, we supplemented these minerals with compositions generated using a synthetic mineral generator (described further in Section 2.3 and Supporting Information S5).

We evaluated the trained model with data from diverse tectonic environments and collected in a wide variety of laboratories by compiling a test dataset, aggregating data from three distinct sources: 1) the GEOROC mineral dataset ($n = 1,112,437$), 2) compilations of secondary standard analyses ($n = 3,692$), and 3) co-collected EDS and Electron Backscatter Diffraction (EBSD) maps acquired at the University of California, Berkeley ($n = 2$ maps each with ∼200,000 pixels). Unlike the curated approach taken with the training and validation dataset, analyses from the test dataset underwent no filter-

ing. While the training dataset emphasizes high-quality, curated analyses that best represent each mineral group, the test dataset was left unfiltered to form an unbiased evaluation of the model's performance and generalization ability.

2.2 Stoichiometric Filters

We developed an object-oriented `Python` framework using `numpy` and `pandas`, built around a `mineralML.BaseMineralCalculator` class, providing a generalized workflow for converting raw oxide weight percentages (mass fractions) into structural mineral formulas. For any given analysis, the `mineralML.BaseMineralCalculator` class performs a sequence of conversions. First, oxide weight percentages are converted into molar compositions, which are used to determine the cation and oxygen proportions. These cation and oxygen proportions are then normalized to the structural atoms per formula unit (apfu) based on the theoretical oxygen basis specific to each mineral. Mineral-specific logic and structural site assignments are implemented through dedicated subclasses, each titled `mineralML.[Subclass]Calculator`. Structural site assignments follow the calculations described in Deer et al. (2013), where possible. Where Fe can exist in both ferric and ferrous form within the mineral structure (e.g., garnet, Fe-Ti oxides), the Fe^{2+} and Fe^{3+} estimations are calculated by the Droop (1987) method. To handle complex solid solutions and site occupancies, `mineralML` implements several additional classification pipelines, described below.

Amphiboles, which exhibit highly complex chemistry and varying site occupancies, are processed using the `mineralML.AmphiboleClassifier`. Structural formulas are calculated on a 23-oxygen and 13-cation basis, with normalizations from both Leake et al. (1997) and Ridolfi (2021) formations for site assignments, with code adapted from `Thermobar` (Wieser et al., 2022). For example, calcic amphibole compositions can be classified by plotting them on a Mg# vs. Si (apfu) diagram from Leake et al. (1997) to return the amphibole sub-classifications of tremolite, actinolite, magnesiohornblende, ferroactinolite, ferrohornblende, and ferrotschermakite.

Feldspar compositions spanning both the plagioclase and alkali feldspar solid solutions are processed using the `mineralML.FeldsparClassifier`. Structural formulas are normalized to an 8-oxygen basis to calculate the molar fractions of the three primary endmembers: anorthite (An), albite (Ab), and orthoclase (Or). These calculated endmember components allow for direct sub-classification and plotting on the standard high-temperature feldspar ternary diagram (Deer et al., 2013). This visualization and classification logic is implemented utilizing `python-ternary` (Harper et al., 2015), building upon compositional boundaries defined by Deer et al. (1963) and Deer et al. (2001) and implemented with code from Wieser et al. (2022). The feldspar ternary returns classifications of anorthite, bytownite, labradorite, andesine, oligoclase, albite, anorthoclase, and sanidine.

Pyroxene compositions (including clinopyroxenes, orthopyroxenes, and sodic pyroxenes) are evaluated using the `mineralML.PyroxeneClassifier`, which calculates structural formulas on a 6-oxygen and 4-cation basis. The classification logic follows the International Mineralogical Association (IMA) nomenclature framework established by Morimoto (1988). The model first discriminates between sodic and non-sodic pyroxenes based on their Q ($Ca + Mg + Fe^{2+}$) and J ($2 \cdot Na$) structural components. Non-sodic pyroxenes are further evaluated on the enstatite-ferrosilite-wollastonite (En-Fs-Wo) quadrilateral. Sodic pyroxenes are similarly evaluated on the wollastonite-jadeite-aegirine (Wo-Jd-Aeg) ternary, utilizing portions of endmember allocation calculations adapted from Walters (2022). `mineralML` automatically calculates and plots both the En-Fs-Wo quadrilateral and the Wo-Jd-Aeg ternary to return sub-classifications. The En-Fs-Wo composition pyroxenes are classified as diopside, hedenbergite, augite, pigeonite, enstatite, and ferrosilite; the Wo-Jd-Aeg composition pyroxenes are classified as omphacite, aegirine-augite, jadeite,

Table 1. Overview of the compiled training and validation dataset. Abbreviated mineral names from the International Mineralogical Association (Whitney & Evans, 2010). PPL and XPL photomicrographs of different minerals from the thin section showcase of Alex Strekeisen (www.alexstrekeisen.it/english/), who kindly provided permission for their use with permissions.

Mineral	n_{total}	$n_{training}$	$n_{validation}$	$n_{synthetic}$	**Thin section**
Amphibole (Amp)	6626	2000	4626	-	
Apatite (Ap)	1788	1000	788	-	
Biotite (Bt)	2933	1000	1933	-	
Carbonate (Cb)	1010	1000	10		
Chlorite (Chl)	1250	1000	250	-	
Epidote (Ep)	1055	1000	55	-	
Feldspar (Fsp)					
Alkali Feldspar (Alk)	6017	1000	5017	-	
Plagioclase Feldspar (Plag)	19 243	1000	18 243	-	
Garnet (Grt)	3459	1000	2459	-	
Glass (Gl)	9366	2000	7366	-	
Kalsilite (Kls)	1500	1000	500	1316	
Leucite (Lct)	1528	1000	528	-	
Melilite (Mll)	1136	1000	136	-	
Muscovite (Ms)	1100	1000	100	-	
Nepheline (Nph)	1020	1000	20	-	
Olivine (Ol)	25 069	1000	22 069	-	
Oxide (Ox)				-	
Rhombohedral Oxide					
Hematite (Hem)	1500	1000	500	1195	
Ilmenite (Ilm)	1957	1000	957	-	
Spinel Group				-	
Magnetite (Mt)	3323	1000	2323	-	
Spinel (Spl)	1717	1000	717	-	
Pyroxene (Px)				-	
Clinopyroxene (Cpx)	12 035	1000	11 035	-	
Orthopyroxene (Opx)	5826	1000	4826	-	
Rutile (Rt)	1718	1000	718	-	
Serpentine (Srp)	2588	1000	1588	-	
SiO_2-polymorphs (Qz)	1500	1000	500	1436	
Titanite (Ttn)	8051	1000	7051	-	
Tourmaline (Tur)	1188	1000	188	-	
Zircon (Zrn)	1500	1000	500	950	

and aegirine. Multiple Fe^{3+} estimation routines are additionally integrated into the pipeline, including charge balance methods from Droop (1987) and Lindsley (1983), along with the formulations of X. Wang et al. (2021).

When training the model, oxide minerals in `mineralML` are subdivided into rhombohedral oxides (Fe-Mg-Mn-Ti oxides, including minerals along the hematite-ilmenite solid solution join) and spinel-group minerals (Fe-Mg-Ti-Al-Cr spinels, including minerals along the ulvöspinel-magnetite solid solution join), following the definition of Ghiorso and Sack (1991). Owing to difficulty distinguishing relatively pure hematite ('rhombohedral oxide') and magnetite ('spinel') in data of variable quality when O is not directly measured but calculated by stoichiometry, we chose to allocate the label of 'oxide' regardless of whether the model allocates a label of 'rhombohedral oxide' or 'spinel group' (see further discussion in 6). We then return the 'rhombohedral oxide' or 'spinel group' labels as the 'Submineral'. For subsequent stoichiometric calculations, we encourage the user to carefully inspect whether these labels are reliable based on their data quality. The separation into rhombohedral oxides and spinel-group minerals is essential prior to further calculations as different assumptions of oxygen- and cation- bases are required for the two groups.

Rhombohedral oxides are processed with the `mineralML.RhombohedralOxideCalculator` on a 3-oxygen and 2-cation basis. The function calculates the site distributions and endmember fractions of hematite, ilmenite, Mn-ilmenite, and geikielite. The rhombohedral oxide calculator returns sub-classifications of ilmenite and hematite, split in the middle along the tie line. Spinel-group minerals—including magnetite, chromite, ulvöspinels—are processed with the `mineralML.SpinelCalculator` on a 4-oxygen and 3-cation basis. For both groups, the calculation of structural Fe^{3+} and Fe^{2+} relies on the stoichiometric charge-balance method of Droop (1987). Hematite-rich samples (defined empirically as $FeO_t >$ 60 wt.%) bypass this correction and are treated as purely ferric. Spinels can be further sub-classified along standard compositional axes (e.g., $Fe^{2+}/(Fe^{2+}+Mg)$ versus $Fe^{3+}/(Fe^{3+}+Al)$) to categorize intermediate solid solutions and return sub-classifications such as magnesioferrite, magnetite, ferrian-spinel, ferrian-pleonaste, Al-magnetite, ferrian-picotite, spinel, pleonaste, and hercynite. Additionally, both mineral groups share structural components (XR^{2+}, XR^{3+}, and XTi) that allow them to be co-plotted and evaluated on the Fe-Ti oxide ternary system within `mineralML`.

2.3 Synthetic Mineral Generator

To address class imbalance in the training dataset for minerals less commonly analyzed by EPMA, we developed a synthetic mineral generator. These less oft analyzed minerals include zircon, kalsilite-nepheline, SiO_2-polymorphs, and hematite. The algorithm simulates natural compositional variability by mixing idealized endmember stoichiometries with stochastic perturbation, while enforcing charge balance and oxygen normalizations. Inputs to the algorithm are endmember compositions (in terms of the counts of cations and anions), oxygen basis, mixing distribution, and noise.

Synthetic base mineral compositions are generated by mixing endmembers using quantities sampled from probability distributions. Single endmember systems are mixed using slight perturbations to their base stoichiometry. Binary systems are mixed with a beta or uniform distribution. Multi-component systems can be mixed with a Dirichlet distribution. Stochasticity is introduced in two parts, to simulate natural variability and analytical uncertainty. Minor and trace elements are introduced with an exponential distribution on these cations, which displaces some charge provided by the major cations. Major cation abundances are then perturbed with a log-normal distribution to simulate natural variability with positive concentrations. Gaussian noise is then added to individual elemental concentrations.

These independent perturbations can result in non-physical formulas, so the total cationic charge is then computed from ion counts and valences. Charge balance between the cationic charge and anionic oxygen charges is then confirmed, and synthetic compositions with charges disagreeing by more than 10% are discarded. From there, cation abundances are re-normalized to the oxygen basis. These cation counts are converted into oxide weight percentages. To validate that synthetic mineral generation reflects empirical distributions, the algorithm compares synthetic elemental distributions against our training dataset with Kolmogorov-Smirnov statistics. The synthetic mineral generator's performance is validated by examining well-known and often analyzed solid solution minerals, olivine and feldspar (Supporting Information S5).

2.4 Imputing and Pre-processing Data

Prior to model training, both the input features of oxide concentrations and the target labels of mineral classes required pre-processing. Oxide concentrations often have missing data (e.g., no Cr_2O_3 measurement), which is unavoidable when aggregating analyses from numerous studies. These missing data values may occur for two primary reasons: 1) the oxide or element was below the detection limit, or 2) the oxide or element was not measured, either to reduce analytical time or because it was not expected in the mineral phase. To maintain consistent features across all mineral groups, missing oxide values were imputed (filled-in or substituted) with zeros. Analyses with fewer than three reported oxides were excluded in the training dataset to avoid samples with insufficient compositional information; this exclusion did not limit minerals such as quartz and hematite, as trace elements are often added to the analytical routine to fully utilize the multiple spectrometers operating simultaneously for EPMA analyses.

Oxide training data were pre-processed with a standard z-score scaler, implemented through `sklearn.preprocessing.StandardScaler`, which normalizes each oxide column (feature) by subtracting the mean and scaling the values to unit variance:

$$\tilde{\mathbf{x}}_\mathbf{i} = (x_i - \mu)/\sigma \tag{1}$$

where μ is the mean, σ is the standard deviation (Pedregosa et al., 2011), and i indexes the oxide. The z-score normalization standardizes the data to a mean of zero with unit variance. This normalization ensures that feature distributions occupy similar ranges and that oxides with higher concentrations do not have outsized influence on the neural network. Oxide data were not normalized to 100 wt.% prior to model training, as in many cases, a deviation in the total from 100 wt.% is a true feature; these deviations may represent OH^- in amphibole, Fe^{3+} in oxides, H_2O in glasses, or C in carbonates.

While Compositional Data Analysis (CoDA) using log-ratio transformations (Aitchison, 1982; Greenacre, 2018) is often recommended to treat mineral oxide data (Boschetty et al., 2022; Lui et al., 2026) given the closure issue, positive values are required. The presence of a large number of missing values in this dataset is a feature, rather than being incidental. Zircon contains primarily ZrO_2 and SiO_2 and the vast majority of analyses have no values reported for Na_2O, K_2O, and Cr_2O_3, while hematite is almost all Fe_2O_3, usually measured as FeO_t, and routinely has no values reported for Na_2O, K_2O, and P_2O_5. In the CoDA framework, these zeros must be imputed to small positive values. Given the phase-dependent sparsity of our dataset, this extensive imputation altered the data structure and introduced artificial variance, producing spurious linear clusters which masked true geochemical signals (Supporting Information S3, Supporting Information Figure S1). The transformed feature space ended up reflecting the imputation strategy rather than the underlying geochemistry, the log-ratios were rendered useless for model training and so we used z-scoring instead.

2.5 Neural Network Training and Architecture

Neural networks are a class of machine learning models loosely inspired by the structure of biological neural systems (i.e., brains), in which stacked computational units transform inputs into progressively more abstract representations. A typical feed-forward network is organized into an input layer consisting of the input data, one or more hidden layers, and an output layer that produces class predictions or a continuous value. Each hidden layer is composed of a set of nodes, each of which computes a weighted sum of the previous layer's outputs and adds a bias, forming the pre-activation passed through a non-linear activation function. Weights control how input features are combined, and biases provide an offset that shifts neuron response. Learning starts by defining a loss function that measures prediction error. The network's learned parameters are then iteratively updated with an optimizer, such as stochastic gradient descent or its variants, to minimize this loss by adjusting parameters in the direction most steeply reducing the loss. The gradient of the loss is computed using the backpropagation algorithm, through the repeated application of the chain rule backwards through the network. A key advantage of neural networks is their modularity. Once a network has learned an intermediate representation (an embedding), downstream components can be swapped or added — such as classifiers, decoders, or bottlenecks — to reuse that representation for different objectives without retraining the full model.

Prior to training the neural network, a few empirical classifications were first applied to three labels of minerals. SiO_2 polymorphs are empirically designated with a filter of $SiO_2 > 90$ wt.%, as these minerals are almost pure SiO_2 with some quantity of trace elements. Carbonates are empirically designated with a filter of $SiO_2 < 5$ wt.% and an oxide total < 70 wt.%. Empirical designations of these minerals also help to minimize skewing of the normalization provided by the z-score scaler. Zircons are also removed from the classification schema to avoid the introduction of the ZrO_2 component to feature space, where every other mineral group would have very small or null values. Analyses of these minerals were subsequently removed from the neural network classification schema.

For the remaining minerals, a two-stage neural network model is used that (i) performs probabilistic mineral-phase classification from oxide compositions and (ii) constructs a lower-dimensional latent representation via a reconstruction objective trained after the classifier is learned (Figure 2). The primary task is to perform probabilistic mineral classification, and the auxiliary task is unsupervised data reconstruction. We employed a sequential, transfer learning approach. The classifier first learns an intermediate feature embedding, $\mathbf{h}$, from the oxide inputs whilst performing classification in the classification branch (Figure 2A-C), followed by a decoder reconstructing the oxide inputs from a 2-D bottleneck in the reconstruction branch (Figure 2D). This architecture diverges from multi-task learning architectures that optimize multiple objective functions simultaneously, as with deep learning applications in seismology (Zhu et al., 2025; Crawshaw, 2020). While concurrent optimization of multiple objective functions is efficient in seismology, it may also force the network to compromise on both the classification and learned latent distribution. `mineralML` does not jointly optimize classification and reconstruction losses in a single training phase, but instead reuses a frozen representation learned for classification before reconstructing the representation. The architecture allows for the latent representation to provide an additional visualization in 2-D, confirming the structure, without directly influencing the classification.

In the first stage through the classification branch, the network passes the normalized oxide inputs, $\tilde{\mathbf{x}}$ through a series of hidden layers (collectively, termed the feature extractor; Figure 2B) followed by a classification head which produces classification outputs $\hat{y}$ (Figure 2C). The layers of neurons structured in hidden layers progressively transform the oxide data into a feature embedding (Figure 2B), which is a multi-dimensional embedding that captures the most relevant emergent features with respect to classifi-

cation. Training proceeds iteratively. At each step, a subset of training analyses (a batch) passes through the network. Label predictions are compared against known labels, and weights are updated accordingly. The following equations describe the transformation applied to each analysis at hidden layer l:

$$\mathbf{a}^{(l)} = \mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)} \tag{2}$$

$$\mathbf{h}^{(l)} = \mathrm{Drop}\left(\phi\left(\mathrm{BN}(\mathbf{a}^{(l)})\right)\right) \tag{3}$$

Equation 2 describes what happens at each hidden layer, where the outputs from the previous layer $\mathbf{h}^{(l-1)}$ are multiplied by the weight matrix $\mathbf{W}^{(l)}$, and shifted by the bias vector $\mathbf{b}^{(l)}$ to produce a pre-activation output $\mathbf{a}^{(l)}$. In the first layer, $\mathbf{h}^{(0)}$ is simply the normalized input oxide data ($\tilde{\mathbf{x}}$), while in subsequent layers, the inputs are the post-activation outputs of the preceding layer. Equation 3 demonstrates the application of three successive operations to $\mathbf{a}^{(l)}$. BN() denotes batch normalization, rescaling the pre-activation values to have zero mean and unit variance in each training batch to stabilize the learning process. $\phi()$ denotes the non-linear activation function aimed at introducing non-linearity — in this case, we use the Leaky Rectified Linear Unit (LeakyReLU) activation function (Maas et al., 2013). Drop() denotes dropout, applied solely in the training process, in which a defined fraction of neuron activations is set to zero at each update step. Dropout aims to minimize reliance through any single pathway through the network, improving generalization and preventing overfitting. We use a dropout rate of 10% during training, meaning that 10% of activations are set to zero after each iteration. No dropout is applied at the time of inference.

The first two hidden layers use fixed point-estimate weights, but the third and final hidden layer of the feature extractor is variational in nature (Figure 2B). Rather than learning a single weight matrix, this layer parameterizes its weights as normal distributions with learned means (μ) and variances (σ^2). During each forward pass, weights are sampled from these learned distributions so the resulting feature embedding $\mathbf{h}^{(L)}$ is itself stochastic. This design follows the Bayes by Backprop framework in which weight uncertainty is maintained and propagated through the network (Blundell et al., 2015). Placing the variational layer prior to forming the embedding means that uncertainty is encoded into the learned features, propagated forth to both classification and reconstruction. The variational layer allows for weight distributions to be assessed at inference through repeated sampling. For analyses resembling training data, the weight distributions are well-constrained so repeated sampling will produce similar logits and subsequently tightly clustered prediction score arrays. For unfamiliar compositions, such as minerals absent from the training data or anomalous compositions, the weight distributions are poorly known so sampled weights will vary more widely (see Section 4). High variance in the logits therefore flags analyses that the model is not well equipped to classify.

The classification head is a standard hidden layer that maps this stochastic feature embedding to raw prediction scores for each mineral class (Figure 2D). These scores, termed logits, are the direct output of the final layer preceding normalization. A large positive logit would indicate that the network associates the input with the class, whilst a large negative value would do the opposite. Near-zero values indicate ambiguity. As logits are unbounded and do not sum to one, they cannot be directly compared across classes, so the softmax function is applied as a non-linear transformation that normalizes logits to a sum of one:

$$p_c = \frac{e^{s_c}}{\sum_{j=1}^{C} e^{s_j}} \tag{4}$$

where p_c is the prediction score for an individual mineral class c (with values closer to 1 indicating stronger association between the composition and the predicted mineral class), s_c is the raw logit for class c, j indexes all classes in the denominator, and C is the total number of classes. Given that the feature embedding is stochastic, each forward pass through the network will produce slightly different logits $\mathbf{s}$ and therefore a different prediction score array $\mathbf{p}$ over the C classes. Running the same oxide compositions through the network S times thus yields S prediction score arrays, each reflecting a different sample of the variational weights with the spread reflecting model confidence.

During training, the prediction score array $\mathbf{p}$ from each forward pass is compared against the known mineral labels using an objective function we denote as the total classification loss ($L_{\text{Classification}}$):

$$L_{\text{Classification}} = L_{\text{CE}} + \lambda_{\text{KL}} L_{\text{KL}}, \tag{5}$$

The first term L_{CE} is the cross-entropy loss, which measures how well the predicted labels match the true labels:

$$L_{\text{CE}} = -\frac{1}{N} \sum_{n=1}^{N} \log(p_{y_n}) \tag{6}$$

where N is the number of samples in the batch, y_n is the true class label of sample n, and p_{y_n} is the prediction score assigned to the true class of sample n. The loss is the negative log of the prediction score assigned to the correct class, averaged over the batch.

Cross-entropy loss penalizes confident and incorrect predictions whilst rewarding confident and correct predictions. The second term L_{KL} is the Kullback-Leibler (KL) divergence, which measures how much the learned weight distributions differ from a reference distribution or prior (here, a standard normal distribution). This term prevents the variational layer from collapsing toward deterministic weights and preserves the uncertainty-quantification properties described above. The KL weight λ_{KL} is increased from zero over the first training epochs — a process called KL annealing, allowing the network to first focus on classification accuracy before the constraint on the shape of the distribution is tightened (Bowman et al., 2016).

Once classification losses are minimized, weights of the feature extractor and classification head are frozen (i.e. fixed) and no longer updated. The learned feature embeddings $\mathbf{h}^{(L)}$ serve as input to the second stage, which trains the autoencoder. An autoencoder is a neural network architecture consisting of two parts: an encoder that compresses its input into a lower-dimensional representation, and a decoder that attempts to reconstruct the original input from that compressed representation (Figure 2D). The network is trained purely on reconstruction quality, as there are no labels involved. Here, the encoder compresses $\mathbf{h}^{(L)}$ into a two-dimensional bottleneck, $\mathbf{z}_2$, using similar hidden layers to those of the classification branch, with layer normalization and LeakyReLU activations. The decoder then maps $\mathbf{z}_2$ back to the original normalized oxide space, $\tilde{\mathbf{x}}$. The reconstruction loss is the mean squared error (MSE) between the original and reconstructed inputs:

$$L_{\text{MSE}} = \frac{1}{N} \sum_{n=1}^{N} ||\tilde{\mathbf{x}}_n - \hat{\mathbf{x}}_n||^2 \tag{7}$$

where $\tilde{\mathbf{x}}_n$ represents the normalized input parameters and $\hat{\mathbf{x}}_n$ represents the model's reconstructed outputs. The two-dimensional bottleneck $\mathbf{z}_2$ can be plotted directly as a

latent space in which the position of each analysis captures geochemical similarity. This makes $\mathbf{z}_2$ comparable to other dimensionality reduction techniques used in petrology, although the sequential architecture of `mineralML` offers advantages over these methods. Linear dimensionality reduction techniques, such as principal component analysis (PCA), are commonly used to capture axes of maximum variance in petrologic data (Boschetty et al., 2022; Shi et al., 2024; Toth et al., 2025) and find axes of maximum compositional variance irrespective of the classifier label. An autoencoder trained directly on raw oxide compositions would behave similarly, compressing data along directions of greatest variance. As the input to our autoencoder is a feature embedding already optimized for phase classifications, the 2-D latent space captures structure that reflects mineralogical distinctions rather than bulk compositional variance (Figures 1C, 2D).

`mineralML` is implemented in `PyTorch` (Paszke et al., 2019) with pre-processing in `scikit-learn`. Model training was performed on a CUDA-enabled NVIDIA RTX 4500 Ada Generation GPU for accelerated computation. Hyperparameters including the learning rate, hidden layer sizes and numbers, and regularizers were tuned by selecting the configuration that minimized loss for each stage: cross-entropy loss with the KL term for classification and MSE loss for reconstruction. The maximum number of epochs, or number of times the model is shown the entire dataset, was capped at 500 with an early stopping mechanism. If validation loss failed to improve for 50 consecutive epochs, the training process was prematurely stopped. The batch size, or the subset of training samples processed before updating the model's weights, was fixed at 128 to introduce stochastic variation into the gradient updates (Keskar et al., 2016). Both training stages were independently optimized with AdamW (Loshchilov & Hutter, 2017), a variant of the Adam (Kingma & Ba, 2014) optimization algorithm with decoupled weight decay.

2.6 Application of the Trained Model

`mineralML` is distributed as an open-source `Python` package with a pre-trained model. At inference, the learned neural network parameters are fixed, except for the variational layer, whose weights are sampled from learned distributions across repeated forward passes to capture prediction uncertainty. `mineralML` returns predicted mineral phases with associated prediction scores with their uncertainties. The model is incredibly fast to run; millions of analyses require 30 seconds to 2 minutes to run on a laptop. To use the model, new oxide compositions from tabular or mapped data are input to obtain mineral classifications, prediction scores, and optional latent-space projections (Figure 1). Both mineral label predictions with prediction scores and latent space projections can be executed in one line of code. Open-source Jupyter notebooks are provided to guide users through the workflow on the ReadTheDocs page `https://mineralml.readthedocs.io/en/latest/`.

Compositions are transformed into phase predictions with their corresponding prediction scores with the function `mineralML.predict_class_prob_nnwr`. Oxide compositions are first screened with the empirical criteria described in Section 2.5 to assign SiO_2 polymorphs, carbonates, and zircons directly from oxide data (Figure 1A). The remaining compositions are z-score normalized with the scaler fit on the training dataset. Compositions are passed through the trained classifier with Monte Carlo inference. The classifier is evaluated $S = 50$ times per composition, with each forward pass sampling from the variational weight distributions to produce a prediction score vector $\mathbf{p}^{(s)}$. Final class prediction scores are reported as the mean across draws:

$$\bar{\mathbf{p}} = \frac{1}{S} \sum_{s=1}^{S} \mathbf{p}^{(s)} \tag{8}$$

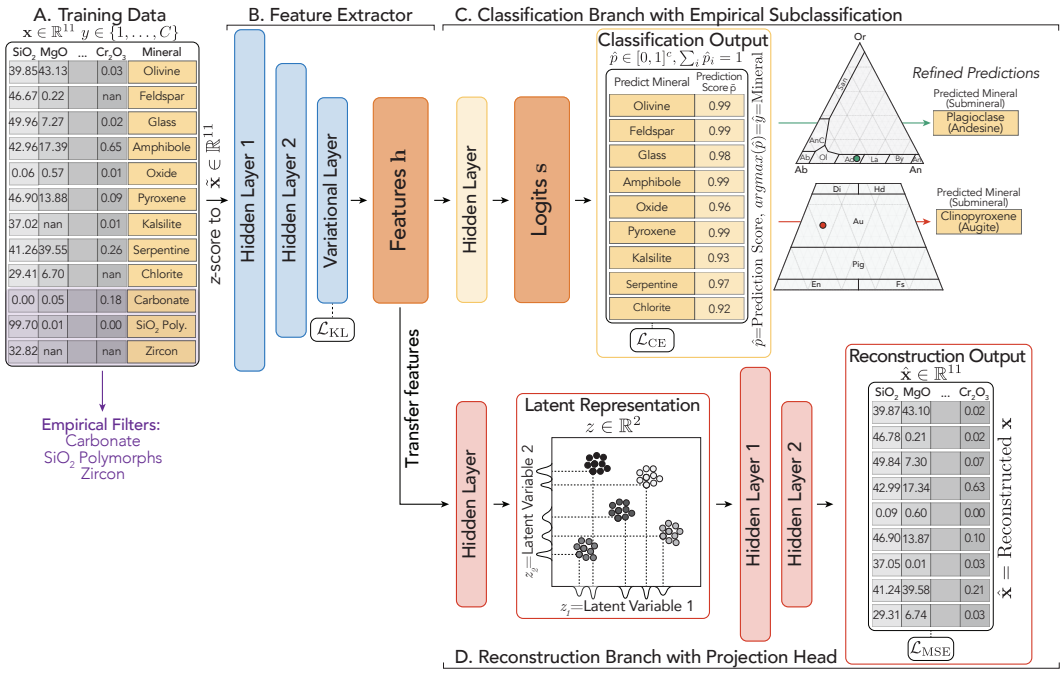

Figure 2. Two-stage neural network architecture for probabilistic mineral classification and 2-D latent space visualization, used in the training stage. The trained model is used for inference. A) Input oxide compositions $\mathbf{x}$ are z-score normalized to $\tilde{\mathbf{x}}$. B) $\tilde{\mathbf{x}}$ is passed through a series of fully-connected hidden layers (the feature extractor), where the final layer is variational. Weights are parameterized as distributions rather than point estimates, regularized by $\mathcal{L}_{\text{KL}}$. The post-activation output of the feature extractor forms a stochastic feature embedding $\mathbf{h}^{(L)}$, whose variance across repeated forward passes reflects model confidence. C) A classification head maps $\mathbf{h}^{(L)}$ to logits and class prediction scores $\hat{\mathbf{p}} \in [0,1]^C$ with softmax. Loss is measured with cross-entropy loss $\mathcal{L}_{\text{CE}}$. Classification output contains predicted mineral, prediction score, and prediction score variance across forward passes. Broader mineral classes (feldspar, pyroxene, oxide) are refined through empirical filters. Once training of the classifier is complete, the feature extractor and classifier's weights are frozen and fixed. D) An autoencoder is trained to compress $\mathbf{h}^{(L)}$ into a two-dimensional latent bottleneck $\mathbf{z}_2$. A decoder reconstructs normalized inputs $\hat{\mathbf{x}}$, optimized with mean squared error loss $\mathcal{L}_{\text{MSE}}$. The reconstructed outputs $\hat{\mathbf{x}}$ are compared against the original inputs $\tilde{\mathbf{x}}$ to evaluate reconstruction.

where $\bar{\mathbf{p}}$ is a vector of prediction scores across all classes considered by `mineralML` and the predicted mineral class is assigned as $\arg\max_c \bar{p}_c$ (Figure 1C). The standard deviation across draws:

$$\sigma_{\bar{\mathbf{p}},c} = \sqrt{\frac{1}{S}\sum_{s=1}^{S}\left(p_c^{(s)} - \bar{p}_c\right)^2} \tag{9}$$

quantifies prediction uncertainty for each class c. Low values of $\sigma_{\bar{\mathbf{p}},c}$ indicate that the variational layer produces consistent predictions across forward passes, suggesting that those compositions are well represented in the training data. High values of $\sigma_{\bar{\mathbf{p}},c}$ indicate poorly constrained compositions, signaling that the classified mineral should be interpreted with caution.

The mineral with the highest $\bar{p}_c$ is returned as the primary classification, along with its prediction score. In the first row of tabular data in Figure 2, for example, if the prediction scores are $\bar{p}_{\text{ol}} = 0.99$ and $\bar{p}_{\text{serpentine}} = 0.009$; the analysis would thus be classified as olivine with serpentine as the second-ranked candidate. If the $\sigma_{\bar{\mathbf{p}},\text{ol}}$ is 0.005, that would indicate that the classification is highly confident.

Following neural-network prediction, an empirical subclassification step is applied to refine mineral assignments within broader groups returned by the classifier (e.g., feldspar, pyroxene, oxides). Users can also project their compositions onto the two-dimensional latent representation with the function `mineralML.plot_z2_overlay` reusing the frozen feature embedding $\mathbf{h}$ and computing $\mathbf{z}_2$ for visualization (Figure 1C). The projection overlays sample compositions onto the two-dimensional latent space of the training and validation dataset. As this step of the latent projection is trained from the feature vector, the neural network shares this feature but the projection does not influence the classification output. The workflow is reusable either for working with tabular point analyses from EPMA, EDS, or a database, or for working with mapped EDS data and linescans.

3 Results

`mineralML` returns a primary mineral classification, a second-ranked mineral classification, and associated prediction scores (along with the standard deviation of the first prediction score across Monte Carlo passes) for each classification. For the four broader groups returned by the neural network classifier—feldspar, pyroxene, rhombohedral oxide, and spinel group—empirical subclassifications refine the prediction to specific minerals. It is these refined classifications that are reported and demonstrated throughout this section.

3.1 Performance on Validation Data

Model performance was evaluated on a validation dataset, withheld from parameter updates through the training process. The validation dataset was subsequently passed through the trained neural network to provide a first understanding of performance. Of the validation dataset's ∼99,000 analyses, 90 predictions of mineral group (0.09%) diverged from the stated mineral. The relatively low percentage of differences in the stated and predicted mineralogy is unsurprising, given that the validation data are drawn from the same underlying datasets as the training data with stratified and/or random sampling. Differences in predictions primarily existed within the stated amphibole, plagioclase, and glass groups. No glaring issues existed within the validation data. One minor difference is attributed to differences between stated glasses and predicted amphiboles or feldspars. Nearly all of these glasses were predicted as amphiboles, although with prediction score $\bar{p}$ of 0.73 and prediction score sigma $\sigma_{\bar{p}}$ of 0.24; these predictions also all have the second most likely predicted label of glass. These compositions lie in the region where there is indeed overlap between glass and amphibole compositions. The trained model may not have the ability to resolve between geochemically near-identical compositions. The reverse is also true, with another eight analyses stated as amphiboles classified as glass. The remaining glasses were predicted as feldspars or within the feldspar miscibility gap with a prediction score mean of 0.87; these analyses may reflect groundmass compositions excitation of feldspars during microanalysis. A more rigorous assessment of model generalization requires evaluation on independently collected data that are not represented in the training dataset. We thus turn to the GEOROC dataset, microprobe standard compilation, and co-collected mapped datasets.

3.2 Performance on GEOROC Data

We tested the model using 15 precompiled GEOROC files for major mineral groups exported on 1 December 2024 (DIGIS Team, 2024, 2025). This dataset consists of ∼1.1 million analyses, which can be processed with `mineralML` in approximately 2 minutes of computational time on a $1000 Windows ThinkPad laptop (Intel I7, 16 GB RAM) or in approximately 30 seconds of computational time on a M2 Macbook Pro with 12 CPU cores. Some minerals in GEOROC are given a generic group name rather than the name of a specific mineral (e.g., both biotite and muscovite may be labeled mica; clinopyroxene is labeled as pyroxene; and SiO_2 polymorphs are labeled as quartz). Using these more generic names, the initial performance of `mineralML` on the unfiltered GEOROC database was very promising: 99% of oxides, 98% of pyroxene and titanite, 97% of feldspar, olivine, quartz, and mica, and 96% of apatite are classified correctly. However, some minerals show poorer initial performance; only 90% of zircons, 88% of garnets, and 86% of amphiboles are classified correctly.

Failures of `mineralML` can be split into two main categories: database errors and classification errors. It is no small task to work out which are which; the 2% of pyroxene misclassifications represent 5673 individual analyses from 730 unique studies! For each mineral, we investigate the studies showing the largest absolute number of incorrect classifications to estimate the prevalence of database versus classification errors. To do this, we make use of the `mineralML` functions to plot analyses in major element space (`mineralML.plot_harker`) and latent space with the calibration dataset (including training and validation data) of each mineral underlain (`mineralML.plot_z2_overlay`).

First, we found that some studies have misclassified minerals (e.g., they analyzed what they thought was an amphibole and labeled it as such in the supporting information, but it was actually a clinopyroxene). Such mistakes are sufficiently widespread that identifying them was one of the key motivations for developing `mineralML`. For example, the 'amphiboles' from Torres-Orozco et al. (2023) show a clear bimodal composition, with one group possessing high Al_2O_3 (10-14 wt.%) and TiO_2 (2-3 wt.%), and low CaO (12 wt.%) and SiO_2 (40-42 wt.%), while a second group have 50-53 wt .% SiO_2, high CaO (20-23 wt.%) and low TiO_2 contents (<1 wt.%). This second group of 59 analyses are classified as 'Cr-actinolites' by the authors in the supplement. However, `mineralML` classifies this high SiO_2 group as pyroxenes, which is supported by the fact that these analyses overlap with the reported pyroxene analyses from the same samples (Figure 3). Torres-Orozco et al. (2023) performed amphibole thermobarometry and calculated melt SiO_2 contents from both compositional groups, yielding pressures of ∼ 400 MPa and SiO_2 contents of 64-65 wt.% for the true amphiboles, but anomalously low pressures (60-80 MPa), high SiO_2 contents, and 'incoherent oxygen fugacity values' for the 'Cr-actinolites'. The authors suggest, at least in part based on the lower pressures and higher SiO_2 contents obtained from the 'Cr-actinolites', that there is a final, shallow storage zone of rhyolitic melts prior to eruption. However, we believe these calculations to be spurious as it is not valid to apply amphibole thermobarometers and chemometers to clinopyroxene. This example shows the power of `mineralML` to filter datasets prior to performing thermobarometry calculations, and also to filter literature datasets when investigating prior interpretations from specific volcanic systems (Figure 3).

While a number of other studies appear to have a small number of true misclassifications (e.g., one or two mislabeled clinopyroxenes in a large group of amphiboles), it is more common that the phase was documented correctly in the supplement but was tabulated incorrectly in GEOROC. For example, `mineralML` classifies 1,067 'clinopyroxene' analyses from Tappe et al. (2009) as biotite. These do indeed plot with the biotite calibration dataset, and the supplement clearly labels them as phlogopite (a Mg-rich biotite), so clearly a mistake was made upon data entry documenting these as clinopyroxene. The dataset of Yamasaki (2022) has 78 amphiboles where GEOROC and `mineralML` agree on the classification, but `mineralML` classifies 35 biotites and 29 plagioclases from

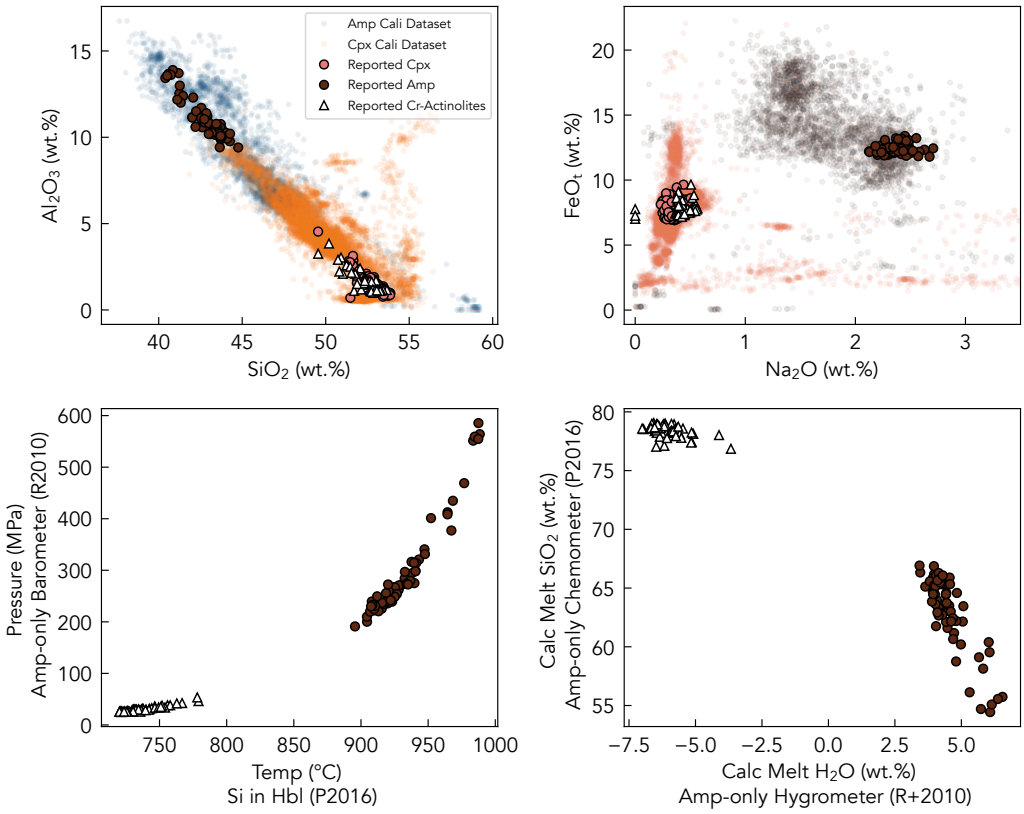

Figure 3. Evaluation of the Torres-Orozco et al. (2023) dataset where `mineralML` predicts that 59 analyses labeled as amphibole (specifically, Cr-Actinolites) are clinopyroxenes. A-B) When plotted in oxide-oxide space, it is apparent that the Cr-actinolites overlap with the major element compositions of reported pyroxenes from the same samples, and are far more aligned with the clinopyroxene calibration dataset in `mineralML` than that of amphiboles. C-D) A comparison of different thermobarometric and chemometric proxies provided by the authors in the supplement show that the misidentification of clinopyroxene as amphibole results in anomalously low storage pressures, temperatures, and melt water contents, and anomalously high predicted melt SiO_2 contents. P2016: Putirka (2016), R2010: Ridolfi et al. (2010).

analyses labeled as amphibole. When plotted in latent space, the plagioclase and biotite analyses sit alongside the calibration data for these phases, and away from that of amphibole (Figure 4). From comparing the datatables in the paper, it is clear that the mineral label for these analyses were mis-typed in GEOROC, likely as a function of having been part of a large in-text datatable. Similarly, GEOROC reports 36 amphiboles from Parman et al. (2003), but `mineralML` classifies 30 of these as clinopyroxene (2 as amphibole and 4 as serpentine). The Parman et al. (2003) paper does not mention amphibole at any point, and catalogs these analyses as clinopyroxene in the datatables, apart from the 4 analyses it labels as 'serpentine/chlorite' that `mineralML` classifies as serpentine. Additionally, of the 85 'olivine' analyses in GEOROC from Thivet et al. (2020), 12 are classified by `mineralML` as clinopyroxene, and 42 as plagioclase. Harker diagrams demonstrate that there are clearly three distinct geochemical groups, which overlap with the calibration dataset of the allocated `mineralML` label. Finally, the 94 'olivine' analyses

in GEOROC from Horn et al. (2022) also plot as three separate groups, and `mineralML` correctly identifies 3 of them as amphibole, and 38 of them as plagioclase.

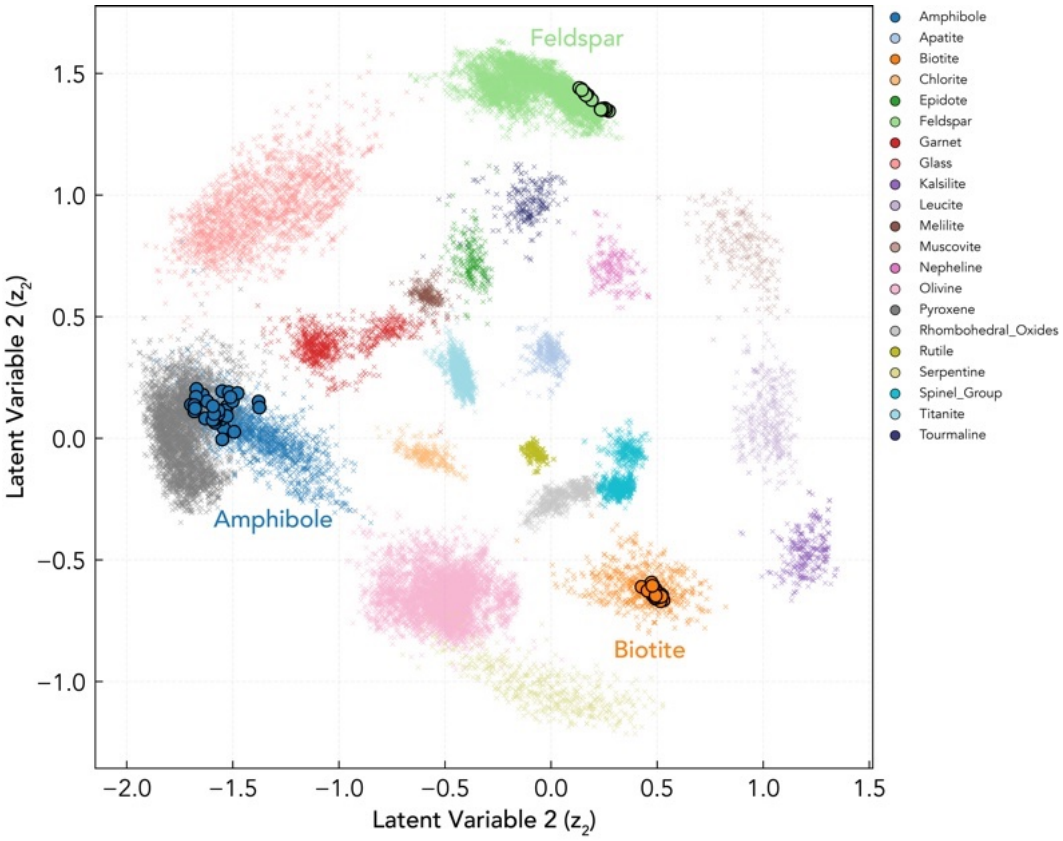

Figure 4. Latent space showing the analyses of Yamasaki (2022) in GEOROC which were labeled as 'amphibole', with the `mineralML` training dataset for each mineral underlain with semi-transparent symbols (marked with crosses). While a subset of these analyses (dark blue circles, labeled towards the center left) are certainly amphiboles, it is clear that the analyses allocated biotite (orange colored circles, labeled on the bottom right) and feldspar (pistachio colored circles, labeled on the upper center) by `mineralML` show far better latent affinity with the calibration data of these labels than with amphibole.

Errors can also be introduced when compiling microprobe data into supporting spreadsheets. In our experience, microprobe software often outputs elements in different orders from different sessions, which can cause issues when stitching together multiple days of analyses. For example, the dataset of Castro et al. (2023) contains 392 analyses labeled as amphibole; `mineralML` agrees with 149 of these, but classifies the other 243 as oxides. It is apparent from plotting a Harker diagram that two clear groups emerge: one plotting alongside the amphibole calibration data, and the other group having extremely low MgO contents, and extremely high MnO contents. In the supporting datatable, the values of MnO and MgO are swapped after row 149 (this was confirmed based on the calculated 'cations pfu' column in the original supporting information). We find that column-swapping was the cause of a surprisingly large number of apparent misclassifications. For example, of the 89 amphiboles in García et al. (2020), `mineralML` classifies 45 as amphiboles, and 44 as pyroxenes (all with low prediction scores, 0.45-0.8). Harker diagrams show that these amphiboles are extremely Cr-rich (∼3-16 wt.% Cr_2O_3), and Al_2O_3-poor.

It is apparent from the supplementary table (and thus GEOROC) that the Cr_2O_3 and Al_2O_3 columns have been swapped for all rows (again, the cation columns are correct). In Larocque and Canil (2010), of the 73 amphiboles, 71 are classified as oxides and 1 is classified as chlorite, with low prediction scores (40-80%). It appears that the MnO and FeO columns were swapped in the supporting information, based on our calculated $Mg\#_{>98}$ not matching those stated in the text ($Mg\#_{77-83}$), and the fact that anomalously high MnO contents are not mentioned in the paper (10-25 wt.%).

In other studies, the columns are correct in the author data table, but have been copied incorrectly into GEOROC. In Didonna et al. (2024), there are 1167 correctly classified olivines which sit alongside the olivine calibration data in major element space. The 1216 'olivine' that `mineralML` classifies as oxides have extremely high Na_2O and MnO contents and very low SiO_2 and FeO_t contents. It is apparent that all the 'misclassifications' arise from table B1 in the supplement, where the following column swaps occurred when data was entered into GEOROC:

- $SiO_2{}^{\text{supplement}} \rightarrow Na_2O^{\text{GEOROC}}$,
- $TiO_2{}^{\text{supplement}} \rightarrow SiO_2{}^{\text{GEOROC}}$,
- $Al_2O_3{}^{\text{supplement}} \rightarrow MgO^{\text{GEOROC}}$,
- $Na_2O^{\text{supplement}} \rightarrow Al_2O_3{}^{\text{GEOROC}}$,
- $MgO^{\text{supplement}} \rightarrow MnO^{\text{GEOROC}}$,
- $MnO^{\text{supplement}} \rightarrow FeO^{\text{GEOROC}}$,
- $FeO^{\text{supplement}} \rightarrow Cr_2O_3{}^{\text{GEOROC}}$.

Similarly, in Repstock et al. (2022), of the 151 pyroxenes, 40 are classified by `mineralML` as nepheline and 111 as oxides. It is apparent that:

- $FeO^{\text{supplement}} \rightarrow Cr_2O_3{}^{\text{GEOROC}}$ and vice versa,
- $CaO^{\text{supplement}} \rightarrow MgO^{\text{GEOROC}}$,
- $Na_2O^{\text{supplement}} \rightarrow CaO^{\text{GEOROC}}$ (among other swaps)

Finally, in Egorova and Latypov (2013), of 477 pyroxenes, 372 are classified correctly, but 99 are misclassified as olivine and 6 as rhombohedral oxide due to the columns FeO and MnO being swapped in GEOROC.

In other cases, GEOROC is missing entire columns that were originally present in the supporting information. In Callegaro et al. (2021), 141 true pyroxenes are classified by `mineralML` as amphibole (88), calcite (3), epidote (9), garnet (33), rhombohedral oxide (3) and tourmaline (5). It is apparent that CaO, Na_2O and Cr_2O_3 data is missing in GEOROC, leading to this misclassification. In Polat et al. (2018), the first 35 entries in GEOROC have no CaO data, causing clinopyroxene to be misclassified as nepheline and glass. Once CaO contents from the supplements are added back, `mineralML` classifies these correctly. Finally, in Carlson and Ionov (2019), 59 clinopyroxenes are misclassified as titanite based on LA-ICPMS data, because FeO or MgO were not measured, so no values were entered in GEOROC.

We also found some column swaps that occurred when data were copied into GEOROC. These column swaps persist in the pre-compiled mineral files, but have been fixed in GEOROC itself. For example, in McCarthy and Müntener (2017), 682 amphiboles are classified correctly, but 88 are incorrectly classified as calcite, and 22 as orthopyroxene. The misclassified ones are missing MgO, CaO, Na_2O, K_2O and Cr_2O_3 data (all zeros) in the pre-compiled datasets, but not in the since-updated GEOROC entries associated with the paper. Similarly, $SiO_2{}^{\text{supplement}} \rightarrow Na_2O^{\text{GEOROC}}$, in Boudoire et al. (2023), resulting in stated misclassifications from `mineralML`. There are additionally anomalously high CaO reported in olivine from Viljoen et al. (2022) and anomalously low SiO_2 contents in olivine in Ionov et al. (2017). The exact root of these issues is unclear.

In contrast to these data errors, the lower prediction rate of amphiboles, garnets, and zircons after the initial round of training represented true model failures; these model failures were fixed with subsequent retraining discussed in the following paragraph. For example, the 655 amphiboles from the Main Ethiopian rift rhyolites measured by Tadesse et al. (2023) were classified as glass and Na-pyroxene. In oxide-oxide space, it was immediately apparent that amphiboles with Mg#<0.5, low Na_2O, CaO and MgO contents, high FeO_t were not represented in the training dataset. This was a recurring problem in very evolved and alkaline systems, in which amphiboles were routinely being classified as glass and to a lesser extent pyroxene (e.g., Teiber et al. (2015); Pontesilli et al. (2022); Park and Kim (2022); Siegel et al. (2017)). There initially were 350 garnets measured by Braunger et al. (2018); Melluso et al. (2022); Whitley et al. (2020), all of which were classified by the neural network as 'epidote'. The high number of differences in classifications revealed the dearth of garnets with high CaO (grossular and andradite) and MnO contents (spessartine) within the training dataset. The lack of representation of these garnets was particularly apparent in more silica-undersaturated (phonolitic) and in carbonatite samples. The 145 zircons in GEOROC which were misclassified lacked ZrO_2 data, making the empirical classification based on ZrO_2 content impossible. There were 89 null values for ZrO_2; the remaining 56 values for ZrO_2 had a maximum value of 49.94 wt.%, which is lower than what is expected based on the stoichiometric assumption. Most of these studies measure U-Pb ages recorded by these zircons, and therefore do not focus on major element chemistry (Ondrejka et al., 2018; C. Wang et al., 2023). The remaining 1368 zircons in GEOROC are predicted correctly. Similarly, `mineralML` initially did a very poor job of predicting olivine with Mg#s<0.2 (i.e. fayalites). While a number of the misclassifications represented database issues including but not limited those described above (e.g., Viljoen et al. (2022); Giuffrida et al. (2018); Didonna et al. (2024)), true fayalite was also being classified as garnet (e.g., Brandt et al. (2021); Swallow et al. (2019)), likely as a consequence of the fact that extremely few olivines had MgO<30 wt.%, FeO_t >50 wt.%, or MnO>0.5 wt.% in the training dataset. In alkaline systems with Na-rich pyroxenes, `mineralML` tended to struggle to distinguish between clinopyroxene and amphibole, typically predicting approximately half of the pyroxenes were amphiboles, with low prediction scores (30-50%, e.g., X. Wang et al. (2021)). It was clear that the calibration dataset did not contain sufficient analyses of Na-rich amphiboles.

To address these issues, we compiled additional mineral data to help address these compositional gaps in our training data. For example, for amphibole, we compiled Fe-rich amphiboles from GEOROC from studies for which `mineralML` had initially performed poorly. We used the data from 14 studies as training data, and held back an additional 12 studies for testing. There were 3566 amphiboles in the studies selected for training, with only a subset used for model training and the remaining compositions were used for model validation. An additional 1802 amphiboles comprised a distinct test dataset. For the 387 fayalites held back as testing data, 100% are now classified correctly. For the 1802 amphiboles, 94% are classified correctly. 2.7% of the remaining failures are amphiboles from (Hidalgo et al., 2011), which `mineralML` classified as orthopyroxene. These flagged compositions are very unusual compared to the other amphiboles analyzed in this study, and their supporting information shows they flag invalid estimates when entered into thermobarometric estimates, so it is possible that the analyses are not amphibole. Apart from the misclassifications of amphibole as glass where `mineralML` is incorrect, most other failures appear to be database issues, where mineral labels were incorrectly entered. From version `mineralML` v0.0.3.2, the model has been updated with this larger test dataset for Fe-rich amphiboles, fayalites, and Ca- (grossular and andradite) and Mn-rich garnets (spessartine). The iterative training process highlights the perpetually evolving nature of `mineralML`. When compositions are underrepresented, or not represented, we can recalibrate the model and build extended versions suitable for more compositions. This also shows the importance of users citing what version of the model they have used in their study.

3.3 Microprobe Standard Compositions

We also tested `mineralML` on the reported composition of microprobe reference materials. These analyses are a useful test in addition to GEOROC, as these materials are well-characterized and as many of these reference materials are close to pure end-members derived from mantle peridotites, and thus push the model outside its calibration dataset, which is focused more on crustal compositions. First, we use the preferred values of the Smithsonian microprobe standards (Jarosewich, 2002). `mineralML` predicts all amphiboles, apatites, carbonates, pyroxenes, garnets, and feldspars correctly. `mineralML` classifies all igneous glass standards correctly, but classifies the Corning glass archaeological reference materials A-B-C-D as Na-pyroxene (prediction score of 0.81), melilite (prediction score of 0.70), muscovite (prediction score of 0.47) and feldspar (prediction score of 0.47). This is not altogether surprising, as these glasses are designed to resemble Ancient Egyptian/Roman plant ash and natron glasses (A, B), high-Ba East Asian glasses (C), and potash-lime glasses of Medieval European compositions (D) rather than volcanic glasses (Brill, 1971).

The compilation also contains several unusual silicate minerals on which `mineralML` was not trained, providing an opportunity to test the ability of the model to recognize when it is uncertain. Benitoite (a Ba-rich cyclosilicate) is classified as tourmaline with only a prediction score of 0.39. Not all mistaken predictions, however, show low confidence. Gahnite ($ZnAl_2O_4$) is classified with the empirical filter as carbonate (as it has a low total), so has no prediction confidence. Meionite (a carbonate-bearing part of the scapolite group) is classified as plagioclase feldspar with a prediction score of 0.95, and osumilite (a K-Na-Fe-Mg-Al cyclosilicate) is classified as Na-pyroxene with a prediction score of 0.97.

The main failure on minerals it was trained on is the incorrect classification of magnetite as a rhombohedral oxide (hematite). We believe this reflects the extremely pure composition of this magnetite compared to the natural magnetites in our training data, which contain appreciable Al_2O_3, SiO_2, TiO_2. Without these minor components, the composition approaches pure Fe_3O_4 making the Droop (1987) recalculation challenging and the distinction from hematite difficult (discussed further in Section 6).

We also compile analyses of microprobe standards from a variety of supporting datasets (including, but not exclusively, standards from the Smithsonian) from Iddon et al. (2019); Gleeson et al. (2020); Loewen et al. (2021) and F. Keller et al. (2024). All amphiboles, pyroxenes, olivines, and feldspars were classified correctly. Of all the glass analyses, 2 were classified as garnet (0.71-0.76 prediction score) and 2 were classified as nepheline (0.62-0.70 prediction score). These misclassified glasses are from Loewen et al. (2023) and are marked in the supplement as 'rejects', showing the ability of `mineralML` to help identify poor quality analyses. Overall, this standard compilation shows that `mineralML` does an excellent job of correctly predicting common silicate minerals, but that care must be taken when applying it to samples that may contain minerals not present in the training datasets.

3.4 Concurrent EDS and EBSD Mapping

A major motivation for developing `mineralML` was to be able to produce phase maps from quantitative Energy Dispersive Spectroscopy (EDS) maps. This application provides a challenging test for the model because the dwell times used for mapping are typically very short (e.g., 1-20 ms per pixel), resulting in larger analytical errors than the minute-length analyses typically used for the point analyses in GEOROC and the secondary standard compilation. To evaluate the success of `mineralML` on noisier mapping data, we co-collected Electron Backscatter Diffraction (EBSD) and Energy Dispersive Spectroscopy (EDS) maps, utilizing the fact that EBSD can independently identify the phases based on their crystallography, using diffraction patterns. We selected a thin sec-

tion of an andesite from Mount Hood, Oregon and a granodiorite from the Porphyritic Half Dome unit of the Tuolumne Intrusive Complex, California (Sample B from Gordon and Wallis (2026); Gordon (2024)), to cover a range of both major rock forming minerals (feldspars, pyroxenes, amphiboles) and accessory minerals (epidote, zircon, apatite, titanite, oxide minerals).

Data were collected using a JEOL IT800HL SEM at the University of California, Berkeley, equipped with an Oxford Instruments Symmetry S3 detector and UltimMax EDS detector. Thin sections were tilted to 70° on the eucentric stage to enable EBSD acquisition. While this geometry is less optimal for EDS, comparison of EDS data collected on the sample when flat and when tilted showed no significant difference in major element composition once EDS data were normalized to 100%. Diffraction patterns were collected and indexed using Oxford Instruments AZtec 6.2 acquisition software to obtain crystallographically-determined phase maps. EBSD patterns are indexed against match units chosen by the user, so knowledge of the sample's mineralogy in advance of EBSD acquisition is highly beneficial. For our samples, mineralogy was determined using transmitted and reflected light microscopy, and from previous studies (Gordon, 2024). Where a good match unit could not be identified during acquisition, patterns were saved and re-indexed in AZtec in post-processing. The match units used for the final maps are available in the EBSD .ctf files (Zenodo DOI will be available upon article acceptance), which can be opened in any text file viewing software. The corresponding EDS data were processed with the QuantMap feature in AZtec using factory calibrations to get elemental weight or oxide percentage data, then classified with `mineralML` to evaluate against EBSD labels. The concurrent EBSD-EDS mapping tests allow for direct pixel-to-pixel validation of the model's phase classification capability on natural compositions and textures. EBSD has a higher spatial resolution than EDS due to its smaller interaction volume at our acquisition conditions, therefore the mapping also allows us to test the performance of `mineralML` at grain boundaries where mixed pixels occur.

In the AZtec acquisition software, EBSD-EDS maps are generally collected using the pixel dwell time selected to optimize the EBSD pattern, which can mean that maps are collected too quickly to obtain high-quality EDS data. The rate at which EBSD detectors can index patterns has experienced a rapid increase since the 1990s, with a notable jump in 2017 due to the introduction of complementary metal oxide semiconductor (CMOS) sensors to replace traditional charge coupled device sensor technology (CCDs) (Oxford Instruments Technical Note, 2026). These analytical advances mean that high-quality EBSD data can now be collected extremely rapidly. At our initial analytical conditions of a working distance of 25 mm, voltage of 20 kV, and beam current of 21 nA, a solvable diffraction pattern was acquired in 2.56 ms (0.00256 seconds). At these conditions, the EDS detector had a 34% deadtime using process time 3, so was only counting X-rays for 1.69 ms. This means that the total number of X-ray signals that arrived at the detector was very low, and thus the counting statistics yield large uncertainties. For example, a clinopyroxene with $\sim$ 51 wt.% SiO_2 would have a 1 sigma absolute error of $\pm$ 9.4 wt.% SiO_2 (an 18% relative error); this means that 68% of pixels (mean$\pm$1 sigma) acquired on a homogeneous clinopyroxene will yield SiO_2 contents spanning 41.8-60.5 wt.%. In contrast, the $\pm$1 sigma range in the calibration dataset for clinopyroxene in `mineralML` is only 48.1-52.2 wt.% SiO_2. EDS maps with short dwell times therefore display poor data quality and speckling (Supporting Information S6, Supporting Information Figure S10).

Analytical errors associated with X-ray counting techniques scale as the square root of the number of counts. As counts are proportional to acquisition time, analytical error scales as $t^{-0.5}$ (Wittke, 2022). To demonstrate this, we collected EDS spots at a variety of total acquisition times (live time/(1-dead time)) on a clinopyroxene grain using a tilted stage, at 25 mm working distance and beam conditions of 20 kV and 21.2 nA (Figure 5A-C). EDS maps were collected under the same conditions, but with pixel ac-

quisition times varying from 1-50 ms. It is apparent from the error curves (Figure 5A-C) that the errors reach 25-55% for 1 ms dwell times, but reduce to <10% for 50 ms. Maps collected with a 1 ms dwell time show lots of checkerboarding in the phase maps calculated by `mineralML`, with pixels of glass appearing in the pyroxene and plagioclase crystals (Figure 5D); the prediction scores are also generally low with a mean of 0.82 (Figure 5G). By 10 ms, the map is greatly improved, but glass pixels still appear in the plagioclase (Figure 5E), and the prediction scores increase to a mean of 0.93 (Figure 5H). By 50 ms, the checkerboarding is almost entirely gone, and the only pixels classified as glass are 'mixels,' or mixed pixels, on the grain boundaries (Figure 5F), and the prediction scores have also increased to 0.97 (Figure 5I). Pixels of clinopyroxene, particularly at phase boundaries, retain their lower prediction scores. This test shows the importance of minimizing errors while EDS mapping to ensure accurate phase identification. We recommend that users perform similar tests for their acquisition conditions of choice, as the errors will be dependent on the beam conditions, the detector hardware, and the elemental concentration in each phase.

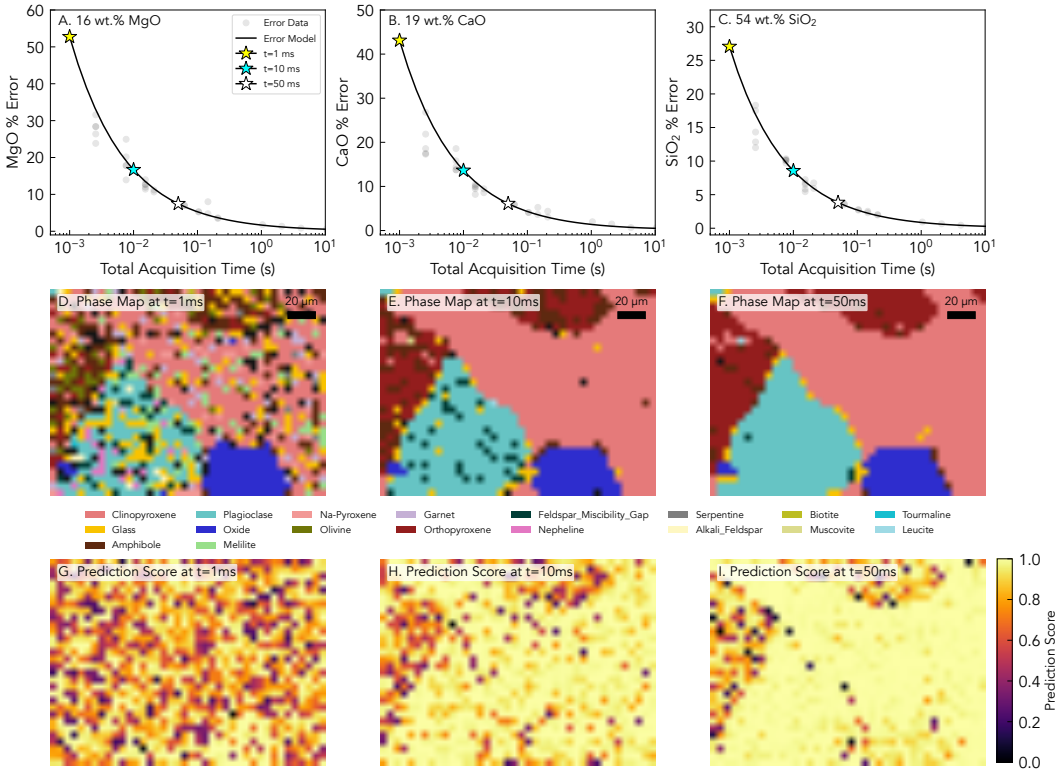

Figure 5. EDS analytical errors and `mineralML` classifications as a function of acquisition time. A-C) EDS acquisitions collected at different total acquisition times (gray dots) fit with a model to predict the error as a function of time. For each element, we calculate the error for t=1 ms, 10 ms and 50 ms (colored stars). D-F) EDS data collected at these acquisition times were run through `mineralML` to produce phase maps, with no prediction score minimum or phase filtering. G-I) Prediction scores associated with `mineralML` classifications. At shorter acquisition times, many different mineral labels are returned, but with low prediction scores of < 0.8. These solutions are improbable and not present in this sample, but reflects noise associated with the short count times. At greater acquisition times, there is a huge improvement in indexing and an increase in prediction scores.

To obtain high-quality co-collected EBSD-EDS mapped data to test the accuracy of `mineralML`'s phase classification, we shortened the working distance, and set the beam current to optimize the EDS detector deadtime. We used a current of 6.9 nA and working distances of 12 mm for the Mount Hood andesite and 15 mm for the Tuolumne granodiorite. At these conditions, 7.65 ms (Mount Hood) and 9.31 ms (Tuolumne) were required to obtain a good EBSD pattern. We set the EBSD frame averaging to 20, yielding total acquisition times of 153 ms and 186 ms respectively. EBSD data were processed in AZtecCrystal using the 'remove wild spikes' function to replace individual incorrectly indexed pixels that did not match their neighbours. Unindexed pixels with at least six matching neighbours were also replaced, and any phases that consisted only of misindexed pixels were removed. `mineralML`-predicted EDS phase maps show high mineralogical fidelity across samples, reproducing the mineralogy provided by EBSD and adding compositional information from EDS to the phase maps (Figures 6, 7).

For the Mount Hood andesite, `mineralML` yields modal proportions of major phases in excellent agreement with those determined by EBSD (Figure 6). Of the major phases, clinopyroxene (23.7% EDS and 22.6% by EBSD) and orthopyroxene (23.7% EDS and 26.3% by EBSD) agree within 3%. Of the accessory phases, oxides (2.3% by EDS and 2.4% by EBSD) and apatite (0.19% by EDS and 0.22% by EBSD) agree within 0.1%. The primary differences between the datasets occur in the classification of feldspars and non-crystalline materials. `mineralML` classifies approximately 3.7% of the map area as falling within the feldspar miscibility gap, which is classified empirically on the high-temperature feldspar ternary from Deer et al. (1963). This classification is likely inaccurate and occurs due to phases mixing within the EDS interaction volume, as it arises primarily at grain boundaries, or in the groundmass where there is a fine-grained intergrowth of plagioclase, alkali feldspar, and glass. EBSD handles these fine-grained regions better due to its smaller interaction volume, and the fact it is extremely good at differentiating between crystalline and non-crystalline groundmass. However, in some cases there is minor mis-indexing between structurally similar feldspars. Additionally, EBSD cannot differentiate between different non-crystalline materials such as glass, epoxy and vesicles, all of which are assigned 'unindexed' due to the lack of a diffraction pattern. `mineralML` is able to differentiate between these materials due to their different X-ray signatures. The combined phase proportion of glass and vesicles from EDS sum to within 3% of the unindexed phase proportion in EBSD. The EDS data also provide information on major element zoning within minerals, which EBSD cannot detect. The plagioclase crystals have irregular zoning with more anorthitic rims, and the clinopyroxene crystals have higher Mg# towards the rims. The in-situ compositional data provided by the EDS map provides a rich source of petrological information and could be used to guide the placement of line profiles, collected at longer dwell times, for diffusion chronometry.

The Tuolumne granodiorite also shows excellent agreement for phase abundances for major and accessory phases. Of the major phases, plagioclase (67.5% by EDS and 65.8% by EBSD), SiO_2 polymorphs (8.7% by EDS and 9.3% by EBSD), titanite (8.1% by EDS and 8.0% by EBSD), and amphibole (7.0% by EDS and EBSD) phase abundances agree within 2%. Accessory phase abundances all agree within 0.5%. This concurrently collected map highlights some issues that emerge with both methods, particularly in the region surrounding the zircon crystal (marked in cyan) on the lower left area of the map, where some pixels failed to index appropriately. Some areas of the EBSD map, particularly of apatite and alkali feldspar, failed to index due to low-quality patterns; these low-quality patterns may emerge due to sample topography, polishing imperfections, or alteration. The diffraction patterns from one Fe-Ti oxide (magnetite) crystal failed to index as such due to orientation, since insufficient high-quality zone axes were present in the patterns; thus, the pixels identified as oxide minerals are more sparse. `mineralML`-classified EDS pixels in this area identify all the Fe-Ti oxide pixels. The large patch of epidote (marked in magenta) failed to index at first pass due to a poor match with the epidote match unit originally selected for indexing. The diffraction patterns were there-

fore saved, and `mineralML` was used to confirm the mineralogy, before an alternative epidote match unit was used to re-index the saved diffraction patterns successfully. This granodiorite sample is fully crystalline, and the pixels identified as glass by `mineralML` mostly arise due to the mixed signals that occur at grain boundaries. The process of generating these maps demonstrates the utility of the paired approach in providing information that other methods lack. `mineralML` may serve as a useful tool for phase identification to facilitate the selection of appropriate match units for indexing EBSD maps.

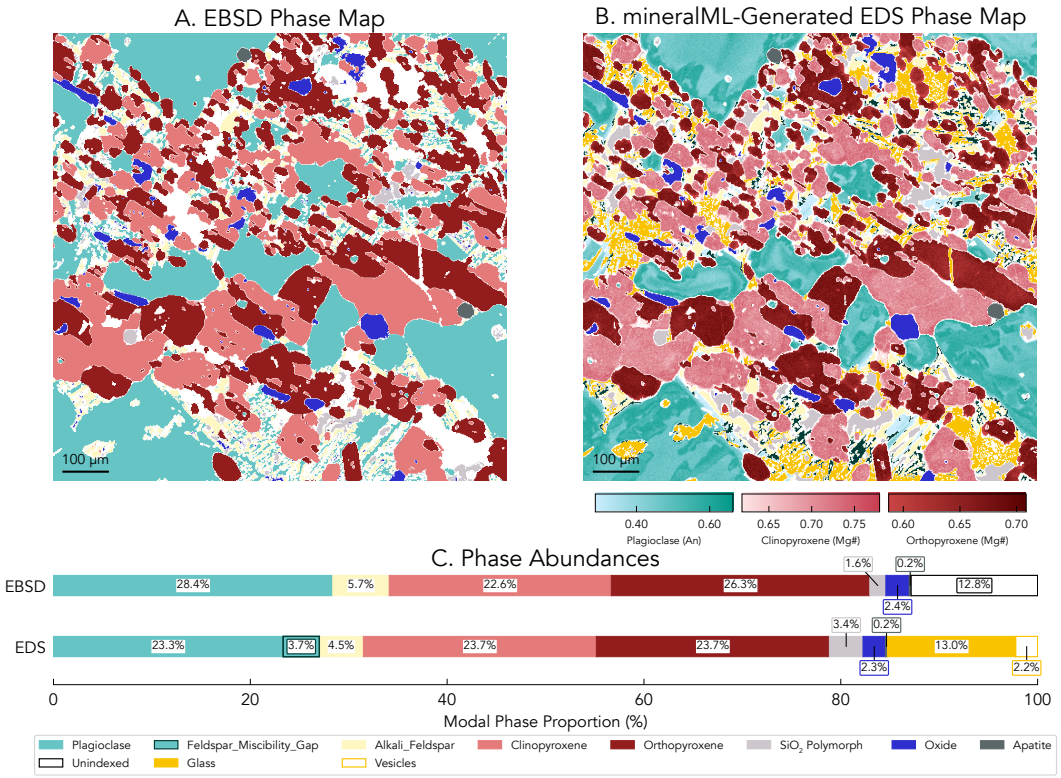

Figure 6. Comparison of EBSD and `mineralML`-generated EDS phase maps for an andesite from Mount Hood, Oregon, collected at a pixel size of 2 µm. A) Electron Backscatter Diffraction (EBSD) provides crystallographically-determined mineral phase labels. As EBSD relies on crystal diffraction, it cannot index non-crystalline phases such as glass or vesicles and leaves these pixels unindexed. B) Concurrently-collected EDS phase map with phase compositions predicted by `mineralML`. EDS mapping correctly identifies the presence of glass in the interstices. Pixels indicated as the feldspar miscibility gap are where there is fine-grained groundmass plagioclase, alkali feldspar, and glass, and the pixels pick up a mixed signal. Major element zoning in plagioclase and the two pyroxenes can be observed, with higher anorthite rims in plagioclase and some higher Mg# rims in the clinopyroxenes. Compositional ranges indicated by the colorbars are centered upon the mean with $\pm$ 2 σ. C) Comparison of modal phase abundances (%) from EBSD and `mineralML`-classified EDS maps, with strong agreement within 2-3% for all minerals.

4 On the Calibration of Prediction Scores

Neural networks are known to produce overconfident predictions, where prediction scores exceed actual classification accuracy (Guo et al., 2017). Ideally, network prediction scores should reflect the true likelihood of a correct classification. Models assign-

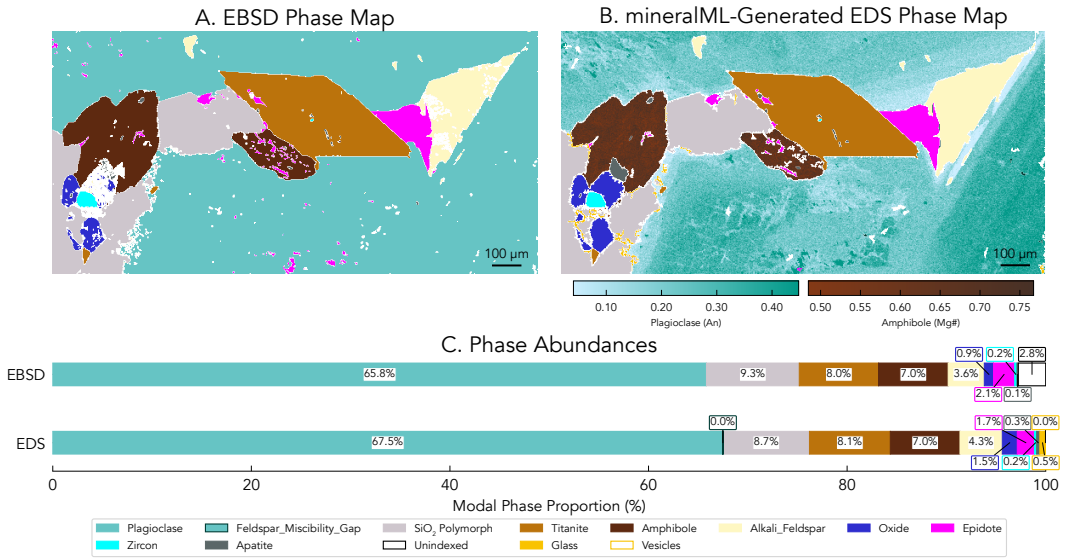

Figure 7. Comparison of EBSD and `mineralML`-generated EDS phase maps for sample Bii, a porphyritic granodiorite from the Tuolumne Intrusive Complex, California, collected at a pixel size of 3 µm. Panel layout, data collection, and colorbar scaling are identical to that of Figure 6. A) EBSD-derived phase map. The lower left region surrounding the zircon has spinels with some pixels that see poor indexing due to their orientation. The epidote did not index well with the original choice of match unit, so `mineralML` was applied to the EDS map to confirm its mineralogy and inform the subsequent selection of a successful match unit. B) Concurrently-collected EDS phase map. Zoning in plagioclase feldspar and amphibole can be observed. C) Comparison of modal phase abundances (%) from EBSD and `mineralML`-classified EDS maps. All phase abundances agree within 2% in this map. `mineralML` identifies additional Fe-Ti oxides that are missed by EBSD, and the EBSD indexing identifies small regions of epidote that are missed by EDS within the amphibole crystal. The two methods work in tandem to provide a unified view of mineralogy and phase abundances.

ing a probability score of 0.90 should be correct 90% of the time. Well-calibrated prediction scores are thus important for ensuring the interpretability and reliability of neural network outputs, particularly as these scores are used downstream. To evaluate whether `mineralML` suffers from this overconfidence, we assessed the calibration of the model's prediction scores using reliability diagrams (DeGroot & Fienberg, 1983; Niculescu-Mizil & Caruana, 2005) and the Expected Calibration Error metric following the workflow outlined by Guo et al. (2017).

We evaluated the calibration of the neural network on the GEOROC test dataset of $\sim$ 1.1 million analyses (Figure 8). The analyses with misclassified minerals ($\sim$ 22,000) identified in Section 3.2 were removed from this analysis. The removal of these analyses addresses some of the most readily identifiable mislabeled data; some residual misclassifications likely remain, but are unavoidable. Reliability diagrams are constructed by grouping `mineralML`-derived classifications into bins by their prediction score. The average accuracy within each bin is compared against the average score (Figure 8A). Perfect calibration would fall along the 1-1 diagonal. Deviations above the diagonal indicate that the accuracy exceeds the prediction score, so the model is underconfident; deviations below the diagonal indicate that the prediction score exceeds the accuracy, so

the model is overconfident. The Expected Calibration Error (ECE) metric summarizes the calibration as a scalar value:

$$\mathrm{ECE} = \sum_{m=1}^{M} \frac{|B_m|}{n} \left|\mathrm{acc}(B_m) - \mathrm{conf}(B_m)\right| \tag{10}$$

where B_m is the set of predictions in bin m, $|B_m|$ is the number of predictions in that bin, $\mathrm{acc}(B_m)$ is the fraction of correct predictions in the bin, $\mathrm{conf}(B_m)$ is the mean prediction score in that bin, M is the number of bins, and n is the total number of predictions. Lower ECE values indicate that there is no significant difference between the accuracy and the prediction score, so the model is well calibrated. The reliability diagram shows that `mineralML`'s prediction scores closely track the 1-1 diagonal across most confidence bins, yielding an ECE of 0.0152 (Figure 8A). Some mild underconfidence is observed in the low-to-mid confidence bins at 0.3-0.4 and 0.4-0.5. The low-confidence bins contain relatively few analyses and thus exhibit noisier calibration estimates. In the high-confidence bins exceeding 0.8, where most predictions are concentrated, the accuracy and prediction scores closely match.

We further examined accuracy as a function of $\sigma_{\bar{p}}$ to determine whether the prediction score uncertainty $\sigma_{\bar{p}}$ (standard deviation across predictions) from Monte Carlo sampling provides a meaningful signal for identifying unreliable classifications (Figure 8B). The prediction accuracy decreases monotonically as $\sigma_{\bar{p}}$ increases. Accuracy nears 100% at the lowest $\sigma_{\bar{p}}$ uncertainties around 0 and decreases to around 84% at $\sigma_{\bar{p}}$ uncertainties exceeding 0.2. The accuracy and standard deviation across Monte Carlo draws, which can be interpreted as the prediction score uncertainty, indicate that $\sigma_{\bar{p}}$ is an informative indicator of prediction reliability. Low values of prediction score uncertainty correspond to well-constrained compositions classified with high accuracy, while elevated values of prediction score uncertainty should be treated with more caution. The distribution of $\sigma_{\bar{p}}$ further supports this interpretation (Figure 8C). Correctly classified analyses are densely concentrated at low $\sigma_{\bar{p}}$, while misclassified analyses exhibit a distribution shifted towards higher values.

The low ECE scalar value and monotonic relationship between accuracy and prediction uncertainty together demonstrate that `mineralML`'s prediction scores are both well-calibrated and that the variational weight sampling provides a reliable measure of classification uncertainty. We reference the softmax-normalized outputs as prediction scores rather than probabilities, as softmax outputs from neural networks are not inherently calibrated (Guo et al., 2017). The calibration analysis presented here, however, indicates that `mineralML`'s prediction scores can be interpreted as meaningful estimates of classification reliability.

5 Applications of mineralML

`mineralML` is designed to serve as a first-pass processing step when working with chemical data from EPMA, EDS, or a geochemical database. The above tests show that `mineralML` can reliably predict mineralogy in a wide variety of igneous rock types. Once analyses are acquired and exported, `mineralML` workflows can provide immediate mineral identifications, identify any discrepancies between the mineralogy hypothesized by the user and the mineralogy predicted by the neural network, and perform all standard mineral component calculations within one Jupyter notebook. All compositions are returned in Excel sheets, with all input data in one tab, and data for each predicted mineral group also split in individual tabs. The workflow supported by this package substantially cuts down on the manual post-processing required after microprobe sessions, whilst reducing opportunities for misclassification errors to propagate into downstream thermobarometric calculations and into databases. Classification is sufficiently rapid to be

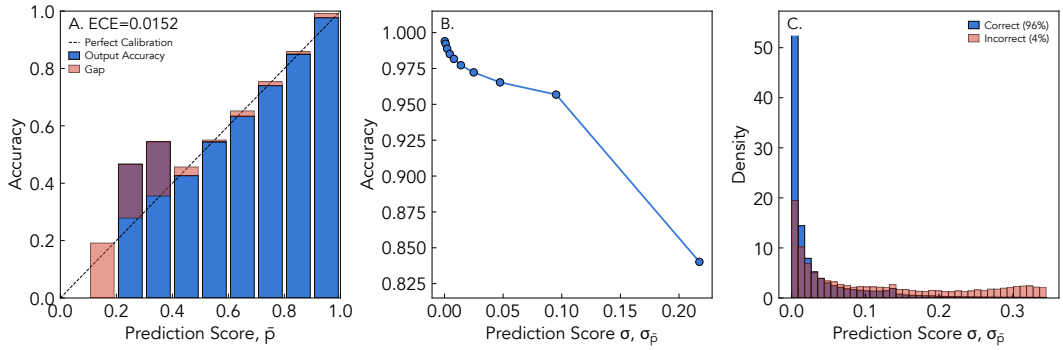

Figure 8. Calibration analysis of the `mineralML` model, evaluated on the GEOROC test dataset (∼1.1 million analyses). A) Reliability diagram comparing the prediction score ($\bar{p}$) and prediction accuracy across 10 equally spaced bins from 0 to 1. The dashed line represents the perfect calibration, where there is no offset between the prediction accuracy and score; the model is well-calibrated if it nears the dashed line. The blue bars indicate the accuracy of the model's predicted outputs, compared against the groundtruth labels. The red bars indicate the calibration gap, or the absolute difference between the prediction accuracy and prediction score for a given bin. Where the red bars extend beyond the blue bars, the model is overconfident (the prediction score exceeds the accuracy), and where the red bars extend below the blue bar, the model is underconfident (the accuracy exceeds the prediction score). The Expected Calibration Error (ECE) of 0.0152 indicates that `mineralML`'s model prediction scores closely approximate true classification accuracy. B) Classification accuracy as a function of $\sigma_{\bar{p}}$, or the standard deviation across predictions from Monte Carlo passes through the variational layer of the network. Analyses are binned by quantiles so each point represents approximately equal numbers of analyses. Accuracy decreases monotonically with increasing $\sigma_{\bar{p}}$, indicating that the variational weight sampling provides a meaningful indicator of prediction reliability. C) Normalized distributions of $\sigma_{\bar{p}}$ for correctly classified (blue) and incorrectly classified (red) analyses. Correct classifications are more densely concentrated at low $\sigma_{\bar{p}}$, while incorrect classifications are more broadly distributed and shifted towards higher $\sigma_{\bar{p}}$.

used in real-time during analytical sessions. Users can export spot analyses, linescans, or preliminary maps during a session, run `mineralML`, and receive mineral identifications to guide subsequent analyses; this is particularly useful when working with rocks with complex mineralogy where phase identifications guide subsequent analytical decisions.

`mineralML` offers a reproducible alternative to the manual phase segmentation and simple chemistry filters or clustering models currently used to process EDS maps. Such processing is time consuming, subjective, and presents a bottleneck in high-throughput workflows. Having a reliable, open-source model to classify mineralogy will allow the fields of volcanology and igneous petrology to fully capitalize on the ability of modern EDS detectors to produce chemical maps at high speed with high spatial resolution. We have developed Jupyter notebooks optimized to load quantitative EDS mapped data exported as .csv files for each element (the standard output from Oxford Instruments), and predict pixel-by-pixel mineralogy using Monte Carlo forward passes through the neural network. The `mineralML.run_map` function loads the path to the .csv files and returns a phase map colored by predicted mineral, a bar chart of mineral phase proportions, and histograms of prediction scores that allow users to evaluate prediction scores at a glance. For phases with well-defined solid solutions such as plagioclase, olivine, clinopyroxene,

and orthopyroxene, `mineralML` automatically computes stoichiometric components (anorthite content, forsterite content, Mg#) from the underlying oxide chemistry and maps these values into arrays that highlight internal crystal zoning, at the resolution of the EDS map. The `mineralML.plot_component_composite` function renders these compositional fields with continuous colors, which are mean centered and with the range of colors set to $\pm$ 2 σ. With both these functions, individual pixels below a specified area threshold can be removed with `remove_islands` and holes within maps enclosed within a single continuous phase can be filled with `fill_phase_holes`. In the case of the compositional map, the filled holes in maps are interpolated from surrounding points. All these compositional data, component calculations, and predicted labels are returned in structured dictionaries and dataframes for further analysis and export. In addition to these EDS workflows, `mineralML` also provides functions to load and visualize EBSD phase maps from .ctf files (`mineralML.plot_ctf_phases`), enabling direct comparison between diffraction- and chemistry-based mineral identification on the same sample — these functions were used for the side-by-side phase diagrams presented in Figures 6 and 7. We welcome collaborations with users of other detectors to develop workflows for other instrument data structures.

Our testing on the GEOROC database highlights the power of `mineralML` to improve the quality of large geochemical databases. Rather than having to look through every single study to check for errors, a difference in the mineralogy reported for the analyses and that allocated by `mineralML` offers an efficient way to identify studies requiring human inspection. These discrepancies can highlight inaccurate phase reporting (both by the authors and during data curation), as well as swapped and missing columns. Improving the quality of large geochemical databases will be absolutely key to the application of big data approaches to geochemical data. We welcome contributions from users who wish to extend `mineralML` to broader ranges of mineral compositions and new phases.

6 Limitations of mineralML

While `mineralML` performs well over a range of igneous compositions, some limitations should be considered, especially when applying it to new datasets that lie outside the calibration range of the model. New data should first be examined in Harker plots (`mineralML.plot_harker`) or in latent space (`mineralML.plot_z2_overlay`), to ensure that compositions lie within regions that have been calibrated. The prediction scores ($\bar{p}$) and their associated standard deviation ($\sigma_{\bar{p}}$) across Monte Carlo draws can also help provide a quantitative flag for misclassifications, as described in Sections 2.5 and 4.

`mineralML` is fundamentally dependent on the quality of the input oxide data. Whilst this is not an issue for most EPMA data or spot EDS analyses with acquisition times $>$ 1 s, it is relevant to EDS maps with short dwell times, including co-collected EBSD-EDS maps. Short pixel dwell times can produce analytical uncertainties that exceed the compositional variability between mineral groups (Figure 5A-C; Supporting Information Figure S10). The classification results are not failures of the model, but instead demonstrate the minimum data quality required for reliable phase identification. The probabilistic framework of `mineralML` communicates this limitation, as low prediction scores flag these classifications.

A related challenge for mapped data is the issue of mixed pixels, or 'mixels,' in which a single pixel analyses samples more than one mineral phase. 'Mixels' are particularly common in rocks with fine-grained groundmass, where pixel analyses may integrate signals from multiple minerals at grain boundaries where the electron beam's interaction volume straddles two phases. In both cases, the resulting analysis returns a composition that is a mixture of the phases. The composition will therefore not align entirely with a single mineral. `mineralML` will either assign a classification with low prediction scores,

or classify the pixel as the mineral whose composition most closely resembles the mixture. These 'mixels' are an artifact of the spatial resolution of the analytical technique. Users should be aware that phase maps of fine-grained or intergrown samples may contain misclassifications within the groundmass, and that individual pixels along grain boundaries may be misclassified.

As `mineralML` classifies minerals solely from chemistry, it cannot distinguish between polymorphs that share the same composition but differ in crystal structure. SiO_2 polymorphs (quartz, coesite, stishovite, tridymite, cristobalite) are grouped under a single label, as are other structurally distinct but chemically identical phases. Users requiring polymorph identification should pair `mineralML` classifications with structural data from techniques such as EBSD or Raman spectroscopy.

One notable limitation is the ability of `mineralML` to distinguish certain oxide phases belonging to the hematite-ilmenite ('Rhombohedral-Oxides') and magnetite-spinel ('Spinel-Group') mineral groups. For studies presenting data for magnetite-ilmenite pairs in intermediate to felsic igneous rocks, `mineralML` correctly subdivides the data into the appropriate mineral groups (Rutherford & Devine, 2008; Pratama et al., 2023; Ji et al., 2023; Blundy et al., 2008; Ghiorso & Evans, 2008), a useful feature for researchers wishing to use these analyses to perform ilmenite-magnetite thermometry and oxybarometry calculations. `mineralML` also successfully identifies Cr-spinel analyses from chromitites as being in the 'Spinel-Group' (Chayka et al., 2023). Where `mineralML` struggles, however, is for oxide phases from iron ores and/or metamorphosed/altered intrusive igneous rocks (Ovalle et al., 2018; Khedr et al., 2024; Mokchah & Mathieu, 2022). In these samples, we find many analyses that the original authors identify as magnetite are classified as 'Rhombohedral-Oxides' (hematite) by mineralML. Closer examination reveals that near-pure magnetite crystals (with low TiO_2), common in iron ore bodies but absent from the predominantly igneous `mineralML` training dataset, are being misclassified as hematite. This misclassification is due to an absence of pure Fe_3O_4 data in the training dataset. As a result, the use of the oxide subclassifications in ore systems is currently not advised.

The training dataset for `mineralML` is currently comprised of minerals from igneous rocks, and the model has primarily been developed and evaluated in igneous contexts. Many minerals classified by `mineralML`—such as olivine, pyroxene, feldspar, amphibole, oxides, and garnet—are common to both igneous and metamorphic assemblages. The model is expected to perform well on certain metamorphic rocks with igneous protoliths, specifically those containing mineral phases whose compositions overlap with the igneous training dataset. However, the model does not yet include the full compositional range of minerals found in metamorphic rocks, particularly in metapelites. Garnet compositions in the training dataset, for example, are dominated by igneous garnets and do not currently span the full range of metamorphic compositions. Several metamorphic minerals, such as kyanite, sillimanite, staurolite, lawsonite, and cordierite, are currently absent from the classification scheme. Extending `mineralML` to both metamorphic and ore systems is a natural direction for future development, requiring the compilation of training datasets with appropriate compositional coverage and new mineral classes.

7 Conclusions

`mineralML` is a novel method for probabilistic mineral classification from oxide chemistry, written in the open-source `Python` language using `PyTorch` (Paszke et al., 2019). We use a sequential, transfer learning approach to build a neural network with a classification and reconstruction branch, providing mineral classifications with uncertainty estimates while visualizing these data in two-dimensional latent spaces. The package further provides stoichiometric calculators for structural formulas, site assignments, and subclassifications across common igneous mineral groups. `mineralML` is designed to be incorporated into the data analysis workflow, following data export from electron beam

instruments. Automating mineral classification will relieve the pressure on data processing as the rate-limiting step in the petrologist's workflow.

Application of `mineralML` to validation datasets and the GEOROC database demonstrates greater than 95% accuracy across most mineral groups. By testing on more than 1.1 million analyses from GEOROC, we find that most misclassifications arise due to database errors such as incorrectly identified mineralogy, swapped columns in supplementary materials, and beyond, rather than true model failure. `mineralML` has the potential to serve as a powerful tool for curating large geochemical databases, identifying errors that may otherwise cascade into thermobarometric models or understandings of magmatic systems. When applied to co-collected EBSD-EDS maps, classified EDS phase maps using `mineralML` reproduce crystallographically-generated EBSD phase maps within 2-3% and provide the additional benefit of quantifying chemical variability within phases. `mineralML`'s development demonstrates how testing the model on large, heterogeneous datasets drives iterative model improvement. Compositional gaps in the training data—for example, Fe-rich amphiboles, fayalitic olivines, Na-rich pyroxenes, and Ca- and Mn-rich garnets—were only identified through evaluation against GEOROC. Addressing these gaps by expanding the training dataset improved classification performance.

Automating mineral classification has the potential to relieve the pressure on data processing as the rate-limiting step in a petrologist's workflow, especially as analytical techniques advance with ever-faster detectors and ever-growing datasets. `mineralML` additionally has the potential to serve as a tool for curating large geochemical databases, identifying errors that otherwise may cascade into our understandings and interpretations of magmatic systems. We anticipate this iterative process continuing as users apply `mineralML` to new rock types and mineral assemblages that are currently unrepresented in the training data. The modular nature of the package can support this growth. New minerals can easily be incorporated into the classification scheme, stoichiometric calculators can be extended and refined, and the training dataset can continue to expand. All code, training data, and Jupyter notebooks for analysis are publicly available on GitHub. We welcome collaboration with users who encounter edge cases or wish to extend `mineralML` to new compositions, to evolve this code package with the needs of the petrology community.

Contributions

SS, PW, KL conceptualized the project. SS developed and documented the code, wrote the manuscript, interpreted the results, and compiled the training, validation, and test datasets. PW compiled the secondary standard test dataset and provided guidance and support. PW and SS performed the testing on the GEOROC dataset. CG collected the EBSD-EDS datasets and processed the EBSD data. NT consulted on machine learning approaches and methods. MG performed model testing on oxides and developed the quantitative EDS mapping routines. PA and KL supported database access and provided advisory input. All authors provided feedback on the results and manuscript and tested the code.

Acknowledgments

We thank Haoyu Wang for discussions on machine learning and on neural network architectures, and Adam Kent for providing the thin section from Mount Hood. We greatly thank Carrie Soderman and Yishen (Eason) Zhang for discussing spinel and oxide calculations, along with the Droop normalization. SS was funded by a University of California, Berkeley Chancellor's Fellowship, and supported by NASA through award #NASA-80NSSC19K1102. PW, SS, and CG were funded by a University of California, Berkeley Heising-Simons Faculty Fellowship. EDS analyses utilized a detector obtained with

NSF Instrumentation and Facilities Grant 2218698 to PW and MG. PA was supported by NSF through grant EAR-2514733.

Conflict of Interest

The authors declare there are no conflicts of interest for this manuscript.

Data Availability

All data and code are available on GitHub (`https://github.com/sarahshi/mineralML`) and are archived on Zenodo (DOI:10.5281/zenodo.18819451). The training and validation data will be archived on EarthChem upon manuscript acceptance. `mineralML` documentation and examples are on Read The Docs (`https://mineralml.readthedocs.io/en/latest/`). The supplement contains information about the code structure and datasets.

References

Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)*, *44*(2), 139–160. doi: 10.1111/j.2517-6161.1982.tb01195.x

Allaz, J. M., Martin, L., Storck, J.-C., Ulmer, P., Reusser, E., Popa, R.-G., ... Bachmann, O. (2026, January). Quantitative analyses by combined energy and wavelength dispersive spectrometry: The best of both worlds. *American Mineralogist*. Retrieved from `http://dx.doi.org/10.2138/am-2025-10014` doi: 10.2138/am-2025-10014

Barber, N. D., Edmonds, M., Jenner, F., Audétat, A., & Williams, H. (2021). Amphibole control on copper systematics in arcs: Insights from the analysis of global datasets. *Geochimica et Cosmochimica Acta*, *307*, 192–211. doi: 10.1016/j.gca.2021.05.034

Blatter, D. L., Sisson, T. W., & Hankins, W. B. (2013). Crystallization of oxidized, moderately hydrous arc basalt at mid-to lower-crustal pressures: implications for andesite genesis. *Contributions to Mineralogy and Petrology*, *166*(3), 861–886. doi: 10.1007/s00410-013-0920-3

Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015, 07–09 Jul). Weight uncertainty in neural network. In F. Bach & D. Blei (Eds.), *Proceedings of the 32nd international conference on machine learning* (Vol. 37, pp. 1613–1622). Lille, France: PMLR. Retrieved from `https://proceedings.mlr.press/v37/blundell15.html`

Blundy, J., & Cashman, K. V. (2008). Petrologic reconstruction of magmatic system variables and processes. *Reviews in Mineralogy and Geochemistry*, *69*(1), 173-239. doi: 10.2138/rmg.2008.69.6

Blundy, J., Cashman, K. V., Berlo, K., Sherrod, D., Scott, W., & Stauffer, P. (2008). Evolving magma storage conditions beneath mount st. helens inferred from chemical variations in melt inclusions from the 1980–1986 and current (2004–2006) eruptions. *US Geological Survey Professional Paper*, *1750*, 755–790. doi: https://doi.org/10.3133/pp175033

Boschetty, F. O., Ferguson, D. J., Cortés, J. A., Morgado, E., Ebmeier, S. K., Morgan, D. J., ... Silva Parejas, C. (2022, April). Insights into magma storage beneath a frequently erupting arc volcano (villarrica, chile) from unsupervised machine learning analysis of mineral compositions. *Geochemistry, Geophysics, Geosystems*, *23*(4). Retrieved from `http://dx.doi.org/10.1029/2022GC010333` doi: 10.1029/2022gc010333

Boudoire, G., Pasdeloup, G., Schiavi, F., Cluzel, N., Rafflin, V., Grassa, F., ... others (2023). Magma storage and degassing beneath the youngest volcanoes of the massif central (france): Lessons for the monitoring of a dor-

mant volcanic province. *Chemical geology*, *634*, 121603. doi: doi:10.1016/j.chemgeo.2023.121603

Bowman, S., Vilnis, L., Vinyals, O., Dai, A., Jozefowicz, R., & Bengio, S. (2016). Generating sentences from a continuous space. In *Proceedings of the 20th signll conference on computational natural language learning* (pp. 10–21). doi: https://doi.org/10.48550/arXiv.1511.06349

Brandt, S., Fassbender, M., Klemd, R., Macauley, C., Felfer, P., & Haase, K. (2021, December). Cumulate olivine: A novel host for heavy rare earth element mineralization. *Geology*, *49*(4), 457–462. Retrieved from http://dx.doi.org/10.1130/G48417.1 doi: 10.1130/g48417.1

Braunger, S., Marks, M. A. W., Walter, B. F., Neubauer, R., Reich, R., Wenzel, T., ... Markl, G. (2018, August). The petrology of the kaiserstuhl volcanic complex, sw germany: The importance of metasomatized and oxidized lithospheric mantle for carbonatite generation. *Journal of Petrology*, *59*(9), 1731–1762. Retrieved from http://dx.doi.org/10.1093/petrology/egy078 doi: 10.1093/petrology/egy078

Brill, R. H. (1971). A chemical-analytical round-robin on four synthetic ancient glasses. In *Int. congr. glass, artistic hist. commun, 9th* (pp. 93–110).

Callegaro, S., Svensen, H. H., Neumann, E. R., Polozov, A., Jerram, D. A., Deegan, F., ... Melnikov, N. (2021). Geochemistry of deep tunguska basin sills, siberian traps: correlations and potential implications for the end-permian environmental crisis. *Contributions to Mineralogy and Petrology*, *176*(7), 49. doi: doi:10.1007/s00410-021-01807-3

Carlson, R. W., & Ionov, D. A. (2019). Compositional characteristics of the morb mantle and bulk silicate earth based on spinel peridotites from the tariat region, mongolia. *Geochimica et Cosmochimica Acta*, *257*, 206–223. doi: doi:10.1016/j.gca.2019.05.010

Cashman, K. V., & Blundy, J. (2013). Petrological cannibalism: the chemical and textural consequences of incremental magma body growth. *Contributions to Mineralogy and Petrology*, *166*(3), 703-729. doi: 10.1007/s00410-013-0895-0

Castro, A., Rodríguez, C., Gutiérrez-Alonso, G., & de la Rosa, J. D. (2023). A post-collisional batholith from southern iberia rooted in the earth's mantle: Los pedroches batholith. *Lithos*, *454*, 107245. doi: https://doi.org/10.1016/j.lithos.2023.107245

Chayka, I. F., Baykov, N. I., Kamenetsky, V. S., Kutyrev, A. V., Pushkarev, E. V., Abersteiner, A., & Shcherbakov, V. D. (2023, January). Volcano–plutonic complex of the tumrok range (eastern kamchatka): An example of the ural-alaskan type intrusion and related volcanic series. *Minerals*, *13*(1), 126. Retrieved from http://dx.doi.org/10.3390/min13010126 doi: 10.3390/min13010126

Cheng, L., Costa, F., & Bergantz, G. (2020). Linking fluid dynamics and olivine crystal scale zoning during simulated magma intrusion. *Contributions to Mineralogy and Petrology*, *175*(6), 1–14. doi: 10.1007/s00410-020-01691-3

Crawshaw, M. (2020). Multi-task learning with deep neural networks: A survey. Retrieved from https://arxiv.org/abs/2009.09796 doi: 10.48550/ARXIV.2009.09796

Deer, W., Howie, R., & Zussman, J. (1963). *Rock-forming minerals: Framework silicates* (1st ed.). Longmans. Retrieved from https://books.google.com/books?id=0IkiwQEACAAJ

Deer, W., Howie, R., & Zussman, J. (2001). *Rock-forming minerals. framework silicates: Feldspars, volume 4A* (2nd ed.). Bath, England: Geological Society.

Deer, W., Howie, R., & Zussman, J. (2013). *An introduction to the rock-forming minerals*. Mineralogical Society of Great Britain and Ireland. Retrieved from http://dx.doi.org/10.1180/DHZ doi: 10.1180/dhz

DeGroot, M. H., & Fienberg, S. E. (1983, March). The comparison and evaluation of forecasters. *The Statistician*, *32*(1/2), 12. Retrieved from http://dx.doi

.org/10.2307/2987588 doi: 10.2307/2987588

Didonna, R., Handley, H., Albert, H., & Costa, F. (2024). Time scales of olivine storage and transport as revealed by diffusion chronometry at waitomokia volcanic complex, auckland volcanic field, new zealand. *Journal of Volcanology and Geothermal Research*, *451*, 108094. doi: 10.1016/j.jvolgeores.2024.108094

DIGIS Team. (2024). *Georoc compilation: Minerals.* Goettingen Research Online / Data. Retrieved from https://doi.org/10.25625/SGFTFN (Export generated 1 December 2024) doi: 10.25625/SGFTFN

DIGIS Team. (2025). *Georoc compilation: Melt inclusions.* Goettingen Research Online / Data. Retrieved from https://doi.org/10.25625/7JW6XU (Export generated 1 December 2025) doi: 10.25625/7JW6XU

Droop, G. (1987). A general equation for estimating fe3+ concentrations in ferromagnesian silicates and oxides from microprobe analyses, using stoichiometric criteria. *Mineralogical magazine*, *51*(361), 431–435. doi: https://doi.org/10.1180/minmag.1987.051.361.10

Egorova, V., & Latypov, R. (2013). Mafic–ultramafic sills: New insights from m-and s-shaped mineral and whole-rock compositional profiles. *Journal of Petrology*, *54*(10), 2155–2191. doi: https://doi.org/10.1093/petrology/egt045

García, M. T., Calderón, M., de Arellano, C. R., Hervé, F., Opitz, J., Theye, T., ... others (2020). Trace element composition of amphibole and petrogenesis of hornblendites and plutonic suites of cretaceous magmatic arcs developed in the fuegian andes, southernmost south america. *Lithos*, *372*, 105656. doi: 10.1016/j.lithos.2020.105656

Ghiorso, M. S., & Evans, B. W. (2008). Thermodynamics of rhombohedral oxide solid solutions and a revision of the fe-ti two-oxide geothermometer and oxygen-barometer. *American Journal of science*, *308*(9), 957–1039. doi: doi.org/10.2475/09.2008.01

Ghiorso, M. S., & Sack, O. (1991). Fe-Ti oxide geothermometry: thermodynamic formulation and the estimation of intensive variables in silicic magmas. *Contributions to Mineralogy and Petrology*, *108*(4), 485–510. doi: https://doi.org/10.1007/BF00303452

Giuffrida, M., Viccaro, M., & Ottolini, L. (2018). Ultrafast syn-eruptive degassing and ascent trigger high-energy basic eruptions. *Scientific Reports*, *8*(1), 1-7. doi: 10.1038/s41598-017-18580-8

Gleeson, M. L., Gibson, S. A., & Stock, M. J. (2020, October). Upper mantle mush zones beneath low melt flux ocean island volcanoes: Insights from isla floreana, galápagos. *Journal of Petrology*, *61*(11–12). Retrieved from http://dx.doi.org/10.1093/petrology/egaa094 doi: 10.1093/petrology/egaa094

Gleeson, M. L., Wieser, P. E., DeVitre, C. L., Shi, S. C., Millet, M.-A., Muir, D. D., ... Lissenberg, J. (2025). Persistent high-pressure magma storage beneath a near-ridge ocean island volcano (isla floreana, galápagos). *Journal of Petrology*, *66*(5), egaf031. doi: doi.org/10.1093/petrology/egaf031

Gordon, C. (2024). *Insights into magmatic processes from the crystal orientation mapping of igneous microstructures* (Doctoral dissertation). doi: 10.17863/CAM.114340

Gordon, C., & Wallis, D. (2026, March). Polymineralic synneusis in silicic magmas revealed by megacryst-inclusion orientation relationships. *EarthArXiv*. Retrieved from http://dx.doi.org/10.31223/X5HZ0Q doi: 10.31223/x5hz0q

Gottlieb, P., Wilkie, G., Sutherland, D., Ho-Tun, E., Suthers, S., Perera, K., ... Rayner, J. (2000). Using Quantitative Electron Microscopy for Process Mineralogy Applications. *Journal of The Minerals, Metals, and Materials Society*, *52*, 24-25. doi: 10.1007/s11837-000-0126-9

Graham, S., Brough, C., & Cropp, A. (2015). An introduction to ZEISS mineralogic mining and the correlation of light microscopy with automated mineralogy: A case study using BMS and PGM analysis of samples from a PGE-bearing

chromite prospect. *Proceedings of the Precious Metals '15, Falmouth, UK*, *11*.

Greenacre, M. (2018). *Compositional data analysis in practice*. Chapman and Hall/CRC. Retrieved from http://dx.doi.org/10.1201/9780429455537 doi: 10.1201/9780429455537

Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. Retrieved from https://arxiv.org/abs/1706.04599 doi: 10.48550/ARXIV.1706.04599

Harper, M., Weinstein, B., Woodcock, T., Simon, C., Morgan, W., Knight, V., ... Zuidhof, G. (2015). python-ternary: Ternary plots in python. *Zenodo 10.5281/zenodo.594435*. Retrieved from https://github.com/marcharper/python-ternary doi: 10.5281/zenodo.594435

Hidalgo, P. J., Vogel, T. A., Rooney, T. O., Currier, R. M., & Layer, P. W. (2011). Origin of silicic volcanism in the panamanian arc: evidence for a two-stage fractionation process at el valle volcano. *Contributions to Mineralogy and Petrology*, *162*(6), 1115–1138.

Hirschmann, M., Ghiorso, M., Davis, F., Gordon, S., Mukherjee, S., Grove, T., ... Till, C. (2008). Library of experimental phase relations (lepr): A database and web portal for experimental magmatic phase equilibria data. *Geochemistry, Geophysics, Geosystems*, *9*(3). doi: 10.1029/2007GC001894

Horn, E. L., Taylor, R. N., Gernon, T. M., Stock, M. J., & Farley, E. R. (2022). Composition and petrology of a mush-bearing magma reservoir beneath tenerife. *Journal of Petrology*, *63*(10), egac095. doi: https://doi.org/10.1093/petrology/egac095

Iddon, F., Jackson, C., Hutchison, W., Fontijn, K., Pyle, D. M., Mather, T. A., ... Edmonds, M. (2019, January). Mixing and crystal scavenging in the main ethiopian rift revealed by trace element systematics in feldspars and glasses. *Geochemistry, Geophysics, Geosystems*, *20*(1), 230–259. Retrieved from http://dx.doi.org/10.1029/2018GC007836 doi: 10.1029/2018gc007836

Ionov, D. A., Doucet, L. S., von Strandmann, P. A. P., Golovin, A. V., & Korsakov, A. V. (2017). Links between deformation, chemical enrichments and li-isotope compositions in the lithospheric mantle of the central siberian craton. *Chemical Geology*, *475*, 105–121. doi: https://doi.org/10.1016/j.chemgeo.2017.10.038

Jarosewich, E. (2002). Smithsonian microbeam standards. *Journal of research of the National Institute of Standards and Technology*, *107*(6), 681. doi: 10.6028/jres.107.054

Ji, G.-Y., Jiang, S.-H., Liu, Y.-F., & Zhang, L.-S. (2023). Petrogenesis of the alubaogeshan granite in the maodeng mo-bi-sn-cu deposit, southern great xing'an range, ne china: Constraints from apatite, plagioclase, magnetite and ilmenite geochemistry. *Ore Geology Reviews*, *155*, 105355. Retrieved from https://www.sciencedirect.com/science/article/pii/S0169136823000707 doi: https://doi.org/10.1016/j.oregeorev.2023.105355

Jones, R., Hiscock, M., Jiang, H., & Stavropoulou, A. (2020). Automated mineralogy and quantitative compositional analysis of geological samples on a multi-functional scanning electron microscope. *Microscopy and Microanalysis*, *26*(S2), 2858–2859.

Keller, C. B., & Schoene, B. (2012). Statistical geochemistry reveals disruption in secular lithospheric evolution about 2.5 gyr ago. *Nature*, *485*(7399), 490–493. doi: 10.1038/nature11024

Keller, F., Wanke, M., Kueter, N., Guillong, M., & Bachmann, O. (2024, January). An amphibole perspective on the recent magmatic evolution of mount st. helens. *Journal of Petrology*, *65*(1). Retrieved from http://dx.doi.org/10.1093/petrology/egad093 doi: 10.1093/petrology/egad093

Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., & Tang, P. T. P. (2016). On large-batch training for deep learning: Generalization gap and

sharp minima. Retrieved from https://arxiv.org/abs/1609.04836 doi: 10.48550/ARXIV.1609.04836

Khedr, M. Z., Moftah, A., El-Shibiny, N. H., Tamura, A., Tan, W., Ichiyama, Y., ... Abdelrahman, K. (2024, June). Mineralogy and geochemistry of titaniferous iron ores in el-baroud layered gabbros: Fe-ti ore genesis and tectono-metallogenetic setting. *Minerals*, *14*(7), 679. Retrieved from http://dx.doi.org/10.3390/min14070679 doi: 10.3390/min14070679

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. doi: 10.48550/arXiv.1412.6980

Klöcking, M., Wyborn, L., Lehnert, K. A., Ware, B., Prent, A. M., Profeta, L., ... others (2023). Community recommendations for geochemical data, services and analytical capabilities in the 21st century. *Geochimica et Cosmochimica Acta*. doi: 10.1016/j.gca.2023.04.024

Lanari, P., & Engi, M. (2017, February). Local bulk composition effects on metamorphic mineral assemblages. *Reviews in Mineralogy and Geochemistry*, *83*(1), 55–102. Retrieved from http://dx.doi.org/10.2138/rmg.2017.83.3 doi: 10.2138/rmg.2017.83.3

Lanari, P., & Tedeschi, M. (2025, February). Chemical map classification in xmaptools. *Applied Computing and Geosciences*, *25*, 100230. Retrieved from http://dx.doi.org/10.1016/j.acags.2025.100230 doi: 10.1016/j.acags.2025.100230

Larocque, J., & Canil, D. (2010). The role of amphibole in the evolution of arc magmas and crust: the case from the jurassic bonanza arc section, vancouver island, canada. *Contributions to Mineralogy and Petrology*, *159*(4), 475–492. doi: 10.1007/s00410-009-0436-z

Leake, B. E., Woolley, A. R., Arps, C. E., Birch, W. D., Gilbert, M. C., Grice, J. D., ... others (1997). Nomenclature of amphiboles; report of the subcommittee on amphiboles of the international mineralogical association, commission on new minerals and mineral names. *The Canadian Mineralogist*, *35*(1), 219–246. doi: doi.org/10.1180/minmag.1997.061.405.13

Le Bas, M., Le Maitre, R., & Woolley, A. (1992). The construction of the total alkali-silica chemical classification of volcanic rocks. *Mineralogy and petrology*, *46*(1), 1–22. doi: https://doi.org/10.1007/BF01160698

Lehnert, K., Su, Y., Langmuir, C., Sarbas, B., & Nohl, U. (2000). A global geochemical database structure for rocks. *Geochemistry, Geophysics, Geosystems*, *1*(5). doi: 10.1029/1999GC000026

Lindsley, D. (1983). Pyroxene thermometry. *American Mineralogist*, *68*, 477-493.

Lloyd, S. (1982). Least squares quantization in PCM. *IEEE transactions on information theory*, *28*(2), 129–137. doi: 10.1109/TIT.1982.1056489

Loewen, M. W., Dietterich, H. R., Graham, N., & Izbekov, P. (2021, October). Evolution in eruptive style of the 2018 eruption of veniaminof volcano, alaska, reflected in groundmass textures and remote sensing. *Bulletin of Volcanology*, *83*(11). Retrieved from http://dx.doi.org/10.1007/s00445-021-01489-6 doi: 10.1007/s00445-021-01489-6

Loewen, M. W., Wallace, K., Lubbers, J., Ruth, D., Izbekov, P., Larsen, J., & Graham, N. (2023). *Glass electron microprobe analyses methods, precision and accuracy for tephra studies in alaska.* Retrieved from http://dx.doi.org/10.14509/31045 doi: 10.14509/31045

Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*. doi: https://doi.org/10.48550/arXiv.1711.05101

Lui, T. C., Prabhu, A., Bidgood, A., Morrison, S. M., & Caers, J. (2026, February). Geochemnet: An interactive tool for visualizing and interpreting outliers in geochemical data using networks. *Applied Geochemistry*, *199*, 106712. Retrieved from http://dx.doi.org/10.1016/j.apgeochem.2026.106712 doi: 10.1016/j.apgeochem.2026.106712

Maas, A. L., Hannun, A. Y., Ng, A. Y., et al. (2013). Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml* (Vol. 30, p. 3).

McCarthy, A., & Müntener, O. (2017). Mineral growth in melt conduits as a mechanism for igneous layering in shallow arc plutons: mineral chemistry of fisher lake orbicules and comb layers (sierra nevada, usa). *Contributions to Mineralogy and Petrology*, *172*(7), 55.

Melluso, L., Scarpati, C., Zanetti, A., Sparice, D., & de' Gennaro, R. (2022, November). The petrogenesis of chemically zoned, phonolitic, plinian and sub-plinian eruptions of somma-vesuvius, italy: Role of accessory phase removal, independently filled magma reservoirs with time, and transition from slightly to highly silica undersaturated magmatic series in an ultrapotassic stratovolcano. *Lithos*, *430–431*, 106854. Retrieved from `http://dx.doi.org/10.1016/j.lithos.2022.106854` doi: 10.1016/j.lithos.2022.106854

Mokchah, N., & Mathieu, L. (2022, February). Origin and evolution of the iron-rich upper unit and fe–ti–v mineralization of the neoarchean lac doré layered intrusion, chibougamau, québec. *Journal of Petrology*, *63*(3). Retrieved from `http://dx.doi.org/10.1093/petrology/egac006` doi: 10.1093/petrology/egac006

Morimoto, N. (1988). Nomenclature of pyroxenes. *Mineralogy and Petrology*, *39*(1), 55–76. doi: https://doi.org/10.1007/BF01226262

Mutch, E. J., Maclennan, J., Holland, T. J., & Buisman, I. (2019). Millennial storage of near-Moho magma. *Science*, *365*(6450), 260-264. doi: 10.1126/science.aax4092

Mutch, E. J., Maclennan, J., & Madden-Nadeau, A. L. (2022). The dichotomous nature of Mg-in-plagioclase partitioning: Implications for diffusion chronometry. *Elsevier*. doi: 10.1016/j.gca.2022.10.035

Neave, D. A., Buisman, I., & Maclennan, J. (2017). Continuous mush disaggregation during the long-lasting Laki fissure eruption, Iceland. *American Mineralogist*, *102*, 2007-2021. doi: 10.2138/am-2017-6015CCBY

Neave, D. A., & Maclennan, J. (2020). Clinopyroxene dissolution records rapid magma ascent. *Frontiers in Earth Science*, *8*, 188. doi: 10.3389/feart.2020.00188

Neave, D. A., Maclennan, J., Hartley, M. E., Edmonds, M., & Thordarson, T. (2014). Crystal storage and transfer in basaltic systems: The Skuggafjöll eruption, Iceland. *Journal of Petrology*, *55*(12), 2311–2346. doi: 10.1093/petrology/egu058

Neave, D. A., Passmore, E., Maclennan, J., Fitton, G., & Thordarson, T. (2013). Crystal-Melt Relationships and the Record of Deep Mixing and Crystallization in the AD 1783 Laki Eruption, Iceland. *Journal of Petrology*, *54*(8). doi: 10.1093/petrology/egt027

Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on machine learning - icml '05* (pp. 625–632). ACM Press. Retrieved from `http://dx.doi.org/10.1145/1102351.1102430` doi: 10.1145/1102351.1102430

Ondrejka, M., Li, X.-H., Vojtko, R., Putiš, M., Uher, P., & Sobockỳ, T. (2018). Permian a-type rhyolites of the muráň nappe, inner western carpathians, slovakia: in-situ zircon u–pb sims ages and tectonic setting. *Geologica Carpathica*, *69*(2), 187–198. doi: 10.1515/geoca-2018-0011

Ovalle, J. T., La Cruz, N. L., Reich, M., Barra, F., Simon, A. C., Konecke, B. A., ... Morata, D. (2018, October). Formation of massive iron deposits linked to explosive volcanic eruptions. *Scientific Reports*, *8*(1). Retrieved from `http://dx.doi.org/10.1038/s41598-018-33206-3` doi: 10.1038/s41598-018-33206-3

Oxford Instruments Technical Note. (2026). *Ebsd detectors.* `https://www.ebsd.com/ebsd-techniques/ebsd-detectors`. (Accessed: 2026-03-03)

Park, B. J., & Kim, H. S. (2022). P–t–x co 2–bulk rock composition modeling of garnet decomposition in amphibolite and mafic granulite: tectonometamorphic insights into the permian–triassic orogeny on the eastern margin of the korean peninsula. *Contributions to Mineralogy and Petrology*, *177*(9), 89. doi: 10.1007/s00410-022-01952-3

Parman, S. W., Shimizu, N., Grove, T. L., & Dann, J. C. (2003). Constraints on the pre-metamorphic trace element composition of barberton komatiites from ion probe analyses of preserved clinopyroxene. *Contributions to Mineralogy and Petrology*, *144*(4), 383–396. doi: 10.1007/s00410-002-0406-1

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., ... others (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, *32*.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... others (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, *12*, 2825–2830. doi: 10.48550/arXiv.1201.0490

Pirrie, D., Butcher, A. R., Power, M. R., Gottlieb, P., & Miller, G. L. (2004). Rapid quantitative mineral and phase analysis using automated scanning electron microscopy (QEMSCAN); Potential applications in forensic geoscience. *Geological Society, London, Special Publications*, *232*(1), 123-136. doi: 10.1144/GSL.SP.2004.232.01.

Pitcher, B. W., & Kent, A. J. (2019). Statistics and segmentation: using big data to assess Cascades arc compositional variability. *Geochimica et cosmochimica acta*, *265*, 443–467. doi: 10.1016/j.gca.2019.08.035

Polat, A., Frei, R., Longstaffe, F. J., Thorkelson, D. J., & Friedman, E. (2018). Petrology and geochemistry of the tasse mantle xenoliths of the canadian cordillera: a record of archean to quaternary mantle growth, metasomatism, removal, and melting. *Tectonophysics*, *737*, 1–26. doi: 10.1016/j.tecto.2018.04.014

Pontesilli, A., Brenna, M., Mollo, S., Masotta, M., Nazzari, M., Le Roux, P., & Scarlato, P. (2022). Trachyte-phonolite transition at dunedin volcano: Fingerprints of magma plumbing system maturity and mush evolution. *Lithos*, *408*, 106545. doi: https://doi.org/10.1016/j.lithos.2021.106545

Pratama, A., Nurfiani, D., Suryanata, P. B., Ismail, T., Bunga Naen, G. N. R., Abdurrachman, M., ... Setiawan, N. I. (2023, October). Magma storage conditions beneath krakatau, indonesia: insight from geochemistry and rock magnetism studies. *Frontiers in Earth Science*, *11*. Retrieved from `http://dx.doi.org/10.3389/feart.2023.1128798` doi: 10.3389/feart.2023.1128798

Putirka, K. D. (2008). Thermometers and Barometers for Volcanic Systems. *Reviews in Mineralogy and Geochemistry*, *69*(1), 61-120. doi: 10.2138/rmg.2008.69.3

Putirka, K. D. (2016). Amphibole thermometers and barometers for igneous systems and some implications for eruption mechanisms of felsic magmas at arc volcanoes. *American Mineralogist*, *101*(4), 841–858. doi: https://doi.org/10.2138/am-2016-5506

Rasmussen, D. J., Plank, T. A., Roman, D. C., & Zimmer, M. M. (2022). Magmatic water content controls the pre-eruptive depth of arc magmas. *Science*, *375*(6585), 1169–1172.

Repstock, A., Casas-García, R., Zeug, M., Breitkreuz, C., Schulz, B., Gevorgyan, H., ... Lapp, M. (2022). The monotonous intermediate magma system of the permian wurzen caldera, germany: Magma dynamics and petrogenetic constraints for a supereruption. *Journal of Volcanology and Geothermal Research*, *429*, 107596. doi: https://doi.org/10.1016/j.jvolgeores.2022.107596

Reubi, O., & Blundy, J. (2009). A dearth of intermediate melts at subduction zone volcanoes and the petrogenesis of arc andesites. *Nature*, *461*(7268), 1269–1273. doi: doi.org/10.1038/nature08510

Ridolfi, F. (2021). Amp-tb2: an updated model for calcic amphibole thermobarome-

try. *Minerals*, *11*(3), 324. doi: doi.org/10.3390/min11030324

Ridolfi, F., Renzulli, A., & Puerini, M. (2010). Stability and chemical equilibrium of amphibole in calc-alkaline magmas: an overview, new thermobarometric formulations and application to subduction-related volcanoes. *Contributions to mineralogy and petrology*, *160*(1), 45–66. doi: https://doi.org/10.1007/s00410-009-0465-7

Rutherford, M. J., & Devine, J. D. (2008). *Magmatic conditions and processes in the storage zone of the 2004-2006 mount st. helens dacite* (D. R. Sherrod, W. E. Scott, & P. H. Stauffer, Eds.). Retrieved from `http://dx.doi.org/10.3133/pp175031` doi: 10.3133/pp175031

Schoene, B., Condon, D. J., Morgan, L., & McLean, N. (2013). Precision and accuracy in geochronology. *Elements*, *9*(1), 19–24. doi: https://doi.org/10.2113/gselements.9.1.19

Shi, S., Towbin, W. H., Plank, T., Barth, A., Rasmussen, D., Moussallam, Y., ... Menke, W. (2024). Pyiroglass: An open-source, bayesian mcmc algorithm for fitting baselines to ftir spectra of basaltic-andesitic glasses. *Volcanica*, *7*(2), 471–501. doi: 10.30909/vol.07.02.471501

Siegel, K., Williams-Jones, A. E., & van Hinsberg, V. J. (2017). The amphiboles of the ree-rich a-type peralkaline strange lake pluton–fingerprints of magma evolution. *Lithos*, *288*, 156–174. doi: https://doi.org/10.1016/j.lithos.2017.07.012

Spear, F. S., Pattison, D. R., & Cheney, J. T. (2017). The metamorphosis of metamorphic petrology. In *The web of geological sciences: Advances, impacts, and interactions ii.* Geological Society of America. Retrieved from `http://dx.doi.org/10.1130/2016.2523(02)` doi: 10.1130/2016.2523(02)

Stall, S., Yarmey, L., Cutcher-Gershenfeld, J., Hanson, B., Lehnert, K., Nosek, B., ... Wyborn, L. (2019). Make scientific data FAIR. *Nature*, *570*(7759), 27–29. doi: 10.1038/d41586-019-01720-7

Stracke, A., Willig, M., Genske, F., Béguelin, P., & Todd, E. (2022). Chemical geodynamics insights from a machine learning approach. *Geochemistry, Geophysics, Geosystems*, *23*(10), e2022GC010606. doi: 10.1029/2022GC010606

Swallow, E. J., Wilson, C. J. N., Charlier, B. L. A., & Gamble, J. A. (2019, June). The huckleberry ridge tuff, yellowstone: evacuation of multiple magmatic systems in a complex episodic eruption. *Journal of Petrology*, *60*(7), 1371–1426. Retrieved from `http://dx.doi.org/10.1093/petrology/egz034` doi: 10.1093/petrology/egz034

Tadesse, A. Z., Fontijn, K., Caricchi, L., Begue, F., Gudbrandsson, S., Smith, V., ... others (2023). Pre-eruptive storage conditions and magmatic evolution of the bora-baricha-tullu moye volcanic system, main ethiopian rift. *Lithos*, *442*, 107088. doi: https://doi.org/10.1016/j.lithos.2023.107088

Tappe, S., Steenfelt, A., Heaman, L. M., & Simonetti, A. (2009). The newly discovered jurassic tikiusaaq carbonatite-aillikite occurrence, west greenland, and some remarks on carbonatite–kimberlite relationships. *Lithos*, *112*, 385–399. doi: 10.1016/j.lithos.2009.03.002

Taylor, R. J., Hill, E., & Andrew, M. (2024). A step forward in quantitative automated mineralogy in 2d and 3d. *Geostandards and Geoanalytical Research*, *48*(3), 579–593. doi: doi.org/10.1111/ggr.12552

Teiber, H., Scharrer, M., Marks, M. A., Arzamastsev, A. A., Wenzel, T., & Markl, G. (2015). Equilibrium partitioning and subsequent re-distribution of halogens among apatite–biotite–amphibole assemblages from mantle-derived plutonic rocks: Complexities revealed. *Lithos*, *220*, 221–237. doi: https://doi.org/10.1016/j.lithos.2015.02.015

Thivet, S., Gurioli, L., & Di Muro, A. (2020). Basaltic dyke eruptions at piton de la fournaise: characterization of the eruptive products with implications for reservoir conditions, conduit processes and eruptive dynamics. *Contributions*

to Mineralogy and Petrology, *175*(3), 26. doi: 10.1007/s00410-020-1664-5

Torres-Orozco, R., Cronin, S. J., Pardo, N., Kósik, S., Ukstins, I., Heinrich, M., & Lee, P. D. (2023). Complex decompression and fragmentation of mingled andesite magmas driving multi-phase plinian eruptions at mt. taranaki, new zealand. *Journal of Volcanology and Geothermal Research*, *433*, 107728. doi: 10.1016/j.jvolgeores.2022.107728

Toth, N., Shi, S., Maclennan, J., & Tung, P.-Y. (2025). Eds analysis for petrology: a probabilistic framework with gpyeds. *Journal of Geophysical Research: Machine Learning and Computation*. doi: doi.org/10.22541/au.174595460.03942520/v1

Tung, P.-Y., Sheikh, H. A., Ball, M., Nabiei, F., & Harrison, R. J. (2022). Sigma: Spectral interpretation using gaussian mixtures and autoencoder. *Geochemistry, Geophysics, Geosystems*, e2022GC010530. doi: 10.1029/2022GC010530

Ubide, T., Gale, C., Larrea, P., Arranz, E., Lago, M., & Tierz, P. (2014). The relevance of crystal transfer to magma mixing: a case study in composite dykes from the central pyrenees. *Journal of Petrology*, *55*(8), 1535–1559. doi: https://doi.org/10.1093/petrology/egu033

Viljoen, A., Howarth, G. H., Giuliani, A., Fitzpayne, A., & Costin, G. (2022). Correlations between olivine composition and groundmass mineralogy in sierra leone kimberlites provide constraints on craton-specific melt-lithosphere interactions. *Lithos*, *430*, 106846. doi: https://doi.org/10.1016/j.lithos.2022.106846

Voigt, A., Cassidy, M., Castro, J. M., Pyle, D. M., Mather, T. A., Helo, C., ... Kurniawan, I. A. (2022). Experimental investigation of trachydacite magma storage prior to the 1257 eruption of mt samalas. *Journal of Petrology*, *63*(8), egac066. doi: https://doi.org/10.1093/petrology/egac066

Walters, J. B. (2022). Minplot: A mineral formula recalculation and plotting program for electron probe microanalysis. *Mineralogia*, *53*(1), 51–66.

Wang, C., Shao, Y.-J., Cawood, P. A., Chen, J.-F., Xiong, Y.-Q., & Wang, Y.-J. (2023, June). Regional zoning of a li-cs-ta pegmatite field: Insights from monazite-cheralite chemistry, u-th-pb and sm-nd isotopes. *Journal of Petrology*, *64*(7). Retrieved from `http://dx.doi.org/10.1093/petrology/egad044` doi: 10.1093/petrology/egad044

Wang, X., Hou, T., Wang, M., Zhang, C., Zhang, Z., Pan, R., ... Zhang, H. (2021). A new clinopyroxene thermobarometer for mafic to intermediate magmatic systems. *European Journal of Mineralogy*, *33*(5), 621–637.

Waters, L. E., & Lange, R. A. (2015). An updated calibration of the plagioclase-liquid hygrometer-thermometer applicable to basalts through rhyolites. *American mineralogist*, *100*(10), 2172–2184. doi: https://doi.org/10.2138/am-2015-5232

Weller, O. M., St-Onge, M. R., Waters, D. J., Rayner, N., Searle, M. P., Chung, S., ... Xu, X. (2013, September). Quantifying barrovian metamorphism in the danba structural culmination of eastern tibet. *Journal of Metamorphic Geology*, *31*(9), 909–935. Retrieved from `http://dx.doi.org/10.1111/jmg.12050` doi: 10.1111/jmg.12050

Whitley, S., Halama, R., Gertisser, R., Preece, K., Deegan, F. M., & Troll, V. R. (2020, April). Magmatic and metasomatic effects of magma–carbonate interaction recorded in calc-silicate xenoliths from merapi volcano (indonesia). *Journal of Petrology*, *61*(4). Retrieved from `http://dx.doi.org/10.1093/petrology/egaa048` doi: 10.1093/petrology/egaa048

Whitney, D. L., & Evans, B. W. (2010). Abbreviations for names of rock-forming minerals. *American mineralogist*, *95*(1), 185–187. doi: 10.2138/am.2010.3371

Wieser, P. E., Gleeson, M. L., Rangel, B., DeVitre, C. L., Bearden, A. T., Lynn, K. J., ... Monteleone, B. (2025). Fluid inclusion constraints on the geometry of the magmatic plumbing system beneath mauna loa—part 2: Xenoliths. *Bulletin of Volcanology*, *87*(10), 1–24. doi: https://doi.org/10.1007/

s00445-025-01869-2

Wieser, P. E., Petrelli, M., Lubbers, J., Wieser, E., Ozaydin, S., Kent, A. J., & Till, C. (2022). Thermobar: An open-source Python3 tool for thermobarometry and hygrometry. *Volcanica*, *5*(2), 349-384. doi: 10.30909/vol.05.02.349384

Williams, M. J., Schoneveld, L., Mao, Y., Klump, J., Gosses, J., Dalton, H., ... Barnes, S. (2020). pyrolite: Python for geochemistry. *Journal of Open Source Software*, *5*(50), 2314. doi: https://doi.org/10.21105/joss.02314

Wittke, J. H. (2022). *Electron microprobe and scanning-electron analysis.* (Course lecture notes)

Yamasaki, T. (2022). Magma hybridization and crystallization in coexisting gabbroic and granitic bodies in the mid-crust, akechi district, central japan. *Mineralogy and Petrology*, *116*(3), 189–228. doi: https://doi.org/10.1007/s00710-022-00775-1

Zhu, W., Song, J., Wang, H., & Münchmeyer, J. (2025). Towards end-to-end earthquake monitoring using a multitask deep learning model. Retrieved from `https://arxiv.org/abs/2506.06939` doi: 10.48550/ARXIV.2506.06939

Supporting Information for mineralML: Leveraging Machine Learning for Probabilistic Mineral Classification

Sarah C. Shi [1,2], Penny Wieser [1], Charlotte Gordon [1], Norbert Toth [3], Paula Antoshechkina [4], Matthew L.M. Gleeson [1], Kerstin Lehnert [2]

[1]Department of Earth and Planetary Sciences, University of California, Berkeley. Berkeley, CA USA

[2]Lamont-Doherty Earth Observatory, Columbia University. New York, NY USA

[3]University of Cambridge. Cambridge, Cambridgeshire UK

[4]California Institute of Technology. Pasadena, CA USA

Contents of this file

1. Text S1 to S6
2. Figures S1 to S14
3. Tables S1

Additional Supporting Information (Files uploaded separately)

1. Dataset S1

Corresponding author: Sarah C. Shi, Department of Earth and Planetary Sciences, University of California, Berkeley, McCone Hall 355, Berkeley, CA 94720, USA. (sarahshi@berkeley.edu)

March 20, 2026, 9:00 am

Introduction

This supporting information provides additional details on the data, methods, and validation experiments described in the main text. Text S1 summarizes the datasets used for training, validation, and testing, and provides an overview of the GitHub/Zenodo repository structure. Text S2 describes the Compositional Data Analysis (CoDA) preprocessing tests and justifies the adoption of z-score normalization. Text S3 motivates the choice of neural network architecture over alternative machine learning techniques. Text S4 validates the synthetic mineral generator against natural mineral compositions. Text S5 documents the co-collected EBSD-EDS mapping conditions and evaluates the effect of acquisition time on classification accuracy. Text S6 provides captions for the supporting dataset. Figures S1–S14, Table S1, and Data Set S1 accompany these texts.

Data Set S1

The training dataset used for `mineralML` has been uploaded as a Supporting Information Dataset. All other datasets will be published on GitHub and Zenodo upon manuscript acceptance.

Text S1: Data Overview

Over a million analyses of minerals were used to train, validate, and test `mineralML`. The training/validation dataset is composed of ∼128,000 analyses. The training dataset is provided in the Supporting Information as a .csv. The test dataset is composed of > 1,100,000 analyses from GEOROC, a reference material dataset is composed of ∼ 3,000 analyses of Smithsonian reference materials, and co-collected mapped EBSD/EDS maps.

All these datasets will be published on GitHub and Zenodo upon manuscript acceptance.

Text S2: GitHub/Zenodo Overview

The directory forest presented provides an overview of the file structure contained within the GitHub and Zenodo.

```
mineralML
├─docs - Documentation for ReadtheDocs
├─src/mineralML - Source code
│ ├─mineral_classes_nn_v0030.npz - Numerical and string representations of minerals
│ ├─nnwr_best_model_v0030.pt - PyTorch file storing the best model for use at inference
│ ├─nnwr_latent_data_v0030.npz - Training latent space embedding for plotting overlay
│ ├─scaler_nn_v0030.npz - Standard scaler (mean and std) used for z-score normalization
│ ├─__init__.py - Initializes package for mineralML
│ ├─_version.py - States version for PyPI
│ ├─confusion_matrix.py - Code generating confusion matrix for visualization
│ ├─constants.py - Constant values for molar masses
│ ├─core.py - Core code for mineralML
│ ├─hybrid.py - Code containing neural network model for mineralML
│ ├─mapping.py - EDS map processing code for mineralML
│ ├─stoichiometry.py - Stoichiometric calculators built into mineralML
│ └─synthetic_minerals.py - Synthetic mineral generator code for mineralML
├─.readthedocs.yaml - ReadtheDocs configuration file
├─LICENSE.txt - GNU GPLv3
├─README.md - README file with key information on mineralML
├─environment.yml - YML file with Python environment information
├─setup.cfg - Setup information for mineralML
└─setup.py - Setup information for mineralML
```

Text S3: Data Preprocessing

As discussed in Section 2.4, we tested Compositional Data Analysis (CoDA) using log-ratio transformations as an alternative to z-score normalization. Compositional data are subject to the closure problem. Oxide weight percentages are constrained to sum to 100 wt.%, meaning that an increase in one component forces a decrease in another component. This introduces spurious correlations that may influence statistical analysis, recognized

by Pearson (1901) and Chayes (1960). Log-ratio transforms are often recommended in the CoDA framework to map closed compositions in simplex space to unconstrained real space where standard statistical methods can be applied (Aitchison, 1982; Greenacre, 2018).

Log-ratio transformations, however, require strictly positive values. In our dataset, missing values are pervasive. Many oxides are often not measured for certain mineral groups. Since zircon contains primarily ZrO_2 and SiO_2, the majority of analyses have no values reported for Na_2O, K_2O, and Cr_2O_3. Hematite is almost all Fe_2O_3, usually measured as FeO_t, and routinely has no values reported for Na_2O, K_2O, and P_2O_5. These missing values reflect the fact that these elements are not expected in the mineral structure, so are often not included in the analytical routine. The pattern of missing data is thus phase dependent, as different mineral groups have different subsets of oxides that are missing.

In the CoDA framework, these zeros must be replaced with small positive values prior to transformation. We tested imputation strategies with replacement of these missing values with small nonzero values, such as 10^{-6}-10^{-2} wt.%. In all cases, imputation introduced artificial structure into the transformed feature space with points plotted following the training dataset mineral label (Figure S1). Minerals that have natural scatter in z-scored oxide-oxide space (Figure S1A) are stretched into linear streaks in bivariate CLR-transformed space (Figure S1B). Minerals sharing the same pattern of missing values (and thus the same set of imputed values) are stretched along parallel lines, while minerals with different missing value patterns are oriented along different axes.

The effect continues into the PCA projections. In z-scored PCA space (Figure S1C), minerals form tight and distinct clusters, albeit with some overlap, with the first two

principal components capturing 25.3% and 17.2% of the variance respectively. In CLR-transformed PCA space (Figure S1C), minerals form spotty clusters that are displaced along different axes that overlap, with the first two principal components capturing 46.4% and 14.5% of the variance respectively. The principal components capture artificial variance introduced by the linear streaking. Mineral groups that form tight clusters in z-scored space are stretched into distributed clusters in CLR space. Minerals sharing the same zero pattern are smeared along parallel streaks even if their true compositions differ. Many of these streaks overlap in the central region of the plot, making it difficult to distinguish between mineral groups that are well-separated in z-scored space. The natural geochemical variation within each mineral group is compressed relative to the imputation-induced elongation, producing a feature space organized primarily by missing-data rather than by composition. The CLR-transformed space was used to train some neural network models, and resulted in poorer performance than with z-score normalization.

We therefore adopted the z-score normalization for all model training and inference. While z-scoring does not directly address the closure constraint, the neural network was able to learn the compositional relationships directly from the data. The classification accuracies reported in Section 3 validate this approach.

Text S4: Comparison with Alternative Machine Learning Techniques

The choice of a sequential, transfer-learning neural network over alternative machine learning techniques was motivated by three primary goals: 1) probabilistic classification with uncertainty estimates, 2) a latent representation for visualization, and 3) robust

performance across minerals with overlapping compositional boundaries. The architecture of the neural network is further discussed in Section 2.5.

Linear dimensionality reduction techniques in the Singular Value Decomposition (SVD) family, such as Principal Component Analysis (PCA), identify axes of maximum variance in the data. For mineral classification, these axes of greatest compositional variance do not necessarily correspond to the axes that best discriminate between mineral groups. PCA applied to z-scored oxide data (Figure S1C) produces a space in which some mineral groups are well-separated, but others with overlapping major element chemistry—such as clinopyroxene, amphibole, and glass—overlap substantially. These overlapping boundaries are non-linear in oxide space, limiting the utility of any classification or clustering method that operates on linear projections of the data. Autoencoders address this limitation by learning non-linear mappings from the input space to a lower-dimensional representation. Because the encoder and decoder are composed of non-linear activation functions, the resulting latent space can capture compositional structure that linear methods collapse. In `mineralML`, this advantage is compounded by the sequential architecture: rather than compressing raw oxide data (which, like PCA, would organize the latent space by bulk compositional variance), the autoencoder compresses a feature embedding already optimized for mineral classification, producing a two-dimensional space that reflects mineralogical distinctions (Section 2.5).

Text S5: Synthetic Minerals

As described in Section 2.3, synthetic mineral compositions were generated for mineral groups with insufficient analyses to achieve class balance in the training dataset (Table **??**). The synthetic mineral generator produces synthetic compositions by mixing idealized endmembers stoichiometries with stochastic perturbation. The method enforces charge balance.

Here we validate that the synthetic compositions reproduce the compositional distributions of the natural training data for well-understood solid solution minerals, olivine and plagioclase feldspar. We compare the synthetic and natural datasets in two ways: (1) bivariate oxide and component space to verify that synthetic compositions span the same compositional trends as the natural data, and (2) violin plots comparing the distributions of cation counts (per formula unit) and oxide weight percentages, evaluated with Kolmogorov-Smirnov (KS) statistics.

Synthetic olivine compositions span the full forsterite-fayalite solid solution (Figure S2). Synthetic olivines reproduce the tight FeO-MgO anticorrelation and the SiO_2-MgO trend of the natural data. The M-site occupancy additionally clusters around the expected value of 2.0 across the full range of XFo. The violin distributions confirm close agreement between natural and synthetic cation counts and oxide concentrations for all major elements (Figure S3). Synthetic data show slightly broader distributions, reflecting the broader compositions sampled with the generator.

Synthetic plagioclase compositions also span the anorthite-albite solution space (Figure S4). The synthetic data reproduce the Na_2O-CaO and Al_2O_3-SiO_2 substitution trends, and plot along the An-Ab tie line. The violin distributions show close agreement for Si, Al, Ca, Na, and K cation counts (Figure S5).

Kalsilite is a feldspathoid with a somewhat restricted compositional range. The synthetic compositions reproduce the Si, Al, Na, and K cation distributions of the natural data (Figure S6). Zircon is dominated by Si and Zr, with trace Fe and Mg. The synthetic compositions reproduce the near-stoichiometric Si and Zr cation counts, with appropriately small Fe and Mg impurities added for stochastic noise (Figure S7).

SiO_2-polymorphs are nearly pure SiO_2 with only trace impurities. The synthetic compositions reproduce this near-unity Si cation count with appropriately small concentrations of Ti and Al (Figure S8). Hematite is similarly nearly pure Fe_2O_3 with only trace impurities. The synthetic compositions reproduce the Fe^{3+} cation counts (Figure S9).

In all cases, the synthetically-generated mineral compositions closely match the natural data confirming the applicability of this method in augmenting the training dataset.

Text S6: EBSD/EDS Mapping

Co-collected EBSD-EDS maps were collected on a JEOL IT800HL SEM at the University of California, Berkeley, equipped with an Oxford Instruments Symmetry S3 EBSD detector and UltimMax EDS detector. In the AZtec acquisition software, the pixel dwell time is set to optimize the EBSD diffraction pattern. EDS X-ray counts are thus governed by the amount of time required by the EBSD detector. Modern complementary metal oxide semiconductor (CMOS) sensors can acquire indexable patterns in just a few milliseconds, far faster than the dwell times typically used for quantitative EDS mapping.

The initial Mount Hood andesite map was collected at a working distance of 25 mm, voltage of 20 kV, and beam current of 21 nA (Figure S10). At these conditions, a solv-

able diffraction pattern was acquired in 2.56 ms (0.00256 seconds). These conditions yielded excellent EBSD indexing, but prevented the EDS detector from accumulating sufficient X-ray counts per pixels, resulting in large uncertainties. The resulting EDS maps exhibited substantial noise, with analytical uncertainties on major oxides exceeding the compositional variability between mineral groups — at 2.56 ms acquisition time, the MgO % error is 27%, the CaO % error is 22%, and the SiO_2 % error is 13%. Agreement in phase proportions is very poor compared to the optimized map (Figure 6C) and further incorporates many spurious mineral phases, reflecting the impact of noisy EDS data on classification accuracy. This map demonstrates that EBSD-optimized conditions for data collection do not extend to EDS, and highlights the minimum data quality required for reliable phase identification.

To systematically evaluate the relationship between acquisition time and `mineralML` classifications, we collected maps on a limited region of the Mount Hood andesite at a range of acquisition times (live time/(1-dead time)) ranging from 1 ms to 50 ms, at 25 mm working distance and beam conditions of 20 kV and 21.2 nA. The resulting phase maps (Figure S11 show a progressive improvement with increasing dwell time. At 1-2 ms, the maps show many white pixels representing all other mineral classifications. Broad compositional ranges are analyzed within these minerals. By 10-20 ms, coherent phase boundaries emerge. At 50 ms, the map is well classified. Glass pixels are confined to grain boundaries. The colorbar ranges for plagioclase An content and pyroxene Mg# narrow substantially, converging on ranges consistent with those observed in the optimized full map (Figure 6B).

The prediction scores returned by `mineralML` highlight this same improvement in data quality (Figures S12, S13). At 1 ms, prediction scores are uniformly low and most prediction scores are below 0.8. The histogram of prediction scores is broad and has a hump at around 0.5. As acquisition times increase, the maps increasingly consist of high prediction score pixels and the histograms shift towards 1.0. At 50 ms, most pixel prediction scores exceed 0.9, with only a small tail of lower confidence pixels corresponding to grain boundaries.

Per-phase prediction score histograms (Figure S14) further illustrate how the classification evolves with data quality. At short dwell times, several phases that are not truly present in this sub-region, such as alkali feldspar and large proportions of glass, appear due to analytical noise. As acquisition times increase, the spurious phases proportions decrease. Glass modal proportions decrease from 12.3% of the map at 1 ms to 2.2% at 50 ms, and alkali feldspar effectively disappears by 10 ms. The proportions of dominant phases also converge on stable values, with clinopyroxene increasing from 28.7% at 1 ms to 50.5% at 50 ms. Noisy pixels previously misclassified are now assigned the correct label with increasing confidence.

To obtain the high-quality co-collected EBSD-EDS maps presented in the main text (Figures 6 and 7), we shortened the working distance and adjusted the beam current to optimize EDS collection while increasing the EBSD frame averaging to 20 frames per pixel to effectively increase the total acquisition time. The optimized conditions are summarized in Table S1. At total acquisition times of 153 ms (Mount Hood, MH0811) and 186 ms (Tuolumne, Bii), the EDS data quality is sufficient for reliable mineral classification and compositional mapping, and `mineralML`-classified phase maps reproduce the

crystallographically-determined EBSD maps within 2–3% for all major phases.

References

Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)*, *44*(2), 139–160. doi: 10.1111/j.2517-6161.1982.tb01195.x

Chayes, F. (1960). On correlation between variables of constant sum. *Journal of Geophysical research*, *65*(12), 4185–4193.

Greenacre, M. (2018). *Compositional data analysis in practice*. Chapman and Hall/CRC. Retrieved from `http://dx.doi.org/10.1201/9780429455537` doi: 10.1201/9780429455537

Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, *2*, 559-572. doi: 10.1080/14786440109462720

do not specify file extension

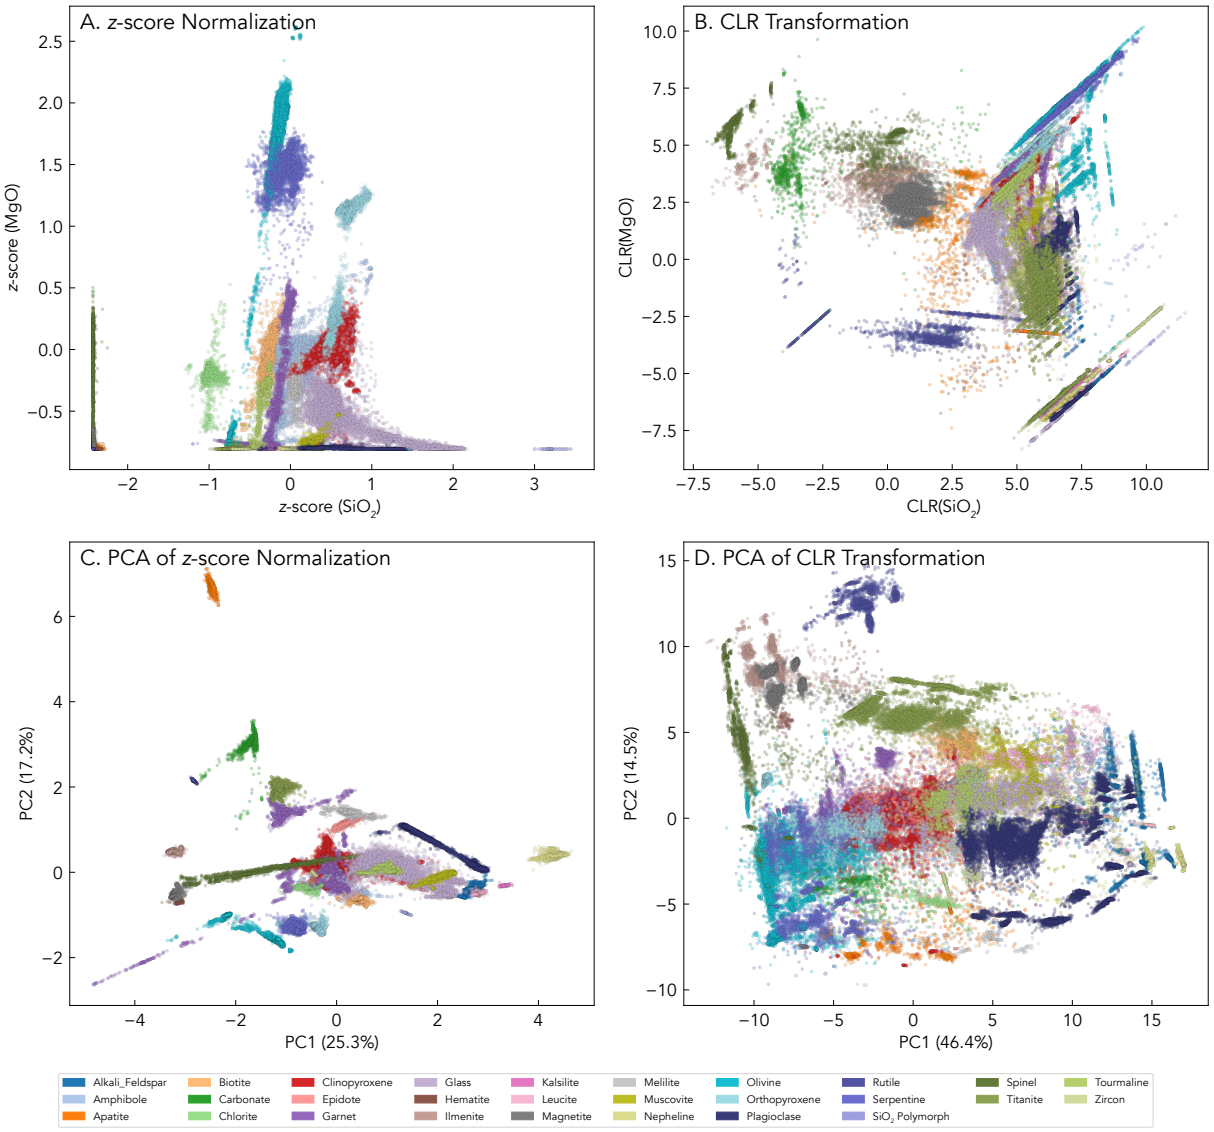

Figure S1. Comparison of feature space resulting from z-score normalization (left column) and centered log-ratio (CLR) transformation (right column) with imputation of missing values with 10^{-4} wt.%. The top row shows the principal component analysis (PCA) projections of the transformed feature space, and the bottom row shows the transformed feature space. A) The z-scored feature PCA space shows how different minerals do form distinct clusters, reflecting their composition. The mineral clusters radiate outwards from the 0, 0 region of principal component space. Many minerals, however, overlap in this 0, 0 region of principal component space, highlighting how simple dimensionality re-

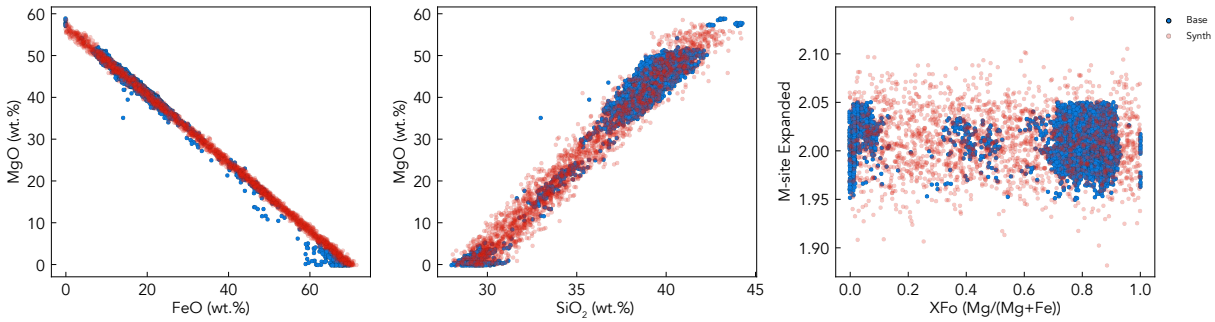

Figure S2. Comparison of natural (blue) and synthetic (red) olivine compositions. The left panel plots FeO and MgO, showing the continuous forsterite-fayalite solid solution. The center panel plots SiO_2 and MgO. The right panel plots M-site occupancy (atoms per formula unit) vs. forsterite content (XFo). Synthetic olivines span the same compositional range as the natural data and cluster around the expected M-site occupancy of 2.0.

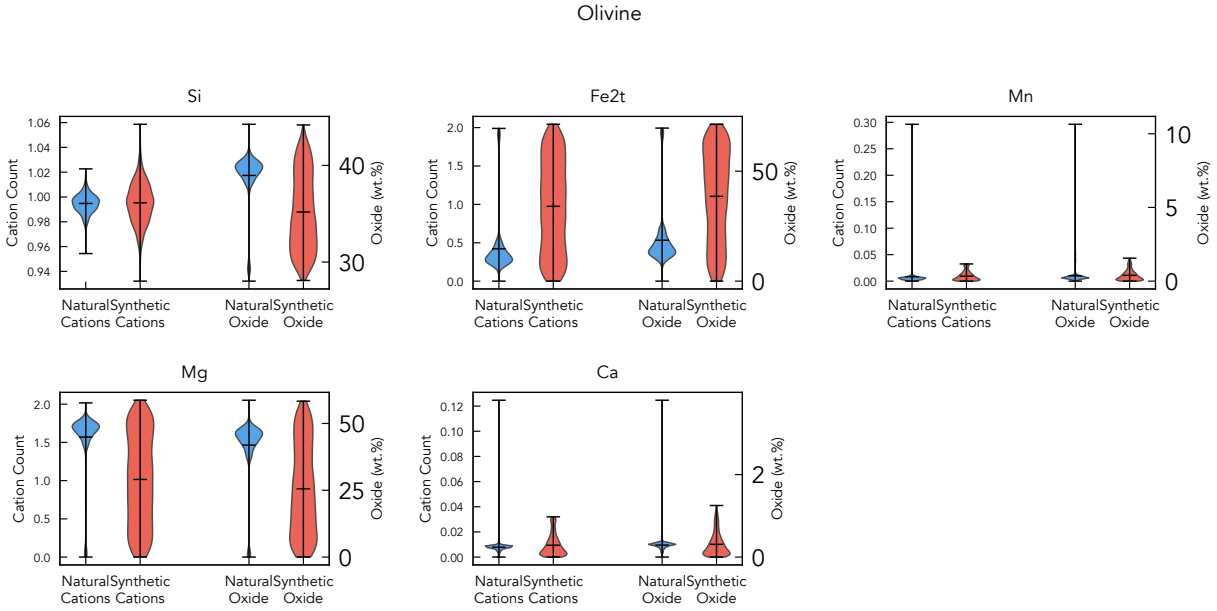

Figure S3. Violin plots comparing the distributions of cation counts (blue violins, left axis) and oxide weight percentages (red violins, right axis) for natural and synthetic olivine. Black horizontal lines indicate the mean of each distribution. Natural and synthetic distributions are in close agreement for all major cations.

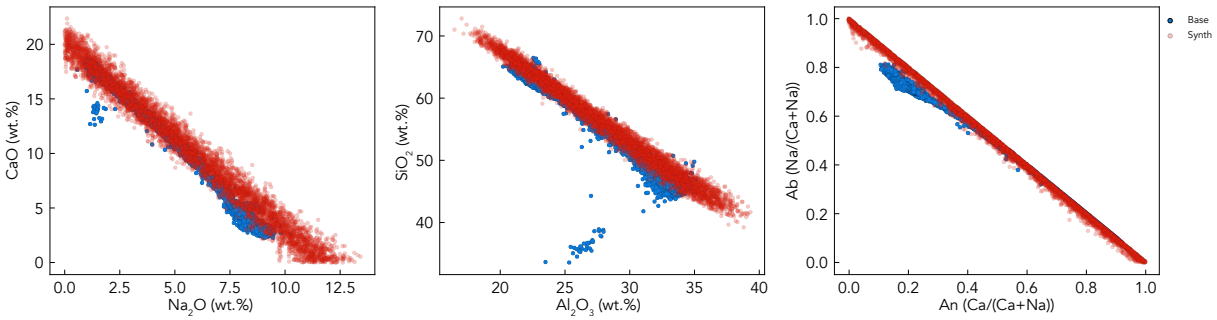

Figure S4. Comparison of natural (blue) and synthetic (red) plagioclase compositions. Left: Na_2O vs. CaO, showing the continuous albite-anorthite substitution. Center: Al_2O_3 vs. SiO_2, reflecting the coupled Si-Al substitution. Right: An vs. Ab molar fractions, plotting along the expected An-Ab tie line.

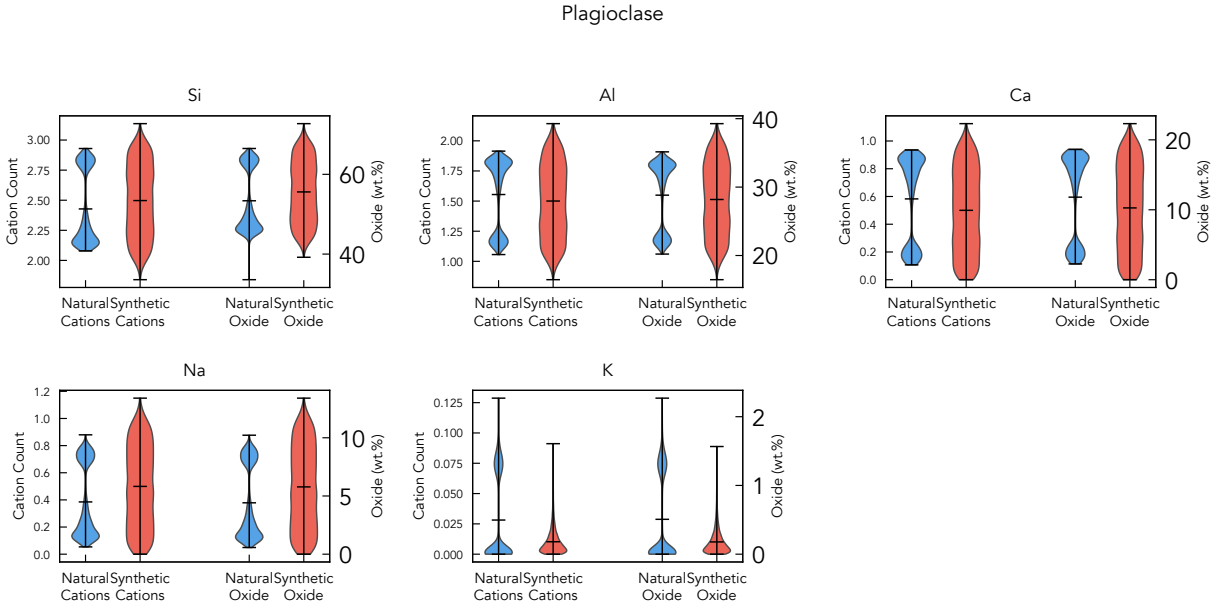

Figure S5. Violin plots comparing the distributions of cation counts and oxide weight percentages for natural and synthetic plagioclase. The figure panel follows the structure outlined in Figure S3. The synthetic generator captures the full range of the An-Ab solid solution, with close agreement for all major cations.

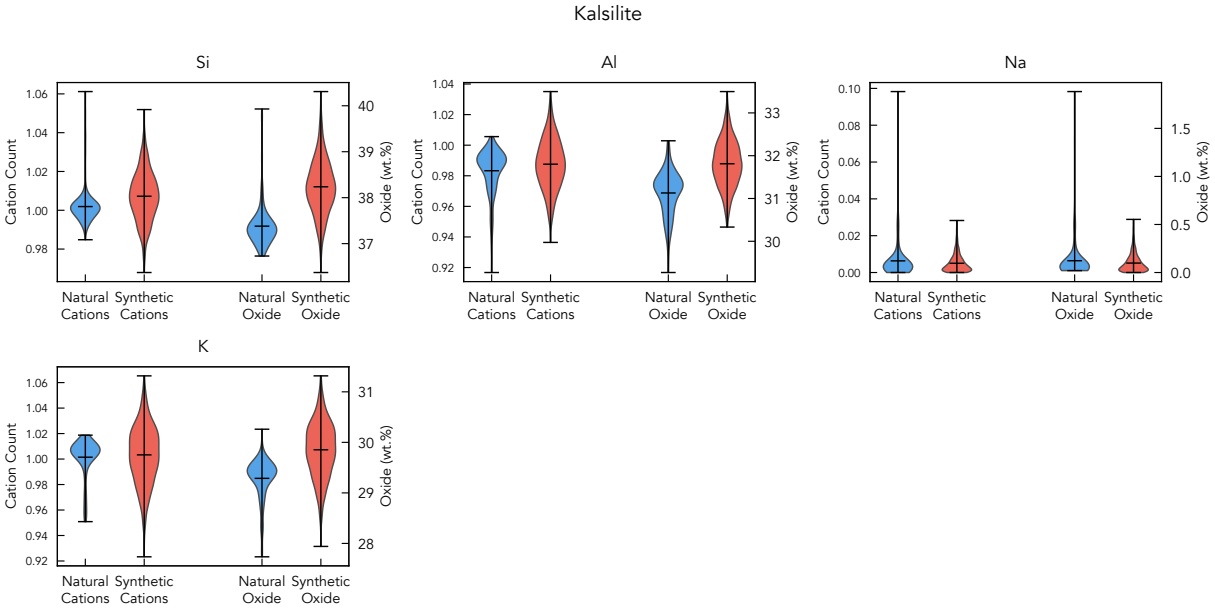

Figure S6. Violin plots comparing the distributions of cation counts and oxide weight percentages for natural and synthetic kalsilite. The synthetic generator reproduces the compositions of this feldspathoid with close agreement for Si, Al, Na, and K.

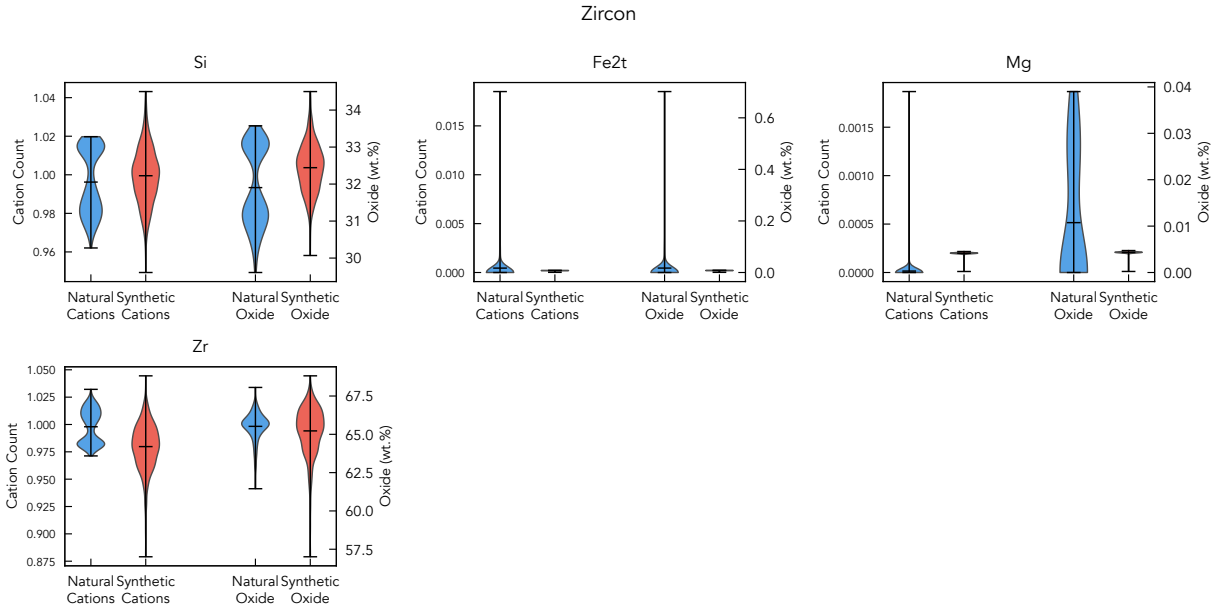

Figure S7. Violin plots comparing the distributions of cation counts and oxide weight percentages for natural and synthetic zircon. Si and Zr cation counts are tightly clustered near their stoichiometric values of 1.0, with small Fe and Mg trace impurities.

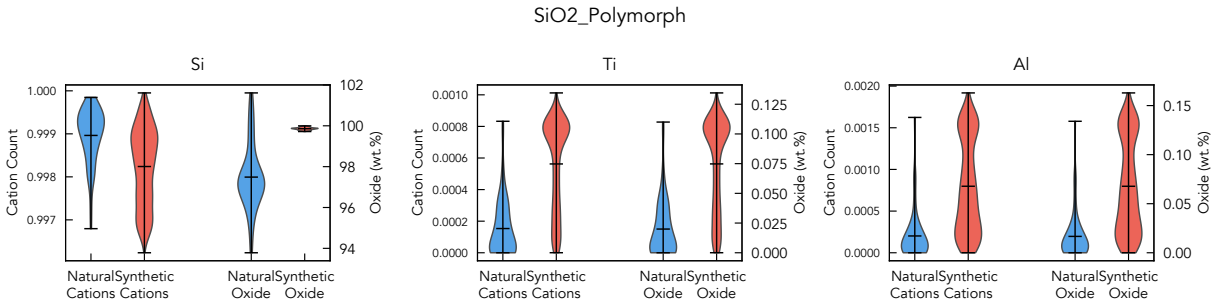

Figure S8. Violin plots comparing the distributions of cation counts and oxide weight percentages for natural and synthetic SiO_2-polymorphs. Si cation counts are tightly clustered near 1.0 for both datasets, with only trace Ti and Al impurities.

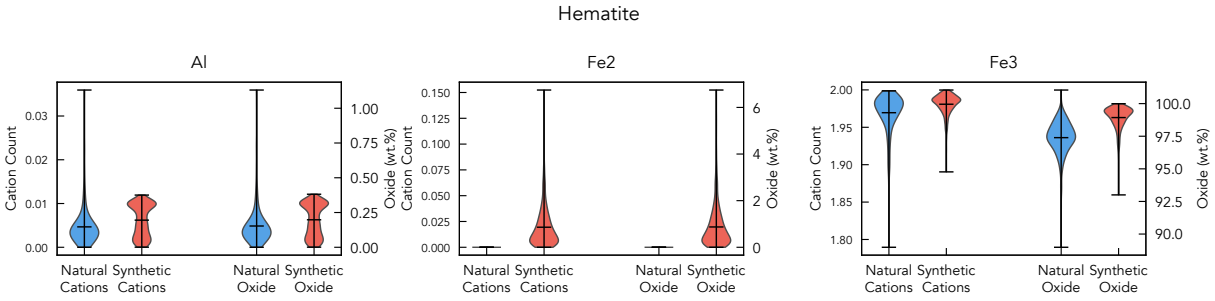

Figure S9. Violin plots comparing the distributions of cation counts and oxide weight percentages for natural and synthetic hematite. Fe^{3+} cation counts are tightly clustered near 2.0 for both datasets, with trace Fe2+ and Al impurities.

Table S1. Acquisition conditions for co-collected EBSD-EDS maps.

	Mount Hood (MH0811) (initial)	**Mount Hood (MH0811) (optimized)**	**Tuolumne (Bii) (optimized)**
Working distance (mm)	25	12	15
Beam current (nA)	20.5	6.9	6.9
Voltage (kV)	20	20	20
Frame live time (ms)	2.56	7.65	9.31
Frame averaging	1	20	20
Total acquisition time (ms)	2.56	153	186

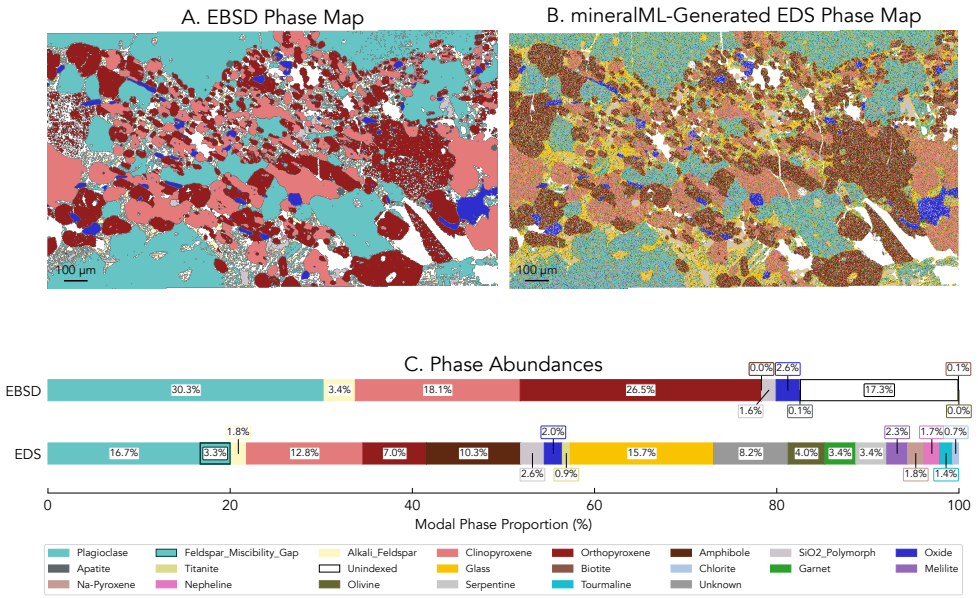

Figure S10. Comparison of EBSD and `mineralML`-generated EDS phase maps for the Mount Hood andesite (MH0811) also shown in Figure 6. The map was collected under initial acquisition conditions of working distance 25 mm, beam current of 20.5 nA, and total dwell time of 2.56 ms per pixel, shown in the initial column of Table S1. A) EBSD phase map, which is largely unaffected by short dwell times. B) Co-collected EDS phase map classified by `mineralML`. The short dwell times contribute to large uncertainties, producing a phase map with many spurious pixels within the major rock forming minerals of plagioclase feldspar, clinopyroxene, and orthopyroxene. These large uncertainties contribute to the checkerboarding observed across minerals. C) Comparison of modal phase abundances (%) from EBSD and `mineralML`-classified EDS maps. Agreement is poor between the EBSD- and EDS-derived modal phase proportions.

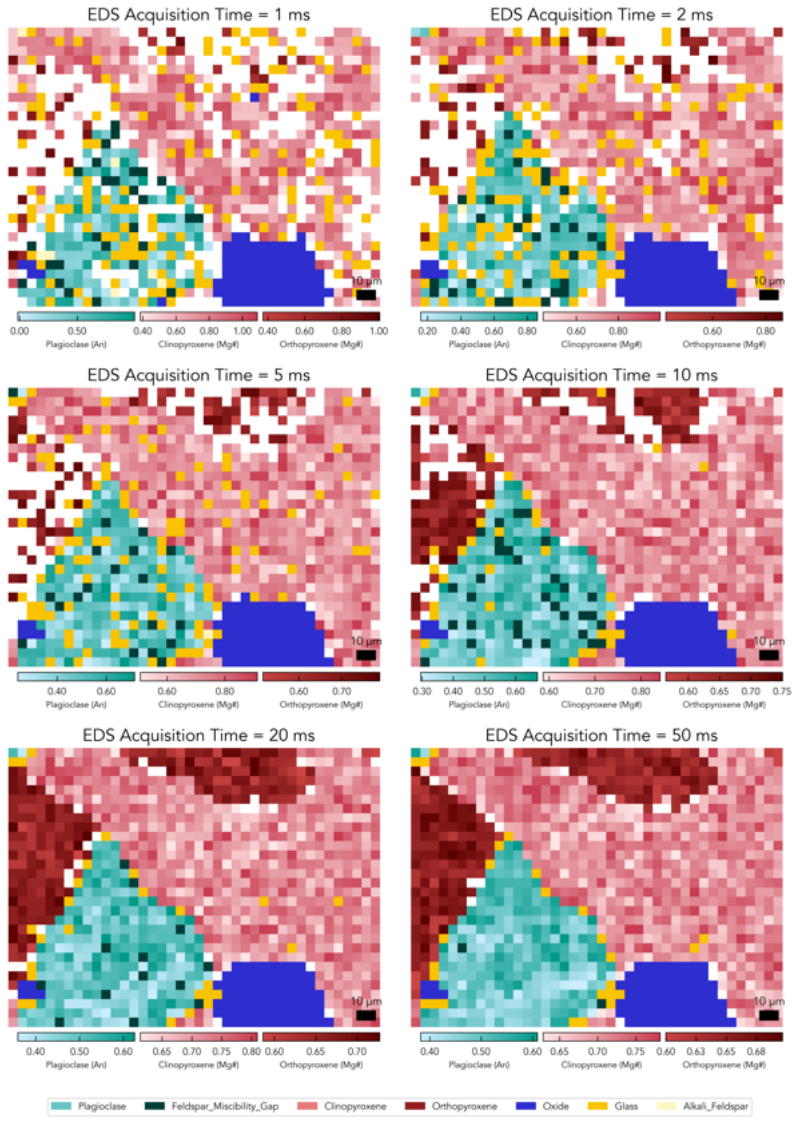

Figure S11. `mineralML`-generated EDS phase maps for a sub-region of the Mount Hood andesite at pixel acquisition times ranging from 1-50 ms. Categorical phases are shown in solid colors; solid-solution compositions (plagioclase An content, pyroxene Mg#) are shown with colorbars beneath each panel, with colors centered on the mean with a range of $\pm$ 2σ. At short dwell times, the maps show spurious phase labels (marked in white) and many pixels of glass interspersed within mineral crystals. Compositional ranges are extremely broad, reflecting large analytical uncertainty. As acquisition time increases, the checkerboarding diminishes and the colorbar ranges for the solid solution components decreases and converge. By 50 ms, the phase map is free of spurious labels and glass pixels are confined to grain boundaries where mixed signals emerge.

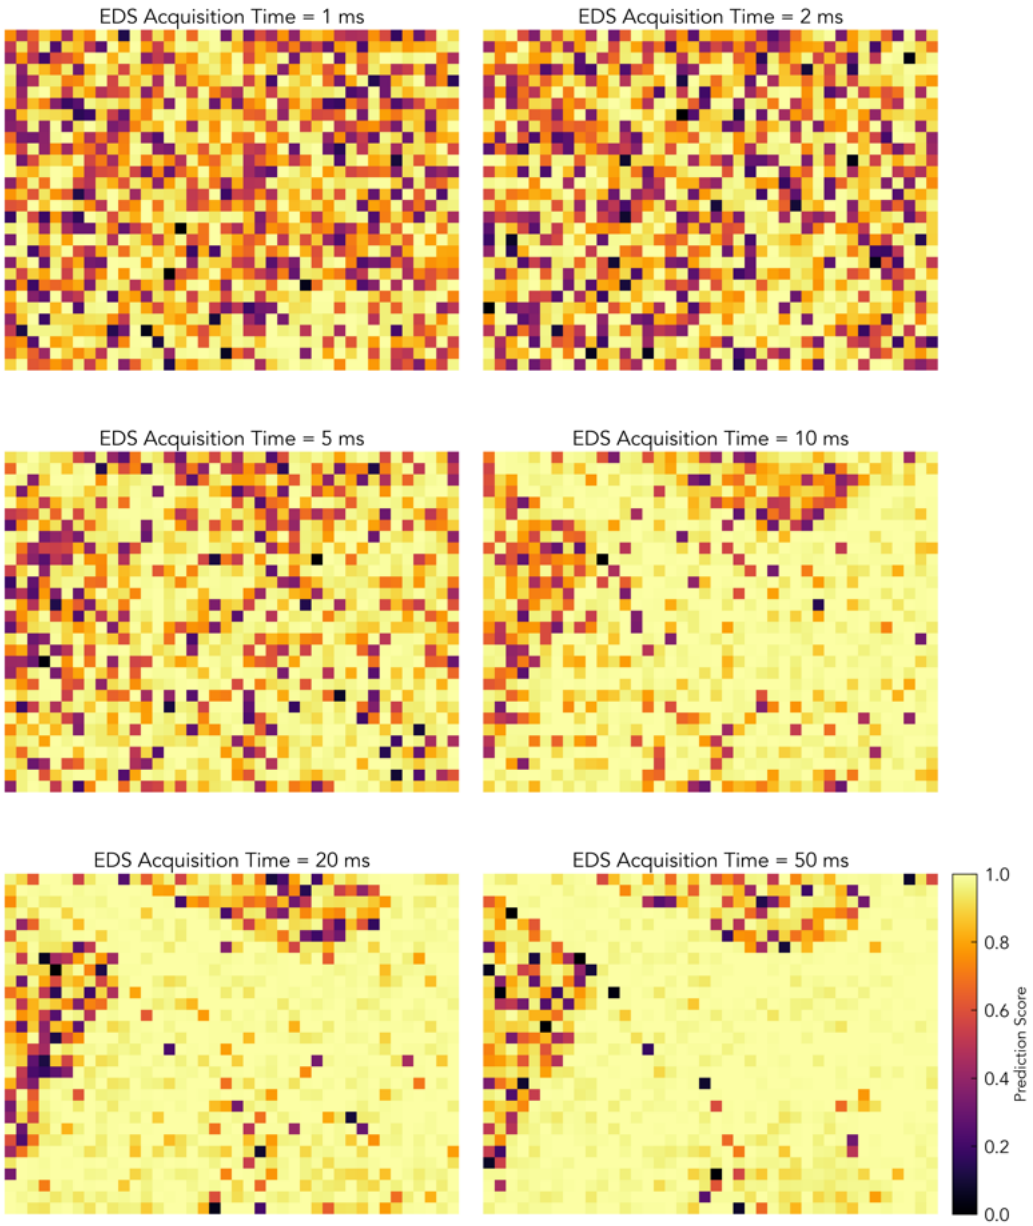

Figure S12. `mineralML`-derived prediction scores associated with the predictions from Figure S11, for a sub-region of the Mount Hood andesite at pixel acquisition times ranging from 1-50 ms. At 1 ms, prediction scores are uniformly low across the map. Most pixels have a prediction score of less than 0.8, reflecting the large uncertainties at this dwell time, where the compositional ranges for these phases all overlap. As acquisition times increase, the prediction scores rise. At 50 ms, most prediction scores exceed 0.95. The increase in prediction scores with increasing acquisition time reflects how `mineralML` is able to communicate data quality limitations to the user through the probabilistic framework.

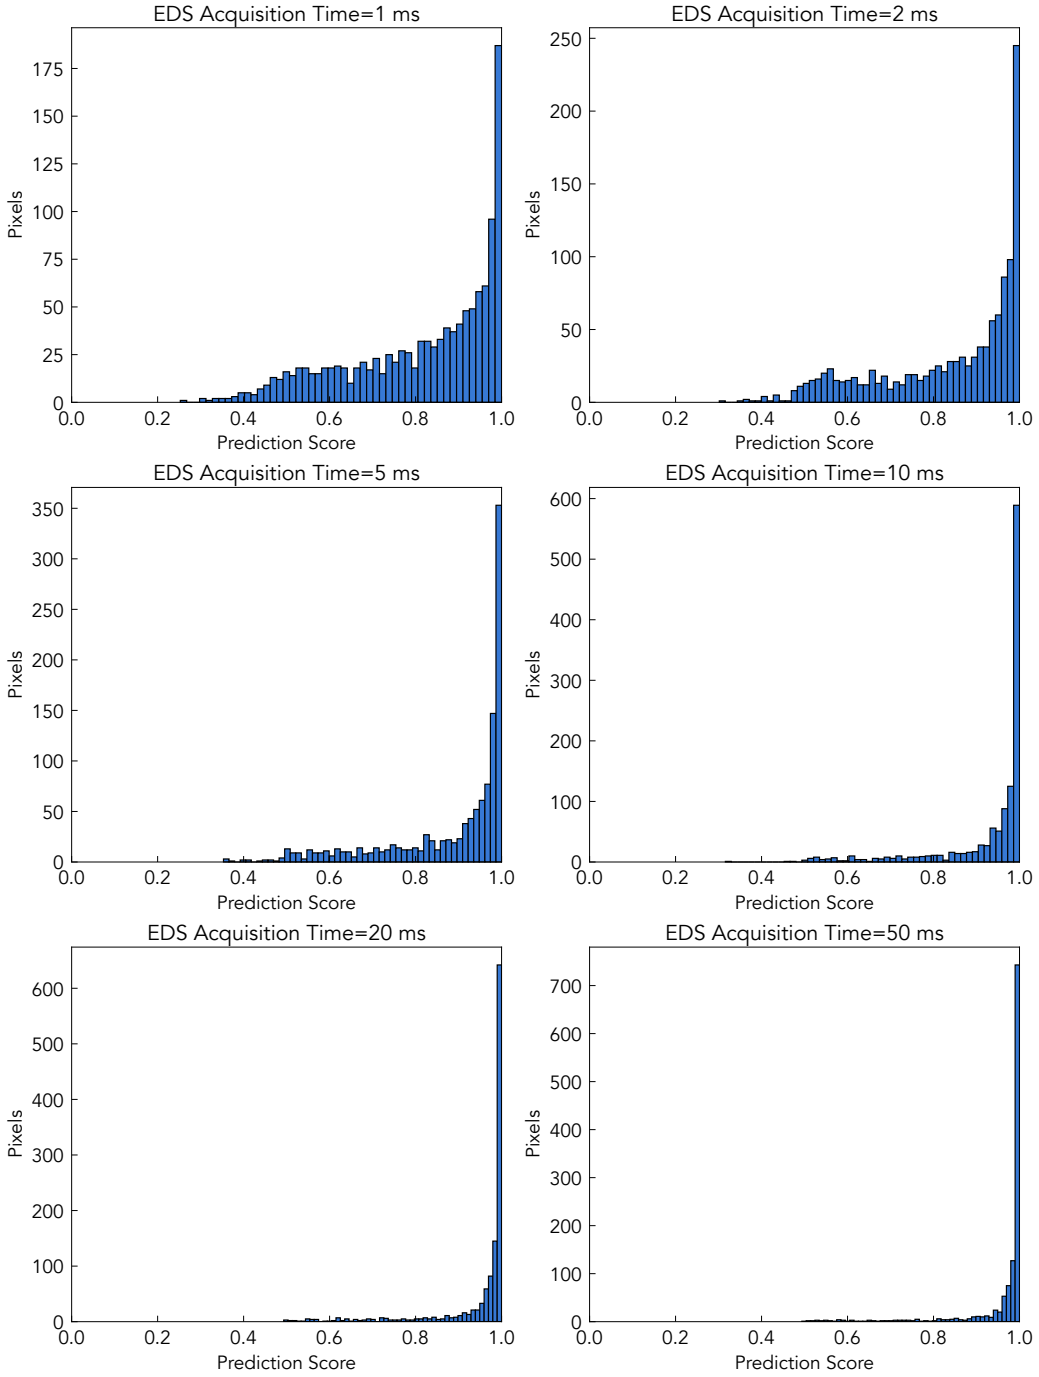

Figure S13. Histograms of `mineralML` prediction scores shown spatially in Figure S12, for a sub-region of the Mount Hood andesite at pixel acquisition times ranging from 1-50 ms. At 1 ms, the distribution of prediction scores is broad and ranges from around 0.25-1, with a hump around 0.5. With increasing acquisition time, the distributions shift towards higher prediction scores and are increasingly concentrated towards 1.0. At 50 ms, the distribution is strongly centered close to 1.0 with a small tail of lower-confidence pixels corresponding to grain boundaries classified as glass.

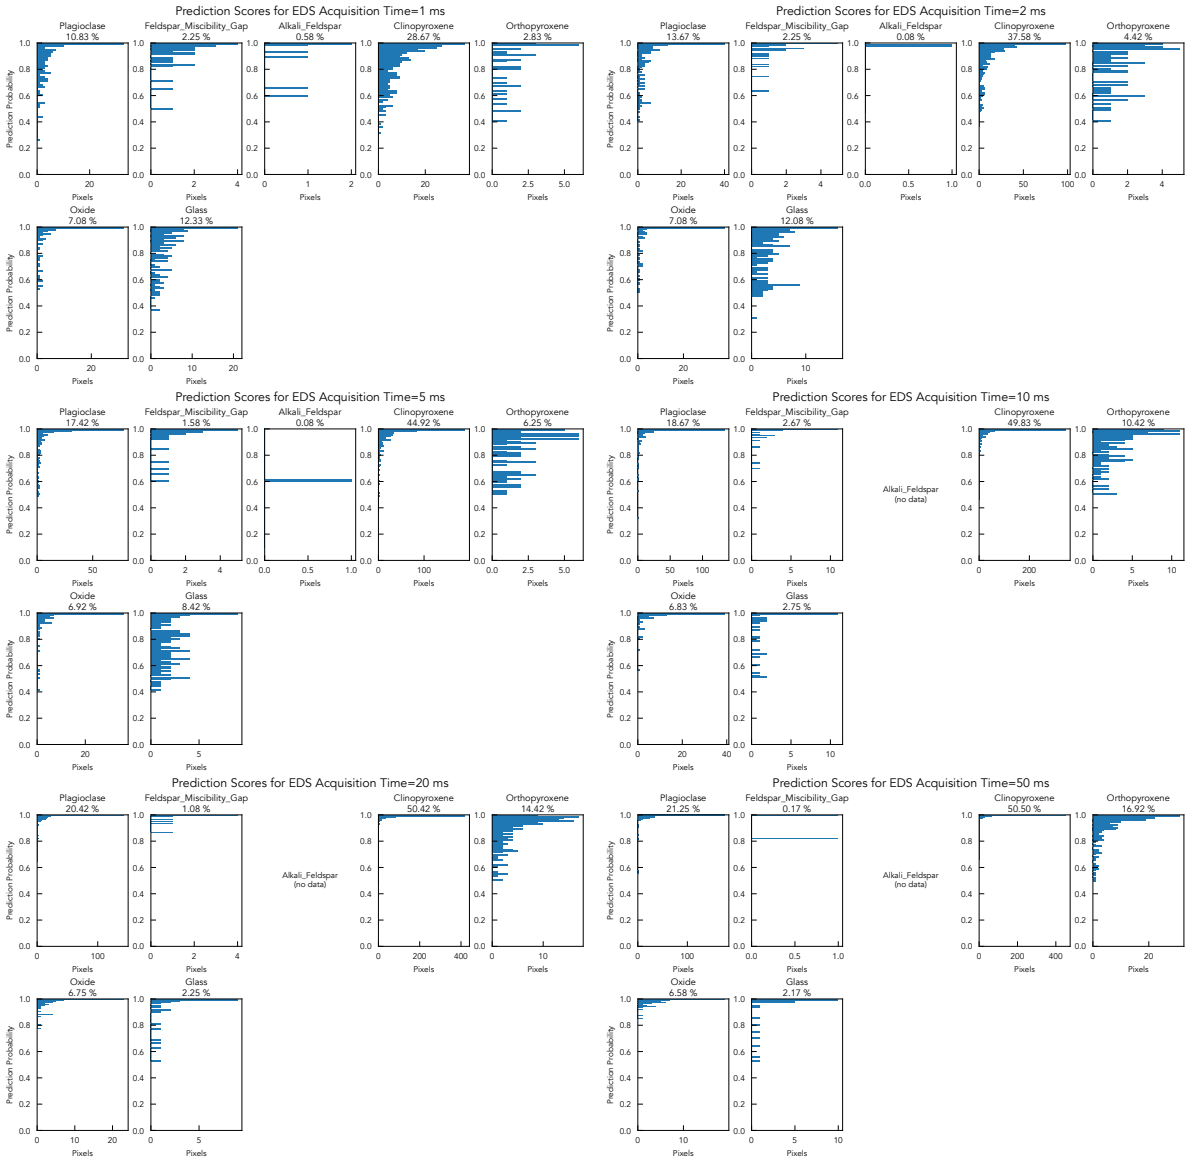

Figure S14. Individual phase-related histograms of `mineralML` prediction scores, for a sub-region of the Mount Hood andesite at pixel acquisition times ranging from 1-50 ms. Each row corresponds to one acquisition time, and each column shows the prediction score distribution for a single mineral phase, with the percentage of the map classified as that phase indicated below the phase name. At 1 ms, prediction scores are broadly distributed for all phases. As acquisition times increase, the prediction scores shift towards higher values and the proportion of some phases decrease. Glass and alkali feldspar proportions decrease, as `mineralML` converges upon more likely mineral classifications. These labels are artifacts of high uncertainties.