

From STEM-EDXS data to phase separation and quantification using physics-guided NMF

Adrien Teurtrie

Unité Matériaux et Transformations, UMR-CNRS 8207,

Université de Lille, Lille, France and

Electron Spectrometry and Microscopy Laboratory (LSME), Institute of Physics (IPHYS),

École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Nathanaël Perraudin and Thomas Holvoet

Swiss Data Science Center (SDSC), ETHZ, Zürich, Switzerland

Hui Chen, Duncan T.L. Alexander, and Cécile Hébert

Electron Spectrometry and Microscopy Laboratory (LSME), Institute of Physics (IPHYS),

École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Guillaume Obozinski

Swiss Data Science Center (SDS), ETHZ, Zürich, Switzerland

(Dated: May 6, 2024)

Abstract

We present the development of a new algorithm which combines state-of-the-art energy-dispersive X-ray (EDX) spectroscopy theory and a suitable machine learning formulation for the hyperspectral unmixing of scanning transmission electron microscope EDX spectrum images. The algorithm is based on non-negative matrix factorization (NMF) incorporating a physics-guided factorization model. It optimizes a Poisson likelihood, under additional simplex constraint together with user-chosen sparsity-inducing and smoothing regularizations, and is based on iterative multiplicative updates. The fluorescence of X-rays is fully modeled thanks to state-of-the-art theoretical work. It is shown that the output of the algorithm can be used for a direct chemical quantification. With this approach, it is straightforward to include a priori knowledge on the specimen such as the presence or absence of certain chemical elements in some of its phases. This work is implemented within two open-source Python packages, **espm** and **emtables**, which are used here for data simulation, data analysis and quantification. Using simulated data, we demonstrate that incorporating physical modeling in the decomposition helps retrieve meaningful components from spatially and spectrally mixed phases, even when the data are very noisy. For synthetic data with a higher signal, the regularizations yield a tenfold increase in the quality of the reconstructed abundance maps compared to standard NMF. Our approach is further validated on experimental data with a known ground truth, where state-of-the-art results are achieved by using prior knowledge about the sample. Our model can be generalized to any other scanning spectroscopy techniques where underlying physical modeling can be linearized.

I. INTRODUCTION

EDX spectroscopy (EDXS) is widely used in analytical electron microscopy. EDXS is a core-level spectroscopy measuring the X-rays emitted by the atoms when they return to ground state after an excitation induced by the incident electron beam in the scanning transmission electron microscope (STEM). In STEM, EDXS performed at high spatial resolution enables determination of the local chemical composition below the nanometric scale [1]. Nevertheless, analyzing materials with complex microstructures and containing multiple phases poses challenges for extracting meaningful information from STEM-EDXS datasets. Three key issues arise when determining chemical composition in such scenarios: noisy spectra

due to the low cross-sections for fluorescence yields of X-rays; limitations on exposure time and/or electron beam current before irreversible damage occurs on the sample; and spatial and spectral mixing. Spectral mixing occurs whenever the X-ray peaks of different elements overlap (e.g., Si-K α and Sr-L α lines). There can be spatial mixing along the beam direction whenever the typical scale of the material microstructure is below the sample thickness (i.e. thickness of tens or hundreds of nanometer). It is then a challenging task to disentangle the contributions of each phase in the spectra.

To overcome this challenge, the STEM-EDXS data can be modeled as a product of two matrices: one matrix containing the pure spectra of each phase and a second matrix containing their associated spatial distributions. This follows from the principle that each pixel spectrum of the STEM-EDXS datacube can be approximated as a linear combination of the pure phase spectra of the studied material [2]. This linear mixing model applies to other fields such as satellite hyperspectral imaging [3]. Within the latter field, many algorithms have been developed for the decomposition of hyperspectral datacubes, such as non-negative matrix factorisation (NMF) [4], independent component analysis (ICA) [5], vertex component analysis (VCA) [6] or multivariate curve analysis [7]. These methods produce a decomposition of the data which makes them adaptable to many different problems, such that, in either direct or modified forms, they have been successfully applied to STEM-EDXS experimental data [7–12]. However, these algorithms do not benefit from prior knowledge about the shape of the solutions. This can result in a lack of direct interpretability for their solutions, and they have a tendency to overfit and hence fail to remove the noise present in the data.

To solve these issues and improve the decomposition of STEM-EDX spectrum images, here we present a machine learning formulation that uses a physics basis to guide the solutions. The first main contribution of this work is a novel factorization model that integrates the physical modeling of EDX spectra within an NMF framework. This approach has two main additional benefits: the NMF directly outputs a chemical quantification of each phase, and it is possible to introduce prior knowledge on the factors to guide their estimation. Contrary to the spectra, the spatial distribution of phases is not modeled, but it is expected to be smooth, or piecewise smooth, so we also add regularizations in the NMF formulation to avoid overfitting.

The second main contribution of this paper is the continued development of two open-

source Python packages. Previously, we introduced the **emtables** package for creating X-ray cross-section tables, together with the **espm** package for the fast simulation of STEM-EDXS datasets [2]. Now, we augment **espm** with a dedicated NMF algorithm for the decomposition of STEM-EDXS data that incorporates designed constraints, regularization, and a physics-informed factorization model that makes recourse to the **emtables**-generated X-ray cross-sections. **espm** is thus updated into a versatile toolbox for the analysis of EDX spectrum image. These Python packages are freely available and provide the user with advanced control over the functionality of the algorithm and simulations.

In the following, we first describe the physics-informed factorization model. Then, we present an efficient algorithm to estimate pure phase spectra from STEM-EDXS datasets. The effectiveness of our method is tested on both synthetic and experimental datasets. This is described in the subsequent parts of this paper, where we detail the parameters of dataset simulation, the experimental methods, and the results obtained on synthetic and experimental datasets.

II. THE FACTORIZATION MODEL

In STEM-EDX spectrum imaging, the electron beam is scanned over an area of the sample, which leads to the emission of X-rays. The raw data comprises X-ray counts measured by the detector at each location of the beam and at various energy levels. The scan is usually performed in a grid pattern and each position of the beam corresponds to a pixel in the data. The count distribution follows a Poisson distribution with the mean (and rate) equal to the magnitude of the spectrum at that specific energy level and location. We start from the premise that the spectrum of a column of material corresponds to the weighted sum of pure phase spectra. This leads to the formulation of the following factorization model for the data matrix $\mathbf{Y} = (y_{\ell,p}) \in \mathbb{R}^{L \times P}$:

$$\mathbf{Y} \approx \mathbf{D}\mathbf{H}, \quad (1)$$

where “ $\approx$ ” comes from the Poisson noise. Here, $\mathbf{D} = (\delta_{\ell,k}) \in \mathbb{R}^{L \times K}$ is the matrix whose columns are the spectra of pure phases and $\mathbf{H} = (h_{k,p}) \in \mathbb{R}^{K \times P}$ is the matrix whose columns are the spatial abundances (i.e. proportion of each phase) at each pixel. L is the number of energy channels, P is the number of pixels and K is the expected number of phases in the

observed sample. This linear mixing model has been widely utilized as a basis to process EDXS data [3].

Without additional constraint, there is an infinite number of pairs $(\mathbf{D}, \mathbf{H})$ that are solutions to equation (1), and in most cases a significant fraction of the elements of $\mathbf{D}$ and $\mathbf{H}$ take negative values. Hence, in general, the obtained $\mathbf{D}$ cannot be directly interpreted as a collection of spectra with physical meaning.

To obtain physically meaningful solutions, some formulations further constrain $\mathbf{D}$ and $\mathbf{H}$. For instance, NMF constrains $\mathbf{D}$ and $\mathbf{H}$ to be non-negative and ICA constrains the rows of $\mathbf{H}$ to be statistically independent. While these formulations are applicable to a broad range of problems, they still lack sufficient constraints to guarantee that the resulting decomposition matches the permissible spectral features dictated by the physics of EDXS.

We therefore propose a formulation which is much more structured and constrained. First, we propose to further decompose $\mathbf{D}$ as the product $\mathbf{GW}$, so that the factorization model becomes:

$$\mathbf{Y} \approx \mathbf{GWH} \quad (2)$$

where $\mathbf{G} = (g_{\ell,m}) \in \mathbb{R}^{L \times M}$ is a pre-computed design matrix whose columns correspond to the modeled X-ray emissions. M is the maximal number of pure spectral features in modeling each EDX spectrum. $\mathbf{W} = (w_{m,k}) \in \mathbb{R}^{M \times K}$ is a matrix of learnable parameters whose k^{th} column represents the amount of $\mathbf{G}$'s spectral feature in the k^{th} phase. We note that in (2), the product $\mathbf{GWH}$ is only an approximation to $\mathbf{Y}$. First, the observed X-ray emissions conform to a Poisson distribution parametrized by $\mathbf{GWH}$. This is detailed in Section III. Second, the model relies on the thin foil approximation for the calculation of the absorption [2], which can lead to errors of up to a few percent [13] in typical samples ~ 100 nm thick. The matrix $\mathbf{H}$ plays the same role here as in (1).

We now provide details on the design of the matrix $\mathbf{G}$ so that the product $\mathbf{GW}$ models the shape of the EDX spectra of pure phases. The modeling of X-ray emission mainly applies to STEM-EDXS and is based on our earlier work [2]. The EDX spectrum is the sum of two contributions:

- The characteristic X-ray emissions result from discrete transitions originating from inner-shell ionization of the probed atoms by the electron beam. The energies of such transitions are specific to each chemical element. For a given element Z , the X-ray

emissions in the ℓ^{th} energy channel can be modeled as:

$$\kappa_{\ell,Z} = \sum_{ij} x_{ij}(Z) \frac{1}{\Delta(\varepsilon_{ij}(Z))\sqrt{2\pi}} \exp^{-\frac{1}{2}\left(\frac{\varepsilon_{\ell}-\varepsilon_{ij}(Z)}{\Delta(\varepsilon_{ij}(Z))}\right)^2} \quad (3)$$

In this formulation, ε is the energy of the measured X-ray, Δ is the energy broadening of the detector, and x_{ij} is the X-ray emission cross-section of the transition between the i^{th} sub-shell and the j^{th} sub-shell. The values of Δ should be calibrated by the experimentalist. The effects of absorption and detection efficiency are included in the calculation of x_{ij} , as described in our previous publication [2] along with more details about our modeling.

- The bremsstrahlung X-ray emissions result from a momentum loss of the electrons within the electric field of the probed material and produce a continuous spectrum. This radiation is commonly parametrized as a polynomial of the energy [14]. However, such polynomials can potentially take negative values, which is not possible physically. To address this issue, we propose a parameterization of an appropriate form of non-negative quadratic polynomial, via non-negative coefficients which guarantee that the polynomial remains positive, as follows:

$$\begin{aligned} b_{\ell} &= \left(\gamma_0 \frac{\tilde{\varepsilon}_{\ell}}{\varepsilon_{\ell}} (1 - \tilde{\varepsilon}_{\ell}) + \gamma_1 \frac{\tilde{\varepsilon}_{\ell}^2}{\varepsilon_{\ell}} \right) d_{\ell}(\varepsilon_{\ell}) c_{\ell}(\varepsilon_{\ell}) \\ &= \gamma_0 b_{\ell}^{(0)} + \gamma_1 b_{\ell}^{(1)} \\ \text{with } \tilde{\varepsilon}_{\ell} &= \frac{e_0 - \varepsilon_{\ell}}{e_0}, \end{aligned} \quad (4)$$

where γ_0 and γ_1 are the parameters of the model, e_0 is the energy of the beam and $d_{\ell}(\varepsilon_{\ell})$ and $c_{\ell}(\varepsilon_{\ell})$ are the detection efficiency and absorption correction, as calculated in [2], respectively. See Appendix A for more details. With this parameterization, the bremsstrahlung can be modeled by two numbers (γ_0 and γ_1) for each phase.

The composition of a single phase can be expressed with vector $\mathbf{c} = (c_z)$ in atomic percentages where $\mathbf{z} \in \{Z_1, \dots, Z_{M'}\}$ are the corresponding atomic numbers, with $M' = M - 2$. The intensity at the ℓ^{th} energy channel $\phi_{\ell k}$ of the k^{th} phase can thus be written as

[15]:

$$\begin{aligned}
\phi_{\ell k} &= N_A N_e I_e \tau \Omega \rho_k^* \sum_z c_{zk} \kappa_{\ell,z} + b_{\ell k} \\
&= \xi_k \left(\sum_z c_{zk} \kappa_{\ell,z} + \frac{b_{\ell k}}{\xi_k} \right),
\end{aligned} \tag{5}$$

where N_A is the Avogadro number, N_e is the number of electrons in a unit of electric charge, I_e is the beam current, τ is the acquisition time, Ω is the geometric efficiency of the detector, and ρ_k^* is the mass-thickness of the compound. We use the pre-factor ξ_k to isolate the elemental concentrations; it contains the mass-thickness of the k^{th} phase and the other experimental parameters. Eq. (5) can thus be factorized as

$$\begin{aligned}
\phi_k &= \xi_k \begin{pmatrix} \kappa_{1,Z_1} & \dots & \kappa_{1,Z_{M'}} \\ \vdots & \ddots & \vdots \\ \kappa_{L,Z_1} & \dots & \kappa_{L,Z_{M'}} \end{pmatrix} \begin{pmatrix} c_{Z_1,k} \\ \vdots \\ c_{Z_{M'},k} \end{pmatrix} \\
&\quad + \xi_k \begin{pmatrix} b_1^{(0)} & b_1^{(1)} \\ \vdots & \vdots \\ b_L^{(0)} & b_L^{(1)} \end{pmatrix} \begin{pmatrix} \gamma_{0,k}/\xi_k \\ \gamma_{1,k}/\xi_k \end{pmatrix} \\
&= \xi_k (\mathbf{K} \mathbf{c}_k + \mathbf{B} \boldsymbol{\gamma}_k).
\end{aligned}$$

Often, in samples studied using STEM-EDXS, several phases coexist in the acquired spectrum image and their proportions vary across the scanned area. To reflect this, we further develop Eq. (5) by introducing the abundances of each phase at each pixel of the scan $\eta_{k,p}$:

$$\begin{aligned}
\Phi &= \left[\begin{pmatrix} \kappa_{1,Z_1} & \dots & \kappa_{1,Z_{M'}} \\ \vdots & \ddots & \vdots \\ \kappa_{L,Z_1} & \dots & \kappa_{L,Z_{M'}} \end{pmatrix} \begin{pmatrix} c_{Z_1,1} & \dots & c_{Z_1,K} \\ \vdots & \ddots & \vdots \\ c_{Z_{M'},1} & \dots & c_{Z_{M'},K} \end{pmatrix} \right. \\
&\quad \left. + \begin{pmatrix} b_1^{(0)} & b_1^{(1)} \\ \vdots & \vdots \\ b_L^{(0)} & b_L^{(1)} \end{pmatrix} \begin{pmatrix} \gamma_{0,1}/\xi_1 & \dots & \gamma_{0,K}/\xi_K \\ \gamma_{1,1}/\xi_1 & \dots & \gamma_{1,K}/\xi_K \end{pmatrix} \right] \\
&\quad \times \begin{pmatrix} \xi_1 \eta_{1,1} & \dots & \xi_1 \eta_{1,P} \\ \vdots & \ddots & \vdots \\ \xi_K \eta_{K,1} & \dots & \xi_K \eta_{K,P} \end{pmatrix} \\
&= [\mathbf{K} \mathbf{C} + \mathbf{B} \boldsymbol{\Gamma}] \mathbf{H}.
\end{aligned}$$

The equation above can be further simplified as a single matrix product:

$$\Phi = (\mathbf{K}\mathbf{C} + \mathbf{B}\mathbf{\Gamma}) \mathbf{H} = \left(\begin{bmatrix} \mathbf{K} & \mathbf{B} \end{bmatrix} \begin{bmatrix} \mathbf{C} \\ \mathbf{\Gamma} \end{bmatrix} \right) \mathbf{H} = \mathbf{G}\mathbf{W}\mathbf{H}. \quad (6)$$

This factorization model has several advantages. Thanks to the modeling of the spectral features, it both decomposes the analysed dataset and denoises its spectral components in one go. Additionally, the decomposition directly provides a chemical quantification from the resulting $\mathbf{C}$ matrix (i.e. a part of the learned $\mathbf{W}$ matrix). It should be noted that the quantification is only valid in the case where

$$\forall k \quad \sum_{m=1}^{M-2} w_{mk} = 1, \quad (7)$$

which is enforced during the learning of $\mathbf{W}$. Since there is an arbitrary scaling between $\mathbf{W}$ and $\mathbf{H}$ when solving Eq. (2), the constraint (7) makes the first $M - 2$ rows of $\mathbf{W}$ correspond to elemental fractions. The last 2 rows of $\mathbf{W}$, that correspond to the learning of the bremsstrahlung model, are let free. Therefore, the formalism used in this work makes a direct bridge between the mixed raw data and a chemical quantification, phase by phase. One should note that a more physically accurate model of the EDX spectra would include the spatial variations of the mass-thickness ($\xi_{k,p}$), but this would break the linearity of Eq. (6), which would in turn greatly complicate the decomposition of experimental datasets.

Another advantage of this model resides in the possibility of imposing physical constraints on the analysis. This can be done by keeping some elements of $\mathbf{W}$ constant during the decomposition. For example, the contribution $w_{\mathbf{z},k}$ of some elements ($\mathbf{z}$) of some phases (k) can be kept to $w_{\mathbf{z},k} = 0$. This is particularly useful when it is known that specific element(s) are absent in some phase(s) of the sample but present in other phase(s).

In summary, the factorization model presented here is a tool to improve the NMF decomposition of STEM-EDX spectrum images, such that its results are denoised and quantified. Thanks to this model, physical constraints, based on prior knowledge about the studied material, can be added to further limit the solutions of the unmixing problem. The following section details the optimization process to obtain $\mathbf{W}$ and $\mathbf{H}$ given $\mathbf{Y}$.

III. THE OPTIMIZATION PROBLEM

The goal of the NMF algorithm developed in this work is to find $\dot{\mathbf{W}}$ and $\dot{\mathbf{H}}$ such that $\mathbf{Y} \approx \mathbf{G}\dot{\mathbf{W}}\dot{\mathbf{H}}$.

$\dot{\mathbf{W}}$ and $\dot{\mathbf{H}}$ are determined through the optimization of a loss function. The loss function is here to measure the discrepancies between $\mathbf{Y}$ and the product $\mathbf{G}\dot{\mathbf{W}}\dot{\mathbf{H}}$. The goal is then to find $\dot{\mathbf{W}}$ and $\dot{\mathbf{H}}$ that optimize the value of the loss function and therefore minimise the discrepancies between the data and the model.

In the most basic formulation of NMF, the matrices $\mathbf{W}$ and $\mathbf{H}$ are determined by minimizing the square loss function:

$$L_2(\mathbf{W}, \mathbf{H}) = \|\mathbf{Y} - \mathbf{G}\mathbf{W}\mathbf{H}\|^2, \quad (8)$$

while preserving the following positivity inequalities:

$$\forall k, m, p, \quad w_{mk} \geq 0, \quad h_{kp} \geq 0 \quad (9)$$

The optimization of L_2 is well-suited when the statistical distribution of the data follows a Gaussian law. In STEM, the inelastic scattering events that lead to X-ray emission, and hence the X-ray emissions themselves, instead follow a Poisson distribution [16]. Given the X-ray detector properties, these statistics are further preserved at the detection level. Thus, the STEM-EDXS data follow a distribution with, according to our model, a mean of:

$$\Phi = (\phi_{\ell,p}), \quad (10)$$

where Φ' corresponds to the modeled spectrum-image associated to the true distribution of materials within the observed sample. Thus, the probability of the observed data ($\pi_{\phi_{\ell,p}}(y_{\ell,p})$) follow the expression:

$$\pi_{\phi_{\ell,p}}(y_{\ell,p}) = \frac{(\phi_{\ell,p})^{y_{\ell,p}}}{y_{\ell,p}!} e^{-\phi_{\ell,p}}. \quad (11)$$

Given that this probability of observing the data is parameterized by $\mathbf{G}\mathbf{W}\mathbf{H}$, by the maximum likelihood principle a good way to estimate those parameters is to maximize this Poisson likelihood (L'_P) with respect to the parameters $\mathbf{W}$ and $\mathbf{H}$:

$$L'_P = \prod_{\ell,p} e^{-(\mathbf{G}\mathbf{W}\mathbf{H})_{\ell p}} \frac{(\mathbf{G}\mathbf{W}\mathbf{H})_{\ell p}^{\mathbf{Y}_{\ell p}}}{\mathbf{Y}_{\ell p}!} \quad (12)$$

This maximization is equivalent to the minimization of the negative log likelihood. We consider the slightly modified objective:

$$L_P(\mathbf{W}, \mathbf{H}) = \sum_{\ell,p}^{L,P} [-\mathbf{Y}_{\ell p} \log(\mathbf{G}\mathbf{W}\mathbf{H})_{\ell p} + (\mathbf{G}\mathbf{W}\mathbf{H})_{\ell p}] \quad (13)$$

Lee and Seung proposed efficient optimization algorithms for NMF [4, 17] using alternating multiplicative updates on $\mathbf{W}$ and $\mathbf{H}$. If the optimization is begun with non-negative $\mathbf{Y}$, $\mathbf{W}$ and $\mathbf{H}$, non-negativity is preserved by the algorithm updates.

In our formulation of the problem, we additionally enforce both the sum-to-one (7) and non-negativity (9) constraints. The linear constraint of Eq. (7) is known as a simplex constraint. Furthermore, the simplex and non-negativity constraints on $\mathbf{W}$ together induce sparsity (i.e. induce solutions which have a certain fraction of entries that are exactly equal to zero) in the columns of $\mathbf{W}$ [18] and, as a consequence, tend to reduce the presence of mixing between the different phases.

The physical model presented above does not put any additional constraint on $\mathbf{H}$, except for non-negativity, even though its rows are high dimensional. The fact that the phase composition can be estimated independently at each pixel (i.e. each sample column) but only from very low counts data would necessarily result in an overfitted and noisy $\hat{\mathbf{H}}$ matrix. However, we can exploit two properties of the samples studied in electron microscopy. First, their phase composition will mostly vary smoothly within space. Then, because the samples are very thin, only a very few different phases can be measured at each pixel, with the vast majority of pixels corresponding to only one phase. Thus, the columns of $\mathbf{H}$ are typically sparse. In consequence, we propose adding two regularization terms to L_P . The first term is known as the Laplacian regularization:

$$r_{\Delta} = \frac{\lambda}{2} \text{tr}(\mathbf{H}\mathbf{\Delta}\mathbf{H}^T) = \lambda \sum_{k=1}^K \sum_{(n,m) \in E} (h_{kn} - h_{km})^2, \quad (14)$$

where E is the set of pairs of (vertically or horizontally) adjacent pixels. In the expression above, $\mathbf{\Delta}$ is the $P \times P$ graph Laplacian matrix [19]. The $\text{tr}(\mathbf{H}\mathbf{\Delta}\mathbf{H}^T)$ is the sum of square differences between neighboring pixels of $\mathbf{H}$.

The more the columns of $\mathbf{H}$ are different from their neighbours, the more the value of r_Δ grows. The competition between the value of r_Δ and L_P will make $\mathbf{H}$ smoother and thus reduce some of the noise.

The second term, which we call logarithmic regularization, is expressed as:

$$r_{log} = \mu \sum_{k,p}^{K,P} \log(h_{kp} + \alpha) \quad (15)$$

This regularization aims at inducing sparsity in $\mathbf{H}$. The sparsity inducing property of the logarithm can be understood as follows. The logarithm is a strictly increasing function, with a steeper slope for small positive values and a lower slope for higher positive values. Therefore, the cost of increasing small $\mathbf{H}$ values is higher than the cost of increasing higher values of $\mathbf{H}$. This tends to increase the proportion of the predominant phase at each pixel, and thus reduce the mixing between phases. While μ controls the strength of this regularisation, α controls its slope at $h_{kp} = 0$.

The combination of smoothness and sparsity using both the Laplacian and logarithm regularizations tends to produce smooth maps and should yield more accurate decompositions. Taking into account the above regularizations and the two constraints (9) and (7), the optimization problem we consider is now:

$$\begin{aligned} & \underset{\mathbf{W}, \mathbf{H}}{\text{minimise}} \quad L_f(\mathbf{W}, \mathbf{H}) \\ & \text{such that} \quad \mathbf{W} \geq \epsilon, \mathbf{H} \geq \epsilon, \sum_{m \leq M-2} \mathbf{W}_m = \mathbf{1} \\ & \text{and} \quad L_f(\mathbf{W}, \mathbf{H}) = L_P(\mathbf{W}, \mathbf{H}) + r_\Delta(\mathbf{H}) + r_{log}(\mathbf{H}), \end{aligned} \quad (16)$$

where $\epsilon > 0$ is a small constant close to the machine precision. This constant ensures that the loss function is well-defined across the entire domain. In practical terms, when the value of ϵ is sufficiently small, the outcomes closely resemble those obtained under a positivity constraint.

Solving the problem formulated above leads to a decomposition of $\mathbf{Y}$ that includes the EDX spectroscopy modeling, yielding a non-negative $\mathbf{W}$ matrix and a non-negative, row-wise smooth and column-wise sparse $\mathbf{H}$ matrix. It is possible to find solutions $\hat{\mathbf{W}}$ and $\hat{\mathbf{H}}$ for problem (16) via an iterative algorithm, as detailed in the next section.

IV. SOLVING THE OPTIMIZATION PROBLEM

Our method to solve the optimization problem (16) is explained in detail in the technical report by Perraudin et al. [20]. The algorithm consists of iterative alternating updates for $\mathbf{W}$ and $\mathbf{H}$. These updates are a generalisation of the multiplicative update from Lee and Sung [17]. Generally, the multiplicative updates present the advantage of being fast [21] while naturally preserving the non-negativity constraint. Indeed, it is sufficient to initialize the algorithm with positive $\mathbf{W}$ and $\mathbf{H}$ for constraint (9) to be respected throughout the optimization process. The two multiplicative updates derived in this work come from Proposition 1 and eq. (22) in [20]. We refer the reader to Appendix B for details on how these updates are derived for this contribution. The resulting updates are formulated as follows:

$$\begin{aligned}\mathbf{W}^{t+1} &= \mathcal{P}_\epsilon \left(\mathbf{W}^t \odot \left(\mathbf{G}^T (\mathbf{Y} \oslash \mathbf{G} \mathbf{W}^t \mathbf{H}^t) \mathbf{H}^{tT} \right) \right. \\ &\quad \left. \oslash \left(\mathbf{G}^T \mathbf{1}^{L,P} \mathbf{H}^{tT} + \mathbf{u} \boldsymbol{\nu}^T \right) \right) \\ &= \mathcal{P}_\epsilon \left(\mathbf{W}^t \odot \boldsymbol{\Sigma}^t \oslash (\mathbf{T}^t + \mathbf{u} \boldsymbol{\nu}^T) \right)\end{aligned}\tag{17}$$

$$\begin{aligned}\mathbf{H}^{t+1} &= \mathcal{P}_\epsilon \left(\mathbf{H}^t \right. \\ &\quad \odot \left(\mathbf{W}^{tT} \mathbf{G}^T (\mathbf{Y} \oslash (\mathbf{G} \mathbf{W}^t \mathbf{H}^t)) + 8\lambda \max(\mathbf{H}) \right) \\ &\quad \oslash \left(\mathbf{W}^{tT} (\mathbf{G}^T \mathbf{1}^{L,P}) + \mu \oslash (\mathbf{H}^t + \alpha) \right. \\ &\quad \left. \left. + 8\lambda \max(\mathbf{H}) + \lambda \mathbf{H} \boldsymbol{\Delta} \right) \right),\end{aligned}\tag{18}$$

where the $\oslash$ and $\odot$ symbols represent element-wise matrix division and multiplication, respectively. The function $\mathcal{P}_\epsilon(x) = \max(x, \epsilon)$ corresponds to the projection on the convex set $x \geq \epsilon$. For each element x of $\mathbf{H}$ or $\mathbf{W}$, it assigns ϵ if $x < \epsilon$ and x otherwise. $\mathbf{u} = (u_m) \in \mathbb{N}^M$ is a vector with $u_m = 1, m \in 1, \dots, M', u_{M-1} = 0$ and $u_M = 0$. To simplify the notation, we define $\boldsymbol{\Sigma} = \left(\mathbf{G}^T (\mathbf{Y} \oslash \mathbf{G} \mathbf{W}^t \mathbf{H}^t) \mathbf{H}^{tT} \right) = (\sigma_{mk}) \in \mathbb{R}^{M \times K}$ and $\mathbf{T} = \mathbf{G}^T \mathbf{1}^{L,P} \mathbf{H}^{tT} = (\tau_{mk}) \in \mathbb{R}^{M \times K}$ as the numerator and a component of the denominator in eq. (17), respectively. The vector $\boldsymbol{\nu} = (\nu_k) \in \mathbb{R}^K$ is the Lagrange multiplier used to enforce the sum-to-one constraint (eq. 7).

After the initialisation of the $\mathbf{W}^0$ and $\mathbf{H}^0$ matrices, they are updated alternately using the equations (17) and (18). For the update of $\mathbf{W}$, an additional step is required to determine

$\boldsymbol{\nu}$ and therefore enforce the sum-to-one constraint. The values of $\boldsymbol{\nu}$ are calculated by solving the following non-linear equation with respect to ν_k :

$$\sum_m^{M'} \max \left(w_{mk}^t \frac{\sigma_{mk}}{\tau_{mk} + \nu_k}, \epsilon \right) = 1. \quad (19)$$

Given that the left-hand-side of the equation above is an increasing function of ν_k , the equation can be solved to high-accuracy by dichotomy.

The successive updates are repeated until a convergence criterion is met: either the maximum number of iterations is reached, or the relative variation of the loss function is small enough. A summary of the operation of the algorithm is shown below in Algorithm 1. A detailed proof that the loss function (16) decreases under the multiplicative updates (17) and (18)) can be found in [20].

Algorithm 1: ESpm-NMF(μ, λ, α)

input : $K, \mathbf{G}, \mu, \lambda, \alpha, tol$

init $L_f^0, H^0 \geq 0, W^0 \geq 0$

while $(L_f^{t+1} - L_f^t) > tol$ **do**

update $\mathbf{W}$

$\forall m, k$, calculate $\sigma_{m,k}$ and $\tau_{m,k}$

dichotomy

$\boldsymbol{\nu} \leftarrow$ equation 19

$\mathbf{W}^{t+1} \leftarrow$ equation 17

update $\mathbf{H}$

$\mathbf{H}^{t+1} \leftarrow$ equation 18

end

output: $\mathbf{W}, \mathbf{H}$

In the following, we denote our physics-guided factorization algorithm, that is implemented in the **espm** package, as ESpm-NMF. The parameters μ , ϵ and λ are considered as hyperparameters of the decomposition; in our notation, they are specified in the order ESpm-NMF(μ, λ, α).

V. MATERIALS AND METHODS

Two Python libraries were developed. The first, `emtables`[22], calculates the characteristic X-ray emission cross-sections. It produces the tables for the data simulations. The second library, `espm`[23], creates the simulated data, the matrix $\mathbf{G}$ and runs the decomposition algorithm. The structure of the algorithm is based on the scikit-learn framework [24, 25], while the user interface relies on the hyperspy library [26].

To test the effectiveness of the ESpm-NMF(μ, λ, α) algorithm – both against other factorization algorithms and in function of chosen hyperparameters and constraints – simulated data were produced using our previously-developed data simulation method [2]. Given that this provides a ground truth to compare against, this therefore allows algorithmic performance to be quantified. The advantages of using our framework to simulate data are twofold: the shape of the ground truth exactly matches with the columns of $\mathbf{G}$, and a full control of the simulation is kept.

The simulated data are inspired by a contemporary geology problem: the mineralogy of Earth’s deep mantle, where the phases Bridgmanite (Brg), Ferropericlasite (Fp), and Ca-perovskite (Ca-Pv) are believed to predominate. The spectra associated to each phase of this problem are the same as the ones that were simulated in our previous work [2]. The abundances of the data in [2] were based on chemical maps of experimental STEM-EDX spectrum images. In the present study, we have taken a different approach, employing small spherical nanoparticles to represent the abundances of Fp and Ca-Pv, while the abundance of the Brg is set to be the matrix that surrounds the particles. The goal of this distribution of phases is to simplify the interpretation of the results of the decomposition. The maximum proportion of the nanoparticles is 0.5 so that there is always a mixing between Brg and the other minerals. This way, it remains challenging to unmix the phases.

All the produced datasets are of dimension $P_x \times P_y \times L = 128 \times 128 \times 1980$. The signal-to-noise ratio (SNR) for a signal that follows a Poisson distribution is $\sqrt{N}$ where N is the intensity of that signal. In this work, we therefore use the total number of counts per spectrum as a simplified indicator of the SNR of the simulated datasets. Two types of datasets are produced: low SNR datasets and high SNR datasets with 18 and 293 counts per spectrum, respectively. The values of N are chosen so that they correspond to a data acquisition of 10 minutes with 500 counts.s⁻¹ or 8000 counts.s⁻¹ on the detector. For each

SNR value, four datacubes were created, each having an identical ground truth but with a different Poisson sampling.

We test the performance of ESpm-NMF against three other algorithms, that themselves are implemented outside of `espm`: NMF, ICA and VCA-SUNSAL. The scikit-learn implementations of both NMF and ICA are used. For NMF, the optimization is done using multiplicative updates to minimise the negative log likelihood loss function [17]. For ICA, a principal component analysis (PCA) pre-processing step is necessary. PCA is performed using the normalization of Keenan et al. [27] to account for the Poisson noise. Then, the data reconstructed with the first three components are fed to the ICA algorithm. The ICA is applied on the data using the fastICA algorithm [5]. For VCA-SUNSAL, the reconstructed spectra are obtained using VCA [6], while the maps are determined using SUNSAL [28] based on the results of VCA. The different algorithms are each applied to all the low and high SNR simulated datasets. For the decomposition of high SNR data, the ESpm-NMF(μ, λ, α) and NMF algorithms are initialized using non-negative double singular value decomposition [29] with zero values filled with small random numbers. For low SNR data, the ESpm-NMF(μ, λ, α) and NMF algorithms are initialized at random. These initializations were chosen to obtain the best optimization results. For ICA, ESpm-NMF(μ, λ, α) and NMF, the same random seed is chosen for the initialization. The version of VCA that is tested in this paper has its own non-random initialization method.

For the results of these algorithms to be compared to the ground truth, their output is first renormalized. The normalized $\mathbf{D}^*$ and $\mathbf{H}^*$ are obtained as follows:

$$\begin{aligned}\mathbf{D}^* &= \mathbf{D}\mathbf{S}^{-1} \\ \mathbf{H}^* &= \mathbf{S}\mathbf{H},\end{aligned}$$

where $\mathbf{S} \in \mathbb{R}^{K \times K}$ is a diagonal matrix that scales $\mathbf{D}$ and $\mathbf{H}$ so that $\sum_k \mathbf{D}^* \approx 1$.

Although visual inspection of the spectra produced by the decomposition may be useful, a metric is necessary to quantify the results. In this paper, two metrics are used to measure the quality of the reconstructed spectra and abundances. For spectra, the angle between the decomposition results and the ground truth is calculated as:

$$\alpha(\mathbf{d}_1, \mathbf{d}_2) = \arccos \left(\frac{\mathbf{d}_1 \cdot \mathbf{d}_2}{\|\mathbf{d}_1\| \|\mathbf{d}_2\|} \right), \quad (20)$$

where $\mathbf{d}_1, \mathbf{d}_2 \in \mathbb{R}^L$. These vectors correspond to columns of $\mathbf{D}$, where typically $\mathbf{d}_1$ and $\mathbf{d}_2$ would correspond to the spectrum of a true phase and to the spectrum of a reconstructed phase, respectively. An angle of 0 deg represents a perfect agreement between the two spectra. For abundance maps, the mean squared errors (MSE) between the reconstructed maps and the true maps are calculated. For each decomposition, the results are quantitatively compared with the ground truth using the mean MSE for the phase maps and the mean angles for the phase spectra.

Following the tests on synthetic data, we demonstrate the validity of ESpm-NMF(μ, λ, α) on experimental data by applying it to data collected by Rossouw et al. [11]. In their work, they did two STEM-EDXS experiments: one on core-shell nanoparticles with an FePt core and a Fe₃O₄ shell and a second one on the FePt bare cores. They were able to show a successful unmixing of the core from the shell using ICA, and they could verify their unmixing procedure using the results on the bare cores. In our work, we use both the bare cores and core-shell data, that are publicly available in the **hyperspy-demos** repository [30] [31] to test the ESpm-NMF(μ, λ, α) algorithm.

VI. RESULTS AND DISCUSSION

We now test the effectiveness and accuracy of the ESpm-NMF(μ, λ, α) algorithm on simulated data, and also compare the quality of decomposition results to those from NMF, ICA and VCA. The series of tests on high SNR data is used to verify whether it is possible to recover the ground truth. On the series with a low SNR, the limits of the decomposition algorithms presented here are tested.

A total of six analyses are performed on each dataset. The first three decompositions are made using VCA, ICA and NMF, which together represent the current state of the art. The remaining decompositions are all made using ESpm-NMF(μ, λ, α), applying different parameters to investigate the effects of the EDXS physical modeling, additional constraints, and regularizations on the decomposition. The first test is performed without additional hyperparameters, and so is labeled “ESpm-NMF(0, 0, 0)”. In some experimental cases, the absence of certain elements in some phases can be deduced from the physical conditions under which the sample was obtained. This is for instance the case for the Fp phase that we simulate here, which cannot contain Si. In this case, the relevant elements of $\mathbf{W}$ can be set

to 0.0 for each phase; this configuration is labeled “ESpM-NMF(0,0,0) fW(0conc)”. The ESpM-NMF(μ, λ, α) algorithm is also tested without prior knowledge but using predefined regularizations; these tests are labeled “ESpM-NMF(0.004, 1, 0.01)”.

Figure 1 plots the mean MSE of the maps versus the mean angles of the spectra for all the analyses. In these plots, the bottom left corner represents the best results, closely matching the ground truth; it is the opposite case for the top right corner.

The results on the high SNR data (figure 1 a)) are discussed first. Both ICA and VCA perform systematically worse than the other algorithms, giving mean angles above 20 deg and mean MSE above 0.2. “Standard” NMF produces more accurate results, with mean angles between 4.5 and 6.7 deg and mean MSE between 0.038 and 0.054. The results of ESpM-NMF(0,0,0) are very similar, with mean angles between 4.1 and 5.3 deg and mean MSE between 0.040 and 0.046. This behaviour can be understood by the fact that NMF and ESpM-NMF(0,0,0) share very similar optimization methods, except for the sum-to-one constraint of the latter on $\mathbf{W}$. It also implies that the incorporation of physical modeling into the decomposition has little influence on average spectral angles in this high SNR case. Nevertheless, later we show that it does help eliminate artifact noise peaks in individual phase spectra solutions. Adding constraints by applying the ESpM-NMF(0,0,0) fW(0conc) configuration leads to an even better agreement, with the mean angles ranging from 3.5 to 3.8 deg and mean MSE ranging from 0.035 to 0.036. The added prior knowledge also tends to stabilize the solutions, since there is a lower spread between the data points of the different synthetic datasets. The ESpM-NMF(0.004, 1, 0.01) configuration exhibits the best results, with mean angles between 3.0 and 3.8 deg and with a tenfold improvement in mean MSE, that ranges from 0.002 to 0.003.

To look more closely at this effect of regularization, Figure 2 shows a comparison between the ground truth and the NMF and ESpM-NMF(0.004, 1, 0.01) decompositions for high SNR data. It can be clearly seen from Figure 2 d) and e) that, even though the three phases are identified by “standard” NMF, there is still some mixing, mainly between the Brg and Fp. On the contrary, for ESpM-NMF(0.004, 1, 0.01), the Laplacian regularization produces uniform maps and the sparsity induced by the log regularization greatly reduces the level of mixing between the phases, as shown in Figure 2 g) and h). The NMF reconstructed spectrum of Figure 2 k) exhibits a noise-related artifact at 1.5 keV, as highlighted in the inset of Figure 2 k). This problem is non-existent for the ESpM-NMF(0.004, 1, 0.01) reconstructed




FIG. 1. a) mean MSE as a function of the mean angle for all the tested algorithms for high SNR datacubes. b) mean MSE as a function of the mean angle for all the tested algorithms for low SNR datacubes. For each algorithm there are four data points, one for each Poisson sampling.

spectrum. Thus, the effect of the physical modeling is similar to an additional noise-reduction filter.

With the weak signal of the low SNR data, it is clear from the angles and MSE in Figure 1 that no algorithm achieves results as close to the ground truth as for high SNR data. Both ICA and VCA show high mean angles above 30 deg. Paradoxically, their mean MSE are low, with values below 0.1. However, these low mean MSE hide the fact that the ICA and VCA algorithms do not actually identify the presence of the three discrete phases. Rather, they each give one abundance map that contains all three phases, with ones everywhere, and two other abundance maps with zeroes everywhere. In comparison, the NMF analysis gives a cluster of mean angles just above 30 deg, and mean MSE above 0.2. Nonetheless, the use of NMF does not result in a precise identification of the phases, as the abundance maps are predominantly randomly populated.

Compared to ICA, VCA and NMF, Figure 1 b) shows that ESpm-NMF gives a markedly improved performance on the low SNR data. All three variants exhibit approximately the same performance as each other, with comparatively low mean angles ranging from 11 deg




FIG. 2. Typical results on a high SNR sample. a), b), c) Ground truth values of the abundance maps of Brg, Fp and Ca-Pv, respectively. d), e), f) Abundance maps resulting from the NMF decomposition. g), h), i) Abundance maps resulting from the ESpM-NMF(0.004,1,0.01) decomposition. j), k) and l) Comparison between the ground truth spectra of Brg, Fp and Ca-Pv and the results of the unmixing algorithms, respectively. In k), the top and bottom insets correspond to the same zoom around the peak of Al-K α as reconstructed using NMF and using ESpM-NMF(0.004,1,0.01), respectively.

to 20 deg and mean MSE between 0.11 and 0.18. ESpM-NMF(0,0,0) shows one outlier at 33.6 deg; we believe that this derives from the random initialisation not being fully reliable. Here, we choose to emphasize the effects of physical modeling on the decomposition results and resulting discrepancies, which we do by comparing the results of NMF and ESpM-NMF(0,0,0) in Figure 3. This figure shows that, for the low SNR data, it is extremely challenging to identify the presence of several phases using NMF. The mixing between the different phases in the three components is so strong that there is no unique way to match them to the ground truth phases. On the contrary, using the ESpM-NMF(0,0,0) algorithm, it is directly possible to recognize the three phases, even without the help of adding prior




FIG. 3. Typical results on a low SNR sample. a), b), c) Ground truth values of the abundance maps of Brg, Fp and Ca-Pv, respectively. d), e), f) Abundance maps resulting from the NMF decomposition. g), h), i) Abundance maps resulting from the ESpm-NMF(0,0,0) decomposition. j), k) and l) Comparison between the ground truth spectra of Brg, Fp and Ca-Pv and the results of the unmixing algorithms, respectively.

knowledge or regularizations to the algorithm.

Since the only two differences between ESpm-NMF(0,0,0) and NMF are the sum-to-one constraint and the physical modeling, this improvement must derive from one or both of these factors. To find out which, the ESpm-NMF(0,0,0) decomposition was tested with $\mathbf{G}$ equals the identity matrix, thus removing the modeling of X-rays. A solution very similar to that of NMF was then obtained. The results therefore demonstrate that including physics in the machine learning process is an effective strategy for analysing low SNR STEM-EDXS data.

While special attention is paid to simulating EDXS data as accurately as possible, the simulation cannot fully encompass all of the features of a real experimental dataset. Therefore, we also test the ESpm-NMF(0,0,0) algorithm on experimental data. We choose to analyze data on FePt@Fe₃O₄ nanoparticles collected by Rossouw et al. [11], since they represent an interesting unmixing problem. Assigning the correct quantities of Fe to the cores and shells is not straightforward in this case. Furthermore, thanks to their analysis of the bare cores, a ground truth is available in Rossouw et al.’s work. This enables us to determine whether our algorithm converges towards physically accurate results, similar to those found in their work that they obtained using ICA.

In their study, Rossouw et al. [11] determined that the bare cores consist, on average, of 18 at. % of Fe and 82 at. % of Pt, with a margin of error of ± 3 at. %. Additionally, they established that the FePt@Fe₃O₄ nanoparticles exhibit a core-shell configuration. This configuration implies that, everywhere a particle exists, the proportion of the shell never drops to 0. A successful decomposition using ESpm-NMF(0,0,0) should achieve two key results: an accurate composition of the cores and a correct spatial distribution of the shells.

Here, we first perform a decomposition without the use of prior knowledge on the chemical composition of the sample. Its outcomes are shown in Figure 4 under the label “No constraint”. While the support (Figure 4 a)) is well separated from the particles, the shells (Figure 4 c)) appear perforated. Around the core, some Fe belonging to the shells is instead attributed to the FePt cores, leading to the latter having an overestimated spatial extent and overestimated Fe content.

This effect can be seen in Figure 4 h), where the Fe K α peak is very intense. Hence, the built-in quantification of the ESpm-NMF(0,0,0) decomposition yields a composition of the core of 45 at. % of Fe and 55 at. % of Pt. Our method, without external help from the experimentalist, fails to accurately retrieve the true distribution of phases within this dataset. In this configuration, the performance of the ESpm-NMF(0,0,0) algorithm is very similar to that of NMF, which also fails to converge towards the correct solution.

Nevertheless, thanks to the modeling presented in section II, we can add constraints on the values of $\mathbf{W}$ that will guide the decomposition process. To produce our own prior knowledge of the sample, we analyze the bare cores dataset with our method and find that the cores are composed of 16 at. % of Fe and 84 at. % of Pt. These proportions are very close to the ratio measured by Rossouw et al. [11] on both the bare cores and on the core-




FIG. 4. Results of the ESpM-NMF(0,0,0) decomposition on experimental data. a), b), c) Abundance maps obtained using the unconstrained ESpM-NMF(0,0,0) of the support, cores and shells, respectively. d), e), f) Abundance maps obtained when introducing prior knowledge about the sample, of the support, cores and shells, respectively. g), h), and i) Comparison between the spectra resulting from the constrained (black line) and unconstrained (blue dots) ESpM-NMF(0,0,0) decomposition. j) Comparison between the mean of the experimental spectra (blue line) and the mean of the reconstructed spectra (red dots).

shell nanoparticles. All the calculations were performed with tabulated theoretical X-ray emission cross sections [2] which should ideally be tuned using standards. Small differences in cross section values between our work and those of Rossouw et al. [11] may well explain the slight discrepancy between our bare core composition and theirs.

Then, for the next decomposition of the core-shell dataset, we set the values of a part of the elements of $\mathbf{W}$ so that the phase corresponding to the cores has a Fe/Pt ratio corresponding to that obtained previously on the bare cores. The results are displayed in Figure 4 under the label “constraint”. The absence of holes in the calculated spatial distribution of the shells indicates a successful unmixing of the FePt@Fe₃O₄ nanoparticles dataset. This is confirmed by the reduced intensity of the Fe K α peak in Figure 4 h). To further assess the quality of the decomposition, we compute the chemical composition of the shells. We find a composition of 40 at. % of Fe and 60 at. % of O which is close to the expected 43 at. % of Fe and 57 at. % of O of Fe₃O₄. We obtain this result by adjusting the mass thickness (ρ^*) of the model with a thickness of 30 nm and a density of 5.2 g.cm⁻³.

Finally, to ensure that our constrained modeling approach has not led to any important features of the data being cut during decomposition, we compare the sum of experimental spectra to the sum of the reconstructed spectra. The results are displayed in Figure 4 j), and exhibit a good agreement between the model and the experiment. Therefore, while standard NMF and ESpm-NMF(0, 0, 0) fail to provide a satisfactory unmixing of the FePt@Fe₃O₄ nanoparticles, by incorporating prior knowledge into the ESpm-NMF(0, 0, 0) decomposition, it is possible to accurately separate the cores from the shells. As well as demonstrating how prior knowledge can greatly enhance the fidelity of decomposition, and together with the algorithm testing on simulated datasets, this result emphasizes the importance of **espm** providing a versatile decomposition toolbox, whose parameters can be appropriately tuned or optimized for the problem of interest.

VII. CONCLUSION

An open source Python software tool called “**espm**” has been developed for the simulation and analysis of STEM-EDXS data. The part of **espm** responsible for the simulation was described in our previous work [2]. In this work, we develop a physics-guided NMF-based algorithm to decompose experimental STEM-EDXS data into phases and their spatial dis-

tributions. First, we designed a linear factorization model of STEM-EDXS spectrum images that integrates the physics modeling into NMF analysis. An additional advantage of using this model is that a direct elemental quantification of each phase is inherently made during the decomposition process. We presented here the optimization method to perform the model-based NMF, as derived from the alternating multiplicative updates algorithm of Lee and Seung [4]. We also added the possibility for the user to apply regularizations to further improve the analysis.

The effectiveness of the developed algorithm was tested on both simulated and experimental data. On the simulated data, it was shown that our method quantitatively outperforms the other tested decomposition algorithms. On data with a high SNR, the ESpm-NMF(0,0,0) algorithm shows similar results as the common implementation of NMF but, with the help of regularization, the ESpm-NMF(0.004,1,0.01) results match the ground truth almost perfectly. On low SNR data, while standard NMF fails to separate the data into several phases, thanks to the modeling it is possible to identify the different phases using ESpm-NMF. Concerning experimental data, by using constraints on the model, we could successfully reproduce the results of Roussouw et al. [11].

NMF is a statistical decomposition method which makes it very flexible, such that it can be applied on a very broad range of subjects. However, since the result is purely statistical, its use requires an interpretation stage. For example, for STEM-EDXS data, standard NMF does not usually produce EDX spectra directly, and at least one quantification step is required. Our approach is to linearize the underlying physical model of EDX and then integrate this linear model into the NMF decomposition. Thanks to this modeling, the decomposition produces directly interpretable spectra, and the model can be constrained to obtain results closer to physical reality. We believe that this linearization method is adaptable to other areas of physics, and hence can be applied to a wide range of decomposition problems.

VIII. ACKNOWLEDGMENTS

This project was funded by the Swiss Data Science Center under the SDSC grant C19-07.

Appendix A: Parametrization of the breaking radiation.

The modeling of bremsstrahlung has been investigated by several authors; a review of the proposed phenomenological models can be found in the work of Small et al. [32]. All of the proposed models give a decreasing bremsstrahlung intensity with increasing energy, either with a linear, quadratic or exponential shape. The most recent theoretical work shows that a second order polynomial is sufficient to describe the behaviour of bremsstrahlung for energies above 1 keV [14, 33]. To integrate the bremsstrahlung in the factorization model developed in this work, its parameterization has to be linear and non-negative.

In the following, we show how the bremsstrahlung equation from [34]:

$$b'(\varepsilon_\ell) = \gamma'_0 \frac{e_0 - \varepsilon_\ell}{\varepsilon_\ell} + \gamma'_1 \frac{(e_0 - \varepsilon_\ell)^2}{e_0 \varepsilon_\ell} \quad (\text{A1})$$

can be reparametrized for a new model ($b(\varepsilon_\ell)$) with only positive coefficients γ_0, γ_1 as used in (4). As the bremsstrahlung is positive (physical constraint), we need to satisfy:

$$\forall \ell, b'(\varepsilon_\ell) \geq 0 \quad \text{and} \quad 0 \leq \varepsilon_\ell \leq e_0, \quad (\text{A2})$$

which puts some constraints on γ'_0, γ'_1 . To formulate these constraints, we simplify (A1). First, since $\varepsilon_\ell \geq 0$, we can multiply $b'(\varepsilon_\ell)$ by ε_ℓ without changing its sign. Second, we introduce the change of variable $\varepsilon'_\ell = e_0 - \varepsilon_\ell$. We note that $\varepsilon'_\ell \geq 0$ for $\varepsilon_\ell \in [0, e_0]$. After factorizing $e_0 - \varepsilon_\ell$, we obtain:

$$\dot{b}(\varepsilon'_\ell) := \frac{\varepsilon_\ell}{\varepsilon'_\ell} b'(\varepsilon_\ell) = \gamma'_0 + \frac{\gamma'_1}{e_0} \varepsilon'_\ell \geq 0 \quad \forall \ell, 0 \leq \varepsilon'_\ell \leq e_0. \quad (\text{A3})$$

Since $\dot{b}$ is linear in ε'_ℓ , ensuring that $\dot{b}(0) \geq 0, \dot{b}(e_0) \geq 0$ forces $\dot{b}(\varepsilon'_\ell) \geq 0$ for any energy ε'_ℓ between 0 and e_0 . Therefore, we finally obtain two constraints:

$$\gamma'_0 \geq 0 \quad \text{and} \quad \gamma'_0 + \gamma'_1 \geq 0. \quad (\text{A4})$$

from $\dot{b}(0) = \gamma'_0 \geq 0$ and $\dot{b}(e_0) = \gamma'_0 + \gamma'_1 \geq 0$.

Our next step is to search for another parameterisation γ_0, γ_1 such that $\gamma'_0 \geq 0, \gamma'_1 \geq 0$ for any value of γ'_0, γ'_1 respecting the constraint of (A4). Given (A4), a trivial solution is

$$\gamma_1 = \gamma'_0 + \gamma'_1 \geq 0 \quad \text{and} \quad \gamma_0 = \gamma'_0 \geq 0.$$

In this new parameterisation, we have $\gamma'_1 = \gamma_1 - \gamma_0$, which allow us to rewrite (A1) as

$$b'(\varepsilon_\ell) = \gamma_0 \left(\frac{e_0 - \varepsilon_\ell}{\varepsilon_\ell} - \frac{(e_0 - \varepsilon_\ell)^2}{\varepsilon_\ell e_0} \right) + \gamma_1 \frac{(e_0 - \varepsilon_\ell)^2}{\varepsilon_\ell e_0}$$

Without loss of modeling power, we further transform $b'(\varepsilon_\ell)$ into:

$$\begin{aligned} b(\varepsilon_\ell) &= \frac{b'(\varepsilon_\ell)}{e_0} \\ &= \gamma_0 \frac{e_0 - \varepsilon_\ell}{\varepsilon_\ell e_0} \left(1 - \frac{e_0 - \varepsilon_\ell}{e_0} \right) + \gamma_1 \frac{(e_0 - \varepsilon_\ell)^2}{e_0^2 \varepsilon_\ell}, \end{aligned} \quad (\text{A5})$$

which gives us the two terms of equation (4).

For simplicity, we did not address the fact that $\lim_{\varepsilon_\ell \rightarrow 0} b(\varepsilon_\ell) = \infty$, the previous reasoning could still be established rigorously taking this fact into account.

Appendix B: Updates computation

In their work, Perraudin et al. [20] introduce novel efficient algorithms for solving the regularized Poisson non-negative factorization problem. The key approach involves the utilization of the block successive upper minimization (BSUM) algorithm [35]. This algorithm operates by successively minimizing upper bounds of the objective function for each block of variables. Specifically, in the context of regularized Poisson NMF, the two blocks of variables involved are denoted as $\mathbf{W}$ and $\mathbf{H}$. Ensuring convergence requires that these upper bounds satisfy a crucial property: they must be majorizing functions. Additional details can be found in [20, 35]. To formulate a concrete algorithm, Perraudin et al. [20] construct appropriate majorizing functions tailored to the regularizers employed in their contribution. Subsequently, they compute the minimiser of the majorizing function efficiently, providing a formula for the iterative updates [20, Proposition 1 and eq. 22].

The formulation of [20, eq. 10] proposes a minimisation of a general loss function with the three different type of regularisations, Lipchitz, relatively smooth and concave functions (noted s_R, s_L and s_C), as well as linear constraint (noted $\mathbf{e}^T \mathbf{x} = \mathbf{1}$). In the present work, the overall loss function we aim to minimise is expressed as follows:

$$\begin{aligned} L_f(\mathbf{W}, \mathbf{H}) &= L_P(\mathbf{W}, \mathbf{H}) + r_\Delta(\mathbf{H}) + r_{\log}(\mathbf{H}) \\ &= \sum_{\ell, p}^{L, P} [-\mathbf{Y}_{\ell p} \log(\mathbf{G}\mathbf{W}\mathbf{H})_{\ell p} + (\mathbf{G}\mathbf{W}\mathbf{H})_{\ell p}] \\ &\quad + \frac{\lambda}{2} \text{tr}(\mathbf{H}\Delta\mathbf{H}^T) \\ &\quad + \mu \sum_{k, p}^{K, P} \log(h_{kp} + \alpha) \end{aligned} \quad (\text{B1})$$

Following the nomenclature and colors proposed in [20], we categorize the loss into three distinct parts corresponding to the Poisson loss (in green), a Lipschitz function (in orange), and a strictly concave function (in purple). The Lipschitz and strictly concave functions correspond to s_L and s_C respectively in [20, eq. 10]. The ESpm-NMF algorithm can be effectively customised with any regularisation that follows the assumption presented in [20].

According to [20, eq. 16,17,18 of Proposition 1], we derive an update for $\mathbf{H}$ as follows:

$$\begin{aligned} \mathbf{H}^{t+1} = & \mathbf{H}^t \\ & \odot \left(\mathbf{W}^{tT} \mathbf{G}^T (\mathbf{Y} \oslash (\mathbf{G} \mathbf{W}^t \mathbf{H}^t)) + 8\lambda \text{max}(\mathbf{H}) \right) \\ & \oslash \left(\mathbf{W}^{tT} (\mathbf{G}^T \mathbf{1}^{L,P}) + \mu \oslash (\mathbf{H}^t + \alpha) + \right. \\ & \left. 8\lambda \text{max}(\mathbf{H}) + \lambda \mathbf{H} \Delta \right). \end{aligned} \tag{B2}$$

To derive this update, we use the fact that $\nabla r_\Delta(\mathbf{H}) = \lambda \Delta \mathbf{H}$ and $\sigma_L = 4$ since it is the maximum eigenvalue of the two dimensional Laplacian matrix Δ . Additionally, we note that $\frac{\partial}{\partial h_{kp}} \mu \log(h_{kp} + \alpha) = \frac{\mu}{h_{kp} + \alpha}$. Since there is no simplex constraint, to enforce the positivity constraint $\mathbf{H} \geq \epsilon$, we use [20, eq. 22] with $\mathbf{e} = 0$ and obtain the update (18).

Similarly, for $\mathbf{W}$, we obtain from [20, eq. 16,17,18 of Proposition 1]:

$$\begin{aligned} \mathbf{W}^{t+1} = & \mathbf{W}^t \odot \left(\mathbf{G}^T (\mathbf{Y} \oslash \mathbf{G} \mathbf{W}^t \mathbf{H}^t) \mathbf{H}^{tT} \right) \\ & \oslash \left(\mathbf{G}^T \mathbf{1}^{L,P} \mathbf{H}^{tT} \right) \end{aligned} \tag{B3}$$

The simplex constraint from (16) reads $\sum_{m \leq M-2} \mathbf{W}_m = \mathbf{1} = \mathbf{u}^T \mathbf{W}$ where $\mathbf{u} = (u_m) \in \mathbb{N}^M$ is defined as a vector with $u_m = 1, m \in 1, \dots, M', u_{M-1} = 0$, and $u_M = 0$. We obtain the update (17) by using [20, eq. 22] with $\mathbf{e} = \mathbf{u}$. Following the methodology of [20], we obtain the value of the Lagrange multiplier $\boldsymbol{\nu}$ using dichotomy to solve Eq. (19).

Finally, following the iterative application of these updates in [20, Algorithm 2] leads to our proposed Algorithm 1.

-
- [1] A. J. D'Alfonso, B. Freitag, D. Klenov, and L. J. Allen, Atomic-resolution chemical mapping using energy-dispersive x-ray spectroscopy, *Physical Review B* **81**, 100101 (2010).
 - [2] A. Teurtrie, N. Perraudin, T. Holvoet, H. Chen, D. T. Alexander, G. Obozinski, and C. Hébert,

- espm: A python library for the simulation of stem-edxs datasets, *Ultramicroscopy* **249**, 113719 (2023).
- [3] J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot, Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **5**, 354 (2012).
 - [4] D. D. Lee and H. S. Seung, Learning the parts of objects by non-negative matrix factorization, *Nature* **401**, 788 (1999).
 - [5] A. Hyvärinen and E. Oja, Independent component analysis: algorithms and applications, *Neural Networks* **13**, 411 (2000).
 - [6] J. Nascimento and J. Dias, Vertex component analysis: a fast algorithm to unmix hyperspectral data, *IEEE Transactions on Geoscience and Remote Sensing* **43**, 898 (2005).
 - [7] F. B. Lavoie, N. Braidý, and R. Gosselin, Including noise characteristics in MCR to improve mapping and component extraction from spectral images, *Chemometrics and Intelligent Laboratory Systems* **153**, 40 (2016).
 - [8] S. Cacovich, F. Matteocci, M. Abdi-Jalebi, S. D. Stranks, A. Di Carlo, C. Ducati, and G. Divitini, Unveiling the chemical composition of halide perovskite films using multivariate statistical analyses, *ACS Applied Energy Materials* **1**, 7174 (2018).
 - [9] L. Henry, N. Bernier, M. Jacob, G. Navarro, L. Clément, J.-L. Rouvière, and E. Robin, Studying phase change memory devices by coupling scanning precession electron diffraction and energy dispersive x-ray analysis, *Acta Materialia* **201**, 72 (2020).
 - [10] B. R. Jany, A. Janas, and F. Krok, Retrieving the quantitative chemical information at nanoscale from scanning electron microscope energy dispersive x-ray measurements by machine learning, *Nano Letters* **17**, 6520 (2017).
 - [11] D. Rossouw, P. Burdet, F. de la Peña, C. Ducati, B. R. Knappett, A. E. H. Wheatley, and P. A. Midgley, Multicomponent signal unmixing from nanoheterostructures: Overcoming the traditional challenges of nanoscale x-ray analysis via machine learning, *Nano Letters* **15**, 10.1021/acs.nanolett.5b00449 (2015).
 - [12] M. Shiga, K. Tatsumi, S. Muto, K. Tsuda, Y. Yamamoto, T. Mori, and T. Tanji, Sparse modeling of EELS and EDX spectral imaging data by nonnegative matrix factorization, *Ultramicroscopy* **170**, 43 (2016).

- [13] C.-Y. Hsu, J. Stodolna, P. Todeschini, F. Delabrouille, B. Radiguet, and F. Christien, Accurate quantification of phosphorus intergranular segregation in iron by stem-edx, *Micron* **153**, 103175 (2022).
- [14] J. N. Chapman, W. A. P. Nicholson, and P. A. Crozier, Understanding thin film x-ray spectra, *Journal of Microscopy* **136**, 179 (1984).
- [15] M. Watanabe and D. B. Williams, The quantitative analysis of thin specimens: a review of progress from the cliff-lorimer to the new zeta-factor methods, *Journal of Microscopy* **221**, 89 (2006).
- [16] R. F. Egerton, *Electron energy-loss spectroscopy in the electron microscope* (Springer Science & Business Media, 2011).
- [17] D. D. Lee and H. S. Seung, Algorithms for non-negative matrix factorization, in *Advances in Neural Information Processing Systems*, Vol. 13 (MIT Press, 2001).
- [18] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, Optimization with sparsity-inducing penalties, *Foundations and Trends in Machine Learning* **4**, 1 (2011).
- [19] F. R. Chung and F. C. Graham, *Spectral graph theory*, 92 (American Mathematical Soc., 1997).
- [20] N. Perraudin, A. Teutrie, C. Hébert, and G. Obozinski, Efficient algorithms for regularized poisson non-negative matrix factorization (2024), arXiv:2404.16505 [cs.LG].
- [21] L. T. K. Hien and N. Gillis, Algorithms for nonnegative matrix factorization with the kullback-leibler divergence, *Journal of Scientific Computing* **87**, 93 (2021), 2010.01935.
- [22] Publicly available at <https://github.com/adriente/emtables>.
- [23] Publicly available at <https://espm.readthedocs.io/en/latest/>.
- [24] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. Vanderplas, A. Joly, B. Holt, and G. Varoquaux, API design for machine learning software: experiences from the scikit-learn project, arXiv:1309.0238 [cs] (2013), 1309.0238.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine learning in python, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [26] F. De La Peña, E. Prestat, V. T. Fauske, P. Burdet, J. Lähnemann, T. Furnival,

- P. Jokubauskas, M. Nord, T. Ostasevicius, K. E. MacArthur, D. N. Johnstone, M. Sarahan, T. Aarholt, J. Taillon, Pquinn-Dls, V. Migunov, A. Eljarrat, J. Caron, T. Poon, S. Mazzucco, C. Francis, B. Martineau, Actions-User, S. Somnath, T. Slater, N. Tappy, M. Walls, N. Cautaerts, F. Winkler, and DENSmerijn, hyperspy/hyperspy: Release v1.6.5 (2021).
- [27] M. R. Keenan and P. G. Kotula, Accounting for poisson noise in the multivariate analysis of ToF-SIMS spectrum images, *Surface and Interface Analysis* **36**, 203 (2004).
- [28] J. M. Bioucas-Dias and M. A. T. Figueiredo, Alternating direction algorithms for constrained sparse regression: Application to hyperspectral unmixing, in *2010 2nd Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing* (IEEE, 2010) pp. 1–4.
- [29] C. Boutsidis and E. Gallopoulos, SVD based initialization: A head start for nonnegative matrix factorization, *Pattern Recognition* **41**, 1350 (2008).
- [30] Link to the github repository: <https://github.com/hyperspy/hyperspy-demos>.
- [31] Direct dropbox link to the data:
https://www.dropbox.com/s/ecdlgwqxjq04m5mx/HyperSpy_demos_EDS_TEM_files.zip?raw=1.
- [32] J. A. Small, S. D. Leigh, D. E. Newbury, and R. L. Myklebust, Modeling of the bremsstrahlung radiation produced in pure-element targets by 10-40 keV electrons, *Journal of Applied Physics* **61**, 459 (1987).
- [33] J. N. Chapman, C. C. Gray, B. W. Robertson, and W. A. P. Nicholson, X-ray production in thin films by electrons with energies between 40 and 100 keV. 1-bremsstrahlung cross-sections, *X-Ray Spectrometry* **12**, 153 (1983).
- [34] E. Lifshin, The use of solid state x-ray detectors for obtaining fundamental x-ray data, in *Proc. 9th Annual Conf. Microbeam Analysis Society*, Vol. 53 (1974).
- [35] M. Razaviyayn, M. Hong, and Z.-Q. Luo, A unified convergence analysis of block successive minimization methods for nonsmooth optimization, *SIAM Journal on Optimization* **23**, 1126 (2013).