

In T. A. Bubba, R. Gaburro, S. Gazzola, K. Papafitsoros, M. Pereyra, C.-B. Schönlieb (Eds.): Scale Space and Variational Methods in Computer Vision. Part II. Lecture Notes in Computer Science, Vol. 15668, Springer, Cham, 324-336, 2025.

Efficient Representations of the Diffusion Echo

Daniel Gaa✉, Joachim Weickert, Iva Farag*, and Özgün Çiçek*

Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science,
Campus E1.7, Saarland University, 66041 Saarbrücken, Germany.
{gaa, weickert, farag, cicek}@mia.uni-saarland.de

Abstract. Diffusion echoes are a fundamental concept for understanding the behaviour of nonlinear diffusion filters. They describe the accumulated data exchange during a diffusion process and are given by the columns or rows of the corresponding state transition matrix. Unfortunately, they involve a prohibitively large amount of data. Therefore, we propose the first compression strategy to efficiently represent and reconstruct diffusion echoes. Using a truncated singular value decomposition (SVD) we can reduce the storage requirements substantially, while still obtaining very accurate reconstructions. The SVD works on all echoes jointly and captures the redundancy between them. To approximate the singular value decomposition, we use a powerful probabilistic approach: the randomised subspace iteration. We show on a test case with a difficult, rapidly decaying diffusivity that we can reduce the storage requirements by a factor of 20, without creating visually noticeable errors. Furthermore, our compressed data representation enables an efficient reconstruction, which allows a fast and detailed echo investigation. This paves the way for various future applications of the diffusion echo that have been prevented by its huge amount of data so far.

Keywords: Diffusion Echoes · Impulse Response · Compression · Singular Value Decomposition.

1 Introduction

The world of linear, shift-invariant (LSI) filtering enjoys beautiful simplicity and transparency: The action of a filter has an explicit representation via a convolution kernel. It is given by the impulse response: the result of filtering an image with a unit impulse in the origin (delta distribution in the continuous setting). A prominent example is Gaussian convolution. Considering the full family of Gaussians with increasing variances σ^2 creates the Gaussian scale-space

* Iva Farag and Özgün Çiçek were affiliated with the Mathematical Image Analysis Group during the time they worked on their student theses.

[13]. It is well-known to be equivalent to a homogeneous diffusion evolution with time $t = \frac{1}{2}\sigma^2$. While its impulse response is time-dependent, it remains space-invariant.

Nonlinear diffusion filters, on the other hand, are different by design: They preserve or even enhance structures of interest such as edges by adapting their diffusivities or diffusion tensors to the evolving local image structure [4,19,22]. Thus, their behaviour is more than just time-variant: It is also space-variant and image-dependent. Nevertheless, it is still desirable to have an insightful representation of the full accumulated filter action at a specified location and time. This has been proposed by Dam and Nielsen under the name *diffusion echo* [6]. They introduce the (source) echo in a pixel as the result of applying the nonlinear diffusion process to a unit impulse located in that pixel.

In theory, diffusion echoes are highly attractive: The entirety of the diffusion echoes in all pixels fully characterises the underlying process: The filtered result can be reconstructed from the echoes and the original image. More importantly, the echoes allow to gain insights into the real effects of a filter, to investigate its behaviour on multiple scales, to explore the deep structure of the nonlinear scale-space in a generic way, or to compare the effects of different parameters, different diffusivities or even of different discretisations. Thus, diffusion echoes offer an intuitive, useful and complete understanding of the nonlinear filter.

Unfortunately, there is a fundamental problem: To benefit from these advantages in practice, fast access to a subset or even all of the echoes may be required. However, the large number of diffusion echoes adds up to a huge amount of data, so storing them becomes very expensive. An image of size $n_x \times n_y$ exhibits $N := n_x n_y$ echoes, each of size N . For an image of size 256×256 this already requires to store roughly $4 \cdot 10^9$ floating point numbers, i.e. 16 gigabytes of data! Alternatively, computing the echoes at the moment they are called for can be expensive, as each single echo requires a full diffusion evolution.

Therefore, it is desirable to have lossy representations of diffusion echoes that require much lower storage while still allowing high-quality reconstructions. Ideally, the reconstruction should be efficient, such that an inspection in short time is possible. Interestingly, examining the echoes of diffusion filters shows that neighbouring echoes are often very similar. This redundancy offers a substantial potential for compression that has not been exploited so far.

Our Contributions. The goal of our work is to address this problem by creating a representation that (a) has low storage requirements and (b) allows for a fast reconstruction with high quality. To this end, we propose the first compression approach for diffusion echoes. Its components are specifically selected to benefit from key properties of the echoes.

We adjust and combine two successful concepts for efficient data representation: a truncated singular value decomposition (SVD) [10] and a powerful probabilistic method for estimating it [11,17,20]. Both concepts are novel in the context of diffusion echoes and require careful engineering. The reconstruction of a decompressed echo from the SVD data is achieved rapidly, since it only re-

quires a single matrix-vector multiplication. We show that our method produces high-quality diffusion echo reconstructions even at high compression rates.

While we describe our method for a prototypical isotropic nonlinear filter with a scalar-valued diffusivity function, it can be used for linear space-variant diffusion filters or anisotropic ones with a diffusion tensor in the same manner.

Related Work. Inpainting-based image compression exploits an idea that resembles the diffusion echo: The so-called *inpainting echo* is used to optimise the grey values of the stored data [14]. It describes the influence of these pixels on the inpainting result. In contrast to the diffusion echo, however, its shape depends on neighbouring data pixels, even for homogeneous linear diffusion.

Several adaptive filters with space-variant smoothing kernels exist that were not initially designed to approximate nonlinear diffusion echo representations. These include e.g. the Nitzberg–Shiota filter [16], the bilateral filter (see e.g. [21]), and nonlocal linear diffusion scale-spaces [3]. Establishing some relations to nonlinear diffusion filters requires specific scaling limits and iterative applications.

Two smoothing methods with adaptive kernels shall be mentioned that intend to approximate nonlinear diffusion filters: Fischl and Schwartz [9] learn a space-variant integral kernel. Dam et al. [7] approximate nonlinear diffusion echoes with adaptively shaped Gaussians, but report quality deteriorations. Both methods do not offer quantifiable and numerically consistent approximation results.

It should be emphasised that the diffusion echo gives an *exact* representation of the nonlinear diffusion result. It remains error-free under a principle component analysis (PCA) [18] or the SVD [10], as long as no components are discarded. By removing components to increase efficiency, one introduces errors in a controllable and quantifiable way. We follow this strategy and combine the truncated SVD [10] with a randomised subspace iteration [11,17,20]. This enables us to benefit from fundamental tools for low-rank matrix approximation problems with well-understood qualities.

There are two unpublished student theses by co-authors of this paper [2,5] that represent diffusion echoes in terms of a PCA. While the PCA is related to the SVD, these theses work on a subset of the diffusion echoes or split the image domain to create smaller subproblems. Here we do not require such adjustments and can work efficiently on the full data.

Paper Organisation. We review essential diffusion concepts in Section 2 and discuss the concept of the diffusion echo in Section 3. Section 4 is concerned with the singular value decomposition. In Section 5 we formulate our compression problem and introduce our approach for the compression of diffusion echoes. Experiments are presented in Section 6, and we conclude our work in Section 7.

2 Nonlinear Diffusion in Image Processing

Given a rectangular, two-dimensional image domain $\Omega \subset \mathbb{R}^2$, we consider a grey value image $f : \Omega \rightarrow \mathbb{R}$. The Perona–Malik filter is a prototype for an isotropic nonlinear diffusion process with a scalar-valued diffusivity function [19]. It regards $f(\mathbf{x})$ as initial state of a diffusion process that yields a filtered image

$u(\mathbf{x}, t)$ as solution at some time $t > 0$. It solves the initial boundary value problem

$$\partial_t u(\mathbf{x}, t) = \operatorname{div}(g \nabla u(\mathbf{x}, t)) \quad \text{for } \mathbf{x} \in \Omega, t \in [0, \infty), \quad (1)$$

$$u(\mathbf{x}, 0) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega, \quad (2)$$

$$\partial_{\mathbf{n}} u(\mathbf{x}, t) = 0 \quad \text{for } \mathbf{x} \in \partial\Omega, t \in (0, \infty), \quad (3)$$

where $\mathbf{n} \in \mathbb{R}^2$ is the outer normal at the boundary $\partial\Omega$, $\nabla = (\partial_x, \partial_y)^\top$ is the spatial gradient, and $\operatorname{div}(\mathbf{v}) = \partial_x v_x + \partial_y v_y$ is the spatial divergence. The function $g = g(|\nabla u(\mathbf{x}, t)|^2)$ is the *diffusivity*, a decreasing function, which takes values in $(0, 1]$ and reduces the smoothing at locations of large gradient magnitude [19]. An example is the rapidly decreasing Weickert diffusivity [22]

$$g(|\nabla u(\mathbf{x}, t)|^2) = 1.0 - \exp\left(\frac{-3.3148}{|\nabla u(\mathbf{x}, t)|^8 / \lambda^8}\right), \quad (4)$$

with a contrast parameter $\lambda > 0$. The scalar diffusivity can be extended to a diffusion tensor, which allows for direction-dependent (*anisotropic*) smoothing [22], but we do not consider such models in this work, as they do not add further aspects in our context. Furthermore, it is standard practice to use a Gaussian-smoothed version u_σ of the image in the calculation of the diffusivity [4]. Here, σ is the standard deviation of the Gaussian convolution kernel.

2.1 Discretisation

As the images we work with are digital, we discretise our diffusion models on a pixel grid of size $n_x \times n_y$, leading to a total number of $N := n_x n_y$ pixels. We stack our images into vectors, e.g. $\mathbf{u}, \mathbf{f} \in \mathbb{R}^N$. The space-discrete version of (1) to (3) is then written as

$$\frac{d\mathbf{u}(t)}{dt} = \mathbf{A}(\mathbf{u}(t))\mathbf{u}(t), \quad (5)$$

$$\mathbf{u}(0) = \mathbf{f}. \quad (6)$$

The matrix $\mathbf{A}(\mathbf{u}(t))$ arises from appropriate discretisations of the differential operators and incorporates the boundary condition (3). We only consider diffusion processes that satisfy the conditions (S1)–(S5) from [22]. These five conditions on $\mathbf{A}(\mathbf{u}(t))$ are Lipschitz-continuity in \mathbf{u} , symmetry, vanishing row sums, nonnegative off-diagonals and irreducibility. Their fulfilment ensures important theoretical properties of the diffusion process. For details we refer to [22].

For practical implementations, we further discretise the time variable t . We use a semi-implicit scheme with time step size τ . This leads to the fully discrete evolution

$$\frac{\mathbf{u}^{k+1} - \mathbf{u}^k}{\tau} = \mathbf{A}(\mathbf{u}^k) \mathbf{u}^{k+1}, \quad (7)$$

$$\mathbf{u}^0 = \mathbf{f}, \quad (8)$$

where \mathbf{u}^k denotes the resulting image at time step k . Obtaining the new iterate \mathbf{u}^{k+1} requires to solve a linear system of equations. Its solution is given by

$$\mathbf{u}^{k+1} = \underbrace{(\mathbf{I} - \tau \mathbf{A}(\mathbf{u}^k))^{-1}}_{\mathbf{P}(\mathbf{u}^k, \tau)} \mathbf{u}^k. \quad (9)$$

Analogously to the semi-discrete case, there are requirements (D1)–(D6) on the matrix $(\mathbf{P}(\mathbf{u}^k, \tau))$: continuity, symmetry, unit row sums, nonnegativity, irreducibility, and positive diagonal entries. One can show that the semi-implicit scheme (7) satisfies these for all τ , given that $\mathbf{A}(\mathbf{u}^k)$ fulfils (S1)–(S5) [22].

We can write the result after n steps as a function of the initial image $\mathbf{u}^0 = \mathbf{f}$:

$$\mathbf{u}^n = \underbrace{\mathbf{P}(\mathbf{u}^{n-1}, \tau) \cdots \mathbf{P}(\mathbf{u}^1, \tau) \mathbf{P}(\mathbf{f}, \tau)}_{\mathbf{S}(\mathbf{f}, \tau, n)} \mathbf{f}. \quad (10)$$

The matrix $\mathbf{S}(\mathbf{f}, \tau, n)$ is the *state transition matrix* of the diffusion process. It maps the initial image \mathbf{f} to the filtered result \mathbf{u}^n at time $t = n\tau$.

3 The Diffusion Echo

3.1 Space-Discrete Diffusion Echo

We start by introducing the diffusion echo in the time-continuous, space-discrete setting. Dam and Nielsen [6] define the space-discrete *source diffusion echo* $\mathbf{s}_i(t)$ at pixel i and time $t > 0$ via the diffusion evolution of a discrete impulse function \mathbf{e}_i located at i (i.e. \mathbf{e}_i is 1 at location i , and 0 elsewhere). The evolution of the echo is governed by the matrix $\mathbf{A}(\mathbf{u}(t))$ from the image evolution (5). That is, the space-discrete evolution of the diffusion echo is given by

$$\frac{d\mathbf{s}_i(t)}{dt} = \mathbf{A}(\mathbf{u}(t)) \mathbf{s}_i(t), \quad (11)$$

$$\mathbf{s}_i(0) = \mathbf{e}_i. \quad (12)$$

The diffusion echo is the nonlinear diffusion analogue of the impulse response of an LSI filter. However, in contrast to the impulse response it is a *space-variant* function of the initial image \mathbf{f} .

3.2 Fully Discrete Diffusion Echo

Following the definition of the diffusion echo (11) as well as the notation from (10), we obtain the discrete echo at pixel i after n steps of size τ as

$$\mathbf{s}_i^n = \mathbf{S}(\mathbf{f}, \tau, n) \mathbf{e}_i. \quad (13)$$

Note that this implies that the fully discrete source diffusion echo is the i -th column of the state transition matrix of the underlying process.

3.3 Properties of the State Transition Matrix

The state transition matrix $\mathbf{S}(\mathbf{f}, \tau, n)$ is the product of matrices of type $\mathbf{P}(\mathbf{u}^k, \tau)$; see (10). (D1)–(D6) from Section 2.1 imply that they are doubly stochastic, i.e. they have nonnegative entries and unit column and row sums. From this it follows that $\mathbf{S}(\mathbf{f}, \tau, n)$ is doubly stochastic as well.

The source echo \mathbf{s}_i^n at pixel i describes how the the pixel’s original grey value intensity is distributed by the diffusion process. In addition, Dam and Nielsen [6] introduced the concept of the *drain echo* \mathbf{d}_j^n at pixel j . It describes from where the intensities at pixel j in the filtered image originated. It is given by the j -th row of the state transition matrix:

$$\mathbf{d}_j^n = \mathbf{e}_j^\top \mathbf{S}(\mathbf{f}, \tau, n) = (\mathbf{S}(\mathbf{f}, \tau, n)^\top \mathbf{e}_j)^\top. \quad (14)$$

In the diffusion case, where each $\mathbf{P}(\mathbf{u}^k, \tau)$ is symmetric, the transposed state transition matrix can be easily applied by performing the diffusion steps in reverse order. This does not require us to compute the $\mathbf{P}(\mathbf{u}^k, \tau)$ explicitly. It suffices to store the intermediate results \mathbf{u}^k or the corresponding diffusivities.

A symmetric state transition matrix $\mathbf{S}(\mathbf{f}, \tau, n)$ would imply equivalence of the source and the drain echoes. While symmetry is guaranteed for homogeneous diffusion, experiments show that it is usually violated for nonlinear diffusion, even though the individual matrices $\mathbf{P}(\mathbf{u}^k, \tau)$ are symmetric.

4 Singular Value Decomposition

For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $p = \min\{m, n\}$, we can find orthogonal matrices $\mathbf{U} \in \mathbb{R}^{m \times m}$, $\mathbf{V} \in \mathbb{R}^{n \times n}$, and a rectangular diagonal matrix $\mathbf{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_p) \in \mathbb{R}^{m \times n}$ with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$, s.t.

$$\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top. \quad (15)$$

This is called the *full singular value decomposition* (SVD) [10]. The columns of \mathbf{U} and \mathbf{V} are the left and right singular vectors, respectively, and the diagonal entries of $\mathbf{\Sigma}$ are the singular values.

If the full matrices \mathbf{U} and \mathbf{V} are not required, one can equivalently represent \mathbf{A} with the *thin SVD*, which only uses the first p columns of \mathbf{U} and \mathbf{V} . This yields the decomposition

$$\mathbf{A} = \mathbf{U}_p \mathbf{\Sigma}_p \mathbf{V}_p^\top = \sum_{i=1}^p \sigma_i \mathbf{u}_i \mathbf{v}_i^\top, \quad (16)$$

with $\mathbf{U}_p \in \mathbb{R}^{m \times p}$, $\mathbf{V}_p \in \mathbb{R}^{n \times p}$ and $\mathbf{\Sigma}_p \in \mathbb{R}^{p \times p}$. If $r = \text{rank}(\mathbf{A}) < p$, we only have r nonzero singular values, and can analogously use the *compact SVD* with $\mathbf{U}_r \in \mathbb{R}^{m \times r}$, $\mathbf{V}_r \in \mathbb{R}^{n \times r}$ and $\mathbf{\Sigma}_r \in \mathbb{R}^{r \times r}$.

4.1 Low-Rank Matrix Approximation

The Eckart–Young theorem [8] states that the matrix \mathbf{B} of rank k , $k < r$, which minimises the error to an original matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ in the Frobenius norm $\|\cdot\|_F$, is given in terms of the *truncated SVD*, i.e.

$$\operatorname{argmin}_{\mathbf{B} \in \mathbb{R}^{m \times n}, \operatorname{rank}(\mathbf{B})=k} \|\mathbf{A} - \mathbf{B}\|_F = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^\top = \mathbf{U}_k \boldsymbol{\Sigma}_k \mathbf{V}_k^\top =: \mathbf{B}_k^*, \quad (17)$$

where $\mathbf{A} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top$ is the SVD of \mathbf{A} . Furthermore, we have for \mathbf{B}_k^* :

$$\|\mathbf{A} - \mathbf{B}_k^*\|_F = \left\| \sum_{i=k+1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^\top \right\|_F = \sqrt{\sum_{i=k+1}^r \sigma_i^2}. \quad (18)$$

4.2 Randomised Approaches for the SVD

Calculating an SVD is expensive, especially for large matrices. This is why, if not all singular values are required, one often resorts to probabilistic approximations. *Randomised Singular Value Decomposition (RSVD)* [17] and *Randomised Subspace Iteration (RSI)* [20] are the best-known algorithms from this field. We sketch these methods in the following.

Consider a given matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, with $m \geq n$ and rank r . The first step is to project \mathbf{A} onto the range of a semi-orthogonal matrix $\mathbf{Q} \in \mathbb{R}^{m \times k}$, with $k < r$:

$$\hat{\mathbf{A}} = \mathbf{Q} \mathbf{Q}^\top \mathbf{A}. \quad (19)$$

This is a product of two low-rank matrices $\mathbf{Q} \in \mathbb{R}^{m \times k}$ and $\mathbf{Q}^\top \mathbf{A} \in \mathbb{R}^{k \times n}$.

The main idea is to construct the compact SVD of $\hat{\mathbf{A}}$ as an approximation to the truncated SVD of \mathbf{A} . To this end, one first calculates the SVD

$$\mathbf{Q}^\top \mathbf{A} = \tilde{\mathbf{U}}_k \tilde{\boldsymbol{\Sigma}}_k \mathbf{V}_k^\top, \quad (20)$$

which is much cheaper than the SVD of \mathbf{A} . Computing $\mathbf{U}_k = \mathbf{Q} \tilde{\mathbf{U}}_k$ yields the desired SVD $\hat{\mathbf{A}} = \mathbf{U}_k \boldsymbol{\Sigma}_k \mathbf{V}_k^\top$, with $\mathbf{U}_k \in \mathbb{R}^{m \times k}$, $\boldsymbol{\Sigma}_k \in \mathbb{R}^{k \times k}$ and $\mathbf{V}_k \in \mathbb{R}^{n \times k}$.

To ensure an accurate approximation, one must construct \mathbf{Q} s.t. $\hat{\mathbf{A}} \approx \mathbf{A}$. RSVD and RSI aim to achieve this based on a matrix $\mathbf{G} \in \mathbb{R}^{n \times k}$, with k normally distributed random vectors as columns, i.e. $\mathbf{g}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. They compute \mathbf{Q} as

$$\mathbf{Q} = \operatorname{orth}((\mathbf{A} \mathbf{A}^\top)^{q-1} \mathbf{A} \mathbf{G}), \quad (21)$$

with a positive integer q . The *orth* operator creates a matrix with columns that form an orthonormal basis for the range of its argument. This can be achieved with a pivoted QR decomposition [10]. The resulting method is called RSVD for $q = 1$ and RSI for $q > 1$.

In practice, one can introduce an oversampling parameter ℓ and construct the random matrix $\mathbf{G} \in \mathbb{R}^{n \times (k+\ell)}$. This makes the methods more robust, which is

important if the singular values decay slowly. Eventually, one can simply discard the ℓ smallest singular values and the corresponding singular vectors. For details on the approximation accuracy of the methods (given e.g. the choice of the parameters k , q and ℓ) we refer to [11]. In this work, we assume that k is fixed in advance. Alternatively, one can also select a target precision and iteratively increase the size of the matrix \mathbf{Q} until the desired error tolerance is met [11].

5 Our Compression Approach

Our compression approach combines the discussed concepts and adapts them to the diffusion echo application. Let $\mathbf{S} \in \mathbb{R}^{N \times N}$ be the state transition matrix of a nonlinear diffusion process. Our goal is to compress the diffusion echoes as efficiently and as accurately as possible. As is often done for image data, we aim for a reconstruction with small mean squared error (MSE). This results in the goal of finding an approximation $\hat{\mathbf{S}}$ that minimises $\|\mathbf{S} - \hat{\mathbf{S}}\|_F$.

To this end, we apply RSI to calculate an approximation to the truncated SVD of the state transition matrix \mathbf{S} . As the iterative diffusion evolution is equivalent to the state transition matrix, we can perform every matrix-vector multiplication by evolving the corresponding vector. Multiplication with the transposed state transition matrix is achieved by applying the symmetric matrices $\mathbf{P}(\mathbf{u}^k, \tau)$ in reverse order. The desired compression ratio (and the oversampling parameter) dictates the size of the matrices in the RSI method and of the final truncated SVD. The computational cost of the RSI method is dominated by the multiplications with \mathbf{A} and \mathbf{A}^\top [15]. In our setting, $2q$ multiplications with matrices of size $N \times (k + \ell)$ are required. This translates to a total of $2q(k + \ell)$ (transposed) diffusion evolutions, while a naïve computation of all echoes would require N evolutions. Although the focus of this work lies on the storage requirements, our approach might therefore also reduce the computational burden.

We write the matrices $\mathbf{U}_k \in \mathbb{R}^{N \times k}$ and $\mathbf{V}_k \boldsymbol{\Sigma}_k \in \mathbb{R}^{N \times k}$ to disk. This representation reduces the storage cost from N^2 to $2Nk$ floats and also ensures optimality guarantees under all rank- k approximations. We can reconstruct each echo by a single matrix-vector multiplication: For source echoes we have $\mathbf{s}_i^n = \mathbf{U}_k \boldsymbol{\Sigma}_k \mathbf{V}_k^\top \mathbf{e}_i = \mathbf{U}_k [(\mathbf{V}_k \boldsymbol{\Sigma}_k)^\top]_i$, and for drain echoes $\mathbf{d}_i^n = \mathbf{e}_i^\top \mathbf{U}_k \boldsymbol{\Sigma}_k \mathbf{V}_k^\top = (\mathbf{V}_k \boldsymbol{\Sigma}_k \mathbf{U}_k^\top \mathbf{e}_i)^\top = (\mathbf{V}_k \boldsymbol{\Sigma}_k [\mathbf{U}_k^\top]_i)^\top$.

Our approach does not work on the echoes individually, but instead processes all N echoes jointly. This allows it to capture the redundancy between the echoes.

6 Results

To evaluate our approach, we compress the echoes of isotropic nonlinear diffusion with the presented Weickert diffusivity [22] from Section 2. We use the 256×256 test image *head* with a grey value range of $[0, 255]$. We set the diffusion time to $t = 25000$, the contrast parameter to $\lambda = 5$ and the pre-smoothing parameter to $\sigma = 0.5$. We use a semi-implicit scheme and solve the occurring linear systems




Fig. 1. The original test image *head*, the filtered result ($t = 25000$, $\lambda = 5$, $\sigma = 0.5$), and two example source echoes. The red dot marks the echo location.

Table 1. Errors for $t = 25000$, $\lambda = 5$ and $\sigma = 0.5$. Estimated Frobenius norm, and MSE of the reconstructed filtered image.

percentage of singular values	1.25 %	2.5 %	5.0 %
Frobenius norm	26.298	9.948	0.014
MSE	2507.986	95.264	0.002

with the conjugate gradient (CG) method. The resulting image and two source echoes are displayed in Figure 1. The rapidly decreasing diffusivity leads to sharp edges and segmentation-like results, implying inherently redundant echoes.

To compute the truncated SVD, we apply RSI with two iterations ($q = 3$) and ten oversampling vectors ($\ell = 10$). We use LAPACK [1] for the necessary matrix decompositions in the RSI algorithm. For the evaluation, we approximate the Frobenius norm of the error matrix $\|\mathbf{S} - \hat{\mathbf{S}}\|_F$ with a stochastic trace estimator [12], using that $\|\mathbf{A}\|_F^2 = \text{trace}(\mathbf{A}^T \mathbf{A})$ for a real matrix \mathbf{A} . Additionally, we perform visual comparisons of the source echoes to intuitively judge the accuracy of the reconstructed echoes. Lastly, we use the decompressed source echoes to reconstruct an approximation to the true filtered image, and evaluate its quality using the MSE. All the displayed echoes are rescaled such that their maximal value is 255.

The results in Table 1 and Figure 2 show that neither 1.25 % nor 2.5 % of the singular values suffice for adequate reconstructions, while 5.0 % yield almost perfect results. When inspecting the original echoes in closer detail, one can notice that there are a number of echoes that are concentrated almost entirely in a single pixel – at the echo location itself. Such echoes lie within edges, where the local contrast is so large that barely any smoothing is performed. These impulse-like echoes lead to large singular values and the corresponding singular vectors induce pointwise errors, if they are discarded in the reconstruction. This effect can be observed in Figure 2 (b) and (c).

As a remedy, we propose to localise impulse-like echoes prior to the computation of the SVD, and to remove the corresponding columns and rows from the state transition matrix. We consider all echoes where the value in the echo location (i.e. the diagonal elements of the state transition matrix) is larger than some threshold $1 - \epsilon$. We then perform the truncated SVD as before. For each of




Fig. 2. The source echoes from Figure 1 (c) and (d) and their reconstructions using 1.25 %, 2.5 % and 5.0 % of the singular values ($t = 25000$, $\lambda = 5$, $\sigma = 0.5$). Note that the brightness differences are due to the rescaling.

Table 2. Errors for $t = 25000$, $\lambda = 5$ and $\sigma = 0.5$ using the proposed exclusion mechanism with $\epsilon \in \{0.05, 0.1\}$. Estimated Frobenius norm, and MSE of the reconstructed filtered image.

	percentage of singular values	1.25 %	2.5 %	5.0 %
$\epsilon = 0.05$	Frobenius norm	11.741	0.673	0.667
	MSE	184.138	0.013	0.010
$\epsilon = 0.1$	Frobenius norm	6.997	1.401	1.401
	MSE	33.457	0.042	0.042

these impulse-like echoes, we only need to store the location of the echo, making this approach very efficient. Note that the size of the singular vectors is also reduced by the number of such echoes. While this remedy induces a small error at those locations, it leads to much better results and does not increase the storage cost. For reconstruction, we simply compute the echoes from the SVD and add back the pixels that were previously excluded.

Table 2 shows the results when we remove impulse-like echoes from the SVD computation. In the experiments, we test $\epsilon = 0.05$ and $\epsilon = 0.1$. We see that our remedy greatly improves the results for high compression ratios. On the other hand, we introduce a lower bound on the reconstruction error, due to the errors we induce at the excluded locations. Compared to Table 1, this leads to worse results at 5.0 % of the singular values. A smaller ϵ improves this bound, at the cost of increasing the error at high compression ratios. Therefore, ϵ should be tuned depending on a desired compression ratio or error.

The error of the reconstruction depends on the singular values that have been truncated; see (18). Therefore, we now investigate the singular value spectra




Fig. 3. The singular value spectra from the experiments. The vertical lines depict 1.25%, 2.5% and 5.0% of the singular values.

from our experiments. Of course, we only have access to the first 5.0% of the singular values, and those are approximations by the RSI method. Nonetheless, the shapes and ranges of the spectra help to understand the differences in the results, so we plot them in Figure 3. We see that our test case, nonlinear diffusion with the Weickert diffusivity, leads to many large singular values, before a fast decay occurs. This explains why the results on the full matrix are rather poor for the two highest compression ratios, before they become almost perfect for 5.0%. Our exclusion approach shifts the spectrum, such that the singular value decay happens earlier, which leads to better results at high compression ratios. Once the largest singular values are included in the compact representation, the error becomes very small, and can barely be improved by adding more singular values.

It should be noted that the Weickert diffusivity is a particularly difficult test case. Less rapidly decreasing diffusivities, such as the Perona–Malik ones [19], suffer less from impulse-like echoes and have fewer large singular values (see Figure 4). This enables our method to produce even better results at high compression ratios.

7 Conclusions

Tailoring an SVD-based approximation and reconstruction method to our requirements, we have presented the first compression approach that makes the concept of diffusion echoes practical. We have shown that the redundancy between echoes for nonlinear diffusion filters can be perfectly exploited by the SVD to obtain inexpensive and accurate representations. The fast reconstruction renders the approach viable for visual inspections, making it possible to fully benefit from the fundamental idea of diffusion echoes. The substantially reduced stor-




Fig. 4. The singular value spectrum of the Weickert diffusivity with the same parameters as before and a prototypical spectrum of the Perona–Malik diffusivity ($t = 200$, $\lambda = 3$, $\sigma = 0.5$).

age requirements and the rapid access to all echoes may pave the way for future applications that have been prohibited so far by hardware constraints.

While our approach may be challenged by the huge size of modern images, we have seen that the echoes are typically fairly localised. Thus, it will be possible to apply it to smaller patches that still cover the full filter action. This is part of our ongoing work. Last but not least, our encouraging results motivate us to extend the idea of a nonlinear filter echo to processes beyond diffusion filtering.

Acknowledgments. We thank Vassilen Chizhov for fruitful discussions. We also thank the anonymous reviewers for their valuable comments.

This version of the contribution has been accepted for publication in T. A. Bubba, R. Gaburro, S. Gazzola, K. Papafitsoros, M. Pereyra, C.-B. Schönlieb (Eds.): *Scale Space and Variational Methods in Computer Vision. Part II. Lecture Notes in Computer Science*, Vol. 15668, Springer, Cham, 324-336, 2025, after peer review but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: https://doi.org/10.1007/978-3-031-92369-2_25

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D.: *LAPACK Users' Guide*. SIAM, Philadelphia, PA, third edn. (1999)
2. Baykova, I.: *PCA-based Representation of Diffusion Echoes*. Bachelor thesis, Department of Computer Science, Saarland University, Saarbrücken, Germany (2016)

3. Cárdenas, G., Weickert, J., Schäffer, S.: A linear scale-space theory for continuous nonlocal evolutions. In: Aujol, J., Nikolova, M., Papadakis, N. (eds.) *Scale Space and Variational Methods in Computer Vision*, Lecture Notes in Computer Science, vol. 9087, pp. 103–114. Springer, Berlin (2015)
4. Catté, F., Lions, P.L., Morel, J.M., Coll, T.: Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis* **29**(1), 182–193 (1992)
5. Çiçek, Ö.: Efficient Computation and Representation of the Diffusion Echo. Master thesis, Department of Computer Science, Saarland University, Saarbrücken, Germany (2014)
6. Dam, E., Nielsen, M.: Exploring non-linear diffusion: The diffusion echo. In: Kerckhove, M. (ed.) *Scale-Space and Morphology in Computer Vision*, Lecture Notes in Computer Science, vol. 2106, pp. 264–272. Springer, Berlin (2001)
7. Dam, E., Olsen, O.F., Nielsen, M.: Approximating non-linear diffusion. In: Griffin, L., Lillholm, M. (eds.) *Scale Space Methods in Computer Vision*, Lecture Notes in Computer Science, vol. 2695, pp. 117–131. Springer, Berlin (2003)
8. Eckart, C., Young, G.: The approximation of one matrix by another of lower rank. *Psychometrika* **1**(3), 211–218 (1936)
9. Fischl, B., Schwartz, E.: Learning an integral equation approximation to nonlinear anisotropic diffusion in image processing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(4), 342–352 (1997)
10. Golub, G.H., Van Loan, C.F.: *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, third edn. (1996)
11. Halko, N., Martinsson, P.G., Tropp, J.A.: Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* **53**(2), 217–288 (2011)
12. Hutchinson, M.: A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics - Simulation and Computation* **19**(2), 433–450 (1990)
13. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bulletin of the Electrotechnical Laboratory* **26**, 368–388 (1962), in Japanese
14. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. In: Bruckstein, A., ter Haar Romeny, B., Bronstein, A., Bronstein, M. (eds.) *Scale Space and Variational Methods in Computer Vision*, Lecture Notes in Computer Science, vol. 6667, pp. 26–37. Springer, Berlin (2011)
15. Martinsson, P.G., Tropp, J.A.: Randomized numerical linear algebra: Foundations and algorithms. *Acta Numerica* **29**, 403–572 (2020)
16. Nitzberg, M., Shiota, T.: Nonlinear image filtering with edge and corner enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**, 826–833 (1992)
17. Papadimitriou, C.H., Raghavan, P., Tamaki, H., Vempala, S.: Latent semantic indexing: A probabilistic analysis. *Journal of Computer and Systems Sciences* **61**(2), 217–235 (2000)
18. Pearson, K.: LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2**(11), 559–572 (1901)
19. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**, 629–639 (1990)

20. Rokhlin, V., Szlam, A., Tygert, M.: A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications* **31**(3), 1100–1124 (2010)
21. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: *Proc. Sixth International Conference on Computer Vision*. pp. 839–846. Narosa Publishing House, Bombay, India (1998)
22. Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)