Scaling Non-Negative Matrix Factorisation for Electron-Microscopy

Nathanaël Perraudin, Adrien Teurtrie and Guillaume Obozinski

January 2022

1 Summary

Electron microscopy is a useful tool to probe the chemistry of materials from the micrometer scale down to the atomic scale. Such information are key for the design of new materials. It is especially true for complex materials composed of several phases with different chemistry.

In electron microscopy, the chemistry of materials is measured using a technique called spectrum imaging (see Figure 1). The electron beam is scanned over the sample and an X-ray spectrum is acquired at each position of the scan. Each spectrum conveys some information about the local chemical composition in the sample. The collection of spectra forms a 3D datacube $\boldsymbol{Y}$ with two spatial axes and a spectral one. Two main problems are faced here:

- Each spectrum is crippled by the measurement noise

- The chemical information of the different phases composing the probed material are mixed.

In this project, the aim is to solve the unmixing problem i.e. retrieve the chemical composition of each pure phase and their respective spatial maps. This inverse problems consists of finding $\boldsymbol{W}$ and $\boldsymbol{H}$ such that:

$$\boldsymbol{W}\boldsymbol{H} \approx \boldsymbol{Y} \tag{1}$$

Where $\boldsymbol{W}$ is the matrix which columns are the spectra of each pure phase and $\boldsymbol{H}$ is the corresponding maps.

To recover $\boldsymbol{W}$ and $\boldsymbol{H}$ from $\boldsymbol{Y}$, we have recently developed an algorithm outperforming state of the art algorithms. However, our initial version is still too slow for daily operation preventing wide adoption by the community. The main goal of this project is to accelerate the algorithm, allowing easier routine usage for the experimentalist. This project is research oriented and can lead to a scientific publication.

2 Project description

As described above, the recovery of the different maps $\boldsymbol{H}$ and their spectra $\boldsymbol{W}$ from the measurement $\boldsymbol{Y}$ involve solving a particular inverse problem known as Non-Negative Matrix Factorisation (NMF) [1]. In the context of this project, we are interested in a regularised version of the Poisson NMF problem. More specifically, the problem is formulated as

$$\dot{\boldsymbol{W}}, \dot{\boldsymbol{H}} = \underset{\substack{\boldsymbol{W} \geq 0, \boldsymbol{H} \geq 0 \\ \mathbf{1}_K^T \boldsymbol{H} = \mathbf{1}_P^T}}{\arg\min} F(\boldsymbol{W}, \boldsymbol{H}, \boldsymbol{Y}) + R_1(\boldsymbol{H}, \lambda) + R_2(\boldsymbol{H}, \mu) \tag{2}$$
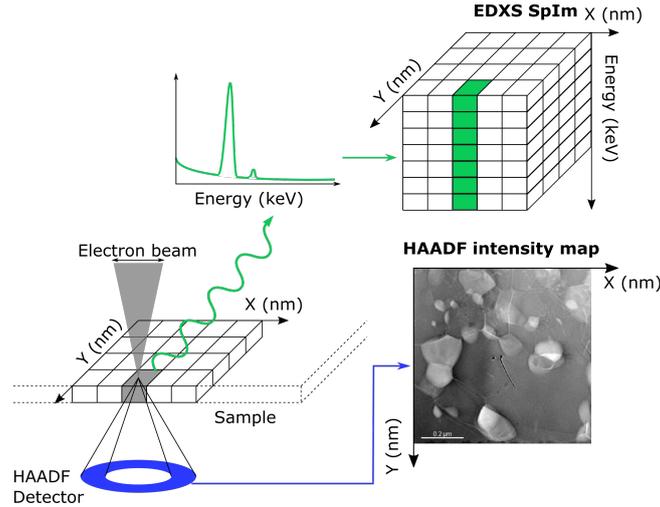
Figure 1: *Overview of the acquisition process of electron microscopy data. An electron beam is scanned over the observed sample. At each pixel of the scan a spectrum (of EDX spectroscopy) is measured. The full collection of spectra forms a 3D datacube. An image (called HAADF) of the sample is also acquired simultaneously.*

where the solution $\dot{\boldsymbol{W}}, \dot{\boldsymbol{H}}$ corresponds to the recovered phase and spectra. The loss can be decomposed in three terms. The fidelity term

$$F(\boldsymbol{W}, \boldsymbol{H}) = \sum_{\ell}^{L} \sum_{p}^{P} -\boldsymbol{Y}_{\ell,p} \cdot \log \left([\boldsymbol{G}\boldsymbol{W}\boldsymbol{H}]_{\ell,p} + \eta\right) + [\boldsymbol{G}\boldsymbol{W}\boldsymbol{H}]_{\ell,p}$$

ensures that $\boldsymbol{W}\boldsymbol{H} \approx \boldsymbol{Y}$ with the assumption that the noise follow a Poisson distribution. The first regularisation term $R_1(\boldsymbol{H}, \lambda) = \frac{\lambda}{2}\text{tr}\left(\boldsymbol{H}^T \Delta \boldsymbol{H}\right)$, where Δ is the Laplacian, ensures that the obtained maps varies smoothly. The second regularisation term favor large values over small one and push the solution to be sparse $R_2(\boldsymbol{H}, \mu) = \mu^T \log(\boldsymbol{H} + \epsilon)\boldsymbol{1}^p$. Using both regularisation terms R_1 and R_2 simultaneously has been instrumental into reducing the noise and has lead to results significantly better than the state of the art.

To obtain these good results, the hyper-parameters μ and λ still needs to be manually tuned. This process is costly both in human labor and in computation. Therefore, the goal of this project is to address this challenge by 1. reducing the overall computation of solving the optimisation problem (2), and 2. searching for automatized techniques to set the regularisation parameters μ and λ. To achieve this goal, we plan to reduce the data dimension and solve a smaller version of the problem. This allows greatly reduce the computation time allowing for easier parameter tuning. Once a good solution is found, we upscale it back to the original dimension and perform some final fine-tuning operation. Figure 2 illustrates the general pipeline of the acceleration.

3 Additional information

- **Difficulty of the project:** Challenging (but we guide you all along)!

- **Is this project theoretical:** If you wish so. As we use traditional optimisation techniques, there is plenty of theory to be used and developed. However, this project is also very practical and theoretical development are completely optional.
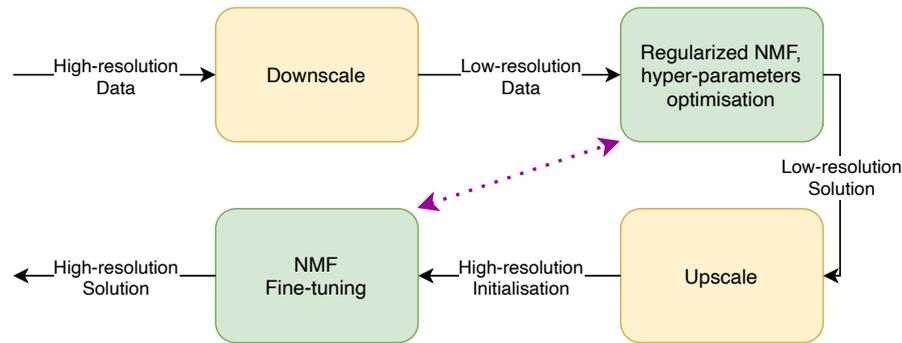
Figure 2: *Proposed approach. 1. The high-resolution data is first downscaled to a lower dimension reducing both the noise and the amount of data to be processed. 2. The regularized NMF algorithm is applied allowing on the low resolution data allowing for easy parameter tuning. 3. The low-resolution is upsampled back to the original dimension providing a good initialisation point for the NMF algorithm. 4. The final solution is obtained with some fine-tuning.*

- **What will you learn?** Machine Learning, optimisation, non negative matrix factorisation, good scientific research practices, literature research.

- **Requirements:** Machine Learning fundamentals, linear algebra, optimization fundamentals, good Python skills, experience with git, *motivation*.

- **Supervisors:**

  - Nathanaël Perraudin (nathanael.perraudin@sdsc.ethz.ch),
  - Adrien Teurtrie (adrien.teurtrie@epfl.ch),
  - Guillaume Obozinski (guillaume.obozinski@epfl.ch).

Bibliography

[1] D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Advances in Neural Information Processing Systems*, T. Leen, T. Dietterich, and V. Tresp, Eds., vol. 13, MIT Press, 2001, pp. 556–562.