



Computing Lagrangian means

Hossein A. Kafiabad^{1,†} and Jacques Vanneste²

¹Department of Mathematical Sciences, Durham University, Durham DH1 3LE, UK

²School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, Edinburgh EH9 3FD, UK

(Received 2 August 2022; revised 10 March 2023; accepted 14 March 2023)

Lagrangian averaging plays an important role in the analysis of wave–mean-flow interactions and other multiscale fluid phenomena. The numerical computation of Lagrangian means, e.g. from simulation data, is, however, challenging. Typical implementations require tracking a large number of particles to construct Lagrangian time series, which are then averaged using a low-pass filter. This has drawbacks that include large memory demands, particle clustering and complications of parallelisation. We develop a novel approach in which the Lagrangian means of various fields (including particle positions) are computed by solving partial differential equations (PDEs) that are integrated over successive averaging time intervals. We propose two strategies, distinguished by their spatial independent variables. The first, which generalises the algorithm of Kafiabad (*J. Fluid Mech.*, vol. 940, 2022, A2), uses end-of-interval particle positions; the second uses directly the Lagrangian mean positions. The PDEs can be discretised in a variety of ways, e.g. using the same discretisation as that employed for the governing dynamical equations, and solved on-the-fly to minimise the memory footprint. We illustrate the new approach with a pseudo-spectral implementation for the rotating shallow-water model. Two applications to flows that combine vortical turbulence and Poincaré waves demonstrate the superiority of Lagrangian averaging over Eulerian averaging for wave–vortex separation.

Key words: internal waves, computational methods, waves in rotating fluids

1. Introduction

Time averaging is a basic yet essential tool in fluid dynamics because of the ubiquity of phenomena involving multiple time scales. Atmospheric and oceanic flows, for

† Email address for correspondence: hossein.amini-kafiabad@durham.ac.uk

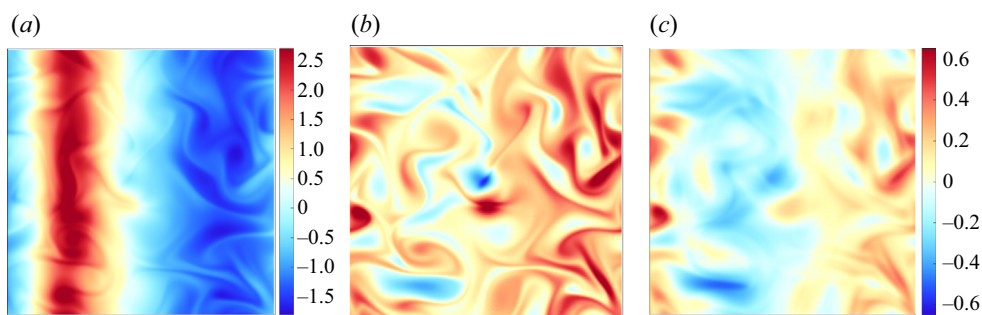


Figure 1. Vorticity field and its Lagrangian and Eulerian means in the shallow-water simulation of § 4.2.1: (a) instantaneous vorticity, (b) Lagrangian mean vorticity, and (c) Eulerian mean vorticity. Both mean fields share the same averaging period, corresponding to 3.6 wave periods. Panels (b) and (c) share the same colour bar, shown to the right of (c).

instance, can often be decomposed usefully into fast and slow components (e.g. Vanneste 2013; Vallis 2017). Time averaging is then used to interpret numerical simulation and observational data, and to assimilate the latter in ocean and weather models. More broadly, it can be argued that all numerical simulations of time-dependent flows rely implicitly on time averaging, since they do not resolve phenomena on time scales shorter than the time step.

Time averaging can be performed in different ways. The most straightforward way is to average the time series of flow variables at fixed spatial positions to obtain the so-called Eulerian mean. An alternative is to average flow variables along particle trajectories, that is, at fixed particle labels instead of fixed positions, to obtain the Lagrangian mean. There are several practical and conceptual reasons to view Lagrangian averaging as superior to Eulerian averaging.

From a practical viewpoint, the advection of fast motion (often consisting of waves) by slowly evolving flow and vice versa adversely affect their Eulerian decomposition. For instance, a strong, slowly evolving flow ‘Doppler shifts’ the frequency of fast waves. This can lead to wave frequencies observed at fixed spatial locations that are much smaller than the intrinsic frequency. It therefore obscures the time scale separation between waves and flow. Lagrangian averaging resolves this issue, as demonstrated in several studies (Nagai *et al.* 2015; Shakespeare & Hogg 2017; Bachman *et al.* 2020; Shakespeare *et al.* 2021; Jones *et al.* 2022). Conversely, the advection of a slow background flow by waves leads to flow features that are blurred by Eulerian averaging but not by Lagrangian averaging. We illustrate this in figure 1, showing the vorticity field in a simulation of a turbulent rotating shallow-water flow interacting with a mode-1 Poincaré wave. The instantaneous vorticity in figure 1(a) is dominated by the high-amplitude wave. Both the Lagrangian and Eulerian means (figures 1b,c) filter out the wave, but the Eulerian mean blurs out fine vorticity structures that are well resolved by the Lagrangian mean (details of this simulation are presented in § 4.2.1).

From a conceptual viewpoint, Lagrangian averaging provides a powerful tool to study wave–mean-flow interactions. This is because the material conservation of key fields (scalar concentrations, vorticity vector, circulation, potential vorticity) is inherited naturally by the corresponding Lagrangian mean fields. As a result, the Lagrangian mean of the dynamical equations is often simpler and more meaningful than the Eulerian mean (Bretherton 1971; Soward 1972; Andrews & McIntyre 1978; Bühler 2014; Gilbert

& Vanneste 2018). Relatedly, the Lagrangian mean emerges naturally in the asymptotic derivation of wave-averaged models (Grimshaw 1975; Wagner & Young 2015). A striking example of the dynamical relevance of the Lagrangian mean is provided by the observation that geostrophic balance, the dominant balance in rapidly rotating flows, continues to hold in the presence of strong waves provided that it is formulated in terms of Lagrangian mean velocity and pressure instead of their instantaneous or Eulerian mean values (Moore 1970; Bühler & McIntyre 1998; Kafiabad, Vanneste & Young 2021; Kafiabad 2022).

Despite its advantages, Lagrangian averaging is not used widely as a practical tool, mainly because Lagrangian means are difficult to compute numerically. Most numerical models are intrinsically Eulerian and provide the fields of interest at fixed spatial locations, typically grid points. The standard approach for the computation of Lagrangian means is then to seed a large number of passive particles in the flow, track them (forwards or backwards in time) using interpolated velocities, and apply time averaging to the resulting Lagrangian time series (e.g. Nagai *et al.* 2015; Shakespeare & Hogg 2017, 2018, 2019; Bachman *et al.* 2020; Shakespeare *et al.* 2021; Jones *et al.* 2022). This has a high computational cost, requires a large memory allocation, suffers from possible particle clustering, and, as discussed in Kafiabad (2022), is difficult to parallelise efficiently (see Shakespeare *et al.* (2021) for a parallel implementation).

To circumvent the difficulties of particle tracking, Kafiabad (2022) developed a grid-based method that computes the Lagrangian mean directly on an Eulerian grid, building the mean through a time step iteration carried out over successive averaging intervals. By eliminating the need to compute explicit particle trajectories, the method reduces memory demands and simplifies integration into parallelised numerical models. The present paper starts with the recognition that the algorithm of Kafiabad (2022) is a particular discretisation of a partial differential equation (PDE) governing the evolution of what we term a partial Lagrangian mean, that is, the mean carried out only up to some intermediate time in the averaging interval. We formulate this PDE using the position of particles at the intermediate time as independent spatial variable, as in Kafiabad (2022). The (total) Lagrangian mean is then obtained by taking the intermediate time to be the end of the averaging interval.

In this form, the Lagrangian mean does not match the Andrews & McIntyre (1978) definition of the generalised Lagrangian mean (GLM): this requires the mean fields to be expressed as functions of the mean position of fluid particles. To achieve this, it is necessary to relate the mean positions of particles to their positions at the end of the averaging interval, and to carry out a remapping of the Lagrangian mean fields. This constitutes our strategy 1 for the computation of GLMs. We show that the algorithm of Kafiabad (2022) amounts to a semi-Lagrangian discretisation of the PDEs of strategy 1. We propose an alternative strategy, strategy 2, which formulates PDEs directly for the partial Lagrangian means using the mean position as independent spatial variable. The PDEs involved in both strategies can be solved by broad classes of numerical methods: finite differences, finite volumes, finite elements or spectral methods. We illustrate this with a pseudo-spectral Fourier implementation for a shallow-water flow in a doubly periodic domain.

The paper is structured as follows. We introduce notation and define the Lagrangian means in § 2. We derive the PDEs of the two strategies in § 3. We discuss their numerical implementation and present two applications to shallow-water flows in § 4. The choice of strategies, their advantages and costs are discussed in § 5. Technical aspects, including the averaging of tensorial fields, are relegated to appendices.

2. Formulation

We consider fluid motion in a two- or three-dimensional Euclidean space. We denote the flow map by φ , with $\varphi(\mathbf{a}, t) \in \mathbb{R}^2$ or \mathbb{R}^3 the position at time t of a particle identified by its label \mathbf{a} (which can be taken as the position at $t = 0$). The flow map and velocity field are related by

$$\partial_t \varphi(\mathbf{a}, t) = \mathbf{u}(\varphi(\mathbf{a}, t), t). \tag{2.1}$$

Lagrangian averaging is averaging at fixed particle label \mathbf{a} , in contrast with Eulerian averaging, which fixes the spatial position. Both can involve different types of means: temporal, spatial or – as often used in theoretical work – ensemble mean. Here we focus on a straightforward time average, of the form

$$\bar{g}(\tau) = \frac{1}{T} \int_{\tau}^{\tau+T} g(s) \, ds, \tag{2.2}$$

when applied to a function $g(t)$ that depends on time only. Equation (2.2) introduces the notation τ for the time at which the averaging is carried out, and T for the averaging period. Usually, the middle of the averaging interval $\tau + T/2$ is used as an argument for the mean function, but we prefer to adopt τ (the beginning of the averaging interval) to simplify the notation in the upcoming derivations. A simple shift in time switches from one convention to the other. A weight function could be inserted into the integrand of (2.2) to generalise the definition of the average; this would lead to minimal changes in what follows.

The Lagrangian mean trajectory associated with (2.2) is represented by the Lagrangian mean map $\bar{\varphi}^L$ defined by

$$\bar{\varphi}^L(\mathbf{a}, \tau) := \frac{1}{T} \int_{\tau}^{\tau+T} \varphi(\mathbf{a}, s) \, ds. \tag{2.3}$$

Thus $\bar{\varphi}^L(\mathbf{a}, \tau)$ returns the mean position from τ to $\tau + T$ of the particle labelled by \mathbf{a} . The definition (2.3) makes sense in \mathbb{R}^n , when $\bar{\varphi}^L$ can be interpreted as a vector and averaged componentwise, but not on other manifolds where more complicated definitions are necessary (Gilbert & Vanneste 2018). The (generalised) Lagrangian mean of a scalar function $f(\mathbf{x}, t)$ is then defined by

$$\bar{f}^L(\bar{\varphi}^L(\mathbf{a}, \tau), \tau) := \frac{1}{T} \int_{\tau}^{\tau+T} f(\varphi(\mathbf{a}, s), s) \, ds. \tag{2.4}$$

Hence $\bar{f}^L(\mathbf{x}, \tau)$ is the average of f along the trajectory of the fluid parcel, regarded as a function of the Lagrangian mean position \mathbf{x} and time τ .

Our aim is the development of an efficient numerical method for the computation of \bar{f}^L that relies on solving PDEs, which can be discretised in a variety of ways, rather than on tracking ensembles of particle trajectories. We propose two strategies and derive the corresponding PDEs in the next section.

3. Two strategies

Following Kafiabad (2022), we introduce another representation of the Lagrangian mean via tilde functions defined by

$$\begin{aligned} \tilde{f}^L(\varphi(\mathbf{a}, \tau + T), \tau) &:= \frac{1}{T} \int_{\tau}^{\tau+T} f(\varphi(\mathbf{a}, s), s) \, ds \\ &= \bar{f}^L(\bar{\varphi}^L(\mathbf{a}, \tau), \tau). \end{aligned} \tag{3.1}$$

Computing Lagrangian means

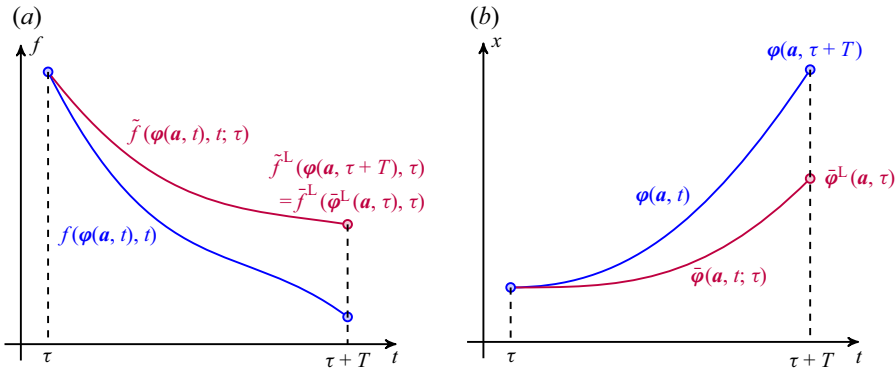


Figure 2. Lagrangian means and partial Lagrangian means in an averaging interval $(\tau, \tau + T)$: (a) means of a function f evaluated along the trajectory $\varphi(\mathbf{a}, t)$ of a fluid parcel labelled by \mathbf{a} ; and (b) means of the position, shown here along a single coordinate axis.

Comparing this definition with (2.4) shows that the overbar indicates a mean along a trajectory identified by the Lagrangian mean position, whereas the tilde indicates the same mean but with the trajectory identified by the actual position at $t = \tau + T$, that is, at the end of the averaging. This is illustrated in figure 2. We also introduce the partial mean versions of (2.3), (2.4) and (3.1), namely

$$\bar{\varphi}(\mathbf{a}, t; \tau) := \frac{1}{t - \tau} \int_{\tau}^t \varphi(\mathbf{a}, s) \, ds, \quad (3.2a)$$

$$\bar{f}(\bar{\varphi}(\mathbf{a}, t; \tau), t; \tau) := \frac{1}{t - \tau} \int_{\tau}^t f(\varphi(\mathbf{a}, s), s) \, ds, \quad (3.2b)$$

$$\tilde{f}(\varphi(\mathbf{a}, t), t; \tau) := \frac{1}{t - \tau} \int_{\tau}^t f(\varphi(\mathbf{a}, s), s) \, ds, \quad (3.2c)$$

as in Kafiabad (2022); see figure 2. Clearly, the partial means give the total means when evaluated at $t = \tau + T$. We emphasise that all means used in the paper are Lagrangian means, and that we indicate this explicitly by a superscript L only for the total means, to distinguish them from the partial means, which are undecorated. The counterpart of (3.1) holds for the partial means:

$$\tilde{f}(\varphi(\mathbf{a}, t), t; \tau) = \bar{f}(\bar{\varphi}(\mathbf{a}, t; \tau), t; \tau). \quad (3.3)$$

Since time-averaged quantities vary over time scales larger than the averaging period T , it is neither necessary nor desirable to compute Lagrangian means at each of the times at which the velocity \mathbf{u} and scalar field f are known, typically discrete times separated by a small time step. Rather, we think of the averaging time τ as a slow variable and propose to compute the Lagrangian means only at $\tau = \tau_n = nT$ for $n = 0, 1, 2, \dots$. We can therefore carry out independent computations for each t_n , each involving only the fields for $t \in (\tau_n, \tau_n + T)$. We now focus on one such interval, and to lighten the notation, drop the parametric dependence on τ from the partial means in (3.2). For instance, we use $\bar{f}(\bar{\varphi}(\mathbf{a}, t), t)$ instead of $\bar{f}(\bar{\varphi}(\mathbf{a}, t; \tau), t; \tau)$ in the following derivations, keeping in mind the now implicit dependence of f and $\bar{\varphi}$ on τ . A warning about notation might be useful: in what follows, we use the symbol x as a generic dummy variable, without attributing it the specific meaning of either an actual or a Lagrangian mean position. This meaning

is determined by the function in which x appears. Thus x in $\tilde{f}(x, t)$ is interpreted as an actual position at time t , whereas in $\tilde{f}^L(x, t)$ it is interpreted as a (partial) Lagrangian mean position.

We now formulate two distinct strategies for the computation of the Lagrangian mean $\tilde{f}^L(x, \tau)$.

3.1. Strategy I

Our first strategy parallels that of Kafiabad (2022) and consists in solving a PDE for $\tilde{f}(x, t)$, evaluating the result at $t = \tau + T$ to obtain $\tilde{f}^L(x, \tau)$, then deducing $\tilde{f}^L(x, \tau)$ by applying a suitable remapping. To derive the PDE for $\tilde{f}(x, t)$, we take the time derivative of (3.2c) at fixed label a and use the chain rule and (2.1) to find

$$\partial_t \tilde{f}(\varphi(a, t), t) + \mathbf{u}(\varphi(a, t), t) \cdot \nabla \tilde{f}(\varphi(a, t), t) = \frac{1}{t - \tau} (f(\varphi(a, t), t) - \tilde{f}(\varphi(a, t), t)). \tag{3.4}$$

Here and in what follows, the gradient ∇ is taken with respect to the first argument of \tilde{f} . Replacing $\varphi(a, t)$ by x as independent variable yields the sought PDE:

$$\partial_t \tilde{f}(x, t) + \mathbf{u}(x, t) \cdot \nabla \tilde{f}(x, t) = \frac{f(x, t) - \tilde{f}(x, t)}{t - \tau}, \tag{3.5}$$

which can be integrated from τ to $t = \tau + T$ to find the total mean $\tilde{f}^L(x, \tau) = \tilde{f}(x, \tau + T; \tau)$. This is a forced advection equation in which the forcing can be interpreted as a time-dependent relaxation of \tilde{f} to f . In a bounded domain, the solution of (3.5) requires no boundary conditions since the differentiation $\mathbf{u}(x, t) \cdot \nabla$ is along the boundary. The initial condition is that $\tilde{f}(x, \tau) = f(x, \tau)$ so that the right-hand side is finite.

Computing \tilde{f}^L may be all that is needed for applications in which the spatial distribution of the Lagrangian mean is not important. For example, wave-averaged geostrophy – the modified form of geostrophic balance expressed in terms of Lagrangian mean quantities – can be validated by comparing Lagrangian mean velocity and pressure gradient at the same position (equivalently, the same label) regardless of where the position is located in physical space (Kafiabad *et al.* 2021; Kafiabad 2022).

However, in many other applications, it is necessary to compute the (generalised) Lagrangian mean field $\tilde{f}^L(x, t)$ for specified Lagrangian mean positions x . This (generalised) Lagrangian mean field can be deduced from $\tilde{f}^L(x, t)$ by considering the lift map $\mathcal{E}(x, t)$, which returns the position at time t of the particle with x as partial mean position from \bar{t} to t , that is,

$$\mathcal{E}(\bar{\varphi}(a, t), t) = \varphi(a, t) \tag{3.6}$$

(Andrews & McIntyre 1978; Bühler 2014). Its inverse $\mathcal{E}^{-1}(x, t)$ returns the partial mean position of the particle that passes through x at t :

$$\mathcal{E}^{-1}(\varphi(a, t), t) = \bar{\varphi}(a, t) = \frac{1}{t - \tau} \int_{\tau}^t \varphi(a, s) ds, \tag{3.7}$$

where we used the definition (3.2a) of the mean position. The map \mathcal{E} and its inverse \mathcal{E}^{-1} are depicted in figure 3. The relation (3.1) between \tilde{f}^L and f^L can then be written in terms of \mathcal{E}^{-1} as

$$\tilde{f}^L(x, \tau) = f^L(\mathcal{E}^{-1}(x, \tau + T), \tau). \tag{3.8}$$

Computing Lagrangian means

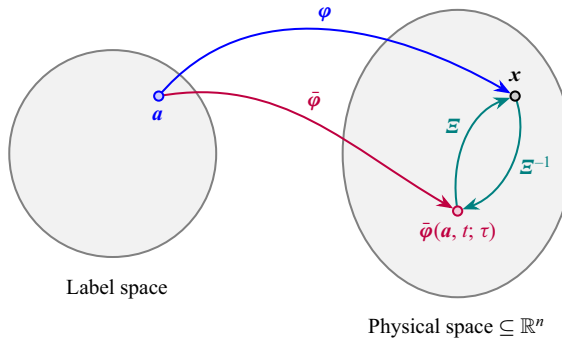


Figure 3. Flow map φ , Lagrangian (partial) mean map $\bar{\varphi}$, lift map Ξ and its inverse Ξ^{-1} .

Now, comparing (3.7) with (3.2c) makes it clear that the components of Ξ^{-1} can be viewed as instances of functions \tilde{f} with $f(x) = x_i$. Thus we can rewrite (3.5) for this special case to obtain

$$\partial_t \Xi^{-1}(x, t) + \mathbf{u}(x, t) \cdot \nabla \Xi^{-1}(x, t) = \frac{x - \Xi^{-1}(x, t)}{t - \tau}. \quad (3.9)$$

Alternatively, we can take the time derivative of (3.7) and replace $\varphi(a, t)$ by x to arrive at (3.9). Integrating (3.9) provides the means to effect the remapping (3.8) between \tilde{f}^L and \tilde{f}^L .

To recapitulate, our first strategy consists in solving the PDEs (3.5) and (3.9) from τ to $\tau + T$ to obtain $\tilde{f}^L(x, \tau) = \tilde{f}(x, \tau + T; \tau)$ and $\Xi(x, \tau + T; \tau)$, then using (3.8) to compute \tilde{f}^L by interpolation. The algorithm proposed by Kafiabad (2022) turns out to be a particular discretisation of this strategy (see Appendix B).

3.2. Strategy 2

Our second strategy bypasses the use of \tilde{f} and is instead based on PDEs for \bar{f} and Ξ . To derive these, we first note that taking the time derivative of (3.2a) gives

$$\partial_t \bar{\varphi}(a, t) = \frac{\varphi(a, t) - \bar{\varphi}(a, t)}{t - \tau} =: \bar{\mathbf{u}}(\bar{\varphi}(a, t), t). \quad (3.10)$$

The second equality defines the auxiliary velocity field $\bar{\mathbf{u}}$ as the time derivative of the partial Lagrangian mean position. Using (3.6), this velocity field can be written in terms of the lift map as

$$\bar{\mathbf{u}}(x, t) = \frac{\Xi(x, t) - x}{t - \tau}, \quad (3.11)$$

where the dummy variable x can be thought of as the partial mean position. We emphasise that $\bar{\mathbf{u}}(\cdot, t)$, like $\Xi(\cdot, t)$, depends implicitly on τ , and warn that it should not be interpreted as the partial mean of the Lagrangian velocity: as discussed in Appendix A, its value at the end of the averaging interval, for $t = \tau + T$, differs from the usual Lagrangian mean velocity, that is, the time derivative of $\bar{\varphi}^L(a, \tau)$ with respect to τ .

Now, differentiating (3.2b) with respect to t at fixed label \mathbf{a} and using (3.10) leads to

$$\partial_t \bar{f}(\bar{\varphi}(\mathbf{a}, t), t) + \bar{\mathbf{u}}(\bar{\varphi}(\mathbf{a}, t), t) \cdot \nabla \bar{f}(\bar{\varphi}(\mathbf{a}, t), t) = \frac{f(\boldsymbol{\varphi}(\mathbf{a}, t), t) - \bar{f}(\bar{\varphi}(\mathbf{a}, t), t)}{t - \tau}. \quad (3.12)$$

We obtain the desired PDE for $\bar{f}(\mathbf{x}, t)$ by using (3.6) and replacing $\bar{\varphi}(\mathbf{a}, t)$ by the independent variable \mathbf{x} to write

$$\partial_t \bar{f}(\mathbf{x}, t) + \bar{\mathbf{u}}(\mathbf{x}, t) \cdot \nabla \bar{f}(\mathbf{x}, t) = \frac{f(\boldsymbol{\Xi}(\mathbf{x}, t), t) - \bar{f}(\mathbf{x}, t)}{t - \tau}. \quad (3.13)$$

This is a forced advection equation, analogous to the PDE (3.5) governing $\tilde{f}(\mathbf{x}, t)$. However, unlike (3.5), it is not closed since it involves $\boldsymbol{\Xi}(\mathbf{x}, t)$, explicitly on the right-hand side and implicitly through $\bar{\mathbf{u}}(\mathbf{x}, t)$ on the left-hand side. It needs to be solved along an equation for $\boldsymbol{\Xi}(\mathbf{x}, t)$. We derive this equation by taking the time derivative of (3.6) at fixed \mathbf{a} and using (3.10) and (2.1) to obtain

$$\partial_t \boldsymbol{\Xi}(\bar{\varphi}(\mathbf{a}, t), t) + \bar{\mathbf{u}}(\bar{\varphi}(\mathbf{a}, t), t) \cdot \nabla \boldsymbol{\Xi}(\bar{\varphi}(\mathbf{a}, t), t) = \mathbf{u}(\boldsymbol{\varphi}(\mathbf{a}, t), t). \quad (3.14)$$

Hence, replacing $\bar{\varphi}(\mathbf{a}, t)$ by \mathbf{x} and using (3.6),

$$\partial_t \boldsymbol{\Xi}(\mathbf{x}, t) + \bar{\mathbf{u}}(\mathbf{x}, t) \cdot \nabla \boldsymbol{\Xi}(\mathbf{x}, t) = \mathbf{u}(\boldsymbol{\Xi}(\mathbf{x}, t), t). \quad (3.15)$$

Strategy 2 consists in solving (3.13) and (3.15), with $\bar{\mathbf{u}}(\mathbf{x}, t)$ defined by (3.11), for $t \in (\tau, \tau + T)$, then deducing the Lagrangian mean of f as $\bar{f}^L(\mathbf{x}, \tau) = \bar{f}(\mathbf{x}, \tau + T; \tau)$.

The initial conditions for (3.13) and (3.15) are that $\bar{f}(\mathbf{x}, \tau) = f(\mathbf{x}, \tau)$ and $\boldsymbol{\Xi}(\mathbf{x}, \tau) = \mathbf{x}$. The boundary conditions are non-trivial: in bounded domains, $\bar{f}(\mathbf{x}, t)$ and $\boldsymbol{\Xi}(\mathbf{x}, t)$ are defined on the image of the label space by the Lagrangian mean map $\bar{\varphi}$ (equivalently, the image of the fluid domain by $\boldsymbol{\Xi}^{-1}$). Thus the problem in principle involves a boundary moving with velocity $\bar{\mathbf{u}}$, and can therefore be difficult to discretise. The common situation where the physical domain has boundaries that coincide with constant-coordinate surfaces (curves in two dimensions) is straightforward, however, because the componentwise definition of $\bar{\varphi}$ in (2.3) ensures that it maps such boundaries to themselves so the domain remains fixed. The case of periodic domains is also straightforward.

4. Numerical implementation

The set of equations for each strategy of the previous section can be discretised in a variety of ways. Here we focus on a pseudo-spectral discretisation, which we apply to the computation of Lagrangian means in a turbulent shallow-water flow interacting with a Poincaré wave. We make general remarks about the choice of strategy and numerical discretisation, but leave a more complete analysis of numerical error and convergence for future studies.

4.1. Strategy 1

To solve (3.5), it is convenient to introduce $\tilde{g}(\mathbf{x}, t) = (t - \tau)\tilde{f}(\mathbf{x}, t)$, leading to

$$(\partial_t + \mathbf{u}(\mathbf{x}, t) \cdot \nabla) \tilde{g}(\mathbf{x}, t) = f(\mathbf{x}, t). \quad (4.1)$$

Integrating this equation over the averaging period then yields the Lagrangian mean $\bar{f}^L(\mathbf{x}, \tau) = \tilde{g}(\mathbf{x}, \tau + T)/T$. For simple geometries, periodic in particular, standard pseudo-spectral methods provide efficient solvers for (4.1), and if the remapping to

Lagrangian mean positions is desired, (3.9). This is particularly convenient if the fluid's governing equations are also solved pseudo-spectrally, because f is then available on spectral grid points to evaluate the right-hand sides of (4.1) and (3.9), and on physical grid points to evaluate the nonlinear terms. An alternative is a semi-Lagrangian discretisation, which leads to the algorithm of Kafiabad (2022) as detailed in Appendix B.

To investigate the validity of numerical solutions of (3.5), we consider a two-dimensional incompressible inviscid flow for which the vorticity, ζ say, is conserved materially. This implies that

$$\zeta(\mathbf{x}, \tau + T) = \tilde{\zeta}^L(\mathbf{x}, \tau). \quad (4.2)$$

Hence we can calculate $\tilde{\zeta}^L$ by integrating (3.5) (or (4.1)) from τ to $\tau + T$, and compare it with the instantaneous vorticity at $\tau + T$ to study the accuracy of the computed Lagrangian mean. Note that this is simply a test for (3.5) using the material conservation of ζ as opposed to an application of the Lagrangian mean. As mentioned earlier, (3.5) should usually be solved in tandem with (3.9) to get a meaningful spatial distribution of Lagrangian mean quantities.

We perform a numerical simulation of this two-dimensional flow without viscosity in a doubly periodic, $[0, 2\pi]^2$ domain, using a standard pseudo-spectral discretisation and 2/3 de-aliasing with 128^2 grid points. We start the simulation with the vorticity

$$\zeta(x, y, t = 0) = e^{-(x-\pi+0.1)^2-(y-\pi+\pi/3)^2} + e^{-(x-\pi-0.1)^2-(y-\pi-\pi/3)^2}, \quad (4.3)$$

corresponding to two like-signed vortices that subsequently merge. We use Heun's method for the time integration of the governing vorticity equation and for (4.1), with time step $\Delta t = 0.005$. Figure 4 displays the instantaneous vorticity ζ at $t = 25$, and $\tilde{\zeta}^L$ obtained for $\tau = 0$ and $T = 25$. As expected from (4.2), the Lagrangian mean $\tilde{\zeta}^L$ matches the instantaneous vorticity ζ . The pseudo-spectral solution for $\tilde{\zeta}^L$, shown in figure 4(d), in particular, shows an excellent agreement with ζ in figure 4(a). The results of the semi-Lagrangian algorithm of Kafiabad (2022) with, respectively, linear and cubic interpolations, are shown in figures 4(b,c) (see Appendix B). These show a poorer agreement with figure 4(a), especially with the linear interpolation, because of an accumulation of interpolation errors. The computation reported in figure 4 is, however, rather extreme in both the coarseness of the resolution and the length of the averaging interval. We have confirmed that the three numerical solutions for $\tilde{\zeta}^L$ converge to each other and to ζ as the spatial resolution increases or the length of the averaging interval decreases. (The interested reader will find a demonstration of this convergence in the supplementary material available at <https://doi.org/10.1017/jfm.2023.228>.) Since we solve the dynamical and Lagrangian mean equations without explicit dissipation, both require small time steps. Our investigation for this particular example shows that the Lagrangian mean equations are less restrictive than the dynamical equations for the size of time step.

The pseudo-spectral method leads to more accurate results, but it is not as stable as its semi-Lagrangian counterpart and therefore requires smaller time steps. The difference arises because the implicit time integration and numerical smoothing due to interpolation that are inherent to semi-Lagrangian methods have a stabilising effect.

4.2. Strategy 2

We now implement strategy 2, which uses the Lagrangian mean position \mathbf{x} as independent spatial variable. In the periodic domain that we consider, there are no

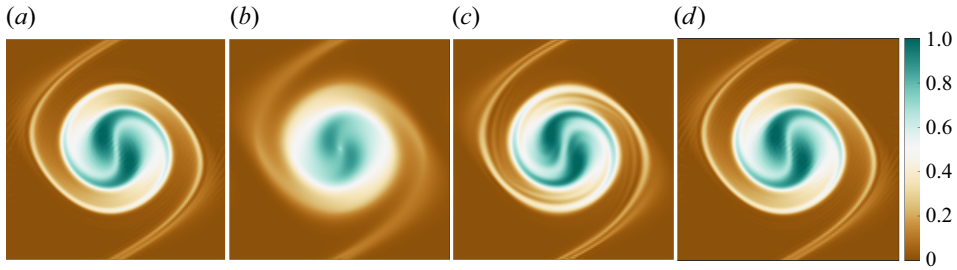


Figure 4. Vorticity field and its Lagrangian mean for a two-dimensional incompressible inviscid flow: (a) instantaneous vorticity ζ at $t = 25$. (b–d) Computation of ζ^L by time integration of (3.5) from $\tau = 0$ to $T = 25$, using (b) a semi-Lagrangian discretisation with linear interpolation, (c) a semi-Lagrangian discretisation with cubic interpolation, and (d) pseudo-spectral discretisation.

difficulties associated with moving boundaries, and a pseudo-spectral discretisation is straightforward. It is convenient to rewrite the PDEs to be integrated, (3.13) and (3.15), in terms of the displacement map

$$\xi(\mathbf{x}, t) = \mathcal{E}(\mathbf{x}, t) - \mathbf{x}, \tag{4.4}$$

since ξ is periodic, unlike \mathcal{E} . (This is the partial mean analogue of the displacement introduced by Andrews & McIntyre 1978.) As in strategy 1, it is also convenient to solve for $\bar{g}(\mathbf{x}, t) = (t - \tau)\bar{f}(\mathbf{x}, t)$ instead of \bar{f} . With these transformations, $\bar{\mathbf{u}} = \xi/(t - \tau)$, and (3.13) and (3.15) are rewritten as

$$\partial_t \xi(\mathbf{x}, t) + \frac{\xi(\mathbf{x}, t)}{t - \tau} \cdot \nabla \xi(\mathbf{x}, t) = \mathbf{u}(\mathbf{x} + \xi(\mathbf{x}, t), t) - \frac{\xi(\mathbf{x}, t)}{t - \tau}, \tag{4.5a}$$

$$\partial_t \bar{g}(\mathbf{x}, t) + \frac{\xi(\mathbf{x}, t)}{t - \tau} \cdot \nabla \bar{g}(\mathbf{x}, t) = f(\mathbf{x} + \xi(\mathbf{x}, t), t). \tag{4.5b}$$

These are the PDEs that we solve numerically. When discretising in time, we found it beneficial for stability to first update ξ using (4.5a), then use the updated ξ for the time integration of (4.5b).

Below, we apply strategy 2 to two examples of rotating shallow-water flows. We write the governing equations in a non-dimensional form, using a characteristic length L , characteristic velocity U , time L/U and mean height H for scaling, leading to

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \frac{1}{Ro} \hat{\mathbf{z}} \times \mathbf{u} = -\frac{1}{Fr^2} \nabla h + \frac{1}{Re} \nabla^2 \mathbf{u}, \tag{4.6a}$$

$$\frac{\partial h}{\partial t} + \nabla \cdot (h\mathbf{u}) = 0, \tag{4.6b}$$

where we introduce the standard dimensionless numbers (e.g. Vallis 2017)

$$Ro = \frac{U}{fL}, \quad Fr = \frac{U}{\sqrt{gH}} \quad \text{and} \quad Re = \frac{UL}{\nu}. \tag{4.7a-c}$$

In the above, g is the gravitational acceleration, ν is the kinematic viscosity, f is the Coriolis parameter, and $\hat{\mathbf{z}}$ is the vertical unit vector.

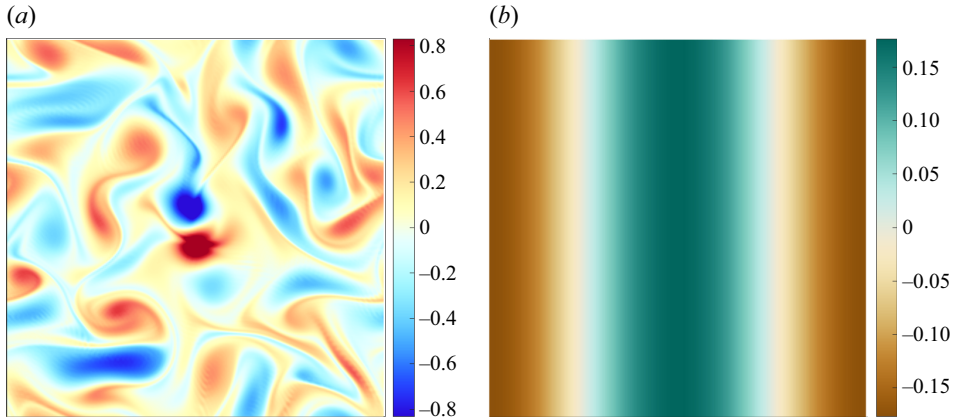


Figure 5. Initial condition of the shallow-water simulation: (a) vertical vorticity of the geostrophic flow; and (b) height field of the Poincaré wave.

4.2.1. Interaction of a turbulent geostrophic flow with a strong wave

In the first example, we compute Lagrangian means in a simulation of a turbulent flow interacting with a Poincaré wave in rotating shallow water. We initialise our simulation with a turbulent flow that is initially in geostrophic balance, with vorticity $\zeta = \partial_x v - \partial_x u$ shown in figure 5(a), and superimpose a mode-1 Poincaré wave, with the height field shown in figure 5(b). The initial geostrophic flow is generated from the output of an incompressible two-dimensional Navier–Stokes simulation, which has reached a fully-developed turbulent state, and the height field that is in geostrophic balance with this vortical flow. We use the root-mean-square velocity of the geostrophic flow as the characteristic velocity U , and choose the length scale of the first Fourier mode as characteristic length L . This makes the dimensionless doubly periodic domain $[0, 2\pi]^2$, which we discretise with 256×256 grid points. The right-travelling mode-1 Poincaré wave has the form

$$u' = a \cos(x - \omega t), \quad v' = \frac{a}{\omega Ro} \sin(x - \omega t), \quad h' = \frac{a}{\omega} \cos(x - \omega t), \quad (4.8a-c)$$

where the intrinsic frequency is $\omega = (Ro^{-2} + Fr^{-2})^{1/2}$, and the velocity amplitude a is a constant, taken as $a = -1.8$ in our simulation. This implies wave velocities that are almost twice as large as the geostrophic velocities. It is in this sense that we regard the wave as strong. We set the dimensionless parameters to

$$Ro = 0.1, \quad Fr = 0.5 \quad \text{and} \quad Re = 3.84 \times 10^3, \quad (4.9a-c)$$

which results in $\omega = 10.2$.

We evaluate (4.8a–c) at $t = 0$ and add the wave fields to the geostrophic field to form the initial condition. We solve the dynamical equations (4.6) in tandem with the Lagrangian mean equations (4.5) over a single averaging time interval, taken to be $T = 2.2$, corresponding to approximately 3.6 wave periods. We use a de-aliased pseudo-spectral discretisation and a forward Euler integrator, with time step 1.25×10^{-4} for (4.6), and 2.5×10^{-4} for (4.5). A bilinear interpolation is used to evaluate \mathbf{u} and f at $\mathbf{x} + \boldsymbol{\xi}(\mathbf{x}, t)$ in the right-hand sides of (4.5). The link for the scripts and data used to produce the results of this section is provided in the ‘Data availability statement’ at the end of this paper.

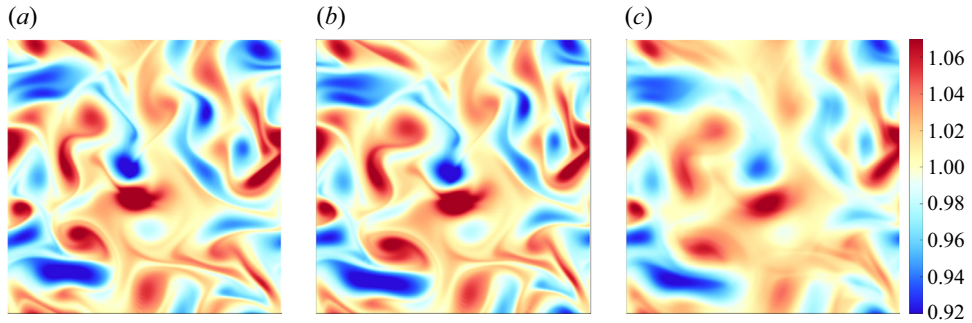


Figure 6. Potential vorticity (PV) q and its Lagrangian and Eulerian means in the shallow-water simulation: (a) instantaneous PV q at $t = 1.1$, (b) Lagrangian mean \bar{q}^L , and (c) Eulerian mean PV. The mean fields are averaged from $\tau = 0$ to $T = 2.2$. All panels share the same colour bar.

Figure 1, discussed briefly in § 1, displays the vorticity field ζ at $t = T/2$ (figure 1a) and its Lagrangian and Eulerian means (figures 1b,c). The strong wave dominates the instantaneous vorticity field. It is filtered successfully by time averaging. Clearly, the Lagrangian mean captures small-scale structures in the vorticity that are blurred by the Eulerian mean. This blurring is the result of the advection of the vorticity by the velocity field associated with the wave, which causes the vorticity structures to oscillate with the wave period. By construction, the Lagrangian mean removes these oscillations, leading to a sharper definition of the flow features.

It is interesting to examine the effect of Lagrangian and Eulerian averaging on the potential vorticity (PV) $q = (1 + Ro \zeta)/h$. In the absence of dissipation ($Re \rightarrow \infty$), this is a materially conserved field, meaning that $q(\boldsymbol{\varphi}(\boldsymbol{a}, t), t) = q_0(\boldsymbol{a})$, with q_0 determined by the initial condition. By definition (2.4), the Lagrangian mean PV then satisfies $\bar{q}^L(\bar{\boldsymbol{\varphi}}(\boldsymbol{a}, t), t) = q_0(\boldsymbol{a})$. Thus both q and \bar{q}^L are (smooth) rearrangements of the initial PV and hence rearrangements of one another; specifically, $\bar{q}^L(\boldsymbol{x}, t) = q(\boldsymbol{\Xi}(\boldsymbol{x}, t), t)$. This imposes constraints such as the two fields sharing the same values for their local extrema. Because $\bar{\boldsymbol{\varphi}}$ is not area-preserving, the distribution functions of q and \bar{q}^L (measuring the areas of regions where the fields are below specified values) do not coincide. Figure 6 shows q at $t = T/2$ and \bar{q}^L as well as the Eulerian mean PV. The Lagrangian mean PV \bar{q}^L appears as a slight deformation of q , consistent with it being rearrangement by a map $\boldsymbol{\Xi}$ that is close to the identity. In contrast, the Eulerian mean PV, which is not transported materially, shows blurred features, with in particular extrema that are substantially reduced compared with those of q and \bar{q}^L . There is a strong argument that the study of wave–mean-flow interactions, in the shallow-water model and more broadly, would benefit from the systematic analysis of Lagrangian mean fields such as the ones displayed in figures 1(b) and 6(b).

4.2.2. Moderate-Rossby-number flow initialised in geostrophic balance

In the second example, we initialise the shallow-water simulation with the same geostrophic flow as in § 4.2.1, but without any waves. The set-up is identical, but with resolution 128^2 and an increased Rossby number $Ro = 0.6$. As the flow evolves at this relatively high Ro , small-scale imbalance is generated, as shown in the instantaneous vorticity of figure 7(a) and the associated enstrophy spectrum in figure 7(d). Figures 7(b,e) show the Lagrangian mean vorticity field and the corresponding perturbation field,

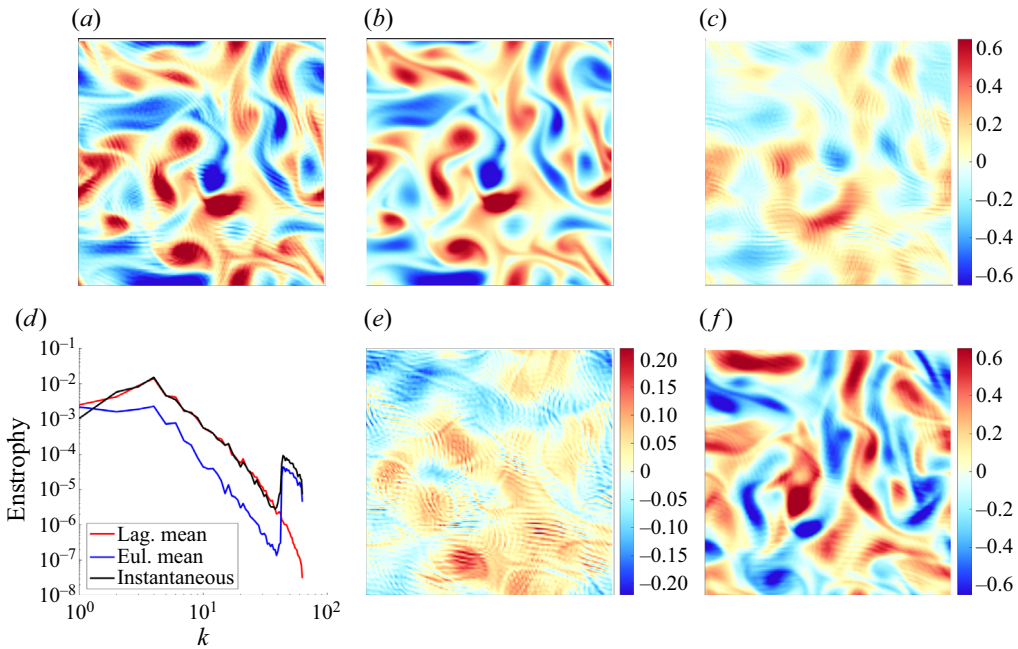


Figure 7. (a) Instantaneous vorticity $\zeta(x, t)$ at $t = 1.1$. (b) Lagrangian mean vorticity $\bar{\zeta}^L(x, \tau = 0)$ corresponding to averaging from $t = 0$ to $T = 2.2$. (c) Eulerian mean vorticity averaged from $t = 0$ to 2.2 . (d) Enstrophy spectra of instantaneous velocity at $t = 2.2$ and the Lagrangian and Eulerian mean fields. (e) Lagrangian perturbation field. (f) Eulerian perturbation field.

respectively. The perturbation field is computed as $\bar{\zeta}^L(x, \tau) - \zeta(x + \xi(x, \tau + T), \tau + T)$ with, in this case, $\tau = 0$ and $T = 2.2$. The Lagrangian average preserves the large-scale balanced flow so that the corresponding perturbation extracts the small-scale imbalance. This is evident in the enstrophy spectra of figure 7(d), where the instantaneous and Lagrangian mean spectra overlap at small wavenumbers, and the jump in the tail of the instantaneous spectrum is removed in the Lagrangian mean spectrum. In contrast, the mean flow is weakened in the Eulerian average (figure 7c) because at this moderate Ro , the vortices are advected substantially during the averaging period. This is corroborated in the Eulerian mean enstrophy spectrum, which has lower values for small wavenumbers, and follows the jump of the instantaneous spectrum at large wavenumbers. Consequently, balanced flow features dominate the Eulerian perturbation field (figure 7f). The results of figure 7 are similar to those of the synthetic flow considered by Shakespeare *et al.* (2021).

5. Discussion

This paper presents a novel approach for the numerical computation of Lagrangian means that relies on solving PDEs rather than tracking particles. We propose two strategies, each leading to a separate set of PDEs. Both strategies are based on the derivation of equations governing the evolution of partial means with respect to t . These partial means are defined as averages over a subset (τ, t) of each averaging interval $(\tau, \tau + T)$, and yield the (total) means for $t = \tau + T$. Strategy 1 uses the position of particles at time t as independent spatial variable. Hence it requires a map from the positions at the final time $t = \tau + T$ to the Lagrangian mean positions to ultimately present the results in terms of the latter, as is

standard in GLM theory. Strategy 2 computes directly the Lagrangian means in terms of Lagrangian mean positions.

A natural question is which of the two strategies should be preferred. There is no definitive answer: each strategy has pros and cons. If the spatial distribution of Lagrangian means does not matter in an application, then it suffices to solve (3.5) of strategy 1. When the spatial distribution is needed, strategy 1 requires the remapping (3.8), which can be affected by clustering: the mean positions $\mathcal{E}^{-1}(x, \tau + T)$ obtained from (3.9) for x on a regular grid may have a highly non-uniform distribution. This can lead to large numerical errors in the interpolation required to discretise (3.8). Strategy 2 circumvents this issue, as the generalised Lagrangian mean is computed directly on the desired spatial grid points. However, this advantage comes at the computational cost of evaluating more complicated right-hand sides in (3.13) and (3.15), which require interpolation at each averaging time step. Furthermore, strategy 2 leads to PDEs posed on a moving domain, unless the domain is periodic or has boundaries that correspond to constant coordinates.

As discussed in Kafiabad (2022), the full potential of our approach in saving memory and reducing computational cost is realised when the PDEs for the Lagrangian mean fields are solved on the fly, together with the dynamical model (as opposed to offline, using saved model outputs). In this case, it is beneficial to solve the Lagrangian mean PDEs using a numerical scheme that matches closely that of the dynamical model, because the instantaneous values of f and \mathbf{u} (required to solve the Lagrangian mean PDEs) are readily available at the same (physical or spectral) grid points. Moreover, the reasons that led to a particular choice of numerical discretisation for the dynamical equations – such as the type of boundary conditions – typically also apply to the Lagrangian mean PDEs.

In the main body of the paper, we restrict our attention to the Lagrangian averaging of a scalar function $f(\mathbf{x}, t)$. The averaging of vectors, differential forms and more general tensors is, however, of interest in applications. For instance, the Lagrangian means of the momentum 1-form $\mathbf{u} \cdot d\mathbf{x}$ (the integrand in Kelvin's circulation) and of the magnetic flux 2-form play crucial roles in the theory of wave–mean–flow interactions in fluid dynamics and magnetohydrodynamics (Soward 1972; Andrews & McIntyre 1978; Holm 2002; Gilbert & Vanneste 2018, 2021). The derivations in § 3 generalise straightforwardly to tensors when the language of push-forwards, pull-backs and Lie derivatives is employed. We illustrate this in Appendix C by generalising (3.13) of strategy 2 for the partial Lagrangian mean of $f(\mathbf{x}, t)$ to a tensor field $\sigma(\mathbf{x}, t)$.

We conclude by noting that practical averages such as the time average in (2.2)–(2.4) do not satisfy exactly the axioms of the more abstract averages assumed in the development of GLM and similar theories. In particular, the basic requirement that averaging leaves mean quantities unchanged, that is, $\overline{\bar{g}} = \bar{g}$, fails for (2.2), though the difference is small if there is a clear time scale separation between mean flow and perturbation. Interpreting numerically computed Lagrangian mean fields in light of these theories will therefore require us to understand how theoretical predictions are affected by the precise nature of the average.

Supplementary material. Supplementary material is available at <https://doi.org/10.1017/jfm.2023.228>.

Funding. H.A.K. and J.V. were supported by the UK Engineering & Physical Sciences Research Council (grant EP/W007436/1). J.V. was also supported by the UK Natural Environment Research Council (grant NE/W002876/1).

Declaration of interests. The authors report no conflict of interest.

Data availability statement. The data and scripts used for the simulations of § 4.2 are available at <https://github.com/turbulencelover/ComputingLagMeans>.

Author ORCID*s*.

 Hossein A. Kafiabad <https://orcid.org/0000-0002-8791-9217>;

 Jacques Vanneste <https://orcid.org/0000-0002-0319-589X>.

Appendix A. Partial and total Lagrangian mean velocities

We show that the partial mean velocity $\bar{\mathbf{u}}$ in (3.11) is not related to the Lagrangian mean velocity $\bar{\mathbf{u}}^L$ as might be expected naively: $\bar{\mathbf{u}}^L(\mathbf{x}, \tau) \neq \bar{\mathbf{u}}(\mathbf{x}, \tau + T; \tau)$, where we have reinstated the parametric dependence on τ for clarity. The definition of $\bar{\mathbf{u}}^L$ itself is not without ambiguity: the most natural definition is as the derivative of the total Lagrangian mean map with respect to the slow time, i.e.

$$\partial_\tau \bar{\varphi}^L(\mathbf{a}, \tau) = \bar{\mathbf{u}}^L(\bar{\varphi}^L(\mathbf{a}, \tau), \tau). \tag{A1}$$

Taking the derivative of the definition (2.3) of $\bar{\varphi}^L$ gives the explicit form

$$\bar{\mathbf{u}}^L(\bar{\varphi}^L(\mathbf{a}, \tau), \tau) = \frac{1}{T} (\varphi(\mathbf{a}, \tau + T) - \varphi(\mathbf{a}, \tau)). \tag{A2}$$

In contrast, the partial mean velocity $\bar{\mathbf{u}}$ is defined via a derivative with respect to the fast time t as seen from the definition (3.10), which takes the form

$$\partial_t \bar{\varphi}(\mathbf{a}, t; \tau) =: \bar{\mathbf{u}}(\bar{\varphi}(\mathbf{a}, t), t; \tau), \tag{A3}$$

on reinstating the parametric dependence on τ .

The difference between $\bar{\mathbf{u}}(\mathbf{x}, \tau + T; \tau)$ and $\bar{\mathbf{u}}^L(\mathbf{x}, \tau)$ should not come as a surprise since $\bar{\mathbf{u}}(\mathbf{x}, \tau + T; \tau)$ is constructed independently for each value of τ and hence cannot capture the change of $\bar{\varphi}^L(\cdot, \tau)$ measured by $\bar{\mathbf{u}}^L$. It can be made explicit by deducing from (A2) that

$$\begin{aligned} \bar{\mathbf{u}}^L(\mathbf{x}, \tau) &= \frac{1}{T} (\varphi(\bar{\varphi}^{-1}(\mathbf{x}, \tau + T; \tau), \tau + T) - \varphi(\bar{\varphi}^{-1}(\mathbf{x}, \tau + T; \tau), \tau)) \\ &= \frac{1}{T} (\bar{\boldsymbol{\Xi}}(\mathbf{x}, t; \tau) - \varphi(\bar{\varphi}^{-1}(\mathbf{x}, \tau + T; \tau), \tau)), \end{aligned} \tag{A4}$$

which differs from $\bar{\mathbf{u}}(\mathbf{x}, \tau + T; \tau) = (\bar{\boldsymbol{\Xi}}(\mathbf{x}, \tau + T; \tau) - \mathbf{x})/T$ since $\varphi(\bar{\varphi}^{-1}(\mathbf{x}, \tau + T; \tau), \tau) \neq \mathbf{x}$.

Note that in GLM, $\bar{\mathbf{u}}^L$ is defined as the Lagrangian mean of the components of \mathbf{u} instead of via (A1). In Euclidean space, the two definitions are equivalent since

$$\frac{1}{T} \int_\tau^{\tau+T} \mathbf{u}(\varphi(\mathbf{a}, s), s) ds = \frac{1}{T} \int_\tau^{\tau+T} \partial_s \varphi(\mathbf{a}, s) ds = \frac{\varphi(\mathbf{a}, \tau + T) - \varphi(\mathbf{a}, \tau)}{T}, \tag{A5}$$

using (2.1). This equivalence does not hold for definitions of the Lagrangian mean flow that have been proposed as geometric alternatives to GLM valid on any manifold (Soward & Roberts 2010; Gilbert & Vanneste 2018; Vanneste & Young 2022).

Appendix B. Semi-Lagrangian discretisation of strategy 1

We show that the algorithm developed by Kafiabad (2022) for the ‘grid-based calculation of Lagrangian average’ is equivalent to a semi-Lagrangian discretisation of (3.5) and (3.9). We denote the grid points by X_i , with i a multi-index in dimension more than 1, and the set of all grid points by $\{X_\alpha\}$. The semi-Lagrangian discretisation (3.5) gives the value $\tilde{f}(X_i, t_{n+1})$ of \tilde{f} at each grid point at t_{n+1} , the end of a time step, in terms of the set

of values $\tilde{f}(\{X_\alpha\}, t_n)$ at t_n , the start of the time step. It is based on a semi-Lagrangian discretisation of the equivalent PDE (4.1), namely

$$\tilde{g}(X_i, t_{n+1}) = \tilde{g}(Y_i, t_n) + (t_{n+1} - t_n)f(X_i, t_{n+1}), \tag{B1}$$

where we use a backward Euler step for the time stepping. Here, Y_i denotes the position at t_n of the particle that passes through X_i at t_{n+1} ; it can be evaluated using

$$Y_i = X_i - (t_{n+1} - t_n) \mathbf{u}(X_i, t_{n+1}), \tag{B2}$$

which follows from approximating the derivative in (2.1) by a backward finite difference. Rewriting (B1) in terms of \tilde{f} and rearranging, we obtain

$$\tilde{f}(X_i, t_{n+1}) = \frac{(t_n - \tau)\tilde{f}(Y_i, t_n) + (t_{n+1} - t_n)f(X_i, t_{n+1})}{t_{n+1} - \tau}, \tag{B3}$$

which is the same as (2.6) in Kafiabad (2022) when the last term is replaced by their (2.7). Using the trapezoidal rule instead of backward Euler leads to

$$\tilde{f}(X_i, t_{n+1}) = \frac{(t_n - \tau)\tilde{f}(Y_i, t_n) + (t_{n+1} - t_n)[f(Y_i, t_n) + f(X_i, t_{n+1})]/2}{t_{n+1} - \tau}, \tag{B4}$$

which is (2.8) of Kafiabad (2022) and can provide more accurate results. Note that the evaluation of $\tilde{f}(Y_i, t_n)$ requires an interpolation since Y_i is not on the grid.

If we consider the particle position as a special case of f , (B4) turns into

$$\Xi^{-1}(X_i, t_{n+1}) = \frac{(t_n - \tau) \Xi^{-1}(Y_i, t_n) + (t_{n+1} - t_n)(X_i + Y_i)/2}{t_{n+1} - \tau}, \tag{B5}$$

which parallels (2.9) in Kafiabad (2022). The complete algorithm iterates (B3)–(B5) from $t_0 = \tau$ to $t_{N+1} = \tau + T$ to obtain $\tilde{f}^L(x, \tau) = \tilde{f}(x, \tau + T)$ and $\Xi^{-1}(x, t + T)$. In the final step, $\Xi^{-1}(x, t + T)$ is employed to apply the remapping (3.8) to calculate $\tilde{f}^L(x, \tau)$.

Appendix C. Lagrangian mean of tensors

We show how (3.13) for the Lagrangian mean of scalar functions generalises readily to tensors when phrased in the language of push-forwards, pull-backs and Lie derivatives (e.g. Frankel 2004). The definition (3.2b) of the partial Lagrangian mean generalises as

$$\bar{\varphi}_\tau^{L*} \bar{\sigma}_\tau^L := \frac{1}{T} \int_\tau^{\tau+T} \varphi_s^* \sigma_s \, ds. \tag{C1}$$

Here, we make the time variable appear explicitly as a subscript, and we use $\bar{\varphi}_t^{L*}$ and φ_t^* to denote the pull-backs by, respectively, the Lagrangian mean map and flow map. Pull-backs are interpreted differently depending on the nature of σ_t : in particular, for a scalar, $(\varphi_t^* f)(x) = f(\varphi_t(x))$, while for a 1-form, $(\varphi_t^*(u \cdot dx))(x) = u_j(\varphi_t(x)) \partial_i \varphi_{tj}(x) dx_i$. The definition (C1) is a natural, geometrically intrinsic definition of the Lagrangian mean because it ensures that the tensors σ_s at different times s are transported to the same position in label space before the averaging is carried out (see Gilbert & Vanneste 2018). The partial mean of σ_t associated with (C1) reads

$$\bar{\varphi}_t^* \bar{\sigma}_t := \frac{1}{t - \tau} \int_\tau^t \varphi_s^* \sigma_s \, ds, \tag{C2}$$

and is understood to depend parametrically on σ .

Computing Lagrangian means

Differentiating (C2) with respect to t and using that $\partial_t \bar{\varphi}_t^* \bar{\sigma}_t = \bar{\varphi}_t^* (\partial_t + \mathcal{L}_{\bar{u}}) \bar{\sigma}_t$, where $\mathcal{L}_{\bar{u}}$ is the Lie derivative, gives (Gilbert & Vanneste 2018)

$$\bar{\varphi}_t^* (\partial_t + \mathcal{L}_{\bar{u}}) \bar{\sigma}_t = \frac{-\bar{\varphi}_t^* \bar{\sigma}_t + \varphi_t^* \bar{\sigma}_t}{t - \tau}. \quad (\text{C3})$$

Pushing forward by $\bar{\varphi}_{t*}$ reduces this to

$$(\partial_t + \mathcal{L}_{\bar{u}}) \bar{\sigma}_t = \frac{\Xi_t^* \sigma_t - \sigma_t}{t - \tau}, \quad (\text{C4})$$

on using that (3.6), here in the form $\Xi_t = \varphi_t \circ \bar{\varphi}_t^{-1}$, implies that $\Xi_t^* = \bar{\varphi}_{t*} \varphi_t^*$. Equation (C4) generalises (3.13) to tensors. Its coordinate form depends on the type of tensor because of the different forms taken by $\mathcal{L}_{\bar{u}}$ and Ξ_t^* . Note that (3.9) and (3.15) for Ξ^{-1} and Ξ do not have a tensorial nature: because of the Cartesian definition of the Lagrangian mean map (2.3), Ξ is simply a triple of scalar functions rather than a vector. See Gilbert & Vanneste (2018) for geometrically intrinsic definitions of the Lagrangian mean map alternative to (2.3).

REFERENCES

- ANDREWS, D.G. & MCINTYRE, M.E. 1978 An exact theory of nonlinear waves on a Lagrangian-mean flow. *J. Fluid Mech.* **89**, 609–646.
- BACHMAN, S.D., SHAKESPEARE, C.J., KLEYPAS, J., CASTRUCCIO, F.S. & CURCHITSER, E. 2020 Particle-based Lagrangian filtering for locating wave-generated thermal refugia for coral reefs. *J. Geophys. Res. Oceans* **125** (7), e2020JC016106.
- BREThERTON, F.P. 1971 The general linearized theory of wave propagation. In *Mathematical Problems in the Geophysical Sciences*, Lectures in Applied Mathematics, vol. 13, pp. 61–102. American Mathematical Society.
- BÜHLER, O. 2014 *Waves and Mean Flows*, 2nd edn. Cambridge University Press.
- BÜHLER, O. & MCINTYRE, M.E. 1998 On non-dissipative wave–mean interactions in the atmosphere or oceans. *J. Fluid Mech.* **354**, 301–343.
- FRANKEL, T. 2004 *The Geometry of Physics*, 2nd edn. Cambridge University Press.
- GILBERT, A.D. & VANNESTE, J. 2018 Geometric generalised Lagrangian-mean theories. *J. Fluid Mech.* **839**, 95–134.
- GILBERT, A.D. & VANNESTE, J. 2021 A geometric look at MHD and the Braginsky dynamo. *Geophys. Astrophys. Fluid Dyn.* **115**, 436–471.
- GRIMSHAW, R. 1975 Nonlinear internal gravity waves in a rotating fluid. *J. Fluid Mech.* **71** (3), 497–512.
- HOLM, D.D. 2002 Lagrangian averages, averaged Lagrangians, and the mean effects of fluctuations in fluid dynamics. *Chaos* **12**, 518–530.
- JONES, C.S., XIAO, Q., ABERNATHEY, R. & SMITH, K.S. 2022 Separating balanced and unbalanced flow at the surface of the Agulhas region using Lagrangian filtering. doi:10.31223/X5D352.
- KAFIABAD, H.A. 2022 Grid-based calculation of the Lagrangian mean. *J. Fluid Mech.* **940**, A21.
- KAFIABAD, H.A., VANNESTE, J. & YOUNG, W.R. 2021 Wave-averaged balance: a simple example. *J. Fluid Mech.* **911**, R1.
- MOORE, D. 1970 The mass transport velocity induced by free oscillations at a single frequency. *Geophys. Fluid Dyn.* **1**, 237–247.
- NAGAI, T., TANDON, A., KUNZE, E. & MAHADEVAN, A. 2015 Spontaneous generation of near-inertial waves by the Kuroshio front. *J. Phys. Oceanogr.* **45** (9), 2381–2406.
- SHAKESPEARE, C.J., GIBSON, A.H., HOGG, A.M.C., BACHMAN, S.D., KEATING, S.R. & VELZEBOER, N. 2021 A new open source implementation of Lagrangian filtering: a method to identify internal waves in high-resolution simulations. *J. Adv. Model. Earth Syst.* **13**, e2021MS002616.
- SHAKESPEARE, C.J. & HOGG, A.M.C. 2017 Spontaneous surface generation and interior amplification of internal waves in a regional-scale ocean model. *J. Phys. Oceanogr.* **47**, 811–826.
- SHAKESPEARE, C.J. & HOGG, A.M.C. 2018 The life cycle of spontaneously generated internal waves. *J. Phys. Oceanogr.* **48**, 343–359.
- SHAKESPEARE, C.J. & HOGG, A.M.C. 2019 On the momentum flux of internal tides. *J. Phys. Oceanogr.* **49**, 993–1013.

- SOWARD, A.M. 1972 A kinematic theory of large magnetic Reynolds number dynamos. *Phil. Trans. R. Soc. Lond.* **A272**, 431–462.
- SOWARD, A.M. & ROBERTS, P.H. 2010 The hybrid Euler–Lagrange procedure using an extension of Moffatt’s method. *J. Fluid Mech.* **661**, 45–72.
- VALLIS, G.K. 2017 *Atmospheric and Oceanic Fluid Dynamics*, 2nd edn. Cambridge University Press.
- VANNESTE, J. 2013 Balance and spontaneous wave generation in geophysical flows. *Annu. Rev. Fluid Mech.* **45**, 147–172.
- VANNESTE, J. & YOUNG, W.R. 2022 Stokes drift and its discontents. *Phil. Trans. R. Soc. A* **380**, 20210032.
- WAGNER, G.L. & YOUNG, W.R. 2015 Available potential vorticity and wave-averaged quasi-geostrophic flow. *J. Fluid Mech.* **785**, 401–424.