



Connections Between Numerical Algorithms for PDEs and Neural Networks

Tobias Alt¹ · Karl Schrader¹ · Matthias Augustin¹ · Pascal Peter¹ · Joachim Weickert¹

Received: 30 July 2021 / Accepted: 17 May 2022 / Published online: 24 June 2022
© The Author(s) 2022

Abstract

We investigate numerous structural connections between numerical algorithms for partial differential equations (PDEs) and neural architectures. Our goal is to transfer the rich set of mathematical foundations from the world of PDEs to neural networks. Besides structural insights, we provide concrete examples and experimental evaluations of the resulting architectures. Using the example of generalised nonlinear diffusion in 1D, we consider explicit schemes, acceleration strategies thereof, implicit schemes, and multigrid approaches. We connect these concepts to residual networks, recurrent neural networks, and U-net architectures. Our findings inspire a symmetric residual network design with provable stability guarantees and justify the effectiveness of skip connections in neural networks from a numerical perspective. Moreover, we present U-net architectures that implement multigrid techniques for learning efficient solutions of partial differential equation models, and motivate uncommon design choices such as trainable nonmonotone activation functions. Experimental evaluations show that the proposed architectures save half of the trainable parameters and can thus outperform standard ones with the same model complexity. Our considerations serve as a basis for explaining the success of popular neural architectures and provide a blueprint for developing new mathematically well-founded neural building blocks.

Keywords Numerical algorithms · Partial differential equations · Neural networks · Nonlinear diffusion · Stability

1 Introduction

Partial differential equations (PDEs) have been a central component of many successful models for signal and image processing in the last three decades; for instance, the monographs [6,16,105] provide a good overview.

PDE-based models are compact, transparent, and benefit from a rich set of mathematical foundations. As a consequence, they offer valuable theoretical guarantees such as stability and well-posedness. This makes PDE models and their numerical solution strategies easy to control, implement, and apply to a plethora of tasks.

Convolutional neural networks (CNNs) and deep learning [39,64,65,99], on the other hand, have revolutionised the field of image processing in recent years. Still, this success was mostly of empirical nature. Many modern CNN archi-

tectures do not provide solid mathematical foundations. They often suffer from undesirable side effects such as a sensitivity against adversarial attacks [40].

More recently, researchers have started to analyse the behaviour and mathematical foundations of CNNs. One strategy is to interpret trained CNNs as approximations of continuous partial or ordinary differential equations (ODEs) [19,20,29]. In this interpretation, the trainable parameters specify the nonlinear dynamics of an evolution equation.

This strategy can be challenging for several reasons. Firstly, finding a compact differential equation is hard, given that typical CNNs learn millions of parameters. Secondly, the reduction of a discrete CNN to a continuous differential equation involves ambiguous limit assumptions, as the same discrete model can approximate multiple evolution equations with different orders of consistency. Last but not least, this strategy is only analytic: It analyses existing networks rather than inspiring novel, well-founded building blocks.

We address these problems by pursuing the opposite direction: We translate successful concepts from the world of PDEs into neural components. This translation justifies neural architectures from a mathematical perspective and

✉ Tobias Alt
alt@mia.uni-saarland.de

¹ Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science, Saarland University, Campus E1.7, 66041 Saarbrücken, Germany

provides novel design criteria for well-founded networks. In addition, these networks are naturally more compact with less trainable parameters.

Our concepts of choice are numerical algorithms instead of continuous differential equations. Similar to untrained neural networks, numerical algorithms can be applied to a multitude of problems in a general purpose fashion. The model at hand is then specified by the differential equation, which is approximated by training the neural network. Thus, we believe that the design principles of modern neural networks realise a small but powerful set of numerical strategies at their core.

1.1 Our Contributions

We investigate what can be learned from translating numerical methods for PDEs into their neural counterparts. This inspires novel building blocks for designing mathematically well-founded neural networks.

The present paper advances our conference publication [1], in which we translated explicit schemes, acceleration strategies [51], implicit schemes, and linear multigrid approaches [11,12] into their neural counterparts. In this extended version, we additionally investigate networks that realise Du Fort–Frankel schemes [28] as a representative for absolutely stable schemes which are still explicit. Moreover, we extend the translation of multigrid approaches into U-nets to the nonlinear setting by considering full approximation schemes (FAS). Last but not least, we supplement our findings with an experimental evaluation of the proposed network architectures for denoising and inpainting tasks. Therein, we demonstrate the effectiveness of the architectures together with nonmonotone activation functions in practice.

As a starting point of our considerations, we consider a generalised nonlinear diffusion equation. We restrict ourselves to the 1D setting for didactic purposes only, since all important concepts can already be translated in this simple setting. We show that an explicit discretisation of this diffusion model can be connected to residual networks (ResNets) [54], which are among the most popular network architectures to date. Their skip connection can be interpreted as the result of a temporal discretisation, which allows to connect them to explicit diffusion schemes.

This connection inspires a novel ResNet architecture. It follows a symmetric structure which saves half the amount of network parameters and additionally allows to derive a theory for guaranteeing stability in the Euclidean norm. Moreover, by identifying the diffusion flux with the activation function our translation motivates the use of nonmonotone activation functions, which are atypical in the CNN world. In a series of denoising experiments, we validate the effectiveness of such activation functions. We choose the denoising problem since

it is the prototypical representative of a well-posed problem, where the result depends continuously on the input data.

By considering acceleration strategies for explicit schemes and solution strategies for implicit schemes, we justify the effectiveness of skip connections in neural networks. We show that Du Fort–Frankel schemes [28], fast semi-iterative (FSI) schemes [51], and fixed point iterations for implicit schemes motivate different architectural designs which all rely on skip connections as a foundation of their efficiency.

Finally, we consider the rich class of multigrid approaches [11,12]. We show that a nonlinear full approximation scheme (FAS) can be cast in the form of the popular U-net [87] architecture. We think that at their core, U-nets realise a multigrid strategy, and we support this claim by proposing a U-net which realises a full multigrid strategy for an inpainting task.

Our findings do not only inspire new design criteria for stable neural architectures and show that uncommon design choices can perform well in practice. They also provide structural insights into the success of popular CNN architectures from the perspective of numerical algorithms.

1.2 Related Work

In recent years, the connection between neural networks and the world of PDEs and variational methods has become an active area of research.

CNNs are used to learn PDEs from data [68,82,90,96], or to solve them efficiently [21,30,83]. Moreover, various model-based approaches have been augmented with trainable parameters to improve their performance [2,5,20,60,61]. Another line of research is concerned with the expressive power of networks [24,44,63,79,86,101] and their robustness properties [23,34,66].

The concept of neural ordinary equations [19] as a continuous time extension of ResNets [54] has gained considerable attention. However, recent works [46,77] suggest that these architectures suffer from a strong dependency between the model and the numerical solver. This supports our motivation to regard the numerical solver as an inherent basis of a neural architecture.

In contrast, our philosophy of translating numerical concepts into neural architectures is shared only by few works [10,67,69,78,111]. They motivate additional or modified skip connections based on numerical schemes for ODEs, such as Runge–Kutta methods or implicit Euler schemes. Our work provides additional motivations for such skip connections based on several numerical strategies for PDEs.

The stability of ResNets has been analysed in several works [17,47,48,88,93,110]. A common result is that ResNets with a symmetric filter structure can be shown to be stable in the Euclidean norm. We motivate this result from a novel viewpoint based on diffusion processes. In contrast to previous results, this unique starting point allows us to

present our stability result independently of the monotonicity of the activation function, inspiring the use of nonmonotone and trainable activation functions.

Nonmonotone activation functions are rarely found in standard CNNs, with some notable exceptions [20,38,76]. Recently, the so-called Swish activation [84] and modifications thereof [72,112] have been found to empirically boost the classification performance of CNNs. While these activations are modifications of the ReLU activation which are nonmonotone around the zero position, the activations that arise from our diffusion interpretation are odd and nonmonotone. Such functions have been analysed before the advent of deep learning [25,71], but have not found their way into current CNN architectures. Our experiments, however, suggest that these activations can be advantageous in practice.

Multigrid ideas have been combined with CNNs already in the early years of neural network research [7,8]. Current works use inspiration from multigrid concepts to learn restriction and prolongation operators of multigrid solvers [43,58], to couple channels for parameter reduction [31] or to boost training performance [45,49].

However, to the best of our knowledge, the only architectures that consequently implement a trainable multigrid approach are presented by He and Xu [53] and Hartmann et al. [52]. However, both works do not draw any connections to the popular U-net architecture [87], whereas we directly link both concepts.

1.3 Organisation of the Paper

In Sect. 2, we review nonlinear diffusion and residual networks. We connect both models in Sect. 3 and analyse the implications in terms of stability and novel activation functions. Afterwards, we motivate skip connections from different numerical algorithms in Sect. 4. We review multigrid approaches and U-nets in Sect. 5 before connecting both worlds in Sect. 6. Finally, we experimentally evaluate the proposed architectures in Sect. 7 and present a discussion and our conclusions in Sect. 8.

2 Review: Diffusion and Residual Networks

In this section, we review generalised diffusion filters in 1D and residual networks as the basic models for our first translation. We restrict ourselves to the 1D setting only for didactic reasons, as already this simple setting allows to translate all necessary concepts.

2.1 Generalised Nonlinear Diffusion

We start by considering a generalised one-dimensional diffusion PDE of arbitrary high order. It produces signals

$u(x, t) : (a, b) \times [0, \infty) \rightarrow \mathbb{R}$ evolving over time from an initial signal $f(x)$ on a domain $(a, b) \subset \mathbb{R}$ according to

$$\partial_t u = -\mathcal{D}^* \left(g(|\mathcal{D}u|^2) \mathcal{D}u \right), \tag{1}$$

with reflecting (homogeneous Neumann) boundary conditions. We use a general differential operator

$$\mathcal{D} = \sum_{m=0}^M \alpha_m \partial_x^m \tag{2}$$

and its adjoint

$$\mathcal{D}^* = \sum_{m=0}^M (-1)^m \alpha_m \partial_x^m. \tag{3}$$

The operators consist of weighted derivatives up to order M with weights α_m of arbitrary sign, yielding a PDE of order $2M$.

Choosing e.g. $M = 1$ yields the second order PDE of Perona and Malik [80], while $M = 2$ leads to a one-dimensional version of the fourth order model of You and Kaveh [109].

The evolution simplifies the input signal f over time. This process is mainly controlled by the scalar *diffusivity* function $g(s^2)$. For example, the Perona–Malik diffusivity [80]

$$g(s^2) = \frac{1}{1 + \frac{s^2}{\lambda^2}} \tag{4}$$

preserves discontinuities which are larger than a contrast parameter λ .

The diffusion PDE (1) is the gradient flow which minimises the energy functional

$$E(u) = \int_a^b \Psi(|\mathcal{D}u|^2) \, dx, \tag{5}$$

where the penaliser Ψ can be linked to the diffusivity with $g = \Psi'$ [97]. The penaliser must be increasing, but not necessarily convex. In Sect. 3.4, we show that this inspires novel activation functions. Their discretisations are stable, despite arising from a nonconvex energy.

2.2 Residual Networks

Residual networks (ResNets) [54] belong to the most popular CNN architectures to date. Their main contribution is the introduction of so-called *skip connections* which facilitate training of very deep networks.

A residual network consists of residual blocks. A single block computes an output signal u from an input f as

$$u = \sigma_2(f + W_2 \sigma_1(W_1 f + b_1) + b_2). \tag{6}$$

One first applies an inner convolution W_1 with a bias vector b_1 to the input signal and passes the result into an inner activation function σ_1 .

Typically, CNNs prescribe simple, monotone activation functions such as the rectified linear unit (ReLU) [73] function

$$\text{ReLU}(s) = \max(0, s) \tag{7}$$

which is a linear function truncated at 0.

Afterwards, an outer convolution W_2 with a bias vector b_2 is applied to the output of the activation.

The result of this convolution is added back to the original signal f . This *skip connection* is the crucial novelty of ResNets over feed-forward networks. It is the key to efficiently train deep networks with large amounts of layers, without suffering from the vanishing gradient problem. This phenomenon appears when backpropagation gradients approach zero for very deep networks, bringing the training process to a halt [9].

Lastly, one applies an outer activation function σ_2 to obtain the output signal u .

3 From Diffusion to Symmetric Residual Networks

We are now in the position to show that explicit diffusion schemes realise a ResNet architecture with a symmetric filter structure. To this end, we rewrite the generalised nonlinear diffusion equation (1) with the help of the flux function

$$\Phi(s) = g(s^2) s \tag{8}$$

as

$$\partial_t u = -\mathcal{D}^* \Phi(\mathcal{D}u). \tag{9}$$

Now we discretise this equation by means of a standard finite difference scheme. To obtain discrete signals u, f , we sample the continuous signals u, f with distance h . We discretise the temporal derivative by a forward difference with time step size τ . The spatial derivative operator \mathcal{D} is implemented by a convolution matrix K . Consequently, the adjoint operator \mathcal{D}^* is realised by a transposed convolution matrix K^\top . The matrix transposition corresponds mirroring the corresponding discrete convolution kernel.

This yields the discrete evolution equation

$$\frac{u^{k+1} - u^k}{\tau} = -K^\top \Phi(Ku^k), \tag{10}$$

where we indicate old and new time levels by superscripts k and $k + 1$, respectively. Solving this expression for the new

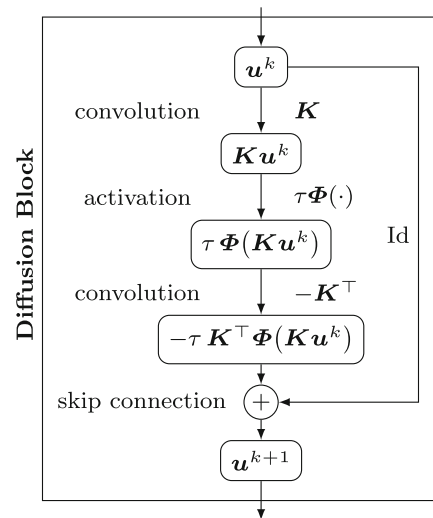


Fig. 1 Diffusion block for an explicit diffusion step (11) with flux function Φ , time step size τ , and a discrete derivative operator K

signal u^{k+1} yields the explicit scheme

$$u^{k+1} = u^k - \tau K^\top \Phi(Ku^k). \tag{11}$$

The explicit diffusion scheme (11) is closely connected to a specific residual block. To this end, we consider a residual block with input u^k , output u^{k+1} , and without bias terms, which reads

$$u^{k+1} = \sigma_2(u^k + W_2 \sigma_1(W_1 u^k)). \tag{12}$$

Now, we can directly identify the explicit scheme with a residual block as follows.

Theorem 1 (Diffusion-inspired ResNets) *An explicit step (11) of the generalised higher order diffusion scheme (1) can be expressed as a residual block (6) by*

$$\sigma_1 = \tau \Phi, \quad \sigma_2 = \text{Id}, \quad W_1 = K, \quad W_2 = -K^\top, \tag{13}$$

with the bias vectors b_1, b_2 set to $\mathbf{0}$.

We call a ResNet block of this form a *diffusion block*. Figure 1 visualises such a block in the form of a graph. Nodes contain the current state of the signal, while edges describe the operations to proceed from one node to the next.

The connection between explicit diffusion schemes and ResNets yields three key structural insights:

1. The rescaled flux function $\tau \Phi$ serves as the sole activation function σ_1 . This motivates us to investigate popular diffusion flux functions in Sect. 3.4. They have not been considered as CNN activations so far.

2. The skip connection naturally arises from the discretisation of the temporal derivative. This is one numerical justification for skip connections in neural networks. We investigate several other motivations of skip connections in Sect. 4.
3. Lastly, the filters exhibit a negated symmetric filter structure $W_2 = -W_1^T$. This is a natural consequence of the gradient flow structure of the diffusion process, and leads to provable stability guarantees for ResNets with such a filter structure, as we show in the following.

3.1 Stability for Symmetric ResNets

The structural connection between explicit schemes and ResNets allows us to transfer classical results for stability [26] and well-posedness [105] of diffusion evolutions to a specific residual network architecture.

To this end, we consider ResNets which chain diffusion blocks. Since we show that a key to stability of these networks is the symmetric filter structure, we refer to these architectures as *symmetric residual networks (SymResNets)*, following [93,110].

For these networks, we prove Euclidean stability and well-posedness. Euclidean stability guarantees that the Euclidean norm of the signal is nonincreasing in each iteration, i.e. $\|u^{k+1}\|_2 \leq \|u^k\|_2$. Well-posedness ensures that the network output is a continuous function of the input data.

Theorem 2 (Euclidean Stability of Symmetric Residual Networks) *Consider a symmetric residual network chaining any number of diffusion blocks (11) with convolutions represented by a convolution matrix K and activation function $\tau\Phi$. Moreover, assume that the activation function can be expressed as a diffusion flux function $\Phi(s) = g(s^2)s$ and has a finite Lipschitz constant L . Then the symmetric residual network is well-posed and stable in the Euclidean norm if*

$$\tau \leq \frac{2}{L\|K\|_2^2}. \tag{14}$$

Here, $\|\cdot\|_2$ denotes the spectral norm which is induced by the Euclidean norm.

Proof The activation function $\sigma(s)$ can be expressed in terms of a diffusivity function by

$$\sigma(s) = \tau\Phi(s) = \tau g(s^2)s. \tag{15}$$

Thus, its application is equivalent to a rescaling with a diagonal matrix $G(u^k)$ with $g((Ku^k)_i^2)$ as i -th diagonal element. Therefore, we can write (11) as

$$u^{k+1} = (I - \tau K^T G(u^k) K) u^k. \tag{16}$$

At this point, well-posedness follows directly from the continuity of the operator $I - \tau K^T G(u^k) K$, as the diffusivity g is assumed to be smooth [105].

We now show that the time step size restriction (14) guarantees that the eigenvalues of the operator always lie in the interval $[-1, 1]$. Then, the explicit step (11) constitutes a contraction mapping which in turn guarantees Euclidean stability.

As the spectral norm is sub-multiplicative, we can estimate the eigenvalues of $K^T G(u^k) K$ for each matrix separately. Since g is nonnegative, the diagonal matrix G is positive semidefinite. The maximal eigenvalue of G is then given by the supremum of g , which can be bounded by the Lipschitz constant L of Φ :

$$L = \sup_s |\Phi'(s)| = \sup_s |g(s^2) + 2s^2 g'(s^2)| \geq \sup_s |g(s^2)|. \tag{17}$$

Consequently, the eigenvalues of $K^T G(u^k) K$ lie in the interval $[0, \tau L \|K\|_2^2]$.

Then, the operator $I - \tau K^T G(u^k) K$ has eigenvalues in $[1 - \tau L \|K\|_2^2, 1]$, and the condition

$$1 - \tau L \|K\|_2^2 \geq -1 \tag{18}$$

leads to the bound (14). □

Similar results have been obtained recently in [88,93,110], albeit with alternative justifications. In Sect. 3.4, we show that our unique diffusion interpretation additionally suggests novel design concepts for CNNs such as nonmonotone activation functions.

3.2 How General is Our Stability Result?

While our focus on explicit diffusion schemes appears restrictive at first glance, our stability result is more general.

The fact that we use discrete differential operators as convolutions is no restriction, since any convolution matrix can be expressed as a weighted combination of discrete differential operators. Moreover, our proof does not even require a convolutional matrix structure.

A key requirement for stability is the symmetric structure $W_2 = -W_1^T$.

The symmetric convolution structure is an important structural difference to the original ResNet formulation [54]. It does not only yield a stable network, but also allows to reduce the amount of trainable parameters by 50%, since inner convolution and outer convolution share their weights.

Moreover, the requirement of using a flux function as an activation function can be relaxed. As we have shown, one only requires the diagonal matrix \mathbf{G} to be positive semidefinite. While this is naturally fulfilled for a diffusion flux function, other activations also adhere to this constraint. For example, the ReLU function multiplies positive arguments with 1 and negative ones with 0, yielding a binary positive semidefinite matrix \mathbf{G} . Thus, using the ReLU instead of a diffusion flux does not affect stability. This shows that diffusion algorithms inspire general, sufficient design criteria for stable networks.

In particular, we do not require any assumptions on the monotonicity of the activation function, in contrast to the results of Ruthotto and Haber [93]. This motivates us to investigate typical diffusivities and their flux functions in Sect. 3.4.

3.3 Enforcing Stability in Practice

While the stability criterion (14) can be computed on the fly already during the training process of the network, evaluating the spectral radius of the operator \mathbf{K} is costly. To this end, we suggest a simple rescaling to turn the stability bound (14) into an a priori criterion.

For a symmetric residual network with a single channel, one can directly use Gershgorin's circle theorem [35] to bound the maximum absolute eigenvalue of \mathbf{K} . More precisely, the eigenvalues of \mathbf{K} lie in the union of circles around the diagonal entries k_{ii} with radii $r_i = \sum_{j \neq i} |k_{ij}|$ corresponding to the absolute sums of the off-diagonal values. Thus, the maximal absolute eigenvalue of \mathbf{K} is bounded by the largest absolute row sum of \mathbf{K} . If we simply rescale both inner and outer convolutions by this sum, we can guarantee $\|\mathbf{K}\|_2^2 \leq 1$. Then the stability condition (14) transforms into $\tau \leq \frac{2}{L}$. Since the Lipschitz constant L of the activation is known a priori, this simple rescaling allows to constrain the time step size to a fixed value, while not affecting the expressive power of the network.

However, most networks in practice are not concerned with only a single channel. To this end, we extend our stability result to symmetric ResNets with multiple channels.

For a diffusion block operating on a signal with C channels, the matrix \mathbf{K} is a $C \times C$ block convolution matrix. As long as the transposed structure is realised, this is not problematic for the stability proof.

An extension of Gershgorin's circle theorem to block matrices [102] states that the eigenvalues of \mathbf{K} lie in the union of circles which are centred around the eigenvalues of the diagonal blocks. The radii of the circles are given by the sum of the spectral norms of the off-diagonal blocks. If we rescale each block matrix as in the single channel case, we simply need to additionally divide the operator \mathbf{K} by \sqrt{C} to ensure that $\|\mathbf{K}\|_2^2 \leq 1$. With this, we obtain the same a priori criterion as in the single channel case.

This strategy constitutes an instance of the popular weight normalisation technique [95], and related spectral normalisations have shown to be successful for improving the performance and convergence speed of the training process [15,22,42,113].

3.4 Nonmonotone Activation Functions

Our connection between diffusivity $g(s^2)$ and activation function $\sigma(s) = \tau \Phi(s)$ with the diffusion flux $\Phi(s)$ refreshes an old idea of neural network design [25,71].

In Fig. 2, we present three activation functions: The ReLU activation [73], along with two activations resulting from popular diffusivities. These activations are *odd* functions, which is natural in the diffusion case, where the argument of the flux function consists of signal derivatives. It reflects the invariance axiom that signal negation and filtering are commutative.

The Charbonnier diffusivity [18], which stems from a convex energy and can be seen as a rescaled regularised total variation diffusivity [4,89], yields a monotone activation function. Similarly shaped activation functions such as the hyperbolic tangent have been used in early neural networks, before being superseded by ReLU activations.

The rational Perona–Malik diffusivity (4) [80], however, results in a nonmonotone activation function. The associated energy functional is nonconvex. Nevertheless, discretisations of diffusion processes using nonmonotone flux functions can be shown to be well-posed, despite acting contrast-enhancing [106]. For more activation functions inspired by diffusivities, we refer to [3].

The concept of nonmonotone activation functions is unusual in the CNN world. Although there have been a few early proposals in the neural network literature arguing in favour of nonmonotone activations [25,71], they are rarely used in modern CNNs. In practice, they often fix the activation to simple, monotone functions such as the rectified linear unit (ReLU). From a PDE perspective, this appears restrictive. The diffusion interpretation suggests that activation functions should be learned in the same manner as convolution weights and biases.

In practice, this hardly happens apart from a few notable exceptions such as [20,38,76]. Recently, activation functions which are slightly nonmonotone variants of the ReLU proved successful for image classification tasks [72,84,112]. As nonmonotone flux functions outperform monotone ones, e.g. for diffusion-based denoising [80], it appears promising to incorporate them into CNNs.

Our translation of explicit schemes is an example of a simple, direct correspondence which in turn allows for multiple novel insights. In the following, we explore variants of explicit schemes as well as implicit schemes which inspire

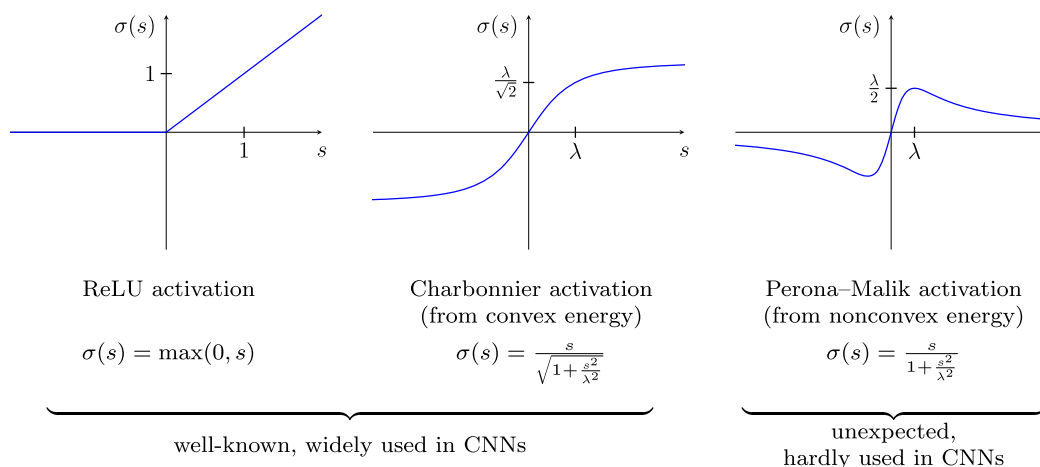


Fig. 2 Visualisation of three activation functions. The ReLU (left) is a standard activation function. The Charbonnier and Perona–Malik activations represent two representatives of diffusion-inspired activation functions. Both are dependent on a contrast parameter λ . The Charbon-

nier activation is monotone, as it arises from a convex energy. On the contrary, the Perona–Malik activation is nonmonotone, and its associated energy is nonconvex. In this visualisation, we set $\tau = 1$ for the sake of simplicity

changes to the skip connections of the symmetric ResNets, leading to more efficient architectures.

4 The Value of Skip Connections

So far, we have seen that the temporal discretisation of an explicit scheme naturally leads to skip connections. This, however, is just one of the many justifications for their use. Since their proposal, skip connections [54] have been adapted into numerically inspired networks in many different forms, see e.g. [56,67,69,78,111].

This motivates us to explore several other numerical algorithms which justify several types of skip connections from a numerical perspective. We explore unconditionally stable schemes, acceleration strategies for explicit schemes, and fixed point iterations for implicit schemes.

4.1 Du Fort–Frankel Schemes

While the classical explicit scheme (11) is only conditionally stable, there exist absolutely stable schemes which are still explicit. These schemes are not that popular in practice since they trade unconditional stability for conditional consistency. However, we will see that this is not problematic from the perspective of learning.

Du Fort and Frankel [28] propose to change the temporal discretisation of the explicit scheme (11) to a central difference and introduce a stabilisation term on the right hand side, corresponding to an approximation of $\partial_{tt}u$. A Du Fort–Frankel scheme for the generalised diffusion (1) evolution

can be written as

$$\frac{\mathbf{u}^{k+1} - \mathbf{u}^{k-1}}{2\tau} = -\mathbf{K}^\top \Phi(\mathbf{K}\mathbf{u}^k) - \alpha(\mathbf{u}^{k+1} - 2\mathbf{u}^k + \mathbf{u}^{k-1}), \tag{19}$$

where a positive constant α controls the influence of the stabilisation term.

Solving this scheme for \mathbf{u}^{k+1} yields

$$\mathbf{u}^{k+1} = \frac{4\tau\alpha}{1 + 2\tau\alpha} \left(\mathbf{u}^k - \frac{1}{2\alpha} \mathbf{K}^\top \Phi(\mathbf{K}\mathbf{u}^k) \right) + \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} \mathbf{u}^{k-1}. \tag{20}$$

For $\tau\alpha = \frac{1}{2}$, one obtains the explicit scheme (11).

The scheme involves the signals \mathbf{u}^k and \mathbf{u}^{k-1} at the current and the previous time level. The first term is nothing else than a rescaled diffusion block, where $\frac{1}{2\alpha}$ takes the role of the original time step size. Since the scalar factors $\frac{4\tau\alpha}{1+2\tau\alpha}$ and $\frac{1-2\tau\alpha}{1+2\tau\alpha}$ add up to 1, this is simply an extrapolation of the result of an explicit step based on the signal at time level $k - 1$.

If α is large enough, this scheme is unconditionally stable. Thus, one does not need to obey any stability condition, in contrast to the explicit case. Whereas classical proofs such as [41] consider only the linear case and typically work in the Fourier space, we are not aware of any proofs for the stability of nonlinear Du Fort–Frankel schemes. To this end, we prove stability of the nonlinear case in Appendix 1.

However, this scheme is not unconditionally consistent. If the time step size τ is too large, the scheme (20) approximates a different PDE [28], namely a nonlinear variant of

the telegrapher’s equation. Such PDEs have also been used in image processing, see e.g. [85].

In the trainable setting, the conditional consistency is not an issue, but can even present a chance. It allows the network to learn a more suitable PDE for the problem at hand. In our experiments, we show that indeed, the unconditional stability of the Du Fort–Frankel scheme can help to achieve better results when only few residual blocks are available.

This scheme can be realised with a small change in the original diffusion block from Fig. 1 by adding an additional skip connection. The two skip connections are weighted by $\frac{4\tau\alpha}{1+2\tau\alpha}$ and $\frac{1-2\tau\alpha}{1+2\tau\alpha}$, respectively.

4.2 Fast Semi-Iterative Schemes

Another numerical scheme which also leads to the same concept from a different motivation is based on acceleration strategies for explicit schemes. Hafner et al. [51] introduced fast semi-iterative (FSI) schemes to accelerate explicit schemes for diffusion processes.

In a similar manner as the Du Fort–Frankel schemes, FSI extrapolates the diffusion result at a fractional time step $k + \frac{\ell}{L}$ with the previous fractional time step $k + \frac{\ell-1}{L}$ and a weight α_ℓ . For the explicit diffusion scheme (11), an FSI acceleration with cycle length L reads

$$\mathbf{u}^{k+\frac{\ell+1}{L}} = \alpha_\ell \left(\mathbf{u}^{k+\frac{\ell}{L}} - \tau \mathbf{K}^\top \Phi \left(\mathbf{K} \mathbf{u}^{k+\frac{\ell}{L}} \right) \right) + (1 - \alpha_\ell) \mathbf{u}^{k+\frac{\ell-1}{L}} \tag{21}$$

with fractional time steps $\ell = 0, \dots, L-1$ and extrapolation weights $\alpha_\ell := (4\ell + 2)/(2\ell + 3)$. One formally initialises with $\mathbf{u}^{k-\frac{1}{L}} := \mathbf{u}^k$.

The crucial difference to Du Fort–Frankel schemes is that FSI schemes use time-varying extrapolation coefficients instead of fixed ones. These coefficients are motivated by a box filter factorisation and allow a cycle to realise a super time step of size $\frac{L(L+1)}{3}\tau$. Thus, with one cycle involving L steps, one reaches a super step size of $\mathcal{O}(L^2)$ rather than $\mathcal{O}(L)$. This explains its remarkable efficiency [51].

Even though Du Fort–Frankel and FSI schemes have fundamentally different motivations, they lead to the same architectural changes, where additional weighted skip connections realise acceleration strategies. This is in line with observations in the CNN literature, see e.g. [69,78]. We visualise this concept at the example of an FSI architecture in Fig. 3.

FSI and Du Fort–Frankel schemes are just two representatives of a large class of extrapolation strategies, see e.g. [74,81,103]. The ongoing success of using momentum methods for training [91,100] and constructing [69,78,111] neural networks warrants an extensive investigation of these strategies in both worlds.




Fig. 3 FSI block realising the acceleration of an explicit diffusion step (21) with time-varying extrapolation parameters α_ℓ . A similar architecture with differently weighted skip connections arises for a Du Fort–Frankel scheme (20)

4.3 Implicit Schemes

So far, we have investigated variants of explicit schemes and their neural counterparts. However, implicit discretisations constitute another important solver class. We now show that such a discretisation of the generalised diffusion equation can be connected to a recurrent neural network (RNN). RNNs are classical neural network architectures, see e.g. [55]. At the same time, this translation inspires yet another way of leveraging skip connections.

A fully implicit discretisation of the diffusion equation (1) is given by

$$\mathbf{u}^{k+1} = \mathbf{u}^k - \tau \mathbf{K}^\top \Phi \left(\mathbf{K} \mathbf{u}^{k+1} \right). \tag{22}$$

The crucial difference as opposed to the explicit scheme lies in using the new signal \mathbf{u}^{k+1} within the flux function Φ . This yields a nonlinear system of equations, which we solve by means of a fixed point iteration with a cycle of length L :

$$\mathbf{u}^{k+\frac{\ell+1}{L}} = \mathbf{u}^k - \tau \mathbf{K}^\top \Phi \left(\mathbf{K} \mathbf{u}^{k+\frac{\ell}{L}} \right), \tag{23}$$

where $\ell = 0, \dots, L-1$, and where we assume that τ is sufficiently small to yield a contraction mapping.

For $L = 1$, we obtain the explicit scheme (11) with its ResNet interpretation. For larger L , however, different skip connections arise. They connect the layer at time step k with all subsequent layers at steps $k + \frac{\ell}{L}$ with $\ell = 0, \dots, L-1$.

There are two possible ways of interpreting this type of connection. One option is to regard this as a consequent extension of the extrapolation idea of FSI and Du Fort–Frankel schemes. Instead of only connecting a node to its two successors, the fixed point iteration above connects a node to L of its successors. Similar ideas have been used in the popular DenseNet architecture [56], where each layer is connected to all subsequent ones.

Another option is to interpret the repeated connection to u^k as a feedback loop, which in turn is closely connected to a recurrent neural network (RNN) architecture [55].

In their trainable nonlinear diffusion model, Chen and Pock [20] proposed a similar architecture. However, they explicitly supplement the diffusion process with an additional reaction term which results from the data term of the energy. Our feedback term is a pure numerical phenomenon of the fixed point solver.

We see that skip connections can implement a number of successful numerical concepts: forward difference approximations of the temporal derivative in explicit schemes, extrapolation steps to accelerate them, e.g. via FSI or Du Fort–Frankel schemes, and recurrent connections within fixed point solvers for implicit schemes.

5 Review: Multigrid Solvers and U-nets

Although the previously investigated numerical strategies and their neural counterparts can be efficient, they work on a single scale: The signal is considered at its original resolution at all points in time. However, using different signal resolutions in a clever combination can yield even higher efficiency.

This is the core idea of the large class of multigrid approaches [11,12,50]. They belong to the most efficient numerical methods for PDE-related problems and have been successfully applied to various tasks such as image denoising [13], inpainting [70], video compression [62], and image sequence analysis [14].

On the CNN side, architectures that work on multiple resolutions of the signal have become very successful. In particular, the shape of the popular U-net architecture [87] suggests that there is a structural connection between multigrid and CNN concepts.

By translating multigrid solvers into a U-Net architecture, we show that this is indeed the case, which serves as a basis for explaining the remarkable success of U-nets. Since both underlying concepts are not self-explanatory, we review them in the following before connecting them in the next section.

5.1 Multigrid Solvers for Nonlinear Systems

Multigrid methods [11,12,50] are designed to accelerate the convergence speed of standard numerical solvers such as the

Jacobi or the Gauß–Seidel method [94]. These solvers attenuate high-frequent components of the residual error very quickly, while low-frequent error components are damped slowly. This causes a considerable drop in convergence speed after a few iterations.

Multigrid methods remedy this effect by transferring the low-frequent error to a coarser grid, transforming them into high-frequent components. This allows a coarse grid solver to attenuate them more efficiently. By correcting the fine grid approximation with coarse grid results, convergence speed can be significantly improved.

In the following, we review the so-called full approximation scheme (FAS) [11] for a nonlinear system of equations. We consider a two-grid cycle as the basic building block of more complex multigrid solvers.

We are interested in solving a nonlinear system of equations of the form

$$A(x) = b, \tag{24}$$

with a nonlinear operator A and a right hand side vector b for an unknown coefficient vector x .

The two-grid FAS involves two grids with different step sizes: a fine grid of size h , and a coarse grid of size $H > h$. We denote the respective grid by superscripts. The following six steps describe the two-grid FAS:

1. *Presmoothing Relaxation* A standard solver is applied to the fine grid system $A^h(x^h) = b^h$. Given an initialisation x_0^h , it produces an approximation \tilde{x}^h to the solution with a reduced high frequency error.
2. *Restriction* In order to approximate low-frequent components of the error more efficiently, one transfers the problem to the coarse grid with the help of a restriction operator $R^{h \rightarrow H}$. One restricts both the residual $r^h = A^h(x^h) - b^h$ as well as the current approximation \tilde{x}^h to the coarse grid. One obtains two parts of the right hand side for the coarse grid problem: One part $b^H = R^{h \rightarrow H} r^h$ which is used directly, and a second one which we denote by $y^H = R^{h \rightarrow H} \tilde{x}^h$ serving as the argument for the nonlinear operator A^H . The coarse grid problem then reads

$$A^H(x^H) = A^H(y^H) + b^H. \tag{25}$$

If we express the desired solution x^H in terms of the error e^H by $x^H = y^H + e^H$, then we see that this equation is solved for the full approximation rather than the error alone, in contrast to a linear multigrid scheme. Hence, this scheme is called the full approximation scheme. If the operator A is linear, FAS reduces to a linear multigrid scheme.

A standard choice for $\mathbf{R}^{h \rightarrow H}$ is a simple averaging. However, finding suitable restriction operators is a difficult task, which motivates researchers to even learn such operators, see e.g. [43,58].

3. *Coarse Grid Computation* Solving the coarse grid problem with a standard solver produces an error approximation $\tilde{\mathbf{x}}^H$.
4. *Prolongation* The approximation on the coarse grid needs to be transferred to the fine grid again. To this end, one applies a prolongation operator $\mathbf{P}^{H \rightarrow h}$. Since a coarse grid solution $\tilde{\mathbf{x}}^H$ is a full approximation, we need to compute the approximation to the error by $\tilde{\mathbf{x}}^H - \mathbf{y}^H$. This error approximation is then transferred to the fine grid via $\mathbf{P}^{H \rightarrow h}$. A standard choice for $\mathbf{P}^{H \rightarrow h}$ is a nearest neighbour interpolation, but as for the restriction operator, finding a good prolongation operator is not easy.
5. *Correction* The fine grid approximation $\tilde{\mathbf{x}}^h$ is corrected with the upsampled coarse grid error approximation $\mathbf{P}^{H \rightarrow h}(\tilde{\mathbf{x}}^H - \mathbf{y}^H)$ to produce a new approximation

$$\tilde{\mathbf{x}}_{\text{new}}^h = \tilde{\mathbf{x}}^h + \mathbf{P}^{H \rightarrow h}(\tilde{\mathbf{x}}^H - \mathbf{y}^H). \tag{26}$$

6. *Postsmoothing Relaxation* Finally, one applies another solver on the fine grid to smooth high frequent errors which have been introduced by the correction step.

The two-grid FAS will serve as the starting point for translating multigrid concepts into the a U-net formulation, in an extension to the linear connections from our conference publication [1].

5.2 Review: U-nets

U-nets [87] are another popular neural architecture. They process information on multiple scales by repeatedly down- and upsampling the input data, interleaved with a series of convolutional network layers. This multiscale analysis makes them well-suited for medical image analysis tasks such as segmentation [27,87], but also for pose estimation [75] and shape generation [32].

A two-level U-net with fine grid size h and coarse grid size H has the following structure:

1. An input signal \mathbf{f}^h is fed into a series of general convolutional layers which we denote by $\mathbf{C}_1^h(\cdot)$. The resulting output signal is denoted by $\tilde{\mathbf{f}}^h = \mathbf{C}_1^h(\mathbf{f}^h)$. Originally, these layers are assumed to be feed-forward convolutional layers, but they can also be replaced by any other suitable layer type such as residual layers.

2. The fine grid signal $\tilde{\mathbf{f}}^h$ is transferred to a coarser grid with a restriction operator $\mathbf{R}^{h \rightarrow H}$, yielding a coarse grid signal $\mathbf{f}^H = \mathbf{R}^{h \rightarrow H} \tilde{\mathbf{f}}^h$.
3. On the coarse grid, another series of convolutional layers $\mathbf{C}^H(\cdot)$ is applied to the signal, yielding a modified coarse grid signal $\tilde{\mathbf{f}}^H = \mathbf{C}^H(\mathbf{f}^H)$.
4. The modified coarse grid signal is upsampled with a prolongation operator $\mathbf{P}^{H \rightarrow h}$.
5. With the help of a skip connection, the modified fine grid signal $\tilde{\mathbf{f}}^h$ and the upsampled coarse grid signal $\mathbf{P}^{H \rightarrow h} \tilde{\mathbf{f}}^H$ are added together. This produces a new fine grid signal $\tilde{\mathbf{f}}_{\text{new}}^h$. While the original U-net formulation [87] suggests to concatenate both signals, other works such as [75] simply add the signals. For our following discussion of connections between U-nets and multigrid schemes, we focus on the latter variant.
6. Lastly, another series of convolutional layers $\mathbf{C}_2^h(\cdot)$ is applied to the new fine grid signal, producing the final output signal $\hat{\mathbf{f}}^h$.

We visualise this architecture in Fig. 4a.

6 From Multigrid to U-nets


Now we show how one can express FAS in terms of a U-net architecture. For our U-net, we use multiple network channels which carry the variables required by FAS. Even though not all variables are used at each point in the network, we keep the channel number consistent for the sake of simplicity.

We track the FAS variables in dedicated channels only for didactic reasons, as a direct translation shows that a U-net architecture is sufficient for representing FAS. When practically implementing FAS, this overhead can be spared.


Firstly, let us assume that we are given suitable solvers $\mathbf{S}_A^h(\cdot)$, $\mathbf{S}_A^H(\cdot)$ for the nonlinear operators on the fine and coarse grid, respectively. To be able to use the two-grid cycle as a recursive building block, we assume that all solvers approximate solutions for nonlinear systems of the form $\mathbf{A}(\mathbf{x}) = \mathbf{A}(\mathbf{y}) + \mathbf{b}$, regardless of the grid.

To this end, we always keep track of the iteration variable \mathbf{x} , the nonlinear right hand side \mathbf{y} , and the linear right hand side \mathbf{b} . In addition, we track the residual \mathbf{r} . By appropriately modifying these variables, we can ensure that the solvers always act on the desired system, despite having a common specification.

1. *Presmoothing Relaxation* The first instance of the fine grid solver obtains an initial iteration variable \mathbf{x}_0^h , which can be $\mathbf{0}$ or a more sophisticated guess. Since the first solver is supposed to solve $\mathbf{A}^h(\mathbf{x}^h) = \mathbf{b}^h$, we simply set $\mathbf{y}^h = \mathbf{0}$.



(a) U-net architecture for an input f^h .



(b) FAS two-grid cycle in the form of a U-net utilising four network channels. Besides the iteration variable \mathbf{x} , the network tracks variables \mathbf{b} and \mathbf{y} as linear and nonlinear parts of the system’s right hand side, as well as the residual \mathbf{r} .

Fig. 4 Architectures for a general U-net (a) and an FAS two-grid cycle (b)

In addition, we provide the linear right hand side \mathbf{b}^h . A residual is not needed as an input.

The solver produces a preliminary approximation $\tilde{\mathbf{x}}^h$, and passes the right hand side components \mathbf{y}^h and \mathbf{b}^h through without changes. It also computes the residual $\mathbf{r}^h = \mathbf{b}^h - \mathbf{A}^h(\tilde{\mathbf{x}}^h)$ as an additional output.

2. *Restriction* As the downsampling is now explicitly concerned with four channels, the corresponding operator in our U-net is a 4×4 block matrix. We apply the multigrid restriction operator only to certain channels.

The coarse grid initialisation \mathbf{x}_0^H can be set to $\mathbf{0}$, taking no information from the fine grid. The coarse, nonlinear right hand side $\mathbf{y}^H = \mathbf{R}^{h \rightarrow H} \tilde{\mathbf{x}}^h$ is given by the downsampled fine grid approximation. The corresponding linear right hand side $\mathbf{b}^H = \mathbf{R}^{h \rightarrow H} \mathbf{r}^h$ is the downsampled residual. In contrast to our linear correspondences in [1], the restriction step in FAS fits a U-net interpretation even better, since the approximation $\tilde{\mathbf{x}}^h$ itself is restricted, as is the case in the U-net.

3. *Coarse Grid Computation* The coarse solver follows the same specification as the fine grid one. However, since \mathbf{y}^H is not set to $\mathbf{0}$ at this point, the coarse solver actually solves the desired system $\mathbf{A}^H(\mathbf{x}^H) = \mathbf{A}^H(\mathbf{y}^H) + \mathbf{b}^H$. It produces a coarse approximation $\tilde{\mathbf{x}}^H$ and a residual \mathbf{r}^H , while leaving the right hand side components unchanged.

4. *Prolongation* The upsampling step allows to prepare the coarse grid variables in such a way that the skip connection automatically performs the correct additions.

The first row of the matrix operator ensures that we upsample the correction $\mathbf{P}^{H \rightarrow h} \tilde{\mathbf{x}}^H - \mathbf{P}^{H \rightarrow h} \mathbf{y}^H$. Note that this is equivalent to the FAS formulation $\mathbf{P}^{H \rightarrow h}(\tilde{\mathbf{x}}^H - \mathbf{y}^H)$ if the prolongation operator is linear. This is no limitation,

however, since for the nonlinear case we can require the solvers to directly output the difference $\tilde{\mathbf{x}}^H - \mathbf{y}^H$.

The right hand side components \mathbf{y}^H and \mathbf{b}^H are not used in the upsampling, as the fine grid right hand side is supposed to be passed on. The same holds for the residual, as it is not relevant to the second fine grid solver. It is only needed in case one adds another coarser level to the cycle.

5. *Correction* In the correction step, the fine approximation $\tilde{\mathbf{x}}^h$ is appropriately corrected, and the fine grid right hand side components \mathbf{y}^h and \mathbf{b}^h are forwarded.

6. *Postsmoothing Relaxation* Another instance of the fine grid solver solves the problem $\mathbf{A}^h(\mathbf{x}^h) = \mathbf{b}^h$. The nonlinear part \mathbf{y}^h of the right hand side is still set to $\mathbf{0}$, ensuring that the correct system is solved.

The resulting architecture is visualised in Fig. 4b. This shows that U-nets share essential structural properties with multigrid methods. In particular, employing multiple image resolutions connected through pooling and upsampling operations, as well as horizontal skip connections which realise correction steps are the keys for the success of both methods. This leads us to believe that at their core, U-nets realise a sophisticated multigrid strategy.

6.1 V-Cycles, W-Cycles and Full Multigrid

Our connections between two-grid FAS and U-nets are the basic building block for more advanced multigrid strategies.

So-called V-cycles arise from recursively stacking the two-grid FAS. Moreover, W-cycles can be built by concatenating several V-cycles. Optimising the depth and length of

these cycles can lead to vast efficiency gains over direct solution strategies.

On the CNN side, the corresponding concept of U-nets with more levels as well as concatenations thereof is successful in practice: Typical U-nets work on multiple resolutions [87], and so-called stacked hourglass models [75] arise by concatenating multiple V-cycle architectures.

A full multigrid (FMG) strategy solves a problem on multiple grids by successively concatenating V- and W-cycles, usually starting at the coarsest grid and progressing towards the finest one. In our experiments in Sect. 7.2, we will construct a trainable FMG model based on the two-grid FAS network to approximate the solution of an inpainting problem. This shows that our model reduction of the full U-net is successful in practice and inspires new design strategies for U-nets.

7 Experimental Evaluations

Let us now show that our findings are also of practical relevance. Our experiments are divided into two parts. First, we evaluate the proposed symmetric ResNet architectures, along with their variations and nonmonotone activation functions for a denoising problem.

In a second experiment, we make use of our connections between multigrid and U-nets to learn an efficient solver for diffusion-based sparse inpainting, based on a trainable FAS architecture.

7.1 Symmetric ResNets and Nonmonotone Activations

Since we motivate our network designs through numerical algorithms for a diffusion problem, we start with an elementary comparison on a denoising problem. We deliberately choose a denoising problem, since it is a prime example of a well-posed problem, for which the presented numerical schemes can be easily applied.

In a second step, we refine the simple network structures to more and more complex ones, approaching the standard neural network design. This shows the extent to which our networks can compete with off-the-shelf ResNets.

7.1.1 Experimental Setup

We compare symmetric ResNets and their Du Fort–Frankel and FSI extensions with the original ResNet architecture [54]. As activation functions we allow the ReLU [73], Charbonnier [18], and Perona–Malik [80] activation functions.

The original ResNets train two filter kernels per ResNet block, along with two bias terms. The symmetric ResNets on the other hand only train one filter kernel per block, without

any bias terms. We only consider kernels of width three. For maximal transparency, we do not use any additional optimisation layers such as batch normalisation.

When using Charbonnier and Perona–Malik activations, we always train the corresponding contrast parameter λ . The Du Fort–Frankel networks also learn the extrapolation parameter α , and the FSI networks train individual extrapolation parameters of each block.

In addition, all models train their numerical parameters such as time step size and extrapolation parameters. We restrict the time step size τ to our stability condition (14) to obtain a stable symmetric ResNet model. In the case of the Du Fort–Frankel extension we restrict the extrapolation parameter α to the bound in Appendix 1, thus also yielding a stable scheme. For FSI, we restrict the extrapolation parameters α_ℓ to the range $[0, 2]$. This preserves the extrapolation character of the scheme. However, no stability theory is available in the case of learned extrapolation parameters.

We evaluate the networks on a synthetic dataset of 1D signals which are piecewise affine, with jumps between the segments. This design highlights the ability of the different approaches to preserve signal discontinuities. The signals are of length 256 and are composed of linear segments that span between $\frac{1}{10}$ and $\frac{1}{2}$ of the signal length. Their values lie within the interval $[0, 255]$.

Finally, we add Gaussian noise of standard deviation $\sigma = 10$ to the signals, without clipping out of bounds values. This yields pairs of corrupted and ground truth signals. The training dataset contains 10000 such pairs, and the test and validation datasets contain 1000 pairs each. As a measure of denoising quality, we choose the peak-signal-to-noise ratio (PSNR), where higher values indicate better denoising performance.

For a fair comparison, we train all network configurations in the same fashion. We use the Adam optimiser [59] with a learning rate of 0.001 for at most 2000 training epochs, and choose the average mean square error (MSE) over the training dataset as an optimisation objective.

The filter weights are initialised according to a uniform random distribution with a range of $[-0.1, 0.1]$. The contrast parameters λ , the time step sizes τ , and the extrapolation weights α are initialised with fixed values of 15, 1.0, and 1.0, respectively. Out of several random initialisations, we choose the best performing one.

7.1.2 Evaluation of Model Components

We first evaluate the potential of the proposed network blocks on an individual level. To this end, we train the architectures for varying amounts of residual blocks. However, all blocks share their weights, and we also use only a single network channel.




Fig. 5 Denoising quality of network architectures with varying depth. We use a single channel, and weights between all blocks are shared. Each plot is concerned with a different activation function. Architectures

This configuration is closest to the interpretation of explicit schemes, and it allows us to investigate the approximation qualities of the different architectures within a tightly controlled frame.

Figure 5 presents the denoising quality of the architectures in dependence of the number of residual blocks. Each plot is concerned with a different activation function.

Firstly, we compare the different network architectures. As the symmetric ResNet is guaranteeing stability and uses less than half of the parameters of the standard ResNet, it performs slightly worse. This is not surprising, since there is a natural tradeoff between performance (high approximation quality) and stability, as is well known in the field of numerical analysis. Nevertheless, when enough blocks are provided, the symmetric ResNet catches up to the standard one.

The acceleration methods of Du Fort-Frankel and FSI outperform the symmetric ResNet and yield comparable performance to the standard ResNet. The trainable extrapolation parameters help both methods to achieve better quality especially when not enough residual blocks are provided to reach a sufficient denoising result. This is in full accordance with our expectations. When enough steps are provided and no extrapolation is required, both methods are on par with the standard and symmetric ResNets.

A side-by-side comparison yields insights into the performance of different activation functions. We observe that the performance of ReLU networks is only slightly better than classical linear diffusion [57]. This shows that the ReLU is not suited for our denoising problem, regardless of the network architecture. After as few as three network blocks, the improvement of deeper networks is only marginal.

In contrast, both the Charbonnier and the Perona-Malik activations are much more suitable. The nonmonotone Perona-Malik activation function yields the best denoising performance, as our diffusion interpretation suggests. When using the original ResNet with a diffusion-inspired activa-

tion function, tremendous performance gains in comparison with the ReLU activation can be achieved. This shows that in this experimental setting, the activation is the key to a good performance.

7.1.3 Optimality of Diffusion Processes

Interestingly, the standard ResNet with a diffusion activation naturally learns a symmetric filter structure with biases close to 0.

For example, the ResNet with Perona-Malik activation, a grid size of $h = 1$, and 20 shared residual blocks learns an inner kernel $k_1 = (0.922, -0.917, 0.006)^T$, an outer kernel $k_2 = (0.051, 0.437, -0.489)^T$ and biases $b_1 = 1.6 \cdot 10^{-1}$ and $b_2 = 1.2 \cdot 10^{-5}$. If we factor out the time step size limit of $\tau = 0.5$ for this setting from the outer kernel k_2 , we see that it transforms into $\tilde{k}_2 = (0.102, 0.874, -0.978)^T$. It becomes apparent that the kernels approximately fulfil the negated symmetric filter structure.

Moreover, the kernels closely resemble rescaled standard forward and backward difference discretisations. This is surprising, as a kernel of width three allows to learn derivative operators of second order, but a first order operator appears to yield already optimal quality.

This shows that in this constrained setting, second order diffusion processes are an optimal model which is naturally learned by a residual network.

7.1.4 Time Dynamic Case

In a practical setting, the residual blocks typically do not share their weights, but train them independently. If the parameters evolve smoothly over the blocks, we can interpret this as an approximation of a time dynamic PDE model.

To investigate the performance of the proposed architectures in this setting, we train the parameters of each block individually, but enforce a certain smoothness between them.




Fig. 6 Denoising quality of network architectures with varying depth and a single channel. The parameters are smoothly changing between the residual blocks. Each plot is concerned with a different activa-

tion function. The proposed architectures can outperform the standard ResNet by saving a large amount of parameters

If the parameter vector of a block at time level k is given by θ^k , we add a regulariser

$$\beta \sum_{k=1}^K \tau \left(\frac{\theta^k - \theta^{k-1}}{\tau} \right)^2 \tag{27}$$

to the loss function. Here, a smoothness parameter β controls the amount of smoothness between the blocks, with higher values of β leading to smoother evolutions. This expression approximates the continuous temporal regulariser

$$\beta \int_0^T (\partial_t \theta(t))^2 dt, \tag{28}$$

which enforces smoothness of the continuous parameter evolution $\theta(t)$. The regularisation ensures that the learned filters change smoothly throughout the layers. This is essential for the numerical scheme to be consistent with the continuous limit case where the step size τ tends towards 0, see also [93].

For the residual network, where no time step size τ is learned explicitly, we set the time step size to the inverse of the number of blocks. This requires to use a different smoothness parameter β . We tune the smoothness parameters for all architectures such that their parameters exhibit similarly smooth evolutions over time. Numerical parameters such as time step size and extrapolation parameters are not affected by the regulariser.

Figure 6 presents the performance of time dynamic architectures with a single channel. We use $\beta = 5$ for the standard ResNets, and $\beta = 10$ for all other architectures. In contrast to the previous comparisons, we now compare the denoising quality against the number of network parameters. This allows us to measure performance against model complexity.

For the symmetric ResNet, we observe the same behaviour as in the setting without a temporal dynamic. The overall best performance is still on par with the respective classi-

cal diffusion process for the Charbonnier and Perona-Malik activation functions.

However, the time dynamic allows this model to achieve better denoising quality for a fewer number of blocks. For example, the symmetric ResNet with Perona-Malik activation and seven residual blocks can achieve a denoising quality of 36.51 dB if weights are shared, but already 37.08 dB with a regularised temporal dynamic. Similar observations for classical diffusion processes with a time dynamic diffusivity function can be found in [36]. Yet, this effectiveness comes at the price of additional parameters.

The Du Fort-Frankel and FSI networks allow for higher efficiency at the cost of more network parameters. Especially in the case of FSI, the trainable extrapolation parameters help to achieve significantly better performance when more residual blocks are provided. This is in accordance with observations in the literature [69].

The proposed architectures outperform the standard ResNet for the same amount of parameters when using Charbonnier and Perona-Malik activations. This shows that the model reduction to a symmetric convolution structure is indeed fruitful. Moreover, the ranking of activation functions remains the same in most cases.

None of the architectures significantly outperform the classical Perona-Malik diffusion process. This supports our claim that these tightly constrained networks realise a numerical algorithm at their core. Different architectures can solve the problem with varying efficiency, but they converge towards the same result. This will only change when we allow for more flexibility within the network architecture, e.g. by utilising multiple network channels.

7.1.5 Towards Larger Networks

So far, we have only considered architectures with a single network channel. However, typical CNNs show their full potential when using multiple channels. In this case, the sim-




Fig. 7 Denoising quality of network architectures with varying depth and $C = 16$ network channels. Each plot is concerned with a different activation function. The proposed architectures outperform the standard

ResNet for the same amount of parameters. In this setting, the acceleration strategies are on par with the symmetric ResNet, and the margin between activation functions becomes smaller

plicity of the ReLU activation is compensated by a rich set of convolutions between channels.

We now extend our evaluations to architectures with multiple network channels. For simplicity, we leave the number of channels constant throughout the network. To this end, we copy the input signal into C channels. The output signal is computed as the average over the individual channel results. In between, we employ the proposed symmetric residual blocks which now use $C \times C$ block convolution matrices with the appropriate stability constraints. These convolutions can be interpreted as ensembles of differential operators which are applied to the signal. The channels are then activated individually and convolved again with the adjoint counterparts of the differential operators.

To allow for maximum flexibility and performance, we remove the temporal parameter regularisation in this experiment.

The denoising performance of the networks is visualised in Fig. 7 for the case of $C = 16$ channels.

This experiment is the first instance where all architectures significantly outperform their respective classical diffusion counterparts. The multichannel architecture allows to approximate a more sophisticated denoising model, as information in the various channels is exchanged by means of multichannel convolution operators.

All proposed architectures can still outperform the standard ResNet for the same amount of parameters. In this case, the extrapolation methods are on par with the symmetric ResNet. Interestingly, the ranking of activation functions remains the same, albeit with a much smaller margin. The symmetric ResNet with 20 residual blocks yields PSNR values of 40.04 dB, 40.44 dB and 40.88 dB for the ReLU, Charbonnier and Perona–Malik activations, respectively. We conclude that the more complex the network, the less the activation function matters for performance. On the contrary, this means that networks might be drastically reduced in size

when trading network size for sophisticated activation function design.

7.2 Learning a Multigrid Solver for Inpainting

So far, it is not clear if our interpretation of the two-grid FAS is a reasonable model reduction of a full U-net. To prove that this interpretation is indeed of practical relevance, we show how it can be used to learn a multigrid solver for edge-enhancing diffusion inpainting of images [33,104,107].

Diffusion-based inpainting aims to restore an image from a sparse set of known data points [33,107]. Diffusion processes allow for a high reconstruction quality even for extremely sparse known data, making them an interesting tool for image compression applications, see e.g. [33,98]. Of particular interest is the edge-enhancing diffusion (EED) operator [104] as it allows to reconstruct discontinuous image data such as edges.

As a proof of concept, we construct a network implementing a full multigrid structure. We replace the prescribed nonlinear solvers by trainable feed-forward layers that learn to approximate the PDE at hand.

7.2.1 Edge-Enhancing Diffusion Inpainting

The EED inpainting problem can be formulated as follows. Given a set of known image data $f : K \rightarrow \mathbb{R}$ on a subset K of the image domain $\Omega \subset \mathbb{R}^2$, the goal is to compute a reconstruction u as the solution of the PDE

$$(1 - c(\mathbf{x}))\nabla^T(D\nabla u) - c(\mathbf{x})(u - f) = 0. \tag{29}$$

Here, $c(\mathbf{x})$ is a binary confidence function indicating whether the data at position \mathbf{x} is known or not. The case $c(\mathbf{x}) = 1$ indicates known data, yielding $u = f$. The case $c(\mathbf{x}) = 0$ indicates that the data need to be reconstructed by EED

inpainting. Consequently, $c(x)$ is the characteristic function of the inpainting mask K .

The EED operator $\nabla^\top(D\nabla u)$ uses a diffusion tensor $D = g(\nabla u_\sigma \nabla u_\sigma^\top)$ based on a Gaussian smoothed gradient ∇_σ and a nonlinear diffusivity function g . It is a 2×2 positive semidefinite matrix, which is designed to propagate information along locally dominant structures [104]. As a diffusivity function, we use the Charbonnier diffusivity [18], which relies on a contrast parameter λ .

7.2.2 Experimental Setup

Our trainable FMG architecture is designed as follows: Instead of prescribing nonlinear solvers on each grid, we employ a series of convolutional layers with trainable weights. The remainder of the architecture is fixed: We set the restriction operators to a simple averaging over a 2×2 pixel neighbourhood, and the prolongation operators to nearest neighbour interpolation.

Since FMG employs the same solver on each grid, we realise this idea also in our network by sharing the weights between all solvers for a specific grid. This drastically reduces the amount of parameters and incites that an iterated application of the solvers performs the correct computations.

Instead of training our network by minimising a Euclidean loss between ground truth data and inpainting reconstructions, we use the absolute residual of a discretisation of the inpainting equation (29) as a loss function. This is closely related to the idea of deep energies [37], where one chooses a variational energy as a loss function. Since we do not have such an energy available for EED inpainting, we resort to minimising the absolute residual of the associated Euler–Lagrange equation, which is given by (29). This guarantees that the trained architecture realises EED inpainting as efficiently as possible. To discretise the Euler–Lagrange equation, we employ the standard discretisation from [108].

Whereas classical solution methods for the inpainting problem specify the known data $u = f$ on $\Omega \setminus K$ by means of Dirichlet boundary conditions, we leave it to the network




Fig. 8 Visualisation of the full multigrid strategy which we employ in our experiments. Dashed horizontal lines denote the three grids, and grids become coarser from top to bottom. Each circle denotes an instance of a solver

to reproduce also the known data. We have found that this leads to a better approximation quality.

The inpainting masks consist of randomly sampled pixels with a density d as a percentage of the number of image pixels. Since the masks are also required on coarser grids to compute the residual within the FMG architecture, we down-sample them by putting defining a coarse pixel as known, if at least one pixel in the 2×2 cell on the fine grid is known.

We train the architecture on a subset of 1000 images of the ImageNet dataset [92] with the Adam optimiser [59] with standard settings.

7.2.3 Evaluation of the Full Multigrid Network

We construct a full multigrid network using three grids of size h , $2h$ and, $4h$. The order in which the problem is solved on different grids is given by $[4h, 2h, 4h, 2h, h, 2h, 4h, 2h, 4h, 2h, h]$. This is the simplest FMG strategy that can be employed in a setting involving three grids and serves as a proof-of-concept architecture. We visualise this strategy in Fig. 8. Thus, we employ 11 solvers, each using 12 feed-forward convolutional layers with 20 channels and ReLU activations. Weights are shared for solvers on the same grid.

We train this network on an EED inpainting problem with random masks of 20% density. The EED parameters $\lambda = 0.93$, $\sigma = 0.97$ have been optimised for inpainting quality with a simple grid search.

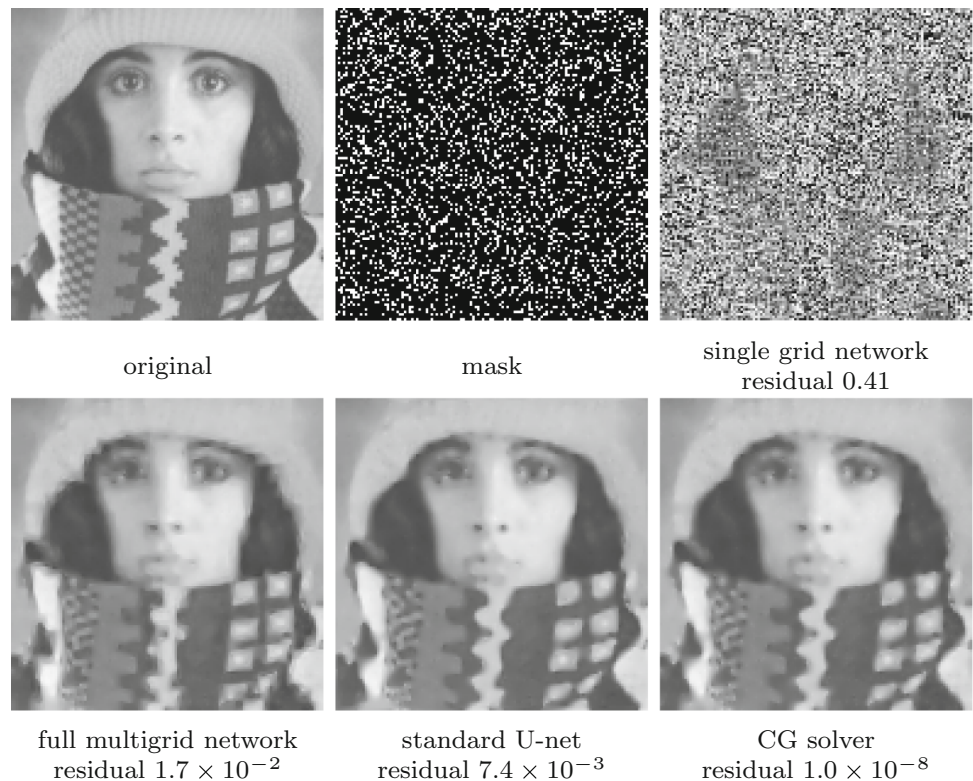
To show how the FMG network can benefit from the multigrid structure, we compare it against two networks with the same amount of parameters. One network solves the problem only on a single grid by using 25 layers and 24 channels. Moreover, we compare our FMG network to a standard U-net with addition with three scales, 17 channels and 2 layers per scale. All three models contain 1.2×10^5 trainable parameters.

Figure 9 shows the inpainting results for the networks at the example of the image *trui*. Moreover, we present the true inpainting result obtained from a conjugate gradient (CG) solver for EED inpainting.

In contrast to the single grid network, the FMG network and the U-net are able to approximate the EED inpainting result. The FMG and U-net results are visually comparable, while the residual of the U-net is slightly better. This does not only show that the multigrid structure is an adequate network design, but also that a standard U-net is not able to obtain a much better solution in this case. From the perspective of numerical algorithms, this is expected, since we know that multigrid methods are highly efficient for these problems.

The advantage of the FMG network is that adding further solvers on the three scales will not inflate the number of parameters as these solvers are shared. This is in contrast to the U-net, where any addition increases the trainable parameter set. The architectural design of the FMG network

Fig. 9 Reconstruction quality of a single grid network, a full multigrid network, a standard U-net, and a classical CG solver for the EED inpainting problem. All results use $\lambda = 0.93$, $\sigma = 0.97$ and the same random mask with 20% density. Both the full multigrid network and the standard U-net approximate an EED inpainting result, while a single grid network fails to do so



suggests that U-nets should also be constructed in a similar way. In practice, already concatenating multiple U-nets [75] is a successful idea. Instead of a single down- and upsampling pass, multiple alternating computations on different resolutions should be beneficial.

8 Discussion and Conclusions

We have shown that numerical algorithms for diffusion evolutions share structural connections to CNN architectures and inspire novel design concepts.

Explicit diffusion schemes yield a specific form of residual networks with a symmetric filter structure, for which one can prove Euclidean stability. Moreover, this architecture saves half of the network parameters, and its stability constraint is easy to implement without affecting its performance. In addition, our connection suggests the use of a diffusion flux function as an activation, revitalising the idea of nonmonotone activation functions. We have shown that these activations perform well for a denoising task, even when using them within standard architectures in a plug-and-play fashion.

By investigating accelerated explicit schemes and implicit schemes, we have justified the effectiveness of skip connections in neural networks. They realise time discretisations in explicit schemes, extrapolation terms to increase their efficiency, and recurrent connections in implicit schemes with fixed point structure. In practice, the resulting architectures are particularly useful when training networks with small amounts of layers.

Lastly, our connection between multigrid concepts and U-net architectures serves as a basis for explaining their success. We have shown that a U-net architecture is able to implement a full multigrid strategy, which allows to learn efficient solutions for PDEs which are typically hard to solve. By directly using the residual norm as a loss function, we can guarantee that the network approximates the PDE at hand. This suggests to extend the standard U-net architecture in a full multigrid fashion.

Our philosophy of identifying numerical concepts as core building blocks of neural architectures has proven to be fruitful. Our direct translation has yielded structural insights into popular neural networks and inspired well-founded neural building blocks with practical relevance. An overview of the detailed connections that we have encountered is presented in Table 1.

Our numerical perspective on neural networks differs from most viewpoints in the literature. However, it provides a blueprint for directly translating a plethora of numerical

Table 1 Overview of the connections between numerical and neural concepts which we have encountered

Numerical Concept	Neural Concept	
numerical algorithm	neural network architecture	} interpretation
evolution equation	trained neural network	
specification of nonlinear dynamics	training	
explicit scheme	residual network	} coarse connections
implicit scheme	recurrent network	
multigrid techniques	U-net architectures	
time level	residual block	} detailed connections
flux function	activation function	
spatial derivative	convolution kernel	
temporal derivative	skip connection	
acceleration strategies	weighted skip connections	

strategies into well-founded and practically relevant neural building components. We hope that this line of research leads to a closer connection of both worlds and to hybrid methods that unite the stability and efficiency of modern numerical algorithms with the performance of neural networks.

Funding Open Access funding enabled and organized by Projekt DEAL. This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 741215, ERC Advanced Grant INCOVID).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Stability Proof for Generalised Du Fort–Frankel Schemes

In this section, we extend the stability proof for linear Du Fort–Frankel Schemes of [41] to the nonlinear setting.

First we rewrite the scheme (20) as a multistep method, obtaining

$$\begin{pmatrix} \mathbf{u}^{k+1} \\ \mathbf{u}^k \end{pmatrix} = \begin{pmatrix} \frac{4\tau\alpha}{1+2\tau\alpha}\mathbf{I} - \frac{2\tau}{1+2\tau\alpha}\mathbf{A}(\mathbf{u}^k) & \frac{1-2\tau\alpha}{1+2\tau\alpha}\mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{u}^k \\ \mathbf{u}^{k-1} \end{pmatrix}. \tag{30}$$

Here, we have abbreviated $\mathbf{A}(\mathbf{u}^k) = \mathbf{K}^\top \mathbf{G}(\mathbf{u}^k) \mathbf{K}$.

To analyse the stability of the Du Fort–Frankel scheme, we have to show that all eigenvalues of the matrix of the multistep method (30) have an absolute value less than or equal to one. Note that this matrix is not symmetric such that it might have complex eigenvalues.

Let us first define as short-hand notations

$$\begin{aligned} \mathbf{Q} &:= \frac{4\tau\alpha}{1+2\tau\alpha}\mathbf{I} - \frac{2\tau}{1+2\tau\alpha}\mathbf{A}(\mathbf{u}^k), \\ \mathbf{B} &:= \begin{pmatrix} \mathbf{Q} & \frac{1-2\tau\alpha}{1+2\tau\alpha}\mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{pmatrix}. \end{aligned} \tag{31}$$

We start with the naive approach to compute eigenvalues of \mathbf{B} : $\mu \in \mathbb{C}$ is an eigenvalue of \mathbf{B} if

$$\det(\mathbf{B} - \mu\mathbf{I}) = \det\left(\begin{pmatrix} \mathbf{Q} - \mu\mathbf{I} & \frac{1-2\tau\alpha}{1+2\tau\alpha}\mathbf{I} \\ \mathbf{I} & -\mu\mathbf{I} \end{pmatrix}\right) = 0. \tag{32}$$

As $\mathbf{B} - \mu\mathbf{I}$ is a block matrix containing square blocks of the same shape where the lower two blocks commute, we have

$$\begin{aligned} \det(\mathbf{B} - \mu\mathbf{I}) &= \det\left(\left(\mathbf{Q} - \mu\mathbf{I}\right)\left(-\mu\mathbf{I}\right) - \frac{1-2\tau\alpha}{1+2\tau\alpha}\mathbf{I}\right) \\ &= \det\left(\mu^2\mathbf{I} - \mu\mathbf{Q} - \frac{1-2\tau\alpha}{1+2\tau\alpha}\mathbf{I}\right). \end{aligned} \tag{33}$$

To proceed from here, it is reasonable to involve the eigenvalues of the matrix \mathbf{Q} . As \mathbf{Q} is real-valued and symmetric, there exist an orthogonal matrix \mathbf{V} of eigenvectors of \mathbf{Q} and a diagonal matrix $\mathbf{\Gamma}$ with the eigenvalues γ of \mathbf{Q} on its diagonal such that

$$\mathbf{Q} = \mathbf{V}\mathbf{\Gamma}\mathbf{V}^\top. \tag{34}$$

As \mathbf{V} is an orthogonal matrix, it holds that

$$\mathbf{I} = \mathbf{V}\mathbf{V}^\top. \tag{35}$$

Plugging both of these relations into (32) yields

$$\begin{aligned} \det(\mathbf{B} - \mu \mathbf{I}) &= \det\left(\mu^2 \mathbf{V} \mathbf{V}^T - \mu \mathbf{V} \mathbf{\Gamma} \mathbf{V}^T - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} \mathbf{V} \mathbf{V}^T\right) \\ &= \det\left(\mathbf{V} \left(\mu^2 \mathbf{I} - \mu \mathbf{\Gamma} - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} \mathbf{I}\right) \mathbf{V}^T\right) \\ &= \det(\mathbf{V}) \det\left(\mu^2 \mathbf{I} - \mu \mathbf{\Gamma} - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} \mathbf{I}\right) \det(\mathbf{V}^T) \\ &= \det\left(\mu^2 \mathbf{I} - \mu \mathbf{\Gamma} - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} \mathbf{I}\right), \end{aligned} \tag{36}$$

since the determinants of the orthogonal matrices \mathbf{V} and \mathbf{V}^T are both 1. The remaining determinant is concerned with a diagonal matrix. Thus, it is equal to the product of the diagonal elements, i.e.

$$\det(\mathbf{B} - \mu \mathbf{I}) = \prod_{j=1}^N \left(\mu^2 - \mu \gamma_j - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha}\right). \tag{37}$$

For μ to be an eigenvalue of \mathbf{B} , we need that this product vanishes. This is exactly the case if one or more factors in the product vanish. Hence, we get that for a fixed eigenvalue γ of \mathbf{Q} , an eigenvalue μ of \mathbf{B} satisfies

$$\mu = \frac{\gamma}{2} \pm \sqrt{\frac{\gamma^2}{4} + \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha}}. \tag{38}$$

For stability of (30), we need that $|\mu| \leq 1$ for all solutions μ for all eigenvalues γ of \mathbf{Q} . To proceed further, we consider the discriminant $\frac{\gamma^2}{4} + \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha}$. Since we know that \mathbf{Q} is real-valued and symmetric, it follows that γ is a real number. Thus, γ^2 is positive. We also know that τ and α are positive, such that

$$-1 < \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} < 1. \tag{39}$$

Therefore, it is possible for the discriminant to have negative values, which results in complex eigenvalues μ . Let us therefore distinguish the three cases of negative discriminant, vanishing discriminant, and positive discriminant. We will see that the last one is the only case which introduces a lower bound on α for unconditional stability.

Vanishing Discriminant First, we consider a vanishing discriminant, which can only happen if $2\tau\alpha \geq 1$. This yields

$$\gamma = \pm 2 \sqrt{\frac{2\tau\alpha - 1}{2\tau\alpha + 1}}. \tag{40}$$

Thus, the eigenvalues of \mathbf{B} are given by

$$\mu = \frac{\gamma}{2} = \pm \sqrt{\frac{2\tau\alpha - 1}{2\tau\alpha + 1}}. \tag{41}$$

Since the fraction takes values between 0 and 1, the same is true for the square root. Therefore, we have $|\mu| < 1$.

Negative Discriminant If the discriminant is negative, the corresponding values of μ are complex and we can write

$$\mu = \frac{\gamma}{2} \pm i \sqrt{-\frac{\gamma^2}{4} - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha}}. \tag{42}$$

Then we get for the squared absolute value of μ

$$|\mu|^2 = \frac{\gamma^2}{4} + \left(-\frac{\gamma^2}{4} - \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha}\right) = \frac{2\tau\alpha - 1}{2\tau\alpha + 1}. \tag{43}$$

Surprisingly, this does not depend on γ and we recover once more the condition

$$-1 < \frac{2\tau\alpha - 1}{2\tau\alpha + 1} < 1. \tag{44}$$

Thus, we again have $|\mu| < 1$.

Positive Discriminant In this case, the eigenvalues μ are real-valued. Thus, we get the condition

$$-1 < \frac{\gamma}{2} \pm \sqrt{\frac{\gamma^2}{4} + \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha}} < 1. \tag{45}$$

If $\tau > 0$ and $\alpha > 0$, this system has the following solutions for γ :

$$\begin{aligned} |\gamma| &< \frac{4\tau\alpha}{1 + 2\tau\alpha}, \quad \text{if } 0 \leq \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} < 1, \\ \sqrt{\frac{8\tau\alpha - 4}{2\tau\alpha + 1}} \leq |\gamma| &< \frac{4\tau\alpha}{1 + 2\tau\alpha}, \quad \text{if } -1 < \frac{1 - 2\tau\alpha}{1 + 2\tau\alpha} < 0. \end{aligned} \tag{46}$$

If we treat γ as a complex number for the moment, the first condition means that γ has to be inside a disc of radius $\frac{4\tau\alpha}{1 + 2\tau\alpha}$ around the origin and the second condition means that γ has to be inside that disc, but outside or on the boundary of a second disc of radius $\sqrt{\frac{8\tau\alpha - 4}{2\tau\alpha + 1}}$. Hence, the second condition is more restrictive.

It remains to determine conditions on τ and α such that (46) is always satisfied for all eigenvalues γ of the matrix \mathbf{Q} . As the matrix $\mathbf{A}(\mathbf{u}^k)$ is real-valued and symmetric, we can diagonalise it in the same fashion as \mathbf{Q} : There exist an orthogonal matrix \mathbf{W} and a diagonal matrix $\mathbf{\Lambda}$ with the eigenvalues λ on its diagonal such that

$$\mathbf{A}(\mathbf{u}^k) = \mathbf{W} \mathbf{\Lambda} \mathbf{W}^T. \tag{47}$$

With the help of this representation, we can write

$$\begin{aligned} \mathbf{Q} &= \frac{4\tau\alpha}{1+2\tau\alpha}\mathbf{I} - \frac{2\tau}{1+2\tau\alpha}\mathbf{A}(\mathbf{u}^k) \\ &= \frac{4\tau\alpha}{1+2\tau\alpha}\mathbf{W}\mathbf{W}^T - \frac{2\tau}{1+2\tau\alpha}\mathbf{W}\mathbf{\Lambda}\mathbf{W}^T \\ &= \mathbf{W}\left(\frac{4\tau\alpha}{1+2\tau\alpha}\mathbf{I} - \frac{2\tau}{1+2\tau\alpha}\mathbf{\Lambda}\right)\mathbf{W}^T. \end{aligned} \quad (48)$$

Hence, the eigenvalues of \mathbf{Q} are given by

$$\gamma = \frac{4\tau\alpha}{1+2\tau\alpha} - \frac{2\tau}{1+2\tau\alpha}\lambda, \quad (49)$$

where λ is an eigenvalue of $\mathbf{A}(\mathbf{u}^k)$.

With this formula, the first case of (46) reduces to

$$0 < \lambda < 4\alpha. \quad (50)$$

This condition is similar to the stability condition of the explicit scheme, however now for α instead of $\frac{1}{\tau}$. The estimate on the left-hand side is always fulfilled if $\mathbf{A}(\mathbf{u}^k)$ is positive definite. For now, we assume that $\mathbf{A}(\mathbf{u}^k)$ is positive definite and consider the case of a zero eigenvalue afterwards.

The inequality (50) has to hold for every eigenvalue of $\mathbf{A}(\mathbf{u}^k)$. If $\mathbf{A}(\mathbf{u}^k)$ is positive definite, we can replace λ by the spectral radius $\rho(\mathbf{A}(\mathbf{u}^k))$ as an upper bound. Hence, we arrive at

$$\alpha > \frac{\rho(\mathbf{A}(\mathbf{u}^k))}{4}. \quad (51)$$

This is in line with the result that has been derived for the linear case [41].

For the second case in (46), plugging in (49) and simplifying yields

$$\begin{aligned} 0 < \lambda \leq 2\alpha - \sqrt{4\alpha^2 - \frac{1}{\tau^2}} \quad \text{or} \\ 2\alpha + \sqrt{4\alpha^2 - \frac{1}{\tau^2}} \leq \lambda < 4\alpha \end{aligned} \quad (52)$$

As we are in the case in which $-1 < \frac{1-2\tau\alpha}{1+2\tau\alpha} < 0$, the square root is a nonnegative real number.

Moreover, we have to investigate what happens if

$$\begin{aligned} -1 \leq \frac{1-2\tau\alpha}{1+2\tau\alpha} < 0 \quad \text{and} \\ 2\alpha - \frac{\sqrt{4\alpha^2\tau^2 - 1}}{\tau^2} < \lambda < 2\alpha + \frac{\sqrt{4\alpha^2\tau^2 - 1}}{\tau^2}. \end{aligned} \quad (53)$$

With some tedious but straight-forward computations, one can show that this case corresponds exactly to the case with

a negative discriminant, which did not introduce any new stability conditions.

Lastly, we consider the case where $\mathbf{A}(\mathbf{u}^k)$ is only positive semidefinite. If $2\tau\alpha < 1$, then $\lambda = 0$ lies in the range given for λ in (53). For this case, we have already shown that $|\mu| < 1$ and the scheme is stable.

If $2\tau\alpha > 1$, we obtain from (49):

$$\gamma = \frac{4\tau\alpha}{1+2\tau\alpha}. \quad (54)$$

We can plug this value of γ into (40) to see that $\mu = 1$ is always a solution in this case. We have to ensure that all other solutions for μ fulfil $|\mu| < 1$. The other solution of (40) for $\lambda = 0$ is

$$\mu = \frac{2\tau\alpha - 1}{2\tau\alpha + 1}, \quad (55)$$

which yields $|\mu| < 1$ since $\tau > 0$ and $\alpha > 0$. All other eigenvalues of \mathbf{B} have an absolute value of less than one by the considerations above. Thus, we can conclude that also in this case, the Du Fort–Frankel scheme is stable.

Consequently, the only stability condition on α is the one in (51). In a similar manner as for the symmetric ResNet, we can transform this condition into an a priori constraint

$$\alpha \geq \frac{L}{4}. \quad (56)$$

References

- Alt, T., Peter, P., Weickert, J., Schrader, K.: Translating numerical concepts for PDEs into neural architectures. In: Elmoataz, A., Fadili, J., Quéau, Y., Rabin, J., Simon, L. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 12679, pp. 294–306. Springer, Cham (2021)
- Alt, T., Weickert, J.: Learning integrodifferential models for denoising. In: *Proc. 2021 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 2045–2049. IEEE Computer Society Press, Toronto, Canada (2021)
- Alt, T., Weickert, J., Peter, P.: Translating diffusion, wavelets, and regularisation into residual networks. [arXiv:2002.02753v3](https://arxiv.org/abs/2002.02753v3) [cs.LG] (2020)
- Andreu, F., Ballester, C., Caselles, V., Mazón, J.M.: Minimizing total variation flow. *Differ. Integral Equations* **14**(3), 321–360 (2001)
- Arridge, S., Hauptmann, A.: Networks for nonlinear diffusion problems in imaging. *J. Math. Imaging Vis.* **62**, 471–487 (2020)
- Aubert, G., Kornprobst, P.: *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*. Applied Mathematical Sciences, vol. 147, 2nd edn. Springer, New York (2006)
- Bäker, M.: Another look at neural multigrid. *Int. J. Mod. Phys. C* **8**(2), 191–205 (1997)
- Bäker, M., Mack, G., Speh, M.: Multigrid meets neural nets. *Nucl. Phys. B Proc. Suppl.* **30**, 269–272 (1993)

9. Bengio, Y., Simard, P., Frasconi, P.: Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Netw.* **5**(2), 157–166 (1994)
10. Benning, M., Celledoni, E., Erhardt, M.J., Owren, B., Schönlieb, C.: Deep learning as optimal control problems: models and numerical methods. *IFAC-PapersOnline* **54**(9), 620–623 (2021)
11. Brandt, A.: Multi-level adaptive solutions to boundary-value problems. *Math. Comput.* **31**(138), 333–390 (1977)
12. Briggs, W.L., Henson, V.E., McCormick, S.F.: *A Multigrid Tutorial*, 2nd edn. SIAM, Philadelphia (2000)
13. Brito-Loeza, C., Chen, K.: Multigrid algorithm for high order denoising. *SIAM J. Imaging Sci.* **3**(3), 363–389 (2010)
14. Bruhn, A., Weickert, J., Kohlberger, T., Schnörr, C.: A multigrid platform for real-time motion computation with discontinuity-preserving variational methods. *Int. J. Comput. Vis.* **70**(3), 257–277 (2006)
15. Bungert, L., Raab, R., Roith, T., Schwinn, L., Tenbrinck, D.: CLIP: Cheap Lipschitz training of neural networks. In: Elmoataz, A., Fadili, J., Quéau, Y., Rabin, J., Simon, L. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 12679, pp. 307–319. Springer, Cham (2021)
16. Chan, T.F., Shen, J.: *Image Processing and Analysis: Variational, PDE, Wavelet, and Stochastic Methods*. SIAM, Philadelphia (2005)
17. Chang, B., Meng, L., Haber, E., Ruthotto, L., Begert, D., Holtham, E.: Reversible architectures for arbitrarily deep residual neural networks. In: Proc. 32nd AAAI Conference on Artificial Intelligence, pp. 2811–2818. New Orleans, LA (2018)
18. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Two deterministic half-quadratic regularization algorithms for computed imaging. In: Proc. 1994 IEEE International Conference on Image Processing, vol. 2, pp. 168–172. IEEE Computer Society Press, Austin, TX (1994)
19. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Proc. 32nd International Conference on Neural Information Processing Systems, vol. 31, pp. 6571–6583. Montréal, Canada (2018)
20. Chen, Y., Pock, T.: Trainable nonlinear reaction diffusion: a flexible framework for fast and effective image restoration. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(6), 1256–1272 (2016)
21. Combettes, P.L., Pesquet, J.: Deep neural network structures solving variational inequalities. *Set-Valued Var. Anal.* **28**(3), 491–518 (2020)
22. Combettes, P.L., Pesquet, J.: Lipschitz certificates for layered network structures driven by averaged activation operators. *SIAM J. Math. Data Sci.* **2**(2), 529–557 (2020)
23. Croce, F., Andriushchenko, M., Hein, M.: Provable robustness of ReLU networks via maximization of linear regions. In: Chaudhuri, K., Sugiyama, M. (eds.) Proc. 22nd International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research, vol. 89, pp. 2057–2066. Okinawa, Japan (2019)
24. Daubechies, I., DeVore, R., Foucart, S., Hanin, B., Petrova, G.: Nonlinear approximation and (deep) ReLU networks. *Constructive Approximation* (2021). Online first
25. De Felice, P., Marangi, C., Nardulli, G., Pasquariello, G., Tedesco, L.: Dynamics of neural networks with non-monotone activation function. *Netw. Comput. Neural Syst.* **4**(1), 1–9 (1993)
26. Didas, S., Weickert, J., Burgeth, B.: Properties of higher order nonlinear diffusion filtering. *J. Math. Imaging Vis.* **35**, 208–226 (2009)
27. Dong, H., Yang, G., Liu, F., Mo, Y., Guo, Y.: Automatic brain tumor detection and segmentation using U-Net based fully convolutional networks. In: Hernández, M.V., González-Castro, V. (eds.) *Medical Image Understanding and Analysis: MIUA 2017, Communications in Computer and Information Science*, vol. 723, pp. 506–517. Springer, Cham (2017)
28. Du Fort, E.C., Frankel, S.P.: Stability conditions in the numerical treatment of parabolic differential equations. *Math. Tables Other Aids Comput.* **7**, 135–152 (1953)
29. Duits, R., Smets, B., Bekkers, E., Portegies, J.: Equivariant deep learning via morphological and linear scale space PDEs on the space of positions and orientations. In: Elmoataz, A., Fadili, J., Quéau, Y., Rabin, J., Simon, L. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 12679, pp. 27–39. Springer, Cham (2021)
30. E, W., Han, J., Jentzen, A.: Algorithms for solving high dimensional PDEs: from nonlinear Monte Carlo to machine learning. [arXiv:2008.13333v2](https://arxiv.org/abs/2008.13333v2) [math.NA] (2020)
31. Eliasof, M., Ephrath, J., Ruthotto, R., Treister, E.: Multigrid-in-channels neural network architectures. [arXiv:2011.09128v2](https://arxiv.org/abs/2011.09128v2) [cs.CV] (2020)
32. Esser, P., Sutter, E., Ommer, B.: A variational U-Net for conditional appearance and shape generation. In: Proc. 2018 IEEE Conference on Computer Vision and Pattern Recognition, pp. 8857–8866. IEEE Computer Society Press, Salt Lake City, UT (2018)
33. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.P.: Image compression with anisotropic diffusion. *J. Math. Imaging Vis.* **31**(2–3), 255–269 (2008)
34. Genzel, M., Macdonald, J., März, M.: Solving inverse problems with deep neural networks—robustness included? [arXiv:2011.04268v1](https://arxiv.org/abs/2011.04268v1) [cs.LG] (2020)
35. Gerschgorin, S.: Fehlerabschätzung für das differenzenverfahren zur Lösung partieller differentialgleichungen. *Z. Angew. Math. Mech.* **10**, 373–382 (1930)
36. Gilboa, G., Zeevi, Y., Sochen, N.: Image enhancement segmentation and denoising by time dependent nonlinear diffusion processes. In: Proc. 2001 IEEE International Conference on Image Processing, vol. 3, pp. 134–137. IEEE Computer Society Press, Thessaloniki, Greece (2001)
37. Golts, A., Freedman, D., Elad, M.: Deep energy: task driven training of deep neural networks. *IEEE J. Sel. Top. Signal Process.* **15**(2), 324–338 (2021)
38. Goodfellow, I., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y.: Maxout networks. In: Dasgupta, S., McAllester, D. (eds.) Proc. 30th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 28, pp. 1319–1327. Atlanta, GA (2013)
39. Goodfellow, I.J., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, Cambridge (2016)
40. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Bengio, Y., LeCun, Y. (eds.) Proc. 3rd International Conference on Learning Representations. San Diego, CA (2015)
41. Gottlieb, D., Gustafsson, B.: Generalized Du Fort-Frankel methods for parabolic initial-boundary value problems. *SIAM J. Numer. Anal.* **13**(1), 129–144 (1975)
42. Gouk, H., Frank, E., Pfahringer, B., Cree, M.J.: Regularisation of neural networks by enforcing Lipschitz continuity. *Mach. Learn.* **110**, 393–416 (2021)
43. Greenfeld, D., Galun, M., Kimmel, R., Yavneh, I., Basri, R.: Learning to optimize multigrid PDE solvers. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proc. 36th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 97, pp. 2415–2423. Long Beach, CA (2019)
44. Gribonval, R., Kutyniok, G., Nielsen, M., Voigtlaender, F.: Approximation spaces of deep neural networks. *Constructive Approximation* (2021). Online first

45. Günther, S., Ruthotto, L., Schroder, J.B., Cyr, E.C., Gauger, N.R.: Layer-parallel training of deep residual neural networks. *SIAM J. Math. Data Sci.* **2**(1), 1–23 (2020)
46. Gusak, J., Katrutsa, A., Daulbaev, T., Cichocki, A., Oseledets, I.: Meta-solver for neural ordinary differential equations. [arXiv:2103.08561v1](https://arxiv.org/abs/2103.08561v1) [cs.LG] (2021)
47. Haber, E., Lensink, K., Treister, E., Ruthotto, L.: IMEXnet a forward stable deep neural network. In: Chaudhuri, K., Salakhutdinov, R. (eds.) *Proc. 36th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 97, pp. 2525–2534. Long Beach, CA (2019)
48. Haber, E., Ruthotto, L.: Stable architectures for deep neural networks. *Inverse Probl.* **34**(1), 014004 (2017)
49. Haber, E., Ruthotto, L., Holtham, E., Jun, S.H.: Learning across scales—multiscale methods for convolution neural networks. In: *Proc. 32nd AAAI Conference on Artificial Intelligence*, pp. 2811–2818. New Orleans, LA (2018)
50. Hackbusch, W.: *Multigrid Methods and Applications*. Springer, New York (1985)
51. Hafner, D., Ochs, P., Weickert, J., Reißel, M., Grewenig, S.: FSI schemes: fast semi-iterative solvers for PDEs and optimisation methods. In: Rosenhahn, B., Andres, B. (eds.) *Pattern Recognition. Lecture Notes in Computer Science*, vol. 9796, pp. 91–102. Springer, Cham (2016)
52. Hartmann, D., Lessig, C., Margenberg, N., Richter, T.: A neural network multigrid solver for the Navier-Stokes equations. [arXiv:2008.11520v1](https://arxiv.org/abs/2008.11520v1) [physics.comp-ph] (2020)
53. He, J., Xu, J.: MgNet: a unified framework of multigrid and convolutional neural network. *Sci. China Math.* **62**, 1331–1354 (2019)
54. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778. IEEE Computer Society Press, Las Vegas, NV (2016)
55. Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.* **79**(8), 2554–2558 (1982)
56. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proc. 2017 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4700–4708. IEEE Computer Society Press, Honolulu, HI (2017)
57. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotech. Lab.* **26**, 368–388 (1962). (In Japanese)
58. Katrutsa, A., Daulbaev, T., Oseledets, I.: Black-box learning of multigrid parameters. *J. Comput. Appl. Math.* **368**, 112524 (2020)
59. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. [arXiv:1412.6980v1](https://arxiv.org/abs/1412.6980v1) [cs.LG] (2014)
60. Kobler, E., Efland, A., Kunisch, K., Pock, T.: Total deep variation for linear inverse problems. In: *Proc. 2020 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 7549–7558. IEEE Computer Society Press, Seattle, WA (2020)
61. Kobler, E., Klatzer, T., Hammernik, K., Pock, T.: Variational networks: connecting variational methods and deep learning. In: Roth, V., Vetter, T. (eds.) *Pattern Recognition. Lecture Notes in Computer Science*, vol. 10496, pp. 281–293. Springer, Cham (2017)
62. Köstler, H., Stürmer, M., Freundl, C., Rüdiger, U.: PDE based video compression in real time. *Tech. Rep. 07-11*, Lehrstuhl für Informatik 10, Univ. Erlangen–Nürnberg, Germany (2007)
63. Kutyniok, G., Petersen, P., Raslan, M., Schneider, R.: A theoretical analysis of deep neural networks and parametric PDEs. *Constructive Approximation* (2021). Online first
64. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–444 (2015)
65. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
66. Leino, K., Wang, Z., Fredrikson, M.: Globally-robust neural networks. [arXiv:2102.08452v1](https://arxiv.org/abs/2102.08452v1) [cs.LG] (2021)
67. Li, M., He, L., Lin, Z.: Implicit Euler skip connections: Enhancing adversarial robustness via numerical stability. In: Daumé, H., III., Singh, A. (eds.) *Proc. 37th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 119, pp. 5874–5883. Austria, Vienna (2020)
68. Long, Z., Lu, Y., Dong, B.: PDE-Net 2.0: learning PDEs from data with a numeric-symbolic hybrid deep network. *J. Comput. Phys.* **399**(2197), 108925 (2019)
69. Lu, Y., Zhong, A., Li, Q., Dong, B.: Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In: Dy, J., Krause, A. (eds.) *Proc. 35th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 80, pp. 3276–3285. Stockholm, Sweden (2018)
70. 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.M., ter Haar Romeny, B., Bronstein, A.M., Bronstein, M.M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 6667, pp. 26–37. Springer, Berlin (2012)
71. Meilijson, I., Ruppig, E.: Optimal signalling in attractor neural networks. In: Tesauro, G., Touretzky, D., Leen, T. (eds.) *Proc. 7th International Conference on Neural Information Processing Systems, Advances in Neural Information Processing Systems*, vol. 7, pp. 485–492. Denver, CO (1994)
72. Mish, D.: Mish: A self regularized non-monotonic activation function. [arXiv:1908.08681v3](https://arxiv.org/abs/1908.08681v3) [cs.LG] (2020)
73. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: *Proc. 27th International Conference on Machine Learning*, pp. 807–814. Haifa, Israel (2010)
74. Nesterov, Y.: A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Sov. Math. Dokl.* **4**, 1035–1038 (1963)
75. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *Computer Vision—ECCV 2016. Lecture Notes in Computer Science*, vol. 9912, pp. 483–499. Springer, Cham (2016)
76. Ochs, P., Meinhardt, T., Leal-Taixe, L., Möller, M.: Lifting layers: analysis and applications. In: Ferrari, V., Herbert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision—ECCV 2018. Lecture Notes in Computer Science*, vol. 11205, pp. 53–68. Springer, Cham (2018)
77. Ott, K., Katiyar, P., Hennig, P., Tiemann, M.: ResNet after all? Neural ODEs and their numerical solution. In: *Proc. 9th International Conference on Learning Representations*. Vienna, Austria (2021)
78. Ouala, S., Pascual, A., Fablet, R.: Residual integration neural network. In: *Proc. 2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3622–3626. IEEE Computer Society Press, Brighton, UK (2019)
79. Parhi, R., Nowak, R.D.: What kinds of functions do deep neural networks learn? Insights from variational spline theory. [arXiv:2105.03361v1](https://arxiv.org/abs/2105.03361v1) [stat.ML] (2021)
80. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 629–639 (1990)
81. Polyak, B.T.: Some methods of speeding up the convergence of iteration methods. *USSR Comput. Math. Math. Phys.* **4**(5), 1–17 (1964)

82. Rackauckas, C., Ma, Y., Martensen, J., Warnter, C., Zubov, K., Supekar, R., Skinner, D., Ramadhan, A., Edelman, A.: Universal differential equations for scientific machine learning. [arXiv:2001.04385v3](https://arxiv.org/abs/2001.04385v3) [cs.LG] (2020)
83. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019)
84. Ramachandran, P., Zoph, B., Le, Q.V.: Searching for activation functions. [arXiv:1710.05941v2](https://arxiv.org/abs/1710.05941v2) [cs.NE] (2017)
85. Ratner, V., Zeevi, Y.Y.: The dynamics of image processing viewed as damped elastic deformation. In: Proc. 17th European Signal Processing Conference, pp. 45–49. IEEE Computer Society Press, Glasgow, UK (2009)
86. Rolnick, D., Tegmark, M.: The power of deeper networks for expressing natural functions. In: Proc. 6th International Conference on Learning Representations. Vancouver, Canada (2018)
87. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds.) Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015. Lecture Notes in Computer Science, vol. 9351, pp. 234–241. Springer, Cham (2015)
88. Rousseau, F., Drumetz, L., Fablet, R.: Residual networks as flows of diffeomorphisms. *J. Math. Imaging Vis.* **62**, 365–375 (2020)
89. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D* **60**(1–4), 259–268 (1992)
90. Rudy, S.H., Brunton, S.L., Proctor, J.L., Kutz, J.N.: Data-driven discovery of partial differential equations. *Sci. Adv.* **3**(4), e1602614 (2017)
91. Rumelhart, D.E., McClelland, J.L.: Parallel Distributed Processing: Explorations in the Microstructure of Cognition. MIT Press, Cambridge, MA (1986)
92. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**, 211–252 (2015)
93. Ruthotto, L., Haber, E.: Deep neural networks motivated by partial differential equations. *J. Math. Imaging Vis.* **62**, 352–364 (2020)
94. Saad, Y.: Iterative Methods for Sparse Linear Systems, 2nd edn. SIAM, Philadelphia (2003)
95. Salimans, T., Kingma, D.P.: Weight normalization: a simple reparameterization to accelerate training of deep neural networks. [arXiv:1602.07868v3](https://arxiv.org/abs/1602.07868v3) [cs.LG] (2016)
96. Schaeffer, H.: Learning partial differential equations via data discovery and sparse optimization. *Proc. R. Soc. Lond. Ser. A* **473**(2197), 20160446 (2017)
97. Scherzer, O., Weickert, J.: Relations between regularization and diffusion filtering. *J. Math. Imaging Vis.* **12**(1), 43–63 (2000)
98. Schmaltz, C., Peter, P., Mainberger, M., Ebel, F., Weickert, J., Bruhn, A.: Understanding, optimising, and extending data compression with anisotropic diffusion. *Int. J. Comput. Vis.* **108**(3), 222–240 (2014)
99. Schmidhuber, J.: Deep learning in neural networks: an overview. *Neural Netw.* **61**, 85–117 (2015)
100. Sutskever, I., Martens, J., Dahl, G., Hinton, G.: On the importance of initialization and momentum in deep learning. In: Dasgupta, S., McAllester, D. (eds.) Proc. 30th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 28, pp. 1139–1147. Atlanta, GA (2013)
101. Thorpe, M., van Gennip, Y.: Deep limits of residual neural networks. [arXiv:1810.11741v2](https://arxiv.org/abs/1810.11741v2) [math.CA] (2019)
102. Tretter, C.: Spectral Theory of Block Operator Matrices and Applications. Imperial College Press, London (2008)
103. van Der Houwen, P.J., Sommeijer, B.P.: On the internal stability of explicit, m-stage Runge-Kutta methods for large m-values. *Z. Angew. Math. Mech.* **60**(10), 479–485 (1980)
104. Weickert, J.: Theoretical foundations of anisotropic diffusion in image processing. *Comput. Suppl.* **11**, 221–236 (1996)
105. Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner, Stuttgart (1998)
106. Weickert, J., Benhamouda, B.: A semidiscrete nonlinear scale-space theory and its relation to the Perona-Malik paradox. In: Solina, F., Kropatsch, W.G., Klette, R., Bajcsy, R. (eds.) Advances in Computer Vision, pp. 1–10. Springer, Wien (1997)
107. Weickert, J., Welk, M.: Tensor field interpolation with PDEs. In: Weickert, J., Hagen, H. (eds.) Visualization and Processing of Tensor Fields, pp. 315–325. Springer, Berlin (2006)
108. Weickert, J., Welk, M., Wickert, M.: L^2 -stable nonstandard finite differences for anisotropic diffusion. In: Kuijper, A., Bredies, K., Pock, T., Bischof, H. (eds.) Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, vol. 7893, pp. 390–391. Springer, Berlin (2013)
109. You, Y.L., Kaveh, M.: Fourth-order partial differential equations for noise removal. *IEEE Trans. Image Process.* **9**(10), 1723–1730 (2000)
110. Zhang, L., Schaeffer, H.: Forward stability of ResNet and its variants. *J. Math. Imaging Vis.* **62**, 328–351 (2020)
111. Zhu, M., Chang, B., Fu, C.: Convolutional neural networks combined with Runge-Kutta methods. In: Proc. 7th International Conference on Learning Representations. New Orleans, LA (2019)
112. Zhu, M., Min, W., Wang, Q., Zou, S., Chen, X.: PFLU and FPFLU: two novel non-monotonic activation functions in convolutional neural networks. *Neurocomputing* **429**, 110–117 (2021)
113. Zou, D., Balan, R., Singh, M.: On Lipschitz bounds of general convolutional neural networks. *IEEE Trans. Inf. Theory* **66**(3), 1738–1759 (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Tobias Alt is a Ph.D. student with the Mathematical Image Analysis group at Saarland University (Saarbrücken, Germany), where he also received his B.Sc. degree in computer science in 2016. His research interests include model-based learning approaches, the mathematical foundations of neural networks, and their applications to inpainting-based image compression.



Karl Schrader is a Ph.D. student with the Mathematical Image Analysis Group at Saarland University (Saarbrücken, Germany). He received two B.Sc. degrees in the fields of mathematics and computer science from Humboldt-University of Berlin in 2018 and his M.Sc. degree in computer science from Saarland University in 2020. His research interests include the mathematical foundations of deep learning and connections between partial differential equations and neural networks.



Matthias Augustin received his diploma degrees in Physics and Mathematics from Saarland University (Saarbrücken, Germany) in 2009 and 2010, respectively. He obtained his Ph.D. in Geomathematics from the University of Kaiserslautern (Germany) in 2014 under the supervision of Willi Freeden. From 2015 to 2021, he worked as a post-doctoral researcher with the Mathematical Image Analysis Group at Saarland University. His research interests include fluid dynamics and fluid–

solid interaction in porous media, sparse inpainting methods, radial basis functions, modelling with partial differential equations, and incorporations of such models in machine learning strategies.



Pascal Peter is a member of the Mathematical Image Analysis Group and a tenured lecturer at Saarland University (Saarbrücken, Germany). At the same institution, he graduated in computer science, mathematics, and pedagogy for high school teachers in 2011 and became a post-doctoral researcher after obtaining his Ph.D. degree in computer science in 2016. His research activities focus on the intersection of computer science and mathematics. His current research interests

include image inpainting, compression, and the mathematical foundations of deep learning.



Joachim Weickert is professor of Mathematics and Computer Science at Saarland University (Saarbrücken, Germany), where he heads the Mathematical Image Analysis Group. He graduated and obtained his Ph.D. from the University of Kaiserslautern (Germany) in 1991 and 1996. He worked as post-doctoral researcher at the University Hospital of Utrecht (The Netherlands) and the University of Copenhagen (Denmark), and as assistant professor at the University of

Mannheim (Germany). Joachim Weickert has developed many models and efficient algorithms for image processing and computer vision using partial differential equations and variational methods. His scientific work covers more than 330 refereed publications. He has served in the editorial boards of ten international journals or book series and is Editor-in-Chief of the Journal of Mathematical Imaging and Vision. In 2010, he has received a Gottfried Wilhelm Leibniz Prize and in 2017 an ERC Advanced Grant.