

Consistency and convergence of simulation schemes in information field dynamics

Martin Dupont and Torsten Enßlin

Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

(Received 26 June 2018; published 22 October 2018)

We explore a new simulation scheme for partial differential equations (PDE's) called information field dynamics (IFD). Information field dynamics is a probabilistic numerics method that seeks to preserve the maximum amount of information about the field being simulated. It rests on Bayesian field inference and therefore allows the incorporation of prior knowledge on the field. This makes IFD attractive to address the closure problem of simulations—how to incorporate knowledge about subgrid dynamics into a scheme on a grid with limited resolution. Here we analytically prove that a restricted subset of simulation schemes in IFD are consistent and thus deliver valid predictions in the limit of high resolutions. This has not previously been done for any IFD schemes. This restricted subset is roughly analogous to traditional fixed-grid numerical PDE solvers, given the additional restriction of translational symmetry. Furthermore, given an arbitrary IFD scheme modeling a PDE, it is *a priori* not obvious to what order the scheme is accurate in space and time. For this subset of models, we also derive an easy rule of thumb for determining the order of accuracy of the simulation. As with all analytic consistency analysis, an analysis for nontrivial systems is intractable; thus these results are intended as a general indicator of the validity of the approach, and it is hoped that the results will generalize.

DOI: [10.1103/PhysRevE.98.043307](https://doi.org/10.1103/PhysRevE.98.043307)**I. INTRODUCTION****A. Probabilistic numerics**

Information field dynamics (IFD) is an alternative framework for constructing numerical simulation schemes for partial differential equations (PDE's) that can be cast into the form

$$\partial_t \phi = F[\phi] \quad (1)$$

with ϕ being some field and F some linear or nonlinear endomorphic operator in the Hilbert space of field configurations. IFD was first proposed in Ref. [1] and further developed in Refs. [2] and [3]. IFD addresses the problem of constructing numerical simulation schemes probabilistically. It therefore belongs to the emerging field of *probabilistic numerics* [4].¹

In classical numerics, a point estimate of the quantity of interest is provided. In contrast to this, in probabilistic numerics, a probability distribution over possible values of the quantity of interest is constructed and investigated. This permits the quantification of uncertainties on the results as well as the incorporation of domain knowledge into their estimation.

Classical numerical simulation schemes therefore follow a representation of the evolving field in time. This is done by updating some data in computer memory, which, for example, specifies the field values within the voxels of a discretized field domain. The field configuration and its discretized representation are the central elements of classical simulation schemes. In probabilistic numerics, the central object is not the field itself but the probability of the field having a specific configuration. The basic idea of IFD is therefore to follow

the evolution of the probability distribution of possible field states. The data in computer memory of a probabilistic simulation scheme therefore parametrize this probability distribution. This distribution can be used to ask questions, such as What is the most probable field configuration, What is the mean field, or What is its uncertainty covariance?

Some of the earliest works in the field of probabilistic numerics are by the authors of Refs. [5–7], who treated the problem of function interpolation as a statistical inference problem. For the study of differential equations, Ref. [8] was the first to propose treating ordinary differential equations (ODE's) as a Bayesian inference problem. References [9,10] were early examples of using randomized Monte Carlo methods for the solution of ODE's. More recently, Ref. [11] used Bayesian uncertainty to quantify and reduce errors in numerical integration, and Ref. [12] develops a Gaussian process solver which generates a probability distribution over a set of solutions to an ODE, which are centered around a Runge-Kutta solution. Reference [13] is particularly relevant for our work, as it analyzes the convergence properties and errors of probabilistic solvers for PDE's.

Our proposed scheme is somewhat different to previous works, as our algorithms are deterministic, although they nonetheless incorporate a notion of prior belief, measurement, and uncertainty. One paper which proposes a scheme similar to IFD is Ref. [14], which solves stochastic differential equations by incorporating prior beliefs, and a notion of observations, and, like our work, derives its equations of motion using the criteria of minimal Kullback-Leibler divergence.

B. Minimal information loss

The dynamical equation that governs the evolution of the field ϕ determines how the probability distribution of field

¹See also <http://probabilistic-numerics.org>.

configurations should evolve. IFD attempts to follow the evolution of this probability distribution as a whole. Since the representation of this distribution is parametrized and therefore cannot perfectly represent any resulting distribution, an approximation scheme is required to map the evolved distribution back into the space of distributions that can be represented by the computer data. We take it as a principle that this approximation should conserve as much of the information of the full distribution as possible.

An appropriate measure of the amount of information lost in this approximation is the Kullback-Leibler (KL) divergence of the approximated distribution $\mathcal{Q}(\phi)$ with respect to the more accurate distribution (or measure) $\mathcal{P}(\phi)$ [15],

$$D_{\text{KL}}(\mathcal{P}||\mathcal{Q}) = \int \mathcal{D}\phi \mathcal{P}(\phi) \ln \left[\frac{\mathcal{P}(\phi)}{\mathcal{Q}(\phi)} \right]. \quad (2)$$

Here $\mathcal{P}(\phi)$ stands for the time-evolved probability of field configurations, and $\mathcal{Q}(\phi)$ for the parametrization chosen to represent this approximatively. The more mathematically inclined reader should read the path integral $\int \mathcal{D}\phi \mathcal{P}(\phi)$ as $\int d\mathcal{P}(\phi)$, an average over ϕ with the measure $\mathcal{P}(\phi)$.

The KL is now minimized with respect to the parameters of \mathcal{P} . This then provides an update rule for these parameters in computer memory, which then represents the desired simulation scheme.

This scheme evolves a probability distribution via its chosen representation in computer memory. The action principle for this is given by the requirement of minimal information loss, ensuring that the simulation is as accurate as possible. The precise details of the derivation will be spelled out in the following sections.

C. Information field theory

In principle, any suitable parametrization of the field probability distribution could be chosen and the field dynamic mapped to it via the principle of minimal information loss. In IFT [16–18], a canonical representation of field probabilities as a function of some data already exists.

IFT was designed to address the problem of field inference from measurement data. It turns the data d into a posterior probability $\mathcal{P}(\phi|d)$ on the field ϕ , which was measured. Thus, IFT provides us with a convenient parametrization of field probabilities, which depend on data. Now, IFD uses this parametrization and therefore just regards the data in computer memory as the result of a virtual measurement process (more on this later). The process can be chosen arbitrarily; however, it should ideally provide an analytically tractable posterior distribution $\mathcal{P}(\phi|d)$.

For this reason, a simple linear measurement equation of the form

$$d = R\phi + n \quad (3)$$

will be chosen for this paper. Here R represents the measurement response, an operator that maps the continuous field into a finite-dimensional data vector, and n is some field-independent Gaussian random noise vector with known covariance $N = \langle n n^\dagger \rangle_{(n)}$.

As the data are finite, but the field has infinitely many degrees of freedom, field inference is usually an ill-posed

problem that requires regularization, i.e., the removal or suppression of implausible solutions that are otherwise allowed by the data. This is provided by the field prior, which for convenience assumes the field to be drawn from a zero-centered Gaussian process with covariance $\Phi = \langle \phi \phi^\dagger \rangle_{(\phi)}$,

$$\mathcal{P}(\phi) = \mathcal{G}(\phi, \Phi) \equiv \frac{1}{|2\pi\Phi|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}\phi^\dagger\Phi^{-1}\phi\right), \quad (4)$$

where $\phi^\dagger\psi \equiv \int dx \overline{\phi(x)}\psi(x)$ denotes the canonical scalar product of the Hilbert space.² Bayes's theorem then provides the field posterior,

$$\mathcal{P}(\phi|d) = \frac{\mathcal{P}(d|\phi)\mathcal{P}(\phi)}{\mathcal{P}(d)}, \quad (5)$$

where $\mathcal{P}(d|\phi) = \mathcal{G}(d - R\phi, N)$ is the likelihood; the probability of the obtained data d given a field configuration ϕ . This field posterior turns out to be a Gaussian

$$\mathcal{P}(\phi|d) = \mathcal{G}(\phi - m, D) \quad (6)$$

under the simplifying assumptions made here [17]. The posterior mean

$$m = W d \equiv D R^\dagger N^{-1} d \quad (7)$$

is a linear function of the data, as $m = m(d)$. In contrast to this, the posterior uncertainty dispersion

$$D = (\Phi^{-1} + R N^{-1} R)^{-1} \quad (8)$$

is independent of the data for this linear and Gaussian field estimation problem. The operator W turning the data into the posterior mean field is called Wiener filter in signal processing and D is also called the Wiener covariance. In this paper, we will often express this operator in its so-called data space version,³

$$W = \Phi R^\dagger (R \Phi R^\dagger + N)^{-1}, \quad (9)$$

as this form allows us to take the no-noise limit $N \rightarrow 0$, in which $W \rightarrow \Phi R^\dagger (R \Phi R^\dagger)^{-1}$. IFT also extends this linear signal inference to nonlinear and non-Gaussian problems. However, for IFD as developed so far, this is not needed and the Wiener filter theory presented here is sufficient.

D. Information field dynamics

IFD regards the data as being in some sense a measurement of the field being simulated. This idea can be taken literally or not. For example, using the formalism one could take the initial data to be the result of a literal measurement, and IFD would prescribe a way of simulating how future measurements of the time-evolved field would appear.

However, one can also regard the response R as simply being a mathematical object which creates a finite-dimensional representation of the continuous field being simulated, and the noise just represents some degree of uncertainty in this

²In this paper we will refer to this Hilbert space as *field space* and the vector space in which the data resides will be called *data space*.

³This is so named, since here the operator inversion happens in data space.

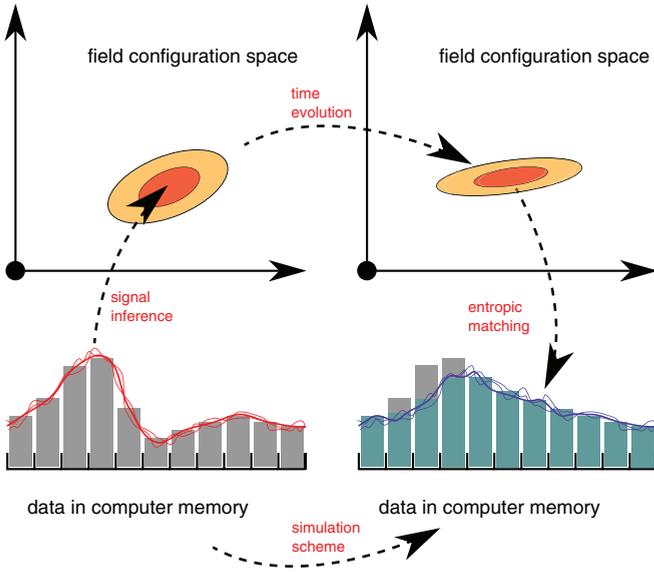


FIG. 1. Schematic representation of the IFD concept. Data in computer memory (gray bars) imply via Bayesian signal inference a posterior probability distribution in field space (orange density contours on the left). For each field configuration (each point of this space) the time evolution assigns it a new location. The time-evolved probability distribution (orange contours on the right) needs to be reexpressed in terms of new data. This is done via entropically matching this (or minimizing the KL) to a new posterior distribution (blue contours) expressed by new, time-evolved data (blue bars). The set of implied operations on the data represents the IFD simulation scheme. It incorporates the field dynamics, prior knowledge on the field (e.g., on subgrid scales), and tries to conserve as much information on the field as possible.

description. If, for example, one defines a response that takes samples of the field at particular points and sets the noise to zero, then the data become identical to the gridpoints of a typical finite-difference scheme. Thus, the idea of representing data as a virtual measurement of a field being simulated should not be too unusual. The point is that in order to run a simulation, one needs to express a field using a finite amount of information, and the formalism of IFT provides us with a convenient way of doing so.

In an IFD simulation, there are many parameters that may be updated in time to best capture the time evolution of the posterior distribution. These include the data d but potentially also the properties of the measurement equations, the response R , and the noise covariance N , as well as the field prior covariance Φ .

Changing the data d while keeping all other parameters constant corresponds most closely to a typical finite-difference simulation scheme; the data points represent samples of a field, and these change in time as the simulation progresses. A response R which changes in time would be analogous to changing the coordinate system during a simulation in order to best capture the behavior of the field under consideration. In this paper, we will only consider schemes where only the data are updated.

We now present an abridged derivation of the IFD simulation scheme, which is shown diagrammatically in Fig. 1.

We restrict ourselves to linear dynamics:

$$\partial_t \phi = L \phi, \tag{10}$$

with L some linear, time-independent, endomorphic operator. This equation has a formal solution given by $\phi(t) = U(t)(\phi_0)$ for $U(t) = \exp(tL)$.

The scheme is then as follows: It is assumed that there are some data d_i taken at some point in time, t_i , which is interpreted as being some coarse-grained representation of the true field $\phi(t_i, x)$, which is obtained by some linear measurement as in Eq. (3). To run the simulation, i.e., obtain the data at t_{i+1} from that at t_i , IFT is used to reconstruct the posterior probability distribution of the field, $\mathcal{P}(\phi(t_i, x)|d_i)$, given the initial data, d_i .

This posterior distribution is then evolved from t_i to t_{i+1} using the equations of motion for the field. This would formally be done using $U(\Delta t)$. However, to achieve a practical simulation scheme, the time evolution U must be truncated to some finite order,⁴ which we denote by $\bar{U} = \sum_{k=0}^{\alpha} (\Delta t L)^k / k!$ for some order α , which corresponds to a choice of the desired time-order accuracy of the simulation.

To obtain the data at the next timestep, a second measurement, R_{i+1} , is postulated, which is used to construct a second posterior distribution. The new data are then chosen as to minimize the information loss between the evolved and unevolved posterior distributions, using the Kullback-Leibler divergence (KL divergence). With linear dynamics, measurements and a Gaussian prior,⁵ the resulting finite-difference equation becomes particularly transparent [3]:

$$d_{i+1} = (R_{i+1} W_{i+1})^{-1} R_{i+1} \bar{U} W_i d_i. \tag{11}$$

The subscripts denote time indices, as the response and prior covariance are allowed to vary between timesteps. The above equation is in the form of a matrix equation, and the matrix will be referred to as the difference operator and will be denoted by $T_i = (R_{i+1} W_{i+1})^{-1} R_{i+1} \bar{U} W_i$.

A new development, which we present in this paper, is that this difference equation may be simplified further in the linear case, because repeated applications of the difference operator during the simulation will result in certain useful cancellations, the result of which is that we may always assume that we are operating in the no-noise limit.

Lemma 1. The finite-difference equations for linear IFD are independent of the noise up to a simple equivalence.

Proof. For a simulation scheme with timesteps t_i for $i \in \{1, \dots, n\}$, responses $\{R_i\}$, priors $\{\Phi_i\}$, noises $\{N_i\}$, Wiener filters $\{W_i = \Phi_i R_i^\dagger (R_i \Phi_i R_i^\dagger + N_i)^{-1}\}$, and linear time evolution operators $\bar{U}_i = \mathbb{1} + \Delta t L_i + \dots$, the finite-difference equation is given by

$$\begin{aligned} d_{i+1} &= (R_{i+1} W_{i+1})^{-1} R_{i+1} \bar{U}_i W_i d_i \\ &= [R_{i+1} \Phi_{i+1} R_{i+1}^\dagger (R_{i+1} \Phi_{i+1} R_{i+1}^\dagger + N_{i+1})^{-1}]^{-1} \\ &\quad \times R_{i+1} \bar{U}_i \Phi_i R_i^\dagger (R_i \Phi_i R_i^\dagger + N_i)^{-1} d_i. \end{aligned} \tag{12}$$

⁴If we could write down a closed-form expression for U , then we would not need to run the simulation.

⁵We refer to this as the *linear case* of IFD.

The second line is obtained by inserting the data-space definition of the Wiener filter. We rename the terms: $(R_i \Phi_i R_i^\dagger + N_i) = C_i$, $(R_i \Phi_i R_i^\dagger) = B_i$, and $R_{i+1} \bar{U}_i \Phi_i R_i^\dagger = A_i$, yielding

$$d_{i+1} = (B_{i+1} C_{i+1}^{-1})^{-1} A_i C_i^{-1} d_i = C_{i+1} B_{i+1}^{-1} A_i C_i^{-1} d_i. \quad (13)$$

The difference equations are then iterated n times. With every matrix multiplication, there is a C_i which cancels with a C_i^{-1} , yielding

$$d_n = C_n \left(\prod_{i=0}^{n-1} B_{i+1}^{-1} A_i \right) C_0^{-1} d_0. \quad (14)$$

The only noise-dependent terms were the C terms and therefore, up to a change of basis at the beginning and end of the simulation, the difference equations are independent of the noise. In the infinite-noise limit, $C \rightarrow N$, and in the zero noise limit $C \rightarrow B$. ■

Given the equivalence, from here onward we will always work in the no-noise limit, and the symbol N will be used to denote number quantities. In this limit, the difference operator becomes

$$T_i = R_{i+1} \bar{U}_i \Phi_i R_i (R_i \Phi_i R_i^\dagger)^{-1}. \quad (15)$$

Although these incarnations of IFD schemes, Eqs. (11) and (15), might intuitively make sense, it still is not guaranteed that they lead to consistent and convergent numerical simulations.

E. Numerical consistency and error

The major goal of this paper is to show that in a restricted setting, the IFD equations are *consistent*. This is a valuable goal, for the Lax equivalence theorem [19] states that if a scheme is consistent, then it converges to the true solution if and only if it is *stable*. We state the (paraphrased) definition of consistency:

Definition 2 (Consistency). For an operator $T(\Delta t, \Delta x)$ which approximates $U(t)$, with $U(t)$ being the analytic time evolution operator corresponding to $L(t)$, the approximation is said to be consistent if for some set of genuine⁶ solutions Ω to the differential equation, then for any $\phi \in \Omega$,

$$\lim_{\Delta t, \Delta x \rightarrow 0} \|(T(\Delta t, \Delta x) - U(\Delta t))\phi(t, x)\| = 0 \quad (16)$$

uniformly in t .

Note that the above definition involves comparing operators which are defined on different spaces: $T(\Delta t, \Delta x)$ acts on a discrete space, yet $U(\Delta t)$ acts on a continuous space. Reference [19] assumes that there is some sufficient level of smoothness such that Taylor series expansions or smooth interpolation, etc., may be used to approximate the norm. We discuss a way of comparing these two operators in IFD later, once a more concrete expression for T has been found.

The other goal of this paper is to analyze the numerical error of IFD schemes and how such error scales as the spatial and temporal resolutions Δx and Δt become arbitrarily fine. IFD is a nominally information-theoretic framework, so it

is conceivable that one could try and use some information-theoretic notion of error. However, this would limit our ability to compare the performance of IFD schemes to standard finite-difference schemes. Thus, in this paper, we follow a standard approach and analyze the *local truncation error* or the *one-step error* [20, p. 593]. We do this by analyzing the distance in the operator norm of the difference operator and the true analytic time evolution operator. This distance provides a bound on the error which can accumulate during a single timestep of the simulation.

II. ARGUMENTATION

A. Groundwork

We now begin the work of proving consistency by defining the type of models we will be working on. We restrict our focus to PDE's for which L is translation invariant. This case can already be solved analytically by Fourier analysis. However, this practice is entirely normal in numerical methods, as many advanced simulation schemes are too complicated to permit an analytic analysis [21, ch. 7]. Such is the case in IFD, as the codes are typically nonlocal, meaning the algebraic equations tend to be dependent on the global geometry of the simulation domain. Thus, the best that one can do is prove convergence for the analytically solvable case and then hope that these conclusions hold in the nonanalytically solvable case. For pedagogical clarity, the results presented in this section are for one spatial dimension only, although it is argued later that they generalize trivially. Furthermore, as the equations we present are somewhat opaque, we provide an Appendix which shows a worked example of a simple real system for interested readers.

We first restrict ourselves to the case where the response R and prior Φ do not change in time, i.e., the coordinate system is static and our prior belief about the system will not change during the simulation. Because these quantities are now static, there is no need to subscript them to denote the timestep in question. This allows us to free up the subscripts for other purposes.

We now select the field and data spaces. The simulated space inside the computer must always be of finite extent. For this reason, we choose the field space to be $\mathcal{L}^2([0, l])$. We apply periodic boundary conditions to render the analytic equations tractable.

Now the prior must be selected. If the PDE under consideration is translation invariant, then one should choose a prior belief which is also translation invariant. Thus the prior covariance will have a diagonal representation in Fourier space. The positivity and self-adjointness conditions on the prior covariance ensure that the eigenvalues in momentum space will be everywhere positive and greater than zero and symmetric about the origin. Priors of this form are generally referred to as smoothness priors. Using k to denote momentum, a prior Φ_{kk} whose values fall to zero as $k \rightarrow \infty$ essentially states that rapid oscillations in the signal are deemed unlikely; the field is smooth. For notational convenience, we will often denote the diagonal entries of the prior, Φ_{kk} as $\Phi(k)$.

Simple examples of a prior include power laws in momentum, i.e., $|k|^{-\beta}$ for some integer β , often supplemented by a regularizing mass term: $\Phi(k) = 1/(|k|^\beta + m^\beta)$.

⁶See the original publication for a definition of a genuine solution.

We pick the responses by assuming that we have N spatial points which will be labeled with the index j . The responses are chosen to be constant in time, and the subscripts R_j now denote spatial indices. The most natural and naive response is to choose the index j to label a regular grid of positions. We define $\Delta x = l/N$. We let the response be any response which measures the field by integrating over some function $B(x)$ on $\mathcal{L}^2([0, l])$ localized at the point x_j :

$$(R\phi)_j = \int_0^l dx B(x - x_j)\phi(x), \quad (17)$$

where x_j is the x position of the j th gridpoint, i.e., $x_j = \Delta x \cdot j$. A simple example of such a function could be the box function:

$$B(x) = \begin{cases} 1/\Delta x & 0 \leq x \leq \Delta x \\ 0 & \text{otherwise} \end{cases}. \quad (18)$$

The response is then an average of the field around that point. If the x_j 's are evenly spaced, then we refer to any response of this form as a *translation-invariant response*. The $B(x)$ functions will be referred to as the *response bins* or just *bins*.

We now begin to calculate the difference operator, starting with the computation of $(R\Phi R^\dagger)^{-1}$. Since both the responses and prior covariance are invariant under translations of multiples of Δx , we can make a very general statement:

Lemma 3. Given a field space of the form $\mathcal{L}^2([0, l])$ with periodic boundary conditions, a translation invariant response R_j whose bin function $B(x)$ has a Fourier series representation, as well as a prior covariance Φ which is diagonal in momentum space, $(R\Phi R^\dagger)_{jl}$ will be of the form:

$$\sum_k \Phi(k) |\widehat{B}(k)|^2 e^{ik(x_j - x_l)}, \quad (19)$$

where $\widehat{B}(k)$ is the Fourier coefficient of $B(x)$.

Proof. By the shift property of the Fourier transform, $\widehat{R}_{j,k} = e^{-ikx_j} \widehat{R}_{0,k} = e^{-ikx_j} \widehat{B}(k)$. Therefore $(R\Phi R^\dagger)_{jl}$ is

$$\begin{aligned} (R\Phi R^\dagger)_{jl} &= \sum_k \sum_q e^{ikx_j} \widehat{B}(k) \Phi_{kq} \widehat{B}^*(q) e^{-iqx_l} \\ &= \sum_k \Phi(k) |\widehat{B}(k)|^2 e^{ik(x_j - x_l)} \end{aligned} \quad (20)$$

as desired. ■

It must be stressed that we are not demanding that the simulation is carried out in Fourier space, we are rather stating that the operator will always have such a representation. From now onward, any simulation scheme which satisfies the criteria of the previous lemma, and in addition has a translation-invariant time evolution operator U , will be referred to as a *translation-invariant scheme*.

The $R\Phi R^\dagger$ matrix must now be inverted; however, the inverse is not equal to the inverse of the Fourier coefficients. Observe that the spatial gridpoints are both finite and discrete, which means that terms such as $\sum_j e^{ix_j(k-q)}$ do not form Kronecker deltas δ_{kq} . The sum equals N , not only when $k = q$ but also when $\frac{l}{2\pi N}(p - q)$ is an integer.

The reason for this is that data space is a discrete periodic interval which has a discrete Fourier transform (DFT). For a DFT, the momentum values k are the same as those for the

continuous interval, albeit with a highest uniquely resolvable frequency known as the *Nyquist frequency*, which is equal to half of the sampling frequency. In this case, the Nyquist frequency is $\frac{\pi}{\Delta x}$ and is denoted by f_N . Given that the matrix is indeed translation invariant in data space, it must have some diagonal representation in the discrete Fourier transform, i.e., some scalar function of k , for k now less than f_N . This representation can be found by resumming over multiples of the Nyquist frequency:

Lemma 4. Given a regular, discrete grid of points $\{x_j\}$ for $j \in \{1, \dots, N\}$ on a periodic interval, and a matrix of the form:

$$A_{lj} = \sum_{k=-\infty}^{\infty} f(k) e^{ik(x_l - x_j)} \quad (21)$$

for $f(k)$ some function of k , it has a diagonal representation in the DFT Fourier space, given by

$$A_{lj} = \sum_{|k|}^{f_N} \underbrace{\left(\sum_{b \in 2f_N \mathbb{Z}} f(k+b) \right)}_{\equiv g(k)} e^{ik(x_l - x_j)} = \sum_{|k|}^{f_N} g(k) e^{ik(x_l - x_j)}. \quad (22)$$

Proof. We partition the infinite sum over k in Eq. (21) into smaller sums shifted by multiples of the Nyquist frequency.⁷ For any x_i and x_j separated by a multiple of Δx and $b = 2\pi n/\Delta x$, we have $(k+b)(x_i - x_j) = k(x_i - x_j) + 2\pi n$. This factor of 2π then disappears in the complex exponential, yielding the desired result:

$$A_{lj} = \sum_{|k|}^{<f_N} \sum_{b \in 2f_N \mathbb{Z}} f(k+b) e^{i(k+b)(x_l - x_j)}, \quad (23)$$

$$\stackrel{\text{Nyquist}}{=} \sum_{|k|}^{f_N} \left(\sum_{b \in 2f_N \mathbb{Z}} f(k+b) \right) e^{ik(x_l - x_j)}. \quad (24)$$

This resummed function is a diagonal function of the DFT frequencies $k < f_N$ and so must be the desired operator. ■

Due to the physical analogy with Brillouin zones, we refer to the procedure of summing over multiples of the Nyquist frequency as *the sum over Brillouin zones*.

Now that we have obtained a representation of the operator which is diagonal in the DFT space, inverting follows easily by taking the inverse of the DFT Fourier coefficients:

$$(R\Phi R^\dagger)_{lj}^{-1} = \frac{1}{N} \sum_{|k|}^{f_N} \frac{e^{ik(x_l - x_j)}}{\sum_{b \in 2f_N \mathbb{Z}} \Phi(k+b) |\widehat{B}(k+b)|^2}. \quad (25)$$

⁷Note that depending on whether the number of data points is even or odd, the domain of $|k| < f_N$ changes. For odd N we use the convention that $k \in [-(N-1)/2, (N-1)/2]$ and if it is even we use $k \in [-N/2, N/2 - 1]$.

The factor of N comes from the different normalizations of the DFT and the regular Fourier transform. Fourier modes in the DFT are normalized as $\frac{1}{\sqrt{N}}e^{-ikx_j}$, as opposed to $\frac{1}{\sqrt{l}}e^{-ikx}$ for the continuous Fourier series.

It is now time to compute the second part of the difference operator, $R\bar{U}\Phi R^\dagger$. Given that L is assumed to be translationally invariant, \bar{U} will have a diagonal representation in Fourier space. Thus, the previous lemma (3) applies, and the operator will also be diagonal in the DFT space, with a sum over Brillouin zones. With this information, we may now write down the general form of the difference operator $T = R\bar{U}\Phi R^\dagger(R\Phi R^\dagger)^{-1}$ for translation-invariant systems:

$$T_{ij} = \sum_{|k|} \frac{\sum_{b \in 2f_N\mathbb{Z}} \bar{U}(k+b)\Phi(k+b)|\widehat{B}(k+b)|^2}{\sum_{\hat{b} \in 2f_N\mathbb{Z}} \Phi(k+\hat{b})|\widehat{B}(k+\hat{b})|^2} e^{ik(x_i-x_j)}. \quad (26)$$

The factor of $1/N$ is canceled by a factor of N coming from the sum over spatial indices.

B. Consistency

The main objectives of this paper are concerned with the behavior of the difference operator in the limit of high spatial and temporal resolutions. We wish to show that in this limit, the finite-difference operator approaches the analytic time-evolution operator. As the proof is lengthy, a brief outline is in order. First, because IFD does not necessarily rely on a notion of gridpoints, we define how to take a limit of high spatial resolutions in a sensible way. Once this is achieved, we place some necessary restrictions on the responses, priors, and equations of motion, which are required to guarantee convergence. Finally, we take the Fourier-space representation of the finite-difference operator and take the limit of high resolution, which requires some straightforward yet tedious mathematics. The crux of the proof lies in that in the limit, the higher terms in the sum over Brillouin zones can be ignored, which we will prove using uniform convergence.

In the translation-invariant schemes we have defined, the spatial points are not just localized samples of the field, but rather the outputs of a response operator R , so it is not immediately clear how to take a limit of spatial resolution. The number of points x_i could be increased, although as the number of points approaches infinity, the data would not begin to look like the true field, as the shape of the bin functions has not changed. Thus, we propose that the correct way to increase the spatial resolution is to increase the number of bins and decreasing the width of the bin functions such that they approach something resembling δ functions in the limit. Rigorously formulated, the process is as follows: Given some initial resolution Δx_0 for which the points x_i are evenly spaced in $[0, l]$, pick an integer λ ranging from 1 to infinity and then set $\Delta x = \Delta x_0/\lambda$. Then, given the initial bin function $B(x)$, replace it with a function $B_\lambda(x) \equiv \lambda B(\lambda x)$. The limit of $\Delta x \rightarrow 0$ is then taken by letting $\lambda \rightarrow \infty$. This process guarantees that the data remain finite and well behaved in the limit.

To prove consistency, we ask if $T \rightarrow U$ in the limit of high resolution. However, as stated earlier, the two operators act on different spaces. Fortunately, in the Fourier representation, comparing the action of T and U is simple. Observing Eq. (26), one sees that the Fourier space representation of the difference operator, $T(k)$, is defined for all $k < f_N$, whereas $U(k)$ is defined for all k . As the spatial resolution increases, eventually the Nyquist frequency ($f_N = \pi/\Delta x$) will become greater than any fixed k . Thus in the limit, the domain of definition of $T(k)$ approaches that of $U(k)$, and the two operators may be compared. We will show that $T(k) \rightarrow \bar{U}(k)$ as $\Delta x \rightarrow 0$, and since $\bar{U}(k) \rightarrow U(k)$ as $\Delta t \rightarrow 0$, this will prove that $T(k) \rightarrow U(k)$ in the joint limit.

We can prove consistency in the translation-invariant case by adding some light restrictions: the response bins $B(x)$ are compactly supported on some strict subset S of $[0, l]$, with bounded Fourier transform, and $\widehat{B}(0) \neq 0$. We also require that $U(k)\Phi(k) \rightarrow 0$ as $k \rightarrow \pm\infty$.

The restrictions placed on the response bins are not too severe. The bounded Fourier transform requirement will almost always be true for any reasonable response bins. In particular, the *Paley-Wiener theorem* [22] states that this requirement will hold for all smooth, compactly supported functions. The example box-response introduced in Eq. (18) satisfies the restrictions, as will any smooth bump function.

We now seek a formula for $T(k)$ as a function of λ , which we call $T_\lambda(k)$. The compact support property of the bins allows us to exploit the fact that up to a normalization constant, the coefficients $\widehat{B}(k)$ of the discrete values of k in the Fourier series of the bins are the same as the values at k in the continuous Fourier transform of $B(x)$:

$$\int_0^l dx B(x)e^{-ikx} = \int_S dx B(x)e^{-ikx} = \int_{\mathbb{R}} dx B(x)e^{-ikx}. \quad (27)$$

This in turn allows us to exploit the scaling property of the Fourier transform to obtain a convenient form for $\widehat{B}_\lambda(k)$:

$$\widehat{B}_\lambda(k) = \lambda \widehat{B}(\lambda x) = \lambda \frac{1}{\lambda} \widehat{B}(k/\lambda) = \widehat{B}(k/\lambda). \quad (28)$$

Now observe the sum over the Brillouin zones in Eq. (26). The formula contains a sum over $b \in 2f_N\mathbb{Z}$ where $f_N = \pi/\Delta x$ and thus $f_N^\lambda = \pi\lambda/\Delta x_0$ and $b^\lambda = 2\pi n\lambda/\Delta x_0$ for $n \in \mathbb{Z}$. We now insert the definitions of $\widehat{B}_\lambda(k)$ and b^λ into Eq. (26) to yield a formula for $T_\lambda(k)$:

$$\frac{\sum_{n \in \mathbb{Z}} \bar{U}(k + \frac{2\pi n\lambda}{\Delta x_0})\Phi(k + \frac{2\pi n\lambda}{\Delta x_0})|\widehat{B}_\lambda(\frac{1}{\lambda}[k + \frac{2\pi n\lambda}{\Delta x_0}])|^2}{\sum_{m \in \mathbb{Z}} \Phi(k + \frac{2\pi m\lambda}{\Delta x_0})|\widehat{B}_\lambda(\frac{1}{\lambda}[k + \frac{2\pi m\lambda}{\Delta x_0}])|^2}. \quad (29)$$

The λ term inside \widehat{B} can be absorbed to give:

$$\left| \widehat{B}_\lambda\left(\frac{1}{\lambda}\left[k + \frac{2\pi n\lambda}{\Delta x_0}\right]\right) \right|^2 = \left| \widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right) \right|^2. \quad (30)$$

We expect that in the limit of $\lambda \rightarrow \infty$, the higher terms in the sums vanish, leaving only terms in the first Brillouin zone.

That is, we can express the numerator of Eq. (29) as:

$$\bar{U}(k)\Phi(k)|\widehat{B}\left(\frac{k}{\lambda}\right)|^2 + \underbrace{\sum_{n \neq 0} \bar{U}\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\Phi\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right)|^2}_{\equiv \delta(k, \lambda) \rightarrow 0}, \quad (31)$$

where we rename the sum $\delta(k, \lambda)$ to denote that the term (hopefully) vanishes as $\lambda \rightarrow \infty$. The denominator is expanded similarly and its vanishing term is denoted by $\delta'(k, \lambda)$. We wish to prove that the $\delta(k, \lambda)$ and $\delta'(k, \lambda)$ terms actually do vanish, as this would give us the desired result:

$$\begin{aligned} \lim_{\lambda \rightarrow \infty} T_\lambda(k) &= \frac{\lim_{\lambda \rightarrow \infty} \sum_{n \in \mathbb{Z}} \bar{U}\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\Phi\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right)|^2}{\lim_{\lambda \rightarrow \infty} \sum_{m \in \mathbb{Z}} \Phi\left(k + \frac{2\pi m\lambda}{\Delta x_0}\right)|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi m}{\Delta x_0}\right)|^2} \\ &= \frac{\lim_{\lambda \rightarrow \infty} \bar{U}(k)\Phi(k)|\widehat{B}\left(\frac{k}{\lambda}\right)|^2 + \delta(k, \lambda)}{\lim_{\lambda \rightarrow \infty} \Phi(k)|\widehat{B}\left(\frac{k}{\lambda}\right)|^2 + \delta'(k, \lambda)} = \frac{\bar{U}(k)\Phi(k)|\widehat{B}(0)|^2}{\Phi(k)|\widehat{B}(0)|^2} = \bar{U}(k). \end{aligned} \quad (32)$$

This works provided $\widehat{B}(0) \neq 0$, so that the above denominator remains nonzero, and the equation remains well defined. For δ and δ' , each individual term in the sum [which we denote by $\delta_n(k, \lambda)$] approaches zero in the limit of $\lambda \rightarrow \infty$, because $\bar{U}(k)\Phi(k)$ and $\Phi(k)$ go to zero at large $|k|$ by assumption. Therefore we want to swap the limit and the infinite sum.

Elementary functional analysis states that this is possible if and only if the sequence of δ_n terms converges uniformly to zero in n . We remind the reader that a sequence of functions δ_n converges uniformly to zero if, for any positive ϵ , there is an N such that $\forall n \geq N, |\delta_n(\lambda)| < \epsilon$ for all values of λ .

We prove uniform convergence for δ , and δ' follows trivially. We consider the positive- n half of the sum first, and the negative- n half will also follow trivially. In this case,

$$\begin{aligned} \delta_n(\lambda) &= \bar{U}\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\Phi\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\left|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right)\right|^2 \\ &\leq \left|\bar{U}\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\Phi\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\right| C, \end{aligned} \quad (33)$$

where we bounded the function $|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right)|^2 < C$ for some constant C , which we may do by assumption.⁸

We now use the condition $\Phi(k)\bar{U}(k) \rightarrow 0$ as $k \rightarrow \infty$ to show convergence. This condition means that for $|k|$ large enough $\bar{U}(k)\Phi(k)$ can be bounded by some monotonically decreasing function of $|k|$, call it $g(|k|)$. We start by finding a bound for $\lambda = 1$ and then show that this bound holds for all λ . For $\lambda = 1$, and the desired ϵ bound, we can pick some n large enough such that we are in this decreasing regime, hence

$$\begin{aligned} \delta_n &\leq \left|\bar{U}\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\Phi\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right)\right| C \\ &< g\left(k + \frac{2\pi n\lambda}{\Delta x_0}\right) < \epsilon. \end{aligned} \quad (34)$$

⁸To understand why this bound is necessary, notice that the bin terms do not vanish in the limit of large λ . Intuitively, this is because as the bins become narrower, their Fourier transforms widen out, at the exact same rate as the Nyquist frequency is increasing.

For higher λ and large n , $|k + \frac{2\pi n}{\Delta x_0}| < |k + \frac{2\pi n\lambda}{\Delta x_0}|$, and since we have taken n to be large enough that we are in the decreasing regime, the $g(k)$ bound also holds. Thus the bound holds for all λ . The sequence of functions is therefore uniformly convergent, and Eq. (32) holds. We can now state:

Theorem 5. For a one-dimensional translationally invariant IFD scheme, whose response bins $B(x)$ are compactly supported on a strict subset $S \subset [0, l]$ with bounded Fourier transform and $\widehat{B}(0) \neq 0$, and some time-order approximation $\bar{U}(k)$ to $U(k)$, then the scheme is consistent provided $\lim_{k \rightarrow \infty} \bar{U}(k)\Phi(k) = 0$.

Important to note is that we only require $\bar{U}(k)\Phi(k) \rightarrow 0$, not $U(k)\Phi(k) \rightarrow 0$. For derivative operators s.t. $U = \exp(\Delta t \partial_x) = \exp(i \Delta t k)$ or similar, this would require that the prior covariance, $\Phi(x)$, is infinitely differentiable (smooth). Using the approximated time expansion, the prior covariance only needs to be as many times differentiable as the order of the expansion dictates.

C. Error scaling

We seek an estimate of the one-step error E by calculating the difference in the operator norm:

$$E \propto \|T(\Delta t, \Delta x) - U(\Delta t)\| \quad (35)$$

and analyzing the rate of convergence in terms of $O(\Delta x)$ and $O(\Delta t)$. We calculate the error for a fixed value of k , thus $E = E(k)$, although it will be seen shortly that the scaling of the error (which is the quantity of interest) is independent of k .

Our approach is as follows: We must consider how the error scales in the limit of high resolutions, and we accomplish this by repeating the construction with λ from the previous proof. Then, to be able to put any fixed bound on the error, we need to know some properties of the prior and the L operator, namely how they scale in proportion to k . These properties are then used to put a bound on the error in the limit $\lambda \rightarrow \infty$.

To begin, we insert the difference operator in Eq. (29) into the definition of the error, yielding:

$$E(k) = \left| \frac{\sum_{n \in \mathbb{Z}} \bar{U}\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \Phi\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \left|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right)\right|^2}{\sum_{m \in \mathbb{Z}} \Phi\left(k + \frac{2\pi m \lambda}{\Delta x_0}\right) \left|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi m}{\Delta x_0}\right)\right|^2} - U(k) \right|. \quad (36)$$

We use the expansion $\bar{U} = \sum_{p=0}^{\alpha} (\Delta t L)^p / p!$ to find the error in terms of powers of L ,

$$E(k) \leq \sum_{p=0}^{\alpha} \frac{\Delta t^p}{p!} \left| \frac{\sum_{n \in \mathbb{Z}} L^p\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \Phi\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \left|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right)\right|^2}{\sum_{m \in \mathbb{Z}} \Phi\left(k + \frac{2\pi m \lambda}{\Delta x_0}\right) \left|\widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi m}{\Delta x_0}\right)\right|^2} - L^p(k) \right|. \quad (37)$$

For each of the Δt^p terms, we will analyze the scaling of the fraction inside the absolute value and then later find an estimate of the total error scaling.

In the limit of high resolutions, we can expand the numerator of said fraction in the same way that we did in Eq. (31): $L^p(k) \Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2 + \epsilon(k, \lambda)$ for some function ϵ , which goes to zero as $\lambda \rightarrow \infty$. We expand the denominator as $\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2 + \delta(k, \lambda)$, with δ being some other small vanishing function.

The strategy is then to find an expression for the fraction in terms of ϵ and δ and bound each term and analyze how fast they approach zero. We use the Taylor expansion for $1/(1-x) \approx 1+x+x^2+\dots$ to expand the denominator in Eq. (37) into

$$\begin{aligned} & \frac{1}{\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2 + \delta(k, \lambda)} \\ &= \frac{1}{\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2} - \frac{\delta(k, \lambda)}{(\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2)^2} + \dots \end{aligned} \quad (38)$$

We then multiply the numerator by the denominator, which gives

$$\begin{aligned} & \left[L^p(k) \Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2 + \epsilon(k, \lambda) \right] \\ & \times \left[\frac{1}{\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2} - \frac{\delta(k, \lambda)}{(\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2)^2} + \dots \right] \\ &= L^p(k) + \frac{\epsilon(k, \lambda)}{\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2} - \frac{L^p(k) \delta(k, \lambda)}{\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2} + \dots \end{aligned} \quad (39)$$

Calculating the power-law scaling of the above terms is complicated by the fact that each has a $|\widehat{B}(\frac{k}{\lambda})|^2$ in the denominator, which has its own scaling with respect to λ . Exploiting the property $\widehat{B}(0) \neq 0$ allows us to write out each of these as a Taylor series and then reuse the $1/(1-x)$ expansion:

$$\begin{aligned} & \frac{1}{\Phi(k) \left|\widehat{B}\left(\frac{k}{\lambda}\right)\right|^2} = \frac{1}{\Phi(k) |\widehat{B}(0)|^2 + O(1/\lambda) + \dots} \\ &= \frac{1}{\Phi(k) |\widehat{B}(0)|^2} + O(1/\lambda) + \dots \end{aligned} \quad (40)$$

We then see that however fast $\epsilon(k, \lambda)$ goes to zero, $\epsilon(k, \lambda) O(1/\lambda)$ goes to zero faster. Since only the slowest-

converging terms are of interest, we can replace $\frac{1}{\Phi(k) |\widehat{B}(\frac{k}{\lambda})|^2}$ with $\frac{1}{\Phi(k) |\widehat{B}(0)|^2}$ without any adverse effects.

The scaling of the ϵ and δ terms can only be estimated if the scaling behavior of the prior and $L(k)$ are known. To this end, suppose that as $|k|$ becomes large, $\Phi(k)$ can be bounded by some decreasing power law in k , $|k|^{-\beta}$ for β positive. We also assume that $L(k)$ can be bounded by some $|k|^\gamma$ for γ positive, as L will typically be a derivative operator, with $\partial_x^n = (ik)^n$. Then $L^p(k)$ will be bounded by $|k|^{p\gamma}$. There will be constants of proportionality, but they are irrelevant with respect to the scaling.

Using the uniform bound C from before, we can bound the ϵ term by

$$\begin{aligned} |\epsilon(k, \lambda)| &= \left| \sum_{n \neq 0} L^p\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \Phi\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \right. \\ & \times \left. \left| \widehat{B}\left(\frac{k}{\lambda} + \frac{2\pi n}{\Delta x_0}\right) \right|^2 \right| \\ &\leq C^2 \sum_{n \neq 0} \left| L^p\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \Phi\left(k + \frac{2\pi n \lambda}{\Delta x_0}\right) \right| \\ &\leq C^2 \sum_{n \neq 0} \left| \frac{2\pi n \lambda}{\Delta x_0} \right|^{p\gamma - \beta} \\ &= \lambda^{p\gamma - \beta} C^2 \sum_{n \neq 0} \left| \frac{2\pi n}{\Delta x_0} \right|^{p\gamma - \beta}. \end{aligned} \quad (41)$$

The term inside the sum is independent of λ . Therefore, this bound scales as $O(\lambda^{p\gamma - \beta})$, which we identify with $O(\Delta x^{\beta - p\gamma})$, since $\Delta x = \Delta x_0 / \lambda$. We repeat the argument with δ and obtain a term of order $O(\Delta x^\beta)$. Thus Eq. (39) scales as

$$L^p(k) + O(\Delta x^\beta) + O(\Delta x^{\beta - p\gamma}) = L^p(k) + \boxed{O(\Delta x^{\beta - p\gamma})}, \quad (42)$$

yielding a total time and space error scaling of $O(\Delta t^p) O(\Delta x^{\beta - p\gamma})$. The other O term vanishes because only the term with the worst scaling (lowest power) contributes. The total error scaling in Eq. (43) is determined by the sum of the individual p terms:

$$E \propto \sum_{p=0}^{\alpha} \boxed{O(\Delta t^p \Delta x^{\beta - p\gamma})}, \quad (43)$$

although the error will be bounded by the worst scaling of any of the individual terms. We see from this formula that taking higher orders in Δt decreases the spatial order. This is fine for $L = \partial_x$, because the total order remains the same, but for higher derivatives, the spatial order decreases faster in p than the time order increases. If Δx and $\Delta t \rightarrow 0$ at a proportional rate, then this will decrease the total order, making the overall error scaling worse.

This can be thought of in the following way: If the prior covariance only drops off as some power β , then it is only β times differentiable, so it is not smooth. Taking higher orders in the expansion $\bar{U} = \sum_{p=0}^{\alpha} (\Delta t L)^p / p!$ involves taking derivatives of ever-higher order, and thus at some point the $L\Phi(x)$ term in the difference operator can no longer be calculated. The bin functions do not appear in the above expression, because in the limit of high resolutions, they tend to approximate δ functions, and their exact form becomes irrelevant.

The consequences of this formula deserve some thought, particularly the troubling implication that going to higher orders in time can in fact decrease the quality of the simulation. First, it should be noted that higher-order schemes are not necessarily better, depending on the task. For example, according to the *Godunov theorem* [21, p. 280], higher-order schemes have a tendency to develop spurious oscillations around shocks. It should also be noted that the above formula applies in the high-resolution (and thus high- k) limit. One could conceivably introduce a prior covariance which has a cutoff at high k or perhaps one whose value drops of exponentially with k . An exponentially falling prior covariance would then raise the prospect of a finite-difference scheme with intermediate error scaling; however, the implications of such a scheme are not yet clear.

D. Generalization to higher dimensions

The previous derivation was only presented for the one-dimensional case for the sake of pedagogical clarity. If we extend to the M -dimensional case, then x and k become vectors \vec{x} , \vec{k} , and the simulation domain becomes $\prod_i^M [0, l_i]$. Equation (26) becomes a sum over vectors \vec{k} less than \vec{f}_N where the Nyquist frequency is now a vector due to the (possibly) differing resolutions along each grid direction, and likewise the sum over Brillouin zones is also vector valued.

To prove consistency, the assumptions do not need to be tightened, except that we must now specify $U(\vec{k})\Phi(\vec{k}) \rightarrow 0$ as $\|\vec{k}\| \rightarrow \infty$. We also need the resolution to be increased in all spatial dimensions at the same rate, so Δx becomes $\Delta \vec{x} = \Delta \vec{x}_0 / \lambda$. The proof then proceeds as before.

In order to show that the error scaling formula (43) holds in higher dimensions, one needs to put a new bound on $L(\vec{k})$ such that it is bounded by $\|\vec{k}\|^\gamma$. This property is easily fulfilled by many differential operators, such as $L = \partial_x^2 + \partial_y^2 + \partial_z^2$, for example. Likewise, we assume that $\Phi(k)$ can now be bounded by some $\|\vec{k}\|^\beta$, and the proof proceeds as before.

III. CONCLUSIONS

We have now proved consistency, and found an estimate of the error scaling for IFD schemes, using a set of strong

simplifying assumptions, which we grouped together under the name of a *translation-invariant scheme*. These assumptions were as follows:

- (i) ‘‘Linear case’’ of IFD: linear differential equation, linear measurements with additive noise, and Gaussian prior distribution of the fields.
- (ii) Translation and time invariance of all the above quantities.
- (iii) Box-shaped simulation space with periodic boundary conditions.
- (iv) A response R which integrates the field ϕ against an evenly spaced grid of bin functions.
- (v) A bin functions $B(x)$ which is compactly supported on a strict subset of the simulation space, has bounded Fourier transform, and whose Fourier transform $\widehat{B}(\vec{k})$ has $\widehat{B}(0) \neq 0$.
- (vi) $\bar{U}(\vec{k})\Phi(\vec{k}) \rightarrow 0$ as $\|\vec{k}\| \rightarrow \infty$.

To obtain an estimate of the error scaling, we needed to assume the following:

- (vii) The operators $L(\vec{k})$ and $\Phi(\vec{k})$ may be bounded by power-laws $\|\vec{k}\|^\gamma$ and $\|\vec{k}\|^{-\beta}$ for $\beta, \gamma > 0$ respectively, at large values of $\|\vec{k}\|$.

These restrictions mean that the results in this paper are only directly applicable to a very small subset of the simulation schemes that may be constructed using IFD. Given the immense amount of freedom inherent in the IFD framework, it is doubtful that a general analytic proof of consistency will be achievable. This paper should be instead taken as a general indication that IFD is at least a sensible methodology.

That being said, it is expected that the above assumptions could be weakened in order to obtain a stronger result. Most obviously, the fact that the difference operator can be expressed using a *sum over Brillouin zones* immediately suggests that these results could be extended to a simulation over any periodic lattice of data points, not just rectangular domains.

The restrictions on the bin functions are relatively weak. The compact support requirement simply ensures that the response corresponds to some sort of local measurement of the field. The requirement that $\widehat{B}(0) \neq 0$ deserves some discussion, however. This requirement, rather than being physically motivated, was inserted solely to avoid the occurrence of 0/0 terms in the limit of high resolutions. It may, however, reflect a physical requirement. Take, for example, a bin function $B(x)$ which is everywhere positive and is symmetric and peaked about zero. It will satisfy $\widehat{B}(0) \neq 0$, and in the limit of high resolutions, this bin will approach a δ function and will represent a sample of the field value at that point. In contrast, take $x B(x)$; this function is now odd, and in the limit of high resolutions, this will approach something that represents a point sample of the *derivative* of the field about that point. Attempting to apply IFD to reconstructions of the derivative of the field may give nonsensical results, which is what the $\widehat{B}(0) \neq 0$ requirement may be implying.

Removing the translation-invariance requirements would be extremely desirable but much more difficult. The main reason the Fourier approach was necessary was to allow the inversion of the $(R\Phi R^\dagger)_{ij}$ matrices. With a reasonable smoothness prior, these matrices tend to be relatively local in the spatial indices. However, inversion is a nonlocal problem, which makes the inverses of these matrices dependent on the global geometry of the simulation domain and makes

them very difficult to study analytically. The use of periodic boundary conditions allowed us to sidestep this consideration. Any proof seeking to show consistency and convergence in the nontranslationally invariant case would probably have to use a different approach to what we have done here.

Finally, an extension of these results outside of the linear regime is self-evidently desirable but may present some significant challenges. In particular, IFT inference problems including nonlinear responses and non-Gaussian priors on the fields typically result in a need to calculate Feynmann diagrams. Integrating these into an analytic proof of consistency will be challenging to say the least.

APPENDIX: ILLUSTRATIVE EXAMPLE

The previous derivations using Fourier representations and infinite limits may be too abstract for some readers, so we present an illustrative example which shows how to estimate the error for an example simulation scheme using the formula we derived. We will not show how to actually implement such a scheme in code, as this is already discussed in depth in Refs. [1] and [3].

The simplest possible equation of motion that we can implement is the diffusion equation, $\partial_t \phi(x, t) = D \partial_x^2 \phi(x, t)$, and thus $L = D \partial_x^2$, with D being the diffusion constant. We take a scheme that is first order in time, which means that we have approximated the analytic time-evolution operator U by a first-order expansion $\bar{U} = \mathbb{1} + \Delta t L = \mathbb{1} + D \Delta t \partial_x^2$. A convenient prior for this situation would be $\Phi(k) = 1/(k^4 + m^4)$. This prior suppresses high-frequency modes and has a (somewhat) arbitrary regularization factor m to ensure that the prior is well defined at 0. A simple response would be the box function which we introduced in Eq. (18). We take a one-dimensional periodic interval of length l as the simulation domain.

The first step is to calculate the Fourier transform of the box function:

$$\begin{aligned} \widehat{B}(k) &= \frac{1}{\sqrt{l}} \int_0^l dx B(x) e^{ikx} = \frac{1}{\Delta x \sqrt{l}} \int_0^{\Delta x} dx e^{ikx} \\ &= \frac{1}{\Delta x \sqrt{l}} \frac{(e^{ik\Delta x} - 1)}{ik}. \end{aligned} \quad (\text{A1})$$

Hence,

$$|\widehat{B}(k)|^2 = \frac{2}{\Delta x^2 l} \frac{1 - \cos(k\Delta x)}{k^2}, \quad (\text{A2})$$

which goes to 1 in the limit $\lambda \rightarrow \infty$.

The momentum-space version of the finite-difference operator can now be obtained by substituting all of the quantities we defined into Eq. (26):

$$\begin{aligned} T(k) &= \frac{\sum_{b \in 2f_n \mathbb{Z}} [1 + D(k+b)^2] \frac{1}{(k+b)^4 + m^4} \frac{1 - \cos[(k+b)\Delta x]}{(k+b)^2}}{\sum_{\hat{b} \in 2f_n \mathbb{Z}} \frac{1}{(k+\hat{b})^4} \frac{1 - \cos[(k+\hat{b})\Delta x]}{(k+\hat{b})^2}} \\ &= \mathbb{1} + D \frac{\sum_{b \in 2f_n \mathbb{Z}} \frac{1}{(k+b)^4 + m^4}}{\sum_{\hat{b} \in 2f_n \mathbb{Z}} \frac{1}{[(k+\hat{b})^4 + m^4](k+\hat{b})^2}}. \end{aligned} \quad (\text{A3})$$

We were able to cancel the $1 - \cos(k\Delta x)$ terms by virtue of the fact that $\cos((k+b)\Delta x) = \cos(k\Delta x)$ if b is a multiple of the Nyquist frequency. Thus, they could be pulled outside of the sum over Brillouin zones and canceled in the numerator and denominator.

Now it is time to obtain an estimate of the error scaling using Eq. (43). L is the second derivative, which is k^2 in Fourier space, so $\gamma = 2$. The prior scales with k^{-4} , so $\beta = 4$, and U was expanded to first order in time, so $p = 1$. Inserting these numbers into the formula tells us that the error scales proportionally to $O(\Delta t \Delta x^{4-2}) = O(\Delta t \Delta x^2)$, so the scheme is first order in time and second order in space.

-
- [1] T. A. Enßlin, Information field dynamics for simulation scheme construction, *Phys. Rev. E* **87**, 013308 (2013).
- [2] C. Münch, Mathematical foundation of information field dynamics (revised version) (2014) [arXiv:1412.1226](https://arxiv.org/abs/1412.1226) [math.DS].
- [3] R. H. Leike and T. A. Enßlin, Towards information-optimal simulation of partial differential equations, *Phys. Rev. E* **97**, 033314 (2018).
- [4] P. Hennig, M. A. Osborne, and M. Girolami, Probabilistic numerics and uncertainty in computations, *Proc. R. Soc. Lond. A* **471** (2015).
- [5] P. Diaconis, *Bayesian numerical analysis*, in *Statistical Decision Theory and Related Topics IV*, edited by S. S. Gupta and J. O. Berger (Springer-Verlag, New York, 1988), Vol. 1, pp. 163–175.
- [6] A. O'Hagan, Some bayesian numerical analysis, *Bayesian Stat.* **4**, 4 (1992).
- [7] F. M. Larkin, Gaussian measure in hilbert space and applications in numerical analysis, *Rocky Mount. J. Math.* **2**, 379 (1972).
- [8] J. Skilling, Bayesian solution of ordinary differential equations, in *Maximum Entropy and Bayesian Methods: Seattle, 1991*, edited by C. R. Smith, G. J. Erickson, and P. O. Neudorfer (Springer, Netherlands, Dordrecht, 1992), pp. 23–37.
- [9] I. Coulibaly and C. Lécot, C.: A quasi-randomized runge-kutta method, *Math. Comput.* **68**, 651 (1999).
- [10] G. Stengle, Error analysis of a randomized numerical method, *Num. Math.* **70**, 119 (1995).
- [11] F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic, Probabilistic integration: A role in statistical computation? (2015) [arXiv:1512.00933](https://arxiv.org/abs/1512.00933) [stat.ML].
- [12] M. Schober, D. K. Duvenaud, and P. Hennig, Probabilistic ode solvers with runge-kutta means, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., New York, 2014), pp. 739–747.
- [13] P. R. Conrad, M. Girolami, S. Särkkä, A. Stuart, and K. Zygalakis, Probability measures for numerical solutions of differential equations (2015) [arXiv:1506.04592](https://arxiv.org/abs/1506.04592) [stat.ME].
- [14] C. Archambeau, D. Cornford, M. Opper, and J. Shawe-Taylor, Gaussian process approximations of stochastic differential equations, in *Gaussian Processes in Practice, Proceedings*

- of Machine Learning Research*, Vol. 1, edited by Neil D. Lawrence, Anton Schwaighofer, and Joaquin Quiñero Candela (PMLR, Bletchley Park, UK, 2007), pp. 1–16.
- [15] R. H. Leike and T. A. Enßlin, Optimal belief approximation, *Entropy* **19** (2017).
- [16] T. A. Enßlin, Information theory for fields, (2018) [arXiv:1804.03350](https://arxiv.org/abs/1804.03350).
- [17] T. A. Enßlin, M. Frommert, and F. S. Kitaura, Information field theory for cosmological perturbation reconstruction and nonlinear signal analysis, *Phys. Rev. D* **80**, 105005 (2009).
- [18] J. C. Lemm, Bayesian field theory: Nonparametric approaches to density estimation, regression, classification, and inverse quantum problems, (1999) [arXiv:physics/9912005](https://arxiv.org/abs/physics/9912005).
- [19] P. D. Lax and R. D. Richtmyer, Survey of the stability of linear finite difference equations, *Commun. Pure Appl. Math.* **9**, 267 (1956).
- [20] R. Corless and N. Fillion, *A Graduate Introduction to Numerical Methods: From the Viewpoint of Backward Error Analysis*, (Springer, New York, 2013).
- [21] C. Hirsch, *Numerical Computation of Internal and External Flows: Fundamentals of Computational Fluid Dynamics* (Elsevier/Butterworth-Heinemann, Amsterdam, 2007).
- [22] M. Reed and B. Simon, *Methods of Modern Mathematical Physics*, Vol. 2 (Academic Press, New York, 1975).