

# Scalable multicomponent spectral analysis for high-throughput data annotation

R. Patrick Xian<sup>1,2\*</sup>   Ralph Ernstorfer<sup>1</sup>   Philipp M. Pelz<sup>3,4</sup>

<sup>1</sup>Fritz Haber Institute of the Max Planck Society, 14195 Berlin, Germany.

<sup>2</sup>Department of Neurobiology, Northwestern University, Evanston 60208, IL, USA.

<sup>3</sup>National Center for Electron Microscopy Facility, Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.

<sup>4</sup>Department of Materials Science and Engineering, University of California, Berkeley, CA 94720, USA.

\*Correspondence authors: xian@fhi-berlin.mpg.de

**Orchestrating parametric fitting of multicomponent spectra at scale is an essential yet underappreciated task in high-throughput quantification of materials and chemical composition. We present a systematic approach compatible with high-performance computing infrastructures using the MapReduce model and task-based parallelization. Our approach is realized in a software, *pesfit*, to enable efficient generation of high-quality data annotation and online spectral analysis as demonstrated using experimental materials characterization datasets.**

## Introduction

Real-time understanding of experimental data in materials and chemical characterization is a pursuit for the coming age of automation [1]. Building machine learning models for experimental techniques that produce spectral or spectrum-like data requires tools for *spectrum annotation*, which pairs data with quantitative information extracted by fitting approximate lineshape or structural models [2]. Achieving computational efficiency in these tasks allows processing

large amounts of data with negligible human intervention as well as execution of online or on-the-fly analysis. In materials characterization measurements, changes in sample and environmental conditions could appear in minutes to hours, while experimental campaigns often run continuously for days or longer. Moreover, the data acquisition process often involves tuning of external variables [3] such as the spatial locations, time, temperature, concentration, etc, resulting in multidimensional datasets of multicomponent spectra (each component being a spectral lineshape or background model). As a realistic example, fitting a thousand 10-component spectra, with 3-4 parameters for each component would result in a large-scale optimization problem involving on the order of  $10^4$  parameters. To assess the experimental outcome on this scale is not uncommon in modern high-throughput experimentation [4], therefore, efficient computational scaling with respect to spectral complexity, as quantified by the number of components, becomes crucial for automating materials characterization.

In this work, we leverage high-performance computing (HPC) to achieve linear scaling for a class of spectrum annotation task – multicomponent spectral analysis (MSA) [5] or fitting. Broadly speaking, the MSA problem encompasses two complementary scenarios: (1) Fitting an unknown number of peaks at fixed spectral locations with varying amplitudes; (2) Fitting a fixed number of peaks at varying spectral locations. Ultimately, MSA extracts information such as the shape and height of spectral components and disentangles signal from background. Mathematically, MSA may be mapped to separable nonlinear least squares problems [6] present in various branches of science and engineering. In analytical chemistry, where scenario (1) is frequently encountered for understanding spectroscopic data, matrix factorization techniques without [7] or with limited prior knowledge [8] of pure spectral components (i.e. spectra from single sources or species) are often adopted for data analysis. In comparison, scenario (2) remains a recurring challenge in upscaling the analysis pipeline of quantitative electron- and X ray-based characterization techniques, when large, annotated, empirical databases are yet to exist for model training. Moreover, the equivalents in scenario (2) of pure spectral components are generally unattainable. Currently, relevant softwares for materials characterization usually only provide the option for sequentially fitting isolated spectra [9] or pattern in the case of diffraction [10]. Therefore, we have developed `pesfit` [11], a software to tackle the MSA problem in (2) at scale to facilitate online and offline analysis and experimental feedback. In the following, we discuss the software architecture, validate the computational scaling on batch spectrum fitting and present use cases of spectrum annotation in materials characterization.

## Results

**Scalable software architecture.** Fitting a multicomponent spectrum requires iterative calculation of its components, which are typically modelled as unimodal probability distributions or analytical functions. The mathematical context is provided in Methods. In practice, a complex spectrum may comprise tens of components or may be divided into segments of a similar size. Maintaining a reasonable computational performance with a large number of spectra with increasing complexities requires both resource and architectural upgrade on almost all steps of single-spectrum fitting that remains the mainstay in existing data analysis workflows [9, 10]. We have realized linear scaling in `pesfit` to satisfy the computational needs and alleviate the




Figure 1: **Schematic of the scalable software architecture.** Each fitting task for a multicomponent spectrum (left) requires iterative evaluation of an approximate spectral lineshape model. A MapReduce approach implemented in `MultipeakModel` (see Methods) is used to execute the calculation of all components. During optimization, an initialization generator may add user-specified perturbations to the initial conditions to refine the outcome, guided by an error function (such as goodness-of-fit metrics) involving the data and the fit. The entire pool of fitting tasks (middle) formed by aggregating individual ones is distributed over the multicore CPUs in a standard computing infrastructure (right) and different task schedulers may be used. Task-based parallelization with the asynchronous scheme shown here (right) allows to balance the computational loads between fitting tasks.

manual task load in (1) multicomponent spectrum computation, (2) batch execution, as well as the associated steps of (3) initial condition generation and (4) fitting result collection. To minimize redundancy in (1), we constructed a generating class (`MultipeakModel`, see Methods) to represent multicomponent spectra, making use of a MapReduce structure [12] for evaluating the spectra given the constituents (see Fig. 1 left). Our approach extends the popular spectral

lineshape fitting framework `lmfit` [13] to allow batch fitting a multitude of spectra with an arbitrary number of components. The MapReduce approach splits up (“map” step) the evaluation of multicomponent spectra into parallelizable computation of the components followed by merging (“reduce” step) of the results (see Fig. 1 left) and allows to generate an arbitrary complex spectral shape from existing lineshape models or user-defined functions without hard-coding each type of multicomponent spectrum separately. Each task thus constructed executes fitting of a single spectrum (see Eq. (2) in Methods). To enable (2), the compartmentalized fitting tasks (see Fig. 1 middle) are executed concurrently through `dask` [14], `parmap` [15], and `torcpy` [16] on an HPC infrastructure with multicore CPUs (see Fig. 1 right). Resource sharing (including the numerical data and the multicomponent spectrum model) between tasks is implemented in distributed fitting (`DistributedFitter`, see Methods) to reduce the memory footprint. To simplify (3), the `pesfit` package includes convenient methods to programmatically generate initial conditions for a large number of fitting parameters (see Methods). The initial conditions are assembled into a nested key-value pairs (Python dictionary) to deliver to the optimizer. In step (4), the fitting results in the form of parameter lists are collected by a `DataFrame` data structure from `pandas`, whose expansive functionality allows further statistical analysis of the outcome [17]. The fitting results and the initial conditions may be exported for reproducible examination of the specific fitting tasks. A schematic of the computational workflow is shown in Supplementary Fig. 1.

**Performance evaluation with use cases.** We demonstrate the use of the batch fitting method on experimental data from two popular materials characterization techniques: photoemission spectroscopy and powder electron diffraction. Both of these techniques generate multidimensional datasets involving multicomponent spectra or spectrum-like signals. Spectrum fitting is carried out with predefined, domain-specific multicomponent models. In momentum-resolved photoemission spectroscopy, extracting electronic band dispersion requires batch fitting of momentum-dependent intensity-valued energy spectra. We use photoemission band mapping data (see Methods) of bulk tungsten diselenide ( $\text{WSe}_2$ ), which was measured for about an hour and contains a maximum of 14 energy bands within the probed energy range of  $\sim 8$  eV, to benchmark the fitting algorithms. The electron diffraction dataset was recorded for polycrystalline platinum (Pt) with a pump-probe scheme – pump light induces electronic and lattice dynamics that are probed by electron diffraction – to obtain time- and fluence-dependent 1D diffraction patterns [18]. The data was acquired continuously over 2 days with up to 10 diffraction peaks

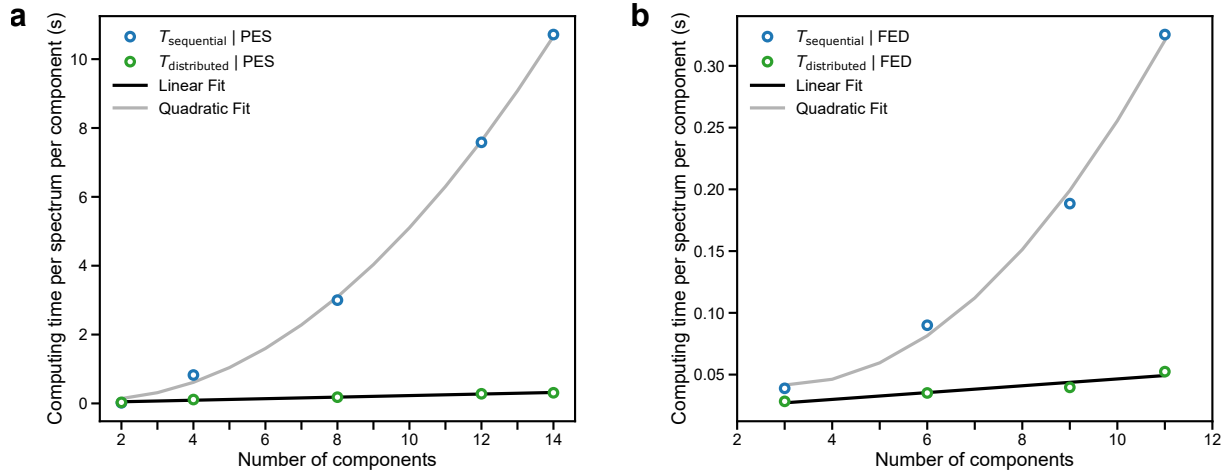


Figure 2: **Computational evaluation on experimental datasets.** Performance comparison between the sequential fitting procedure and the distributed fitting procedure on benchmark datasets: **a**, photoemission spectroscopy (PES) of the valence band (around K point in momentum space) of WSe<sub>2</sub> with up to 14 energy bands considered; **b**, femtosecond electron diffraction (FED) of polycrystalline Pt [18], with up to 10 diffraction peaks and a background profile considered. The computing times ( $T_{\text{sequential}}$  and  $T_{\text{distributed}}$ ), here used as the evaluation metric, are normalized to both the number of components and spectra (or diffraction peaks).

of Pt within the measured reciprocal space. Further details of the experiments are provided in Methods.

In both use cases, the batch fitting is initiated with realistic guesses in their respective domains (see Methods), the optimization processes using both sequential and distributed approaches are timed for comparison. As shown in Fig. 2, the realistic cost metric used for comparison is the computing time normalized to the per-spectrum and per-component level. The distributed approach generally exhibits a linear dependence on the number of components in the data, while the sequential approach, as is common practice in the field, has at least a quadratic scaling. In both cases, distributed fitting of a significant portion of the data is sufficiently faster than the respective experimental duration (see Methods). The polynomial speed-up therefore allows multicomponent spectral analysis on the fly during the experimental measurements, with only the added cost on the computing hardware. Moreover, the tuning of the initialization to improve the fitting outcome may be conducted interactively using the same software, which allows to design custom approaches to deal with complex spectra. Some intuitive approaches to achieve high performance and consistency within the fitting results are discussed in the Supplementary Information and illustrated in Supplementary Fig. 2.

## Discussion

We have described a linear-scaling software architecture for mapping the multicomponent spectral analysis problem to distributed hardware. The modular architecture of the `pesfit` package also allows convenient upgrade of its components individually in the future. For example, it is possible to plug in alternative optimizers for batch fitting, such as brute-force search or variable projection [19], which have respective speed-accuracy tradeoffs in solving large-scale nonlinear least squares problems. Although we have assumed prior knowledge for the number of components, the requirement may be relaxed using numerical estimation [20, 21]. Our approach fully exploits the HPC infrastructure to provide a convenient toolkit for speeding up data annotation in materials characterization without requiring repeated data loading and human intervention. The computational performance demonstrated with benchmark use cases provides a baseline for future algorithm development using such annotation for training supervised learning algorithms. The software may be integrated with existing computational workflows [10, 22], thereby functioning as a building block for closed-loop materials characterization to facilitate experimental optimization and discovery.

## Methods

**Materials characterization datasets.** Both experimental datasets are measured with home-built instruments at the Fritz Haber Institute. The photoemission spectroscopy experiment is conducted using pulsed light source based on high harmonic generation as described previously [23]. The sample under study is single crystalline tungsten diselenide ( $\text{WSe}_2$ ). Extreme UV pulses with an energy of 21.7 eV liberate electrons from the  $\text{WSe}_2$  sample surface, which are subsequently captured as single events in 3D (energy and two momenta) by a time-of-flight delay-line detector (METIS 1000, SPECS GmbH). A custom computational pipeline is used to preprocess the single events to produce the 3D photoemission band mapping data [22]. The femtosecond electron diffraction experiment [18] contains fluence-dependent measurements of polycrystalline Pt diffraction pattern with the laser excitation at 0.70 eV and 70 kV electron bunches generated from a gold photocathode [24]. The 2D diffraction patterns are recorded with a CMOS camera (TemCam-F416, TVIPS GmbH) at a frame integration time of 5 s. The measured diffraction rings are collapsed into 1D diffraction peaks via radial averaging.

**Software details and mathematical context.** The main modules in `pesfit` are `lineshape`, `fitter` and `metrics`. The `lineshape` module contains the `MultipeakModel` class. The `fitter` module contains functions to automate the construction of fitting tasks and initializations as well as simple visualization methods for displaying fitting results. The `metrics` module includes evaluation metrics for assessing the quality of fits. The sequential and distributed fitting are achieved through the `PatchFitter` and `DistributedFitter` classes, respectively, both lie within the `fitter` module. To fit higher-dimensional ( $> 2D$ ) spectral data, the data are first partially transformed to 2D before the fitting tasks are farmed out to processors. The initial conditions can be generated programmatically using `init_generator` and supplied to a multicomponent model via `varsetter`, both are integrated into the `fitter` module classes discussed here, but are also invocable separately and individually. The ordering of fitting tasks is tracked by a spectrum index (`spec_id`) to accommodate asynchronous concurrent execution, which permits index scrambling.

Both sequential and distributed fitting of the benchmark datasets uses the default Levenberg-Marquardt algorithm for minimization with a spectrum-wise least-squares loss function and bound constraints. We formulate the single multicomponent spectral analysis task as

$$\tilde{\mathbf{S}} = \operatorname{argmin} \left\| \left\| I(\omega) - \sum_n f_n(\omega, \mathbf{s}_n) \right\|_{\omega} \right\|^2. \quad (1)$$

Here,  $I$  is the intensity-valued spectrum data with the spectral dimension  $\omega$ .  $\mathbf{s}_n = (s_{n,1}, s_{n,2}, \dots)$  is the set of parameters of the single-component spectral lineshape or background model  $f_n$ , with  $n$  being the index of spectral components. The symbol  $\|\cdot\|_{\omega}$  denotes the Euclidean norm with respect to the spectral dimension.  $\tilde{\mathbf{S}} = [\tilde{\mathbf{s}}_1 \tilde{\mathbf{s}}_2 \dots]$  is the concatenated list of component-wise fitting parameters after optimization. For example, a 10-component spectrum with 4 parameters for each component results in an  $\tilde{\mathbf{S}}$  of size 40 for a single-spectrum fitting task. Based on this notation, batch fitting of multidimensional data is written as

$$\{\tilde{\mathbf{S}}^{\{a_i\}}\} = \operatorname{argmin} \left\| \left\| I(\omega, \{a_i\}) - \sum_n f_n^{\{a_i\}}(\omega, \mathbf{s}_n^{\{a_i\}}) \right\|_{\omega} \right\|^2, \quad (i = 1, 2, \dots). \quad (2)$$

Now,  $I$  becomes multidimensional with the coordinates of other dimensions represented by  $\{a_i\} = (a_1, a_2, \dots)$ . A scalar-valued spectrum index is assigned to each multidimensional coordinate  $\{a_i\}$  according to the row-major order. On the left-hand side of Eq. (2),  $\{\tilde{\mathbf{S}}^{\{a_i\}}\} = \{[\tilde{\mathbf{s}}_1 \tilde{\mathbf{s}}_2 \dots]^{\{a_i\}}\}$  involves a double concatenation (firstly row-wise over the inner brackets  $[\dots]$  and ordered by the spectral components, secondly column-wise over the outer brackets  $\{\dots\}$ )

and ordered by the spectrum index), which represents the entirety of spectrum-wise optimized parameter sets gathered over all  $\{a_i\}$  coordinates. This arrangement turns the fitting results into a 2D array regardless of the dimensionality of the data. The nonlinear least-squares problems in Eqs. (1)-(2) are considered separable in the sense that the spectrum model may be written as a sum of components [6, 19].

**Computational benchmarks.** All batch fitting benchmarks are run on an on-premises computing server (Dell PowerEdge R840), equipped with four Intel Xeon Gold 6150 multicore CPUs. The sequential fitting tasks are run with a single logical core, while the distributed fitting tasks are run using the specified number of processes mapped onto logical cores. The RAM size doesn't pose a limit on the computing time of the benchmark fitting tasks. Fitting of the photoemission data of WSe<sub>2</sub> uses the density functional theory electronic structure calculation available within the NOMAD repository (DOI: [10.17172/NOMAD/2020.03.28-1](https://doi.org/10.17172/NOMAD/2020.03.28-1)) along with a rigid shift of 0.2 eV to all bands as the initial condition for the band positions, while fitting the diffraction peak positions of polycrystalline Pt uses the estimate positions from a reference spectrum. Each spectrum in the photoemission data was fit with up to 14 Voigt lineshapes in combination. For experimental monitoring, band dispersion information from a region around a high-symmetry point of the material suffices. As a time reference, batch fitting of 1600 4-component spectrum (see Supplementary Fig. 2) costs around 6 min. The electron diffraction data was fit with up to 10 Voigt lineshapes and an exponential-decay background model. The entire set of 225 11-component spectrum costs < 3 mins. Since in practice, only a fraction of the data (both in range and in quantity) is needed to monitor the sample condition, the speed boost is sufficient to enable real-time interaction with experiments. Running the distributed fitting consistently at the optimal speed requires tuning of the task parallelization parameters – the number of worker processes and the number of tasks per process. An extended example showing the trade-off between these two parameters is shown in Supplementary Fig. 3. Nevertheless, in practice, only a few trials on subsets of data are needed to find a reasonable trade-off that minimizes the computing time.



## **Acknowledgments**

We thank S. Dong, S. Beaulieu and L. Rettig for providing the photoemission spectroscopy data, D. Zahn for providing the electron diffraction data and helpful discussions. We thank S. Schülke and G. Schnapka at Gemeinsames Netzwerkzentrum (GNZ) and L. Rettig at the Fritz Haber Institute in Berlin for support on the computing infrastructure. R.P.X. thanks S. Oller at Universitätsklinikum Hamburg-Eppendorf (UKE) for helpful discussion on code parallelization. R.P.X. and R.E. acknowledge the support by BiGmax, the Max Planck Society's Research Network on Big-Data-Driven Materials-Science, the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant No. ERC-2015-CoG-682843). P.M.P. acknowledges financial support from STROBE: A National Science Foundation Science & Technology Center under Grant No. DMR 1548924.

## **Data Availability**

The data presented in the paper is available upon reasonable request from the corresponding author.

## **Author contributions**

R.P.X. conceived the project and developed the software with help from P.M.P.. R.P.X. wrote the first draft of the paper. All authors contributed to discussion and revision of the manuscript to its final form.

## **Competing interests**

The authors declare no competing interests in the content of the article.

# Supplementary Information

## Scalable multicomponent spectral analysis for high-throughput data annotation

R. Patrick Xian<sup>1,2\*</sup>   Ralph Ernstorfer<sup>1</sup>   Philipp M. Pelz<sup>3,4</sup>

<sup>1</sup>Fritz Haber Institute of the Max Planck Society, 14195 Berlin, Germany.

<sup>2</sup>Department of Neurobiology, Northwestern University, Evanston 60208, IL, USA.

<sup>3</sup>National Center for Electron Microscopy Facility, Molecular Foundry,  
Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.

<sup>4</sup>Department of Materials Science and Engineering,  
University of California, Berkeley, CA 94720, USA.

\*Correspondence authors: xian@fhi-berlin.mpg.de

### **S1 Modes of operation**


Although the problem mapping and software realization presented in this work offer a scalable and automatable solution for batch fitting of multicomponent spectra, situations will always exist when a subset of the fitting tasks yield unfeasible answers due to experimental resolution, noise in the data, the sloppiness of the model used for fitting [25], etc. This also applies to offline data analysis where a more accurate model is often required for fitting to extract specific parameters, while in online analysis, approximate models with less computational overhead are favored. In the following, we offer here a set of strategies outlined in the computational workflow (see Supplementary Fig. 1) to resolve potential outstanding cases while making use of the existing convenience functionalities offered by `pesfit`. In the demonstrations below, we discuss primarily in the context of photoemission data due to its complexity compared with diffraction data, but a similar line of reasoning also applies to other materials characterization methods involving spectra or spectrum-like data.

## S1.1 Tuning of hyperparameters and initialization

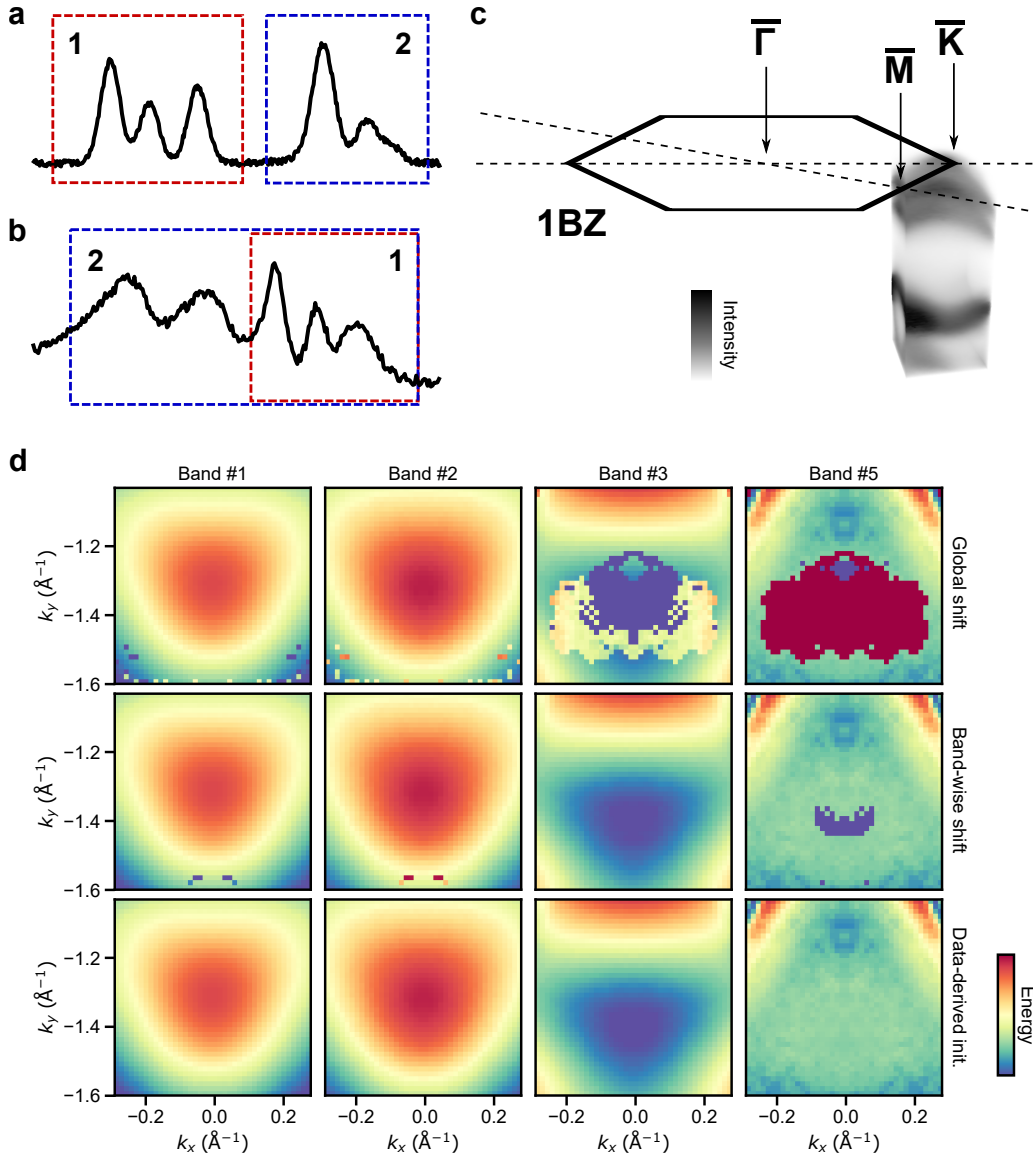
To arrive at a meaningful outcome, batch fitting of multicomponent spectra requires suitable initialization with knowledge from theoretical calculations or educated guesses. We treat rigid shifts applied to theory or reference in initialization as hyperparameters for the fitting procedure, which may be tuned in different ways according to the data characteristics. We discuss here three sets of methods that may be realized using `pesfit`.

**Theory-based approach.** When theoretical calculation or reference spectrum is present, (i) The `random_varshift` method in the `fitter` module of `pesfit` enables automated tuning of initial condition using random search within user-defined values to optimize a goodness-of-fit metric. (ii) Alternatively, if regions of the spectra may be isolated with no overlap with the rest (e.g. when spectral intensity goes to baseline level at the two ends of the segment), a (1) divide-and-conquer tuning approach may be adopted to batch fit and tune the hyperparameters for a segment at a time. This applies to diffraction data and the photoemission data of materials with relatively flat and isolated bands, such as organic molecular crystals or monolayers [26]. If the isolation between segments is less perfect, one may adopt an (2) expanding window tuning approach to batch fit and tune an increasing number of spectral components (thus the “expanding window”), retaining the hyperparameters used for the previous segments while fitting in an increasingly broader spectral window. These two approaches are illustrated in Supplementary Fig. 3a-b. During the tuning process, the fitting outcome may be appraised using either goodness-of-fit metrics or via visual inspection. An example showing the results before and after hyperparameter tuning is shown in Supplementary Fig. 3c-d using photoemission data from  $\text{WSe}_2$ . The results compare that from initialization with a 0.2 eV rigid shift of all bands (3c) and that with a band-wise tuned rigid shift of 0.2 eV, 0.18 eV, 0.22 eV and -0.24 eV for each energy band (3d), respectively.

**Data-driven approach.** When theoretical calculation or a reference spectrum is not available, (iii) one can resort to a data-driven approach. A selected number of distinctive spectra may be fit manually (guess fit) to determine the anchor points as reference for interpolation, which generates the initial conditions for the entire dataset. The data-derived initial conditions may then be used in batch fitting. An example using 16 equally-spaced anchor points to fit 1600 ( $40 \times 40$  grid in  $k_x - k_y$  coordinates) multicomponent spectra is shown in Supplementary Fig. 3d (lower row). Here, we used the Clough-Tocher interpolator [27] implemented in `scipy` [28]




Supplementary Figure 1: **Proposed multicomponent spectral analysis workflow.** The computational workflow starts from the dataset organized into a multidimensional array. Then, the multicomponent spectrum model is constructed using domain knowledge. If theoretical calculation or a reference spectrum is available, the workflow progresses to the stage of tuning task parallelization parameters. Otherwise, an initial guess fit of single spectra at distinct locations is required, along with interpolation of the spectral peak locations to other instances without guess fitting results. Using theory or reference spectrum positions as input, hyperparameters tuning is needed to compensate for the offsets from experimental data, but this is often unnecessary if guess fitting results are available. Finally, the batch fitting routine may be deployed for offline analysis or online monitoring of experimental results, using user-defined metrics constructed with the spectrum fitting outcome.



Supplementary Figure 2: **Semi-automated fine-tuning of hyperparameters.** Tuning of the hyperparameters in batch fitting multicomponent spectra with **a**, the divide-and-conquer approach and **b**, the expanding window approach, illustrated with synthetic spectra. The two segments are numbered in dashed boxes with different colors. Demonstration of the tuning uses **c**, the photoemission dataset (containing 1600 multicomponent spectra) measured near the  $\bar{K}$  point of  $\text{WSe}_2$ . Other important landmarks within the projected first Brillouin zone (1BZ) include the zone center ( $\bar{\Gamma}$ ) and zone edge ( $\bar{K}$ ). The overline signifies the surface projection (e.g.  $\bar{K}$  in photoemission corresponds to  $K$ ). **d**. Comparison of the band dispersions reconstructed using batch fitting without (upper row) and with band-wise fine-tuning (middle row) of the shift hyperparameters of band structure calculation. The lower row shows the batch fitting results using data-derived initialization. The fourth band is not resolved here due to strong overlap with the third band at around  $\bar{K}$ . The band mapping data used for fitting has been symmetrized to balance the matrix elements. The colormap for each energy band is scaled separately for visualization.

to obtain the data-derived initial conditions for all bands, although other multivariate nonlinear interpolators should suffice in general. In comparison, the data-driven approach shows higher consistency among neighboring spectra due to the closer resemblance of data-derived initial condition than the theoretical calculation to the final outcome.



Supplementary Figure 3: **Effects of task parallelization parameters on computing time.** Linear-scale heatmap of the computing time for the distributed multicomponent fitting procedure presented in this work. Each entry represents the computing time (in seconds) elapsed in a numerical experiment, where a distinct pair of parameters (processes and tasks per process) is selected to execute the same number of multicomponent spectral analysis tasks (900 spectra in total, each with 4 components), which uses the synthetic photoemission data benchmark from around the K point of WSe<sub>2</sub>.

## S1.2 Tuning of model complexity

In cases where strictly isolating a group of spectral features using data slicing is challenging, there may exist the situation where a subset of the data requires a more complex model (secondary model) to fit than the rest (primary model). Alternatively, the tuning may also be realized by changing the data range used in fitting. This appears commonly in cases where spectral shift

is present, such as in photoemission data due to the existence of band dispersion and matrix element effects, which results in the dependence of spectral features on momentum coordinate (in momentum-resolved measurements), sample position (in spatially-resolved measurements), etc. Comparatively, in diffraction datasets involving phase transition due to changes in macroscopic parameters or dynamics, the diffraction peak position changes, if any, are relatively small with respect to the peak width. After batch fitting with the primary model, the less well-fit spectra may be selected using goodness-of-fit metrics. These outlier spectra may be batch fit again using a secondary spectrum model with more (or less) components than the primary model.

### S1.3 Tuning of task parallelization parameters

For any batch fitting task, there are generally two key parameters to tune to minimize the overall computing time by balancing the number of worker processes ( $n_W$ ) and the number of tasks per worker process ( $n_T$ ), or *chunk size* in this context [29]. Their tradeoff depends on the complexity of the multicomponent spectrum model and the quality of the initial condition provided to the optimizer. An example is provided in Supplementary Fig. 3. To facilitate the use of our approach, we provide here a list of potential choices for these parameters for reference. From our experience, batch fitting hundreds to thousands of spectra generally requires  $n_W > 10$  for performance optimization, although for low-complexity problems (e.g. with very few spectral components), a smaller number of processes will suffice. It's generally, when there are only 2-3 components in the multicomponent spectrum model,  $n_T$  should be a few hundreds to maximize the computational efficiency. When there are 4-6 components per spectrum,  $n_T$  should be below 100. When there are 8-10 components per spectrum,  $n_T$  should be less than 30-40. Beyond 10 components per spectrum, it is advised not to use  $n_T$  larger than 10-20.

## References

1. Spurgeon, S. R. *et al.* Towards data-driven next-generation transmission electron microscopy. *Nature Materials* (2020).
2. Woodruff, P. *Modern Techniques of Surface Science* 3rd (Cambridge University Press, Cambridge, 2016).
3. Hofmann, P. In operando devices studied by angle-resolved photoemission spectroscopy. *arXiv*, 2011.11490. arXiv: [2011.11490](https://arxiv.org/abs/2011.11490) (2020).
4. Umehara, M. *et al.* Analyzing machine learning models to accelerate generation of fundamental materials insights. *npj Computational Materials* **5**, 34 (2019).

5. Blackburn, J. A. Computer Program for Multicomponent Spectrum Analysis Using Least Squares Method. *Analytical Chemistry* **37**, 1000–1003 (1965).
6. Ruhe, A. & Wedin, P. Å. Algorithms for Separable Nonlinear Least Squares Problems. *SIAM Review* **22**, 318–337 (1980).
7. De Juan, A., Jaumot, J. & Tauler, R. Multivariate Curve Resolution (MCR). Solving the mixture analysis problem. *Anal. Methods* **6**, 4964–4976 (2014).
8. Kriesten, E. *et al.* Identification of unknown pure component spectra by indirect hard modeling. *Chemometrics and Intelligent Laboratory Systems* **93**, 108–119 (2008).
9. Speranza, G. & Canteri, R. RxpsG a new open project for Photoelectron and Electron Spectroscopy data processing. *SoftwareX* **10**, 100282 (2019).
10. René de Cotret, L. P., Otto, M. R., Stern, M. J. & Siwick, B. J. An open-source software ecosystem for the interactive exploration of ultrafast electron scattering data. *Advanced Structural and Chemical Imaging* **4**, 11 (2018).
11. Xian, R. P. *pesfit* <https://github.com/mpes-kit/pesfit>.
12. Lämmel, R. Google’s MapReduce programming model — Revisited. *Science of Computer Programming* **70**, 1–30 (2008).
13. Newville, M. *et al.* *lmfit/lmfit-py 1.0.0* <https://doi.org/10.5281/zenodo.3588521>. 2019.
14. Dask Development Team. *Dask: Library for dynamic task scheduling* (2016).
15. Oller, S. *parmap* <https://github.com/zeehio/parmap>.
16. Hadjidoukas, P. *et al.* torcpy: Supporting task parallelism in Python. *SoftwareX* **12**, 100517 (2020).
17. McKinney, W. *Data Structures for Statistical Computing in Python* in *Proceedings of the 9th Python in Science Conference* (eds van der Walt, S. & Millman, J.) (2010), 56–61.
18. Zahn, D., Seiler, H., Windsor, Y. W. & Ernstorfer, R. Ultrafast lattice dynamics and electron-phonon coupling in laser-excited platinum. *arXiv*, 2012.10428. arXiv: [2012.10428](https://arxiv.org/abs/2012.10428) (2020).
19. Golub, G. & Pereyra, V. Separable nonlinear least squares: the variable projection method and its applications. *Inverse Problems* **19**, R1–R26 (2003).
20. Lindner, R. R. *et al.* Autonomous Gaussian decomposition. *The Astronomical Journal* **149**, 138 (2015).
21. Ament, S. E. *et al.* Multi-component background learning automates signal detection for spectroscopic data. *npj Computational Materials* **5**, 77 (2019).
22. Xian, R. P. *et al.* An open-source, end-to-end workflow for multidimensional photoemission spectroscopy. *Scientific Data* **7**, 442 (2020).
23. Puppin, M. *et al.* Time- and angle-resolved photoemission spectroscopy of solids in the extreme ultraviolet at 500 kHz repetition rate. *Review of Scientific Instruments* **90**, 023104 (2019).



24. Waldecker, L., Bertoni, R. & Ernstorfer, R. Compact femtosecond electron diffractometer with 100 keV electron bunches approaching the single-electron pulse duration limit. *Journal of Applied Physics* **117**, 044903 (2015).
25. Transtrum, M. K. *et al.* Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *The Journal of Chemical Physics* **143**, 010901 (2015).
26. Nakayama, Y., Kera, S. & Ueno, N. Photoelectron spectroscopy on single crystals of organic semiconductors: experimental electronic band structure for optoelectronic properties. *Journal of Materials Chemistry C* **8**, 9090–9132 (2020).
27. Alfeld, P. A trivariate clough—tocher scheme for tetrahedral data. *Computer Aided Geometric Design* **1**, 169–181 (1984).
28. Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* **17**, 261–272 (2020).
29. Singh, N., Browne, L.-M. & Butler, R. Parallel astronomical data processing with Python: Recipes for multicore machines. *Astronomy and Computing* **2**, 1–10 (2013).