

LEARNING FILTERED DISCRETIZATION OPERATORS: NON-INTRUSIVE VERSUS INTRUSIVE APPROACHES

SYVER DØVING AGDESTEN^{1*} AND BENJAMIN SANDERSE^{2*}

¹ e-mail: syver.agdestein@cwil.nl, <https://www.cwi.nl/people/syver-agdestein>

² e-mail: b.sanderse@cwil.nl, <https://www.cwi.nl/~sanderse>

* Centrum Wiskunde & Informatica (CWI)
Science Park 123, 1098XK Amsterdam, Netherlands

Key words: Closure modelling, Filtering, Operator inference, Embedded model learning, Data-driven discretization, Quadrature

Abstract. Simulating multi-scale phenomena such as turbulent fluid flows is typically computationally very expensive. Filtering the smaller scales allows for using coarse discretizations, however, this requires closure models to account for the effects of the unresolved on the resolved scales. The common approach is to filter the continuous equations, but this gives rise to several commutator errors due to nonlinear terms, non-uniform filters, or boundary conditions.

We propose a new approach to filtering, where the equations are discretized first and then filtered. For a non-uniform filter applied to the linear convection equation, we show that the discretely filtered convection operator can be inferred using three methods: intrusive ('explicit reconstruction') or non-intrusive operator inference, either via 'derivative fitting' or 'trajectory fitting' (embedded learning). We show that explicit reconstruction and derivative fitting identify a similar operator and produce small errors, but that trajectory fitting requires significant effort to train to achieve similar performance. However, the explicit reconstruction approach is more prone to instabilities.

1 INTRODUCTION

Turbulent fluid flows are characterized by the presence of high frequencies. Direct numerical simulation (DNS) aims to resolve all turbulent frequencies at the expense of a high computational cost. In contrast, approaches such as large eddy simulation (LES) consist of modelling the behavior of the large features only, i.e. the low frequency components. Mathematically, this implies the use of low-pass filters to obtain equations for the low frequency components. These filters do typically not commute with the differential operators, nonlinear terms, or boundary conditions [18]. Furthermore, requiring that they do commute imposes severe restrictions on the choice of filters, e.g. by requiring vanishing moments or spatial uniformity [18, 19, 12]. This often conflicts with other restrictions, such as requiring the filter to be linked to or directly deduced from the spatial discretization (*implicit* LES) [3, 4].

Traditionally, the equations are filtered on the continuous level, and then discretized. The commutator errors arising from the filter are then modelled in terms of the filtered solution only, through a closure model [15]. Recent efforts have been made to learn the closure model *for a given filter and discretization* using deep neural networks [1, 2], thus closing the equations after discretization.

In this article, we continue on this discrete approach, in which we aim to find discrete representations of filtered operators, constructed by using high-fidelity simulations. We compare three different approaches: (i) explicit construction of the filtered operator using quadrature and reconstruction operators; (ii) least-squares reconstruction of the filtered operator, analogous to operator inference; (iii) reconstruction of the embedded filtered operator, analogous to “solver in the loop” methods [10, 11, 17].

We focus specifically on commutator errors arising from non-uniform filtering of linear operators (in which case filtering and differentiation do not commute), so that the filtered operators that need to be constructed remain linear, and we do not need to use complex function approximators such as neural networks. This allows us to get a clean picture of how these three data-driven methods perform in constructing a filtered operator. In Section 2, we present the discrete problem. In Section 3, we present the three approaches for building the discrete filtered operator. In Section 4, we show the results for the filtered convection equation, comparing the three approaches for two different filters.

2 PROBLEM STATEMENT

In this paper we consider the one-dimensional linear partial differential equation

$$\frac{\partial u}{\partial t}(x, t) = \mathcal{A}u(x, t), \quad x \in \Omega, \quad t > 0, \quad (1)$$

where $\Omega = [0, 1]$ is a periodic domain, \mathcal{A} is a linear differential operator with respect to the spatial variable x , and $u(x, 0) = u_0(x)$ are the initial conditions.

A uniform periodic discretization of Ω is given by $\mathbf{x}^{(N)} = (x_n^{(N)})_{n=1}^N \in \mathbb{R}^N$, where N is the number of grid points, $x_n^{(N)} = n\Delta x^{(N)}$, and $\Delta x^{(N)} = \frac{1}{N}$. A semi-discrete approximation to (1) is then given by

$$\frac{d\mathbf{u}}{dt}(t) = \mathbf{A}\mathbf{u}(t), \quad t > 0, \quad (2)$$

where $\mathbf{u}(t) \in \mathbb{R}^N$ is a semi-discrete solution and $\mathbf{A} \in \mathbb{R}^{N \times N}$ is a discrete approximation of the differential operator \mathcal{A} , taking into account the periodicity. The exact semi-discrete solution is given by $\mathbf{u}(t) = e^{t\mathbf{A}}\mathbf{u}_0$.

We are interested in the case where the spectral content of u is too large to be represented by the discrete solution. A low-pass filter is thus applied to obtain a new set of equations that can be resolved with a reasonably small N . We here consider a general linear continuous kernel filter \mathcal{F} defined by

$$\bar{u}(x) = \mathcal{F}(u)(x) = \int_{\mathbb{R}} G(x, \xi)u(\xi) d\xi \quad \forall x \in \Omega, \quad (3)$$

where \bar{u} is the filtered solution and $G : \Omega \times \mathbb{R} \rightarrow \mathbb{R}$ is the kernel. Note that the function u is extended by periodicity to \mathbb{R} , i.e. $u(x) = u(x \bmod |\Omega|)$ for $x \notin \Omega$, thus allowing for infinitely long filter kernels such as a Gaussian kernel. The filter is non-uniform in space, so it is in general not of convolution type (in which case one would have $G(x, \xi) = G(\xi - x) \forall (x, \xi)$).

In many closure modelling approaches, the equations are filtered on the continuous level:

$$\frac{\partial \bar{u}}{\partial t} = \overline{(\mathcal{A}u)} = \mathcal{A}\bar{u} + \mathcal{C}(u, \bar{u}), \quad (4)$$

where $\mathcal{C}(u, \bar{u}) = \overline{\mathcal{A}u} - \mathcal{A}\bar{u}$ is the commutator error, typically caused by nonlinear term. In the current setting where \mathcal{A} is linear, a commutator error still arises due to the fact that non-uniform filtering and differentiation do not commute. The above equations are *unclosed* in the sense that \mathcal{C} depends on u , and typically has to be modelled as a function of \bar{u} only, i.e. $\mathcal{C}(u, \bar{u}) \approx \mathcal{M}(\bar{u})$, where \mathcal{M} is a *closure model*, for example obtained by truncating the continuous Taylor series expansion of u in the integral (3) [7]. The filtered equations are then discretized on a coarse grid $\mathbf{x}^{(M)}$, $M \leq N$:

$$\frac{d\bar{\mathbf{u}}}{dt}(t) = \mathbf{A}^{(M)}\bar{\mathbf{u}}(t) + \mathbf{M}(\bar{\mathbf{u}}(t)), \quad t > 0, \quad (5)$$

where $\bar{\mathbf{u}}(t) \in \mathbb{R}^M$ is an approximation to $\bar{u}(\mathbf{x}^{(M)}, t)$, and $\mathbf{A}^{(M)} \in \mathbb{R}^{M \times M}$ and $\mathbf{M} : \mathbb{R}^M \rightarrow \mathbb{R}^M$ are coarse-grid approximations of the operators \mathcal{A} and \mathcal{M} respectively. In this approach one thus first commits a continuous closure model error, and then a (coarse scale) discretization error.

We propose to obtain the discrete equations for the filtered solution in a different way, by *discretizing first*, and *then filtering*. We employ a discrete approximation of the filter, represented by a matrix \mathbf{W} built using a *quadrature rule*:

$$\bar{\mathbf{u}} = \mathbf{W}\mathbf{u}. \quad (6)$$

Note that the filtered and unfiltered discrete solutions $\bar{\mathbf{u}}$ and \mathbf{u} are represented on different grids $\mathbf{x}^{(M)}$ and $\mathbf{x}^{(N)}$, with $\bar{\mathbf{u}}(t) \in \mathbb{R}^M$ and $\mathbf{u}(t) \in \mathbb{R}^N$, $M \leq N$. An accurate discrete filter should be such that $\sum_n W_{mn} u(x_n^{(N)}, t) \approx \bar{u}(x_m^{(M)}, t)$, $m \in \{1, \dots, M\}$. It is assumed that all the relevant frequencies of u are resolved on $\mathbf{x}^{(N)}$, but not necessarily on $\mathbf{x}^{(M)}$. To close the filtered equations, the ‘discretize first’ approach requires an approximation of the inverse filtering operator by using a reconstruction rule, represented by a matrix $\mathbf{R} \in \mathbb{R}^{N \times M}$, such that $\mathbf{u} \approx \mathbf{R}\bar{\mathbf{u}}$. The discretely filtered solution $\bar{\mathbf{u}}$ is then obtained by solving the equations

$$\frac{d\bar{\mathbf{u}}}{dt}(t) = \bar{\mathbf{A}}\bar{\mathbf{u}}(t), \quad t > 0, \quad (7)$$

where $\bar{\mathbf{A}} = \mathbf{W}\mathbf{A}\mathbf{R}$ is the discretely filtered differential operator and the filtered initial conditions are given by $\bar{\mathbf{u}}(0) = \mathbf{W}\mathbf{u}_0$. Since the unfiltered solution is reconstructed directly, Equation (7) is closed. In this approach one thus commits a (fine-grid) discretization error, a quadrature error, and a reconstruction error. The choice of $\bar{\mathbf{A}}$ may be seen as an effective closure model, and the resulting coarse-grid commutator error $\bar{\mathbf{A}} - \mathbf{A}^{(M)}$ can be computed, although it is not explicitly modelled (nor is it needed).

The problem that we address in this article is now the following: given a fine-scale discretization \mathbf{A} and a quadrature rule \mathbf{W} , can we infer $\bar{\mathbf{A}}$? To answer this question, we consider two different approaches, that are outlined in the next section.

3 DATA-DRIVEN OPERATOR INFERENCE

The two data-driven approaches to infer the matrix $\bar{\mathbf{A}}$ considered in this article are:

- *Indirectly* building $\bar{\mathbf{A}}$ by first explicitly building \mathbf{W} , \mathbf{A} , and \mathbf{R} and then computing $\bar{\mathbf{A}} = \mathbf{W}\mathbf{A}\mathbf{R}$, see section 3.1. This approach is *intrusive* in the sense that $\mathbf{A} \in \mathbb{R}^{N \times N}$ is required.
- *Directly* building the reduced order operator $\bar{\mathbf{A}} \in \mathbb{R}^{M \times M}$ (see section 3.2), using gradient-free derivative fitting (3.2.1) and gradient-based embedded (3.2.2) approaches. This approach is *non-intrusive* since access to \mathbf{A} is not required.

In the following, we will employ the snapshot matrices $\mathbf{U} \in \mathbb{R}^{N \times d}$, $\dot{\mathbf{U}} \in \mathbb{R}^{N \times d}$, $\bar{\mathbf{U}} \in \mathbb{R}^{M \times d}$, and $\dot{\bar{\mathbf{U}}} \in \mathbb{R}^{M \times d}$, containing unfiltered and filtered discrete solutions and their time derivatives at different time instances. Here, $d = n_{\text{IC}} n_t$ is the number of snapshots, n_{IC} is the number of initial conditions, and n_t is the number of time instances at which the solutions are evaluated.

3.1 Indirect approach: building the reconstruction matrix \mathbf{R}

Using the midpoint rule [8], the quadrature weights \mathbf{W} used for approximating the weighted integral (3) are constructed as follows:

$$W_{mn} = \frac{\tilde{G}(x_m, \xi_n)}{\sum_{i=1}^N \tilde{G}(x_m, \xi_i)}, \quad (8)$$

where $\tilde{G}(x, \xi) = \sum_{z \in \mathbb{Z}} G(x, \xi + z|\Omega|)$ accounts for the periodicity of Ω and potentially infinitely long kernel G . The sum is truncated to $z \in \{-1, 0, 1\}$ in the case of a sufficiently local filter support (e.g. for the Gaussian kernel). The normalization factor ensures that constant functions are preserved upon filtering. Note that higher order quadrature rules may be built by requiring that certain classes of functions are filtered exactly, but the simple rule (8) is sufficient for our purposes (since N is assumed to be large) and has the advantage that it leads to positive weights if the kernel is positive.

Finding the reconstruction matrix \mathbf{R} is not straightforward, as the matrix \mathbf{W} is in general not invertible ($M \neq N$). \mathbf{R} is therefore built by minimizing the expectation of the reconstruction error for a certain class of functions \mathcal{U} :

$$\min_{\mathbf{R} \in \mathbb{R}^{N \times M}} \mathbb{E}_{u \sim \mathcal{U}} L^{\text{fit}}(\mathbf{R}, \bar{u}(\mathbf{x}^{(M)}), u(\mathbf{x}^{(N)})) + \lambda L^{\text{prior}}(\mathbf{R}), \quad (9)$$

where L^{fit} measures the reconstruction error, L^{prior} penalizes deviation from prior assumptions on the form of \mathbf{R} , and λ is a regularization parameter. Setting the accuracy metric as the mean squared error

$$L^{\text{fit}}(\mathbf{R}, \bar{\mathbf{u}}, \mathbf{u}) = \|\mathbf{R}\bar{\mathbf{u}} - \mathbf{u}\|^2 \quad (10)$$

leads to a class of least squares problems. We consider the case where $L^{\text{prior}}(\mathbf{R}) = \|\mathbf{R}\|_F^2 = \sum_{nm} R_{nm}^2$, and after approximating the expectation value by a mean over the d snapshots \mathbf{U} and $\bar{\mathbf{U}}$, we get

$$\mathbf{R} = \operatorname{argmin}_{\mathbf{R} \in \mathbb{R}^{N \times M}} \frac{1}{d} \|\mathbf{R}\bar{\mathbf{U}} - \mathbf{U}\|_F^2 + \lambda \|\mathbf{R}\|_F^2 = \mathbf{U}\bar{\mathbf{U}}^T (\bar{\mathbf{U}}\bar{\mathbf{U}}^T + d\lambda \mathbf{I})^{-1}. \quad (11)$$

The resulting operator $\bar{\mathbf{A}}^{\text{int}} = \mathbf{W}\mathbf{A}\mathbf{R}$ is labelled as *intrusive* since it requires access to the full order model operator \mathbf{A} .

3.2 Direct approach: building the filtered operator $\bar{\mathbf{A}}$

The operator $\bar{\mathbf{A}}$ can also be constructed non-intrusively from filtered data. The solver used for (2) is then considered to be a “black box”, which returns data samples of the solution and its time derivative. Filtered samples are obtained by applying the discrete filter \mathbf{W} . The operator $\bar{\mathbf{A}}$ then follows from a minimization problem of the form

$$\min_{\bar{\mathbf{A}} \in \mathbb{R}^{M \times M}} \mathbb{E}_{u \sim \mathcal{U}} \mathbb{E}_{t \sim \mathcal{T}} L^{\text{fit}}(\bar{\mathbf{A}}, \bar{u}(\mathbf{x}^{(M)}, \cdot), t) + \lambda_{\text{prior}} L^{\text{prior}}(\bar{\mathbf{A}}) + \lambda_{\text{stab}} L^{\text{stab}}(\bar{\mathbf{A}}), \quad (12)$$

where \mathcal{U} is the space of solutions to (1), \mathcal{T} is a distribution of expected time instances, for example a uniform distribution on $[0, T]$ for some time horizon T , L^{prior} penalizes deviation from prior expected properties, and L^{stab} penalizes numerical instabilities that arise from the non-uniform filter (namely anti-diffusion caused by the non-uniformity of the filter). In the following, we set $L^{\text{stab}}(\bar{\mathbf{A}}) = \|\bar{\mathbf{A}} - \mathbf{D}^{(M)}\|_F^2$, where $\mathbf{D}^{(M)}$ is a discretization of the diffusion operator $\frac{\partial^2}{\partial x^2}$ on $\mathbf{x}^{(M)}$. This limits anti-diffusion due to eigenvalues with positive real part. In addition, we set $L^{\text{prior}}(\bar{\mathbf{A}}) = \|\bar{\mathbf{A}} - \mathbf{A}^{(M)}\|_F^2$, motivated by the assumption that the behavior of the filtered solution should be similar to the one of the unfiltered solution (as is the reason for doing LES in the first place), resulting in a commutator error that is small compared to the exact right-hand size, i.e. $\mathcal{O}\left(\frac{\|\bar{\mathbf{A}} - \mathbf{A}^{(M)}\|_F}{\|\mathbf{A}^{(M)}\|_F}\right) = \mathcal{O}\left(\frac{\|\mathcal{C}\|}{\|\mathcal{A}\|}\right) \ll 1$.

L^{fit} is an instantaneous metric evaluating the performance of the operator $\bar{\mathbf{A}}$ with respect to a discrete filtered solution $\bar{\mathbf{u}}$, for which two different choices will be proposed, as explained in the next two sections.

3.2.1 Gradient-free operator inference (derivative fitting)

The first choice for L^{fit} is similar to the operator inference framework presented in [13], where the authors propose to fit the operators describing an ODE to data samples of the time derivative, by using performance metrics of the form

$$L^{\text{fit}}(\bar{\mathbf{A}}, \bar{\mathbf{u}}, t) = \left\| \bar{\mathbf{A}}\bar{\mathbf{u}} - \frac{d\bar{\mathbf{u}}}{dt} \right\|^2(t). \quad (13)$$

We call this approach ‘derivative fitting’. Note that with this performance metric, neither the filtered solution $\bar{\mathbf{u}}$ nor its time derivative $\frac{d\bar{\mathbf{u}}}{dt}$ depend on $\bar{\mathbf{A}}$, and that the dependency on $\bar{\mathbf{A}}$ is quadratic in the loss function. As in Equation (11), approximating the expectation value by a sum over the d snapshots results in the following closed form expression for the operator:

$$\begin{aligned} \bar{\mathbf{A}}^{\text{DF}} &= \underset{\bar{\mathbf{A}} \in \mathbb{R}^{M \times M}}{\text{argmin}} \frac{1}{d} \|\bar{\mathbf{A}}\bar{\mathbf{U}} - \dot{\bar{\mathbf{U}}}\|_F^2 + \lambda_{\text{prior}} \|\bar{\mathbf{A}} - \mathbf{A}^{(M)}\|_F^2 + \lambda_{\text{stab}} \|\bar{\mathbf{A}} - \mathbf{D}^{(M)}\|_F^2 \\ &= \left(\frac{1}{d} \dot{\bar{\mathbf{U}}}\bar{\mathbf{U}}^\top + \lambda_{\text{prior}} \mathbf{A}^{(M)} + \lambda_{\text{stab}} \mathbf{D}^{(M)} \right) \left(\frac{1}{d} \bar{\mathbf{U}}\bar{\mathbf{U}}^\top + (\lambda_{\text{prior}} + \lambda_{\text{stab}}) \mathbf{I} \right)^{-1}, \end{aligned} \quad (14)$$

where DF denotes derivative fitting.

3.2.2 Embedded operator inference (solver-in-the-loop)

The second choice for L^{fit} is the instantaneous performance metric

$$L^{\text{fit}}(\bar{\mathbf{A}}, \bar{\mathbf{u}}, t) = \|\mathbf{S}(\bar{\mathbf{A}}, \bar{\mathbf{u}}(0), t) - \bar{\mathbf{u}}(t)\|^2. \quad (15)$$

In this expression we introduced the notion of the ODE solver $\mathbf{S} : \mathbb{R}^{M \times M} \times \mathbb{R}^M \times \mathbb{R} \rightarrow \mathbb{R}^M$, where $\mathbf{S}(\bar{\mathbf{A}}, \bar{\mathbf{u}}_0, t)$ approximates the solution to (7) for the given operator $\bar{\mathbf{A}}$ and filtered initial conditions $\bar{\mathbf{u}}_0$. In the previous section, the operator $\bar{\mathbf{A}}^{\text{DF}}$ was obtained by fitting known filtered solution trajectories $\bar{\mathbf{u}}$ to their time derivatives $\frac{d\bar{\mathbf{u}}}{dt}$ using least squares. In contrast, in this section the operator $\bar{\mathbf{A}}$ is inferred

by fitting the predicted solution trajectories $\mathbf{S}(\bar{\mathbf{A}}, \bar{\mathbf{u}}(0), \cdot)$ to the known filtered solution trajectories $\bar{\mathbf{u}}$. The operator $\bar{\mathbf{A}}$ is thus *embedded* in the solver \mathbf{S} [10, 11, 17].

Assuming the program \mathbf{S} is differentiable, i.e. we have access to the quantity $\frac{\partial \mathbf{S}}{\partial \bar{\mathbf{A}}}(\bar{\mathbf{A}}, \bar{\mathbf{u}}_0, t) \in \mathbb{R}^{M \times M \times M}$, we will use a gradient-descent based optimization algorithm to identify the correct $\bar{\mathbf{A}}$, using adjoint-mode differentiation with checkpointing to compute the vector-Jacobian products [6, 14].

Note that each evaluation of the loss function requires solving the ODE (7) of size M for the given operator $\bar{\mathbf{A}}$ (solver-in-the-loop). For more computationally expensive problems, e.g. three-dimensional non-linear equations such as the Navier-Stokes equations, evaluating L^{fit} at all initial conditions and time steps in the training dataset can quickly get expensive. We will use *stochastic* gradient descent [5] to reduce this computational cost, by evaluating the loss and its gradient for only a few initial conditions and time instances at a time. In the following, the operator fitted while embedded in the solver will be denoted by $\bar{\mathbf{A}}^{\text{emb}}$.

4 NUMERICAL RESULTS: CONVECTION EQUATION

We consider the convection equation with unit velocity, obtained by setting $\mathcal{A} = -\frac{\partial}{\partial x}$ in (1). To test the effectiveness of the three approaches to infer the filtered operator $\bar{\mathbf{A}}$, we consider two filters: the top-hat filter, for which

$$G(x, \xi) = \begin{cases} \frac{1}{2h(x)}, & |\xi - x| \leq h(x), \\ 0, & \text{otherwise,} \end{cases} \quad (16)$$

and the Gaussian filter, for which

$$G(x, \xi) = \sqrt{\frac{3}{2\pi h^2(x)}} e^{-\frac{3(\xi-x)^2}{2h^2(x)}}, \quad (17)$$

where $h(x)$ is a variable filter radius. For both filters, it is chosen to be

$$h(x) = \left(1 + \frac{1}{3} \sin(2\pi x)\right) h_0 \quad (18)$$

with $h_0 = \frac{1}{50}$. At $x = \frac{1}{4}$, the filter radius is thus twice as large as at $x = \frac{3}{4}$, and filtered waves moving from $x = \frac{1}{4}$ to $x = \frac{3}{4}$ are subject to *anti-diffusion*, where high-frequency components grow. Small numerical errors get amplified on this interval, and finding a stable $\bar{\mathbf{A}}$ is thus not straightforward. Note also that the Gaussian standard deviation $\frac{h}{\sqrt{3}}$ is chosen to give the same Taylor series expansion for the local attenuation factor as a function of kh in spectral space as for the top-hat filter, where k is the frequency.

Given initial conditions, closed form expressions are available for the unfiltered and filtered solutions and their time derivatives. However, we use DNS to generate data samples (by solving Equation (2)), as this is the approach one would take for other equations that do not have closed-form solutions. For this purpose, the continuous convection operator \mathcal{A} is discretized using a sixth order central difference stencil on the fine grid $\mathbf{x}^{(N)}$:

$$\mathbf{A} = \frac{1}{60\Delta x^{(N)}} \text{circ}_N(1, -9, 45, 0, -45, 9, -1), \quad (19)$$




Figure 1: Left: Filter radius h used in the simulations. Middle and right: Unfiltered and filtered initial conditions and their Fourier coefficients for one of the training samples.

and similarly for $\mathbf{A}^{(M)}$, where $\text{circ}_\eta(s_{-\alpha}, \dots, s_\alpha) \in \mathbb{R}^{\eta \times \eta}$ denotes a circulant matrix where the i -th periodically extended superdiagonal is filled with the stencil entry s_i for $i \in \{-\alpha, \dots, \alpha\}$. Similarly, we build $\mathbf{D}^{(M)} = \frac{1}{180(\Delta x^{(M)})^2} \text{circ}_M(2, -27, 270, -490, 270, -27, 2)$ (of order six).

The unfiltered DNS equations (2) are discretized with $N = 1000$ points, and solved using an adaptive fourth order Runge-Kutta time integrator [16] with absolute and relative tolerances $\epsilon_{\text{abs}} = 10^{-10}$ and $\epsilon_{\text{rel}} = 10^{-8}$. The resulting solution is saved at n_t time points. The time derivative snapshots are obtained by evaluating the DNS right hand side $\mathbf{A}\mathbf{u}$. The filters are discretized using the periodically extended quadrature rule (8). The resulting \mathbf{W} is then used to generate the reference solutions $\mathbf{W}\mathbf{u}$ and time derivatives $\mathbf{W}\mathbf{A}\mathbf{u}$ for a given M .

Four datasets $\mathcal{D} \in \{\mathcal{D}^{\text{train}}, \mathcal{D}^{\text{valid}}, \mathcal{D}^{\text{test}}, \mathcal{D}^{\text{long}}\}$ are created. They consist of waves with a maximum frequency of $K = 250$ and $n_{\text{IC}}^{\mathcal{D}}$ different initial conditions $(u_0^{\mathcal{D},i})_{i \in \{1, \dots, n_{\text{IC}}^{\mathcal{D}}\}}$ given by

$$u_0^{\mathcal{D},i}(x) = \sum_{k=0}^K \frac{1 + \epsilon^{\mathcal{D},i}}{(5+k)^2} \cos(2\pi kx + \theta^{\mathcal{D},i}), \quad i \in \{1, \dots, n_{\text{IC}}^{\mathcal{D}}\}, \quad (20)$$

where $\epsilon^{\mathcal{D},i} \sim \mathcal{N}(0, \frac{1}{5})$ and $\theta^{\mathcal{D},i} \sim \mathcal{U}([0, 2\pi])$ are sampled independently for each \mathcal{D} and i . Note that the fine grid $\mathbf{x}^{(N)}$ resolves all the K frequencies. This choice of coefficients creates a profile where the “turbulent” features (high frequencies) are small compared to the large features, but are still visually different from their filtered counterparts. One of the resulting initial conditions and its coefficients are shown for both filters in Figure 1, along with the common filter radius h .

The parameters for the datasets are given in Table 1. $\mathcal{D}^{\text{train}}$ is used to fit the operators. $\mathcal{D}^{\text{valid}}$ is used to choose the hyperparameters λ . $\mathcal{D}^{\text{test}}$ is used to measure the generalization capacity of the inferred models. $\mathcal{D}^{\text{long}}$ is used to measure the long term stability of the operators.

The operators \mathbf{R} and $\bar{\mathbf{A}}^{\text{DF}}$ are both fitted to $\mathcal{D}^{\text{train}}$ using least squares (Equations (11) and (14), respectively), while $\bar{\mathbf{A}}^{\text{emb}}$ is fitted using 10^4 iterations of stochastic gradient descent with the ADAM optimizer [9] with step size 0.001. A grid search over the regularization parameter $\lambda \in \{10^{-12}, 10^{-11}, \dots, 10^0\}$ is performed to obtain the smallest time averaged validation error $\frac{1}{n_t} \sum_{i=1}^{n_t^{\text{valid}}} e_{\bar{\mathbf{A}}}^{\text{valid}}(t_i^{\text{valid}})$, where the

\mathcal{D}	n_{IC}	n_t	T
train	1000	50	0.1
valid	20	10	1.0
test	100	20	1.0
long	50	500	100.0

Table 1: Parameters used to build the different datasets \mathcal{D} : number of initial conditions ($n_{\text{IC}}^{\mathcal{D}}$), number of time points ($n_t^{\mathcal{D}}$), and convection time ($T^{\mathcal{D}}$).




Figure 2: Different operators for $M = 100$. Note that $\mathbf{R} \in \mathbb{R}^{1000 \times 100}$ is non-square, despite the appearance. Here, $\mathbf{A}^{(100)} = \text{circ}_{100}\left(\frac{5}{3}, -15, 75, 0, -75, 15, -\frac{5}{3}\right)$.

error for a given dataset \mathcal{D} at a time t is given by

$$e_{\bar{\mathbf{A}}}^{\mathcal{D}}(t) = \frac{1}{n_{\text{IC}}^{\mathcal{D}}} \sum_{\bar{\mathbf{u}} \in \mathcal{D}} \frac{\|\mathbf{S}(\bar{\mathbf{A}}, \bar{\mathbf{u}}(0), t) - \bar{\mathbf{u}}(t)\|}{\|\bar{\mathbf{u}}(t)\|}. \quad (21)$$

The resulting operators are shown in Figure 2 for both filters (16) and (17) with $M = 100$. Note that this grid is 10 \times coarser than the DNS grid and cannot represent all frequencies. While the top-hat reconstructor and resulting filtered operator are dense, with large off-diagonal elements, the corresponding Gaussian operators admit a sparser structure. One can clearly recognize the filter radius profile h in the different operators, with more blurred and stretched out features on the left half of the domain. Note also that the embedded operator $\bar{\mathbf{A}}^{\text{emb}}$ is closer to the unfiltered operator $\mathbf{A}^{(M)}$ than $\bar{\mathbf{A}}^{\text{int}}$ and $\bar{\mathbf{A}}^{\text{DF}}$, possibly due to a local minimum near the initial guess $\mathbf{A}^{(M)}$.

In Figure 3, the time averaged relative error $\frac{1}{n_t^{\text{test}}} \sum_{i=1}^{n_t^{\text{test}}} e_{\bar{\mathbf{A}}}^{\text{test}}(t_i^{\text{test}})$ is shown for both filters as a function of M . For the unfiltered operator $\mathbf{A}^{(M)}$ the error initially decreases with M , but stabilizes at a point between 3% and 5% after $M = 100$. This is due to the true commutator error \mathcal{C} , confirming that a closure model is indeed necessary to predict the filtered dynamics. For $M < 100$ the discretization




Figure 3: Average relative error on test dataset $\mathcal{D}^{\text{test}}$ for increasing coarse grid precision M for the top-hat filter (left) and Gaussian filter (right). The error measures generalization capacity of the ODE-model for new initial conditions and time points.

error is more important than the commutator error, but this is test-case dependent. The operators $\bar{\mathbf{A}}^{\text{int}}$ and $\bar{\mathbf{A}}^{\text{DF}}$ have very similar error profiles. In fact, they only differ by the regularization terms, as the data fitting terms are identical (up to a linear transformation). For large M , they both admit errors that are orders of magnitude smaller than $\mathbf{A}^{(M)}$. Note that for small M , the algorithm gives a high weight to the prior for $\bar{\mathbf{A}}^{\text{DF}}$, resulting in the initial overlap with $\mathbf{A}^{(M)}$. The embedded operator $\bar{\mathbf{A}}^{\text{emb}}$ performs the best for small M , but the errors seems to stagnate like those of its initial guess $\mathbf{A}^{(M)}$ for large M . A possible reason is that during gradient descent, the gradients quickly become small, possibly due to vanishing gradients in back-propagation through the solver. The operator may also have ended up in a local minimum close to the initial guess. For all the operators, the errors for the Gaussian filters are smaller than for the top-hat filter, as the higher frequencies are filtered more strongly.

To further investigate the long term stability of the inferred operators, the evolution of the mean relative error $e_{\bar{\mathbf{A}}}^{\text{long}}$ is computed on the long term testing data $\mathcal{D}^{\text{long}}$ for $M = 100$. The results are shown in Figure 4. At $t = 0$, all operators have an error of zero, since the initial conditions are given by the filtered DNS solution. The unfiltered coarse-scale operator $\mathbf{A}^{(M)}$ initially has the highest errors, as it does not account for the filter. In addition, the error profile is not fully periodic, but is increasing at a linear rate. This indicates that the filtered solutions contain frequencies not fully resolved by the sixth order central difference convection stencil on the coarse grid. The operators $\bar{\mathbf{A}}^{\text{int}}$ and $\bar{\mathbf{A}}^{\text{DF}}$ initially perform the best and show similar error profiles. However, the intrusive operator is not regularized for stability, and the errors start increasing rapidly after a few convection periods, eventually surpassing the ones of the unfiltered operator. This is likely due to instabilities from anti-diffusion on the parts of the domain where the filter radius h is decreasing (on $[\frac{1}{4}, \frac{3}{4}]$). For $\bar{\mathbf{A}}^{\text{DF}}$ the errors seem to be increasing at the same rate as $\mathbf{A}^{(M)}$. The errors for the embedded operator $\bar{\mathbf{A}}^{\text{emb}}$ stay about one order of magnitude below those of $\mathbf{A}^{(M)}$, and increase at the same rate.

Note that putting the continuously filtered initial conditions \bar{u}_0 into the unfiltered continuous




Figure 4: Evolution of the relative error on $\mathcal{D}^{\text{long}}$ for $M = 100$ using the top-hat filter (left) and Gaussian filter (right). The abscissa show the time elapsed starting from the filtered initial conditions to one hundred convection periods on Ω .

equations (1) would lead to errors as the commutator error \mathcal{C} would be not be accounted for. But the error profile would still be periodic, as the exact solution to (1) is periodic for all initial conditions. In particular, the error would be exactly zero when $t = n$ for all $n \in \mathbb{N}$. This does not seem to be the case for the unfiltered reference operator $\mathbf{A}^{(M)}$. The inferred operators are thus not only learning to account for the explicit filter \mathcal{F} , but also for the under-resolved discretization $\mathbf{x}^{(M)}$.

5 CONCLUSION

In this article we have considered the effect of applying discrete non-uniform filters to a linear partial differential equation. Different discrete closure models were considered, using filtered data samples to infer operators. Firstly, we built an explicit reconstruction model which allows us to use the DNS equations to evolve the system in time. This approach gave satisfactory short-term errors, but was shown to lack long-term stability. It also requires having access to the full order model operator \mathbf{A} on the fine discretization. Secondly, we built the discrete filtered operator directly, using gradient-free and embedded approaches. The former gave similar error profiles to the intrusive approach, while the latter gave less significant improvements for the cases considered. When building the filtered operator directly, we learn not only the correction for the filter, but also possibly improving the finite difference approximations to the unfiltered operator itself.

In the case of the convection equation, the least-squares time-derivative fit on the coarse grid was found to be just as performant as the one involving the DNS-grid, being sufficient for inferring operators with good extrapolation and generalization capacities. On the other hand, the embedded model did not achieve the same performance. It requires the choice of more hyperparameters, such as the initial guess, number of iterations, step size, regularization, and choice of sensitivity algorithm for computing the gradients. Care must be taken to avoid local minima.

In conclusion, this study shows that constructing a discrete closure model with an embedded strategy, even for a linear model problem, leads to a complicated optimization problem with many

choices (such as hyperparameters) and does not easily lead to more accurate or stable results than the more straightforward ‘derivative fitting’ approach. This indicates that more advanced parameterizations, such as neural networks, are also expected to require very careful consideration in order to work in an embedded set-up.

Software availability statement

The source code for the simulations and figures is available at <https://github.com/agdestein/DiscreteFiltering.jl>. All the simulations were run on a recent laptop computer.

REFERENCES

- [1] Y. Bar-Sinai, S. Hoyer, J. Hickey, and M. P. Brenner. Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences of the United States of America*, 116(31):15344–15349, 2019.
- [2] A. Beck and M. Kurz. A perspective on machine learning methods in turbulence modeling. *GAMM Mitteilungen*, 44(1):1–27, 2021.
- [3] J. P. Boris. On large eddy simulation using subgrid turbulence models comment 1. In J. L. Lumley, editor, *Whither Turbulence? Turbulence at the Crossroads*, pages 344–353, Berlin, Heidelberg, 1990. Springer Berlin Heidelberg.
- [4] J. P. Boris, F. F. Grinstein, E. S. Oran, and R. Kolbe. New insights into large eddy simulation. *Fluid Dynamics Research*, 10:199–228, 1992.
- [5] L. Bottou et al. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nimes*, 91(8):12, 1991.
- [6] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural Ordinary Differential Equations. *arxiv:1806.07366*, 2018.
- [7] S. Ghosal and P. Moin. The Basic Equations for the Large Eddy Simulation of Turbulent Flows in Complex Geometry. *Journal of Computational Physics*, 118(1):24–37, apr 1995.
- [8] P. C. Hansen. *Discrete Inverse Problems*. Society for Industrial and Applied Mathematics, jan 2010.
- [9] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arxiv:1412.6980*, 2014.
- [10] D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer. Machine learning-accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21), may 2021.
- [11] B. List, L.-W. Chen, and N. Thuerey. Learned turbulence modelling with differentiable fluid solvers. *arxiv:2202.06988*, 2022.
- [12] A. L. Marsden, O. V. Vasilyev, and P. Moin. Construction of commutative filters for LES on unstructured meshes. *Journal of Computational Physics*, 175(2):584–603, 2002.

- [13] B. Peherstorfer and K. Willcox. Data-driven operator inference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering*, 306:196–215, 2016.
- [14] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, A. Ramadhan, and A. Edelman. Universal differential equations for scientific machine learning. *arxiv:2001.04385*, 2020.
- [15] P. Sagaut. *Large Eddy Simulation for Incompressible Flows*. Scientific Computation. Springer-Verlag, Berlin/Heidelberg, 2006.
- [16] C. Tsitouras. Runge–Kutta pairs of order 5 (4) satisfying only the first column simplifying assumption. *Computers & Mathematics with Applications*, 62(2):770–775, 2011.
- [17] K. Um, R. Brand, Y. R. Fei, P. Holl, and N. Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative PDE-solvers. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6111–6122. Curran Associates, Inc., 2020.
- [18] H. Van Der Ven. A family of large eddy simulation (LES) filters with nonuniform filter widths. *Physics of Fluids*, 7(5):1171–1172, 1995.
- [19] O. V. Vasilyev, T. S. Lund, and P. Moin. A General Class of Commutative Filters for LES in Complex Geometries. *Journal of Computational Physics*, 146(1):82–104, 1998.