



Implementation of the emulator-based component analysis

Anton Vladyka*, Eemeli A. Eronen, Johannes Niskanen

University of Turku, Department of Physics and Astronomy, FI-20014 Turun yliopisto, Finland

ARTICLE INFO

Dataset link: [10.5281/zenodo.10405685](https://zenodo.org/doi/10.5281/zenodo.10405685)

Keywords:

Inverse problems
Machine learning
PyTorch

ABSTRACT

We present a PyTorch-powered implementation of the emulator-based component analysis used for ill-posed numerical non-linear inverse problems, where an approximate emulator for the forward problem is known. This emulator may be a numerical model, an interpolating function, or a fitting function such as a neural network. With the help of the emulator and a data set, the method seeks dimensionality reduction by projection in the variable space so that maximal variance of the target (response) values of the data is covered. The obtained basis set for projection in the variable space defines a subspace of the greatest response for the outcome of the forward problem. The method allows for the reconstruction of the coordinates in this subspace for an approximate solution to the inverse problem. We present an example of using the code provided as a Python class.

1. Introduction

Inverse problems consist of reconstructing input data \mathbf{X} or parameters of a model from a later observation of a dependent response \mathbf{Y} . In addition to their rich heritage from pure mathematics [1–3], many applied inverse problems originate from geophysics [4]. Other modern examples of the applied problems include tomography [5] (as a part of the more general image/signal processing problem) and spectroscopy. Inverse problems can be identified in a wealth of research contexts, and their approximate solutions are traditionally searched for by regularization or by Bayesian procedures [1,6].

For example, given an atomistic structure with a charge state, the resulting electronic spectrum (ultraviolet/visible, X-ray, etc.) is defined by the quantum-mechanical electronic Hamiltonian of the system. The formed map is certainly nonlinear in terms of structure, and not necessarily an injective function. Moreover, some structural degrees of freedom may account for huge spectral responses while some may result little or virtually none. As often the case with inverse problems, full reconstruction of structural parameters is an ill-posed problem, and already identification of the relevant structural degrees of freedom may be a worthwhile scientific result. It may likewise be important to filter out structural degrees of freedom without spectral response to achieve a more well-posed inverse problem of reconstruction of structural information from the respective spectra [7].

Machine learning (ML) opens new possibilities for solving inverse problems. With the backpropagation algorithm [8] and graphics processing units (GPU), training of extensive neural networks (NN) has become feasible. Because the NN is a universal approximator [9–11],

mimicking many forward problems is possible given that enough data is available for the task. The complete solution of a potentially ill-defined inverse problem may still be difficult, but an approximative one may be obtained by first identifying the reconstructable information. To do this, the fast evaluation of the forward problem by NN may be utilized in iterative procedures. We propose the emulator-based component analysis (ECA) [12] as a tool in finding a fast approximate solution for nonlinear inverse problems. The method has been applied to spectroscopic problems from gaseous [12], to condensed amorphous [7] and to liquid [13] phases using the original implementation, and to a molecular liquid using the notably faster presented implementation [14].

The ECA algorithm aims for dimensionality reduction by projection on basis vectors in the input space \mathbf{X} while maximizing the variance of a dependent target variable of the forward problem in space \mathbf{Y} of vectors $\mathbf{y} = \mathbf{y}(\mathbf{x})$. Thus ECA belongs to the family of ‘projection pursuit’ algorithms [15–17] and is closely related to projection pursuit regression (PPR) [18]. However, instead of using an unknown smoothing function (ridge function), ECA utilizes a more general and pre-selected function: a machine learning model (i.e. a NN) acting as an emulator for the forward problem $\mathbf{y}_{\text{emu}}(\mathbf{x}) \approx \mathbf{y}(\mathbf{x})$. For example, in spectroscopy studies, such emulator would predict a spectrum for the corresponding atomistic structure. Approximate solution \mathbf{x}' of the inverse problem $\mathbf{y} \rightarrow \mathbf{x}$ may then be found for the obtained dimensionally-reduced space. In this contribution we present an efficient numerical implementation of the ECA algorithm and its application for the inverse problem together with an example.

* Corresponding author.

E-mail addresses: anton.vladyka@utu.fi (A. Vladyka), johannes.niskanen@utu.fi (J. Niskanen).

<https://doi.org/10.1016/j.jocs.2024.102437>

Received 20 December 2023; Received in revised form 22 August 2024; Accepted 1 September 2024

Available online 18 September 2024

1877-7503/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. The principle of ECA illustrated for 2 ECA components $\{v_1, v_2\}$ in a 3-dimensional X space $\text{span}\{x_1, x_2, x_3\}$ (a). The ECA components (basis vectors) are selected by iterative optimization so that emulated y_{emu} for the projected data points match closest with the known y of the original data (b).

2. Definition

The result of ECA is an orthogonal set of vectors $\mathbf{V} = \{v_j\}_{j=1}^k$ in X space. These vectors are optimized using data $\{(x_i, y_i)\}$ separate to those used for emulator training, so that most variance of vectors y_i is covered by emulator prediction $y_{\text{emu}}(\tilde{x}_i^{(k)})$ subsequent to projection of vectors x_i onto \mathbf{V}

$$\tilde{x}_i^{(k)} = \sum_{j=1}^k \underbrace{(v_j \cdot \tilde{x}_i)}_{=:t_{i,j}} v_j. \quad (1)$$

The algorithm is designed to work with z-standardized X data points \tilde{x}_i (corresponding to x_i) for which Eq. (1) represents a reduction to k dimensions. The coordinates of the data point $\tilde{x}_i^{(k)}$ reduced dimension (t scores) read $t_i = (t_{i,1}, \dots, t_{i,k})$. The idea is depicted in Fig. 1.

As a metric for the optimization of \mathbf{V} , we use a generalized covered variance (R^2 -score). Organizing z-standardized Y data points \tilde{y}_i as row vectors of matrix \mathbf{Y} , and the corresponding predicted data points $y_{\text{emu}}(\tilde{x}_i^{(k)})$ as row vectors of matrix $\mathbf{Y}^{(\text{pred})}$, this is defined as:

$$\rho = 1 - \frac{\text{tr}(\tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}})}{\text{tr}(\mathbf{Y}^T \mathbf{Y})} \quad (2)$$

where $\tilde{\mathbf{Y}} = \mathbf{Y} - \mathbf{Y}^{(\text{pred})}$.

The procedure utilizes iterative optimization of the basis vectors so that ρ is maximized for the used data. The method proceeds one vector v_j at a time where orthonormality is enforced after optimization of each vector. We note that the replacing any of the ECA basis vector v_i with its opposite $-v_i$ does not affect the covered variance, but only changes the sign of the corresponding t_i score.

The expansion $\tilde{x}_i^{(k)}$ gives an approximation of point \tilde{x}_i (and x_i after inverse z-standardization), in the subspace $\text{span}(\mathbf{V})$, which \mathbf{Y} by design shows greatest response to. We approach the inverse problem by deducing the respective coordinates t' for a given y , and \tilde{y} , after which the approximate solution is

$$\tilde{x}' = \sum_{j=1}^k t'_j v_j. \quad (3)$$

This reconstruction utilizes numerical optimization for t' for given y and \mathbf{V} . The approximate solution x' in X space is obtained via z-score inverse transformation.

3. Implementation

The module is written in Python utilizing PyTorch [19] (> 2.0). The use of PyTorch offers some benefits:

1. Runs very efficiently owing to automatic differentiation approach implemented in this package.
2. Can easily be run in parallel over multiple CPU cores or a GPU.
3. It provides more control during the fitting procedure than some other optimization toolkits, (e.g., `scipy.optimize` [20]).

We arranged the search for ECA vectors as a mini-batch optimization using the Adam optimizer [21]. The algorithm utilizes previously found component vectors, and orthogonality of the search is forced by allowing optimization steps only in their orthogonal complement. The workflow of the procedure is presented in Algorithm 1 and the optimization options together with their default values are given in Table 1. The full reference manual and a sample code for the implementation are available in the Appendix.

Algorithm 1 ECA implementation

Require: Trained emulator $y_{\text{emu}} : x \rightarrow y$

Require: Data (X, Y)

Initialize first vector: $v_1 \leftarrow \text{rnd}()$

while v_1 not optimized **do**

 Calculate projection: $X_{\text{proj}} \leftarrow (X \cdot v_1) \cdot v_1$

 Optimize $v_1 : \rho(Y, y_{\text{emu}}(X_{\text{proj}})) \rightarrow \max$

 Normalize $v_1 \leftarrow v_1 / |v_1|$

end while

Initialize the basis set: $V_1 \leftarrow v_1$

for rank $i \leftarrow 2, \dots, k$ **do**

 Initialize i -th component: $v_i \leftarrow \text{rnd}()$

while v_i not optimized **do**

 Calculate projection: $X_{\text{proj}} \leftarrow (X \cdot V) \cdot V + (X \cdot v_i) \cdot v_i$

 Optimize $v_i : (\rho(Y, y_{\text{emu}}(X_{\text{proj}})) \rightarrow \max) \cap (V \cdot v_i = 0) \quad \triangleright v_i$ is

orthogonal to vector space \mathbf{V}

 Normalize $v_i \leftarrow v_i / |v_i|$

end while

 Append optimized component to the basis: $V_i \leftarrow v_i$

end for

Reproducibility of the results is satisfied by setting the seed for the random numbers generator via `seed` parameter of the options or by calling `set_seed()` method. The `fit()` method also allows the manual step by step evaluation and re-evaluation of the components by specifying `keep=n` argument: a positive value is interpreted as ‘start after n th component’, and for negative values the last $|n|$ calculated components are removed, evaluation starts after the remaining ones.

In addition to providing the ECA basis vectors for the data-emulator pair, the presented implementation carries out typically used transformations. These functionalities are presented in Fig. 2.

In this work, we propose to utilize dimensionality reduction achieved via projection of the X data onto ECA space for approximation of the inverse problem solution. The `inverse()` method is searching for the points in ECA space (t scores) for which the emulator provides the best match with given y , while the `reconstruct()` method is performing inverse transformation followed by its expansion to X space. In other words, it returns an x' , which predicts to the best match with given y . The obtained t' scores can be used for studying the differences in the data, e.g., via possible grouping, cluster analysis etc, and interpretability of x' vectors allows to find the relation between x and y in a significantly more lightweight way than an exact solution of the inverse problem would.

The PyTorch implementation, while robust and efficient, requires an emulator implemented in PyTorch. Numerous machine learning libraries exist and provide a similar functionality. A simple feed-forward network can be transferred between the implementations by specifying the matrices, bias vectors and activation functions. We provide a script which allows to convert a trained sklearn-based multilayer perceptron (`sklearn.MLPRegressor` [22]) into an equivalent PyTorch-based neural network for use with the PyTorch-based ECA.

Rudimentary example

ECA is intended to find a few degrees of freedom that potentially cover most of target variance of the forward problem. As the first

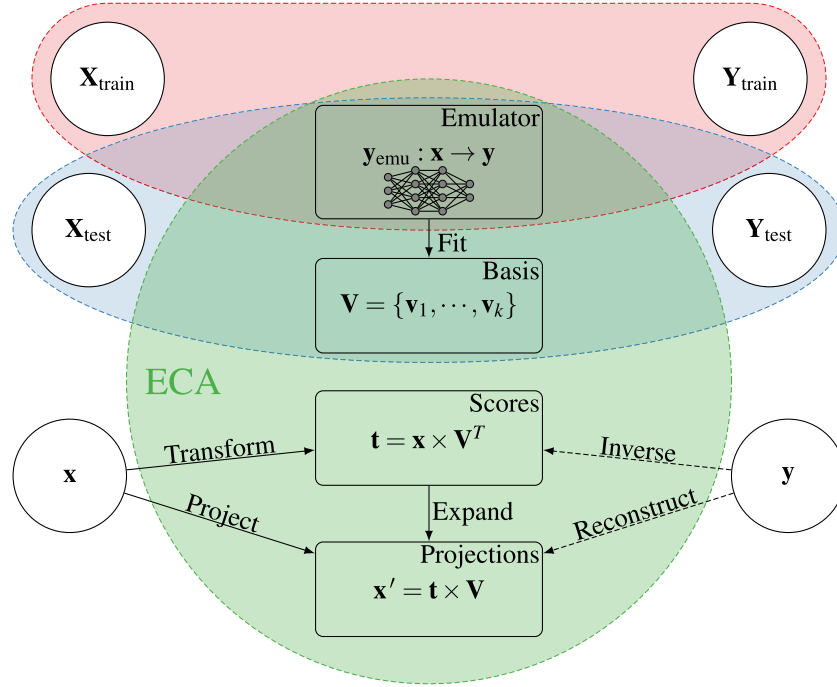


Fig. 2. Schematics of ECA functionality. ECA uses an emulator, trained on $(X, Y)_{train}$, validated on $(X, Y)_{test}$. Fitting of the ECA object on $(X, Y)_{test}$ (or $(X, Y)_{train}$) results in a basis set V . ‘Transform’ and ‘project’ operations are linear transformations of the input vector x . On the contrary, ‘inverse’ and ‘reconstruct’ are implemented using optimization algorithms. These are *not* linear operations and *may* produce ambiguous results for the complex (highly multidimensional) systems.

Table 1
The list of options for both forward and inverse algorithms.

Parameter	Explanation	Default	(inv)
lr	Learning rate for the Adam optimizer	10^{-3}	0.05
betas	Beta parameters for the Adam optimizer (see [21])	(0.9, 0.999)	(0.9, 0.999)
tol	Stopping condition	10^{-4}	10^{-4}
epochs	Max. number of epochs	10^4	10^3
batch_size	Mini-batch size (forward only)	200	–
seed	Seed for random number generator	None	None

example for this we study the problem for the forward function

$$y = y(x) = (v \cdot x)^3 \quad (4)$$

tailored to represent the problem well. Function y of Eq. (4) is dependent on a single degree of freedom in the input space, and in this example we use

$$v = \frac{1}{\sqrt{d}} \sum_{i=1}^d \hat{e}_i, \quad (5)$$

where d is the dimensionality of the input vectors for the forward problem and \hat{e}_i the i th basis vector of this space. This function depends only on vector v which is expected as the result for the first ECA vector v_1 . Equally obviously, at most the scores $t_1 = v_1 \cdot x$ are reconstructable quantities for the inverse problem for this non-bijective map, which manifests partial loss of information about x in the target variable y . We carry out all procedures following the case of an intended study.

We use 20000 data points (x_i, y_i) and vary the dimension d of the input vector space. For simplicity we sample the x data from a normal distribution, allowing us to skip the element-wise z-score standardization of the x_i data, otherwise expected at this point. To bring the target values to the scale of typical application, we z-score standardized the y data. We split the data to 80% for training the emulator (16000 points) and 20% (4000 points) for testing and ECA fitting. For the emulator required by ECA, we chose an NN with four

hidden layers of 16 neurons and ReLU as the activation function. We implemented the network in PyTorch and trained it using Adam as the optimizer with weight decay of 0.001 and initial learning rate of 0.001. The use of this fixed architecture limits the changes of evaluation time to the ECA procedure and to the inevitable differences in the input-layer operations. For each input-dimensionality d we trained the model with 80% of the training data and determined the condition for patient early stopping with the rest. In the end R^2 score for the test set was always at least 0.94.

To study the dependency of the ECA fit time on input dimensionality d , we ran the fit algorithm 25 times for each d with varying initial guesses. We repeated this procedure for both the presented and for the original implementation using identical emulators. The original implementation uses MLPRegressor class of the sklearn [22] package, to which we converted the trained PyTorch emulators. We ran all computations using a single core on a cluster computer (Puhti, CSC, Kajaani, Finland) to minimize hardware-induced variation in the run time. The limited number of trials (25) is due to the slow evaluation time of the original implementation at large d .

Fig. 3a represents the test set data points in the case $d = 2$, while Fig. 3b depicts this data projected on the two ECA vectors. In this case the vector v_2 is given by the orthonormality condition apart from its sign, but in a higher-dimensional case the choice would not be unique due to equal irrelevance of all other degrees of freedom of the input. Fig. 3c depicts the evaluation time for the ECA fit




Fig. 3. Performance of the new implementation on the rudimentary example data. (a) Test set $\mathbf{x} = (x_1, x_2)$ for two-dimensional ($d = 2$) case. Color of each point depicts the corresponding value $y(\mathbf{x})$ (Eq. (4)) (b) Projection scores $\mathbf{t} = (t_1, t_2)$ of the data on the two ECA vectors using the same color scheme. (c) Evaluation times for ECA fit for \mathbf{v}_1 as a function of dimensionality d . We observed a linear dependence of the ECA fitting time on the number of data points for both implementations. (d) Fit success rate measured as fraction of fits with $|\mathbf{v} \cdot \mathbf{v}_1| > 0.95$ as a function of d from 25 independent trials.

with the two implementations as a function of the input dimension d . The original implementation based on optimization by SciPy [20] optimization routines are faster in fewer dimensions but the mini-batch optimization routine of the presented implementation shows no significant dependence on d , surpassing the former in speed at around $d = 30$. We observed a linear dependence of the ECA fitting time on the number of fit data points for both implementations. Fig. 3d depicts the rate of success of the procedure as a function of d , measured as fraction of fits with $|\mathbf{v} \cdot \mathbf{v}_1| > 0.95$. The presented implementation always fits \mathbf{v}_1 successfully while the original implementation is less stable failing with a third of the initial guesses at $d = 512$. We note that when dimensionality d approaches the number of used data points (in this case 4000 points), the fit became increasingly unstable for the presented implementation as well.

The presented implementation carries also reconstruction of scores \mathbf{t} for a given \mathbf{y} by the method `inverse()`. When only few ECA components cover most of the target variance, this procedure reduces the dimensionality of the inverse problem notably, potentially allowing for an approximate solution for it. Having multiple target values (dimensionality of \mathbf{y}) that provide independent information can be beneficial for this procedure, as is the case of this example. We generated a vector-valued $\mathbf{y}(\mathbf{x})$ by adding components to the function of Eqs. (4) and (5) depending only on t_1 as follows

$$\mathbf{y}(\mathbf{x}) = \begin{bmatrix} (\mathbf{v} \cdot \mathbf{x})^3 \\ \sin(0.2(\mathbf{v} \cdot \mathbf{x})) \\ \cos(0.2(\mathbf{v} \cdot \mathbf{x})) \\ \tanh(0.2(\mathbf{v} \cdot \mathbf{x})) \end{bmatrix}. \quad (6)$$

First, we z-score standardized the four features of the generated \mathbf{y} data, and applied the same procedures as in the one-dimensional case to obtain emulators for each d with R^2 score for the test set always better than 0.98. Then, we ran one dimensional ECA fit and projected all the ECA fit data points onto the obtained \mathbf{v}_1 providing us with scores t_1 . Finally, we performed inverse ECA fit to reconstruct the scores well, as shown in Fig. 4 for $d = 2$ and $d = 512$. The R^2 scores for t_1 between the projection and reconstruction are at least 0.98 for all input dimensionalities.

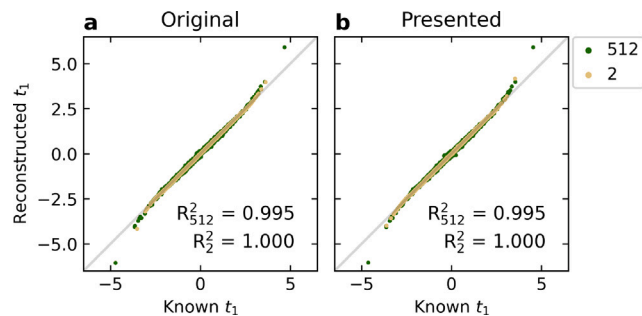


Fig. 4. Reconstructed t_1 scores versus the known (projected) ones for input dimensionalities $d = 2$ and $d = 512$ obtained using (a) the original and (b) the presented implementation. The diagonal gray line indicates a perfect match. The scatter plots show that both the implementations work well in solving the inverse problem for t_1 , further indicated by the high R^2 scores between the reconstructed and known points.

We observed constant inverse fit times regardless of d for both implementations. However, the original implementation was faster by one order of magnitude. This is in agreement with the results in Fig. 3c as the inverse problem is low dimensional and, therefore, does not benefit from the optimization done by PyTorch. Regardless, we consider the presented implementation relatively fast at around 0.07 s per reconstructed data point on the aforementioned cluster computer.

Scientific example

Next, we turn to the analysis of simulated N K-edge near-edge X-ray absorption fine structure (NEXAFS) of aqueous triglycine [13]. This data consists of local structural descriptors (local version of the many-body tensor representation, LMBTR [23], implemented in [24]) and related core-level spectra, available at [25]. The descriptors are built from atomistic structures, consisting of the single triglycine molecule and the water molecules within 3 Å from it. The spectral data are given as intensities of the three regions of interest (ROIs; I–III). The purpose of the analysis is to discover spectrally dominant structural modifications, *i.e.* finding structural degrees the change of which results in the most change in the spectra. ECA was applied as in the original paper: the components were obtained with an ECA fit data set of 3000 data points and the results were validated with ECA validation data set of 2999 data points. The original emulator was trained on 23996 data points, and for this work it was simply converted (by replicating its architecture and weights) to PyTorch. The scripts of our implementation for this data are provided with the codes.

ECA covered variance up to rank five is given in Table 2 and closely resembles the original work. The LMBTR descriptor allows for an interpretation of the deduced basis vectors in terms of the structural features. Fig. 5 presents the first basis vector which describes the interatomic distance distributions (from each active nitrogen site N_{ex}). These results are also in agreement with the previous ones. Moreover, whereas the principal component analysis (PCA) of the structural descriptor shows two distinct regions, we showed that this information is lost in ROIs of NEXAFS, and therefore cannot be recovered by X-ray absorption experiments. The results, available in the scripts provided with this work, were the same with the presented implementation.

4. Discussion

We find the results of the basis vector optimization to sometimes be dependent on the initial guess given for the vectors. Therefore optimization for each component vector should be carried out several (~10–100) times after which the best outcome is to be selected. This behavior is understood by the loss function of the respective minimization problem being wiggly. Our selection of stochastic mini-batch




Fig. 5. The interatomic distance distribution part of the first component vector for the example of aqueous triglycine. The results shown here are in agreement with the original study in [13]. A non-zero value means that the studied spectral regions are sensitive to changes in this structural feature.

Table 2
The covered spectral ROI variances as a function of the rank of the ECA.

Rank	ECA Fit	ECA Validation
1	0.535	0.516
2	0.747	0.710
3	0.850	0.810
4	0.865	0.811
5	0.874	0.817

optimization, although a viable means to reach global minimum, does not guarantee that this minimum is found from a random initial guess. The behavior can likely be made more robust by choice of batch size, learning rate etc. We note that the solution of the inverse problem via inverse transformation did not seem to suffer from the aforementioned instability.

We discovered that the Adam optimizer works faster than the stochastic gradient descent (SGD) algorithm and reduces the dependence of the results from the initial guess (e.g., the random seed). Utilization of the Adam optimizer requires less parameters to be specified, and in our tests with spectroscopic data we even found the default values to work well. We also discovered that for the test cases the learning rate for the inverse transformation must be rather large (≥ 0.05). However, for robustness we would recommend to finely tune the optimizer hyperparameters as well as to check the reproducibility of the fit with, e.g., different initial guesses.

When applied to X-ray spectroscopic data [7,12–14], ECA deduced a subspace in structural data space which defines the majority of the spectral variations, i.e. represents the sensitivity of the spectral changes to the corresponding structural ones. The next step of the analysis is an interpretation of the obtained vectors as meaningful physical quantities. The example shown in this work is not a complete demonstration of the power the ECA — other applications are possible and likely to arise in further studies. For example, these could include visualization and analysis of the possible clustering in the data. Moreover, we have demonstrated that the magnitude of the first ECA component vector may be useful as a feature selection tool [13].

Machine learning models, more specifically simple neural networks have been previously utilized inside variations of the PPR [26–29], but the applied nonlinear projection complicates the interpretation of the results which is essential for further analysis. ECA offers a simple interpretation of the relevant structural features through its component vectors: only a single score is required for each feature (per ECA

component). This is further aided by our observation that often just rank one or two are needed to cover most of the target variance.

Although the schematic in Fig. 2 implies that the emulator is trained and the ECA is fitted with separate data sets, we have found no drawbacks when using the same data for both. After all, the emulator and the ECA can be considered as a single machine learning system.

5. Conclusions

The implementation of the ECA algorithm as a Python class makes the method easily applicable for dimensionality reduction in inverse problems for which paired data of input and output values, and a computationally fast emulator, are available. The original application [12] of the method targeted structural interpretation of X-ray spectra by a dimensionality reduction, which may be drastic [7,13,14]. With the presented PyTorch implementation a robust and fairly fast ECA decomposition is available to any scientific Python user.

CRedit authorship contribution statement

Anton Vladyka: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation. **Eemeli A. Eronen:** Writing – review & editing, Writing – original draft, Validation, Investigation. **Johannes Niskanen:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Anton Vladyka reports financial support was provided by Academy of Finland. Eemeli Eronen reports financial support was provided by Jenny and Antti Wihuri Foundation and Academy of Finland. Johannes Niskanen reports financial support was provided by Academy of Finland.

Data availability

The PyTorch implementation of the ECA and the relevant scripts are available in Zenodo: [10.5281/zenodo.10405685](https://doi.org/10.5281/zenodo.10405685). The data for the example of aqueous Triglycine N K-edge NEXAFS is available at [25].

Acknowledgments

Academy of Finland is acknowledged for funding via project 331234. The authors acknowledge CSC – IT Center for Science, Finland, and the FGCI – Finnish Grid and Cloud Infrastructure for computational resources. E.A.E. acknowledges Jenny and Antti Wihuri Foundation for funding.

Reference manual

`class eca.ECA(emulator)`

Parameters:

emulator: `torch.nn.Module`

Trained PyTorch-based emulator.

Default options:

`DEFAULT_LR = 1E-3`

Learning rate.

`DEFAULT_TOL = 1E-4`

Tolerance level. Optimization stops when the change of the calculated error ('loss') is smaller than tolerance, *i.e.* no more improvement in fit is observed.

`DEFAULT_EPOCHS = 10000`

Maximum number of epochs for optimization. One epoch is an optimization step done for all available X data.

`DEFAULT_BETAS = (0.9, 0.999)`

Beta parameters for the Adam optimizer. These parameters are used for computing running averages of a gradient and its square ('momentum') [21]. See `help(torch.optim.Adam)` for more details.

`DEFAULT_BATCH_SIZE = 200`

Size of mini-batches for optimization.

`DEFAULT_LR_INV = 5E-2`

Learning rate for inverse transformation.

`DEFAULT_TOL_INV = 1E-4`

Tolerance level for inverse transformation.

`DEFAULT_EPOCHS_INV = 1000`

Maximum number of epochs for inverse transformation.

Attributes: (acquired during `fit()`):

V: `torch.Tensor` – calculated ECA components.

y_var: `list[float]` – covered variance of Y data (cumulative).

x_var: `list[float]` – covered variance of X data (cumulative).

Methods:

fit(x, y, n_comp=3, options=None, keep=0, verbose=False)

Performs ECA decomposition.

Parameters:

x: `torch.Tensor`, **y:** `torch.Tensor` – X and Y data, respectively.

n_comp: `int` – number of components to evaluate.

options: `dict`

'lr', `float` – learning rate.

'betas', `tuple(float, float)` – 'betas' parameters of Adam optimizer.

mizer.

'tol', `float` – tolerance level.

'epochs', `int` – maximum number of epochs.

'batch_size', `int` – size of mini-batches.

'seed', `int` – seed for random number generator (affects initial guess and mini-batch splits).

keep: `int` – to keep specified number of components for another fit. *E.g.*, `keep=1` keeps only the first component; `keep=-2` removes two last components.

verbose: `bool` – if True, displays progress of the fit (requires `tqdm` package).

Returns:

V: `torch.Tensor` – calculated ECA components.

y_var: `list[float]` – covered variance of Y data (cumulative).

x_var: `list[float]` – covered variance of X data (cumulative).

inverse(y, n_comp=3, options=None, return_error=False)

Perform inverse transformation from the given **y** to **t'**.

Parameters:

y: `torch.Tensor` – Y data for inverse transformation.

n_comp: `int` – number of **t'** score components to evaluate, *i.e.* the number of V components to use for the inverse transformation. (None: use all available).

options: `dict` – options for inverse transformation: learning rate and maximum number of epochs. See `fit()` for details.

Returns:

t: `torch.Tensor` – **t'** scores for given **y**.

err: `torch.Tensor` – mean-squared error for each evaluated data point.

reconstruct(y, n_comp=3, options=None, return_error=False)

Performs inverse transformation followed by its expansion to X space. `reconstruct(y) ≡ expand(inverse(y))`.

Parameters:

Same as for `inverse()`.

Returns:

x_prime: `torch.Tensor` – **x'** data for given **y**.

err: `torch.Tensor` – mean-squared error for each evaluated data point.

transform(x, n_comp=None)

Performs transformation of **x** data to V space.

Parameters:

x: `torch.Tensor` – given **x** data.

n_comp: `int` – number of V components to use for the transformation (None: use all available).

Returns:

t: `torch.Tensor` – **t** scores for given **x**.

expand(t, n_comp=None)

Expands **t** scores to X space.

Parameters:

t: `torch.Tensor` – given **t** scores.

n_comp: `int` – number of V components to use for the transformation (None: use all available).

Returns:

x_prime: `torch.Tensor` – **x'** data.

project(x, n_comp=None)

Projects given **x** data to ECA space using specified number of V components. `project(y) ≡ expand(transform(y))`.

Parameters:

x: `torch.Tensor` – given **x** data.

n_comp: `int` – number of V components to use for the transformation (None: use all available).

Returns:

x_prime: `torch.Tensor` – **x'** projections.

set_seed(seed)

Set seed for random number generator. Affects the initial guess as well as the mini-batches. Can be set explicitly before running `fit()` or `inverse()`.

Parameters:

seed: `int` – seed.

test(x, y)

For given (**x**, **y**) pair, calculates the covered variance $\rho(y_{\text{pred}}, y)$ between **y** and **y_{pred}**, predicted from the projection **x'** of **x** on ECA V space (see Eq. (2)).

Parameters:

x, y: *torch.Tensor* – given x and y data.

Returns:

rho: *float* – covered variance.

r2loss(y_pred, y_known)

Calculates missing variance between known y_{known} and predicted y_{pred} values defined as $1 - \rho(y_{\text{pred}}, y_{\text{known}})$ (see Eq. (2)).

Parameters:

y_pred, y_known: *torch.Tensor* – given predicted and known values.

Returns:

r2: *torch.Tensor* – covered variance.

Sample code to run ECA decomposition and inverse

```
## import necessary modules to handle data
import numpy as np
from eca import ECA
## Load the data, for example:
data_x = np.loadtxt(...)
data_y = np.loadtxt(...)
# convert data to torch.tensor if necessary:
data_x_t, data_y_t = map(lambda x: torch.tensor(
    x, dtype=torch.float32), [data_x, data_y])
## Initialize an instance of the ECA class with
an emulator
# For pytorch, a torch.nn.Module object can be
called as an emulator as well as its forward
() method
eca = ECA(model)
# define options
options = {
    "tol" : 1E-4,
    "epochs" : 10000,
    "lr" : 1E-3,
    "batch_size": 200,
    "seed" : 123
}
# Carry out ECA decomposition
vectors, y_variance, x_variance = eca.fit(
    data_x_t, data_y_t, n_comp=3, options=
options)

## Example: recalculate the last component with
another seed
# set a new seed
eca.set_seed(42)
# keep the first two (already calculated)
components and redo the third one with a new
seed
vectors, y_variance, x_variance = eca.fit(
    data_x_t, data_y_t, n_comp=3, keep=-1)

# Calculate scores for data_x
t_k = eca.transform(data_x_t)
# Directly project data_x on ECA subspace
data_x_k = eca.project(data_x_t)
# Search iteratively for scores matching the
spectral data points
options_inv = {
    "tol" : 1E-4,
    "epochs" : 1000,
    "lr" : 0.05,
}
t_k_prime = eca.inverse(data_y_t, options=
options_inv)

# Search for approximate data_x_appr matching
the spectral data points.
```

```
data_x_appr = eca.reconstruct(data_y_t)
```

References

- [1] J.P. Kaipio, E. Somersalo, *Statistical and computational inverse problems*, Springer Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, USA, 2005.
- [2] D. Calvetti, E. Somersalo, Inverse problems: From regularization to Bayesian inference, *WIREs Comput. Stat.* 10 (3) (2018) <http://dx.doi.org/10.1002/wics.1427>.
- [3] G. Uhlmann, Inverse problems: Seeing the unseen, *Bull. Math. Sci.* 4 (2) (2014) 209–279, <http://dx.doi.org/10.1007/s13373-014-0051-9>.
- [4] A. Tarantola, B. Valette, Inverse problems=quest for information, *J. Geophys.* 50 (1) (1982) 159–170.
- [5] S.R. Arridge, J.C. Schotland, Optical tomography: Forward and inverse problems, *Inverse Problems* 25 (12) (2009) 123010, <http://dx.doi.org/10.1088/0266-5611/25/12/123010>.
- [6] A. Mohammad-Djafari, Bayesian inference for inverse problems, in: N. Tang (Ed.), *Bayesian Inference*, IntechOpen, Rijeka, 2022, <http://dx.doi.org/10.5772/intechopen.104467>.
- [7] A. Vladyka, C.J. Sahle, J. Niskanen, Towards structural reconstruction from X-Ray spectra, *Phys. Chem. Chem. Phys.* 25 (2023) 6707–6713, <http://dx.doi.org/10.1039/D2CP05420E>.
- [8] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning representations by back-propagating errors, *Nature* 323 (6088) (1986) 533–536, <http://dx.doi.org/10.1038/323533a0>.
- [9] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Netw.* 2 (5) (1989) 359–366, [http://dx.doi.org/10.1016/0893-6080\(89\)90020-8](http://dx.doi.org/10.1016/0893-6080(89)90020-8).
- [10] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Math. Control Signals Systems* 2 (4) (1989) 303–314, <http://dx.doi.org/10.1007/BF02551274>.
- [11] M. Leshno, V.Y. Lin, A. Pinkus, S. Schocken, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, *Neural Netw.* 6 (6) (1993) 861–867, [http://dx.doi.org/10.1016/S0893-6080\(05\)80131-5](http://dx.doi.org/10.1016/S0893-6080(05)80131-5).
- [12] J. Niskanen, A. Vladyka, J. Niemi, C.J. Sahle, Emulator-based decomposition for structural sensitivity of core-level spectra, *R. Soc. Open Sci.* 9 (6) (2022) 220093, <http://dx.doi.org/10.1098/rsos.220093>.
- [13] E.A. Eronen, A. Vladyka, F. Gerbon, C.J. Sahle, J. Niskanen, Information bottleneck in peptide conformation determination by X-ray absorption spectroscopy, *J. Phys. Commun.* 8 (2) (2024) 025001, <http://dx.doi.org/10.1088/2399-6528/ad1f73>.
- [14] E.A. Eronen, A. Vladyka, C.J. Sahle, J. Niskanen, Structural descriptors and information extraction from X-ray emission spectra: Aqueous sulfuric acid, *Phys. Chem. Chem. Phys.* (2024) <http://dx.doi.org/10.1039/D4CP02454K>.
- [15] J.B. Kruskal, Toward a practical method which helps uncover the structure of a set of multivariate observations by finding the linear transformation which optimizes a new “index of condensation”, in: R.C. Milton, J.A. Nelder (Eds.), *Statistical Computation*, Academic Press, 1969, pp. 427–440, <http://dx.doi.org/10.1016/B978-0-12-498150-8.50024-0>.
- [16] J.H. Friedman, J.W. Tukey, A projection pursuit algorithm for exploratory data analysis, *IEEE Trans. Comput. C-23* (9) (1974) 881–890, <http://dx.doi.org/10.1109/T-C.1974.224051>.
- [17] P.J. Huber, Projection pursuit, *Ann. Statist.* 13 (2) (1985) 435–475, <http://dx.doi.org/10.1214/aos/1176349519>.
- [18] J.H. Friedman, M. Jacobson, W. Stuetzle, Projection pursuit regression, *J. Amer. Statist. Assoc.* 76 (1981) 817, <http://dx.doi.org/10.2307/2287576>.
- [19] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, high-performance deep learning library, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, vol. 32, Curran Associates, Inc., 2019, pp. 8024–8035, URL <https://dl.acm.org/doi/10.5555/3454287.3455008>.
- [20] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, Í. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nature Methods* 17 (2020) 261–272, <http://dx.doi.org/10.1038/s41592-019-0686-2>.

- [21] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2017, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [22] 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, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830, URL <https://dl.acm.org/doi/10.5555/1953048.2078195>.
- [23] H. Huo, M. Rupp, Unified representation of molecules and crystals for machine learning, *Mach. Learn.: Sci. Technol.* 3 (4) (2022) 045017, <http://dx.doi.org/10.1088/2632-2153/aca005>.
- [24] L. Himanen, M.O.J. Jäger, E.V. Morooka, F. Federici Canova, Y.S. Ranawat, D.Z. Gao, P. Rinke, A.S. Foster, Dscribe: Library of descriptors for machine learning in materials science, *Comput. Phys. Comm.* 247 (2020) 106949, <http://dx.doi.org/10.1016/j.cpc.2019.106949>.
- [25] E.A. Eronen, A. Vladyka, F. Gerbon, C.J. Sahle, J. Niskanen, Data for submission: Information bottleneck in peptide conformation determination by X-ray absorption spectroscopy, 2023, <http://dx.doi.org/10.5281/zenodo.8239300>.
- [26] O. Lingjærde, K. Liestøl, Generalized projection pursuit regression, *SIAM J. Sci. Comput.* 20 (1999) 844–857, <http://dx.doi.org/10.1137/S1064827595296574>.
- [27] A.R. Barron, R.L. Barron, Statistical learning networks: A unifying view, in: *Symposium on the Interface: Statistics and Computing Science*. April, 1988, pp. 21–23.
- [28] N. Intrator, Combining exploratory projection pursuit and projection pursuit regression with application to neural networks, *Neural Comput.* 5 (3) (1993) 443–455, <http://dx.doi.org/10.1162/neco.1993.5.3.443>.
- [29] Y. Zhao, C.G. Atkeson, Implementing projection pursuit learning, *IEEE Trans. Neural Netw.* 7 (2) (1996) 362–373, <http://dx.doi.org/10.1109/72.485672>.



Anton Vladyka, Ph.D. is a postdoctoral researcher at the Department of Physics and Astronomy at the University of Turku, Turku, Finland.



Eemeli A. Eronen is a graduate student at the Department of Physics and Astronomy at the University of Turku, Turku, Finland.



Johannes Niskanen, Ph.D. is a university research fellow at the Department of Physics and Astronomy at the University of Turku, Turku, Finland.