

Data-driven insights into jet turbulence: Explainable AI approaches

Original

Data-driven insights into jet turbulence: Explainable AI approaches / Amico, E.; Matteucci, L.; Cafiero, G.. - In: PHYSICS OF FLUIDS. - ISSN 1070-6631. - 37:5(2025). [10.1063/5.0268702]

Availability:

This version is available at: 11583/2999933 since: 2025-05-07T10:05:31Z

Publisher:

AIP Publishing

Published

DOI:10.1063/5.0268702

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

RESEARCH ARTICLE | MAY 06 2025

Data-driven insights into jet turbulence: Explainable AI approaches

E. Amico  ; L. Matteucci  ; G. Caferio 



Physics of Fluids 37, 055108 (2025)

<https://doi.org/10.1063/5.0268702>



Articles You May Be Interested In

Advancing transparency of model training: The interplay of SHAP and convolutional neural networks based on Iris dataset

AIP Conf. Proc. (December 2024)


A novel study on explainable artificial intelligence powered churn prediction in businesses

AIP Conf. Proc. (October 2024)

Explaining reaction coordinates of alanine dipeptide isomerization obtained from deep neural networks using Explainable Artificial Intelligence (XAI)

J. Chem. Phys. (April 2022)

21 October 2025 09:21:47



Physics of Fluids
Special Topics
Open for Submissions

[Learn More](#)

Data-driven insights into jet turbulence: Explainable AI approaches

Cite as: Phys. Fluids **37**, 055108 (2025); doi: [10.1063/5.0268702](https://doi.org/10.1063/5.0268702)

Submitted: 3 March 2025 · Accepted: 15 April 2025 ·

Published Online: 6 May 2025



View Online



Export Citation



CrossMark

E. Amico,^{a)}  L. Matteucci,  and G. Cafiero 

AFFILIATIONS

Department of Mechanical and Aerospace Engineering, Politecnico di Torino, Torino, Italy

^{a)} Author to whom correspondence should be addressed: enrico.amico@polito.it

ABSTRACT

In this study, eXplainable Artificial Intelligence (XAI) methods are applied to analyze flow fields obtained through particle image velocimetry measurements of an axisymmetric turbulent jet. A convolutional neural network (U-Net) was trained to predict velocity fields at subsequent time steps. Three XAI methods—SHapley Additive explanations (SHAP), Gradient-SHAP, and Gradient-weighted Class Activation Mapping (Grad-CAM)—were employed to identify the flow field regions relevant for prediction. SHAP requires predefined segmentation of the flow field into relevant regions, while Gradient-SHAP and Grad-CAM avoid this bias by generating gradient-based heatmaps. The results show that the most relevant structures do not necessarily coincide with the regions of maximum vorticity, but rather with those that play a critical role in energy transfer and jet dynamics. Additionally, structures with high turbulent dissipation values are identified as the most significant. Gradient-SHAP and Grad-CAM methods reveal a uniform spatial distribution of relevant regions, emphasizing the contribution of nearly circular structures to turbulent mixing. This study advances the understanding of turbulent dynamics through XAI tools, providing an innovative approach to correlate machine learning models with physical phenomena.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0268702>

I. INTRODUCTION

In recent years, there has been a steady increase in computational power, along with the development of Graphics Processing Unit (GPUs). These advancements have enabled much faster solutions to numerical problems. In recent literature, it is quite common to find studies where direct numerical simulations (DNS) are performed using GPUs, resulting in significantly shorter computation times compared to traditional Central Processing Unit (CPU)-based methods.

At the same time, the past decade has seen the rise of machine learning (ML) and the development of data-driven techniques, also benefiting from the progress of increasingly powerful and efficient GPUs. In various studies, data obtained from DNS have been used to train neural networks (NN) for a range of applications: prediction, reconstruction, data augmentation, simulation acceleration, and data estimation.

While machine learning techniques offer excellent tools for handling large datasets, it is important to consider that they often function as black-box models. As a result, understanding the internal mechanisms that lead to the results obtained after training can be highly complex, and sometimes impossible. One of the key strengths of neural networks lies in their ability to model highly nonlinear phenomena, which is a central focus of ongoing research in the field.

Fluid dynamics is governed by a complex and highly nonlinear system of equations known as the Navier–Stokes equations (NS). Due to this nonlinearity, neural networks represent a valid alternative for modeling the NS. However, it is reasonable to assume that during the training phase with fluid dynamics data—whether numerical (DNS) or experimental—a neural network learns to capture these nonlinearities and, consequently, the underlying physical mechanisms of the phenomenon.

The main challenge is that turbulence is a multi-scale phenomenon in both time and space. Energy is primarily transferred from the largest scales to the smallest, where it dissipates, although there is also a reverse energy transfer. In any engineering flow, the difference in magnitude between these scales can be vast, and this scale separation increases with higher Reynolds numbers.

Jets are characterized by a region of high shear at the jet boundary, where the jet fluid mixes with the surrounding fluid or in the self-similar region. This shear region is unstable and leads to the development of turbulence (coherent structures, ring vortices^{1,2}). Turbulence in a jet is marked by the presence of vortices of various sizes and time scales. These vortices interact with each other, resulting in a complex and chaotic flow pattern. Turbulence in jets is responsible for increased mixing with the surrounding fluid, which is crucial in many engineering applications, such as combustion and cooling.^{1,2}

Despite the chaotic nature of turbulence, jets exhibit coherent turbulent structures. These structures are organized flow patterns that persist for a certain period of time and play a significant role in jet dynamics. One of the most important turbulent structures in jets is the vortex ring.

The process of formation and dispersion of vortex rings is one of the main sources of turbulence in jets.¹

In axisymmetric jets, entrainment occurs when the jet's radial spread and shear layers between the jet and ambient fluid initiate turbulence, drawing in ambient fluid. Critical to this turbulence are small, high-enstrophy structures that persist beyond the flow's characteristic timescales and primarily drive energy dissipation.³

Enhanced computational power has allowed detailed examination of these structures,⁴ which generally align the vorticity vector with the main strain direction and have a typical radius five times the Kolmogorov length scale. The Kelvin–Helmholtz (KH) mechanism explains the evolution of coherent structures in shear layers and jets, representing a spatial instability that grows as it moves downstream in jets. Initially, the KH wave grows but later decays due to flow spreading, helping engineers predict and control flow in many applications.^{3,5}

Streamwise vortices in jets generate streaky structures through the lift-up mechanism, redistributing kinetic energy and promoting mixing. These streaks are elongated regions of alternating high- and low-momentum fluid aligned with the main flow. They extend downstream and interact with smaller-scale eddies, influencing jet turbulence. Spectral analysis shows that streaky structures are prevalent in the flow's low-frequency range, highlighting their large scale and slow evolution relative to smaller turbulence, making them essential to understanding jet mixing and spreading.⁶

Despite extensive research into the contributions of various structures to the dynamics of turbulent flows, their precise roles are not yet fully understood. The development of new data-driven techniques and methodologies makes it possible to extend current knowledge, identifying potential relationships within the data that, in fact, enable a better understanding of the mechanisms underlying turbulence.

Lellep *et al.*⁷ used the SHapley Additive explanations (SHAP) algorithm to identify the flow features relevant to predicting relaminarization, in a wall bounded flow, by analyzing the output of a machine learning (ML) model trained for this purpose. SHAP is an algorithm derived from the game theory, and it quantifies the contribution of each flow feature to the model's final prediction.

In their work, Lellep *et al.*⁷ showed that SHAP is an effective tool for identifying flow features influencing relaminarization and for understanding the underlying dynamic mechanisms in wall-bounded turbulence.

In Cremades *et al.*,⁸ the authors showed that they could predict velocity fields over time using data from turbulent channel flow simulations leveraging a U-net. They then employed the SHAP algorithm, to evaluate the significance of various flow structures in making these predictions. Contrary to previous assumptions, they find that the most important structures for prediction are not necessarily those with the highest contribution to Reynolds shear stress.

Cremades *et al.*⁹ examined four SHAP variants and their applications to turbulence modeling, flow structures dynamics, and heat transfer, emphasizing the crucial role that can be played by explainable AI in fluid mechanics.

Cremades *et al.*¹⁰ presented a methodology using deep learning and explainable AI (XAI) to analyze turbulent flows. A U-net model

and the gradient-SHAP algorithm were employed to identify the most significant regions in a turbulent flow, comparing them with classical coherent structures like Q events, streaks, and vortices.

In the context of convolutional neural networks (CNN), explainability and interpretability Gradient-weighted Class Activation Mapping (Grad-CAM)¹¹ has emerged as a widely used technique to visualize the regions of an image that most influence the model's predictions. This method leverages the gradients flowing through the network to generate a weighted activation map, highlighting the areas most relevant to classification or regression tasks. Due to its simplicity and ability to provide intuitive visual explanations, Grad-CAM has become an essential tool for enhancing the interpretability of neural networks across various applications.^{12–14}

The aim of this work is to apply different XAI approaches to flow fields obtained from particle image velocimetry (PIV) measurements¹⁵ for an axisymmetric turbulent jet, comparing methods already used in the literature (SHAP and Gradient-SHAP) for other types of flow fields^{7,8} with GRAD-CAM, which has not yet been used in this context. This comparison of techniques aims to identify which regions of the flow field are relevant for the development of an axisymmetric turbulent jet.

II. METHODOLOGY

Convolutional neural networks (CNNs) have been widely employed for prediction tasks across various scientific and engineering domains, demonstrating their ability to extract meaningful patterns from complex data.^{16,17} In the context of fluid dynamics, CNNs offer a powerful tool for forecasting velocity fields, learning from past states to predict future flow evolution.¹⁸

The application of AI methods to the modeling of turbulent flows has gained momentum, offering new perspectives for analyzing high-Reynolds-number axisymmetric jets.^{19,20} In particular, deep neural networks have proven promising for capturing the nonlinear interactions among vortices and shear layers, significantly reducing computational costs compared to traditional direct numeric simulation (DNS) or large eddy simulation (LES).²¹ Duraisamy *et al.*¹⁹ have demonstrated that AI can accurately predict the propagation of turbulent structures and identify key mechanisms of mixing and turbulent kinetic energy.

In this work, a CNN is used for a prediction task: let U represent the velocity fields on a plane at a given time t_i , where a neural network is trained to predict the velocity field at the subsequent time t_{i+1} . The network used is a U-Net,¹⁶ a type of architecture that effectively extracts features at multiple spatial scales by employing feature maps of varying dimensions.²²

During training, the neural network minimizes the loss function. If the loss function is sufficiently low at the end of training, the network is capable of accurately predicting the velocity field at the next instant, given the input at the previous instant. This suggests that the network, upon reaching a sufficiently low loss function value, has learned some of the underlying physics of the problem, making it interpretable as a model of the phenomenon itself.²³

A. SHAP

Following^{7,8} the SHAP algorithm was used to interpret and analyze the importance of variables within the adopted predictive model. SHAP, based on Shapley values from cooperative game theory, assigns

a unique contribution to each feature in the model for the prediction, enabling an accurate estimate of its influence. If g is a black box function and z is an input for g , with $z \in \mathbb{R}^M$ where M as number of features, the Shapley values are defined as Φ_m ²⁴ following equation:

$$g(z) = \phi_0 + \sum_{m=1}^M \phi_m \cdot z_m. \quad (1)$$

$z_m \in \{0, 1\}$ indicates whether feature m is “active” (1) or “inactive” (0).

When dealing with interpretability tasks, it is essential to introduce the concept of a coalition: a coalition represents a subset of features considered in constructing the importance of each variable with respect to the model’s prediction. Coalitions are formed by including or excluding different features to simulate how the presence or absence of each one influences the outcome. To calculate the Shapley value, we consider all possible features’ permutations. A permutation identifies a coalition, and it determines each feature’s marginal contribution to every possible coalition they can join. The Shapley value for each feature is then the average of these marginal contributions over all permutations.

The Shapley values can be computed through the following equation:

$$\phi(i, v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} [v(S \cup \{i\}) - v(S)], \quad (2)$$

where $\phi(i, v)$ is the SHAP value due to the i -th feature, N is the set of coalition participants, S is a subset of features excluding i , $v(S \cup \{i\}) - v(S)$ is the marginal contribution term for the i -th feature, and $\frac{|S|!(|N| - |S| - 1)!}{|N|!}$ is the probability that the i -th feature will be part of the coalition.²⁴

The exact calculation of Shapley values is not immediate; therefore, the typical approach is to approximate them with the kernel-SHAP method,^{24,25} which integrates Shapley values with the Local Interpretable Model-agnostic Explanations (LIME) method.²⁶ This can be expressed as follows:

$$\mathcal{L}(f, g, \pi_x) = \sum_{q'} [f(h_x(q')) - g(q')]^2 \pi_x(q'), \quad (3)$$

$$\pi_x(q') = \frac{|Q| - 1}{\binom{|Q|}{|q'|} |q'| (|Q| - |q'|)}, \quad (4)$$

where $|q'|$ is the number of structures, h_x represents a mask function that transforms the binary space of q' into the space of the input of the model. The LIME equation has been solved by means of a linear regression, with the resulting error being $(f - g)^2 \mathcal{O}(10^{-7})$.²⁶

The local accuracy property of Shapley values guarantees that the total of the Shapley values for a given sample z equals the difference between the model’s output for that sample, $g(z)$, and the average prediction of the model, $\mathbb{E}[g(\tilde{z})]$. Thus, the total of all Shapley values is equivalent to the difference between the model output and the average prediction of the model.

While the application of SHAP to regression or classification problems is straightforward,²⁷ in cases like the one analyzed in this work, flow prediction, it becomes more complex because the input and/or output lack predefined features. For example, in Cremades

et al.,⁸ the authors defined the structures present in the field leveraging the approach proposed by Lozano-Durán *et al.*,²⁸ and using each of these structures as features for calculating their importance via SHAP.

In this work, we focus on an axisymmetric jet flow. We have chosen to use the vorticity normal to the symmetry plane, ω_z , as a discriminant to determine input features, implementing a thresholding method.

We used the jet exit diameter D and V_{exit} to normalize ω_z ,

$$\omega_z^* = \frac{\omega_z \cdot D}{V_{exit}}. \quad (5)$$

Specifically, a threshold ω_{thr} is set and applied to the flow field, yielding two distinct masks: for $\omega_z^* > \omega_{thr}^*$ and $\omega_z^* < -\omega_{thr}^*$, respectively. For each mask, the connected regions are labeled, and these regions serve as input features, as sketched in Fig. 1.

B. Gradient-SHAP

Gradient-SHAP is a method for attributing feature importances to the predictions of a machine learning model.^{29,30} It combines the ideas of SHAP with the gradients of a model outputs with respect to the input features. This method aims to explain the outputs of complex models, such as deep neural networks, by providing robust explanations even in highly nonlinear settings.

The choice of \mathbf{x}' plays a critical role in defining the reference from which contributions are measured, aligning with the broader SHAP methodology. Given a model f and an input vector \mathbf{x} , Gradient SHAP estimates the contribution of each feature x_i to the model output $f(\mathbf{x})$ by considering a baseline input \mathbf{x}' ; in the present case, we consider the mean flow velocity. The method interpolates between \mathbf{x} and \mathbf{x}' by generating random samples \mathbf{z} along a linear path, defined as

$$\mathbf{z}_\alpha = \alpha \mathbf{x} + (1 - \alpha) \mathbf{x}', \quad \alpha \in [0, 1]. \quad (6)$$

The contributions are then estimated by sampling multiple values of α and computing the gradients of f with respect to \mathbf{z}_α . The SHAP value for the i -th feature is calculated as

$$\phi_i = \mathbb{E}_{\alpha \sim \text{Uniform}[0,1]} [(\mathbf{x} - \mathbf{x}')_i \cdot \nabla_{\mathbf{z}_\alpha} f(\mathbf{z}_\alpha)], \quad (7)$$

where $\nabla_{\mathbf{z}_\alpha} f(\mathbf{z}_\alpha)$ denotes the gradient of the model output with respect to \mathbf{z}_α .

C. GradCAM

The Grad-CAM algorithm, introduced by Selvaraju *et al.*,¹¹ represents a significant advancement over previous approaches regarding versatility and accuracy in interpretability tasks, but it is applicable exclusively to neural networks. The method can be introduced at a high level in a very simple way: the process begins by taking an input image and constructing a model that is truncated at the layer of the neural network for which we want to create a Grad-CAM heat-map following the algorithm shown below. It is important to point out that the layer against which to truncate the network represents a parameter of the algorithm. For prediction purposes, fully connected layers are added. Subsequently, the input traverses the model, capturing the output and the loss. Then, the gradient of the output of the model layer with respect to the model loss is calculated. Next, we identify the portions of the gradient contributing to the prediction, refining, resizing,

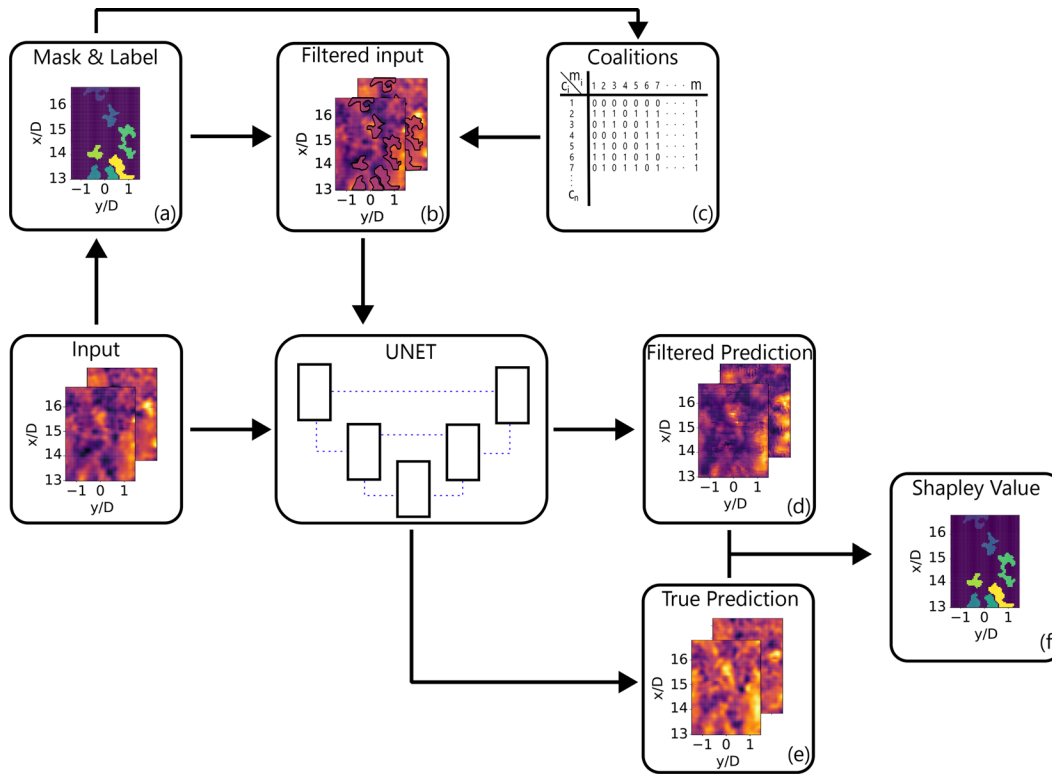


FIG. 1. Schematic representation of the framework used to calculate the Shapley values. For each validation snapshot, a filter is applied based on the vorticity (ω_z^*) to identify the structures in the snapshot (a). For each snapshot, coalitions are then constructed (c), which determine the structures that ought to be removed from the input snapshot (b). The filtered image is then provided as input to the neural network, and the obtained output (d) is compared with the true prediction (e), thereby yielding the Shapley values.

and adjusting their scale to overlay the heatmap onto the original image.

The main idea behind Grad-Cam is that the convolutional layers naturally preserve spatial information that is lost in fully connected layers.^{11,31} Vinogradova *et al.*³¹ used Grad-CAM for a CNN trained for Semantic Segmentation, and they suggest choosing the last convolutional layers to strike the best balance between high-level semantics and detailed spatial information. Neurons in these layers seek out semantic class-specific information in the image. Grad-CAM utilizes the gradient information in input to the last convolutional layer of the CNN to assign importance values to each neuron for a particular decision of interest.

Following the original implementation,¹¹ we consider a CNN with an RGB image as input ($\text{Img} \in \mathbb{R}^{x \times y \times 3}$), where x and y represent the width and height of the image, respectively. The output is a tensor $\in \mathbb{R}^{x \times y \times c}$, where c represents the number of classes.

The first operation involves computing the gradient [Eq. (8)] for each class with respect to the feature map activations A^k of a convolutional layer, where K is the number of feature maps (kernels) in the target layer. These gradients are then averaged across the spatial dimensions of the image to obtain an importance weight α_k^c . The coefficient α_k^c represents a partial linearization of the deep network downstream of A^k and captures the importance of the feature map k for a target class c ,

$$\alpha_k^c = \frac{1}{Z} \sum_m \sum_n \underbrace{\frac{\partial y^c}{\partial A_{mn}^k}}_{\text{Gradients via backprop}}, \quad (8)$$

Global Average Pooling

where $Z = x \times y$ is a normalization factor corresponding to the total number of spatial locations in the feature map.

The last step is the weighted combination of activation maps followed by a Rectified Linear Unit (ReLU) to obtain Eq. (9). Selvaraju *et al.*¹¹ introduced the ReLU in Eq. (9) because they are only interested in the features that have a positive influence on the class of interest, i.e., the pixels whose intensity should be increased to increase y^c ,

$$L_{\text{Grad-CAM}}^c = \text{ReLU} \left(\underbrace{\sum_k \alpha_k^c A^k}_{\text{linear combination}} \right). \quad (9)$$

This methodology can be applied to any differentiable activation function, as reported in Ref. 11.

In Vinogradova *et al.*,³¹ the authors adapted Grad-CAM to an image segmentation problem, by substituting y^c in Eq. (8) with proposed $\sum_{(m,n) \in M} y_{mn}^c$, where M is a set of pixel indices of interest in the output mask. M can denote just a single pixel, or pixels of an object instance, or simply all pixels of the image.




FIG. 2. Time-averaged streamwise velocity component normalized with the jet's exit velocity (V_{exit}) plotted as a function of the spanwise coordinate. The data are taken from the dataset of Roy *et al.*¹⁵.

What is discussed in this paper starts from the approach proposed by Vinogradova *et al.*,³¹ in order to extend Grad-CAM to a flow-prediction problem, replacing flow quantities, i.e., $y = u$.

III. RESULTS

The dataset¹⁵ consists of 8451 instantaneous velocity snapshots, obtained via PIV of an axisymmetric jet. The analyzed jet flow is characterized by a Reynolds number (Re), based on the mass flow rate and the jet diameter, of $Re = 21\,000$, with an exit velocity $V_{exit} = 66$. The spatial mesh grid comprises 67×51 points, spaced in both directions at $\Delta z = \Delta x = 0.268$ mm. The time resolution is $\Delta t = 0.01$ ms. The streamwise velocity profiles along the spanwise direction for $x/D = 13, 14.88, \text{ and } 16.77$ are shown in Fig. 2.

To extract meaningful features from this dataset according methods described in A, B, C, we employ a deep learning approach based on a U-Net architecture.

The implemented U-Net architecture consists of three main components, as sketched in Fig. 3: the contracting path (encoder), the bottleneck, and the expansive path (decoder). This architecture is designed to learn hierarchical representations of input data by combining local and global information through skip connections, as described in Ronneberger *et al.*²²

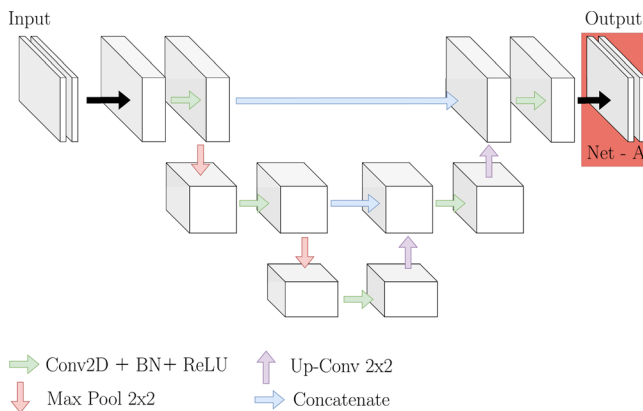


FIG. 3. Scheme of the implemented neural network.

The encoder comprises a series of convolutional blocks, each performing a two-dimensional convolution followed by batch normalization to stabilize and accelerate the learning process. These operations are followed by the ReLU activation function, which introduces non-linearity and enhances the model's expressiveness. At the end of each block, a MaxPooling2D operation with a 2×2 kernel is applied, reducing the spatial dimensions of the feature maps by a factor of 2. This enables the network to extract increasingly non-linear features as it progresses through the layers.

The bottleneck represents the deepest part of the U-Net and serves as a latent space that captures the most compact and meaningful representations of the input data. This stage is crucial for learning high-level non-linear features.

The decoder is designed to progressively reconstruct the spatial dimensions of the image. It uses transposed convolutions with a 2×2 kernel to double the spatial dimensions of the feature maps. In each block of the decoder, the resulting feature maps are concatenated with those from the corresponding block in the contracting path via skip connections, allowing the network to combine high-level semantic details with low-level, high-resolution information. After concatenation, convolutions, batch normalization, and ReLU activations are applied to further refine the reconstructed features.

The final output of the network is obtained by applying a convolution that adjusts the number of channels to match the desired output, ensuring that the resulting tensor has the same spatial dimensions as the input. This modular and flexible structure makes the U-Net particularly effective for image segmentation and reconstruction tasks, preserving important details by integrating local and global information.

The skip connections, in particular, play a critical role in enhancing the network's performance by enabling direct information flow across layers, as highlighted in previous studies.^{22,32}

The NN was trained randomly splitting the dataset into a training (60%) and a validation set (40%).

For the training process, a RMSprop gradient-descent algorithm was used,³³ with a learning rate of 10^{-4} and a momentum of 0.9. We used the early stopping in order to prevent overfitting. At the end of the training, the rms of the error was around 10^{-4} .

Figure 4 shows the percentage error (ϵ) of a velocity field predicted by the trained NN. The error is calculated as the local difference




FIG. 4. Percentage error of the velocity on the prediction of the streamwise velocity component.




FIG. 5. Results of the percolation analysis. The red dotted line highlights the selected value of ω_z^* used for the filter.

with respect to the ground truth of the predicted velocity field, normalized with respect to the ground truth. The error is quite limited to within $\pm 4\%$ and as such is of minimal consequence for the intended application of the neural network.

The prediction task performed by the U-net is implemented in fed framework used to extract the features' importance, as schematically described in Fig. 1. As reported in Sec. II, a thresholding on the value of ω_z is applied to each snapshot to identify the connected regions within each frame. The resulting identified regions are then labeled, as indicated in Fig. 1(a). Coalitions are formed by including or excluding different features to simulate how the presence or absence of each one influences the predictive outcome [see Fig. 1(c)].

For each coalition element, the corresponding filtered input is calculated. This input is identical to the real input, except that elements not in the coalition are set to zero. The filtered input is then provided to the neural network, which performs the prediction. The same process is repeated with the unfiltered input. By comparing the two outputs, it is possible to measure how important the coalition is for the prediction.

The outcome can be sensitive with respect to the choice of ω_{thr}^* , as this obviously influences the identification. Therefore, a percolation analysis was performed to select the optimal value that maximizes the number of identified structures, as shown in Fig. 5. The resulting value of ω_{thr}^* is 0.354.

Figure 6(a) shows the values of ϕ as a function of the area of the structure. As the area increases, ϕ increases accordingly. This is an obvious consequence of the fact that the larger the area of a turbulent structure, the greater the content will be removed from the snapshot, or more specifically the greater the amount of information that will be




FIG. 6. Scatter plot of the Shapley values (ϕ) (a) and their normalization with the area of the structure (b) as a function of the corresponding area of the structure (A_s), normalized with respect to the square of the jet exit diameter (D^2).

removed from the input data. Consequently, the network will produce a less accurate result, and the Shapley value will increase. Cremades *et al.*⁸ already showed that this is an obvious result, which does not add any significant information to the interpretation of the flow.

To avoid this misleading result, a normalization is introduced according to the area occupied by the structure (A_s), specifically defined as

$$\Phi = \frac{|(\phi/A_s)|}{\max(|\phi/A_s|)}. \quad (10)$$

Figure 6(b) shows that the most important structures are not necessarily those with high values of A_s . To better understand what the identified structures represent in relation to their importance, let us consider evaluating an integral value of vorticity on each structure. The integral value of the vorticity calculated over the identified structures, with area A_s , allows us to draw some conclusions on the structures' importance,

$$\langle f \rangle = \frac{1}{A_s} \int f(x, y) d\sigma, \quad (11)$$

where f indicates the generic quantity that is being analyzed.

Figure 7(a) shows the joint probability density function (jPDF) between Φ and $\langle \omega_z \rangle$, which highlights two main clusters, based on the sign of the vorticity.

It is evident that, regardless of the sign of $\langle \omega_z \rangle$, higher values of Φ occur for $0.25 < |\langle \omega_z \rangle| < 0.5$. This suggests, contrary to initial assumptions, that the most relevant regions are not those with the largest vorticity.

Additionally, it is clear that most of the identified structures show values in the range of $0.1 < \Phi < 0.2$. It is important to note that even though Φ is normalized to reach unity, as shown in Fig. 7(b), the probability of $\Phi > 0.5$ is negligible. Therefore, it can be concluded that statistically, the identified structures have $0 < |\langle \omega_z \rangle| < 0.75$. It is highly probable to find moderately important structures ($\Phi \sim 0.15$) with $0.25 < |\langle \omega_z \rangle| < 0.5$.

Figures 8(a) and 8(b) show the conditionally averaged structure corresponding to the two clusters characterized by values of positive and negative vorticity, respectively. Specifically, for each region identified using the ω_{thr}^* filter, the bounding box containing it and its center were identified. In accordance with their cluster, they were re-centered to share a common center and then averaged, thereof defining a conditional vorticity field ω_z^* . The two clusters, respectively, for $\omega_z^* > \omega_{thr}^*$




FIG. 7. Joint PDF of the average integral vorticity ($\langle \omega_z^* \rangle$) and of the Shapley values (Φ) (a); PDF of the averaged Shapley values (b).




FIG. 8. Kernel-SHAP conditional averaged structures, according to the value of $\omega_z^* > 0$ (a) and $\omega_z^* < 0$ (b), with overlaid velocity vectors.

and $\omega_z^* < \omega_{thr}^*$, have a comparable number of structures, which is to be expected, given the axisymmetric flow field. This property is reflected in the conditionally averaged structures, which show a similar topology, as seen in Figs. 8(a) and 8(b), with an inverted rotation direction.

The jPDF between the shape factor (SF), defined as the ratio between the extent of the bounding box containing the structure in the streamwise direction and in the spanwise direction, and Φ is shown in Fig. 9. It is evident that structures with SF in the range between 1 and 1.5 are not only the most common but also have the highest values of Φ .

In addition to the connection with the topology of the flow structures, it is also of interest to understand how these structures and their importance evaluated with the Shapley values can be related to the turbulence properties of the flow. We introduce the turbulence dissipation,

$$\varepsilon = -\nu \overline{\frac{\partial u'_i}{\partial x_j} \frac{\partial u'_i}{\partial x_j}}, \quad (12)$$

where $u' = u - \bar{u}$ is obtained through the Reynolds decomposition and subscript “i” denotes the components.

The average integral, obtained with Eq. (11), of the dissipation function on the structure’s surface $\langle \varepsilon \rangle$ was normalized with the average integral of the dissipation function’s computed on the entire field $\langle \varepsilon_0 \rangle$. From Fig. 10, it can be seen that as $\frac{\langle \varepsilon \rangle}{\langle \varepsilon_0 \rangle}$ increases, Φ also increases. This suggests that the most important structures are those that contribute the most to dissipation.




FIG. 9. Joint PDF of the shape factor (SF) and of the Shapley values (Φ).




FIG. 10. Joint PDF of the normalized dissipation function and of the Shapley values (dashed lines). Joint PDF of the normalized dissipation function and the integral average of the normalized vorticity (solid lines).

Hussain³⁴ points out that dissipation is not concentrated in regions of maximum vorticity but rather in adjacent areas where the interaction between coherent structures occurs, and where longitudinal vortices are stretched and connect to spanwise vortices. Moderate vorticity areas, in this context, refer to the saddle regions between coherent structures, where most turbulence production and dissipation takes place.

The results of the analysis carried out with Grad-CAM and Gradient-SHAP shall now be examined. As already mentioned in Sec. II A, these two methods not require an *a priori* metric to filter the relevant structures in the flow.

The schematic framework of Grad-CAM and GRADIENT-Shap is shown in Fig. 11. These methods give in output heatmaps (Ψ_i) that carry the importance of each input pixel on the final prediction. The subscript *i* for Grad-CAM depends on the number of output channels, while for Gradient-SHAP, it depends on the number of input channels. In the present case, the prediction yields in output the velocity components u, v starting from the same velocity components at the previous step, thereof resulting in $i = 2$ in both cases.

To extract the relevant regions of the flow, a filter is applied to Ψ_i to identify connected regions. The applied filter requires that the heatmap values are such that, at a pixel level, it must be verified that $\sqrt{\Psi_1^2 + \Psi_2^2} > H$. The value of the threshold *H* is chosen through percolation analysis to maximize the number of identified structures, and it is set to 0.87 for Grad-CAM and 0.3 for Gradient-SHAP.

The regions of the flow that were identified with Gradient-SHAP and Grad-CAM according to the criterion defined above are those that are relevant to the flow. In order to compare the results obtained with the three methods, a common metric is defined based on the out-of-plane vorticity. In particular, we focus on the conditionally averaged structures similar to those identified in Fig. 8. Figures 12 and 13 display the average conditional structures for $\omega_z > 0$ [see Figs. 12(a) and 13(a)] and for $\omega_z < 0$ [see Figs. 12(b) and 13(b)]. An initial evident result emerges by comparing these averaged conditional structures with Fig. 8: indeed, Grad-CAM and Gradient-SHAP yield similar results in terms of conditional structures: they are different from those obtained with SHAP. This is not surprising, since the results from SHAP are forced to identify those regions that are characterized by values of $(|\omega_z| > \omega_{zthr})$, which inevitably leads to the identification of the

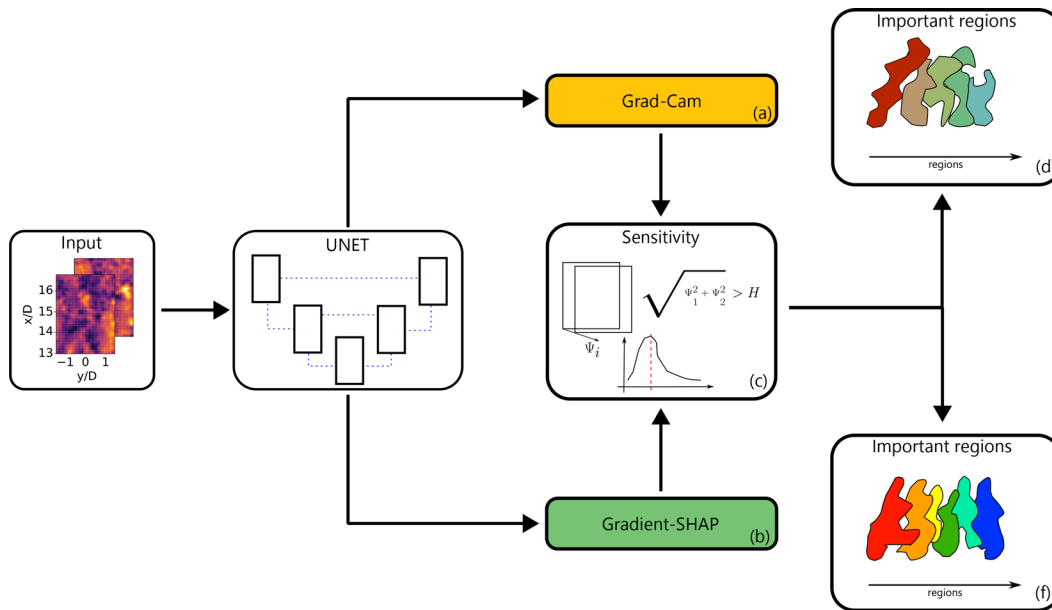


FIG. 11. Schematic representation of the framework used for Grad-CAM (a) and Gradient-SHAP (b). Each of the two algorithms is applied to the trained UNET. After applying Gradient-SHAP and Grad-CAM for each validation snapshot of the network, a percolation analysis is performed on the resulting heatmaps (c) to identify the important regions within the flow field (d) and (f).




FIG. 12. Grad-CAM conditional averaged structures, according to the value of $\omega_z^* > 0$ (a) and $\omega_z^* < 0$ (b), with overlaid velocity vectors.




FIG. 13. Gradient-SHAP conditional averaged structures, according to the value of $\omega_z^* > 0$ (a) and $\omega_z^* < 0$ (b), with overlaid velocity vectors.

vortical structures shown in Fig. 8. This bias could hinder the generality of the results, somehow hiding the physical mechanism at play.

The physics of turbulent jets involves energy transfer across multiple scales, from large vortices such as vortex rings to smaller scales dominated by turbulent dissipation.³⁴ The average conditional structures of Grad-CAM and Gradient-SHAP, respectively, Figs. 12 and 13, highlight regions of moderate vorticity and strain, areas known in the literature to be regions of high dissipation and turbulent kinetic energy.

Grad-CAM and Gradient-SHAP highlight exactly these areas, aligning with other studies,^{3,5,34} which have shown that dissipation is concentrated in regions strain rather than at vorticity peaks. Compared to SHAP, which requires predefined segmentation, Grad-

CAM identifies relevant regions directly from the model gradient. This less constrained approach avoids biases introduced by segmentation and allows for the identification of physically significant regions for jet dynamics, as demonstrated by the results.

IV. CONCLUSION

In this study, three different XAI algorithms were used to identify the important regions for an axisymmetric jet. Starting from time-resolved PIV data for an axisymmetric jet,¹⁵ a neural network was trained to perform a one-step prediction. To apply SHAP, it was necessary to define a criterion for segmenting the flow field in such a way as to identify regions for which the importance could be calculated using SHAP. Regions with high Shapley values indicate areas where the

model focuses to make its predictions. The results show that the most relevant structures do not always correspond to those with high vorticity. This suggests that structures with moderate vorticity (ω_z^* between 0.25 and 0.5) appear to play a key role in energy transfer and jet dynamics. This behavior is consistent with previous studies that highlighted the contribution of intermediate regions to mixing and energy dissipation.^{1,35} Additionally, the analysis shows that the structures with the highest normalized Shapley values coincide with those that contribute the most to the dissipation of turbulent kinetic energy. Regions with high dissipation not only indicate loss of turbulent kinetic energy but also are also closely associated with mixing and the dispersion of energy scales.

In contrast, the use of Grad-CAM and Gradient-SHAP does not require prior recognition of structures, which is advantageous as it avoids any bias. The heatmaps produced by SHAP, Grad-CAM, and Gradient-SHAP differ significantly. The main results show that Grad-CAM and Gradient-SHAP identify regions with a uniform spatial distribution, suggesting a correlation with the transport of momentum by smaller-scale vortices. The identified structures, often with a shape ratio close to 1, suggest the predominance of nearly circular vortices (e.g., vortex rings) in the jet dynamics. Vortex rings represent the dominant structures in energy transfer and turbulent mixing, supporting both jet dispersion and the formation of secondary instabilities.

Regarding the link between the regions of the flow identified through XAI as relevant to the flow and the dissipation, it can be argued that areas of high dissipation in turbulent flows tend to be concentrated in regions with strong velocity gradients (shear fronts) or intense vortical structures (vortex tubes). In these regions, the combined effects of strain (flow deformation) and vorticity give rise to extreme events where turbulent kinetic energy is rapidly converted into heat. Such events often appear intermittent in space and time, and typically occur at small scales, close to the Kolmogorov scale. Multiple studies^{36–38} confirm that the most intense dissipation episodes are found in thin shear layers or coherent vortex tubes, where vortex stretching and high-velocity gradients produce much higher dissipation rates than the flow average.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

E. Amico: Conceptualization (equal); Data curation (equal); Methodology (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **L. Matteucci:** Conceptualization (equal); Data curation (equal); Methodology (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **G. Cafiero:** Conceptualization (equal); Data curation (equal); Methodology (equal); Resources (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- S. C. Crow and F. H. Champagne, “Orderly structure in jet turbulence,” *J. Fluid Mech.* **48**, 547–591 (1971).
- P. O. A. L. Davies, M. J. Fisher, and M. J. Barratt, “The characteristics of the turbulence in the mixing region of a round jet,” *J. Fluid Mech.* **15**, 337–367 (1963).
- J. Jimenez and A. A. Wray, “On the characteristics of vortex filaments in isotropic turbulence,” *J. Fluid Mech.* **373**, 255–285 (1998).
- G. Cafiero, G. Castrillo, and T. Astarita, “Turbulence properties in jets with fractal grid turbulence,” *J. Fluid Mech.* **915**, A12 (2021).
- E. Pickering, G. Rigas, P. A. S. Nogueira, A. V. G. Cavalieri, O. T. Schmidt, and T. Colonius, “Lift-up, Kelvin–Helmholtz and Orr mechanisms in turbulent jets,” *J. Fluid Mech.* **896**, A2 (2020).
- P. Nogueira, A. Cavalieri, P. Jordan, and V. Jaunet, “Large-scale streaky structures in turbulent jets,” *J. Fluid Mech.* **873**, 211–237 (2019).
- M. Lellep, J. Prexl, B. Eckhardt, and M. Linkmann, “Interpreted machine learning in fluid dynamics: Explaining relaminarisation events in wall-bounded shear flows,” *J. Fluid Mech.* **942**, A2 (2022).
- A. Cremades, S. Hoyas, R. Deshpande, P. Quintero, M. Lellep, W. J. Lee, J. P. Monty, N. Hutchins, M. Linkmann, I. Marusic, and R. Vinuesa, “Identifying regions of importance in wall-bounded turbulence through explainable deep learning,” *Nat. Commun.* **15**, 3864 (2024).
- A. Cremades, S. Hoyas, and R. Vinuesa, “Additive-feature-attribution methods: A review on explainable artificial intelligence for fluid dynamics and heat transfer,” *Int. J. Heat Fluid Flow* **112**, 109662 (2025).
- A. Cremades, S. Hoyas, and R. Vinuesa, “Classically studied coherent structures only paint a partial picture of wall-bounded turbulence,” [arXiv:2410.23189](https://arxiv.org/abs/2410.23189) (2024).
- R. R. Selvaraju, M. Cogswell, A. Das *et al.*, “Grad-CAM: Why did you say that? visual explanations from deep networks via gradient-based localization,” *Int. J. Comput. Vis.* **128**, 336–359 (2020).
- H. Panwar, P. Gupta, M. K. Siddiqui, R. Morales-Menendez, P. Bhardwaj, and V. Singh, “A deep learning and grad-CAM based color visualization approach for fast detection of COVID-19 cases using chest X-ray and CT-scan images,” *Chaos, Solitons Fractals* **140**, 110190 (2020).
- S. Majid, F. Alenezi, S. Masood, M. Ahmad, E. S. Gündüz, and K. Polat, “Attention based CNN model for fire detection and localization in real-world images,” *Expert Syst. Appl.* **189**, 116114 (2022).
- V. Jahmunah, E. Ng, R.-S. Tan, S. L. Oh, and U. R. Acharya, “Explainable detection of myocardial infarction using deep learning models with grad-CAM technique on ECG signals,” *Comput. Biol. Med.* **146**, 105550 (2022).
- S. Roy, J. Miller, and G. Gunaratne, “Deviations from Taylor’s frozen hypothesis and scaling laws in inhomogeneous jet flows,” *Commun. Phys.* **4**, 32 (2021).
- Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature* **521**, 436–444 (2015).
- I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
- J. Ling, A. Kurzawski, and J. Templeton, “Machine learning strategies for systems with invariance properties,” *J. Comput. Phys.* **318**, 22–35 (2016).
- K. Duraisamy, G. Iaccarino, and H. Xiao, “Turbulence modeling in the age of data,” *Annu. Rev. Fluid Mech.* **51**, 357–377 (2019).
- S. L. Brunton, B. R. Noack, and P. Koumoutsakos, “Machine learning for fluid mechanics,” *Annu. Rev. Fluid Mech.* **52**, 477–508 (2020).
- P. Jordan and T. Colonius, “Wave packets and turbulent jet noise,” *Annu. Rev. Fluid Mech.* **45**, 173–195 (2013).
- O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (Springer, 2015), pp. 234–241.
- G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-informed machine learning,” *Nat. Rev. Phys.* **3**, 422–440 (2021).
- S. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Curran Associates Inc., 2017), pp. 4768–4777.
- I. Covert and S.-I. Lee, “Improving kernelshap: Practical shapley value estimation via linear regression,” [arXiv:2012.01536](https://arxiv.org/abs/2012.01536) (2021).
- M. T. Ribeiro, S. Singh, and C. Guestrin, “‘‘Why should i trust you?’’: Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM*

- SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16* (Association for Computing Machinery, New York, NY, 2016), pp. 1135–1144.
- ²⁷A. Das and P. Rad, “Opportunities and challenges in explainable artificial intelligence (XAI): A survey,” [arXiv:2006.11371](https://arxiv.org/abs/2006.11371) (2020).
- ²⁸A. Lozano-Durán, O. Flores, and J. Jiménez, “The three-dimensional structure of momentum transfer in turbulent channels,” *J. Fluid Mech.* **694**, 100–130 (2012).
- ²⁹M. Sundararajan, A. Taly, and Q. Yan, “Axiomatic attribution for deep networks,” in *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research* (PMLR, 2017), Vol. 70, pp. 3319–3328.
- ³⁰G. Erion, J. D. Janizek, P. Sturmfels, S. M. Lundberg, and S.-I. Lee, “Improving performance of deep learning models with axiomatic attribution priors and expected gradients,” *Nat. Mach. Intell.* **3**, 620–631 (2021).
- ³¹K. Vinogradova, A. Dibrov, and G. Myers, “Towards interpretable semantic segmentation via gradient-weighted class activation mapping (student abstract),” in *Proceedings of the AAAI Conference on Artificial Intelligence* (Association for the Advancement of Artificial Intelligence (AAAI), 2020), Vol. 34, pp. 13943–13944.
- ³²M. Drozdal, E. Vorontsov, G. Chartrand, S. Kadoury, and C. Pal, “The importance of skip connections in biomedical image segmentation,” in *Deep Learning and Data Labeling for Medical Applications* (Springer, 2016), pp. 179–187.
- ³³F. Zou, L. Shen, Z. Jie, W. Zhang, and W. Liu, “A sufficient condition for convergences of adam and rmsprop,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 11119–11127.
- ³⁴A. K. M. F. Hussain, “Coherent structures and turbulence,” *J. Fluid Mech.* **173**, 303–356 (1986).
- ³⁵W. K. George, “The self-preservation of turbulent flows and its relation to initial conditions and coherent structures,” in *Advances in Turbulence*, edited by W. K. George and R. E. A. Arndt and (Hemisphere, New York, 1989), pp. 39–73.
- ³⁶G. E. Elsinga, T. Ishihara, and J. C. R. Hunt, “Extreme dissipation and intermittency in turbulence at very high Reynolds numbers,” in *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* (Springer, 2020), Vol. 476, p. 20200591.
- ³⁷D. Buaria, A. Pumir, and E. Bodenschatz, “Self-attenuation of extreme events in Navier–Stokes turbulence,” *Nat. Commun.* **11**, 5852 (2020).
- ³⁸E. W. Saw, D. Kuzzay, D. Faranda, A. Guittonneau, F. Daviaud, C. Wiertel-Gasquet, V. Padilla, and B. Dubrulle, “Experimental characterization of extreme events of inertial dissipation in a turbulent swirling flow,” *Nat. Commun.* **7**(1), 1–8 (2016), see <https://www.nature.com/articles/ncomms12466>.