

---

# A Validated LBM Dataset and Pipeline for Surrogate Modeling of Turbulent 3D Obstructed Channel Flows

---

**Lukas Schröder**

Erlangen National High Performance Computing Center (NHR@FAU)  
lukas.ls.schroeder@fau.de

**Shubham Kavane   Harald Köstler**

Chair of Computer Science 10 (System Simulation)  
Friedrich-Alexander-Universität Erlangen-Nürnberg  
{shubham.kavane, harald.koestler}@fau.de

## Abstract

Evaluating neural operators for 3D turbulent flow requires validated datasets with physical benchmarks. We present a reproducible pipeline generating training data for 3D channel flows around generated geometries at  $Re=1,000\text{--}10,000$ . Our lattice Boltzmann solver with cumulant collision operators is rigorously verified against experimental measurements (Strouhal number, drag coefficients, turbulent fluctuations) with comprehensive grid convergence studies at resolution  $1024 \times 512 \times 512$ . Building upon an established framework, this validated pipeline enables standardized surrogate model comparison. We outline planned systematic evaluation of Fourier Neural Operator and U-Net variants on forecasting, super-resolution, and error correction tasks, using physics-informed metrics to assess turbulent energy cascade representation. Future work will compare computational efficiency between numerical solvers and neural surrogates, exploring practical application. We seek community feedback on our validation approach, planned benchmark methodology, and evaluation priorities for neural operators in turbulent flows.

## 1 Introduction & motivation

Computational fluid dynamics simulations of 3D turbulent flows at moderate to high Reynolds numbers remain computationally expensive. Despite ease of parallelization on modern GPUs, time-resolved predictions still require substantial computational resources. This motivates the development of machine learning surrogate models that can distill knowledge from expensive simulations into fast neural network (NN) predictions, potentially reducing computational effort while preserving essential physical accuracy. While databases such as the Johns Hopkins Turbulence Database [1] provide high-quality simulations, they focus primarily on canonical configurations (isotropic turbulence, channel flows) rather than 3D flows around complex geometries relevant for engineering applications.

Recent advances in neural operators, particularly Fourier Neural Operators (FNO) [2] and U-Net [3] architectures, have shown promise for fluid dynamics applications [4, 5]. However, rigorous comparison of state-of-the-art models remains challenging due to the lack of validated, high-quality datasets with consistent physical benchmarks. Fair model assessment requires not only comparable training data but also physically meaningful evaluation metrics, especially for turbulent flows where representation of the energy cascade and high-frequency dynamics are critical [6, 7]. Furthermore, a fair assessment requires not only physically accurate results but also consistent measurement of computational cost from highly optimized implementations of both traditional solvers and neural models.

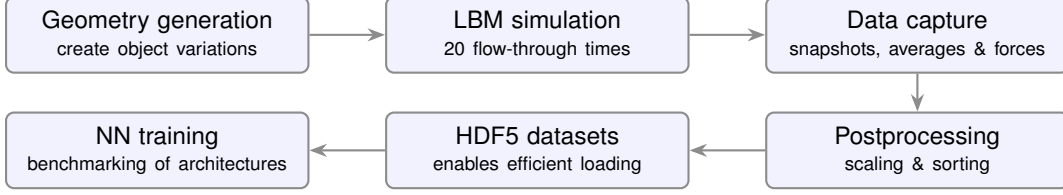


Figure 1: Pipeline workflow from geometry generation for NN training [8]

This work addresses this gap by presenting a reproducible and rigorously validated pipeline for simulating flows around objects across Reynolds numbers from 1000 to 10,000. Building upon an established framework [8], which established the geometry generation and data processing infrastructure, this work introduces a rigorously validated lattice Boltzmann method (LBM) [9] solver configuration with cumulant collision operators and systematic experimental verification. Our LBM-based approach is verified against experimental measurements including Strouhal number, recirculation length, turbulent energy fluctuations, and drag coefficients, with grid convergence studies ensuring numerical accuracy. We provide preliminary baseline results of a Fourier Neural Operator predicting time-averaged flow fields as proof-of-concept, demonstrating the pipeline’s utility. Finally, we outline our planned framework for systematic comparison of U-Net and Fourier-based Neural Operators across multiple prediction tasks. The gained insights on physics-based NNs also contribute to the research on differential solvers.

## 2 Pipeline description

The pipeline in Figure 1 provides code that handles data generation, simulation, and curation for ready-to-use results in NN training. Starting with a module that generates derivative objects from a set of primitives [Cylinder, Rectangle, Sphere, Torus, Wedge], which can also be replaced by user-provided objects. The flow around objects can then be simulated in the defined Reynolds number range [1000, 10000] with varying angles of attack using a LBM simulation (detailed in Section 3.2). The simulation captures the average flow after 15 flow-through times, recording the mean velocity and pressure of the flow, together with 20 consecutive snapshots separated by  $25 \delta t$  in lattice units. Additionally, the pipeline records the accumulated forces on the object in x and y directions and 2D snapshots of the center x-y plane for monitoring the development of the entire simulation. Finally, the compressed files per simulation can be unpacked, scaled, and repacked into HDF5 files with specified resolutions in the range of  $[(512 \times 256 \times 256), (64 \times 32 \times 32)]$ , along with the additional flag field, which captures the spatial occupancy mask or a signed distance field.

## 3 Dataset creation and validation

### 3.1 Generated dataset

The preliminary dataset comprises flows around 42 procedurally generated objects within a computational domain of dimensions  $2 \times 1 \times 1$  (in units of channel height  $L$ ). Each object is positioned at  $x/2L = 0.3$ , centered in the y-z plane and spans  $1/6L$  (86 cells). Reynolds numbers  $Re_c$  from 1000 to 10,000, in the transitional-to-turbulent regime, are achieved by adjusting the lattice viscosity while maintaining constant inlet velocity, with  $Re$  defined based on channel height. The inflow is implemented using a prescribed velocity boundary condition, the outflow employs extrapolation, and the lateral walls enforce no-slip conditions. The domain is resolved at  $1024 \times 512 \times 512$  to ensure quality, as mesh refinement in LBM reduces accuracy due to interpolation artifacts. Data is stored at half resolution; simulations require 3h on 4 Nvidia H100 GPUs.

### 3.2 waLBerla framework and solver setup

Simulations employ the waLBerla framework (licensed under GPLv3), which provides highly optimized code-generated LBM implementations with efficient parallelization capabilities [10]. This framework’s performance characteristics allow for a fair computational effort comparison between it and surrogates later.

The cumulant collision operator with a D3Q27 stencil provides improved numerical stability and accuracy compared to moment-based multiple relaxation time models [11]. A fourth-order advection-diffusion correction scheme with limiting enhances numerical precision and stability for turbulent flows [12]. This approach accurately captures the energy cascade and maintains well-resolved energy spectra essential for turbulence representation [6]. The method’s inherent controlled dissipation effectively models subgrid-scale effects, eliminating the need for explicit subgrid models (SGS) such as Smagorinsky closures. Testing confirmed that adding SGS modeling reduced physical accuracy. Notably, the correction introduces instabilities in laminar flows with elevated kinematic viscosity, so the correction is disabled for  $Re_c < 3000$ . Fluid-solid interaction at object boundaries employs a quadratic bounce-back scheme, which provides second-order spatial accuracy at the cost of computing wall distances [11]. This approach ensures physically accurate boundary conditions critical for turbulent flow development.

### 3.3 Validation

The primary objective of this dataset is providing physically accurate training data that enables rigorous assessment of NN surrogates’ ability to capture turbulent flow physics. Comprehensive validation was therefore conducted across multiple physical parameters representing the dynamics and the average properties of the flow.

Turbulent regime validation uses the experimental measurements of Choi and Park [13] for flow around a sphere at  $Re_s$  1000, relative to sphere diameter  $D$ , corresponding to  $Re_c$  6000 relative to our channel height. This reference provides validation of Strouhal number, recirculation length, and turbulent velocity fluctuation statistics.

To verify dataset validity extends to laminar conditions, drag coefficient validation was performed at the Reynolds number extremes:  $Re_c$  600 and  $Re_c$  12000 even outside our simulated range. Additionally, a grid convergence study [14] on drag coefficients confirms numerical accuracy at our chosen resolution. Validation results and quantitative comparisons are presented in the following, with all the plots and referenced data in the appendix A.1.

**Drag Coefficient & Grid Convergence:** The drag coefficient is an expressive and simple to conduct measurement, because it summarizes the pressure distribution, viscous effects and separation and wake dynamics in one parameter while being sensitive to numerical errors as well. The coefficients show a convergent behavior in Table 1, matching experimental values within  $\pm 6\%$  of Roos and Willmarth [15] with our final spacing of 0.0019 ( $L \sim 512$  cells). The error growth with  $Re_s$  above 1000, which demands more attention in future research.

Table 1: Drag coefficients at different Reynolds numbers and grid resolutions

Reynolds number		drag coefficient at resolution					reference [15]
$Re_s$	$Re_c$	128	256	384	512	1024	
100	600	1.122	1.108	1.106	1.106	—	1.08
500	3000	0.573	0.562	0.563	0.562	—	0.547
1000	6000	0.479	0.460	0.466	0.467	0.472	0.472
2000	12000	0.434	0.407	0.407	0.404	—	0.430

**Average Field with recirculation length:** The normalized mean streamwise velocity field in Figure 3 exhibits a recirculation bubble of  $l_r/D = 1.8$  at  $Re_s$  1000, consistent with the measured range of  $1.76 \pm 0.07$  of Choi and Park [13]. The velocity profiles in Figure 6 and 7 closely match the measured mean flow distribution, confirming accurate capture of the mean flow characteristics.

**Strouhal number:** Spectral analysis of vertical velocity fluctuations in the wake reveals a dominant frequency corresponding to  $St \approx 0.18$  at  $Re_s$  1000 in Figure 5, agreeing with literature values and confirming accurate capture of vortex shedding dynamics.

**Streamwise turbulent intensity:** This serves as an indicator of the scale of the fluctuations relative to the distance and thus is essential in capturing the mechanics of turbulent energy dissipation. Figure 4 shows two peaks 1.5 and 2.4  $x/D$  which is identical to the experiment of Choi and Park [13] with the scale in between their measurements and their referenced sources.

**Turbulent structures:** The instantaneous snapshot in Figure 8 reproduce characteristic turbulent features including hairpin vortices with identifiable heads and trailing legs, validating the solver’s capability to resolve fine-scale turbulent coherent structures.

The validated solver configuration is applied unchanged to diverse geometries, consistent with established CFD surrogate modeling practice, where canonical benchmarking validates numerical frameworks for application across geometrically diverse cases.

## 4 Baseline results

We present preliminary baseline results in Table 2 for predicting the mean streamwise velocity field from a binary geometry mask using factorized FNO [16], original FNO [2], a combination of both, and a U-Net [3] model as a proof-of-concept. These baselines establish reference levels of predictive accuracy and computational efficiency, serving as a foundation for evaluating and comparing future improvements in surrogate modeling. The U-Net outperforms the other base models and thus can be set as the most promising candidate for further research. The settings are listed under appendix A.2.

Table 2: Preliminary baseline results on validation data (details in appendix A.2)

model	layers	MSE	max MAE	NRMSE	MAPE	sample/s
FNO	4	$6.01 \cdot 10^{-4}$	0.0338	0.0359	5.71	20.77
Hybrid F/FNO	4	$6.69 \cdot 10^{-4}$	0.0345	0.0379	5.80	16.36
FFNO	4	$11.13 \cdot 10^{-4}$	0.0464	0.0495	7.50	12.30
U-Net	5	$1.22 \cdot 10^{-4}$	0.0158	0.0159	2.17	32.87

## 5 Planned model comparison

Having established a validated dataset, we now focus on systematic model comparison. The experimental setup covers fundamental surrogate modeling tasks: (1) predicting mean flow fields from geometry and Reynolds number, and (2) next-timestep forecasting from current state. When architectures support it, we will extract drag coefficients ( $c_d$ ), perform super-resolution from low-resolution snapshots, and evaluate temporality through multi-timestep input/output configurations.

From a methodological perspective, we focus on advancements of baseline U-Nets and FNOs. A critical challenge in turbulent flow prediction is spectral bias—the tendency to oversmooth fine-scale structures [17]. Following Oommen et al. [7], we plan to evaluate expanded loss functions combining adversarial and spectral components alongside standard  $L_2$  losses. We will apply these techniques to Fourier-based architectures, U-Net variants, and hybrid combinatory models [18, 5]. Additionally, exploring attention mechanisms integrated as channel/spatial gates in Time-Conditioned U-Nets ([19]) and as linear attention coupling in FNOs ([20]). Finally, rigorous comparison of computational costs between solvers and surrogates will assess practical applicability in 3D turbulent flows.

## 6 Limitations

One of the limitations of our pipeline is the high computational and storage demand of high-resolution turbulent data generation, limiting dataset and resolution scalability. The solver is limited for the transitional-to-turbulent regime ( $Re$  1,000–10,000). The baseline results lack variance estimation or comprehensive hyperparameter tuning, are predicting a significantly lower resolution with reduced turbulent scales and serve only as proof-of-concept demonstrations.

## 7 Conclusion

This work-in-progress presents a validated pipeline for generating physically accurate 3D turbulent flow datasets around complex geometries. We welcome feedback from the community on our planned systematic comparison of U-Net and Fourier Neural Operator variants, particularly regarding methodological priorities and evaluation metrics for flow prediction surrogates. Ultimately, this

validated pipeline aims to advance physics-based NN, also as building block of hybrid differentiable numerical solvers.

## Acknowledgments and Disclosure of Funding

The authors gratefully acknowledge the HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). NHR funding is provided by the German Federal Ministry of Education and Research and the state governments participating on the basis of the resolutions of the GWK for the national high-performance computing at universities ([www.nhr-verein.de/unsere-partner](http://www.nhr-verein.de/unsere-partner)) by federal and Bavarian state authorities. NHR@FAU hardware is partially funded by the German Research Foundation (DFG) – 440719683

## References

- [1] Kalin Kanov, Randal Burns, Cristian Lalescu, and Gregory Eyink. The johns hopkins turbulence databases: An open simulation laboratory for turbulence research. *Computing in Science & Engineering*, 17(5):10–17, 2015.
- [2] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2021.
- [3] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [4] Yunpeng Wang, Zhijie Li, Zelong Yuan, Wenhui Peng, Tianyuan Liu, and Jianchun Wang. Prediction of turbulent channel flow using fourier neural operator-based machine-learning strategy. *Physical Review Fluids*, 9(8):084604, 2024.
- [5] Fernando Gonzalez, François-Xavier Demoulin, and Simon Bernard. Towards long-term predictions of turbulence using neural operators. *arXiv preprint arXiv:2307.13517*, 2023.
- [6] Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.
- [7] Vivek Oommen, Siavash Khodakarami, Aniruddha Bora, Zhicheng Wang, and George Em Karniadakis. Learning turbulent flows with generative models: Super-resolution, forecasting, and sparse flow reconstruction. *arXiv preprint arXiv:2509.08752*, 2025.
- [8] Shubham Kavane, Kajol Kulkarni, and Harald Koestler. Channelflow-tools: A standardized dataset creation pipeline for 3d obstructed channel flows. *arXiv preprint arXiv:2509.15236*, 2025.
- [9] Timm Krüger, Halim Kusumaatmaja, Alexandr Kuzmin, Orest Shardt, Goncalo Silva, and Erlend Magnus Viggen. *The lattice Boltzmann method*, volume 10. Springer, 2017.
- [10] Martin Bauer, Sebastian Eibl, Christian Godenschwager, Nils Kohl, Michael Kuron, Christoph Rettinger, Florian Schornbaum, Christoph Schwarzmeier, Dominik Thönnies, Harald Köstler, et al. walberla: A block-structured high-performance framework for multiphysics simulations. *Computers & Mathematics with Applications*, 81:478–501, 2021.
- [11] Martin Geier, Martin Schönherr, Andrea Pasquali, and Manfred Krafczyk. The cumulant lattice boltzmann equation in three dimensions: Theory and validation. *Computers & Mathematics with Applications*, 70(4):507–547, 2015.
- [12] Martin Geier, Andrea Pasquali, and Martin Schönherr. Parametrization of the cumulant lattice boltzmann method for fourth order accurate diffusion part i: Derivation and validation. *Journal of Computational Physics*, 348:862–888, 2017.
- [13] Daehyun Choi and Hyungmin Park. Flow around in-line sphere array at moderate reynolds number. *Physics of Fluids*, 30(9), 2018.

- [14] Patrick J Roache. Quantification of uncertainty in computational fluid dynamics. *Annual review of fluid Mechanics*, 29(1):123–160, 1997.
- [15] Frederick W Roos and William W Willmarth. Some experimental results on sphere and disk drag. *AIAA journal*, 9(2):285–291, 1971.
- [16] Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural operators. *arXiv preprint arXiv:2111.13802*, 2021.
- [17] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International conference on machine learning*, pages 5301–5310. PMLR, 2019.
- [18] Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M Benson. U-fno—an enhanced fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163:104180, 2022.
- [19] Oded Ovadia, Vivek Oommen, Adar Kahana, Ahmad Peyvan, Eli Turkel, and George Em Karniadakis. Real-time inference and extrapolation with time-conditioned unet: Applications in hypersonic flows, incompressible flows, and global temperature forecasting. *Computer Methods in Applied Mechanics and Engineering*, 441:117982, 2025.
- [20] Wenhui Peng, Zelong Yuan, Zhijie Li, and Jianchun Wang. Linear attention coupled fourier neural operator for simulation of three-dimensional turbulence. *Physics of Fluids*, 35(1), 2023.
- [21] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.

## A Technical Appendices and Supplementary Material

### A.1 Validation plots and references




Figure 2: Reference measurements from Choi et al. of "wake characteristics behind a single sphere at  $Re = 1000$ : (a) normalized mean streamwise velocity; (b) streamwise turbulence intensity along the centerline; (c) power spectral density (PSD) of the vertical velocity" (reference measurements: ■, rest: literature comparisons from Choi and Park [13])




Figure 3: Normalized mean streamwise velocity  $u/U_\infty$  showing recirculation length  $l_r/D = 1.8$  at  $Re_s = 1000$




Figure 4: RMS of streamwise velocity fluctuations  $u'_{rms}/U_\infty$  along centerline at  $Re_s = 1000$ . Two peaks at  $x/D = 1.5$  and  $2.4$  match reference [13].




Figure 5: Power spectral density of vertical velocity showing dominant frequency at  $St \approx 0.18$ , consistent with experimental vortex shedding frequency for  $Re_s = 1000$  [13]




Figure 6: Mean streamwise velocity field  $\langle u \rangle / U_\infty$  at  $Re_s = 1000$  validation. Contour lines show velocity magnitude. (a) LBM simulation with cumulant collision operator. (b) PIV experimental data from Choi & Park [13].

## A.2 Dataset and training details of the baseline results

The preliminary dataset contains flow around 42 geometries each simulated in the Reynolds number range of [2000, ..., 6000, 8000, 9000, 10000]. 3 rotations of  $90^\circ$  around the x-axis per geometry have been added to augment the existing data. The resolution was scaled down to  $128 \times 64 \times 64$  with min-max-scaling applied for each field. The input consists of a binary field and the Re number,



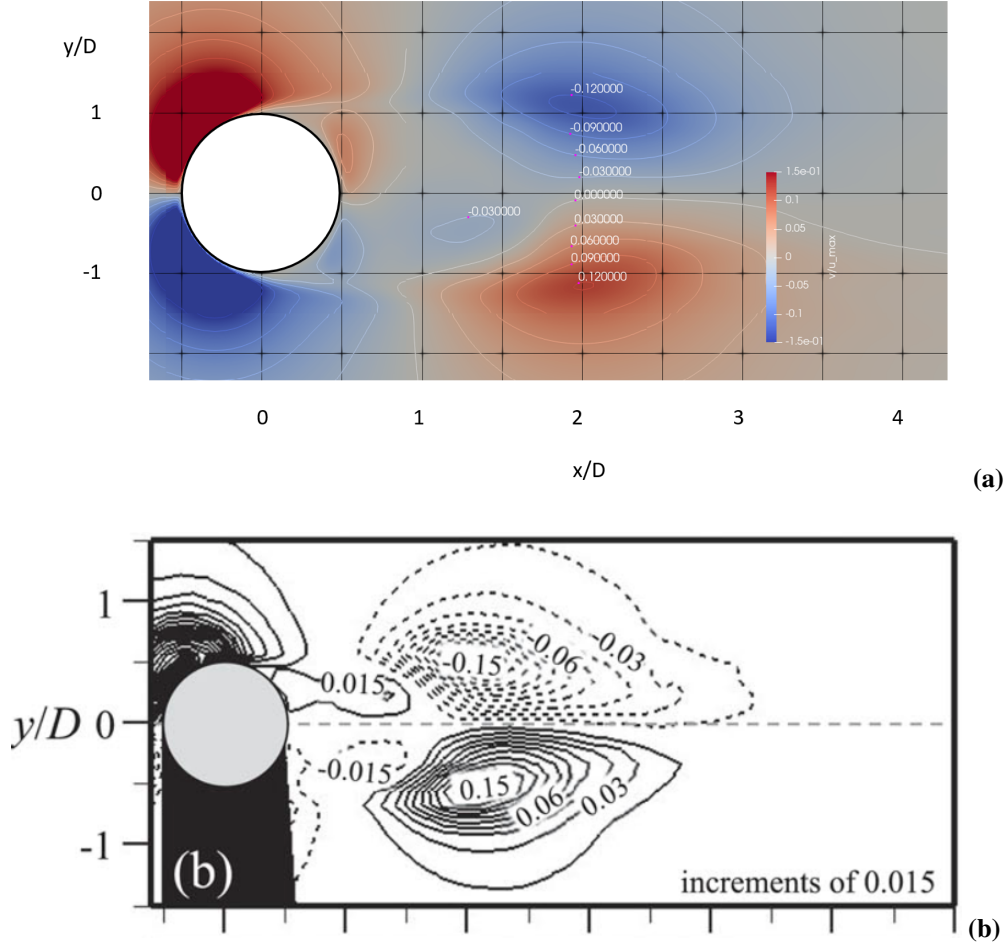


Figure 7: Mean vertical velocity field  $\langle v \rangle / U_\infty$  at  $Re_s = 1000$  validation. Contour lines show velocity magnitude. (a) LBM simulation with cumulant collision operator. (b) PIV experimental data from Choi & Park [13].

the output of the 3-component velocity and the density field. A split 0.8/0.1/0.1 into train, test and validation set was applied. Mean square error with masking, focusing on the important and complex inner region  $([(0.2 \times 0.25 \times 0.25), (1.8 \times 0.75 \times 0.75)] \text{ of } (2 \times 1 \times 1))$ , was chosen as the loss function. The training endured 50 epochs with the Adam optimizer and cosine learning rate scheduling and the metric of samples/s of the training on Nvidia A100 GPUs was reported and the training time for all models was below 2h.

The FNO models represent  $32 \times 16 \times 16$  modes in the frequency space and a depth of 32 channels in the co-domain. The lifting layer was set to a simple feed forward network to expand the input to the 32 channels. The projection layers is chosen as the SIREN from Sitzman et al. [21]. The FFNO follows the structure of the Tran et al. [16] and the F/FNO hybrid maintains the FFNO structure, while replacing the split Fourier transformation with the original 3D, without any loss of practical computational time.

The U-Net follows a standard 3D encoder-decoder architecture with skip connections. The encoder comprises four downsampling stages with channel dimensions [32, 64, 128, 256], each employing two convolutional blocks with batch normalization and ReLU activation, followed by max-pooling. The bottleneck operates at 512 channels, while the decoder mirrors the encoder with upsampling via trilinear interpolation and concatenation with skip connections. A final  $1 \times 1 \times 1$  convolution projects to the output channels with sigmoid activation.




Figure 8: Instantaneous turbulent structures visualized showing characteristic hairpin vortices with identifiable heads and trailing legs at  $Re_s = 1000$ .

Models are evaluated in Table 2 by the mean square error (MSE), the sample of highest mean absolute error (MAE), normalized root MSE (NRMSE) and mean absolute percentage error (MAPE).




Figure 9: Slice plot of the middle plane of x-direction velocity of the baseline UNet model from the validation set