

Microstructure-informed probability-driven point-particle model for hydrodynamic forces and torques in particle-laden flows

Arman Seyed-Ahmadi¹ and Anthony Wachs^{1,2,†}

¹Department of Chemical & Biological Engineering, The University of British Columbia, 2360 East Mall, Vancouver, BC V6T 1Z3, Canada

²Department of Mathematics, The University of British Columbia, 1984 Mathematics Road, Vancouver, BC V6T 1Z2, Canada

(Received 16 December 2019; revised 25 April 2020; accepted 31 May 2020)

We present a novel deterministic model that is capable of predicting particle-to-particle force and torque fluctuations in a fixed bed of randomly distributed monodisperse spheres. First, we generate our dataset by performing particle-resolved direct numerical simulations (PR-DNS) of arrays of stationary spheres in moderately inertial regimes with a Reynolds number range of $2 \leq Re \leq 150$ and a solid volume fraction range of $0.1 \leq \phi \leq 0.4$. The key idea exploited by our model is that, while the arrangement of neighbours around each particle is uniform and random, conditioning forces or torques exerted on a reference sphere to specific ranges of values results in the emergence of significantly non-uniform distributions of neighbouring particles. Based on probabilistic arguments, we take advantage of the statistical information extracted from PR-DNS to construct force/torque-conditioned probability distribution maps, which are ultimately used as basis functions for regression. Given the locations of surrounding particles as input to the model, our results demonstrate that the present probability-driven framework is capable of predicting up to 85 % of the actual observed force and torque variation in the best cases. Since the precise location of each particle is known in an Eulerian–Lagrangian (EL) simulation, our model would be able to estimate the unresolved subgrid force and torque fluctuations reasonably well, and thereby considerably enhance the fidelity of EL simulations via improved interphase coupling.

Key words: particle/fluid flow, computational methods

1. Introduction

The ubiquity of flows where a dispersion of particles is carried by a fluid phase has motivated many theoretical, experimental and numerical studies. Such systems, also referred to as particle-laden flows, are widely encountered in natural and industrial settings. Sediment transport, rain and drop formation, fluidised beds and slurry flows are only a few examples among many where particle-laden flows occur. Analytical treatment of these systems is generally limited to asymptotic cases with very low Reynolds numbers

† Email address for correspondence: wachs@mail.ubc.ca




FIGURE 1. Depiction of the concept of multiscale strategy in modelling particle-laden flows.

and solid volume fractions, whereas the practical interest usually lies in the opposite end of this spectrum. Conversely, the unprecedented availability of massive computational power in the past two decades has substantially promoted the utility of numerical methods for investigation of particle-laden flows.

In a dispersed multiphase system, particles interact with each other locally on the length scale of the particle diameter d , which could lead to the occurrence of close-range phenomena such as drafting–kissing–tumbling (DKT) due to wake attraction in inertial regimes (Fortes, Joseph & Lundgren 1987; Dash & Lee 2015). Interactions of this nature are known to contribute to the formation of particle clusters that extend several diameters (Kajishima & Takiguchi 2002; Uhlmann & Doychev 2014; Fornari, Ardekani & Brandt 2018). These structures, in turn, potentially interact with each other and thus bring about an integral length scale that could be one or two orders of magnitude larger than d . Clustering of water droplets in clouds is an interesting natural example, which is known to greatly enhance coalescence and thereby explaining the growth rate of droplets (Collins & Keswani 2004). On the other hand, clusters have a remarkable effect on particle residence time, and heat and mass transfer in circulating fluidised beds (Louge, Lischer & Chang 1990; Wylie & Koch 2000). A fluidised bed reactor in a fluid catalytic cracking unit is typically 14 m high and 6 m in diameter. Within such a device, fluid–solid interactions and particle collisions in the sub-millimetre scale can directly influence flow structures such as clusters of the order of metres (van der Hoef *et al.* 2008). The enhancement of local concentration of particles in clusters also leads to increased collision rates. Clearly, the time scale of particle collisions is much shorter than convective or diffusive time scale of the flow, creating yet another level of scale separation. The cascade of such multi-scale interactions modifies the character of particle suspensions and significantly affects quantities of interest, such as settling rate and particle velocity fluctuations (Zaidi, Tsuji & Tanaka 2014; Zaidi 2018; Willen & Prosperetti 2019). Due to the presence of a wide range of length scales, a single description of the physics would fail to provide balance between the required complexity or resolution, and the associated computational cost for all involved scales. This is the motivation for development of various numerical methodologies aiming at resolving the flow at three major length scales; namely, the micro-, meso- and macroscale. Such a multi-scale methodology is schematically shown in figure 1.

A micro-scale flow description assumes that particles are at least an order of magnitude larger than fluid grid cells, meaning that the fluid–solid interface is well resolved. Since the hydrodynamic forces and torques can be directly computed from the integration of pressure and viscous stresses available as field variables, transfer of momentum between phases involves no approximation or modelling, and is dictated by satisfying the no-slip condition on the solid boundary. Micro-scale methods (i.e. particle-resolved

direct numerical simulations, PR-DNS) were first introduced in the framework of body-conforming moving-mesh methods (Feng, Hu & Joseph 1994a,b; Hu, Patankar & Zhu 2001) where Navier–Stokes equations are solved on the fluid mesh, and particles are treated as boundaries of the flow. As the system evolves, frequent re-meshings become inevitable so as to adapt the computational grid to the new configuration of the fluid–solid system, rendering such methods inefficient for simulating particle-laden flows. Fixed-grid methods, also collectively called fictitious domain methods (Maxey 2017), were proposed to relieve this burden by extending the fluid domain to include the particles and solving Navier–Stokes equations over the entire fluid–particle domain, thus completely eliminating the need for re-meshing. Immersed boundary methods (IBM) (Peskin 1977; Mittal & Iaccarino 2005) and distributed Lagrange multiplier/fictitious domain (DLM-FD) methods (Glowinski *et al.* 1999; Wachs 2010) are two such techniques which differ in the enforcement of rigid-body motion in the particle domain and in the computation of interaction forces. Another popular class of PR-DNS tools are the lattice Boltzmann methods (LBM) which are based on kinetic theory of gases (Aidun & Clausen 2010). A significant number of fixed-bed studies of particulate flows in the literature are carried out using LBM (Hill, Koch & Ladd 2001a,b; van der Hoef, Beetstra & Kuipers 2005; Beetstra, van der Hoef & Kuipers 2007; Bogner, Mohanty & Rde 2015). Assuming particles of typical size in the range of [200 μm , 1000 μm], the micro-scale approach is suitable for simulating laboratory scale devices with a size of $O(0.01\text{ m})$ (van der Hoef *et al.* 2008), or $O(10^3\text{--}10^4)$ particles in the context of gas–solid fluidisation. Although PR-DNS provides a complete and model-free description of particle-laden flows, handling systems at the intermediate (i.e. meso) scale with a physical size of $O(0.1\text{ m})$ (van der Hoef *et al.* 2008) containing $O(10^5\text{--}10^6)$ well-resolved particles (Subramaniam 2013; Zhong *et al.* 2016) has only recently been feasible with PR-DNS, and requires massively parallel computing resources (Gtz *et al.* 2010; Rettinger *et al.* 2017; Horne & Mahesh 2019; Willen & Sierakowski 2019). The Eulerian–Lagrangian (EL) technique, on the other hand, attempts to reduce computational costs by taking the volume of each fluid cell to be generally an order of magnitude larger than that of an individual particle or at least to be of the order of magnitude of that of an individual particle. Because flow in the immediate vicinity of the particles is not resolved, direct computation of the hydrodynamic interaction force and torque is not possible. Consequently, interphase coupling needs to be established via a suitable force closure model. The fluid phase sees particles only as point sources and sinks of momentum, with the volume of particles appearing solely through the local porosity in the volume-averaged mass and momentum conservation equations and the employed drag force correlation (van der Hoef, van Sint Annaland & Kuipers 2004). In such an approach, the volume-averaged Navier–Stokes equations are solved on an Eulerian grid for the fluid phase, whereas the positions of the particles are tracked using Newton’s equations of motion in a Lagrangian manner. For the case of dense suspensions, this is usually handled by the discrete element method (DEM) with a proper contact model accounting for particle collisions (and hence the name discrete element method-computational fluid dynamics or ‘DEM-CFD’). Since particles are treated as points suspended in fluid cells, the EL method is also referred to as the point-particle (PP) method. This terminology, however, may not be descriptive enough for the newer EL methods that are capable of simulating finite-size particles with a diameter of the order of fluid cells (Capecelatro & Desjardins 2013). Even with EL methods, modelling engineering scale pilot devices with a size of $O(1\text{ m})$ containing $O(10^9)$ particles is impractical at present, since tracking billions of individual particles in a Lagrangian manner poses a serious computational challenge. With that being the case, one could alleviate this issue by accounting for the presence of

particles indirectly. The two-fluid model (TFM) (Gidaspow 1994) is a numerical model employing an Eulerian–Eulerian approach, where the fluid and the solid are both assumed to behave as interpenetrating continua. Since the solid phase is also modelled as a fluid continuum, details of particle–particle interactions are embedded in the effective solid pressure and shear and bulk viscosity closure terms. The two phases, in turn, interact through an appropriate drag model (van der Hoef *et al.* 2004, 2008).

The up-scaling of simulations can be a viable alternative for PR-DNS only if the closure models to be used are sufficiently accurate and faithful to the actual underlying physics. In an EL simulation, the governing equations have to be closed with appropriate terms accounting for the fluid–solid momentum transfer. Multiphase flow modelling enters the stage to bridge the gap between PR-DNS and the EL approach by supplying this missing piece of crucial information. The classical Stokes drag $\mathbf{F} = 3\pi\mu d\mathbf{u}$, with μ being the dynamic viscosity of the fluid and \mathbf{u} the fluid-particle relative velocity, is only valid at vanishing Reynolds numbers, i.e. $Re = 0$, for a steady uniform flow around a fixed sphere. When the macroscale undisturbed flow (the undisturbed flow is defined as the flow that would have existed in the absence of a particle (Squires 2007, Balachandar 2009)) \mathbf{u} is no longer steady, one can use the Basset–Boussinesq–Oseen relation given as

$$\mathbf{F} = \mathbf{F}_{um} + \mathbf{F}_D + \mathbf{F}_{am} + \mathbf{F}_h, \quad (1.1)$$

where the terms on the right-hand side account for the undisturbed flow force, the quasi-steady drag, added-mass and Basset history forces, respectively. The elegant Faxén’s law (Faxen 1923) makes it possible to extend the validity of (1.1) to spatially non-uniform flows by replacing the undisturbed flow \mathbf{u} by $\bar{\mathbf{u}}^s$ and $\bar{\mathbf{u}}^v$, i.e. the average value over the particle’s surface or volume. This formulation was derived by Maxey and Riley (Maxey 1983) and Gatignol (Gatignol 1983) and is rigorously valid in the Stokes limit. When particles are much smaller than the macroscopic length scale of the flow, the particle Reynolds number based on relative slip velocity typically becomes very small. This is the case when a dilute dispersion of particles are suspended in a turbulent flow such that $d/\eta \ll 1$, where η shows the Kolmogorov scale. In this situation, the Maxey–Riley–Gatignol (MRG) equation accurately predicts hydrodynamic forces experienced by the particle. For inertial regimes with finite Reynolds numbers, force contributions in the MRG equation should be modified and empiricism is inevitable due to a lack of theoretical analysis. The standard drag curve of an isolated sphere given as $C_D = (24/Re)(1 + 0.15Re^{0.687})$ (Schiller & Naumann 1933) is such a modification that characterises the Reynolds number dependence of the quasi-steady drag term \mathbf{F}_D in the MRG equation. For $Re > 0$, ambient shear or vorticity in the background flow gives rise to the Saffman lift force (Saffman 1965), which is absent in the MRG equation since there is no such lift force in Stokes flow (Maxey 1983). Furthermore, rotation of a particle induces an excess lift even in uniform flow, which is attributed to the Magnus effect (Bagchi & Balachandar 2002). The Magnus force should as well be included as an additional term alongside the Saffman lift in the MRG equation when inertial regimes are considered. The interested reader is referred to Subramaniam & Balachandar (2018) for a detailed overview of various approaches of modelling dispersed multiphase flow.

The PP approach faces a serious challenge when the size of particles becomes comparable to the scales of the macroscopic flow. When the carrier flow is already turbulent, this occurs when $d/\eta \approx 1$. The other scenario is when the particle concentration increases and the suspension can no longer be regarded as dilute. Even at a solid volume fraction of 1%, the average distance to the closest neighbour is only $3.7d$ (Akiki, Jackson & Balachandar 2016). The likelihood of hydrodynamic interaction hence

increases substantially. The disturbances created by the particles results in the appearance of pseudo-turbulence, i.e. a non-uniform flow that not only varies spatially on the scale of the particle diameter, but is also potentially prone to temporal fluctuations for high enough Re . Extending the PP model to finite-size particles is hampered with significant complications associated with the application of analytical or empirical force relations, the expression of which are all in terms of undisturbed flow. In the context of finite-size particles, the undisturbed flow is difficult to obtain since it requires evaluating the same system but with a specific particle removed. Moreover, now that the particle is relatively large, the undisturbed flow varies on the particle scale. How such a spatially non-homogeneous flow would affect the resulting forces at finite Re remains an open question (Balachandar 2009).

Parametrising the drag law in terms of solid volume fraction, in addition to the Reynolds number, has been the first step towards the prediction of particle-laden flow behaviour where collective effect of particles cannot be neglected. Theoretical studies are limited to the Stokes flow conditions and very dilute suspensions; namely, $Re \rightarrow 0$ and $\phi \rightarrow 0$ where ϕ denotes the solid volume fraction (Hasimoto 1959; Batchelor 1972; Sangani & Acrivos 1982). Proposed correlations by Wen & Yu (1966) for dilute suspensions and Ergun's equation (Ergun 1952) for denser systems are the earliest experimental efforts in this regard. Exponential growth of computing power in the last two decades has made PR-DNS a preferable alternative for developing more accurate drag correlations over wider ranges of Re and ϕ . Contrary to the experimental approach, arbitrary ideal flow conditions can be imposed, and forces are computed directly in PR-DNS (Tang *et al.* 2015) instead of indirect measurement based on the settling velocity (Richardson & Zaki 1954) or pressure drop (Ergun 1952). Among PR-DNS techniques, LBM has been the method of choice in numerous studies on drag correlation. Hill *et al.* (2001a) simulated a fixed bed of spheres with ordered and random arrangements up to the close packing limit at low Re , and later for moderate Re (Hill *et al.* 2001b). Bi-dispersity in random arrays of fixed spheres was also addressed at very low Re using LBM by van der Hoef *et al.* (2005), and subsequently for Re up to 1000 by Beetstra *et al.* (2007). Recently, other drag correlations were proposed by Bogner *et al.* (2015) using LBM, and by Tang *et al.* (2015) and Tenneti, Garg & Subramaniam (2011) using IBM.

Even though the idea of simulating fixed beds of spheres instead of realistic moving particles is justified by drawing analogy with high Stokes number gas–solid flows, this simplification was challenged by Tang *et al.* (2015). They showed that at $\phi = 0.5$, deviation between the actual drag experienced by mobile particles and the drag computed from static bed correlations is significant. A similar observation had been previously made by Kriebitzsch, van der Hoef & Kuipers (2013), that not only are the gas–solid forces underestimated by conventional drag laws in an EL simulation, but there is also a large scatter of the drag data in a PR-DNS of fluidised beds which the EL approach fails to capture. They noted that the higher drag force seems to correlate with local granular temperature or particle agitation, which in turn is an outcome of subgrid force fluctuations. In a recent work, Esteghamatian *et al.* (2018) attempted to alleviate the suppression of granular temperatures by introducing a stochastic component in the drag closure, the parameters of which were extracted from PR-DNS. Comparisons of their stochastic model with conventional DEM-CFD indicated better prediction of granular temperatures in liquid-solid regimes. Ultimately, the key missing component seems to pertain much more conspicuously to the physical fidelity of the drag model rather than the accuracy of fixed-bed drag correlations.

The preceding discussion signifies that unless the physics is properly accounted for, any effort towards further improving EL simulations would be futile regardless of the accuracy

of conventional drag correlations. From all the proposed drag models, we know that the mean drag force experienced by the particles demonstrates a strong correlation with solid volume fraction. In other words, the functional dependence of the drag model on ϕ is able to account for the presence of other particles, but only and strictly in an average sense. It does not matter whether a particle is shielded by an upstream neighbour or exposed entirely to the oncoming flow, the drag law predicts the same force in both situations. Therefore, it would be highly desirable to construct a force model that is capable of accounting for the specific arrangement of surrounding particles; namely,

$$F_i = f(Re, \phi, \{r_{j=1}, \dots, r_{j=M}\}), \quad (1.2)$$

where F_i is the force experienced by particle i and r_j the position vector of neighbour j relative to particle i , while M denotes the number of influential neighbours and f the functional dependence. Akiki *et al.* (2016) showed by analysis of PR-DNS data for $20 \leq Re \leq 180$ and $0.11 \leq \phi \leq 0.44$ that there is a substantial scatter in the hydrodynamic force experienced by individual particles due to the particular arrangement of surrounding spheres, and that the local volume fraction had almost no correlation with the force fluctuations. By utilising a simple anisotropic measure of each particle's neighbourhood, they were able to capture some of the drag variation, whereas results for the lift force were less accurate. The notable work of Akiki, Jackson & Balachandar (2017a) and their pairwise interaction extended point-particle (PIEP) model is the first to systematically account for the effect of neighbouring particles on drag and lift in a deterministic manner. Their model involved linear superposition of perturbations created by each neighbouring particle in a pairwise manner in order to obtain the undisturbed flow, then using Faxén's law to compute various contributions to the total hydrodynamic force from the non-uniform undisturbed flow. The PIEP model was shown to be capable of predicting up to 75 % of the drag force variations for the $(Re, \phi) = (0.1, 38)$, whereas for a denser case of $(Re, \phi) = (0.21, 87)$, approximately 56 % of the variations were captured. For the lift, however, the results were not as promising. Subsequently, they extended their work to include modelling of hydrodynamic torques as well, and also tested their model for sedimentation of 2, 5 and 80 spheres (Akiki, Moore & Balachandar 2017b). Remarkably, they were able to reproduce DKT of two spheres with the PIEP model in contrast to the inability of standard PP approach to do so. Quite recently, the same group (Moore, Balachandar & Akiki 2019) attempted to improve the shortcomings of the PIEP model particularly at high volume fractions by combining a data-driven approach based on nonlinear regression with their original physics-driven model (Akiki *et al.* 2017a,b). The resulting hybrid model was shown to considerably enhance the accuracy of the PIEP model particularly at higher solid volume fractions.

The fast-growing trend of machine learning (ML) and data-driven algorithms has brought new prospects to fluid flow modelling. Data-based methods have already been present in the context of dimensionality reduction techniques and are hence not alien in the fluids community (Kutz 2017). Neural networks are a popular subset of ML techniques that have been applied with ground-breaking success to image (Krizhevsky, Sutskever & Hinton 2012) and speech (Hinton *et al.* 2012) recognition tasks. Neural networks are shown to be universal function approximators (Hornik, Stinchcombe & White 1989), and are thus capable of mapping input features to output variables in complex multidimensional problems rife with strong inherent nonlinearities. The tempting power of ML has motivated a major effort towards its applications in fluid flow simulations, especially turbulent flow modelling. Several works have focused on improving closure terms in the widely used Reynolds-averaged Navier–Stokes models using ML algorithms (Ling & Templeton 2015; Ling, Kurzwski & Templeton 2016; Wu, Xiao & Paterson 2018;

Duraisamy, Iaccarino & Xiao 2019), and fewer on large eddy simulation modelling (LES) (Maulik *et al.* 2018; Beck, Flad & Munz 2019; Xie *et al.* 2019a,b; Yang *et al.* 2019). There has also been attempts in the context of multiphase flows to develop closure models for TFM simulations (Ma, Lu & Tryggvason 2015; Jiang *et al.* 2019). Another direction pursued by some researchers is to take advantage of convolutional neural networks (LeCun, Bengio & Hinton 2015), commonly used in image recognition tasks, for direct approximation of flow field variables (Guo, Li & Iorio 2016; Sekar & Khoo 2019). Despite its seeming success, a crucial fact about ML is that its high accuracy and flexibility are achieved at the expense of supplying large amounts of data for the purpose of training. ML algorithms such as neural networks require large volumes of data to find fitting functions through adjustment of their parameters (i.e. weights and biases). This process, also known as the ‘learning’ or ‘training’ phase, is an iterative optimisation procedure aimed at minimising errors between the real data and those predicted by the algorithm. This approach works successfully for computer vision tasks, for example, due to an abundance of labelled data and the interpolatory nature of the problem (Sun *et al.* 2019). However, the application of off-the-shelf ML techniques to flow dynamics problems (or any other physical phenomena for that matter) inevitably suffers from non-interpretability in terms of governing physics due to ML’s ‘black-box’ nature. Another issue that practically impedes ML application in physical systems such as particle-laden flows is that we can generate, at best, no more than a few thousand or a few tens of thousands of samples (e.g. force and torque on each particle) with PR-DNS. Consider that even for the most idealised case of mono-dispersed spheres, we would still have to sweep the parameter space of Re and ϕ . In such a ‘small data’ regime (Raissi, Perdikaris & Karniadakis 2019), the full power of neural networks may not be exploited unless the physics equations are incorporated, or ‘hard-coded’ in the structure of the algorithm. This can be achieved by directly minimising residuals of the governing equations through loss functions and hence ensuring the physical fidelity of the predictions, giving rise to ‘physics-informed’ deep learning algorithms (Raissi & Karniadakis 2018; Raissi *et al.* 2018, 2019). Even if successful in prediction (within the range of training dataset at best), the initial spark of a theory-blind conventional ML model quickly fades away since even at its peak, it still is a ‘glorified curve-fitting’ procedure (Succi & Coveney 2019) and without guiding theory, such pure empiricism fails to provide knowledge (Coveney, Dougherty & Highfield 2016). A physics-informed ML approach (Raissi & Karniadakis 2018; Raissi *et al.* 2018, 2019) offers an advantage in that respect due to having physical fidelity engraved in its core. Even if an accurate ML model is at hand, there exists another issue in the framework of EL simulations. In the DEM-CFD approach, the hydrodynamic force on each particle has to be evaluated at each time step. A deep neural network typically runs the input through hundreds or even thousands of pre-tuned parameters to output a single prediction, incurring a significant computational cost on the EL simulation.

In the present work, we propose a novel data-driven model that relies on force/torque-conditioned probabilities of particle arrangements extracted from PR-DNS in order to correlate hydrodynamic forces and torques to the unique neighbourhood of each particle. In the remainder of this work, we refer to our model as the microstructure-informed probability-driven point-particle (MPP) model. We will provide probabilistic arguments for the prediction of force/torque fluctuations, and apply our method to the data from PR-DNS of fixed beds of randomly distributed spherical particles at various Reynolds numbers and solid volume fractions. After validating our results with the existing literature, we then evaluate the accuracy of our model’s predictions and demonstrate its performance by providing comparisons with PR-DNS force and torque data.

2. PR-DNS of fixed beds of spheres

2.1. Governing equations

Conservation of momentum and mass for the fluid phase is expressed in terms of incompressible Navier–Stokes and continuity equations for a Newtonian fluid as follows:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u}, \quad (2.1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.2)$$

where \mathbf{u} , p and Re respectively denote the fluid velocity vector, pressure and particle Reynolds number, defined as

$$Re = \frac{\rho U d}{\mu} = \frac{\rho(1 - \phi)u_s d}{\mu}, \quad (2.3)$$

where ρ , μ are the fluid phase density and dynamic viscosity. The particle Reynolds number is defined based on the superficial velocity $U = (1 - \phi)u_s$, and u_s represents the average interstitial fluid velocity. In (2.1) and (2.2) and what follows, all variables are non-dimensionalised with respect to the particle diameter d as the length scale, U as the velocity scale, ρU^2 as the pressure scale and $\rho U^2 d^2$ as the force scale. The hydrodynamic force and torque exerted on each particle denoted by \mathbf{F} and \mathbf{T} are given as

$$\mathbf{F} = \int_S \left[-p \mathbf{I} + \frac{1}{Re} (\nabla \mathbf{u} + \nabla \mathbf{u}^T) \right] \cdot \mathbf{n} \, dS, \quad (2.4a)$$

$$\mathbf{T} = \int_S \mathbf{r} \times \left[-p \mathbf{I} + \frac{1}{Re} (\nabla \mathbf{u} + \nabla \mathbf{u}^T) \right] \cdot \mathbf{n} \, dS, \quad (2.4b)$$

with S denoting the surface enclosing the solid body, \mathbf{I} the identity matrix, $(\cdot)^T$ the matrix transpose, \mathbf{n} the unit vector normal to the boundary of the solid body and \mathbf{r} the position vector relative to the particle centre of mass.

2.2. Numerical method

As our PR-DNS tool, we use PeliGRIFF (Parallel Efficient Library for GRains in Fluid Flow) which is a multiphase flow solver based on the distributed Lagrange multiplier-fictitious domain (DLM-FD) formulation proposed by Glowinski *et al.* (1999). In our implementation, we employ a finite-volume staggered-grid scheme for the fluid conservation equations (Wachs *et al.* 2015). Similar to the IBM, particles are immersed in the fluid domain in the DLM-FD method and rigid-body constraints on the fictitious fluid inside the solid region are enforced through a set of Lagrange multipliers collocated in the particle domain. For a fixed array of particles, the combined momentum and continuity equations in a non-variational form are given as

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} - \boldsymbol{\lambda} \quad \text{in } \mathbb{D}, \quad (2.5a)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \mathbb{D}, \quad (2.5b)$$

$$\mathbf{u} = \mathbf{0} \quad \text{in } \mathbb{P}, \quad (2.5c)$$

where the solid domain is denoted by \mathbb{P} and the fluid–particle domain by \mathbb{D} . Furthermore, $\boldsymbol{\lambda}$ shows the distributed Lagrange multiplier vector which is used to enforce the rigid-body

motion constraint (2.5c). For the temporal discretisation, we employ a first-order Marchuk–Yanenko operator-splitting algorithm. At each time t^{n+1} , we solve:

- (i) A classical projection scheme for the solution of the Navier–Stokes problem: find $\mathbf{u}^{n+1/2}$ and p^{n+1} such that

$$\begin{aligned} \frac{\tilde{\mathbf{u}}^{n+1/2} - \mathbf{u}^n}{\Delta t} - \frac{1}{2Re} \nabla^2 \tilde{\mathbf{u}}^{n+1/2} \\ = -\nabla p^n + \frac{1}{2Re} \nabla^2 \mathbf{u}^n - \frac{1}{2} (3\mathbf{u}^n \cdot \nabla \mathbf{u}^n - \mathbf{u}^{n-1} \cdot \nabla \mathbf{u}^{n-1}) - \lambda^n, \end{aligned} \quad (2.6a)$$

$$\left. \begin{aligned} \nabla^2 \psi^{n+1} &= \frac{1}{\Delta t} \nabla \cdot \tilde{\mathbf{u}}^{n+1/2}, \\ \frac{\partial \psi^{n+1}}{\partial n} &= 0 \quad \text{on } \partial \mathbb{D}, \end{aligned} \right\} \quad (2.6b)$$

$$\mathbf{u}^{n+1/2} = \tilde{\mathbf{u}}^{n+1/2} - \Delta t \nabla \psi^{n+1}, \quad (2.6c)$$

$$p^{n+1} = p^n + \psi^{n+1} - \frac{\Delta t}{2Re} \nabla^2 \psi^{n+1}. \quad (2.6d)$$

- (ii) A fictitious domain problem: find \mathbf{u}^{n+1} and λ^{n+1} such that

$$\frac{\mathbf{u}^{n+1} - \mathbf{u}^{n+1/2}}{\Delta t} + \lambda^{n+1} = \lambda^n, \quad (2.7a)$$

$$\mathbf{u}^{n+1} = \mathbf{0} \quad \text{in } \mathbb{P}, \quad (2.7b)$$

where Δt denotes the time step, ψ the pseudo-pressure and $\partial \mathbb{D}$ the domain boundary. In (2.6a), second order in time Crank–Nicolson and Adams–Bashforth schemes are used to discretise the viscous and advective terms, respectively, and the saddle-point problem in step (ii) is handled by an Uzawa algorithm (Wachs *et al.* 2015). Considering the high-order correction of the pressure, the projection scheme in step (i) is also second-order accurate in time. However, the first-order time discretisation of the fictitious domain sub-problem in step (ii) and the first-order Marchuk–Yanenko method reduce the global time accuracy of our algorithm to first order only. Equations presented in step (i) are spatially discretised with a second-order central scheme for the diffusion term, whereas the advective term is treated with a total variation diminishing scheme combined with a Superbee flux limiter. Despite the second-order discretisation of the flow equations, the accuracy of our method is between first and second order due to the presence of rigid bodies immersed within the domain. It can be shown that with the DLM-FD method, the Lagrange multiplier λ can be directly integrated over the volume of each particle to obtain the hydrodynamic force and torque acting on the particle \mathbb{P}

$$\mathbf{F} = \int_{\mathbb{P}} \lambda \, d\mathbf{x}, \quad (2.8a)$$

$$\mathbf{T} = \int_{\mathbb{P}} \mathbf{r} \times \lambda \, d\mathbf{x}. \quad (2.8b)$$

ϕ	Re	$d/\Delta x$	L	N_p
0.1	2	24	25	2984
0.1	10	24	25	2984
0.1	40	24	25	2984
0.1	150	32	25	2984
0.2	2	24	20	3055
0.2	40	32	20	3055
0.2	150	40	20	3055
0.4	2	32	15	2578
0.4	40	40	15	2578
0.4	150	48	15	2578

TABLE 1. Summary of the parameters used for PR-DNS of fixed beds of spheres.

2.3. Simulation set-up

A summary of all considered cases is given in [table 1](#). For our PR-DNS simulations, we consider triply periodic cubic domains of edge length L containing N_p spherical particles each taking up a volume of $v_p = \pi/6$, corresponding to a solid volume fraction of $\phi = Nv_p/L^3$. A constant flow rate so as to attain the desired Reynolds number is imposed in the x direction using a dynamically adjusted pressure drop. The x direction hence corresponds here to the streamwise direction. As noted by Akiki *et al.* (2016), Tenneti *et al.* (2011) demonstrated that using a domain size of only $2.4d$ guarantees the decorrelation of fluid velocities for $(Re, \phi) = (20, 0.2)$ and $(Re, \phi) = (300, 0.2)$. Our computational domains given in [table 1](#) all extend far beyond $2.4d$ containing ≈ 2500 – 3000 particles each, ensuring both the decay of fluid correlations and statistical reliability. Initialisation of particle locations for cases where $\phi \in \{0.1, 0.2\}$ is performed by distributing spheres randomly in the domain according to a random number generation algorithm without allowing any overlap. For the highly dense cases with $\phi = 0.4$, we start with a structured array of particles where each particle is given a random translational and rotational velocity, and we let the system reach an asymptotic motionless state through dissipative collisions. For all solid volume fractions, pair correlation functions have been obtained and verified to be satisfactorily matching theoretical radial distribution functions of hard spheres (Percus & Yevick 1958; Wertheim 1963). Visualisations of two sample cases in the present study are shown in [figure 2](#). The time step is taken to be smaller than $\Delta t = 2 \times 10^{-3}$ in all cases, ensuring time accuracy of the simulations along with satisfying the Courant–Friedrichs–Lewy (CFL) condition; namely, $CFL < 0.4$ for the spatial resolutions presented in [table 1](#). The simulations have been run until steady state is achieved, and force and torque data are all collected from the steady part. For $Re = 150$ at all volume fractions, the flow becomes time dependent and oscillating. This occurs because particles in close proximity act as a single obstacle and hence increase the effective length scale and the Reynolds number. Data from these cases are collected after a statistically pseudo-stationary state is established by averaging over sufficiently extended time intervals. In terms of computational demand, the mesh resolution and the domain size in each case along with the load per CPU core dictate resource requirements. With each processor core handling 512×10^3 grid cells, the smallest simulations with $\approx 110 \times 10^6$ grid cells were run on 192 cores, whereas the largest ones with 512×10^6




FIGURE 2. Fixed beds of spherical particles at $Re = 40$ and solid volume fractions of $\phi = 0.1$ (a) and $\phi = 0.4$ (b). The streamlines shown are coloured with respect to the fluid velocity magnitude.

grid cells required 960 cores. All computations were carried out on groups of 48-core nodes each equipped with ≈ 187 GB of memory, provided by the Cedar supercomputer as a part of Compute Canada's advanced research computing infrastructure (<http://www.computecanada.ca>).

2.4. Validation

The aim in this section is to provide a comparison of the results of our PR-DNS simulations with the reported drag correlations in the literature (Beetstra *et al.* 2007; Tenneti *et al.* 2011; Bogner *et al.* 2015; Tang *et al.* 2015). In terms of spatial resolution, our PR-DNS is comparable with the work of Bogner *et al.* (2015) and Tenneti *et al.* (2011) for the range of (Re, ϕ) considered here, but generally better resolved than simulations done by Tang *et al.* (2015) and Beetstra *et al.* (2007). Nevertheless, validations are presented merely to establish the accuracy of our code rather than intending to offer benchmarking data. The ensemble-average drag is obtained by summing the drag force acting on each particles in the bed and dividing by the total number of particles N_p

$$\langle F_d \rangle = \frac{1}{N_p} \sum_{i=1}^{N_p} \mathbf{F}_i \cdot \hat{\mathbf{e}}_x. \quad (2.9)$$

Obviously, the normal components of the mean force $\langle F_y \rangle = \langle \mathbf{F} \cdot \hat{\mathbf{e}}_y \rangle$ and $\langle F_z \rangle = \langle \mathbf{F} \cdot \hat{\mathbf{e}}_z \rangle$ are both expected to be vanishingly small. In order to compare our data with the available drag correlations, normalisation of forces is done with respect to the Stokes drag given as

$$F_{st} = 3\pi\mu dU. \quad (2.10)$$

As shown in figure 3, the drag computed for our cases listed in table 1 all lie within the range of existing correlations and generally indicate good accordance. In particular, our data seem to agree well with correlations proposed by Bogner *et al.* (2015) and Tang *et al.* (2016) at $Re = 150$ where significant discrepancies are observed between different correlations. In addition, we have also examined force distributions and the extent of data scatter for each case in table 2. The magnitude of the ensemble-average lift force given as $\langle F_L \rangle = \langle F_y \rangle$ or $\langle F_z \rangle$ is practically zero for all cases, as expected. However, the ratio of the

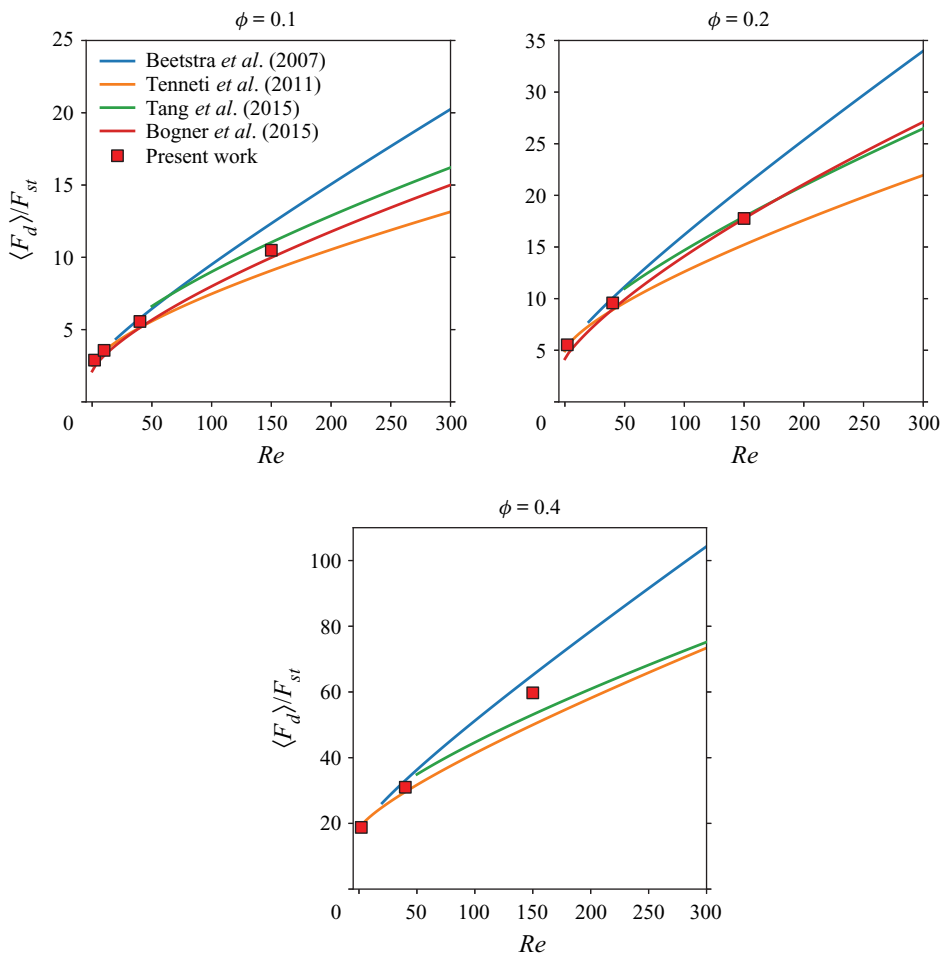


FIGURE 3. Validation of the drag data obtained from PR-DNS of the present work with correlations of Beetstra *et al.* (2007), Tenneti *et al.* (2011), Tang *et al.* (2015) and Bogner *et al.* (2015).

standard deviation of drag and lift shown by $\sigma_{F_d} / \langle F_d \rangle$ and $\sigma_{F_L} / \langle F_d \rangle$ are both significant and $\sigma_{F_d} / \langle F_d \rangle$ turns out to be always greater than $\sigma_{F_L} / \langle F_d \rangle$. In cases where a rough comparison is possible, our results for $\sigma_{F_d} / \langle F_d \rangle$ and $\sigma_{F_L} / \langle F_d \rangle$ in table 2 show a difference of $\approx 2\% - 4\%$ with those obtained by Akiki *et al.* (2016). A contributing feature that might explain the difference is that our sample sizes for each case are significantly larger by a factor of 5–6, rendering the results statistically more converged. This is obvious from the $\langle F_L \rangle$ values in table 2 compared with those reported by Akiki *et al.* (2016), as their mean lift data (while still quite close to zero) are orders of magnitude greater than present results. According to Akiki *et al.* (2016), realisation dependence alone could cause variance in $\sigma_{F_d} / \langle F_d \rangle$ and $\sigma_{F_L} / \langle F_d \rangle$ of up to 2.7% and 1.7%, respectively.

2.5. Dataset construction

For the purpose of analysis and development of our model, we first need to construct datasets corresponding to each simulation. Each dataset contains as many rows as the

ϕ	Re	$\langle F_d \rangle / F_{st}$	$\sigma_{F_d} / \langle F_d \rangle$	$\langle F_L \rangle / F_{st}$	$\sigma_{F_L} / \langle F_d \rangle$
0.1	2	2.89	21.32 %	2.7×10^{-10}	15.14 %
0.1	10	3.57	24.86 %	1.6×10^{-9}	14.78 %
0.1	40	5.56	26.71 %	3.0×10^{-5}	15.19 %
0.1	150	10.49	26.09 %	7.1×10^{-4}	16.48 %
0.2	0.2	5.48	17.66 %	7.5×10^{-5}	15.21 %
0.2	2	5.53	17.56 %	-9.3×10^{-9}	14.97 %
0.2	40	9.59	23.77 %	2.3×10^{-10}	16.33 %
0.2	150	17.76	25.97 %	-4.2×10^{-6}	18.86 %
0.4	2	18.71	22.96 %	3.7×10^{-8}	18.22 %
0.4	40	30.99	26.07 %	5.3×10^{-7}	18.84 %
0.4	150	59.73	27.68 %	2.3×10^{-5}	19.62 %

TABLE 2. Statistics of the drag and lift data for cases presented in table 1.

number of particles N_p (which we call samples), while columns represent the input and output variables. For each particle i , the set of position vectors of the neighbouring spheres denoted by $\{\mathbf{r}_{j=1}, \dots, \mathbf{r}_{j=M}\}$ along with the average fluid velocity $\bar{\mathbf{u}}_i$ around particle i constitute the inputs, whereas the hydrodynamic forces and torques experienced by particle i are the outputs that we aim to model. The position vectors of neighbours are expressed relative to the location of particle i . The first 30 nearest neighbours are identified for each particle i by looping over all other particles in the simulation. These neighbours are then numbered from $j = 1$ to $j = 30$ depending on their relative distance, with $j = 1$ being the closest. It is worthwhile noting that, since our simulations are performed in tri-periodic domains, forces and torques on particles near boundaries could be affected by periodic images of particles on the opposite side of the domain whose positions do not exist in the simulation data. For this reason, prior to performing the neighbour identification loop, we create the periodic image positions manually and explicitly insert them in the data so that the periodic effects are correctly accounted for.

3. Probability-driven model

3.1. Motivation

As pointed out in § 1, we seek to predict the hydrodynamic forces and torques on each particle as a function of the flow conditions, and more importantly, of the unique neighbourhood of each particle as highlighted by (1.2). The hydrodynamic force and torque exerted on each sphere may be decomposed and expressed as follows:

$$\mathbf{F}_i = \langle \mathbf{F}_i \rangle (Re, \phi) + \Delta \mathbf{F}_i (Re, \phi, \{\mathbf{r}_{j=1}, \dots, \mathbf{r}_{j=M}\}), \quad (3.1a)$$

$$\mathbf{T}_i = \Delta \mathbf{T}_i (Re, \phi, \{\mathbf{r}_{j=1}, \dots, \mathbf{r}_{j=M}\}), \quad (3.1b)$$

where $\Delta \mathbf{F}_i$ and $\Delta \mathbf{T}_i$ show the fluctuating contribution to the force and torque due to the specific arrangement of neighbours surrounding particle i . Statistically, the lateral components of $\langle \mathbf{F}_i \rangle$ and all components of $\langle \mathbf{T}_i \rangle$ are close to zero for a sufficiently large number of particles in each simulation. In the streamwise direction, the ensemble-average force $\langle \mathbf{F}_i \rangle$ is identical for all particles by definition and therefore only depends on macroscopic variables, i.e. the Reynolds number and solid volume fraction.

Many correlations exist that describe such a dependence as discussed in § 1. For this reason, we henceforth focus on ΔF_i and ΔT_i in (3.1); namely, the deviations from the mean arising due to the unique local neighbourhood of each particle.

The key question posed here may be expressed as follows: How does changing the location of each neighbour affect the force and torque experienced by a particle? The most obvious and naive strategy would be attempting multiple simulations where the location of a single neighbour particle is changed systematically and the resulting effect on forces and torques are recorded. The outcome would be similar to, for example, the two-dimensional interaction maps given by Zhou & Alam (2016) for two neighbouring cylinders characterising forces for various locations of a second cylinder. Assuming a three-dimensional system, adding only one or two more neighbours to the scenario makes the number of simulations needed for populating the modelling dataset an almost intractable computational task. For their physics-based PIEP model, Akiki *et al.* (2017a) needed to account for 15–40 closest neighbours to obtain good accuracy, which in three dimensions translates into 45–120 input variables, respectively. It is well known that in a data-driven framework, such systems are afflicted with the ‘curse of dimensionality’ (Hastie, Tibshirani & Friedman 2009). It means that constructing a reliable dataset for a regression-type model requires millions of samples for each case, while we are only able to model a few thousand particles with PR-DNS at best. The pairwise interactions and order-invariance approximations utilised by Moore *et al.* (2019) for their data-driven model is a good demonstration of a strategy for reducing the number of input variables to a manageable level.

3.2. Theory

Consider a monodisperse random array of spherical particles with a volume fraction ϕ subject to fluid flow with a Reynolds number Re . Let X and Y be two continuous random variable vectors in a d -dimensional space ($X, Y \in \mathbb{R}^d$) with marginal probability density functions (PDFs) $f_X(\mathbf{x})$ and $f_Y(\mathbf{y})$ where $f_X(\mathbf{x}), f_Y(\mathbf{y}) : \mathbb{R}^d \rightarrow \mathbb{R}$. By definition, PDFs in general are required to satisfy the following conditions:

$$\left. \begin{aligned} f_X(\mathbf{x}) &\geq 0, \\ \int f_X(\mathbf{x}) \, d\mathbf{x} &= 1, \end{aligned} \right\} \tag{3.2}$$

where the integration is performed on all possible values of \mathbf{x} . The value $f_X(\mathbf{x}) \, d\mathbf{x}$ represents the probability of X falling in the infinitesimal volume $d\mathbf{x}$ about \mathbf{x} . Therefore, the probability that the random variable X will happen to be in \mathcal{A} is given by

$$P[X \in \mathcal{A}] = \int_{\mathcal{A}} f_X(\mathbf{x}) \, d\mathbf{x}. \tag{3.3}$$

The same properties similarly hold for $f_Y(\mathbf{y})$. The expected value of the random variable X is the weighted average of all possible values of X , each value being weighted according to its probability of occurrence (Ross 2010). The expected value of X is thus governed by its probability density distribution $f_X(\mathbf{x})$, and is given as

$$E[X] = \langle X \rangle = \int \mathbf{x} f_X(\mathbf{x}) \, d\mathbf{x}. \tag{3.4}$$

The probability associated with a random variable may also be conditioned on another random variable. For this case, $f_{X|Y}(\mathbf{x} | \mathbf{y})$ may be defined as the conditional PDF of X

given that Y assumes a particular fixed value y .

$$E[X | Y = y] = \langle X | Y = y \rangle = \int x f_{x|y}(x | y) dx. \tag{3.5}$$

Referring to (3.1), the three components of the force/torque fluctuations pertaining to particle i are shown by $\Delta F_i = (\Delta F_{i,x}, \Delta F_{i,y}, \Delta F_{i,z})^T$ or $\Delta T_i = (\Delta T_{i,x}, \Delta T_{i,y}, \Delta T_{i,z})^T$. A particular configuration of the neighbourhood surrounding particle i is denoted by $\mathbf{R}_i = \{\mathbf{r}_{j=1}, \dots, \mathbf{r}_{j=M}\}$. Here, we present the analysis only for force fluctuations; namely, ΔF , since the derivation of similar equations is similarly done for torque fluctuations ΔT . From this point onwards, the subscript i is dropped from all variables in order to minimise cluttering in the notations. Upon setting $X = \Delta F$ and $Y = \mathbf{R}$, the following PDF

$$p_{\Delta F|\mathbf{R}}(\Delta F | \mathbf{R}) \tag{3.6}$$

represents the probability distribution of the force/torque fluctuations given the locations of M neighbours. In the above equations, \mathbf{R} shows the set of position vectors $\{\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \dots, \mathbf{r}_M\}$ belonging to the neighbours numbered according to their distance from a reference particle, starting with $j = 1$ being the closest. Each distribution $p(\Delta F | \mathbf{R})$ is to be obtained from collecting ΔF values experienced by particles that happen to have the same neighbour configuration \mathbf{R} , while no constraints are imposed on neighbours $\{\mathbf{r}_j : j > M\}$ located further away. The expected value of the PDF in (3.6) gives a prediction of each fluctuating force/torque component as follows:

$$E[\Delta F | \mathbf{R}] = \langle \Delta F | \mathbf{R} \rangle = \int \xi p_{\Delta F|\mathbf{R}}(\Delta F | \mathbf{R}) d\xi. \tag{3.7}$$

Two extreme cases of M , the number of included neighbours, are worth elaborating on.

- (a) $M = 0$: this means that no conditions are imposed on p , hence $E[\Delta F] = 0$ and the predicted force/torque is the same for all particles in the array. Such a situation corresponds to, for instance, the conventional microstructure-ignorant correlations of the form $\mathbf{F} \cdot \hat{\mathbf{e}}_x = \langle F_x \rangle (Re, \phi)$ in case of the drag force.
- (b) $M \rightarrow N_p - 1$: the positions of all neighbours in the particle array are imposed as a condition on p , hence the variance of p tends to zero and the value of ΔF is uniquely determined.

Since the problem is fully constrained in the latter case, the predicted value of ΔF would be free of uncertainty. Put another way, p would no longer be a distribution *per se*, but rather resembles a Dirac delta function instead:

$$\lim_{M \rightarrow N_p - 1} E[\Delta F | \mathbf{R}] = \int_{-\infty}^{+\infty} \xi \delta_{\Delta F|\mathbf{R}}(\xi - \Delta F_{DNS} | \mathbf{R}) d\xi = \Delta F_{DNS}, \tag{3.8}$$

which is equivalent of having a dataset populated with an infinite number of samples. In other words, this ‘ideal’ dataset would contain samples for each and every combination of neighbour locations. Such a dataset would conceivably provide a full and exact description of the problem, thus representing a solution to the Navier–Stokes equations. Consequently, (3.6) can be viewed as a data-driven description of forces/torques exerted on each particle within an array derived from probabilistic arguments, which is potentially capable of predicting fluctuations with varying degrees of accuracy, depending on how constrained we force p to be. It is crucial to realise that unless $M = N_p - 1$, our best estimate of the

fluctuating component $E[\Delta F | \mathbf{R}]$ would be an average (i.e. the central tendency) of the distribution, implying that for $M < N_p - 1$, influence of neighbours whose positions are not constrained is only accounted for in a statistical sense. Having established the notion and properties of the conditioned force/torque PDFs, the probability-driven prediction of a fluctuating component based on its corresponding distribution is given as

$$\Delta F_{MPP} = E[\Delta F | \mathbf{R}] = \int_{-\infty}^{+\infty} \xi p_{\Delta F|\mathbf{R}}(\xi | \mathbf{R}) d\xi. \tag{3.9}$$

The integral above may be split to cover the positive and negative force/torque contributions separately, yielding

$$\Delta F_{MPP} = \int_{-\infty}^0 \xi p_{\Delta F|\mathbf{R}}(\xi | \mathbf{R}) d\xi + \int_0^{+\infty} \xi p_{\Delta F|\mathbf{R}}(\xi | \mathbf{R}) d\xi. \tag{3.10}$$

The impeding hurdle with (3.10) is that computing ΔF_{MPP} requires *a priori* knowledge of the force/torque distributions for each and every possible configuration \mathbf{R} . Such an approach is clearly not feasible and certainly would not fulfil the goal of constructing a model based on a limited dataset. An elegant alternative is to take advantage of Bayes' theorem for probabilities, which serves to convert a conditional probability problem to its reverse case. According to the Bayes' formula for distributions of random variables (Papoulis, Pillai & Pillai 2002), $f_{X|Y}(\mathbf{x} | \mathbf{y}) = f_{Y|X}(\mathbf{y} | \mathbf{x})f_X(\mathbf{x})/f_Y(\mathbf{y})$. Therefore, it follows that

$$p_{\Delta F|\mathbf{R}}(\Delta F | \mathbf{R}) = \frac{p_{\mathbf{R}|\Delta F}(\mathbf{R} | \Delta F) p_{\Delta F}(\Delta F)}{p_{\mathbf{R}}(\mathbf{R})}. \tag{3.11}$$

Substituting for $p_{\Delta F|\mathbf{R}}(\xi | \mathbf{R})$ in first integral in (3.10) yields

$$\frac{1}{p_{\mathbf{R}}(\mathbf{R})} \int_{-\infty}^0 \xi p_{\mathbf{R}|\Delta F}(\mathbf{R} | \xi) p_{\Delta F}(\xi) d\xi. \tag{3.12}$$

Assuming the continuity of both $p_{\mathbf{R}|\Delta F}$ and $p_{\Delta F}$ and realising that $p_{\mathbf{R}|\Delta F}, p_{\Delta F} \geq 0$, the generalised mean value theorem for integrals can be applied to the above equation to give

$$\frac{\alpha p_{\Delta F}(\alpha)}{p_{\mathbf{R}}(\mathbf{R})} \int_{-\infty}^0 p_{\mathbf{R}|\Delta F}(\mathbf{R} | \xi) d\xi = \left[\frac{\alpha p_{\Delta F}(\alpha)}{p_{\mathbf{R}}(\mathbf{R})} \right] p_{\mathbf{R}|\Delta F}(\mathbf{R} | \Delta F < 0), \tag{3.13}$$

where $\xi = \alpha$ is a constant value that belongs to the interval $(-\infty, 0]$. Following the same operations for the second integral of (3.10), there exists a value $\xi = \beta$ in $[0, +\infty)$ for which $\xi p_{\Delta F}(\xi)$ can be brought out of the integral. Substituting results back in (3.10), we arrive at the following:

$$\Delta F_{MPP} = \left[\frac{\alpha p_{\Delta F}(\alpha)}{p_{\mathbf{R}}(\mathbf{R})} \right] p_{\mathbf{R}|\Delta F}(\mathbf{R} | \Delta F < 0) + \left[\frac{\beta p_{\Delta F}(\beta)}{p_{\mathbf{R}}(\mathbf{R})} \right] p_{\mathbf{R}|\Delta F}(\mathbf{R} | \Delta F > 0). \tag{3.14}$$

Note that, according to (3.5), $p_{\mathbf{R}|\Delta F}(\mathbf{R} | \Delta F < 0)$ gives the PDF of particle positions provided that the force/torque fluctuations obey $\Delta F < 0$ (same applies to $p_{\mathbf{R}|\Delta F}(\mathbf{R} | \Delta F > 0)$ when $\Delta F > 0$). Since particles are assumed to be randomly, yet uniformly distributed, $p_{\mathbf{R}}(\mathbf{R})$ is expected to have an approximately constant distribution. This can be verified by inspection of unconditioned particle location PDFs for

various neighbours. Therefore, both terms inside the square brackets in (3.14) are constant values and can be expressed as two coefficients c_α and c_β . The final form of the force/torque model reads

$$\Delta F_{MPP}(\mathbf{R}, Re, \phi) = c_\alpha \tilde{p}(\mathbf{R} | \Delta F < 0) + c_\beta \tilde{p}(\mathbf{R} | \Delta F > 0), \quad (3.15)$$

where \tilde{p} represents the conditional PDF $p_{\mathbf{R}|\Delta F}$. It should be noted that the \tilde{p} probability maps depend on the Reynolds number and solid volume fraction, which is why ΔF_{MPP} is shown as $\Delta F_{MPP}(\mathbf{R}, Re, \phi)$. Thus far, no approximation was involved in the derivation of ΔF_{MPP} , which means that in case a full knowledge of the aforementioned distributions exists, (3.15) would be theoretically exact. Practically, however, for a dataset with a few thousand samples the probability distributions $\tilde{p}(\mathbf{R} | \Delta F < 0)$ or $\tilde{p}(\mathbf{R} | \Delta F > 0)$ would be extremely sparse, as alluded to in § 3.1. Taking N and D to denote sample size and number of input dimensions, the data density would be proportional to $N^{1/D}$. As D increases, it is straightforward to see that maintaining a constant sample density requires exponentially more data points. In our case, each neighbour adds three inputs to the problem, each representing one component of the position vector $\mathbf{r}_j = (x_j, y_j, z_j)$. If, for instance, $N_1 = 1000$ is deemed a sufficiently dense dataset to form reliable local averages when $D_1 = M \times 3 = 1 \times 3$ (i.e. a single neighbour considered), in case of 15 neighbours $D_2 = 15 \times 3 = 45$, thus $N_2 = N_1^{(D_2/D_1)} = 1000^{15}$. Ordinary PR-DNS-generated simulations of particle arrays are far from being sufficiently dense for creating meaningful, reliable functional forms for \tilde{p} . As a result, the curse of dimensionality precludes the construction of PDFs in (3.15) in their current form. Any such effort is in fact bound to produce a severely over-fitted model with almost no generalisable prediction capability. Consequently, instead of attempting to obtain $\tilde{p}(\mathbf{R} | \Delta F < 0)$ and $\tilde{p}(\mathbf{R} | \Delta F > 0)$ directly, we suggest that

$$\left. \begin{aligned} \tilde{p}(\mathbf{R} | \Delta F < 0) &\approx c_{\alpha,1} \tilde{p}_1(\mathbf{r}_1 | \Delta F < 0) + c_{\alpha,2} \tilde{p}_2(\mathbf{r}_2 | \Delta F < 0) + \dots \\ &= \sum_{j=1}^M c_{\alpha,j} \tilde{p}_j(\mathbf{r}_j | \Delta F < 0), \\ \tilde{p}(\mathbf{R} | \Delta F > 0) &\approx c_{\beta,1} \tilde{p}_1(\mathbf{r}_1 | \Delta F > 0) + c_{\beta,2} \tilde{p}_2(\mathbf{r}_2 | \Delta F > 0) + \dots \\ &= \sum_{j=1}^M c_{\beta,j} \tilde{p}_j(\mathbf{r}_j | \Delta F > 0). \end{aligned} \right\} \quad (3.16)$$

In the above equation, each function $\tilde{p}_j(\mathbf{r}_j | \cdot)$ is equivalent to a marginal probability distribution which is obtained by integrating out the position \mathbf{r}_k of each and every neighbour where $k \neq j$ in $\tilde{p}(\mathbf{R} = \{\mathbf{r}_{k=1}, \dots, \mathbf{r}_{k=M}\} | \cdot)$. In other words, this approximation only considers the effect of a single neighbour on the functional form of \tilde{p} , while accounting for the presence of all other surrounding particles in an average manner. This marginal distribution can be expressed as

$$\tilde{p}_j(\mathbf{r}_j | \cdot) = \int \tilde{p}(\mathbf{R} | \cdot) \, d\mathbf{r}_1 \dots d\mathbf{r}_{j-1} \, d\mathbf{r}_{j+1} \dots d\mathbf{r}_M. \quad (3.17)$$

Equation (3.16) is reminiscent of the pairwise approximation employed by Moore *et al.* (2019) for the purpose of reducing the number of independent variables by accounting for the influence of only one neighbour at a time. In their data-driven model, the authors also invoked the ‘order-invariance’ approximation, which removes the dependence of the model functions on the neighbour number. This was done by weighting the functionals

according to the probability of a particular neighbour being the j th closest neighbour. In the MPP model, the neighbours first need to be sorted based on their distance to the reference particle. Each neighbour's position can be then passed to the appropriate PDF.

Substitution of (3.16) in (3.15) leads to

$$\Delta F_{MPP}(\mathbf{R}, Re, \phi) = \sum_{j=1}^M c_{\alpha,j} \tilde{p}_j(\mathbf{r}_j | \Delta F < 0) + \sum_{j=1}^M c_{\beta,j} \tilde{p}_j(\mathbf{r}_j | \Delta F > 0). \quad (3.18)$$

In (3.18), $c_{\alpha,j}$ and $c_{\beta,j}$ are constant unknown coefficients. With the final form of ΔF_{MPP} , we have assumed an additive nature for the effect that each neighbour might have on the force/torque fluctuation experienced by a reference particle. Notably, both distributions $\tilde{p}_j(\mathbf{r}_j | \cdot)$ (this notation is used to refer to a PDF with any arbitrary force/torque conditioning) are now functions of three variables only; that is, the three components of the position vector $\mathbf{r}_j = (r_{j,x}, r_{j,y}, r_{j,z})$. This means that the functional form of the PDFs in (3.18) can be conveniently inferred with a data-driven approach, as the high-dimensionality of the input space has been evaded. To this end, a dataset can be constructed by running PR-DNS of fixed beds for desired sets of parameters. The discrete estimation of the distribution \tilde{p}_j can be extracted by filtering neighbour positions according to the conditions $\{\Delta F < 0, \Delta F > 0\}$. The discrete form of the distribution is then fitted with a multivariate kernel density estimation (KDE) function (Chacón & Duong 2018) of the form

$$\tilde{p}(\mathbf{r} | \cdot) = \frac{1}{m} \sum_{q=1}^m K_{\mathbf{H}}(\mathbf{r} - \mathbf{r}_q), \quad (3.19)$$

where

$$K_{\mathbf{H}}(\mathbf{r}) = |\mathbf{H}|^{-1/2} K(\mathbf{H}^{-1/2} \mathbf{r}). \quad (3.20)$$

In the above equations, m shows the total number of samples after filtering with the aforementioned conditions, \mathbf{H} the bandwidth matrix and K is a symmetric density function of choice. We select a Gaussian distribution as the kernel function defined as

$$K(\mathbf{r}) = \frac{1}{(2\pi)^{D/2}} \exp\left\{-\frac{1}{2} \mathbf{r}^T \mathbf{r}\right\}. \quad (3.21)$$

We will demonstrate and discuss particular examples of these functions in § 3.3. Now that the appropriate functional forms are identified and known, it remains to decide the values of the constant coefficients $c_{\alpha,j}$ and $c_{\beta,j}$ in (3.18). The values are obtained using the ordinary least-squares method for linear regression

$$\{c_{\alpha,j}, c_{\beta,j}\}_{j=1}^M = \operatorname{argmin} \sum_{i=1}^{N_p} (\Delta F_{MPP,i} - \Delta F_{DNS,i})^2, \quad (3.22)$$

which minimises the residual sum of squares between values of force/torque fluctuations predicted by the MPP model and the true values from PR-DNS. After optimal values for $c_{\alpha,j}$ and $c_{\beta,j}$ are found, (3.18) can be used to make predictions for each particle based on the particular microstructure of the surrounding particles. It is of great importance to realise that in the presented approach, the model relies solely on the unique configuration of neighbours surrounding each particle, and the PR-DNS-generated force/torque data (as opposed to velocity and pressure fields) in order to make predictions. Consequently, the

MPP model accounts for the combined effect of various force contributions such as the undisturbed flow, quasi-steady and added-mass forces. This can be contrasted with the PIEP physics-driven model (Akiki *et al.* 2017a) which uses perturbed velocity and pressure fields together with the Faxén's theorem to obtain hydrodynamic forces experienced by each particle.

The above analysis applies identically to all force and torque fluctuation components; namely, $(\Delta F_x, \Delta F_y, \Delta F_z)$ and $(\Delta T_x, \Delta T_y, \Delta T_z)$. Let C_α denote a $1 \times M$ vector that contains the $c_{\alpha,j}$ coefficients in (3.18)

$$C_\alpha = (c_{\alpha,1} \quad \cdots \quad c_{\alpha,j} \quad \cdots \quad c_{\alpha,M}) \quad (3.23)$$

and let $P_{\Delta F < 0}$ denote an $M \times 1$ vector that contains the PDFs $\tilde{p}_j(\mathbf{r}_j | \Delta F < 0)$ in (3.18)

$$P_{\Delta F < 0} = \begin{pmatrix} \tilde{p}_1(\mathbf{r}_1 | \Delta F < 0) \\ \vdots \\ \tilde{p}_j(\mathbf{r}_j | \Delta F < 0) \\ \vdots \\ \tilde{p}_M(\mathbf{r}_M | \Delta F < 0) \end{pmatrix}. \quad (3.24)$$

Together with their counter-parts, i.e. C_β and $P_{\Delta F > 0}$, the above vectors can be ultimately defined for all force components. The complete model equations for force fluctuations can thus be expressed in the following form:

$$\Delta F_{MPP}(\mathbf{R}, Re, \phi) = \begin{pmatrix} \Delta F_x \\ \Delta F_y \\ \Delta F_z \end{pmatrix} = \begin{pmatrix} C_\alpha^x P_{\Delta F_x < 0} + C_\beta^x P_{\Delta F_x > 0} \\ C_\alpha^y P_{\Delta F_y < 0} + C_\beta^y P_{\Delta F_y > 0} \\ C_\alpha^z P_{\Delta F_z < 0} + C_\beta^z P_{\Delta F_z > 0} \end{pmatrix}, \quad (3.25)$$

where the superscripts of the coefficients indicate the corresponding component. Following the above definitions, torque fluctuations are given as

$$\Delta T_{MPP}(\mathbf{R}, Re, \phi) = \begin{pmatrix} \Delta T_x \\ \Delta T_y \\ \Delta T_z \end{pmatrix} = \begin{pmatrix} D_\alpha^x P_{\Delta T_x < 0} + D_\beta^x P_{\Delta T_x > 0} \\ D_\alpha^y P_{\Delta T_y < 0} + D_\beta^y P_{\Delta T_y > 0} \\ D_\alpha^z P_{\Delta T_z < 0} + D_\beta^z P_{\Delta T_z > 0} \end{pmatrix}, \quad (3.26)$$

where D_α , D_β , $P_{\Delta T > 0}$ and $P_{\Delta T < 0}$ for each torque component are equivalent to C_α , C_β , $P_{\Delta F > 0}$ and $P_{\Delta F < 0}$ in (3.25). Note that the vectors C_α , C_β , D_α , D_β , $P_{\Delta F < 0}$, $P_{\Delta F > 0}$, $P_{\Delta T < 0}$, $P_{\Delta T > 0}$ each contain as many coefficients and PDFs as the number of included neighbours M , the values of which depend on Re and ϕ of each case represented in table 1.

3.3. Probability distribution maps

In this section, we scrutinise a few of the probability distribution maps obtained for different force/torque components and various neighbours. Before proceeding further, we first describe the necessary steps taken for generating the foregoing distributions. In order to construct the maps, those particles that satisfy the desired condition (e.g. $\Delta F_x < 0$ or $|\Delta F_y| > \sigma$, with σ being the standard deviation of the corresponding variable) are identified and filtered by looping over the entire array of spheres. In this filtered subset of the original data, the positions of the chosen neighbours are recorded. These positions

constitute data points to which KDE functions in (3.19)–(3.21) are then fitted, thus giving $\tilde{p}_j(\mathbf{r}_j | \cdot)$. The generated PDFs are functions of the three spatial coordinates, which makes it impossible to visualise these functions directly. Inevitably, we resort to contour plots of the PDFs; however, three-dimensional contour surfaces are also neither efficient nor convenient to visualise. Consequently, we have opted for depicting their projections on two-dimensional planes. In doing so, the contour surfaces are entirely compressed onto an x – y plane. This is why the reference particle will be shown as a blank circle in contrast to a solid disk, since the neighbouring spheres can happen to be located even at $(x, y) = (0, 0)$ when $|z| \geq 0.5$.

In the case of streamwise force, all maps are in fact axisymmetric about the x axis; meaning that the choice of the plane on which the maps are shown does not make any difference, as long as it is parallel to the x axis. As for the lateral forces, the maps are plane symmetric with respect to the same plane whose normal has the same direction as the force component (e.g. $\tilde{p}(\mathbf{r}_j | |\Delta F_y| > \sigma)$ would be symmetric about the y axis). For the hydrodynamic torques, the maps are also plane symmetric but with respect to the plane that contains the torque component. Owing to the symmetry about the streamwise direction, the probability distribution map for $|\Delta F_y| > \sigma$ shown on the x – y plane is essentially almost identical to that for $|\Delta F_z| > \sigma$ on the y – z plane. The same is also true for $|\Delta T_z| > \sigma$ on x – y and $|\Delta T_y| > \sigma$ on y – z . Noting the existence of this similarity, only projections on the x – y plane will be shown in this section. With an ideal densely populated dataset, the PDFs are expected to exhibit perfect symmetry. The limited number of samples in our dataset, however, causes some PDFs to show minor yet noticeable deviations from symmetry. This would be particularly evident in PDFs generated for farther neighbours due to the lower data density compared with closer neighbours. In what follows, statistical symmetry is hence enforced in all PDF depictions.

Examination of the probability maps in all cases remarkably reveals quite distinct, non-uniform and physically meaningful functional forms. Figure 4 shows the distribution of the first (i.e. closest) neighbour positions, conditioned on two different ranges of the streamwise force, while the blank circle in the middle of represents the reference particle which experiences the hydrodynamic force or torque. According to these probability maps, it is considerably more likely for a reference particle to experience higher than the mean drag force when the closest neighbour is located laterally, and slightly downstream of the particle. The right plot in figure 4 appears more intuitive: drag force diminishes if the reference particle is shielded by the closest neighbour directly in front of it, as expected. A somewhat less intuitive observation is that even if the neighbour happens to be immediately downstream of the reference particle, drag force would be lower than average, due to the suppression of the low-pressure region behind the reference particle. It is crucial to stress this point once more that these maps demonstrate the higher likelihood of coming across neighbours at particular regions depending on the condition imposed on force/torque fluctuation. This means that, for instance, even if the closest neighbour occurs to be located laterally, the drag force might be lower than average because the second neighbour happens to directly shield the reference particle. It signifies the fact that while accounting for the influence of one neighbour provides a great deal of valuable information about the fluctuations, effects of other neighbours are indispensable in explaining the force/torque fluctuation accurately.

Figures 5 and 6 show PDFs for the distribution of closest neighbours when conditions are imposed on the lateral force in the y direction and on the lateral torque in the z direction, respectively. Note that for demonstration purposes, we have chosen more extreme conditions in figures 5 and 6, e.g. $\Delta F_y > \sigma$ and $\Delta F_y < -\sigma$, and $\Delta T_z > \sigma$ and




FIGURE 4. (a) Unconditioned PDF of the first closest neighbour position and PDF of the first closest neighbour position when the reference particle experiences a (b) higher than average or (c) lower than average drag force. The PDFs are obtained for the case of $Re = 40$ and $\phi = 0.1$.




FIGURE 5. The PDF of the first closest neighbour position when the reference particle experiences a (a) higher than σ or (b) lower than σ lift, with σ being the standard deviation of the data. (c) The PDF of the first closest neighbour position when the lift is either higher or lower than σ . The PDFs are obtained for the case of $Re = 40$ and $\phi = 0.1$.




FIGURE 6. The PDF of the first closest neighbour position when the reference particle experiences a (a) higher than σ and (b) lower than σ lateral torque, with σ being the standard deviation of the data. (c) The PDF of the first closest neighbour position when the lateral torque is either higher or lower than σ . The PDFs are obtained for the case of $Re = 40$ and $\phi = 0.1$.

$\Delta T_z < -\sigma$, respectively, in order to make the PDFs stand out more prominently. However, the MPP model ((3.25) and (3.26)) is always constructed with the original positive and negative conditioning (e.g. $\Delta F_y > 0$ and $\Delta F_y < 0$ for the lateral force in the y direction). On the leftmost plot of figure 5, the reference particle is seen to experience a positive lateral force in the y direction when the first neighbour is located within the region that

lies mostly in the bottom right quadrant where $x > 0$ and $y < 0$, whereas a negative lateral force results when the first neighbour happens to be on the opposite side where $x, y > 0$ depicted in the middle plot. As pointed out by Prahl *et al.* (2007) for a staggered arrangement of only two spheres, a higher than normal pressure field develops in the gap between the two particles, which causes a repulsive force acting mainly on the upstream particle. Upon increasing the distance between the particles, the fluid can penetrate the gap more easily. The nozzle effect hence creates a low-pressure region in the gap which gives rise to an attraction force. The latter does not appear in figure 5 as the first neighbour is quite close to the reference particle, but regions corresponding to such a condition show up remarkably on the PDFs for the lateral force in case we consider more distant neighbours. The right most plot shows the PDFs conditioned with $|\Delta F| > \sigma$, which is a combination of the foregoing plots with $\{\Delta F < -\sigma, \Delta F > \sigma\}$. These are regions for which the reference particle experiences a non-zero lateral force in the y direction.

In figure 6, we can see essentially the same kind of information as in figure 5. The most critical regions for the modification of the torque are located where the highest amount of vorticity is generated in the flow over a single sphere. As with the lateral force, the higher pressure in the gap decreases the fluid velocity and the resulting vorticity. The vorticity imbalance between the top and bottom regions of the reference particle brings about a net torque exerted on the particle. The interpretation of the PDFs emerging from our analysis is consistent with the observations made with a binary system of two spheres exhibiting the same dependency of the forces on relative positions of the spheres (Prahl *et al.* 2007; Yoon & Yang 2007). Such modification of the forces and torques were also elucidated by Akiki *et al.* (2016). Our results imply that even when the system being modelled is a dense multi-particle system, the forces and torques acting on each particle still vary in the same manner on average, in response to the positioning of the closest neighbour. For instance, when a particle happens to experience less drag force, it is most probably shielded by an upstream neighbour (or supported by a neighbour immediately downstream with a comparatively less probability) as shown on the right plot in figure 4. Although this may not be always true, it is the most likely reason. Similar arguments apply in regards to all other conditional PDFs involving lateral forces and torques as well.

Interestingly, similar patterns are found for neighbours located further away. An example of the PDFs for farther neighbours is shown in figure 7. While the distributions are spread out and have become fainter for surrounding particles located farther away from the reference particle, regions with higher probability are discernible and turn out to be generally in accordance with those obtained for $j = 1$, i.e. the closest neighbour. The reason for the fact that the densities are smaller for $j > 1$ is twofold. First, farther particles are distributed over a larger volume around the reference particle. The second and more important reason is that the more distant particles are comparatively less likely to alter the force/torque fluctuations. The functional forms of \tilde{p} in (3.16) has been kept distinct for each neighbour, meaning that \tilde{p}_1 is a different function compared with \tilde{p}_2 , and they both are different from \tilde{p}_3 and so on. This is also reflected in the probability maps shown in figure 7. A possible simplification can be made by combining all functions \tilde{p}_j into a single function \tilde{p} that still contains the contribution of all included neighbours, except for not making distinction based on the ordering of the neighbours. In this manner, instead of filtering the locations of the j th neighbour, we filter the location of all the included neighbours (up to $j = M$) based on the desired condition. It is to be noted that we still retain the contribution of each neighbour to the net force/torque fluctuation by allowing separate terms for every included neighbour in (3.16). That is to say, we construct only one PDF, namely, \tilde{p} using the entire neighbourhood, while each term $\tilde{p}(r_j | \cdot)$ is assigned a separate unknown coefficient and influences the prediction of fluctuations differently.

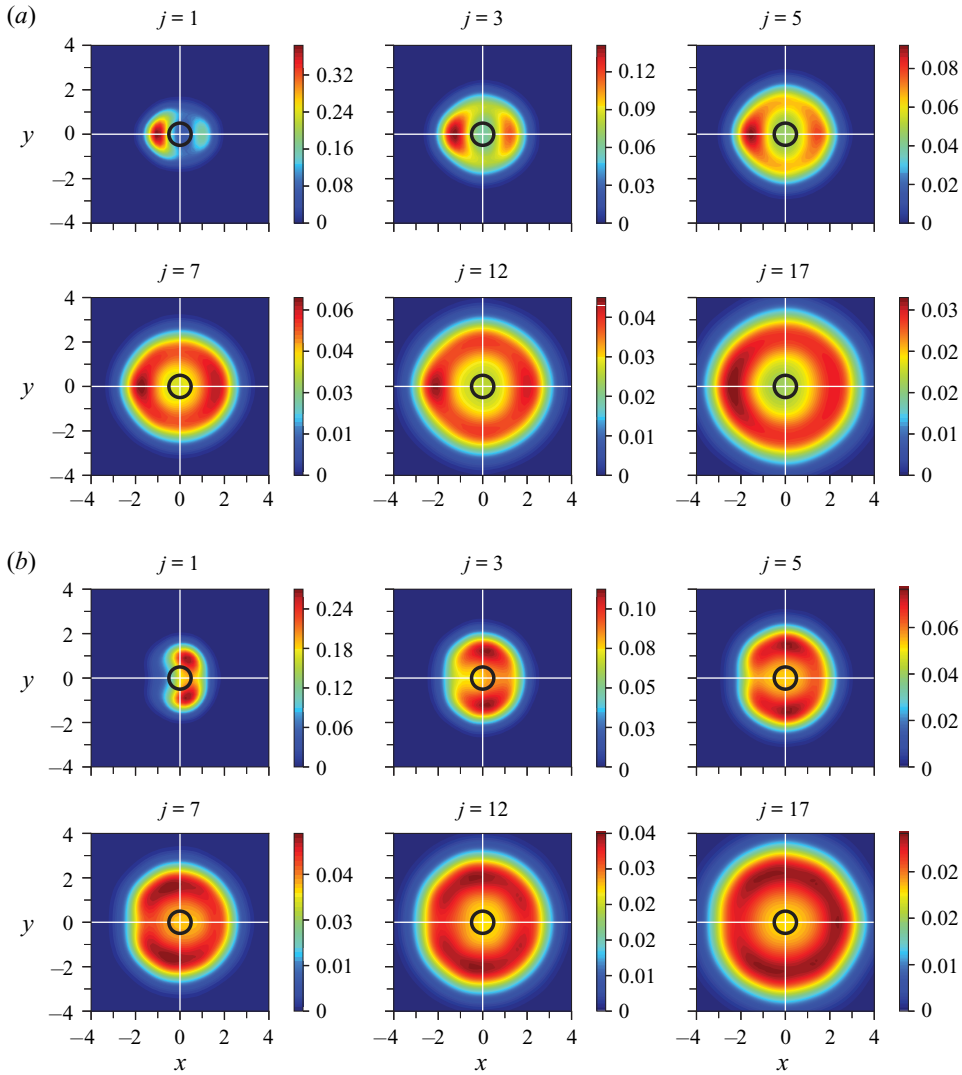


FIGURE 7. The PDFs of various neighbour positions for (a) $\Delta F_x < -\sigma$, and (b) $\Delta F_x > \sigma$, with σ being the standard deviation of the data. Note that the numbering represented by j is based on proximity to the reference particle, where $j = 1$ shows the closest neighbour. The PDFs are obtained for $Re = 40$ and $\phi = 0.1$.

An example of such a combined PDF is shown for the streamwise force in figure 8(a) for $\phi = 0.1$. Comparing with figure 4, it is immediately obvious that the significant region around the reference particle has expanded due to the influence of more distant neighbours. Nevertheless, the general form of the maps that were obtained for $j = 1$ in figure 4 still prevails even when all the neighbours up to $j = M$ are included all at once. Another remarkable observation can be made from figure 8(b), where PDFs similar to those in figure 8(a) have been shown for the highest solid volume fraction of $\phi = 0.4$. The PDFs in figure 8(b) have gained fore-aft symmetry and are shrunk towards the reference particle, reminiscent of the force maps obtained with the data-driven model of Moore *et al.* (2019) at $\phi = 0.45$.




FIGURE 8. PDFs of positions of 15 closest neighbours for two different solid volume fractions. (a) $Re = 40$ and $\phi = 0.1$, (b) $Re = 40$ and $\phi = 0.4$.

4. Results and discussion

4.1. Practical aspects of implementation

4.1.1. Computational efficiency of KDE evaluations

Since the MPP model is to be used in EL simulations, at each time step it might be necessary to compute the forces and torques for several thousand, or perhaps millions of particles at once. The computational effort required for making predictions with the model is thus an important matter. The MPP model summarised in (3.18) consists of a handful of constant coefficients, multiplied by the same number of KDE-estimated PDFs. For m evaluations given n sample points, evaluation of a KDE-estimated PDF by naive kernel summation in (3.19) (which we use for model computations in this work) requires a quadratic $O(mn)$ operations, which may be computationally prohibitive for the practical implementation of the MPP model. However, this issue can be circumvented to a great extent by using efficient approaches that have been proposed over the past years, including data binning with fast Fourier transform, fast sum updating, fast Gauss transform and the dual-tree method (Langrené & Warin 2019). The computational cost can be reduced from quadratic $O(mn)$ operations to linear $O(m + n)$ or $O(m \log m + n \log n)$, resulting in a vast improvement of computational efficiency by orders of magnitude compared with naive kernel summation.

Another possibility for speeding up probability density estimations is to approximate the density distributions with well-known functional forms such as polynomials or exponential functions. As KDE is a non-parametric estimation method, the dataset based on which predictions are made need to be stored in order to obtain KDE-estimated probability densities. Using simpler functional forms obviates the need for storing the dataset, albeit at the expense of potential loss of accuracy. Nevertheless, one needs to study the model degradation due to such approximations versus the gained speed up in computations, and decide how much of a trade-off is acceptable in a particular situation.

4.1.2. Model hyper-parameters

In our model, there are a few parameters that need to be given values before making predictions, and finely tuned to deliver optimal performance. These parameters are to be distinguished from the set of unknown constants of regression in (3.22), which can be deterministically found from the minimisation of sums of errors. Two important hyper-parameters involved in our model are the kernel bandwidth indicated by \mathbf{H} in (3.20) and the number of included neighbours shown by M .

In its most general form the bandwidth is given by a symmetric, positive definite matrix when multidimensional data are concerned. The bandwidth \mathbf{H} significantly influences the obtained KDEs through controlling the orientation and smoothing intensity of a kernel, equivalent to the standard deviation of a Gaussian distribution. A very narrow bandwidth results in a KDE with high variance, whereas an over-smoothed estimation is generated if the bandwidth is too wide. The former tends to capture specificities of a particular dataset resulting in an over-fitted model performing poorly on other datasets, while the latter fails to acquire the crucial patterns in the data, giving an overall weak prediction capability. Since in the present work, the neighbour positions are uniformly distributed in all three dimensions, choosing the same bandwidth for all directions is sensible. In this case, the bandwidth matrix is determined by a scalar h such that $\mathbf{H} = h\mathbf{J}$ with \mathbf{J} being an all-ones matrix. A rule of thumb for estimating an optimal value for the bandwidth is given by Scott (1992) as $h = \sigma N^{-1/(D+4)}$, with N as the number of samples and D as the number of dimensions. In most of the cases in this work, we have used Scott's rule multiplied by a factor of 1–1.8. In most cases, a factor of >1 was necessary to avoid over-fitted results.

The other critical parameter to determine before attempting to make predictions is the number of neighbours to include in the modelling process. In low Reynolds number regimes, the force/torque fluctuations for each particle depend on a large number of surrounding particles (e.g. theoretically depending on each and every particle in the system for Stokes flow) due to the dominance of the elliptic nature of the governing equations at such regimes. In relatively inertial regimes, though, only a limited number of neighbouring particles influence the deviation of the hydrodynamic forces and torques acting on a particle from the average values. Akiki *et al.* (2017a) concluded that ≈ 15 – 40 and ≈ 10 closest neighbours were required for optimal modelling of drag and lift, respectively, and that inclusion of more surrounding particles did not improve the quality of their physics-driven PIEP model in terms of the coefficient of determination R^2 . For their data-driven PIEP model (Moore *et al.* 2019), ≈ 13 neighbours for drag and ≈ 15 neighbours for lift resulted in the best coefficient of determination. We have overall observed a similar trend for our MPP model concerning the number of neighbours to be included. A large number of neighbours have to be accounted for in lower Re for achieving maximum R^2 , while this number decreases with the Reynolds number. Furthermore, we typically require fewer neighbours to include in order to model the lateral forces and torques compared with the streamwise force. Full details of the variation of the model quality with the number of neighbours will be given in the following sections.

Finally, another choice to be made is the type of kernel in (3.20). Different types of kernels such as parabolic (Epanechnikov), cosine, exponential and linear kernels can be alternatively used to give the weighted-average value at each point. As pointed out by Silverman (1986), bandwidth selection is much more of a concern compared with the choice of the kernel, which can be legitimately based on the smoothness requirements or computational effort involved in making evaluations. In line with Silverman (1986), our experimentation with different types of kernels indicates that the choice of the kernel type

does not remarkably affect the performance of the model, given that optimal values of bandwidth are used.

4.1.3. Inclusion of average velocity

The channelling of the flow through the pores of a bed causes the fluid velocity experienced by the particles to fluctuate from particle to particle. This variance of undisturbed fluid velocity is partly dictated by the immediate neighbourhood of each particle, but not entirely. A bundle of a few particles might be collectively fully exposed to the fluid flow, while another group are mostly blocked by upstream neighbour groups. In the former case, the drag on a particle inside the bundle is significantly affected by its neighbours in the group, while in the latter case the particle would have experienced a relatively small or perhaps no drag, with little influence from its neighbours as the drag was not great to begin with. If we were able to include the effects of several neighbours around each particle instead of the approximation employed in (3.16), the channelling effect would theoretically be captured solely from the configuration of the neighbours. In the present work, this information can be directly supplied to the model by including a measure of the velocity seen by each particle. Since (3.18) is an approximation, linearly adding the average velocity seen by the particle may enhance the model. Therefore, (3.18) can be modified as follows:

$$\begin{aligned} \Delta F_{MPP}(\mathbf{R}, Re, \phi) = & \sum_{j=1}^M c_{\alpha,j} \tilde{p}_j(\mathbf{r}_j | \Delta F < 0) \\ & + \sum_{j=1}^M c_{\beta,j} \tilde{p}_j(\mathbf{r}_j | \Delta F > 0) + \sum_{d=1}^3 c_{\gamma,d} \bar{u}_d^v, \end{aligned} \quad (4.1)$$

where each $c_{\gamma,d}$ is an additional unknown coefficient, and \bar{u}_d^v denotes an average velocity component of $\bar{\mathbf{u}}^v$ in x , y or z direction. The estimation of fluid velocity seen by a finite-size particle (i.e. undisturbed fluid velocity) within a multiparticle system is not straight-forward, resulting in several definitions in the literature (Bagchi & Balachandar 2003; Kidanemariam *et al.* 2013). Kidanemariam *et al.* (2013) suggest the average of fluid velocity taken on a spherical shell centred at the particle location as the characteristic fluid velocity; a definition also utilised by Uhlmann & Doychev (2014) in their analysis of settling suspensions of spheres. Here, we perform the phase-averaging in a spherical volume about the location of each particle to obtain the fluid velocity seen by the particle

$$\bar{\mathbf{u}}^v = \sum_{l=1}^{N_l} \Phi_l \mathbf{u}_l \bigg/ \sum_{l=1}^{N_l} \Phi_l, \quad (4.2)$$

where Φ is the phase indicator function, which is utilised in order to avoid sampling velocity data inside the solid region, and N_l the number of grid points falling within the spherical volume. The diameter of the averaging volume is chosen to be $4d$ for $\phi = \{0.1, 0.2\}$ and $3d$ for $\phi = 0.4$. The inclusion of the average velocity $\bar{\mathbf{u}}^v$ improves the predictions in all cases, reflected in the increasing of R^2 by up to 5%. This improvement is more pronounced for $\phi = 0.4$, resulting in an increase of up to 15% in the coefficient of determination.

4.2. Model assessment

As a measure of performance, we will use the coefficient of determination defined as

$$R^2 = 1 - \frac{\sum_{i=1}^N (\Delta F_{MPP,i} - \Delta F_{DNS,i})^2}{\sum_{i=1}^N (\Delta F_{DNS,i} - \langle \Delta F_{DNS,i} \rangle)^2}, \quad (4.3)$$

where N is the number of samples used for the computation of R^2 . Also, we note that $\langle \Delta F_{DNS,i} \rangle = 0$ by definition. The coefficient of determination can be interpreted as the fraction of variations that are explained by the model. As such, $R^2 = 0$ implies that the particle-to-particle force/torque fluctuations are not captured at all and the model does not perform any better than giving a single average value for all particles. On the other hand, $R^2 = 1$ would indicate a perfect fit; i.e. fluctuations predicted by the model exactly match those obtained from PR-DNS.

With every data-driven model, it is necessary to ensure that the model has not only fitted the data in a satisfactory manner, but is also able to generalise its predictions to unseen data. Complex data-driven models in the realm of ML such as neural networks often have numerous unknown parameters which makes them prone to overfitting, particularly when dealing with small datasets. Put another way, an over-fitted model simply memorises the dataset thanks to its huge set of parameters, while it performs poorly when presented with unseen data. One strategy to prevent over-fitting and to assess the generalisation capability of a model is to split the dataset into a training set and a smaller validation set, so as to fit the model using the former and test its performance on the latter. This technique, commonly known as the hold-out method, works well if used with sufficiently large datasets. For datasets limited in size, such as data obtained from PR-DNS here, not only the validation set might not be sufficiently representative, the dataset is only partially utilised to train the model (Witten *et al.* 2016). The K -fold cross-validation is an assessment technique that randomly partitions the dataset into K subsets or ‘folds’, so that one of these folds is held out for testing while the model is trained on the rest of $K - 1$ folds. This process is repeated with every partitioning such that the entire dataset has been used both for training and for validation. Any measure of error or accuracy and also the model’s parameters may be averaged to give a more reliable estimation of the performance. To ensure the prediction reliability of the MPP model, we employ the K -fold cross validation with $K = 8$ where multiple rounds of cross-validation are performed using K subsets of the data. All reported results are thus the averaged values over multiple rounds so that a reliable estimate of the model’s predictive performance is achieved.

4.3. Performance of the MPP model

The results of the MPP model’s predictions and comparison with our PR-DNS data are presented in terms of the coefficient of determination R^2 in table 3. Note that we have deliberately avoided inclusion of data for ΔF_z and ΔT_y , as the results obtained for these variables are equivalent to those for ΔF_y and ΔT_z , respectively. The green and light green cells in table 3 refer to cases for which $R^2 \geq 0.7$ and $R^2 \geq 0.6$, respectively, and cells highlighted in red represents cases for which $R^2 \leq 0.5$. Overall, the MPP model exhibits a remarkable ability to correlate the observed force/torque deviations from the average values to the specific neighbourhood of each particle. Our proposed model is able to

ϕ	Re	ΔF_x	ΔF_y	ΔT_z
0.1	2	0.84	0.78	0.85
0.1	10	0.80	0.76	0.80
0.1	40	0.70	0.69	0.64
0.1	150	0.57	0.53	0.38
0.2	0.2	0.71	0.71	0.85
0.2	2	0.76	0.74	0.86
0.2	40	0.71	0.66	0.70
0.2	150	0.62	0.55	0.48
0.4	2	0.54	0.53	0.69
0.4	40	0.67	0.58	0.64
0.4	150	0.61	0.47	0.52

TABLE 3. Performance of the MPP model represented by the coefficient of determination R^2 for cases considered in this work.

explain up to 84 %, 78 % and 85 % of the drag, lift and torque variations, respectively, in the best cases with a mean coefficient of determination of $R^2 = 0.68$, 0.63 and 0.67 averaged over all cases. Except for three cases (shown in light red in table 3), at least 50 % of the particle-to-particle force/torque variations are captured by the model, while in most cases the coefficient of determination is seen to be greater than 60 %. In figure 9 correlation scatter plots are shown for the stream-wise force or drag, the lateral force or lift and the lateral component of the torque. For each hydrodynamic load, we have selected three cases representing a range of Reynolds numbers and solid volume fractions in increasing order, in order to demonstrate cases where the model exhibits varying degrees of performance. The horizontal coordinate of each point on the plots shows the value of the force/torque deviation obtained from PR-DNS, whereas the vertical coordinate shows the value predicted by the present MPP model for the same sample in the dataset. The red bisector in figure 9 indicates a perfect fit, for which $\Delta F_{MPP} = \Delta F_{DNS}$ and $R^2 = 1$. Notably, the classical microstructure-ignorant drag correlations of the form $F = \langle F \rangle (Re, \phi)$ in § 2.4 would all lie on a horizontal line given by $\Delta F = 0$ in figure 9, since these correlations only give an average value for an entire ensemble of particles and are thus unable to explain any particle-to-particle variation of the drag. Clearly, the MPP model is not perfect in any of the cases, but the improvements over conventional correlations are substantial, reaching up to $R^2 = 0.84$, 0.77 and 0.85 for ΔF_x , ΔF_y and ΔT_z in best cases where $(Re, \phi) = (2, 0.1)$. As both Re and ϕ are increased from left to right in figure 9, the performance of the model suffers to some extent. Remarkably, however, the drag variation is still captured up to 61 % at the extreme case of $(Re, \phi) = (150, 0.4)$, while R^2 remains in the range of 0.66–0.71 for all variables at $(Re, \phi) = (40, 0.2)$.

The approximation in (3.16) that $\tilde{p}(\mathbf{R} | \cdot)$ (i.e. the PDF depending on locations of all included neighbours) may be estimated as a linear combination of $\tilde{p}_j(\mathbf{r}_j | \cdot)$ functions (i.e. marginal PDFs each depending on the location of only a single neighbour) imposes an important limit on how accurate the model can be. For instance, in case the joint distribution $\tilde{p}_{1,2}(\mathbf{r}_1, \mathbf{r}_2 | \cdot)$ is practical to obtain, the functional form of \tilde{p} for the first neighbour would be different depending on where the second neighbour is located. In such a situation, the PDF for the drag in figure 4 would no longer be axisymmetric about the x axis; it would take on different forms depending on where the second closest neighbour

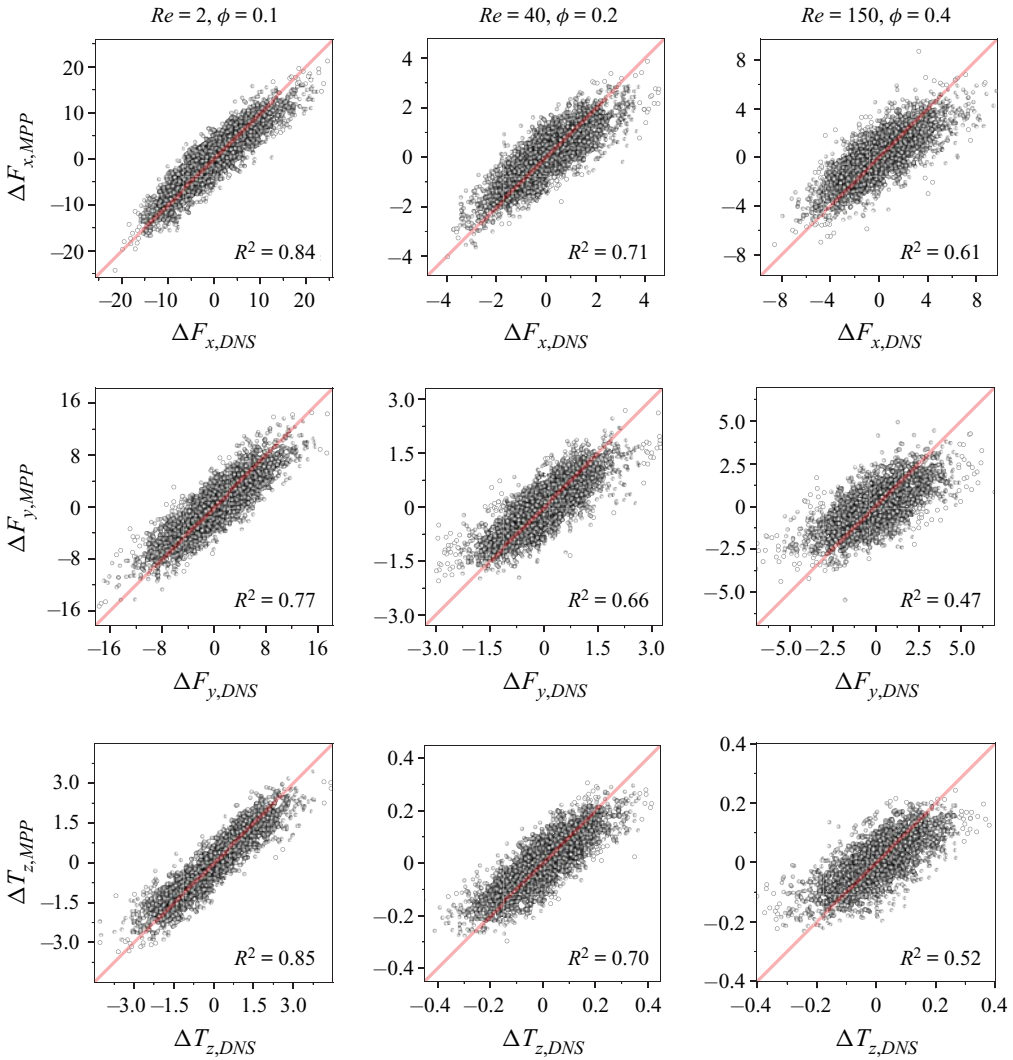


FIGURE 9. Regression plots for the drag, lift and lateral torque for various Reynolds numbers and solid volume fractions. The oblique red line in each plot shows an ideal fit.

is located. An example of such a case is shown on the left-hand side of the PDF in figure 10. On the left plot, we have fixed the location of the second neighbour (shown as the grey circle) to stay within a small radius of $r_2 = A = (-1, 1, z)$ on x - y plane, while the z coordinate is left free. On the right, we have reiterated the right-hand side plot of figure 4, where the PDF is obtained by conditioning only the drag and observing the locations of the first neighbour. Obviously, the presence of the second neighbour has affected the probability distribution of the first neighbour's position which is indicated by pushing the upstream probability density peak downwards and shrinking it simultaneously. This alteration of the PDF is partially, but not merely, due to an exclusion of the space occupied by second neighbour, but also owing to the modification of the flow field induced by the particular positioning of the second neighbour. Notably, the first neighbour can still happen to be located immediately upstream of the reference particle. In short, the second




FIGURE 10. The PDF of the first closest neighbour position (a) having a second neighbour (drawn as a dashed white circle) deliberately fixed on x - y plane at $\mathbf{r}_2 = \mathbf{A} = (-1, 1, z)$, (b) PDF of the first closest neighbour position, same as in figure 4(c). The PDFs are obtained for $Re = 40$, $\phi = 0.1$.

neighbour influences how the first neighbour affects the reference particle. This type of interaction dubbed as ‘ternary effects’, and also higher-order interactions of this type are ingredients of the flow physics sacrificed in (3.16) in exchange for a functional form that is practical to estimate.

In Stokes flow and dilute conditions, binary interactions prevail and high-order interactions are negligible, and the net effect arising from different neighbours can be given as a linear superposition of the influence of each individual neighbour. As the Reynolds number increases, linear addition of effects can no longer account for ternary and higher-order interactions, thus diminishing the validity of the linear combination of effects employed in (3.16). Consider the relatively low solid volume fraction cases with $\phi = 0.1$ in table 3. As the Reynolds number increases, the performance of the MPP model deteriorates as expected as a result of the superposition of each neighbour’s influence in (3.16). The coefficient of determination R^2 , however, remains almost constant for $\phi = 0.2$ and counter-intuitively improves for $\phi = 0.4$ when the Reynolds number increases. When $\phi = 0.4$, the average interparticle distance is considerably smaller than for $\phi = 0.1$ making the PDFs more flattened and uniform, thus reducing the variability of the extracted probability map. The increase of the Reynolds number renders the distribution more anisotropic to some extent (also reflected in the higher standard deviation of higher Re cases for $\phi = 0.4$ in table 2) which leads to an improved R^2 .

In order to be able to practically implement the MPP model, we separated the influence of each neighbour on the force/torque fluctuations of a reference particle by marginalisation of the probability distribution in (3.16), and expressed the net deviation as the linear combination of all effects. Although this assumption resembles the pairwise interaction assumption of Akiki *et al.* (2017a) in their physics-driven PIEP model, it is nevertheless inherently different. The pairwise interaction approximation employed by Akiki *et al.* (2017a) is based on binary interaction maps where there is no incorporated notion of solid volume fraction. In (3.16), however, each functional $\tilde{p}_j(\mathbf{r}_j | \cdot)$ is obtained for neighbour j in the presence of all other neighbouring spheres. In other words, $\tilde{p}_j(\mathbf{r}_j | \cdot)$ appears as a binary interaction in which the effects of all other neighbours are averaged, but still statistically present. This implies that the functional form of each \tilde{p}_j is in fact dependent on the solid volume fraction, which may explain the superior performance of the present MPP model (and the data-driven PIEP model, for that matter) in a higher particle concentrations compared with physics-driven PIEP model. For the purpose of

comparison, with the MPP model $R^2 = 0.67$ and 0.61 for drag in $(Re, \phi) = (40, 0.4)$ and $(Re, \phi) = (150, 0.4)$, respectively; whereas $R^2 = 0.12$ and 0.24 in $(Re, \phi) = (21, 0.45)$ and $(Re, \phi) = (115, 0.45)$, respectively, with the physics-driven PIEP model. Although Re and ϕ are not exactly the same for the cases here and for those considered by Moore *et al.* (2019), a general comparison of results shows that the performance of the MPP model is on par with the hybrid PIEP model which combines the physics-driven and data-driven approaches.

Lastly, it must be pointed out that the present MPP model only considers static arrays of spheres. In a general particle-laden flow situation, however, particles will be in motion and the dynamic evolution of the system will depend on the positions as well as the motion of the particles. With the current formulation, inclusion of the translational and rotational velocity and acceleration of each neighbour in the PDFs of (3.18) will increase their dimensionality, thus rendering them impractical to use in their present form. Future investigations (similar to what has been done by Esteghamatian *et al.* (2017, 2018) for conventional EL techniques) are needed to evaluate the MPP model's performance when implemented in EL simulations through comparisons with corresponding PR-DNS cases. Such an assessment would be greatly beneficial in quantifying the significance of the inclusion of translational and rotational velocity and acceleration of the neighbours, and also to identify situations where such contributions might be more or less important.

5. Summary and conclusion

Numerical simulations are indispensable tools in analysing particle-laden flows, as experimental investigation of this type of flow is both very costly if not entirely impractical, while the extent to which details of these complex flows are available to experimental measurements is also quite limited. Particle transport in microfluidic separation devices or in highly dense particulate flows and heat transfer properties of combustion fluidised beds are a few examples for which resorting to numerical tools is inevitable in order to obtain the details of the physical processes. Most particle-laden flows of practical interest host billions of particles and span a space orders of magnitude larger than the size of individual particles. Even with the exponential growth of computing power, resolving all relevant scales is not within reach in the foreseeable future. In a higher intermediate (i.e. meso) scale, the fluid governing equations are averaged and solved in sub-volumes each containing a few particles, treating the fluid in an Eulerian way while still tracking each individual particle in a Lagrangian manner. Consequently, the computational cost is substantially lowered at the expense of the need to supply closure models for fluid–solid momentum exchange. As field variables are not available at the particle level, hydrodynamic interaction forces and torques cannot be directly computed and should hence be appropriately modelled. In relatively dense particle-laden systems, the flow varies on the scale of the particle dimension due to the pseudo-turbulence created by the neighbours. Conventional point-particle models that are parameterised only in terms of the Reynolds number and average solid volume fraction fail to account for the effects of the complex undisturbed flow on the drag. Microstructure-induced lateral forces and torques, on the other hand, have been neglected entirely in such models by definition. Given the fact that the physical fidelity of EL simulations directly relies on the accuracy of the interphase coupling scheme, developing force/torque models capable of incorporating the influence of local neighbourhood of particles is crucial.

In the present work, we have attempted to develop a deterministic model based on probabilistic arguments for hydrodynamic forces and torques exerted on each individual particle within a random array of fixed spheres. Owing to the unique neighbourhood

of each particle, the flow in its vicinity is modified in a distinct way, giving rise to significant force and torque variations. The principal idea exploited by our MPP model is that conditioning force or torque deviations to positive or negative ranges results in the emergence of particularly interesting, non-uniform distributions of neighbour locations. That is to say, if particles (i.e. samples) that experience particular ranges of force/torque fluctuations are filtered, and in turn the positions of their neighbours are examined, we will find that their surrounding particles are always spatially distributed in remarkably non-uniform manners consistent with our physical understanding of binary interactions between two spheres. For instance, a neighbour located immediately upstream shields a particle from the oncoming flow and eliminates the frontal high pressure region in a binary system. Analysis of the neighbours distribution when the drag force is lower than average clearly demonstrates that this is still true even in dense arrays of sphere, which signifies that invaluable information can be extracted from PR-DNS for the dependence of forces/torques on the local neighbourhood. If a particle experience a low drag force, it is highly likely that the particle is shielded by an upstream neighbour. The same argument holds for other force/torque-conditioned cases as well. It is critically important that while in a binary system these observations are certain events, for an array of particles the occurrence of such events is probabilistic in nature. Upon fixing the location of one neighbour (i.e. $M = 1$), the configuration of all other neighbours is still free to randomly change. Therefore, for a given location of a neighbour, the force/torque fluctuation on the reference particle would be represented by a distribution, not a fixed value unlike in a binary system. This is why the notion of the expected value is invoked in order to provide an average value of the experienced force/torques for a given location of one neighbour. Two conceptual extremes are worth reiterating.

- (a) All neighbours are free ($M = 0$). In this case, the expected value for force/torque deviation is zero, as the distribution is centred on $\Delta F = 0$.
- (b) All neighbours are fixed ($M = N_p - 1$). The distribution of force/torque deviation becomes extremely narrow (Dirac's delta function) as all sources of fluctuation are held fixed. The expected value for force/torque deviation approaches $\Delta F = \Delta F_{DNS}$.

We note that in a binary system, $M = N_p - 1 = 1$, therefore $\Delta F = \Delta F_{DNS}$. Theoretically then, this framework can have varying degrees of accuracy approaching that of PR-DNS; practically, however, obtaining probability distribution maps for $M > 1$ requires exponentially more samples to sufficiently cover the input space of the problem. We recognise the case $M = 1$ as representing first-order effects, while higher-order effects may be captured by fixing the positions of two or more neighbours simultaneously (i.e. $M \geq 2$, which translates into, for instance, obtaining $\tilde{p}(\mathbf{r}_1, \mathbf{r}_2 | \cdot)$ instead of $\tilde{p}(\mathbf{r}_1 | \cdot)$ only). In the present work, we have considered $M = 1$ as a first step of improvement over microstructure-ignorant classical correlations. In order to estimate forces and torques, we establish a framework to obtain the expected value of the force/torque deviations from the mean by taking advantage of the force/torque-conditioned probability distribution functions of neighbour locations. These distributions are approximated using KDEs which serve as basis functions for regression. The effect of each neighbour on the deviations is then linearly combined and the value of unknown coefficients is found through an ordinary least-squares regression method. With $M = 1$, adding more neighbours results in the linear addition of more terms, meaning that the functional form of the PDFs do not change. For $M = 2$, the PDFs themselves would change depending on where the second neighbour is located in addition to the first neighbour, as alluded to in [figure 10](#).

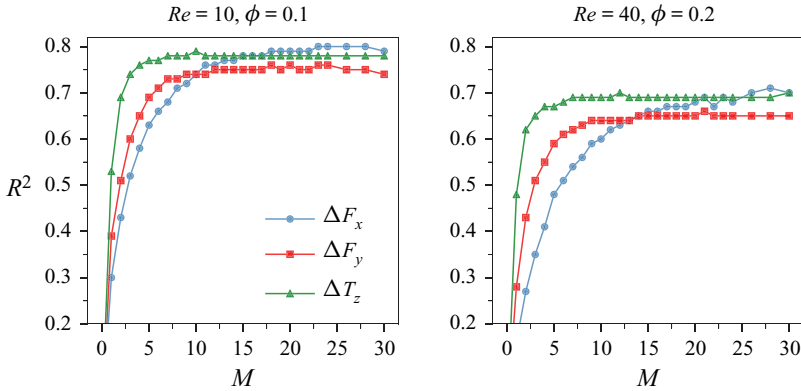


FIGURE 11. Coefficient of determination R^2 as a function of the number of included neighbours M used for construction of the model.

We have generated a dataset consisting of several cases at various Reynolds numbers and solid volume fractions relevant to dense particle-laden systems of interest. Following the aforementioned discussion, the model is not expected to be perfect when $M = 1$. Nevertheless, the MPP model demonstrates remarkable performance by explaining up to 60%–70% of particle-to-particle force/torque variation in most cases, while for a few cases in the low Re and ϕ the percentage of explained variance in the data rises to up to 85%. We experimented with the model to examine the dependence of the model performance on the number of neighbours included in the superposition of effects. (Note that in the present work, the probability distribution functions are all generated by considering only one neighbour, i.e. $M = 1$. The number of included neighbours in figure 11 only determines how many of these PDFs are superposed to make a final prediction.) As shown in figure 11, we found out that inclusion of ≈ 20 neighbours is adequate for gaining optimal performance for the streamwise force. The jump in R^2 is quite steep for lateral forces and torques, as the inclusion of ≈ 5 – 10 neighbours are sufficient to achieve the maximum performance for most cases. Consequently, the MPP model requires two probability distribution functions for negative and positive contributions, and a few constant coefficients (≈ 10 – 40) to predict a significant portion of variations in force/torque components for given Re and ϕ . Computing the probability densities requires evaluation of the KDEs in (3.19), which in turn needs the dataset samples based on which the model is constructed. The samples needed for KDE estimation can be easily stored and utilised along with the constant coefficients to make predictions in an EL simulation.

According to the results obtained by our MPP model and also by the PIEP model (Akiki *et al.* 2017a,b; Moore *et al.* 2019), the performance experiences a significant enhancement by only considering the first-order effects, hinting an asymptotic behaviour with inclusion of higher effects. Even though $M = N_p - 1$ is theoretically needed (in construction of PDFs) to reach the PR-DNS accuracy, we anticipate that even considering the second-order effects properly renders the performance levels satisfactorily high to obviate the need for making the model more complex, as the gain would probably not be significant for $M \geq 3$. On the one hand, accounting for the first-order effects of local microstructure on forces and torques can effectively predict the occurrence of phenomena such as wake attraction or DKT, which is deemed the dominant mechanism in preferential concentration and clustering of suspensions (Yin & Koch 2007, 2008; Uhlmann & Doychev 2014; Zaidi *et al.* 2014). On the other hand, we have previously shown that transverse particle

velocity fluctuations and granular temperature are considerably underestimated by using conventional drag correlations in EL simulations (Esteghamatian *et al.* 2017, 2018) due to the unavailability of microstructure-induced lateral forces. The MPP model can play a promising role in alleviating both challenges in current meso-scale simulation tools.

Acknowledgements

We greatly appreciate the financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC) via their Discovery Grant Program. This research was enabled by support provided by Compute Canada (<http://www.computecanada.ca>) through Anthony Wachs's 2019 Resources for Research Groups allocation qpf-764-ab.

Declaration of interests

The authors report no conflict of interest.

REFERENCES

- AIDUN, C. K. & CLAUSEN, J. R. 2010 Lattice-Boltzmann method for complex flows. *Annu. Rev. Fluid Mech.* **42** (1), 439–472.
- AKIKI, G., JACKSON, T. L. & BALACHANDAR, S. 2016 Force variation within arrays of monodisperse spherical particles. *Phys. Rev. Fluids* **1** (4), 044202.
- AKIKI, G., JACKSON, T. L. & BALACHANDAR, S. 2017a Pairwise interaction extended point-particle model for a random array of monodisperse spheres. *J. Fluid Mech.* **813**, 882–928.
- AKIKI, G., MOORE, W. C. & BALACHANDAR, S. 2017b Pairwise-interaction extended point-particle model for particle-laden flows. *J. Comput. Phys.* **351**, 329–357.
- BAGCHI, P. & BALACHANDAR, S. 2002 Effect of free rotation on the motion of a solid sphere in linear shear flow at moderate *Re*. *Phys. Fluids* **14** (8), 2719–2737.
- BAGCHI, P. & BALACHANDAR, S. 2003 Effect of turbulence on the drag and lift of a particle. *Phys. Fluids* **15** (11), 3496–3513.
- BALACHANDAR, S. 2009 A scaling analysis for point-particle approaches to turbulent multiphase flows. *Intl J. Multiphase Flow* **35** (9), 801–810.
- BATCHELOR, G. K. 1972 Sedimentation in a dilute dispersion of spheres. *J. Fluid Mech.* **52** (02), 245.
- BECK, A., FLAD, D. & MUNZ, C.-D. 2019 Deep neural networks for data-driven LES closure models. *J. Comput. Phys.* **398**, 108910.
- BEETSTRA, R., VAN DER HOEF, M. A. & KUIPERS, J. A. M. 2007 Drag force of intermediate Reynolds number flow past mono- and bidisperse arrays of spheres. *AIChE J.* **53** (2), 489–501.
- BOGNER, S., MOHANTY, S. & RÜDE, U. 2015 Drag correlation for dilute and moderately dense fluid-particle systems using the lattice Boltzmann method. *Intl J. Multiphase Flow* **68**, 71–79.
- CAPECELATRO, J. & DESJARDINS, O. 2013 An Euler-Lagrange strategy for simulating particle-laden flows. *J. Comput. Phys.* **238**, 1–31.
- CHACÓN, J. E. & DUONG, T. 2018 *Multivariate Kernel Smoothing and Its Applications*. CRC.
- COLLINS, L. R. & KESWANI, A. 2004 Reynolds number scaling of particle clustering in turbulent aerosols. *New J. Phys.* **6**, 119–119.
- COVENEY, P. V., DOUGHERTY, E. R. & HIGHFIELD, R. R. 2016 Big data need big theory too. *Phil. Trans. R. Soc. A* **374** (2080), 20160153.
- DASH, S. M. & LEE, T. S. 2015 Two spheres sedimentation dynamics in a viscous liquid column. *Comput. Fluids* **123**, 218–234.
- DURAISAMY, K., IACCARINO, G. & XIAO, H. 2019 Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51** (1), 357–377.
- ERGUN, S. 1952 Fluid flow through packed columns. *Chem. Engng Prog.* **48**, 89–94.
- ESTEGHAMATIAN, A., BERNARD, M., LANCE, M., HAMMOUTI, A. & WACHS, A. 2017 Micro/meso simulation of a fluidized bed in a homogeneous bubbling regime. *Intl J. Multiphase Flow* **92**, 93–111.

- ESTEGHAMATIAN, A., EUZENAT, F., HAMMOUTI, A., LANCE, M. & WACHS, A. 2018 A stochastic formulation for the drag force based on multiscale numerical simulation of fluidized beds. *Intl J. Multiphase Flow* **99**, 363–382.
- FAXEN, H. 1923 Die bewegung einer starren kugel langs der achse eines mit zaher flussigkeit gefullten rohres. *Ark. Mat. Astr. Fys.* **17**, 1–28.
- FENG, J., HU, H. H. & JOSEPH, D. D. 1994a Direct simulation of initial value problems for the motion of solid bodies in a Newtonian fluid. Part 2. Couette and Poiseuille flows. *J. Fluid Mech.* **277**, 271–301.
- FENG, J., HU, H. H. & JOSEPH, D. D. 1994b Direct simulation of initial value problems for the motion of solid bodies in a newtonian fluid. Part 1. Sedimentation. *J. Fluid Mech.* **261**, 95–134.
- FORNARI, W., ARDEKANI, M. N. & BRANDT, L. 2018 Clustering and increased settling speed of oblate particles at finite Reynolds number. *J. Fluid Mech.* **848**, 696–721.
- FORTES, A. F., JOSEPH, D. D. & LUNDGREN, T. S. 1987 Nonlinear mechanics of fluidization of beds of spherical particles. *J. Fluid Mech.* **177**, 467–483.
- GATIGNOL, R. 1983 The faxén formulas for a rigid particle in an unsteady non-uniform stokes-flow. *J. de Méc. Théor. et Appl.* **2** (2), 143–160.
- GIDASPOW, D. 1994 *Multiphase Flow and Fluidization: Continuum and Kinetic Theory Descriptions*. Academic.
- GLOWINSKI, R., PAN, T. W., HESLA, T. I. & JOSEPH, D. D. 1999 A distributed Lagrange multiplier/fictitious domain method for particulate flows. *Intl J. Multiphase Flow* **25** (5), 755–794.
- GÖTZ, J., IGLBERGER, K., STÜRMER, M. & RÜDE, U. 2010 Direct numerical simulation of particulate flows on 294 912 processor cores. In *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–11. IEEE Computer Society.
- GUO, X., LI, W. & IORIO, F. 2016 Convolutional neural networks for steady flow approximation. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 481–490. ACM.
- HASIMOTO, H. 1959 On the periodic fundamental solutions of the stokes equations and their application to viscous flow past a cubic array of spheres. *J. Fluid Mech.* **5** (2), 317–328.
- HASTIE, T., TIBSHIRANI, R. & FRIEDMAN, J. 2009 *The Elements of Statistical Learning*. Springer.
- HILL, R. J., KOCH, D. L. & LADD, A. J. C. 2001a The first effects of fluid inertia on flows in ordered and random arrays of spheres. *J. Fluid Mech.* **448**, 213–241.
- HILL, R. J., KOCH, D. L. & LADD, A. J. C. 2001b Moderate-Reynolds-number flows in ordered and random arrays of spheres. *J. Fluid Mech.* **448**, 243–278.
- HINTON, G., DENG, L., YU, D., DAHL, G., MOHAMED, A., JAITLY, N., SENIOR, A., VANHOUCKE, V., NGUYEN, P., SAINATH, T. & KINGSBURY, B. 2012 Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Signal Process. Mag.* **29** (6), 82–97.
- VAN DER HOEF, M. A., BEETSTRA, R. & KUIPERS, J. A. M. 2005 Lattice-Boltzmann simulations of low-Reynolds-number flow past mono- and bidisperse arrays of spheres: results for the permeability and drag force. *J. Fluid Mech.* **528**, 233–254.
- VAN DER HOEF, M. A., VAN SINT ANNALAND, M., DEEN, N. G. & KUIPERS, J. A. M. 2008 Numerical simulation of dense gas-solid fluidized beds: a multiscale modeling strategy. *Annu. Rev. Fluid Mech.* **40**, 47–70.
- VAN DER HOEF, M. A., VAN SINT ANNALAND, M., KUIPERS, J. A. M. 2004 Computational fluid dynamics for dense gas–solid fluidized beds: a multi-scale modeling strategy. *Chem. Engng Sci.* **59** (22–23), 5157–5165.
- HORNE, W. J. & MAHESH, K. 2019 A massively-parallel, unstructured overset method to simulate moving bodies in turbulent flows. *J. Comput. Phys.* **397**, 108790.
- HORNIK, K., STINCHCOMBE, M. & WHITE, H. 1989 Multilayer feedforward networks are universal approximators. *Neural Networks* **2** (5), 359–366.
- HU, H. H., PATANKAR, N. A. & ZHU, M. Y. 2001 Direct numerical simulations of fluid–solid systems using the arbitrary Lagrangian–Eulerian technique. *J. Comput. Phys.* **169** (2), 427–462.
- JIANG, Y., KOLEHMAINEN, J., GU, Y., KEVREKIDIS, Y. G., OZEL, A. & SUNDARESAN, S. 2019 Neural-network-based filtered drag model for gas-particle flows. *Powder Technol.* **346**, 403–413.

- KAJISHIMA, T. & TAKIGUCHI, S. 2002 Interaction between particle clusters and particle-induced turbulence. *Intl J. Heat Fluid Flow* **23** (5), 639–646.
- KIDANEMARIAM, A. G., CHAN-BRAUN, C., DOYCHEV, T. & UHLMANN, M. 2013 Direct numerical simulation of horizontal open channel flow with finite-size, heavy particles at low solid volume fraction. *New J. Phys.* **15** (2), 025031.
- KRIEBITZSCH, S. H. L., VAN DER HOEF, M. A. & KUIPERS, J. A. M. 2013 Fully resolved simulation of a gas-fluidized bed: a critical test of DEM models. *Chem. Engng Sci.* **91**, 1–4.
- KRIZHEVSKY, A., SUTSKEVER, I. & HINTON, G. E. 2012 Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25* (ed. F. Pereira, C. J. C. Burges, L. Bottou & K. Q. Weinberger), pp. 1097–1105. Curran Associates.
- KUTZ, J. N. 2017 Deep learning in fluid dynamics. *J. Fluid Mech.* **814**, 1–4.
- LANGRENÉ, N. & WARIN, X. 2019 Fast and stable multivariate kernel density estimation by fast sum updating. *J. Comput. Graph. Stat.* **28** (3), 596–608.
- LECUN, Y., BENGIO, Y. & HINTON, G. 2015 Deep learning. *Nature* **521** (7553), 436–444.
- LING, J., KURZAWSKI, A. & TEMPLETON, J. 2016 Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166.
- LING, J. & TEMPLETON, J. 2015 Evaluation of machine learning algorithms for prediction of regions of high Reynolds averaged Navier–Stokes uncertainty. *Phys. Fluids* **27** (8), 085103.
- LOUGE, M., LISCHER, D. J. & CHANG, H. 1990 Measurements of voidage near the wall of a circulating fluidized bed riser. *Powder Technol.* **62** (3), 269–276.
- MA, M., LU, J. & TRYGGVASON, G. 2015 Using statistical learning to close two-fluid multiphase flow equations for a simple bubbly system. *Phys. Fluids* **27** (9), 092101.
- MAULIK, R., SAN, O., RASHEED, A. & VEDULA, P. 2018 Data-driven deconvolution for large eddy simulations of Kraichnan turbulence. *Phys. Fluids* **30** (12), 125109.
- MAXEY, M. 2017 Simulation methods for particulate flows and concentrated suspensions. *Annu. Rev. Fluid Mech.* **49** (1), 171–193.
- MAXEY, M. R. 1983 Equation of motion for a small rigid sphere in a nonuniform flow. *Phys. Fluids* **26** (4), 883.
- MITTAL, R. & IACCARINO, G. 2005 Immersed boundary methods. *Annu. Rev. Fluid Mech.* **37** (1), 239–261.
- MOORE, W. C., BALACHANDAR, S. & AKIKI, G. 2019 A hybrid point-particle force model that combines physical and data-driven approaches. *J. Comput. Phys.* **385**, 187–208.
- PAPOULIS, A., PILLAI, S. U. & PILLAI, S. U. 2002 *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill.
- PERCUS, J. K. & YEVICK, G. J. 1958 Analysis of classical statistical mechanics by means of collective coordinates. *Phys. Rev.* **110** (1), 1–13.
- PESKIN, C. S. 1977 Numerical analysis of blood flow in the heart. *J. Comput. Phys.* **25** (3), 220–252.
- PRAHL, L., HÖLZER, A., ARLOV, D., REVSTEDT, J., SOMMERFELD, M. & FUCHS, L. 2007 On the interaction between two fixed spherical particles. *Intl J. Multiphase Flow* **33** (7), 707–725.
- RAISSI, M. & KARNIADAKIS, G. E. 2018 Hidden physics models: machine learning of nonlinear partial differential equations. *J. Comput. Phys.* **357**, 125–141.
- RAISSI, M., PERDIKARIS, P. & KARNIADAKIS, G. E. 2019 Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707.
- RAISSI, M., WANG, Z., TRIANTAFYLLOU, M. S. & KARNIADAKIS, G. E. 2018 Deep learning of vortex-induced vibrations. *J. Fluid Mech.* **861**, 119–137.
- RETTINGER, C., GODENSCHWAGER, C., EIBL, S., PRECLI, T., SCHRUFF, T., FRINGS, R. & RÜDE, U. 2017 Fully resolved simulations of dune formation in riverbeds. In *International Supercomputing Conference*, pp. 3–21. Springer.
- RICHARDSON, J. F. & ZAKI, W. N. 1954 Sedimentation and fluidisation. Part I. *Trans. Inst. Chem. Engrs* **32**, 35–53.
- ROSS, S. M. 2010 *A First Course in Probability*. Pearson Prentice Hall.
- SAFFMAN, P. G. 1965 The lift on a small sphere in a slow shear flow. *J. Fluid Mech.* **22** (2), 385–400.

- SANGANI, A. S. & ACRIVOS, A. 1982 Slow flow through a periodic array of spheres. *Intl J. Multiphase Flow* **8** (4), 343–360.
- SCHILLER, L. & NAUMANN, A. 1933 Über die grundlegenden berechnungen bei der schwerkraftaufbereitung. *Z. Vereines Deutscher Inge.* **77**, 318–321.
- SCOTT, D. W. 1992 *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley.
- SEKAR, V. & KHOO, B. C. 2019 Fast flow field prediction over airfoils using deep learning approach. *Phys. Fluids* **31** (5), 057103.
- SILVERMAN, B. W. 1986 *Density Estimation for Statistics and Data Analysis*. Taylor & Francis.
- SQUIRES, K. 2007 *Point-Particle Methods for Disperse Flows*, pp. 282–319. Cambridge University Press.
- SUBRAMANIAM, S. 2013 Lagrangian–Eulerian methods for multiphase flows. *Prog. Energy Combust. Sci.* **39** (2-3), 215–245.
- SUBRAMANIAM, S. & BALACHANDAR, S. 2018 *Towards Combined Deterministic and Statistical Approaches to Modeling Dispersed Multiphase Flows*, pp. 7–42. Springer.
- SUCCI, S. & COVENEY, P. V. 2019 Big data: the end of the scientific method? *Phil. Trans. R. Soc. A* **377** (2142), 20180145.
- SUN, L., GAO, H., PAN, S. & WANG, J.-X. 2019 Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Comput. Meth. Appl. Mech. Engng* **361**, 112732.
- TANG, Y., LAU, Y. M., DEEN, N. G., PETERS, E. A. J. F. & KUIPERS, J. A. M. 2016 Direct numerical simulations and experiments of a pseudo-2D gas-fluidized bed. *Chem. Engng Sci.* **143**, 166–180.
- TANG, Y., PETERS, E. A., KUIPERS, J. A., KRIEBITZSCH, S. H., VAN DER HOEF, M. A. 2015 A new drag correlation from a fully resolved simulations of flow past monodisperse static arrays of spheres. *AIChE J.* **61** (2), 688–698.
- TENNETI, S., GARG, R. & SUBRAMANIAM, S. 2011 Drag law for monodisperse gas–solid systems using particle-resolved direct numerical simulation of flow past fixed assemblies of spheres. *Intl J. Multiphase Flow* **37** (9), 1072–1092.
- UHLMANN, M. & DOYCHEV, T. 2014 Sedimentation of a dilute suspension of rigid spheres at intermediate Galileo numbers: the effect of clustering upon the particle motion. *J. Fluid Mech.* **752**, 310–348.
- WACHS, A. 2010 PeliGRIFF, a parallel DEM-DLM/FD direct numerical simulation tool for 3D particulate flows. *J. Engng Maths* **71** (1), 131–155.
- WACHS, A., HAMMOUTI, A., VINAY, G. & RAHMANI, M. 2015 Accuracy of finite volume/staggered grid distributed lagrange multiplier/fictitious domain simulations of particulate flows. *Comput. Fluids* **115**, 154–172.
- WEN, C. Y. & YU, Y. H. 1966 Mechanics of fluidization. *Chem. Engng Prog. Symp. Series* **62**, 100–111.
- WERTHEIM, M. S. 1963 Exact solution of the Percus-Yevick integral equation for hard spheres. *Phys. Rev. Lett.* **10** (8), 321–323.
- WILLEN, D. P. & PROSPERETTI, A. 2019 Resolved simulations of sedimenting suspensions of spheres. *Phys. Rev. Fluids* **4** (1), 014304.
- WILLEN, D. P. & SIERAKOWSKI, A. J. 2019 Resolved particle simulations using the physalis method on many GPUs. *Comput. Phys. Commun.* **250**, 107071.
- WITTEN, I. H., FRANK, E., HALL, M. A. & PAL, C. J. 2016 *Data Mining: Practical Machine Learning Tools and Techniques*. Elsevier Science.
- WU, J.-L., XIAO, H. & PATERSON, E. 2018 Physics-informed machine learning approach for augmenting turbulence models: a comprehensive framework. *Phys. Rev. Fluids* **3** (7), 074602.
- WYLIE, J. J. & KOCH, D. L. 2000 Particle clustering due to hydrodynamic interactions. *Phys. Fluids* **12** (5), 964–970.
- XIE, C., WANG, J., LI, H., WAN, M. & CHEN, S. 2019a Artificial neural network mixed model for large eddy simulation of compressible isotropic turbulence. *Phys. Fluids* **31** (8), 085112.
- XIE, C., WANG, J., LI, K. & MA, C. 2019b Artificial neural network approach to large-eddy simulation of compressible isotropic turbulence. *Phys. Rev. E* **99** (5), 053113.
- YANG, X. I. A., ZAFAR, S., WANG, J.-X. & XIAO, H. 2019 Predictive large-eddy-simulation wall modeling via physics-informed neural networks. *Phys. Rev. Fluids* **4** (3), 034602.
- YIN, X. & KOCH, D. L. 2007 Hindered settling velocity and microstructure in suspensions of solid spheres with moderate Reynolds numbers. *Phys. Fluids* **19** (9), 093302.

- YIN, X. & KOCH, D. L. 2008 Velocity fluctuations and hydrodynamic diffusion in finite-Reynolds-number sedimenting suspensions. *Phys. Fluids* **20** (4), 043305.
- YOON, D.-H. & YANG, K.-S. 2007 Flow-induced forces on two nearby spheres. *Phys. Fluids* **19** (9), 098103.
- ZAIDI, A. A. 2018 Particle velocity distributions and velocity fluctuations of non-Brownian settling particles by particle-resolved direct numerical simulation. *Phys. Rev. E* **98** (5), 053103.
- ZAIDI, A. A., TSUJI, T. & TANAKA, T. 2014 Direct numerical simulation of finite sized particles settling for high Reynolds number and dilute suspension. *Intl J. Heat Fluid Flow* **50**, 330–341.
- ZHONG, W., YU, A., LIU, X., TONG, Z. & ZHANG, H. 2016 DEM/CFD-DEM modelling of non-spherical particulate systems: theoretical developments and applications. *Powder Technol.* **302**, 108–152.
- ZHOU, Y. & ALAM, MD. M. 2016 Wake of two interacting circular cylinders: a review. *Intl J. Heat Fluid Flow* **62**, 510–537.

© The Author(s), 2020. Published by Cambridge University
Press