14th International Symposium on Turbulence and Shear Flow Phenomena (TSFP14)
Heidelberg, Germany, July 28-31, 2026

PREDICTION OF PARTICLE-LADEN TURBULENT FLOWS USING LONG SHORT-TERM MEMORY NEURAL NETWORKS

Lee Mortimer
School of Chemical and Process Engineering
University of Leeds, Leeds, LS2 9JT, UK
l.f.mortimer@leeds.ac.uk

Michael Fairweather
School of Chemical and Process Engineering
University of Leeds, Leeds, LS2 9JT, UK
m.fairweather@leeds.ac.uk

ABSTRACT

A conditional generative long short-term memory (LSTM) hybrid machine-learning model was developed to predict particle-laden turbulent channel flows using training data obtained from direct numerical simulation with Lagrangian particle tracking (DNS-LPT). Particle trajectories from DNS-LPT databases at $Re_\tau = 180$, spanning a range of practical particle diameters and density ratios, were used to train a variational autoencoder-type model. The governing flow and particle phase parameters, together with an initial particle state, condition the generation of synthetic trajectories. In addition to trajectory reconstruction, the training procedure also incorporates statistical constraints via a loss function on the wall-normal concentration profile and the mean and root-mean-square (RMS) particle velocity profiles. Comparisons with DNS-LPT for $d_p^* = 0.0025$ and $\rho_P^* = 2.5, 1111, 2041$ demonstrate that the model reproduces the mean streamwise particle velocity well and that it captures the wall-normal and spanwise velocity fluctuations with encouraging accuracy, although discrepancies remain in the peaks of the streamwise RMS component at higher density ratios. Very strong agreement is also obtained for the particle concentration profiles. Overall, the results demonstrate the potential of conditional generative LSTM-based hybrid models as low-cost surrogates for particle-laden turbulent flows.

INTRODUCTION

Machine learning (ML) has rapidly become an important computational tool for analysis, prediction, and decision-making across both science and industry, with recent applications in diverse fields such as healthcare (Barragán-Montero et al., 2021) and autonomous vehicles (Gupta et al., 2021). Its value in science lies in data-driven models which improve with experience and data. In fluid dynamics, ML now plays a promising role in turbulence modelling and flow optimisation (Brunton et al., 2020; Lino et al., 2023), typically leveraging high-fidelity simulations or experiments to train predictive models and improve accuracy and performance.

Over the past decade, these approaches have shown strong performance in recovering fluid flows, extracting features of turbulence, and improving accuracy while reducing computational cost for industrial applications (de Oliveira Maionchi et al., 2022). Multiphase systems, particularly particle-laden flows, remain especially challenging due to the effects of turbulent fluctuations, emergent phenomena such as turbophoresis and accumulation, complexities that intensify at high concentrations and Stokes numbers (Ferrante and Elghobashi, 2003; Mortimer et al., 2019). Previous hybrid ML studies for turbulent dispersed flows have demonstrated that data-driven models can reproduce key first- and second-order statistics with much lower cost than high-fidelity simulations (Wang et al., 2018; Jain et al., 2021; Li et al., 2023). However, maintaining accuracy across changing flow conditions remains a central challenge, particularly when the model is required to predict particle concentration fields, which are often omitted from results of such studies (Yang et al., 2023). This issue is important if such methods are to become practical predictive tools rather than case-specific regressors.

In this study we develop and demonstrate a predictive framework centred on the application of long short-term memory (LSTM) networks trained on databases generated by direct numerical simulation (DNS) with Lagrangian particle tracking (LPT). The model is trained using DNS-LPT trajectory databases obtained at $Re_\tau = 180$, and includes a large range of practical particle diameters and density ratios. By explicitly modelling temporal dependencies and correlations, LSTMs capture long-range, non-Markovian effects in particle–turbulence interactions (Lino et al., 2023) and are applied here to learn to reproduce particle trajectories in turbulent channel flow. The present work builds upon previous attempts by considering a conditional generative modelling framework for particle-laden turbulent channel flows. Rather than using ML solely to predict the next particle state from either its current state or its history, the objective is now to learn a probabilistic representation of trajectory evolution conditioned both on the governing flow conditions, as well as the particle parameters. In this approach, trajectory segments extracted from DNS-LPT databases are used to train a conditional generative model that receives the Reynolds number, particle diameter, density ratio and an initial particle state as conditioning variables. A latent stochastic variable is also used to represent the variability between trajectories evolving under the same macroscopic bulk conditions. We further aim to assess the model based on its ability to replicate statistical observations at a range of Stokes numbers. These developments facilitate a move toward ML surrogates that reproduce not only low-order statistics, but also the qualitative appearance and temporal coherence of Lagrangian particle trajectories over long times. This is a more demanding requirement than matching mean quantities alone, since realistic particle paths depend on persistence, temporal correlations, and repeated interactions with the turbulent structures of the carrier flow. These capabilities position conditional generative LSTM-based hybrid ML models as promising components for digital-twin technology, where pre-trained ML models can work towards enabling near-real-time simulation and control.

METHODOLOGY

Direct Numerical Simulation

To obtain the Eulerian turbulent channel flow, direct numerical simulations were performed using Nek5000 (Fischer et al., 2008) to obtain a statistically stationary flow field at shear Reynolds number, $Re_\tau = 180$. The spectral element mesh represents a channel flow domain of size $12\delta \times 2\delta \times 6\delta$, corresponding to approximately 3.9 million Gauss-Lobatto-Legendre nodes. Grid spacing was refined in the wall-normal direction (y^*) to accurately resolve near-wall structures while maintaining uniform spacing in the streamwise (x^*) and spanwise (z^*) directions. Periodic boundary conditions were applied in the streamwise and spanwise directions, with no-slip and impermeability conditions imposed at the walls located at $y^* = \pm 1$. The flow was driven by a constant pressure gradient calculated to obtain the desired Reynolds number and advanced in time using a fixed non-dimensional timestep of $\Delta t^* = 0.005$. Note that spatial and temporal variables are non-dimensionalised here and throughout using bulk flow properties, as indicated by symbols with an asterisk superscript (*).

Lagrangian Particle Tracking

For the dispersed phase, solid particles, treated as point-like, impenetrable spheres, were tracked using an in situ Lagrangian particle tracking routine (Mortimer et al., 2019). After each fluid timestep, the local fluid quantities required by the particle equations of motion were interpolated spectrally at each particle location, and the particle translation was then updated concurrently with the flow solution. Particle equations of motion based upon force balances were integrated concurrently with the fluid solver using a fourth-order Runge-Kutta scheme using a fixed time step equivalent to that of the fluid phase solver. Forces included contributions from drag, lift, virtual mass and pressure-gradient effects. Collisions between particles and the channel walls were treated as perfectly elastic, such that upon impact, the wall-normal component of the particle velocity was reversed. In the streamwise and spanwise directions, periodicity was enforced so that particles exiting one side of the domain were reintroduced in the correct location at the opposite boundary.

Particle information was written every 10 fluid timesteps, corresponding to a saved interval of $\Delta t_s^* = 0.05$. Each instantaneous file contained 10,000 particles and stored the particle positions and instantaneous velocity components; therefore the training database consisted of the variables:

$$\left[x^*,\ y^*,\ z^*,\ v_x^*,\ v_y^*,\ v_z^*\right]_{i,t} \qquad (1)$$

where i represents the particle index and t represents the timestep. These were extracted together with the associated simulation labels: Re_τ, d_p^* and ρ_P^* representing the shear Reynolds number; the particle diameter relative to the channel half-height; and the density-ratio relative to that of the fluid respectively.

Conditional Generative LSTM Model

To preprocess the training data, the DNS-LPT particle datasets were reorganised into a suitable form for use with the conditional generative sequence model. In the original simulations, particle information was stored as a series of instantaneous snapshots, each containing the positions and velocities of 10,000 particles at the saved instantaneous time step. These data were first reordered into Lagrangian trajectory sequences by linking each particle across consecutive saved snapshots, after which overlapping windows of fixed length $T_{WIN} = 20$ were extracted from the full particle histories. Each sequence retained the wall-normal particle position and the three particle velocity components, together with the associated simulation parameters. The initial particle state at the beginning of each sequence was also retained explicitly so that the model could generate trajectories from a physically admissible starting condition. All variables were then normalised using statistics obtained from the training subset, and the resulting database divided into training, validation and test sets for subsequent model development and assessment using a 70:20:10 split.

The machine-learning model developed in the present work is a conditional generative sequence model based on a variational autoencoder-type formulation with long short-term memory networks. The use of LSTM layers as opposed to previous attempts was motivated due to particle motion in turbulent flow being inherently history-dependent such that the instantaneous particle response depends not only on its present location, but also on its recent velocity history and accumulated momentum. This is especially important for inertial particles, whose motion remains correlated over multiple timesteps, which increases with Stokes number. The generative formulation was introduced so that the model could produce multiple plausible trajectory realisations for the same macroscopic conditions, rather than a single deterministic prediction, better representing the stochastic nature of such trajectories.

The model also consisted of an encoder and a decoder. During training, the encoder receives a DNS-LPT trajectory sequence and maps it into a low-dimensional latent representation. LSTM layers are then used within the encoder to extract the temporal structure of the input trajectory before projection into the latent space. The latent representation was characterised by a mean and log-variance, and a latent vector was sampled from this distribution during training using standard variational reparameterisation techniques.

The decoder is then used to reconstruct the trajectory sequence from the sampled latent vector together with the condition information, $\mathbf{c} = [Re_\tau,\ d_p^*,\ \rho_P^*]$. In the present formulation, the conditioning variables comprised the simulation parameters, as well as the initial particle state $\mathbf{s}_0$. The conditional generative model therefore learned the following mapping:

$$(\mathbf{c}, \mathbf{s}_0, \mathbf{z}) \longrightarrow \left[y^*(t), v_x^*(t), v_y^*(t), v_z^*(t)\right]_{t=1}^{T}, \qquad (2)$$

where $\mathbf{z}$ is the sampled latent vector.

Training of the model was based not only on reconstruction of the DNS-LPT trajectory sequences, but also on recovery of the first- and second-order dispersed-phase statistics. The total loss function therefore consisted of several components. First, a reconstruction term was used to discourage mismatch between the generated and DNS-LPT trajectory segments. Second, a Kullback-Leibler regularisation term was imposed on the latent space and enabled sampling of new trajectories during generation. These

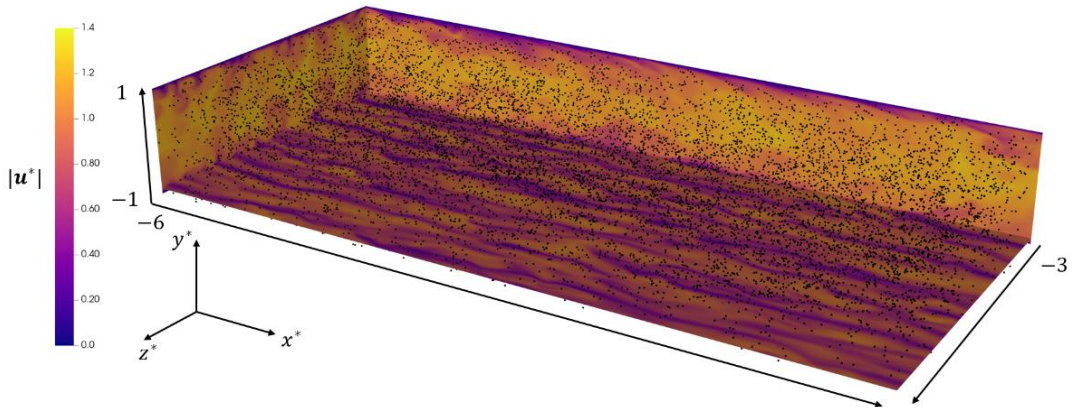

Figure 1: Computational domain used for DNS-LPT of particle-laden turbulent channel flow. Streamwise velocities are based on $Re_\tau = 180$ flow. Black spheres represent the instantaneous position of particles.

two components formed the basic variational objective. Additional statistical loss terms were included to penalise mismatch in the wall-normal concentration profile, the mean particle velocity profiles and the root-mean-square (RMS) of the particle velocity fluctuation profiles. These were evaluated within the training batches as functions of wall-normal position. The concentration term constrained the spatial distribution of particles across the channel, while the mean and RMS terms constrained the first- and second-order particle velocity statistics to remain close to those of the original DNS-LPT database.

The overall loss may therefore be written as:

$$\mathcal{L}_T = \mathcal{L}_\mathrm{r} + \beta\mathcal{L}_\mathrm{KL} + \lambda_\phi\mathcal{L}_\phi + \lambda_u\mathcal{L}_u + \lambda_{u'}\mathcal{L}_{u'} \quad (3)$$

where β, λ_ϕ, λ_u and $\lambda_{u'}$ are weighting coefficients which controlled the relative influence of latent regularisation and statistical matching. The weighting coefficients were set to $\beta = 10^{-4}$, $\lambda_\phi = 1$, $\lambda_u = 1$ and $\lambda_{u'} = 1$, with equal weighting also applied to the reconstruction terms for y^*, v_x^*, $v_y^*$and v_z^*. By combining the sequence reconstruction objective with explicit statistical constraints, the model was trained not only to generate plausible short trajectory segments, but also to reproduce the statistics of the DNS-LPT particle field.

Hybrid ML Technique

Once trained, the model was used to generate synthetic particle trajectories for prescribed values of Re_τ, d_p^* and ρ_P^*. For a given set of conditions, an initial particle state was provided and a latent vector was sampled from the learned latent distribution. The decoder then produced a new particle trajectory segment consistent with those conditions. Since correct wall behaviour was a necessary part of the target physics, the trajectory algorithm incorporated an explicit reflective wall treatment consistent with the underlying DNS-LPT calculations. Whenever a generated particle position crossed either wall at $y^* = \pm 1$, the particle was reflected back into the channel domain, corresponding to an elastic rebound. In this manner, wall penetration was prevented by construction and the generated particle motion remained consistent with the boundary interaction assumed in the Lagrangian simulations. Periodic behaviour in the streamwise and spanwise directions was also implemented to match the underlying DNS-LPT framework. Repeated sampling yielded large synthetic particle ensembles analogous to those obtained from the DNS-LPT.

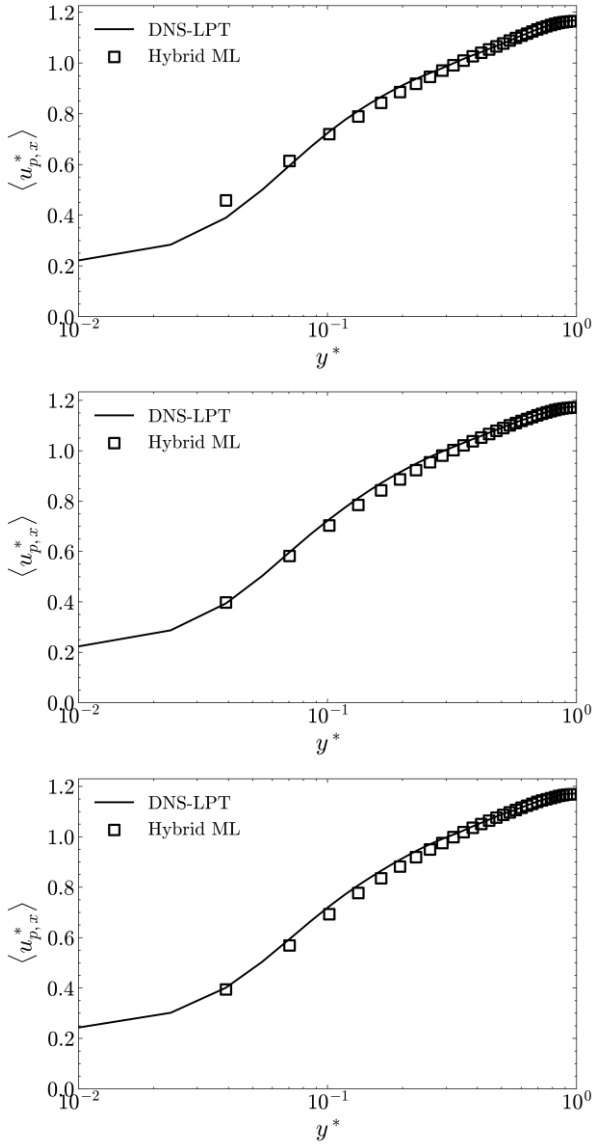

Figure 2: Mean streamwise particle velocity, comparing DNS-LPT against the LSTM hybrid model predictions. Upper: $\rho_P^* = 2.5$, middle: $\rho_P^* = 1111$, lower: $\rho_P^* = 2041$.

RESULTS AND DISCUSSION

The performance of the conditional generative LSTM-based hybrid ML model is assessed here by comparison with DNS-LPT data for the turbulent channel flow at $Re_\tau = 180$. Although the model was trained on a broad range of particle Stokes numbers, the present analysis focuses on the case $d_p^* = 0.0025$ while varying the density ratio, in order to examine the extent to which the method reproduces both the statistical behaviour and qualitative trajectory structure of the dispersed phase.

Figure 1 presents an instantaneous view of the DNS-LPT particle-laden channel flow at $Re_\tau = 180$, illustrating the computational domain together with the dispersed particle point field superimposed on the streamwise velocity distribution. The figure shows the strongly inhomogeneous structure of the near-wall region, with elongated low speed streaks visible in the streamwise direction, consistent with the wall-bounded turbulence field expected for this geometry and flow conditions.

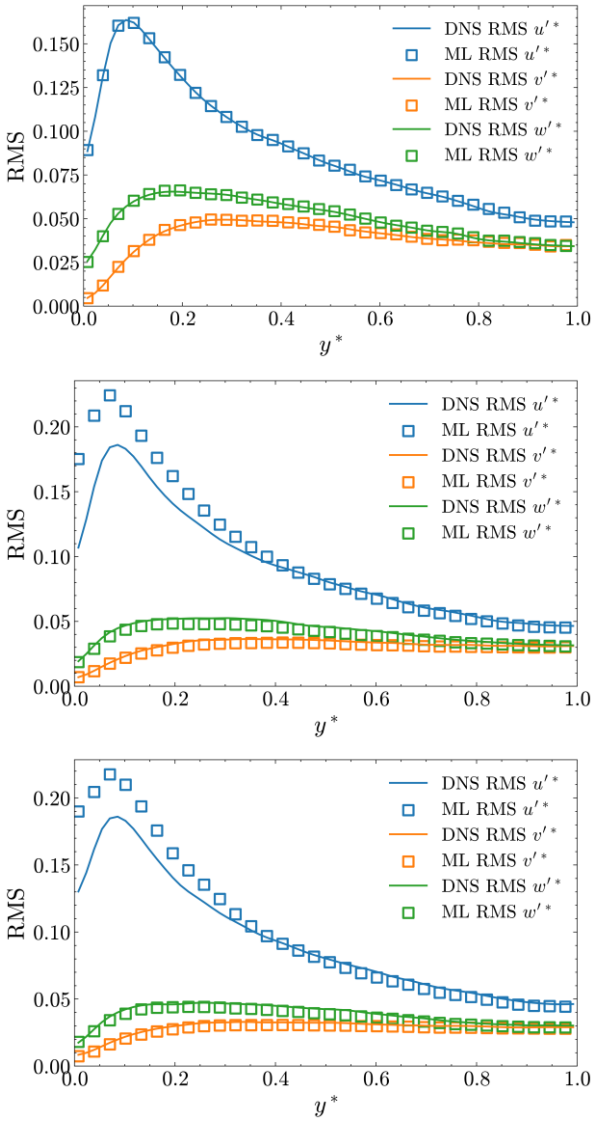

Figure 3: RMS of particle velocity fluctuations, comparing DNS-LPT against the LSTM hybrid model predictions. Upper: $\rho_P^* = 2.5$, middle: $\rho_P^* = 1111$, lower: $\rho_P^* = 2041$.

Figure 2 compares the mean streamwise particle velocity profiles predicted by the conditional generative LSTM hybrid ML model against the DNS-LPT data for $Re_\tau = 180$, $d_p^* = 0.0025$, across the three density ratios considered. In all three cases, the model recovers the overall wall-normal variation in $\langle u_{p,x}^* \rangle$ well, with the expected increase from the near-wall region toward the channel centre captured consistently. The agreement is strongest through the mid-channel and bulk flow regions, where the ML predictions follow the DNS-LPT trend closely and remain only slightly offset from the reference data. These observations indicate that the present conditional generative formulation retains one of the main strengths of earlier approaches: accurate recovery of the first-order streamwise particle statistics. Since the mean streamwise velocity largely reflects the physical interaction dynamics associated with particle inertia, turbulent transport and repeated wall-bounded motion, this agreement is encouraging.

Figure 3 compares the RMS of particle velocity fluctuations from both techniques. It shows that the conditional generative LSTM model offers strongest agreement with DNS-LPT in the wall-normal and spanwise RMS components, while the largest discrepancies occur in the streamwise fluctuations, particularly as the density ratio, and hence particle Stokes number, increases. This is consistent with the physical role of the streamwise motion in wall-bounded turbulence. The streamwise particle response is influenced not only by local fluctuating motion, but also by sustained interaction with the earlier pictured (Figure 1) elongated near-wall streaks, as well as mean shear. In the present framework, the model is trained to recover both trajectory sequences and statistical profiles from relatively short windows, and although the LSTM architecture captures temporal dependence, the streamwise component remains the one most strongly affected by long-range persistence and it is likely that the accumulation of small errors occurs due to this. As particle inertia increases, these effects become more pronounced because the particles respond less directly to the instantaneous carrier-phase motion and retain momentum over longer times, making the accurate reconstruction of $u_{p,x}^{'}$ more demanding than for the weaker cross-stream components. A further reason could be that the wall-normal and spanwise fluctuations are smaller in magnitude and more tightly constrained by the channel geometry and elastic wall treatment, so the model is aided by the imposed physical structure in these directions. At higher density ratios, within the bulk flow, the generated trajectories appear to preserve the overall level and wall-normal variation of the transverse fluctuations reasonably well, while the streamwise RMS peaks in the most turbulent regions become more difficult to match exactly, especially close to the wall where the gradients are strongest and where inertial particles experience repeated interaction with coherent structures. The discrepancy in the streamwise direction therefore likely reflects the greater challenge of reproducing these longer temporally correlated, shear-dominated dynamics within a conditional generative framework that uses larger windows, rather than a failure of the model to capture second-order particle statistics.

Figure 4 is particularly promising, since the conditional generative LSTM model reproduces the wall-normal particle concentration profiles very closely for all three density ratios considered. In each case, the generated concentrations remain essentially uniform across the bulk of the channel, with the

near-wall increase at higher density ratios captured consistently and in good agreement with the DNS-LPT data. This is an important result because concentration is a more difficult to predict quantity than mean velocity alone, it being the outcome of many phenomena such as turbophoresis, particle dispersion, wall and near-wall turbulent structure interaction and long-time trajectory evolution. The close agreement therefore indicates that the model is not simply recovering local kinematic behaviour, but is also preserving the balance of motions towards and away from the wall without introducing any appreciable artificial drift. That this behaviour is maintained across a wide range of density ratios further suggests that the conditional generative formulation is able to retain the correct dispersed-phase spatial distribution even as particle inertia increases, which is highly encouraging for its use as a practical surrogate for DNS-LPT particle transport.

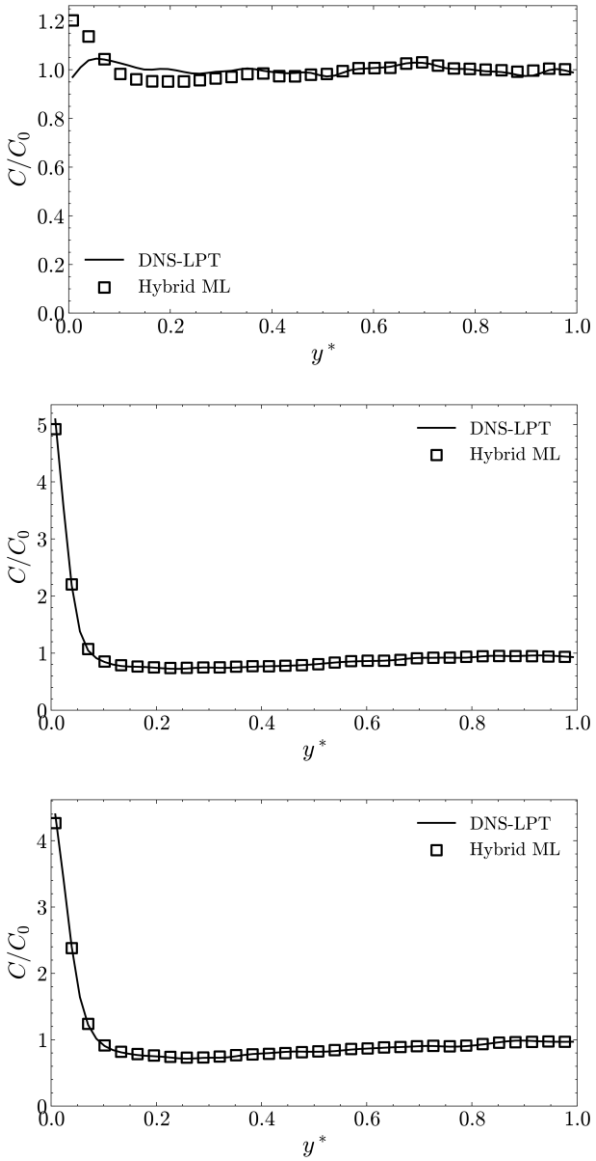

Figure 4: Particle concentrations, C/C_0, comparing DNS-LPT against the LSTM hybrid model predictions. Upper: $\rho_P^* = 2.5$, middle: $\rho_P^* = 1111$, lower: $\rho_P^* = 2041$.

Finally, Figure 5 provides a qualitative assessment of the generated particle motion in the $x^* - y^*$ plane. Here we compare trajectory windows generated from the initial starting point of a particle in the original dataset. In all three cases, the model reproduces the strongly streamwise-aligned nature of the DNS-LPT trajectories, while also retaining the shorter trajectory lengths close to the wall associated with the lower local streamwise velocities in this region. This behaviour is particularly important, since it shows that the model is capturing not only the correct mean advection, but also the reduced particle displacement and associated streamwise velocities near the wall.

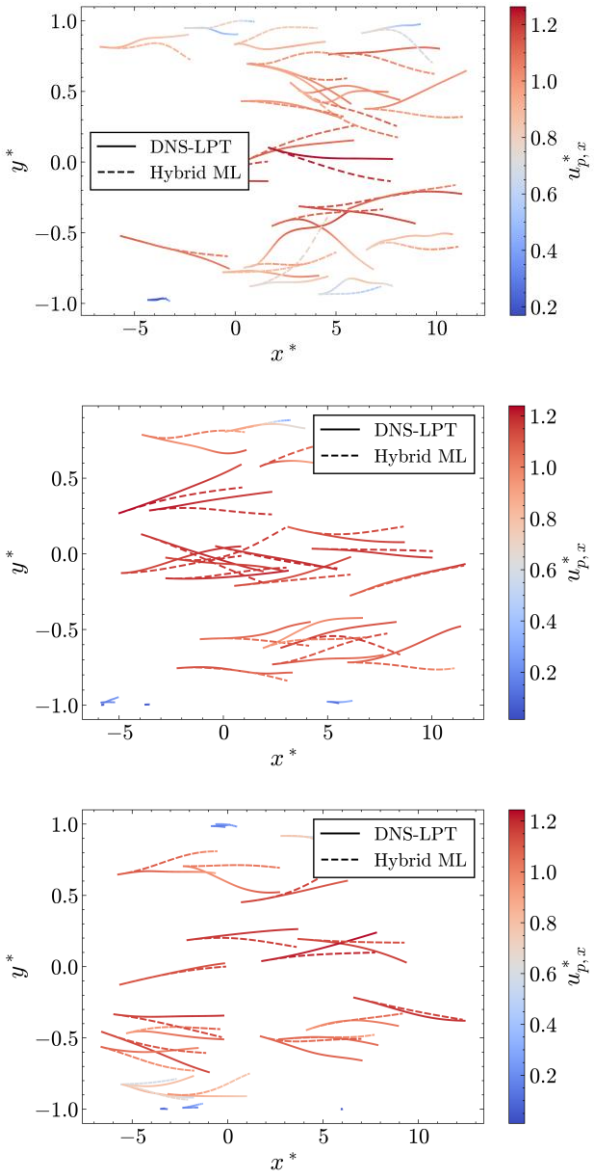

Figure 5: Comparison of particle trajectories in the $x^* - y^*$ plane between DNS-LPT and the LSTM hybrid model predictions. Upper: $\rho_P^* = 2.5$, middle: $\rho_P^* = 1111$, lower: $\rho_P^* = 2041$.

Another encouraging feature is the curvature and undulation frequency of the generated paths, especially at the lower Stokes numbers where the trajectories are wavier, which is consistent with the DNS-LPT behaviour. The particles do

not simply travel in straight lines but instead show the oscillatory wall-normal motions expected due to rapid response to the surrounding turbulent structures. This is qualitatively correct for the low-inertia cases, where particle motion remains strongly influenced by the carrier-phase fluctuations. The overall appearance of the trajectories is promising, since the model preserves the key visual characteristics of the DNS-LPT motion rather than producing overly smooth or obviously artificial particle paths.

CONCLUSIONS

A conditional generative LSTM-based hybrid ML model has been developed and applied to predict particle-laden turbulent channel flows using DNS-LPT training data. When trained on trajectory sequences at $Re_\tau = 180$, spanning a range of particle diameters and density ratios, the model is shown to reproduce the principal dispersed-phase statistics with encouraging accuracy. These include the mean streamwise particle velocity, the wall-normal and spanwise RMS velocity fluctuations, and the near-uniform wall-normal concentration profiles across a range of density ratios. Visual comparison of the trajectories in the $x^* - y^*$ plane further indicates that the model preserves the main qualitative features of the DNS-LPT motion, including the streamwise-aligned nature of the paths, the shorter near-wall trajectories associated with lower local velocities, and the wavy behaviour expected for the low-Stokes-number cases.

The largest remaining discrepancies occur in the streamwise RMS component as particle inertia increases, reflecting the greater difficulty of reproducing the longer-memory, shear-dominated particle dynamics associated with interaction with near-wall streaks. That said, the overall level of agreement obtained here demonstrates that conditional generative LSTM-based hybrid ML models provide a viable method for generating synthetic particle trajectory ensembles which are both statistically and qualitatively consistent with those produced by DNS-LPT. In doing so, the computational cost of obtaining such results is greatly reduced compared to that associated with repeated high-fidelity simulation. The fact that the conditional generative formulation is able to recover the correct overall concentration behaviour whilst also generating trajectories with the appropriate qualitative character suggests that the inclusion of both temporal memory through the LSTM architecture and statistical constraints via the loss function used during training is beneficial. In particular, the present results indicate that these two elements together allow the model to retain the balance between local fluctuating motion, near-wall transport and streamwise advection that governs the behaviour of inertial particles in wall-bounded turbulence.

The present work also highlights a number of clear directions for further development. Most importantly, the remaining discrepancies in the streamwise RMS component at higher density ratios. Future work will therefore focus on extending the approach to a wider range of Stokes numbers and flow conditions, as well as investigating refinements to the conditioning strategy, sequence length and loss formulation in order to improve the recovery of these more demanding trajectory features.

The ML approach represents a useful step towards near real-time predictive tools for particle-laden turbulent flows in practical engineering applications. These capabilities are particularly promising in the context of digital twin development, where pre-trained hybrid ML surrogates of this type could be used to provide rapid prediction of particle transport, concentration and trajectory response under changing operating conditions, without the need to repeatedly perform DNS-LPT calculations.

ACKNOWLEDGEMENTS

The authors would like to thank Sellafield Ltd. for funding.

REFERENCES

Barragán-Montero, A., Javaid, U., Valdés, G., Nguyen, D., Desbordes, P., Macq, B., Willems, S., Vandewinckele, L., Holmström, M., Löfman, F., and Michiels, S., 2021, "Artificial Intelligence and Machine Learning for Medical Imaging: A Technology Review", *Physica Medica*, Vol. 83, pp. 242-256.

Brunton, S. L., Noack, B. R., and Koumoutsakos, P., 2020, "Machine Learning for Fluid Mechanics", *Annual Review of Fluid Mechanics*, Vol. 52, pp. 477-508.

de Oliveira Maionchi, D., Ainstein, L., dos Santos, F. P., and de Souza Júnior, M. B., 2022, "Computational Fluid Dynamics and Machine Learning as Tools for Optimization of Micromixers Geometry", *International Journal of Heat and Mass Transfer*, Vol. 194, 123110.

Ferrante, A., and Elghobashi, S., 2003, "On the Physical Mechanisms of Two-Way Coupling in Particle-Laden Isotropic Turbulence", *Physics of Fluids*, Vol. 15, pp. 315-329.

Fischer, P. F., Lottes, J. W., and Kerkemeier, S. G., 2008, "Nek5000", available at: http://nek5000.mcs.anl.gov/; accessed 1st September (2008).

Gupta, A., Anpalagan, A., Guan, L., and Khwaja, A. S., 2021, "Deep Learning for Object Detection and Scene Perception in Self-Driving Cars: Survey, Challenges, and Open Issues", *Array*, Vol. 10, 100057.

Jain, P., Choudhury, A., Dutta, P., Kalita, K., and Barsocchi, P., 2021, "Random Forest Regression-Based Machine Learning Model for Accurate Estimation of Fluid Flow in Curved Pipes", *Processes*, Vol. 9, 2095.

Li, K., Savari, C., Sheikh, H. A., and Barigou, M., 2023, "A Data-Driven Machine Learning Framework for Modeling of Turbulent Mixing Flows", *Physics of Fluids*, Vol. 35, 015150.

Lino, M., Fotiadis, S., Bharath, A. A., and Cantwell, C. D., 2023, "Current and Emerging Deep-Learning Methods for the Simulation of Fluid Dynamics", *Proceedings of the Royal Society A*, Vol. 479, 2275.

Mortimer, L. F., Njobuenwu, D. O., and Fairweather, M., 2019, "Near-Wall Dynamics of Inertial Particles in Dilute Turbulent Channel Flows", *Physics of Fluids*, Vol. 31, 063302.

Wang, Z., Xiao, D., Fang, F., Govindan, R., Pain, C. C., and Guo, Y., 2018, "Model Identification of Reduced Order Fluid Dynamics Systems Using Deep Learning", *International Journal for Numerical Methods in Fluid*s, Vol. 86, pp. 255-268.

Yang, Z., Li, K., and Barigou, M., 2023, "Experimentally Trained Hybrid Machine Learning Algorithm for Predicting Turbulent Particle-Laden Flows in Pipes", *Physics of Fluids*, Vol. 35, 113309.