

Batch-efficient EigenDecomposition for Small and Medium Matrices

Yue Song^[0000–0003–1573–5643], Nicu Sebe, and Wei Wang

DISI, University of Trento, Trento 38123, Italy
yue.song@unitn.it
<https://github.com/KingJamesSong/BatchED>

Abstract. EigenDecomposition (ED) is at the heart of many computer vision algorithms and applications. One crucial bottleneck limiting its usage is the expensive computation cost, particularly for a mini-batch of matrices in the deep neural networks. In this paper, we propose a QR-based ED method dedicated to the application scenarios of computer vision. Our proposed method performs the ED entirely by batched matrix/vector multiplication, which processes all the matrices simultaneously and thus fully utilizes the power of GPUs. Our technique is based on the explicit QR iterations by Givens rotation with double Wilkinson shifts. With several acceleration techniques, the time complexity of QR iterations is reduced from $O(n^5)$ to $O(n^3)$. The numerical test shows that for small and medium batched matrices (*e.g.*, $dim < 32$) our method can be much faster than the Pytorch SVD function. Experimental results on visual recognition and image generation demonstrate that our methods also achieve competitive performances.

Keywords: Differentiable SVD, Global Covariance Pooling, Universal Style Transfer, Vision Transformer

1 Introduction

The EigenDecomposition (ED) or the Singular Value Decomposition (SVD) explicitly factorize a matrix into the eigenvalue and eigenvector matrix, which serves as a fundamental tool in computer vision and deep learning. Recently, many algorithms integrated the SVD as a meta-layer into their models to perform some desired spectral transformations [33, 32, 30, 5, 22, 9, 44, 23, 8, 13, 46, 45, 37]. The applications vary in global covariance pooling [30, 43, 37], decorrelated Batch Normalization (BN) [22, 44, 23, 38], Perspective-n-Points (PnP) problems [5, 8, 13], and Whitening and Coloring Transform (WCT) [31, 9, 46].

The problem setup of the ED in computer vision is quite different from other fields. In other communities such as scientific computing, batched matrices rarely arise and the ED is usually used to process a single matrix. However, in deep learning and computer vision, the model takes a mini-batch of matrices as the input, which raises the requirement for an ED solver that works for batched matrices efficiently. Moreover, the differentiable ED works as a building block

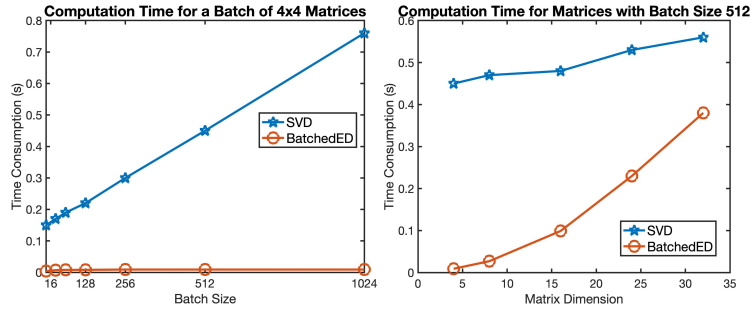


Fig. 1. The speed comparison of our Batched ED against the TORCH.SVD. (*Left*) Time consumption for a mini-batch of 4×4 matrices with different batch sizes. (*Right*) Time consumption for matrices with batch size 512 but in different matrix dimensions.

and needs to process batched matrices millions of times during the training and inference. This poses a great challenge to the efficiency of the ED solver and could even stop people from adding the ED meta-layer in their models due to the huge time consumption (see Fig. 1).

In the current deep learning frameworks such as Pytorch [36] or Tensorflow [1], the ED solvers mainly adopt the SVD implementation from the linear algebra libraries (*e.g.*, LAPACK [3] and Intel MKL [42]). These solvers can efficiently process a single matrix but do not support batched matrices on GPUs well. Most of the implementations are based on the Divide-and-Conquer (DC) algorithm [11, 20]. This algorithm partitions a matrix into multiple small sub-matrices and performs the ED simultaneously for each sub-matrix. Aided by the power of parallel and distributed computing, its speed is only mildly influenced by the matrix dimension and can be very fast for a single matrix. The core of the DC algorithm is the characteristic polynomials $\det(\lambda \mathbf{I} - \mathbf{A}) = 0$, which can be solved by various methods, such as secular equations [20] and spectral division [34]. However, solving the polynomial requires simultaneously localizing all the eigenvalue intervals for each individual matrix. Despite the high efficiency for a single matrix, these DC algorithms do not scale to batched matrices.

Except for the DC algorithm, some ED solvers would use the QR iteration. The QR iteration has many implementation methods and one particular batch-efficient choice is by Givens rotation. The Givens rotation can be implemented via matrix-matrix multiplications, which naturally extends to batched matrices. During the QR iterations, the Givens rotation is applied successively to annihilate the off-diagonal entries until the matrix becomes diagonal. The major drawback limiting the usage of QR iterations is the $O(n^5)$ time cost, which makes this method only applicable to tiny matrices (*e.g.*, $\dim < 9$). To alleviate this issue, modern QR-based ED implementations apply the technique of *deflation* [2, 6, 7], *i.e.*, partition the matrix into many sub-matrices. The deflation technique can greatly improve the speed of the QR iterations but only works for an individual matrix. For the QR iteration, the convergence speed is related with the

adjacent eigenvalue ratio $\frac{\lambda_{i+1}}{\lambda_i}$. For multiple matrices within a mini-batch, the off-diagonal entries of each matrix converge to zero with inconsistent speed and where each matrix can be partitioned is different. Consequently, the deflation technique does not apply to batched matrices either. To give a concrete example, consider 2 matrices of sizes 8×8 in a mini-batch. Suppose that the deflation would split one into two 3×3 and 5×5 matrices, while the other matrix might be partitioned into two 4×4 matrices. In this case, the partitioned matrices cannot be efficiently processed as a mini-batch due to the inconsistent matrix sizes.

To attain a batch-friendly and GPU-efficient ED method dedicated to computer vision field, we propose a QR-based ED algorithm that performs the ED via batched matrix/vector multiplication. Each step of the ED algorithm is carefully motivated for the best batch-efficient and computation-cheap consideration. We first perform a series of batched Householder reflectors to tri-diagonalize the matrix by the batched matrix-vector multiplication. Afterward, the explicit QR iteration by matrix rotation with double Wilkinson shifts [47] is conducted to diagonalize the matrix. The proposed shifts make the last two diagonal entries of the batched matrices have consistent convergence speed. Thereby the convergence is accelerated and the matrix dimension can be progressively shrunk during the QR iterations. Besides the dimension reduction, we also propose some economic computation methods based on the complexity analysis. The time complexity of QR is thus reduced from $O(n^5)$ to $O(n^3)$. The numerical tests demonstrate that, for matrices whose dimensions are smaller than 24, our PyTorch implementation is consistently much faster than the default SVD routine for any batch size. For matrices with larger dimensions (*e.g.*, $dim=32$ or 36), our method could also have an advantage when the batch size is accordingly large (see also Fig. 1). We validate the effectiveness of our method in several applications of differentiable SVD, including decorrelated BN, covariance pooling for vision transformers, and neural style transfer. Our Batched ED achieves competitive performances against the SVD.

The contributions of the paper are summarized threefold:

- We propose an ED algorithm for a mini-batch of small and medium matrices which is dedicated to many application scenarios of computer vision. Each step of ED is carefully motivated and designed for the best batch efficiency.
- We propose dedicated acceleration techniques for our Batched ED algorithm. The progressive dimension shrinkage is proposed to reduce the matrix size during the iterations, while some economic computation methods grounded on the complexity analysis are also developed.
- Our batch-efficient ED algorithm is validated in several applications of differentiable SVD. The experiments on visual recognition and image generation demonstrate that our method achieves very competitive performances against the SVD encapsulated in the current deep learning platforms.

2 Related Work

In this section, we discuss the related work in computing the differentiable ED and its applications.

2.1 Computing the Differentiable ED

To perform the ED, modern deep learning frameworks (*e.g.*, Pytorch and Tensorflow) call the LAPACK’s SVD routine by default. The routine uses the Divide-and-Conquer algorithm [11, 20] to conduct the ED. Assisted by the power of parallel and distributed computing, the divide-and-conquer-based ED can simultaneously process each sub-matrix and achieve high efficiency for a single matrix regardless of the matrix size. However, solving the core characteristic polynomials requires simultaneously finding all the eigenvalue intervals for each individual matrix, which causes this algorithm unable to scale to batched matrices well. There are also some routines that use QR iterations with deflation for performing the ED [6, 7]. Equipped with the deflation technique to partition the matrices, the QR iteration can also have a fast calculation speed. When it comes to a mini-batch of matrices, the off-diagonal entries of each matrix converge to zero with different speeds, and where each matrix can be partitioned is inconsistent. Hence, the deflation technique cannot be applied to batched matrices.

For the back-propagation of the ED, it suffers from the numerical instability caused by the close and repeated eigenvalues. Recently, several methods have been proposed to solve the instability issue [44, 45, 37]. Wei *et al.* [44] propose to use Power Iteration (PI) to approximate the SVD gradients. Song *et al.* [37] propose to use Padé approximants to closely estimate the gradients. Despite the applicability of these methods, a more practical approach is to divide the features $\mathbf{X} \in \mathbb{R}^{C \times BHW}$ into groups $\mathbf{X} \in \mathbb{R}^{G \times \frac{C}{G} \times BHW}$ in the channel dimension and attain a mini-batch of small covariance matrices $\mathbf{X}\mathbf{X}^T \in \mathbb{R}^{G \times \frac{C}{G} \times \frac{C}{G}}$ [22, 35], which can keep more channel statistics and naturally avoid the gradient explosion issue caused by the rank-deficiency. This also raises the need of such an ED solver that works for batched matrices efficiently.

To attain the batch-efficient ED algorithm dedicated to computer vision field, we propose our QR-based algorithm for small and medium batched matrices. We motivate each step of our ED algorithm for the best batch efficiency. Our ED solver integrates double Wilkinson shifts [47] to guarantee that the last two diagonal entries have consistent convergence speed within the mini-batch, and consequently the matrix dimension can be progressively reduced. With several other acceleration technique grounded on the complexity analysis, our solver can be much faster than Pytorch SVD for a mini-batch of small matrices.

2.2 Applications of the Differentiable ED

The need for differentiable ED arises in numerous applications of computer vision. Some methods adopt the end-to-end ED to compute the matrix square root of the global covariance feature before the fully-connected layer [33, 30, 29, 43, 37,

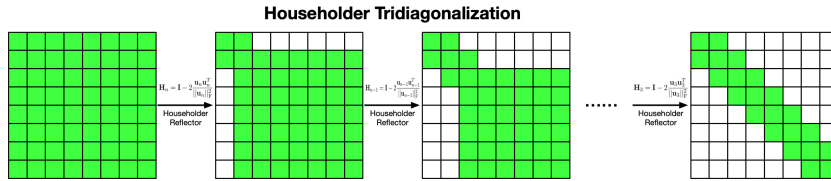


Fig. 2. Visual illustration of our batched Householder tri-diagonalization. After $(n-2)$ designed reflections, the symmetric matrix \mathbf{A} is reduced to a tri-diagonal matrix \mathbf{T} .

48, 16, 39, 40]. Such approaches are termed as Global Covariance Pooling (GCP) methods, and they have achieved state-of-the-art performances on both generic and fine-grained visual recognition. Another line of research uses the ED to perform the decorrelated batch normalization (BN) [22, 25, 35, 23, 38, 24, 38]. The process resembles the ZCA whitening transform to compute the inverse square root for eliminating the correlation between features. The differentiable ED can be also applied in the area of neural style transfer. As pointed out in [17, 18], the feature covariance naturally embeds the style information. Some methods use the differentiable ED to perform successive WCT on the feature covariance for the universal style transfer [31, 9, 10]. In the geometric vision, the ED is often applied to solve the PnP problem and estimate the camera pose [28, 5, 8, 13]. Besides the main usages above, there are some other minor applications [12, 41].

3 Methodology

In this section, we present our method that performs the Batched ED. Our algorithm is implemented via the sequential batched Householder reflectors to tri-diagonalize the matrix and the batched QR iteration to diagonalize the tri-diagonal matrix. Both processes are GPU-friendly and batch-efficient. Now we illustrate each process in detail. Notice that every step is applied on batched matrices for the best efficiency.

3.1 Batched Tri-diagonalization based on Householder Reflection

Given the Hermitian matrix \mathbf{A} , the tri-diagonalization process is defined as:

$$\mathbf{A} = \mathbf{P}\mathbf{T}\mathbf{P}^T \quad (1)$$

where \mathbf{T} is a tri-diagonal matrix, and \mathbf{P} is an orthogonal matrix. To perform such an orthogonal similarity transform, we can decompose \mathbf{P} into $n-2$ Householder reflectors. This leads to the re-formulation:

$$\mathbf{T} = \mathbf{P}^T \mathbf{A} \mathbf{P} = (\mathbf{H}_n \dots \mathbf{H}_4 \mathbf{H}_3)^T \mathbf{A} (\mathbf{H}_n \dots \mathbf{H}_4 \mathbf{H}_3) \quad (2)$$

Each reflector is both orthogonal ($\mathbf{H}\mathbf{H}^T = \mathbf{I}$) and unitary ($\mathbf{H} = \mathbf{H}^T$). The reflector is constructed using an vector:

$$\mathbf{H} = \mathbf{I} - 2 \frac{\mathbf{u}\mathbf{u}^T}{\|\mathbf{u}\|_F^2} \quad (3)$$

The matrix \mathbf{H} reflects the vector \mathbf{u} along the direction that is perpendicular to the hyper-plane orthogonal to \mathbf{u} . This property can be used to tri-diagonalize a symmetric matrix by reflecting each row and column sequentially. A Householder reflection is computed by:

$$\begin{aligned}\mathbf{H}\mathbf{A}\mathbf{H} &= (\mathbf{I} - 2\frac{\mathbf{u}\mathbf{u}^T}{\|\mathbf{u}\|_{\mathbb{F}}^2})\mathbf{A}(\mathbf{I} - 2\frac{\mathbf{u}\mathbf{u}^T}{\|\mathbf{u}\|_{\mathbb{F}}^2}) \\ &= \mathbf{A} - \mathbf{p}\mathbf{u}^T - \mathbf{u}\mathbf{q}^T\end{aligned}\quad (4)$$

where the temporary variables \mathbf{q} , \mathbf{p} , and K are defined as:

$$\mathbf{q} = \mathbf{p} - K\mathbf{u}, \mathbf{p} = \frac{2\mathbf{A}\mathbf{u}}{\|\mathbf{u}\|_{\mathbb{F}}^2}, K = \frac{\mathbf{u}^T\mathbf{p}}{\|\mathbf{u}\|_{\mathbb{F}}^2}\quad (5)$$

As can be seen, Eq. (4) actually defines a symmetric rank-2 update on \mathbf{A} . By some deductions on Eq. (4), each Householder reflector can be designed to introduce zero entries to a row and a column (see Fig. 2). We omit the derivation of the vector for conciseness and give the result here:

$$\begin{aligned}\mathbf{u}_i &= [0, \dots, a_{i,i}, a_{i,i+1}, \dots, a_{i,n-1}, \sigma], \\ \sigma &= \pm\sqrt{a_{i,i}^2 + a_{i,i+1}^2 + \dots + a_{i,n-1}^2}\end{aligned}\quad (6)$$

where $a_{i,j}$ denotes the entry of \mathbf{A} at i -th row and j -th column, and the sign of σ is usually chosen as $\text{sign}(a_{i,n})$ to reduce the round-off error. By such a construction, only $n-2$ reflections are needed to transform the symmetric matrix into the tri-diagonal form. Each householder reflection needs 2 matrix-matrix multiplication, which takes $O(2n^3)$ complexity. However, as indicated in Eq. (4) and Eq. (5), the calculation can be reduced to one matrix-matrix multiplication and two matrix-vector multiplication, which needs the complexity of $O(n^3+2n^2)$.

When the eigenvector is required, we can calculate \mathbf{P} by accumulating the Householder reflectors:

$$\mathbf{P} = \mathbf{H}_n \dots \mathbf{H}_4 \mathbf{H}_3\quad (7)$$

The computation needs $(n-2)$ matrix multiplication where the complexity of each multiplication is $O(n^3)$. We note this step can be further accelerated by:

Theorem 1 (WY representation [4]) *For any accumulation of m Householder matrices $\mathbf{H}_1 \dots \mathbf{H}_m$, there exists $\mathbf{W}, \mathbf{Y} \in \mathbb{R}^{(n-2) \times m}$ such that we have the relation $\mathbf{I} - 2\mathbf{W}\mathbf{Y}^T = \mathbf{H}_1 \dots \mathbf{H}_m$. Computing \mathbf{W} and \mathbf{Y} takes $O((n-2)m)$ time and $(m-1)$ sequential Householder multiplications.*

Relying on this theorem, we can divide the accumulation $\mathbf{H}_n \dots \mathbf{H}_4 \mathbf{H}_3$ into $(n-2)/m$ sub-sequences and compute them in parallel. Each sub-sequence takes $O((m-1)n^3 + (n-2)m)$ time to compute the WY representation and $O((n-2)^2 m)$ time to compute $\mathbf{I} - 2\mathbf{W}\mathbf{Y}^T$. Combining all the sub-sequence needs extra time of $O((n-2)^3/m)$. This can further reduce the complexity of computing \mathbf{P} from $O((n-3)n^3)$ to $O((m-1)n^3 + (n-2)m + (n-2)^2 m + (n-2)^3/m)$. The computation saving would be huge when n is large.

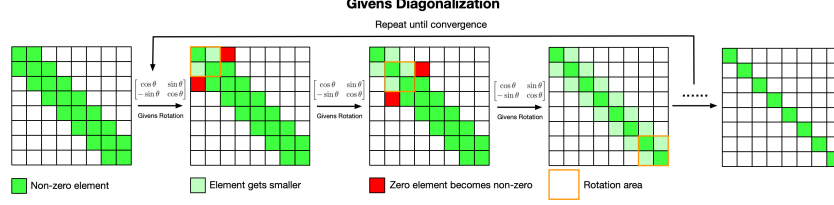


Fig. 3. Visual illustration of the batched Givens diagonalization. For each QR iteration, the Givens rotation is applied from the left top corner to the bottom right corner to reduce the magnitude of the off-diagonal elements. The iteration continues till the off-diagonal entries become zero or below a certain tolerance.

3.2 Batched Diagonalization based on QR Iteration

After obtaining the tri-diagonal matrix \mathbf{T} , we use the Givens rotation to perform the QR iterations, which can be implemented efficiently via batched matrix multiplication. Based on the ordinary QR iteration, we further apply several techniques to speed up the convergence and save the computational budget.

Basic QR Iteration by Givens Rotation. Given the tri-diagonal matrix \mathbf{T} , the QR iteration takes the following iterative update:

$$\mathbf{T}_k = \mathbf{Q}_k \mathbf{R}_k, \quad \mathbf{T}_{k+1} = \mathbf{R}_k \mathbf{Q}_k \quad (8)$$

where \mathbf{Q}_k denotes the orthogonal matrix, and \mathbf{R}_k is the upper-triangular matrix. Replacing \mathbf{R}_k with $\mathbf{Q}_k^T \mathbf{T}_k$ leads to the re-formulation of Eq. (8) :

$$\mathbf{T}_{k+1} = \mathbf{Q}_k^T \mathbf{T}_k \mathbf{Q}_k \quad (9)$$

As can be seen, a single QR iteration is equivalent to performing an orthogonal similarity transform. By performing the iterations, the sub-diagonal and super-diagonal entries are gradually reduced till the matrix becomes diagonal.

For each QR iteration, we construct the orthogonal transform using successive Givens rotations moving from the left top corner to the right bottom corner. The 2×2 Givens Rotation and its $n \times n$ extension are defined by:

$$\mathbf{R}_{2 \times 2} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}, \quad \mathbf{R}_{n \times n} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{2 \times 2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \quad (10)$$

where θ is the rotation angle, and the rotation matrix is orthogonal but not symmetric (*i.e.*, $\mathbf{R}^T \mathbf{R} = \mathbf{I}$ and $\mathbf{R}^T \neq \mathbf{R}$). As shown in Fig. 3, by design of the rotation angle, the successive Givens Rotation applied on \mathbf{T} can keep the tri-diagonal form but reduce the magnitude of off-diagonal elements. For the derivation of Givens rotation, please refer to the supplementary material for detail. The sequential Givens rotations moving along the diagonal form one single QR iteration:

$$\mathbf{T}_{k+1} = (\mathbf{R}_{n-2}^T \dots \mathbf{R}_0^T) \mathbf{T}_k (\mathbf{R}_0 \dots \mathbf{R}_{n-2}) \quad (11)$$

where \mathbf{R}_i denotes the i -th rotation counting from the left top corner. For the orthogonal matrix \mathbf{Q}_i in the i -th QR iteration, we can easily find out

$$\mathbf{Q}_k = \mathbf{R}_0 \dots \mathbf{R}_{n-2} \quad (12)$$

Taking the Householder tri-diagonalization and Givens diagonalization together, our batch-efficient ED algorithm can be formally defined by:

$$\mathbf{A} = (\mathbf{P}\mathbf{Q}_0 \dots \mathbf{Q}_k)\mathbf{\Lambda}(\mathbf{P}\mathbf{Q}_0 \dots \mathbf{Q}_k)^T \quad (13)$$

where k is the iteration times of the QR iteration, $\mathbf{\Lambda}$ is the eigenvalue matrix, and $\mathbf{P}\mathbf{Q}_0 \dots \mathbf{Q}_k$ is the eigenvector matrix. For the convergence, we have:

Theorem 2 (Convergence of QR iteration) *Let \mathbf{T} be the positive definite tri-diagonal matrix with the eigendecomposition $\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$ and assume \mathbf{Q}^T can be LU decomposed. Then the QR iteration of \mathbf{T} will converge to $\mathbf{\Lambda}$.*

We defer the proof to the Supplementary Material. The key results of this theorem is that the convergence speed depends on the adjacent eigenvalue ratio $\frac{\lambda_i}{\lambda_j}$ for $i > j$. The QR iterations usually take $2n$ iterations to make the resultant matrix diagonal [15]. Consider the fact that each iteration takes $(n-1)$ Givens rotation. The computation overhead would be huge. For deriving the eigenvalues, we need $4n(n-1)n^3$ time, while it takes the complexity of $2n(n-2)n^3 + (2n-1)n^3$ to compute the eigenvector. The time complexity of the QR iteration is quintic to the matrix dimension n , which would make this method only applicable to the tiny matrices (< 9). Existing deflation techniques [6, 7] to accelerate the computation cannot be applied to our batched matrices. To resolve this issue, we propose the following techniques:

Double Wilkinson Shift. As indicated in Theorem 2, the convergence speed of QR iteration depends on the ratio $\frac{\lambda_i}{\lambda_j}$, where $i > j$. A natural approach to accelerate the convergence is to shift the matrix by $\mathbf{T} - \mu\mathbf{I}$ such that the convergence speed becomes $\frac{\lambda_i - \mu}{\lambda_j - \mu}$. A preferable shift coefficient should be $u = \lambda_j$, as this can help the matrices to converge quickly: $\frac{\lambda_i - \mu}{\lambda_j - \mu} = \infty$. This is particularly useful for matrices in a mini-batch as the speed can be made consistent by shifting.

Since each Givens rotation will affect the area rotated by the previous one, only the last 2×2 Givens rotation will not be influenced, *i.e.*, the two eigenvalues of the last block can be locally estimated. Thus, we propose to extract the shift coefficients from the 2×2 block on the right bottom corner:

$$\mu_{n-2}, \mu_{n-1} = \text{Wilkinson}(\mathbf{T}_k[n-2:n]) \quad (14)$$

where $\mathbf{T}_k[n-2:n]$ denotes the last 2×2 block of \mathbf{T}_k , and μ_{n-2} and μ_{n-1} are the two eigenvalues computed from this block. These shifting coefficients are referred to as the Wilkinson shift [47], and we give the derivation in the supplementary

material. After attaining the shifts, we can reformulate the QR iterations with double shifts:

$$\begin{aligned}\mathbf{T}_{k+1/2} &= \mathbf{Q}_k^T (\mathbf{T}_k - \mu_{n-1} \mathbf{I}) \mathbf{Q}_k + \mu_{n-1} \mathbf{I} \\ \mathbf{T}_{k+1} &= \mathbf{Q}_k^T (\mathbf{T}_{k+1/2} - \mu_{n-2} \mathbf{I}) \mathbf{Q}_k + \mu_{n-2} \mathbf{I}\end{aligned}\quad (15)$$

With the shifts, the integrated iteration consists of two sequential QR iterations shifted by the eigenvalues μ_{n-2} and μ_{n-1} , respectively.

Progressive Dimension Shrinkage. One direct benefit brought by the Wilkinson shift is that, for all the matrices in a mini-batch, the last two diagonal entries can quickly converge to the corresponding eigenvalues and the off-diagonal elements can converge to zero:

$$\begin{aligned}t_{n-2,n-2} &\rightarrow \lambda_{n-2}, & t_{n-2,n-3} &= t_{n-3,n-2} \rightarrow 0 \\ t_{n-1,n-1} &\rightarrow \lambda_{n-1}, & t_{n-2,n-1} &= t_{n-1,n-2} \rightarrow 0\end{aligned}\quad (16)$$

We can use this property to speed up the computation by gradually reducing the matrix dimension, *i.e.*, shrinking the matrix by $\mathbf{T} \in \mathbb{R}^{n \times n} \rightarrow \mathbf{T} \in \mathbb{R}^{(n-1) \times (n-1)}$ after one iteration. As shown in Fig. 4, when the last sub-diagonal entry is below a given small threshold (*e.g.*, $1e-5$), we could shrink the matrix by removing the last row and column. In doing so, the matrix size is progressively reduced during the QR iterations. With the dimension reduction, one QR iteration would take $(n-1-r)$ Givens rotations, where r is the reduction times.

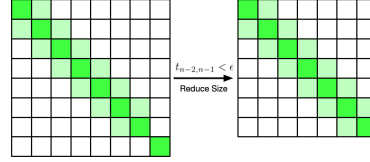


Fig. 4. Illustration of the progressive dimension reduction in the QR iterations. After one iteration, if the last sub-diagonal entry is below a small threshold ϵ , we can remove the last row and column.

Economic Eigenvalue Calculation. For a Givens rotation, it only affects the adjacent 4×4 block. We can save the computation budget by applying the matrix multiplication on the 4×4 rotation region in the neighborhood. This reduces the time of a rotation from $O(2n^3)$ to $O(2 \times 4^3) = O(128)$, which makes each rotation consume a constant time cost. Taking the above dimension reduction into account, the QR iterations need $O(256n(n-1-r))$ time to derive the eigenvalues.

Economic Eigenvector Calculation. Equipped with the progressive dimension reduction, the orthogonal transform \mathbf{Q}_k in a QR iteration is defined by:

$$\mathbf{Q}_k = \mathbf{R}_0 \mathbf{R}_1 \dots \mathbf{R}_{n-2-r} \quad (17)$$

where we need $(n-2-r)$ rotations for each iteration. The computation can be potentially simplified by the theorem:

Theorem 3 (Implicit Q Theorem [19]) *Let \mathbf{B} be an upper Hessenberg and only have positive elements on its first sub-diagonal. Assume there exists a unitary transform $\mathbf{Q}^H \mathbf{A} \mathbf{Q} = \mathbf{B}$. Then \mathbf{Q} and \mathbf{B} are uniquely determined by \mathbf{A} and the first column of \mathbf{Q} .*

We give the proof and some discussion in the supplementary document. This theorem implies that, without the need of explicit QR iteration, the orthogonal transform \mathbf{Q} and the transformed matrix \mathbf{B} can be both implicitly calculated. However, it assumes that the sub-diagonal elements of \mathbf{B} are positive. In our case, the Givens rotation can easily zero out the last two sub-diagonal entries. Consequently, directly using the theorem would cause large round-off errors and data overflow.

Although the theorem cannot be directly applied, it allows us to simplify the eigenvector calculation. As indicated by the theorem, the i -th rotation would only affect the orthogonal matrix \mathbf{Q} on the area after the i -th row and column. We can reduce the computation by involving only part of the matrix and simplify the calculation in Eq. (17) as:

$$\mathbf{Q}_k = \mathbf{R}_0[1:] \mathbf{R}_1[2:] \dots \mathbf{R}_{n-2-r}[n-2-r:] \quad (18)$$

where $[i:]$ denotes part of the matrix that excludes the first i rows and columns. By doing so, the time complexity of calculating \mathbf{Q}_k can be reduced to:

$$\begin{aligned} & (n-2-r)^2 n + (n-3-r)^2 n + \dots + 1^2 n \\ &= \sum_{i=1}^{n-2-r} i^2 n = \frac{(n-2-r)(n-1-r)(2n-3-2r)}{6} n \end{aligned} \quad (19)$$

Compared with the original time cost $O((n-2-r)n^3)$, the saving would be considerable for large n and r .

3.3 Computation Complexity Summary

Table 1. Comparison of time complexity of the basic QR-based ED solver and our ED solver dedicated for batched matrices. Here n denotes the matrix size and r represents the average reduction times during the QR iterations.

Time	Basic QR-based ED Solver	
	Eigenvalue	Eigenvector
Tri-diag.	$2n^3$	$(n-3)n^3$
QR	$4n(n-1)n^3$	$2n(n-2)n^3 + (2n-1)n^3$
Sum	$(4n^2 - 4n + 2)n^3$	$(2n^2 - n - 4)n^3$
Time	Our Batched ED Solver	
	Eigenvalue	Eigenvector
Tri-diag.	$n^3 + 2n^2$	$(m-1)n^3 + (n-2)m + (n-2)^2 m + \frac{(n-2)^3}{m}$
QR	$256(n-1-r)n$	$\frac{(n-2-r)(n-1-r)(2n-3-2r)}{6} 2n^2 + (2n-1)n^3$
Sum	$n^3 + 258n^2 - 256n(1+r)$	$-\left(\frac{3}{2}r^3 + 3r^2 + \frac{13}{3}r + 2 - m\right)n^2 - \left(3m - \frac{3}{m}\right)n + 2m - \frac{6}{m}$

Table 1 summarizes the time complexity of the basic QR-based ED solver and our proposed ED solver dedicated for batched matrices. Taking the highest-order

term for simpler analysis, our ED solver reduces the time from $O(4n^5)$ to $O(n^3)$ for computing the eigenvalues, and saves the time from $O(2n^5)$ to $O(\frac{2}{3}n^5)$ for eigenvectors. Moreover, depending on the reduction times r , the complexity can be further reduced with the term $-256(1+r)n$ for eigenvalues and $-(2r+1)n^4$ for eigenvectors.

3.4 Convergence and Error Bounds

For the tri-diagonalization process, the convergence is guaranteed with $n-2$ Householder reflectors. The error is only related to the machine precision and data precision, which can be sufficiently neglected. For the QR iterations, the convergence mainly depends on the adjacent eigenvalue ratio $\frac{\lambda_{i+1}}{\lambda_i}$ and the shift μ . In certain cases when the two eigenvalue are close ($\frac{\lambda_{i+1}}{\lambda_i} \approx 1$), the convergence speed is slow and the residual term $(\frac{\lambda_{i+1}-\mu}{\lambda_i-\mu})^{2n}$ becomes the error. Another error source comes from the tolerance ϵ for the dimension reduction. Let $\bar{\mathbf{\Lambda}}$ represent the exact eigenvalues and $\mathbf{\Lambda}$ denote the eigenvalues calculated by our ED solver. Then the error is bounded by:

$$\|\bar{\mathbf{\Lambda}} - \mathbf{\Lambda}\|_F \leq \max_i \left(\left(\frac{\lambda_{i+1} - \mu}{\lambda_i - \mu} \right)^{2n} |l_{i+1,i}| \right) + \epsilon \quad (20)$$

where $l_{i+1,i}$ is the entry of \mathbf{L} computed by $\mathbf{Q}^T = \mathbf{L}\mathbf{U}$, and the shift μ changes every QR iteration. Since \mathbf{Q} is orthogonal, the magnitude of $l_{i+1,i}$ is often quite small. Considering the small magnitude of $l_{i+1,i}$ and the additional shifting technique, the accuracy of our method will not get affected.

4 Experiments

In this section, we first perform a numerical test to compare our method with SVD for matrices in different dimensions and batch sizes. Subsequently, we evaluate the effectiveness of the proposed methods in three computer vision applications: decorrelated BN, second-order vision transformer, and neural style transfer. The implementations details are referred to supplementary material.

4.1 Numerical Test

Fig. 5 depicts the computational time of our Batched ED against the SVD for different matrix dimensions and batch sizes. The time cost of the SVD grows almost linearly with the batch size, while the time consumption of our Batched ED only has slight or mild changes against varying batch sizes. For matrices whose dimensions are smaller than 24, our Batched ED is consistently faster than the SVD for any batch size. When the matrix dimension is 32, our method is faster than the SVD from batch size 256 on. The speed of our Batched ED is more advantageous for smaller matrix dimensions and larger batch sizes.

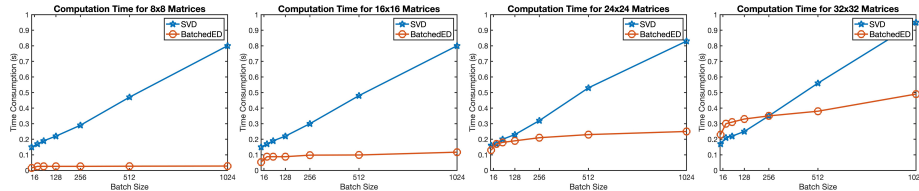


Fig. 5. The speed comparison of our Batched ED against TORCH.SVD for different batch sizes and matrix dimensions. Our implementation is more batch-friendly and the time cost does not vary much against different batch sizes. For matrices in small and moderate sizes, our method can be significantly faster than the Pytorch SVD.

Table 2. Validation error of decorrelated BN on ResNet-18 [21]. The results are reported based on 5 runs, and we measure the time of the forward ED in a single step.

Solver	Group	Size	Time (s)	CIFAR10		CIFAR100	
				mean±std	min	mean±std	min
SVD	16	16×4×4	0.172	4.52±0.09	4.33	21.24±0.17	20.99
Batched ED			0.006	4.37±0.11	4.29	21.25±0.20	20.90
SVD	8	8×8×8	0.170	4.55±0.13	4.34	21.32±0.31	20.88
Batched ED			0.016	4.36±0.11	4.25	20.97±0.27	20.62
SVD	4	4×16×16	0.165	4.52±0.14	4.33	21.30±0.33	20.86
Batched ED			0.075	4.45±0.11	4.32	21.19±0.21	20.98

4.2 Decorrelated BN

Following [38], we first conduct an experiment on the task of ZCA whitening. In the whitening process, the inverse square root of the covariance is multiplied with the feature as $(\mathbf{X}\mathbf{X}^T)^{-\frac{1}{2}}\mathbf{X}$ to eliminate the correlation between each dimension. We insert the ZCA whitening meta-layer into the ResNet-18 [21] architecture and evaluate the validation error on CIFAR10 and CIFAR100 [27]. Table 2 compares the performance of our Batched ED against the SVD. Depending on the number of groups, our method can be 2X faster, 10X faster, and even 28X faster than the SVD. Furthermore, our method outperforms the SVD across all the metrics on CIFAR10. With CIFAR100, the performance is also on par.

4.3 Second-order Vision Transformer

We turn to the experiment on the task of global covariance pooling for the Second-order Vision Transformer (So-ViT) [48]. To leverage the rich semantics embedded in the visual tokens, the covariance square root of the visual tokens $(\mathbf{X}\mathbf{X}^T)^{\frac{1}{2}}$ are used to assist the classification task. Since the global covariance matrices are typically very ill-conditioned [37], this task poses a huge challenge to the stability of the ED algorithm. We choose the So-ViT architecture with different depths and validate the performance on ImageNet [14]. As observed from Table 3, our Batched ED has the competitive performance against the standard SVD. Moreover, our method is about 44% and 27% faster than the SVD for covariance in different sizes.

Table 3. Validation accuracy on ImageNet [14] for the second-order vision transformer with different depths. Here 32 and 36 denote the spatial dimension of visual tokens. We report the time consumption of the forward ED in a single step.

Solver	Size	Time (s)	Architecture	
			So-ViT-7	So-ViT-10
SVD	$768 \times 32 \times 32$	0.767	76.01 / 93.10	77.97 / 94.10
Batched ED	$768 \times 32 \times 32$	0.431	76.04 / 93.05	77.91 / 94.08
SVD	$768 \times 36 \times 36$	0.835	76.10 / 93.14	78.09 / 94.13
Batched ED	$768 \times 36 \times 36$	0.612	76.07 / 93.10	78.11 / 94.19




Fig. 6. Exemplary visual comparison. The red circle/rectangular indicates the region with subtle details. In this example, our method generates sharper images with more coherent style information and less artifacts. Zoom in for a better view.

4.4 Universal Style Transfer

Now we apply our Batched ED in the WCT for neural style transfer. Given the content feature \mathbf{X}_c and the style feature \mathbf{X}_s , the WCT performs successive whitening $((\mathbf{X}_c \mathbf{X}_c)^{-\frac{1}{2}} \mathbf{X}_c)$ and coloring $((\mathbf{X}_s \mathbf{X}_s)^{\frac{1}{2}} \mathbf{X}_c)$ to transfer the target style. We follow [31, 46] to use the LPIPS distance and the user preference as the evaluation metrics. Table 4 presents the quantitative comparison with different groups. Our Batched ED achieves very competitive performance and predominates the speed. To give a concrete example, when the group number is 64, our method is about 35X faster than the default SVD. Fig. 6 displays the exemplary visual comparison. In this specific example, our Batched ED generates images with better visual appeal.

Similar to the finding in [9], we also observe that the number of groups has an impact on the extent of transferred style. As shown in Fig. 7, when more groups are used, the style in the transferred image becomes more distinguishable and the details are better preserved. Since the number of groups determines the number of divided channels and the covariance size, more groups correspond to smaller covariance and this might help to better capture the local structure. Despite this superficial conjecture, giving a more comprehensive and rigorous analysis is worth further research.

Table 4. The LPIPS distance between the transferred image and the content image and the user preference (%) on the Artworks [26] dataset. We report the time consumption of the forward ED that is conducted 10 times to exchange the style and content feature at different network depths. The batch size is set to 4.

Solver	Group	Size	Time (s)	LPIPS [49] (\uparrow)	Preference (\uparrow)
SVD	64	$256 \times 4 \times 4$	3.146	0.5776	48.25
Batched ED			0.089	0.5798	47.75
SVD	32	$128 \times 8 \times 8$	2.306	0.5722	47.75
Batched ED			0.257	0.5700	48.75
SVD	16	$64 \times 16 \times 16$	1.973	0.5614	46.25
Batched ED			0.876	0.5694	47.75




Fig. 7. Visual illustration of the impact of groups. When more groups are used, the strength of the target style is increased and the details are better preserved.

To sum up, our ED solver has demonstrated the superior batch efficiency for small matrices in various real-world experiments and numerical tests. The limitation on large matrices indicates the key difference: *our method is more batch-efficient, while TORCH.EIG/SVD is more dimension-efficient.*

5 Conclusion

In this paper, we propose a batch-efficient QR-based ED algorithm dedicated for batched matrices which are common in the context of computer vision and deep learning. Aided by the proposed acceleration techniques, our solver is much faster than Pytorch SVD function for a mini-batch of small and medium matrices. Our method can directly benefit a wide range of computer vision applications and we showcase this merit in several applications of differentiable ED. Extensive experiments on visual recognition and image generation demonstrate that our method can also achieve very competitive performances.

Acknowledgement This research was supported by the EU H2020 projects AI4Media (No. 951911) and SPRING (No. 871245).

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al.: {TensorFlow}: A system for {Large-Scale} machine learning. In: 12th USENIX symposium on operating systems design and implementation (OSDI 16). pp. 265–283 (2016)
2. Ahues, M., Tisseur, F.: A new deflation criterion for the qr algorithm. LAPACK Working Note **122** (1997)
3. Anderson, E., Bai, Z., Bischof, C., Blackford, L.S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., et al.: LAPACK Users’ guide. SIAM (1999)
4. Bischof, C., Van Loan, C.: The wy representation for products of householder matrices. *SIAM Journal on Scientific and Statistical Computing* **8**(1), s2–s13 (1987)
5. Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C.: Dsac-differentiable ransac for camera localization. In: CVPR (2017)
6. Braman, K., Byers, R., Mathias, R.: The multishift qr algorithm. part i: Maintaining well-focused shifts and level 3 performance. *SIAM Journal on Matrix Analysis and Applications* **23**(4), 929–947 (2002)
7. Braman, K., Byers, R., Mathias, R.: The multishift qr algorithm. part ii: Aggressive early deflation. *SIAM Journal on Matrix Analysis and Applications* **23**(4), 948–973 (2002)
8. Campbell, D., Liu, L., Gould, S.: Solving the blind perspective-n-point problem end-to-end with robust differentiable geometric optimization. In: ECCV (2020)
9. Cho, W., Choi, S., Park, D.K., Shin, I., Choo, J.: Image-to-image translation via group-wise deep whitening-and-coloring transformation. In: CVPR (2019)
10. Choi, S., Jung, S., Yun, H., Kim, J.T., Kim, S., Choo, J.: Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In: CVPR (2021)
11. Cuppen, J.J.: A divide and conquer method for the symmetric tridiagonal eigenproblem. *Numerische Mathematik* **36**(2), 177–195 (1980)
12. Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-order attention network for single image super-resolution. In: CVPR (2019)
13. Dang, Z., Yi, K.M., Hu, Y., Wang, F., Fua, P., Salzmann, M.: Eigendecomposition-free training of deep networks for linear least-square problems. TPAMI (2020)
14. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
15. Francis, J.G.: The qr transformation—part 2. *The Computer Journal* **4**(4), 332–345 (1962)
16. Gao, Z., Wang, Q., Zhang, B., Hu, Q., Li, P.: Temporal-attentive covariance pooling networks for video recognition. In: NeurIPS (2021)
17. Gatys, L., Ecker, A.S., Bethge, M.: Texture synthesis using convolutional neural networks. NeurIPS (2015)
18. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: CVPR (2016)
19. Golub, G.H., Van Loan, C.F.: *Matrix computations*. edition (1996)
20. Gu, M., Eisenstat, S.C.: A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. *SIAM Journal on Matrix Analysis and Applications* **16**(1), 172–191 (1995)
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)

22. Huang, L., Yang, D., Lang, B., Deng, J.: Decorrelated batch normalization. In: CVPR (2018)
23. Huang, L., Zhao, L., Zhou, Y., Zhu, F., Liu, L., Shao, L.: An investigation into the stochasticity of batch whitening. In: CVPR (2020)
24. Huang, L., Zhou, Y., Liu, L., Zhu, F., Shao, L.: Group whitening: Balancing learning efficiency and representational capacity. In: CVPR (2021)
25. Huang, L., Zhou, Y., Zhu, F., Liu, L., Shao, L.: Iterative normalization: Beyond standardization towards efficient whitening. In: CVPR (2019)
26. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
27. Krizhevsky, A.: Learning multiple layers of features from tiny images. Master’s thesis, University of Tront (2009)
28. Lepetit, V., Moreno-Noguer, F., Fua, P.: Epnnp: An accurate $O(n)$ solution to the pnp problem. IJCV (2009)
29. Li, P., Xie, J., Wang, Q., Gao, Z.: Towards faster training of global covariance pooling networks by iterative matrix square root normalization. In: CVPR (2018)
30. Li, P., Xie, J., Wang, Q., Zuo, W.: Is second-order information helpful for large-scale visual recognition? In: ICCV (2017)
31. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In: NeurIPS (2017)
32. Lin, T.Y., Maji, S.: Improved bilinear pooling with cnns. arXiv preprint arXiv:1707.06772 (2017)
33. Lin, T.Y., RoyChowdhury, A., Maji, S.: Bilinear cnn models for fine-grained visual recognition. In: ICCV (2015)
34. Nakatsukasa, Y., Higham, N.J.: Stable and efficient spectral divide and conquer algorithms for the symmetric eigenvalue decomposition and the svd. SIAM Journal on Scientific Computing **35**(3), A1325–A1349 (2013)
35. Pan, X., Zhan, X., Shi, J., Tang, X., Luo, P.: Switchable whitening for deep representation learning. In: ICCV (2019)
36. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. NeurIPS (2019)
37. Song, Y., Sebe, N., Wang, W.: Why approximate matrix square root outperforms accurate svd in global covariance pooling? In: ICCV (2021)
38. Song, Y., Sebe, N., Wang, W.: Fast differentiable matrix square root. In: ICLR (2022)
39. Song, Y., Sebe, N., Wang, W.: Fast differentiable matrix square root and inverse square root. arXiv preprint arXiv:2201.12543 (2022)
40. Song, Y., Sebe, N., Wang, W.: On the eigenvalues of global covariance pooling for fine-grained visual recognition. IEEE TPAMI (2022)
41. Sun, Q., Zhang, Z., Li, P.: Second-order encoding networks for semantic segmentation. Neurocomputing (2021)
42. Wang, E., Zhang, Q., Shen, B., Zhang, G., Lu, X., Wu, Q., Wang, Y.: Intel math kernel library. In: High-Performance Computing on the Intel® Xeon Phi™, pp. 167–188. Springer (2014)
43. Wang, Q., Xie, J., Zuo, W., Zhang, L., Li, P.: Deep cnns meet global covariance pooling: Better representation and generalization. TPAMI (2020)
44. Wang, W., Dang, Z., Hu, Y., Fua, P., Salzmann, M.: Backpropagation-friendly eigendecomposition. In: NeurIPS (2019)
45. Wang, W., Dang, Z., Hu, Y., Fua, P., Salzmann, M.: Robust differentiable svd. TPAMI (2021)

46. Wang, Z., Zhao, L., Chen, H., Qiu, L., Mo, Q., Lin, S., Xing, W., Lu, D.: Diversified arbitrary style transfer via deep feature perturbation. In: CVPR (2020)
47. Wilkinson, J.: The algebraic eigenvalue problem. In: Handbook for Automatic Computation, Volume II, Linear Algebra. Springer-Verlag New York (1971)
48. Xie, J., Zeng, R., Wang, Q., Zhou, Z., Li, P.: So-vit: Mind visual tokens for vision transformer. arXiv preprint arXiv:2104.10935 (2021)
49. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)