



This is the accepted manuscript made available via CHORUS. The article has been published as:

Generative-machine-learning surrogate model of plasma turbulence

B. Clavier, D. Zarzoso, D. del-Castillo-Negrete, and E. Frénod

Phys. Rev. E **111**, L013202 — Published 22 January 2025

DOI: [10.1103/PhysRevE.111.L013202](https://doi.org/10.1103/PhysRevE.111.L013202)

A generative machine learning surrogate model of plasma turbulence.

B. Clavier* and D. Zarzoso

Aix Marseille Univ, CNRS, Centrale Med, M2P2 UMR 7340, Marseille

D. del-Castillo-Negrete

Oak Ridge National Laboratory, Oak Ridge, TN 37831-8071, United States of America

E. Frénod

Université Bretagne Sud, LMBA UMR 6205, Vannes

Generative artificial intelligence methods are employed for the first time to construct a surrogate model for plasma turbulence that enables long time transport simulations. The proposed GAIT (Generative Artificial Intelligence Turbulence) model is based on the coupling of a convolutional variational auto-encoder, that encodes precomputed turbulence data into a reduced latent space, and a recurrent neural network and decoder that generates new turbulence states 400 times faster than the direct numerical integration. The model is applied to the Hasegawa-Wakatani (HW) plasma turbulence model, that is closely related to the quasigeostrophic model used in geophysical fluid dynamics. Very good agreement is found between the GAIT and the HW models in the spatio-temporal Fourier and Proper Orthogonal Decomposition spectra, and the flow topology characterized by the Okubo-Weiss decomposition. The GAIT model also reproduces Lagrangian transport including the probability distribution function of particle displacements and the effective turbulent diffusivity.

Turbulence is ubiquitous in nature and industrial systems. The atmosphere of planets [1], oceanic currents [2], ionized gases in stars [3], the solar tachocline [4], and the solar wind [5] are some examples of systems where turbulence can be encountered. Without doubt, turbulence represents one of the greatest challenges in numerical modeling due to the vast range of spatiotemporal scales involved and the unpredictable behavior.

In this Letter we focus on turbulence in magnetized plasmas of interest to controlled nuclear fusion and astrophysics. The pressing need to understand and predict the role of turbulence in particles and energy transport in fusion plasmas has motivated the development of a large number of computational tools ranging from fluid to gyro-fluid and gyro-kinetic models. However, due to the spatiotemporal multiscale properties of plasma turbulence these computations are very time consuming and in some cases impractical or unaffordable. To overcome this limitation we propose the use of state-of-the-art artificial intelligence (AI) methods to accelerate turbulence simulations. Although our focus is on plasma physics, the turbulence model of interest contains as a special case a model extensively used in geophysical fluid dynamics. Recently, machine learning techniques have been used to accelerate and validate plasma physics simulations as well as to develop data-driven surrogate models [6–12]. Of particular interest is the potential of AI to create new data from existing data using generative deep-learning techniques. These techniques, which have been used in chemistry to automatically design new molecules [13] and fluid mechanics [14], form the basis of our proposed GAIT (Generative Artificial Intelligence Turbulence) surrogate model.

Once computed, a turbulence field can be used to perform transport studies of passive tracers, e.g., impurities. In the simplest setting, the basic idea is to solve the transport equation in a time interval $(0, t_{max})$ for a large ensemble of particles using the electromagnetic fields computed from the numerical solution of a turbulence model (e.g., gyro-kinetics or gyro-fluid). The main limitation encountered in this type of studies is that the t_{max} of relevance for transport studies is typically much larger than the time range for which the turbulence model can be solved, due to limited computational resources. One way to overcome this limitation is by using surrogate models, like the GAIT model proposed here, that once trained in a time range $(0, t_{train})$ can be used to produce, with negligible computational cost, turbulent states in the transport scale range $(0, t_{max})$ where $t_{train} \ll t_{max}$.

The fast production of long-time turbulent simulations are also useful for capturing the correct asymptotic behavior of tracers exhibiting rare events in the Lagrangian statistics. Examples include anomalously large trapping and/or flights [15], and anomalous diffusion transitions beyond intermediate asymptotic regimes [16].

As a tractable example to illustrate and test the GAIT model we consider the extensively used Hasegawa-Wakatani (HW) model [17]

$$\partial_t n + [\phi, n] = C(\phi - n) - \kappa \partial_y \phi + D_c \nabla^2 n \quad (1)$$

$$\partial_t \Omega + [\phi, \Omega] = C(\phi - n) + \nu \nabla^2 \Omega \quad (2)$$

where ϕ is the electrostatic potential, n is the density, and $\Omega = \nabla_\perp^2 \phi$ the vorticity in the (x, y) plane. C is the adiabaticity coefficient, D_c is the density diffusivity, ν is the viscosity, $[A_1, A_2] = \partial_x A_1 \partial_y A_2 - \partial_y A_1 \partial_x A_2$ is the Poisson bracket, and $\kappa = -\partial_x \log n_0$ is the drive of the instability underlying the turbulence with n_0 the background density. The magnetic field is assumed of the

* benoit.clavier@univ-amu.fr

form $\mathbf{B} = B\mathbf{e}_z$ with B constant. Time is normalized to the cyclotron frequency ω_{c0} , and space to the Larmor radius ρ_s . In applications to controlled nuclear fusion in toroidal magnetic confinement devices, this model is used to describe edge turbulence and recently used in [18] for data-driven turbulence modeling.

Assuming $D_c = \nu = 0$, in the limit $C \rightarrow \infty$, $\phi = n$ and Eqs.(1)-(2) reduce to the Charney-Hasegawa-Mima equation describing Rossby waves in geophysical flows [19] and drift waves in inhomogeneous plasmas [20]. In this reduction, the magnetic field corresponds to the Earth's rotation, the potential to the fluid streamfunction, and the plasma density gradient to the variation of the Coriolis force with latitude. This analogy allows the application of methods and ideas in magnetized plasmas to geophysical flows and viceversa, see for example Ref.[21] and references therein. Since our methodology focuses on the construction of a surrogate model for the potential, and not the density, it is directly applicable to the study of turbulence and transport in geophysical flows [22].




FIG. 1: Schematic representation of the architecture of the proposed GAIT (Generative Artificial Intelligence Turbulence) model. [For further details see supplemental material [23]]

To train the GAIT model, Eqs.(1)-(2) with $C = 1$, $\kappa = 1$, $D_c = \nu = 10^{-3}$ and $k_0 = 0.15$ were solved numerically using a pseudo-spectral method with a $(N_x, N_y) = 512 \times 512$ spatial grid and a 4th order Runge-Kutta scheme for the time advance with $\Delta t = 2 \cdot 10^{-2} \omega_{c0}^{-1}$ to guarantee convergence and stability. The physical parameters are within the range of what is typically used in the literature [24] [25], and the resolution guarantees the numerical stability of the simulation [10].

For the size of the training data we considered $N_s = 8000$ snapshots of the potential obtained by saving the data at a rate $\Delta t_{saving} = \omega_{c0}^{-1}$, i.e. every 50 numerical steps. To eliminate transients, snapshots were saved after the turbulence saturation time $\omega_{c0} t \simeq 200$. Information on the dependence of the GAIT model on N_s can be found in the supplemental material [23].

As illustrated in Fig. 1, the proposed GAIT model con-




FIG. 2: Time snapshots of electrostatic potential, ϕ , in (x, y) obtained from the direct numerical integration of the HW (Hasegawa-Wakatani) model (left) and the proposed GAIT model.

sists of a convolutional variational auto-encoder (CVAE) coupled to a recurrent neural network (RNN), a detailed animation can be found in the supplemental material [23]. CVAEs are probabilistic generative machine learning models composed of an encoder, $\mathcal{E}$, and a decoder, $\mathcal{D}$ [26, 27]. A similar architecture has been also used to compress diagnostic measurements in plasmas [7]. In our problem $\mathcal{E}$ is a convolutional neural network (CNN) consisting of five successive 2D convolution layers (with variable stride, kernel sizes, and filters) that map snapshots of the electrostatic potential, ϕ , represented in an $N_x \times N_y$ grid of the (x, y) plane, to points in a reduced N -dimensional latent space of parameters of a variational Gaussian distribution $\mathcal{N}(\mathbf{z}|\mu, \sigma)$, with $\mathbf{z} \in \mathbb{R}^N$, $\mu \in \mathbb{R}^N$ and $\sigma \in \mathbb{R}^N$. For training purposes, the 512×512 turbulence data was down-sampled to a $N_x \times N_y = 64 \times 64$ resolution. For the dimension of the latent space we used $N = 64$, resulting in a encoder data compression factor of $N_x N_y / N = 64$. The decoder, $\mathcal{D}$, is designed as a mirror image of the encoder with transposed convolution layers that progressively unfold N -dimensional latent vectors into $N_x \times N_y$ dimensional images, such that $\mathcal{D}[\mathcal{E}[\phi]] = \hat{\phi} \approx \phi$. The goal of the training is to find an optimal $\mathcal{E}$ that compresses the data preserving the maximum information and an optimal $\mathcal{D}$ that has the minimum reconstruction error. This is achieved by minimizing the loss function

$$L_{CVAE} = \sum_i^{N_s} w_\phi \left\| \phi_i - \hat{\phi}_i \right\|^2 + w_{\nabla\phi} \left\| \nabla\phi_i - \nabla\hat{\phi}_i \right\|^2 + w_{\mathbb{KL}} \mathbb{KL}(\mathcal{N}(\mathbf{z}_i|\mu_i, \sigma_i), \mathcal{N}(\mathbf{z}_i|\mathbf{0}, \mathbf{1})) \quad (3)$$

where $\{\phi_i\}_{i=1}^{N_s}$ and $\{\hat{\phi}_i\}_{i=1}^{N_s}$ are the N_s training and the N_s reconstructed snapshots respectively, $\mathbf{z}_i = \mathcal{E}(\phi_i)$. The first two terms in Eq. (3) are the reconstruction loss of ϕ and $\nabla\phi$ in the L2 norm, weighted by the hyperparameters $w_\phi \geq 0$ and $w_{\nabla\phi} \geq 0$. The third term in Eq.(3), with weight $w_{\mathbb{KL}} \geq 0$, uses the Kullback-Leibler divergence $\mathbb{KL}$ [28] to quantify the difference between the normal distributions of the mapped snapshots in the latent space with mean μ_i and variance σ_i and a normal distribution with zero mean and unit variance. Minimizing this difference forces the encoder to map the different images into a compact set with small variability in the

latent space, a property that is critical for the success of the generation of turbulence for transport studies.

Since L_{CVAE} is defined up to a scale factor, there are only two independent parameter ratios, which for $w_{\text{KL}} \neq 0$ must satisfy $w_\phi/w_{\text{KL}} > 1$ (if $w_\phi \neq 0$) and $w_{\nabla\phi}/w_{\text{KL}} > 1$ (if $w_{\nabla\phi} \neq 0$), so that the KL loss does not dominate over the reconstruction loss. Increasing the value of a hyper-parameter will yield a higher accuracy in the observable(s) associated to it. For example, increasing $w_{\nabla\phi}$ will increase the accuracy of observables like the electric field $\mathbf{E} = -\nabla\phi$ and the $\mathbf{E} \times \mathbf{B}$ velocity. w_{KL} controls the regularity of the latent space. Decreasing this parameter too much leads to noisy Lagrangian dynamics that might artificially increase the diffusivity.

The training of the CVAE was implemented using an Adam optimizer with a learning rate $l_r = 10^{-3}$ and 5000 epochs. Since the numerical simulations are performed in a double periodic domain, the convolution algorithms of TensorFlow were modified so the encoder captures the periodicity of the input data and the decoder generates periodic images from points in the latent space.

Following the CVAE training, the sequence $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{N_s}\}$, corresponding to the evolution in time of the turbulence snapshots $\{\phi_1(x, y), \phi_2(x, y), \dots, \phi_{N_s}(x, y)\}$, is split into training and testing sequences containing 85% and 15% of the data respectively. The training sequence is then used to train a multi-layer RNN (Recurrent Neural Network) with two hidden recurrent layers, using an Adam optimizer with learning rate $l_r = 10^{-3}$, 500 epochs, and loss function $L_{\text{RNN}} = \sum_i \|(\mathbf{z}_{i+1}, \dots, \mathbf{z}_{i+l+1}) - \text{RNN}(\mathbf{z}_i, \dots, \mathbf{z}_{i+l})\|^2$ with $l = 50$ to correctly reproduce the time auto-correlation. Once trained, the RNN is used to advance the latent space sequence from $\{\mathbf{z}_i, \dots, \mathbf{z}_{i+l}\}$ to $\{\mathbf{z}_{i+1}, \dots, \mathbf{z}_{i+l+1}\}$.

The turbulence AI generation process is completed by using the decoder to map the new points in the latent space into new snapshots of turbulence, $\{\tilde{\phi}_{N_s+1}(x, y), \dots, \tilde{\phi}_{N_s+N_g}(x, y)\}$ with $\tilde{\phi}_j(x, y) = \mathcal{D}[\tilde{\mathbf{z}}_j]$. As shown in Fig.2, the images of the AI generated turbulent states are practically indistinguishable from those generated from the direct numerical simulation of the Hasegawa-Wakatani turbulence model in Eq. (1)-(2). However, from the physics perspective it is critical to go beyond the observed qualitative agreement of the images and perform systematic tests using quantitative metrics. The rest of the letter is devoted to this goal.

We present numerical results for two realizations of the GAIT model. *Case 1*, with $(w_\phi, w_{\nabla\phi}, w_{\text{KL}}) = (1, 10, 0.01)$, optimizes the reconstruction of ϕ and $\nabla\phi$, with low latent space regularity. *Case 2*, with $(w_\phi, w_{\nabla\phi}, w_{\text{KL}}) = (0, 10, 0.1)$, optimizes only $\nabla\phi$, with high latent space regularity.

The first test compares the Fourier power spectra in space of the electric field turbulent fluctuations, $|\widehat{\delta E}(k)|$, where $\widehat{\delta E}(k, t)$ is the Fourier transform in space of $\delta E(x, y, t) = E(x, y, t) - \langle E \rangle$, $\langle f \rangle$ denotes spatial aver-

age, $\bar{f}$ denotes time average, E is the magnitude of the electric field, $E = |-\nabla\phi|$, and $k = \sqrt{k_x^2 + k_y^2}$. The second test compares the Fourier power spectra in time of the electric field turbulent fluctuations, $\langle |\widehat{\delta E}(\omega)| \rangle$, where $\widehat{\delta E}(x, \omega)$ is the Fourier transform in time of $\delta E(x, y, t)$. Figure 3 compares these two metrics in the HW and the GAIT models.




FIG. 3: Comparison between the HW and GAIT models electric field turbulent fluctuations power spectra in space (left) and time (right).




FIG. 4: Comparison between the HW and GAIT models proper orthogonal decomposition spectrum of potential, ϕ , (left) and time auto-correlation function (right).

A distinctive feature of 2D turbulence is the spontaneous formation and persistence of coherent structures, e.g. vortices and zonal flows [29]. Understanding these structures is important because the trapping effect of vortices and the long displacements caused by zonal flows can have a critical role in transport, see for example [21] and references therein. Accordingly, an important metric in the evaluation of the proposed AI model is how well the model captures the coherent structures of the flow. To quantify this, we use Proper Orthogonal Decomposition (POD) methods which have been extensively used in fluids [30] and plasmas [31]. Contrary to Fourier spectral analysis, POD is based on data-driven empirical modes providing an optimal representation of the turbulence state in the energy norm. As discussed in [32], the POD mode decomposition and spectra can be used to characterize the spatiotemporal coherence of HW plasma turbulence. Following this idea, the left panel of Fig. 4 shows the POD singular values w_l as a function of the rank, l , where $\phi(x, y, t) = \sum_{l=1}^N w_l u^{(l)}(x, y) v^{(l)}(t)$ is the




FIG. 5: Comparison between the HW and GAIT models flow topology. Top panel: spatial distribution of Okubo-Weiss field at a fixed time in HW (left), GAIT case 1 (center), and GAIT case 2 (right). Bottom panel: probability distribution of Q values in space and time.

singular value decomposition of the electrostatic potential. The close agreement, especially for *case 1*, indicates that the GAIT model reproduces well the coherent structures of the HW model. Another metric of interest is the Eulerian correlation time of the electrostatic potential fluctuations, $\mathcal{C}(\tau) = \langle \phi(x, t) \phi(x, t + \tau) \rangle$. The right panel in Fig. 4 compares auto-correlation functions of the HW and the GAIT models, showing a very good reproduction of the turbulent dynamics.

To further compare the GAIT and the HW turbulence models, we consider the Okubo-Weiss (OW) parameter, used in [33] to characterize the flow topology in HW plasma turbulence. The OW parameter is defined as $Q = s^2 - \Omega^2$, where $\Omega = \partial_x V_y - \partial_y V_x$ is the vorticity, $s^2 = s_1^2 + s_2^2$ is the deformation with $s_1 = \partial_x V_x - \partial_y V_y$ and $s_2 = \partial_x V_y + \partial_y V_x$ [34]. Here $\mathbf{V} = -\nabla\phi \times \mathbf{e}_z = (V_x, V_y) = (-\partial_y\phi, \partial_x\phi)$. Given the threshold, $Q_0 = \sqrt{\langle Q^2 \rangle}$, the turbulent flow can then be partitioned into three topologically distinctive regions: strongly elliptic (vorticity dominated) $Q \leq -Q_0$, strongly hyperbolic (deformation dominated) $Q \geq Q_0$, and intermediate $-Q_0 < Q < Q_0$. From a transport perspective, strongly elliptic regions tend to reduce transport due to particle trapping and strongly hyperbolic regions tend to enhance mixing due to particle dispersion. The agreement between the GAIT and the HW models in the flow topology of the turbulent field is shown in Fig.5. The top panels compare the Okubo-Weiss three levels ($Q < -Q_0$, $-Q_0 < Q < Q_0$, $Q > Q_0$) decomposition of the turbulence field at a given time. The bottom panel compares the probability distribution of Q values in space and time $\{Q_{i,j}(t_k)\}$, where $i = 1, \dots, 64$, $j = 1, \dots, 64$ and $k = 1, \dots, N_s$.

For all the Eulerian metrics shown in Figs.3-5, the agreement between the GAIT model and the HW turbulence model is remarkably good for Case 1. This is to be expected since this case optimizes the reconstruction

of ϕ and $\nabla\phi$. Case 2, optimizes only $\nabla\phi$, and despite a high latent space regularity, yields less accurate results.

As a final test, we study the turbulence-induced transport by solving the guiding center equations $\frac{d\mathbf{x}}{dt} = -\nabla\phi \times \mathbf{e}_z$ for 10^4 tracers initially distributed uniformly in the square domain $\{(x, y) | 0 < x < 2\pi/k_0, 0 < y < 2\pi/k_0\}$. The orbits are integrated using a Runge-Kutta method with a timestep $\Delta t = 0.2$. The data in the electrostatic potential snapshots was interpolated using linear interpolation in time and bilinear interpolation in space. When computing the Lagrangian statistics the particle orbits are considered in the extended $\mathbb{R}^2$, i.e. the values of $(x_i(t), y_i(t))$ were not restricted to the $[0, 2\pi/k_0] \times [0, 2\pi/k_0]$ double periodic domain. The main quantities of interest are the PDF (Probability Distribution Function) of the displacement, $P(R)$, and the evolution of the mean squared displacement $\langle R^2(t) \rangle$ where $R = \|\delta\vec{r} - \langle \delta\vec{r} \rangle\|$, $\delta\vec{r} = \vec{r}(t) - \vec{r}(0)$ is the displacement vector and $\langle \cdot \rangle$ is an ensemble average. Figure 6 shows that both the HW and the GAIT models exhibit diffusive scaling $\langle R^2(t) \rangle \sim Dt$, where D is the diffusivity.

Case 2, that focuses on the optimization of $\nabla\phi$ (i.e., the $\mathbf{E} \times \mathbf{B}$ velocity) and high regularity of the latent space, yields a better description of the Lagrangian transport with a diffusivity $D_{GAIT,2} = 1.19$ very close to the ground truth diffusivity $D_{HW} = 1.27$, along with a very good agreement on the PDFs of R . The smaller value of w_{KL} in Case 1, reduces the latent space regularity and leads to a slightly higher value of the diffusivity, $D_{GAIT,1} = 1.47$. Most importantly, as proof of the fidelity of the turbulence generation process, the GAIT model preserves the transport properties well beyond ($t = 10^5$) the numerical integration time of the HW model ($t = 8 \times 10^3$).




FIG. 6: Comparison of radial displacement statistics of tracers in the HW and GAIT models. $\langle R^2 \rangle$ as function of t in log-log scale (left). Probability distribution function of displacements at different times (right).

The direct numerical integration of the HW model for 10,000 snapshots took 50 hours on a V100 GPU node of the French Jean-Zay supercomputer. On the same node, the training of the CVAE took 1 hour, and the training of the RNN five minutes. The generation of 10,000 snapshots of turbulent states with the GAIT model took 7.5 min, about 50 hrs/7.5 min ~ 400 times faster than the direct numerical integration of the HW model.

The HW model requires the integration of the potential, ϕ , and density, n , coupled equations. An advantage of the GAIT model is that it can provide a surrogate model of ϕ only, for applications not requiring n , like the one discussed here. If needed, a direct generalization of the GAIT model can be used to construct surrogates for n and ϕ .

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily

reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. This work has received financial support from the AIM4EP project (ANR-21-CE30-0018), funded by the French National Research Agency (ANR), and from the Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the US Department of Energy under Contract No. DE-AC05-00OR22725. All the simulations and training of neural networks reported here were performed on HPC resources of IDRIS under the allocations 2021-A0100512455, 2022-AD010512455R1 and 2023-A0140514165 made by GENCI.

-
- [1] J. C. Wyngaard, Atmospheric turbulence, Annual Review of Fluid Mechanics **24**, 205 (1992).
 - [2] A. E. Gargett, Ocean turbulence, Annual Review of Fluid Mechanics **21**, 419 (1989).
 - [3] V. Canuto and J. Christensen-Dalsgaard, Turbulence in astrophysics: Stars, Annual review of fluid mechanics **30**, 167 (1998).
 - [4] E. Spiegel and J.-P. Zahn, The solar tachocline, Astronomy and Astrophysics (ISSN 0004-6361), vol. 265, no. 1, p. 106-114. **265**, 106 (1992).
 - [5] R. Bruno and V. Carbone, The solar wind as a turbulence laboratory, Living Reviews in Solar Physics **10**, 1 (2013).
 - [6] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach, Physical review letters **120**, 024102 (2018).
 - [7] B. Zhu, M. Zhao, H. Bhatia, X.-q. Xu, P.-T. Bremer, W. Meyer, N. Li, and T. Rognlien, Data-driven model for divertor plasma detachment prediction, Journal of Plasma Physics **88**, 895880504 (2022).
 - [8] P. Rodriguez-Fernandez, A. White, A. Creely, M. Greenwald, N. Howard, F. Sciortino, and J. Wright, Vitals: a surrogate-based optimization framework for the accelerated validation of plasma transport codes, Fusion Science and Technology **74**, 65 (2018).
 - [9] R. M. Churchill, M. D. Boyer, and S. C. Cowley, Accelerating fusion energy with ai, in *Artificial Intelligence for Science: A Deep Learning Revolution* (World Scientific, 2023) pp. 271–284.
 - [10] R. Greif, F. Jenko, and N. Thuerey, Physics-preserving ai-accelerated simulations of plasma turbulence, arXiv preprint arXiv:2309.16400 (2023).
 - [11] M. Yang, P. Wang, D. del Castillo-Negrete, Y. Cao, and G. Zhang, A pseudo-reversible normalizing flow for stochastic dynamical systems with various initial distributions, arXiv preprint arXiv:2306.05580 (2023).
 - [12] C. J. McDevitt, A physics-informed deep learning model of the hot tail runaway electron seed, Physics of Plasmas **30** (2023).
 - [13] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik, Automatic chemical design using a data-driven continuous representation of molecules, ACS central science **4**, 268 (2018).
 - [14] Y. Wang, A. Solera-Rico, C. S. Vila, and R. Vinuesa, Towards optimal β -variational autoencoders combined with transformers for reduced-order modelling of turbulent flows, International Journal of Heat and Fluid Flow **105**, 109254 (2024).
 - [15] D. del Castillo-Negrete, Asymmetric transport and non-gaussian statistics of passive scalars in vortices in shear, Physics of Fluids **10**, 576 (1998).
 - [16] Á. Cartea and D. del Castillo-Negrete, Fluid limit of the continuous-time random walk with general lévy jump distribution functions, Physical Review E—Statistical, Nonlinear, and Soft Matter Physics **76**, 041105 (2007).
 - [17] A. Hasegawa and M. Wakatani, Plasma edge turbulence, Physical Review Letters **50**, 682 (1983).
 - [18] C. Gahr, I.-G. Farcas, and F. Jenko, Learning physics-based reduced models from data for the hasegawa-wakatani equations, arXiv preprint arXiv:2401.05972 (2024).
 - [19] J. Pedlosky, *Geophysical fluid dynamics*, 2nd ed. (Springer New York, NY, 1987).
 - [20] A. Hasegawa and K. Mima, Stationary spectrum of strong turbulence in magnetized nonuniform plasma, Physical Review Letters **39**, 205 (1977).
 - [21] D. del Castillo-Negrete, Chaotic transport in zonal flows in analogous geophysical and plasma systems, Physics of Plasmas **7**, 1702 (2000).
 - [22] A. Bracco, J. von Hardenberg, A. Provenzale, J. B. Weiss, and J. C. McWilliams, Dispersion and mixing in quasigeostrophic turbulence, Physical review letters **92**, 084501 (2004).
 - [23] See supplemental Material at [URL] for (1) an animated video showcasing how the GAIT model works. The animation corresponds to the "case 1" set of parameters. The latent space representation is the result of a 2D principal component analysis. (2) A document showing how the model performs depending on the training dataset size.
 - [24] J. Anderson, E.-j. Kim, B. Hnat, and T. Rafiq, Elucidating plasma dynamics in hasegawa-wakatani turbulence by information geometry, Physics of Plasmas **27** (2020).
 - [25] A. Pushkarev, *Self-organization of isotopic and drift-wave turbulence*, Ph.D. thesis, Citeseer (2013).
 - [26] D. P. Kingma and M. Welling, Auto-encoding variational bayes, arXiv preprint arXiv:1312.6114 (2013).
 - [27] M. Cheng, F. Fang, C. Pain, and I. Navon, An advanced

- hybrid deep adversarial autoencoder for parameterized nonlinear fluid flow modelling, *Computer Methods in Applied Mechanics and Engineering* **372**, 113375 (2020).
- [28] S. Kullback and R. A. Leibler, On information and sufficiency, *The annals of mathematical statistics* **22**, 79 (1951).
- [29] J. Sommeria, Experimental study of the two-dimensional inverse energy cascade in a square box, *Journal of fluid mechanics* **170**, 139 (1986).
- [30] G. Berkooz, P. Holmes, and J. L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annual review of fluid mechanics* **25**, 539 (1993).
- [31] P. Beyer, S. Benkadda, and X. Garbet, Proper orthogonal decomposition and galerkin projection for a three-dimensional plasma dynamical system, *Physical Review E* **61**, 813 (2000).
- [32] S. Futatani, S. Benkadda, and D. del Castillo-Negrete, Spatiotemporal multiscaling analysis of impurity transport in plasma turbulence using proper orthogonal decomposition, *Physics of Plasmas* **16**, 042506 (2009), https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/1.3095865/15930324/042506_1_online.pdf.
- [33] B. Kadoch, D. del Castillo-Negrete, W. J. T. Bos, and K. Schneider, Lagrangian conditional statistics and flow topology in edge plasma turbulence, *Physics of Plasmas* **29**, 102301 (2022), https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/5.0098501/16570139/102301_1_online.pdf.
- [34] D. Elhmaïdi, A. Provenzale, and A. Babiano, Elementary topology of two-dimensional turbulence from a lagrangian viewpoint and single-particle dispersion, *Journal of Fluid Mechanics* **257**, 533–558 (1993).