Improving turbulence control through explainable deep learning

Ricardo Vinuesa[1,2]

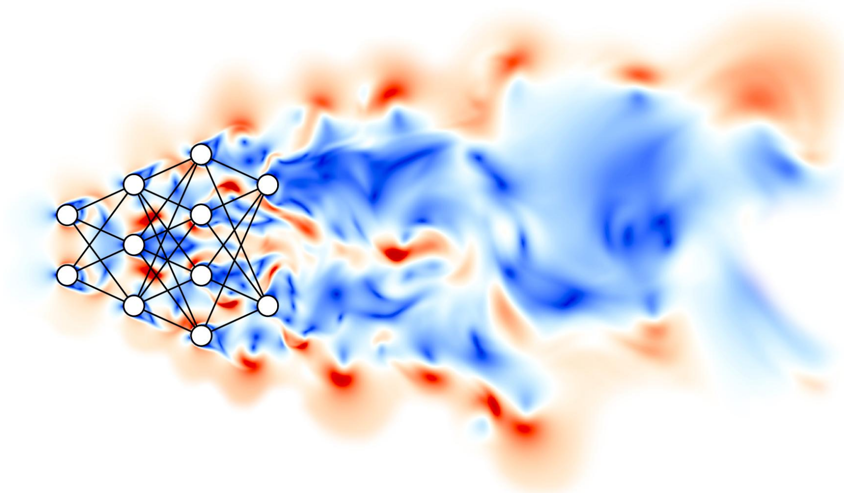

[1]FLOW, KTH Engineering Mechanics, Stockholm, Sweden.

[2]Swedish e-Science Research Centre (SeRC), Stockholm, Sweden.

25 June 2025, 41st IAHR World Congress, Harleman Lecture (Singapore)

The Sustainable Development Goals (SDGs)

SUSTAINABLE DEVELOPMENT GOALS

- **2030 Agenda for Sustainable Development adopted by all United Nations Member States in 2015**

- Shared blueprint for peace and prosperity for people and the planet

- Recognize that **ending poverty and other deprivations must go hand-in-hand** with strategies that improve health and education, reduce inequality, and spur economic growth – all while tackling **climate change** and working to **preserve our oceans and forests**

- **17 different** Sustainable Development Goals (SDGs)**; 169** targets

Motivation

- We want to answer the question: “Is there published evidence of **AI acting as an enabler or an inhibitor for each of the SDG targets**?”

Vinuesa et al., Nature Communications 11, 233 (2020)

Motivation

- We want to answer the question: “Is there published evidence of **AI acting as an enabler or an inhibitor for each of the SDG targets**?”
- We needed to assemble a **multi-disciplinary team** spanning the wide range of required areas of knowledge.

Vinuesa et al., Nature Communications 11, 233 (2020)

Dividing the 17 SDGs into 3 main pillars

– We divided the 17 SDGs into 3 main categories (Stockholm Resilience Center, 2017; United Nations, 2019): **Society, Economy, and Environment.**

FLOW **Vinuesa et al., Nature Communications 11, 233 (2020)**

Impact of AI on each of 169 targets

- We divided the 17 SDGs into 3 main categories (Stockholm Resilience Center, 2017; United Nations (2019): **Society, Economy, and Environment.**
- Percentage of targets where **positive (79%)** or **negative (35%)** impact of AI is documented:

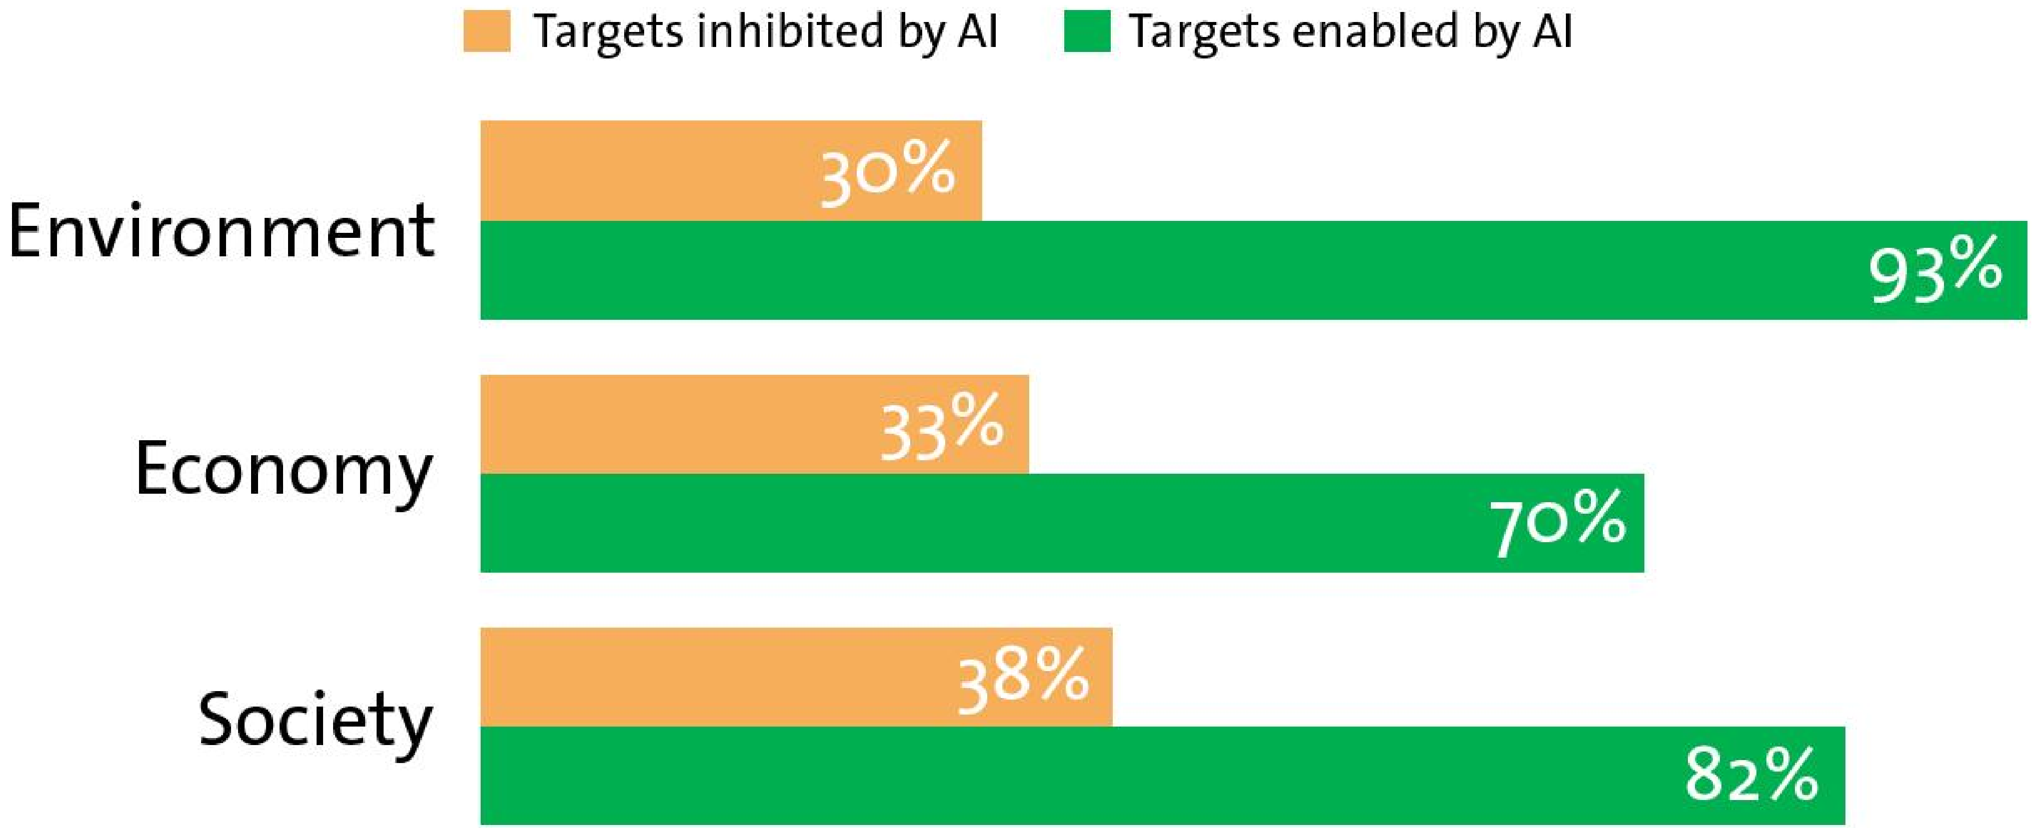

Vinuesa et al., Nature Communications 11, 233 (2020)

Results

– **Environment and Society <u>higher reduction of negative</u>; Economy the opposite.**

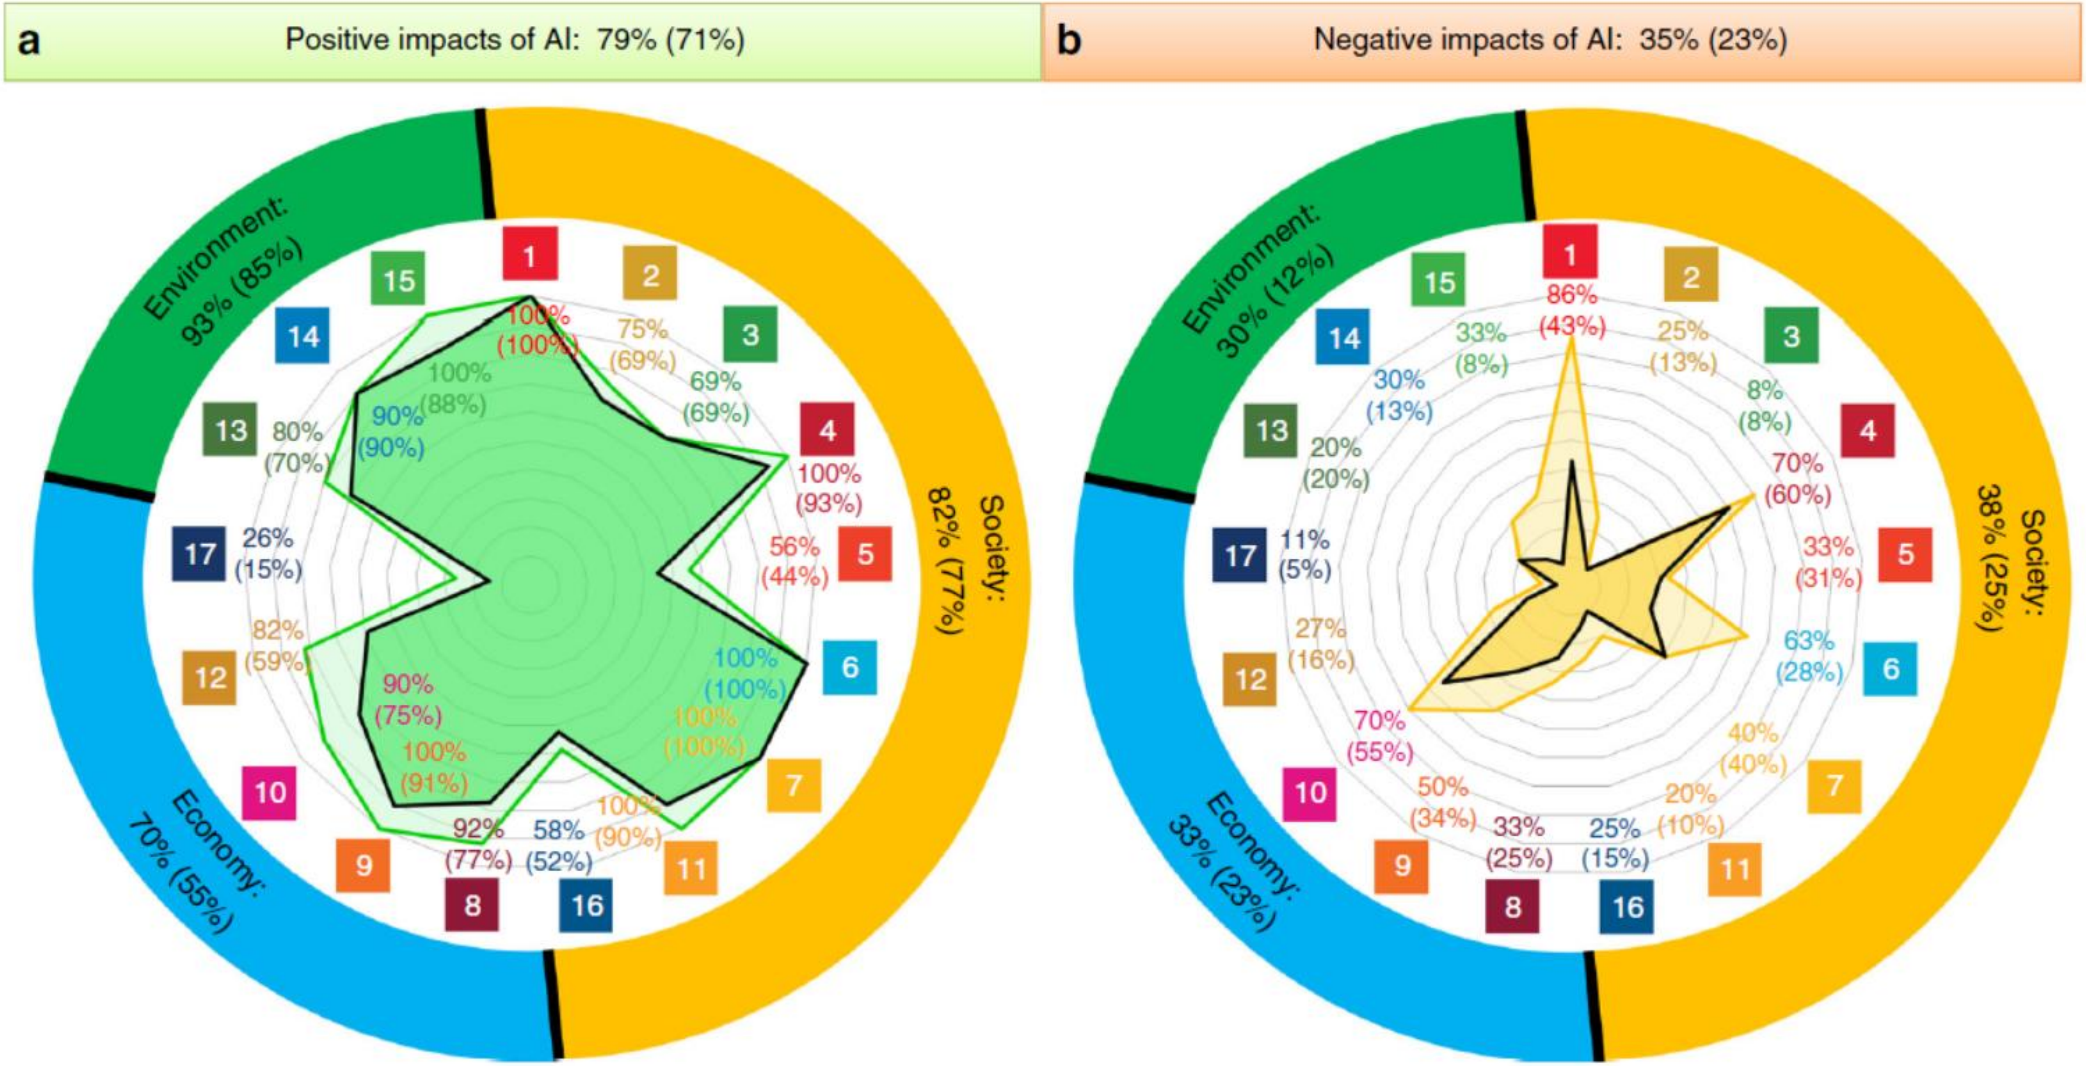

Vinuesa et al., Nature Communications 11, 233 (2020)

Types of evidence

- References using sophisticated tools and data to refer to this particular issue and with the possibility to be generalized are of type (A). **1**
- Studies based on data to refer to this particular issue, but with limited generalizability, are of type (B). **0.75**
- Anecdotal qualitative studies and methods are of type (C). **0.5**
- Purely theoretical or speculative references are of type (D). **0.25**

Vinuesa et al., Nature Communications 11, 233 (2020)

Types of evidence

– **Environment and Society <u>higher reduction of negative</u>; Economy the opposite.**

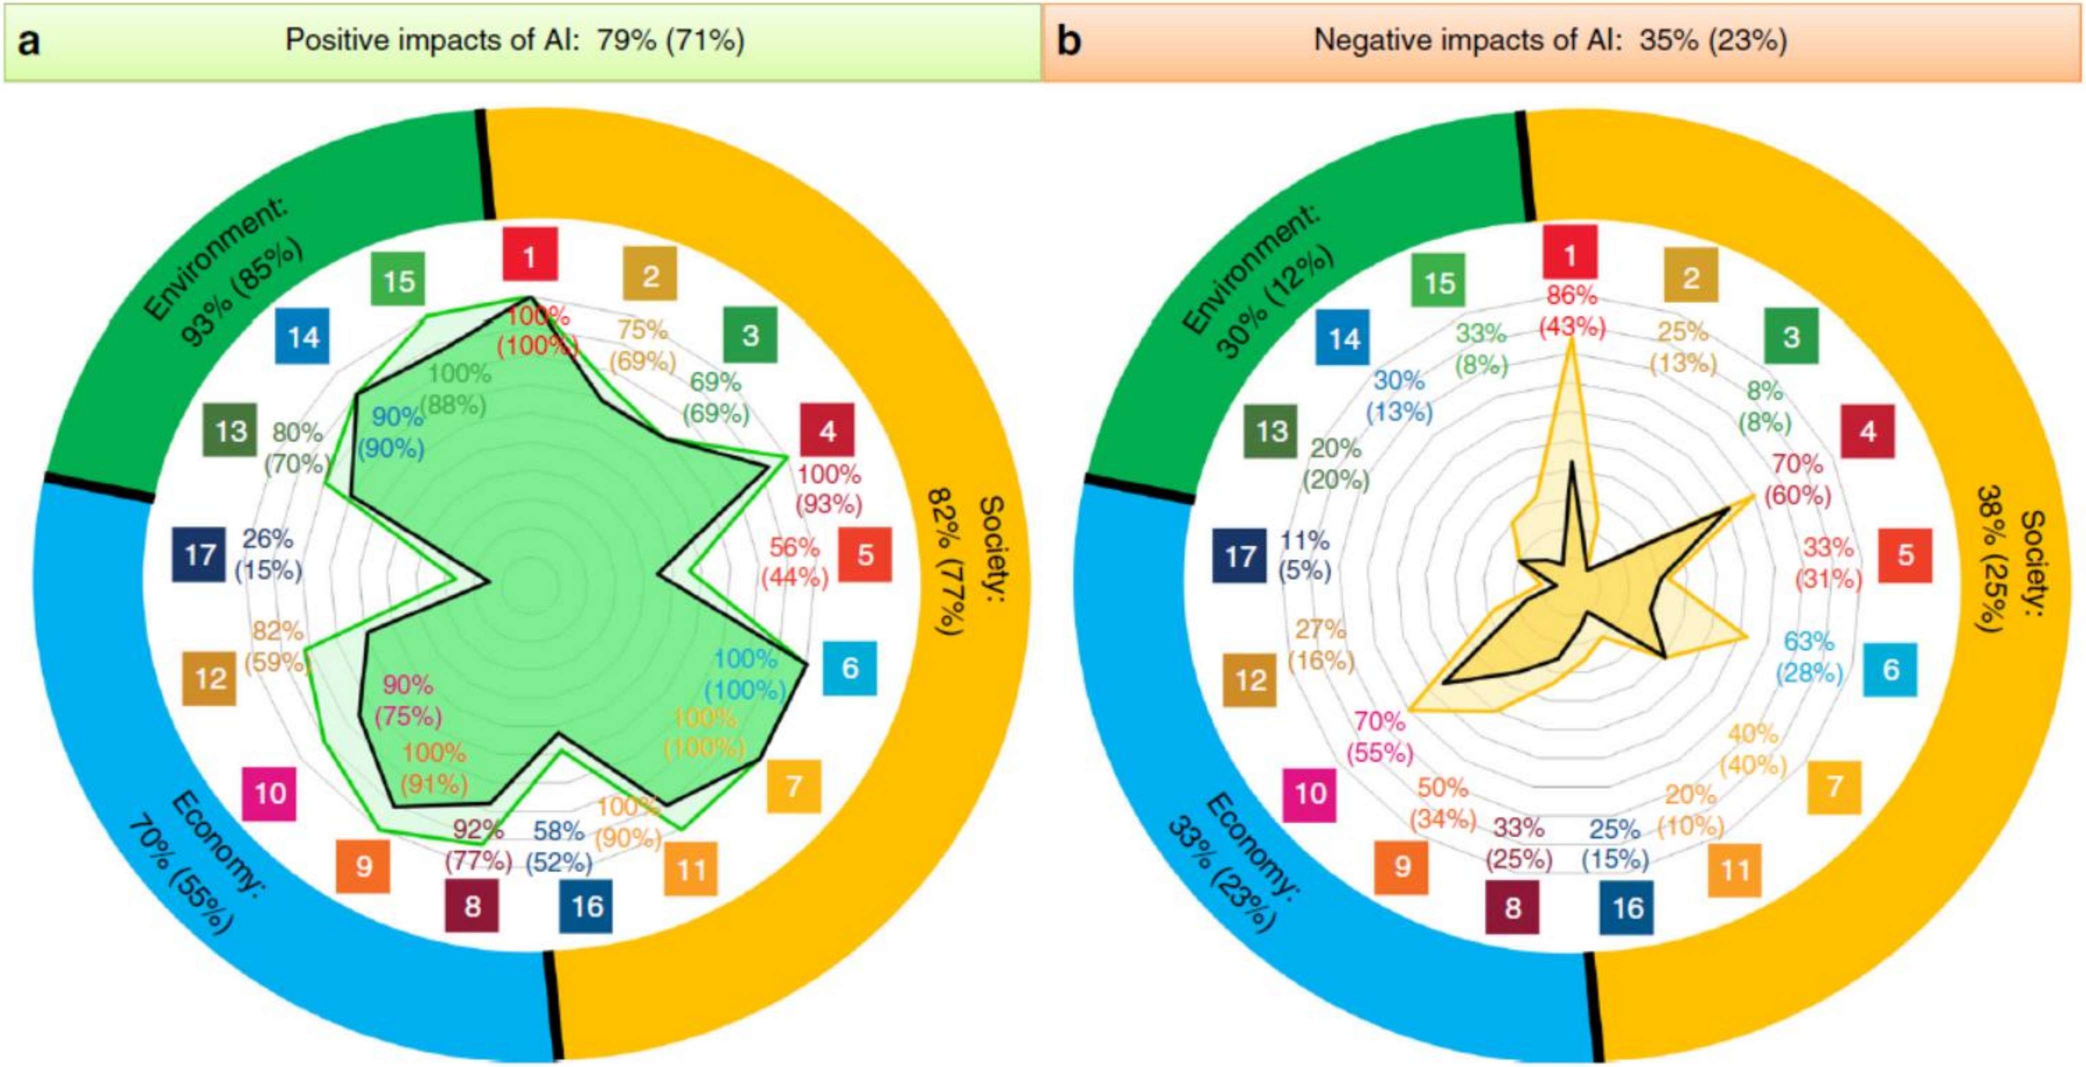

Vinuesa et al., Nature Communications 11, 233 (2020)

Some key results

- **POSITIVE:** AI-enabled technology which may help overcome current barriers (**satellite data** to track poverty, SDG1).

- **NEGATIVE:** Uneven opportunities to access AI resources may end up **increasing inequalities** (SDG 10).

Vinuesa et al., Nature Communications 11, 233 (2020)

One application: SDG11 on sustainable cities

- **POSITIVE:** Positive impact of AI on all 10 targets within SDG 11 on sustainable cities. In particular, AI will be able to help us build more accurate and robust **technology to measure air pollution in cities**, which causes 800,000 deaths each year in Europe alone.

Vinuesa et al., Nature Communications 11, 233 (2020)

Non-intrusive sensing in urban flows

- Using highly detailed simulations, we can reproduce the **flow in complex urban** environments.
- Use AI to **improve pollution measurement**.

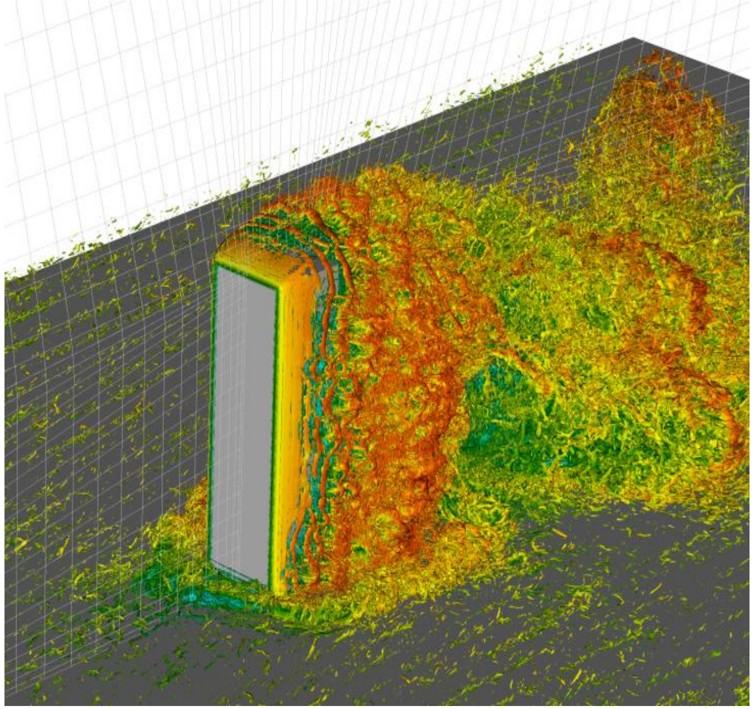

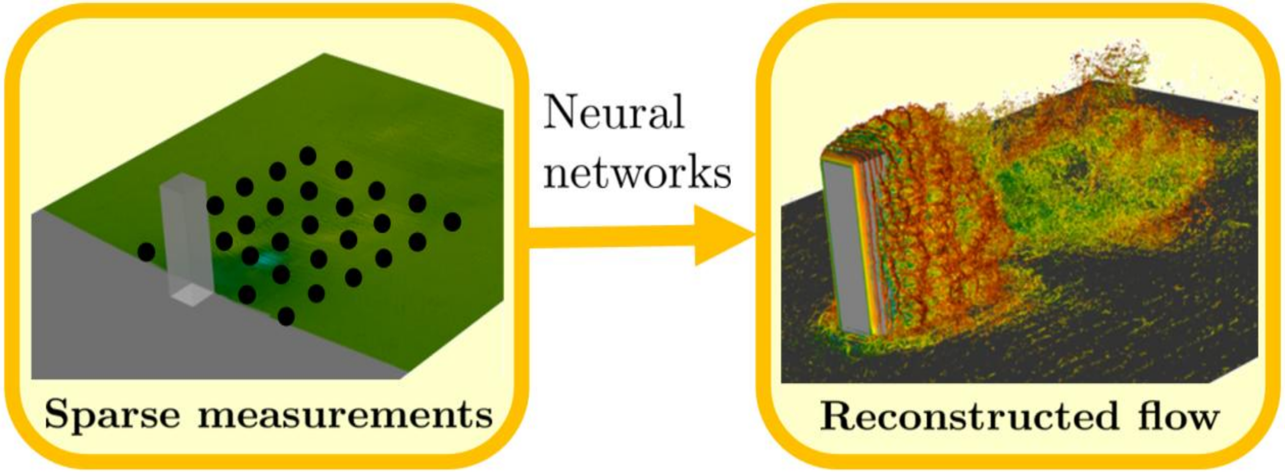

Lazpita et al., Phys. Fluids (2022); Martínez-Sánchez et al., IJHFF (2023)

Non-intrusive sensing in urban flows

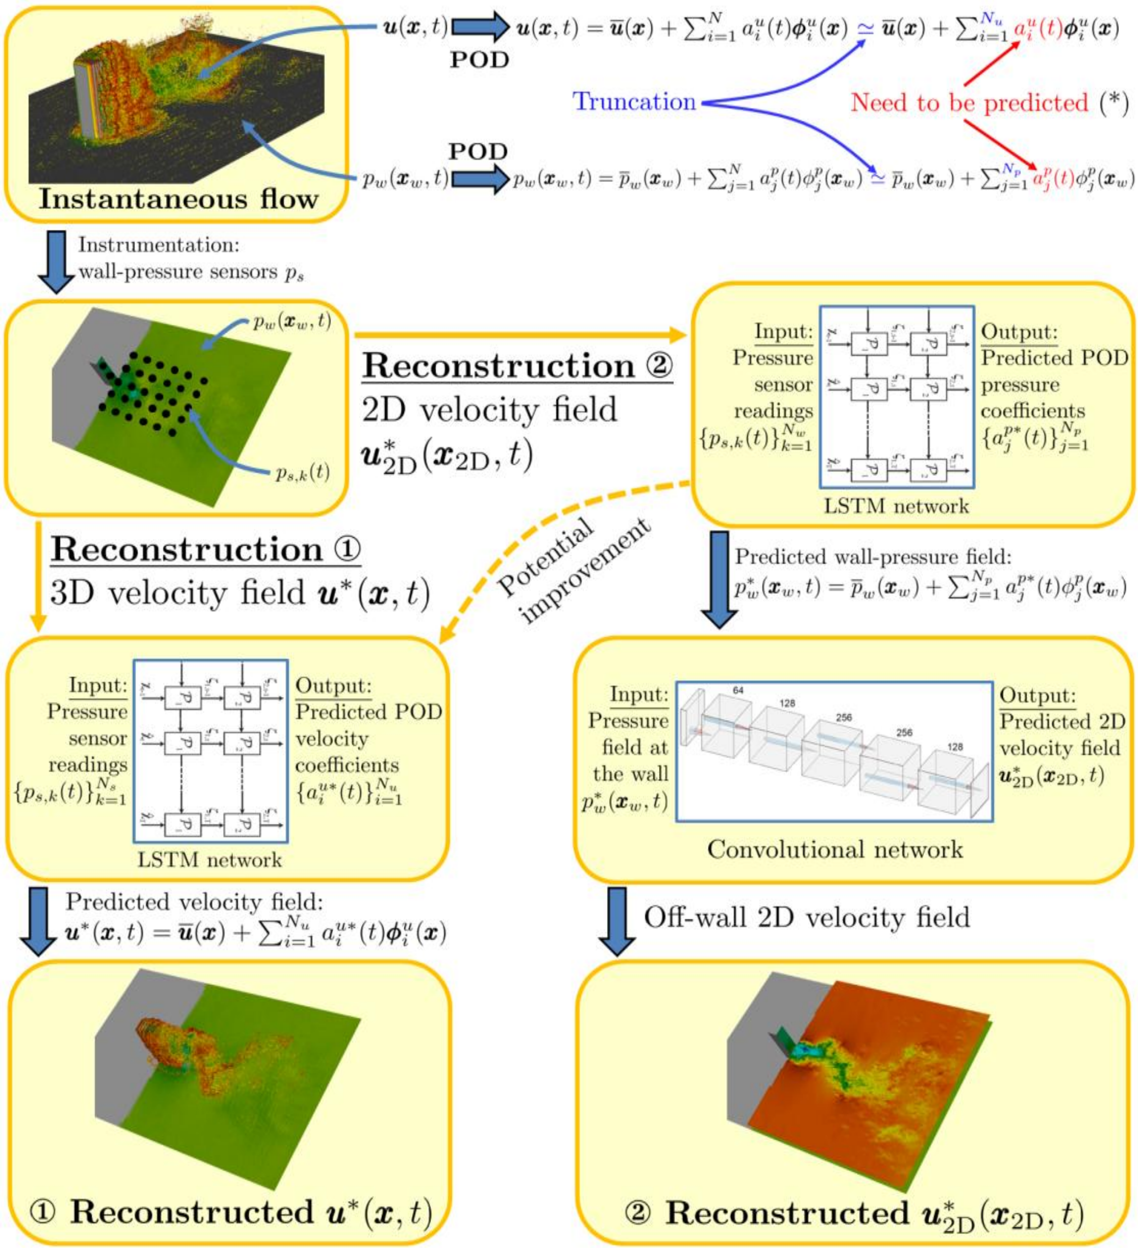

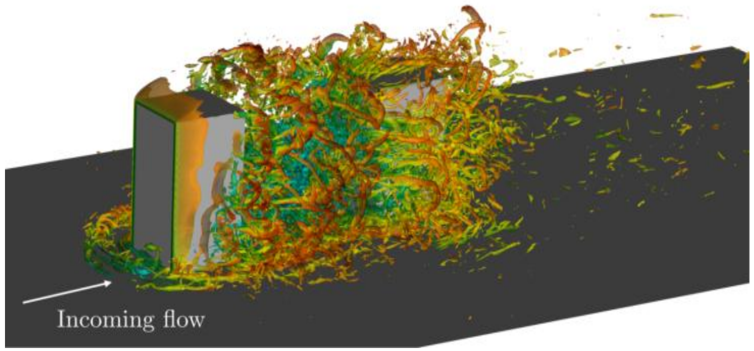

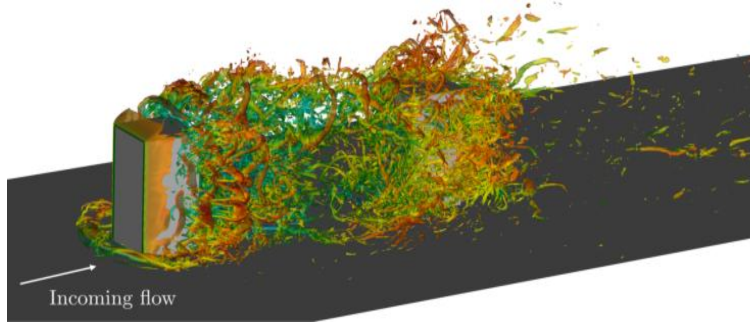

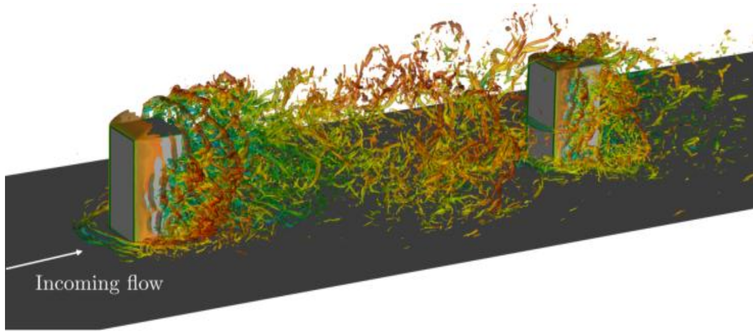

Lazpita et al., Phys. Fluids (2022); Martínez-Sánchez et al., IJHFF (2023)

Reduce pollution, minimize drag…
How to control a turbulent flow?

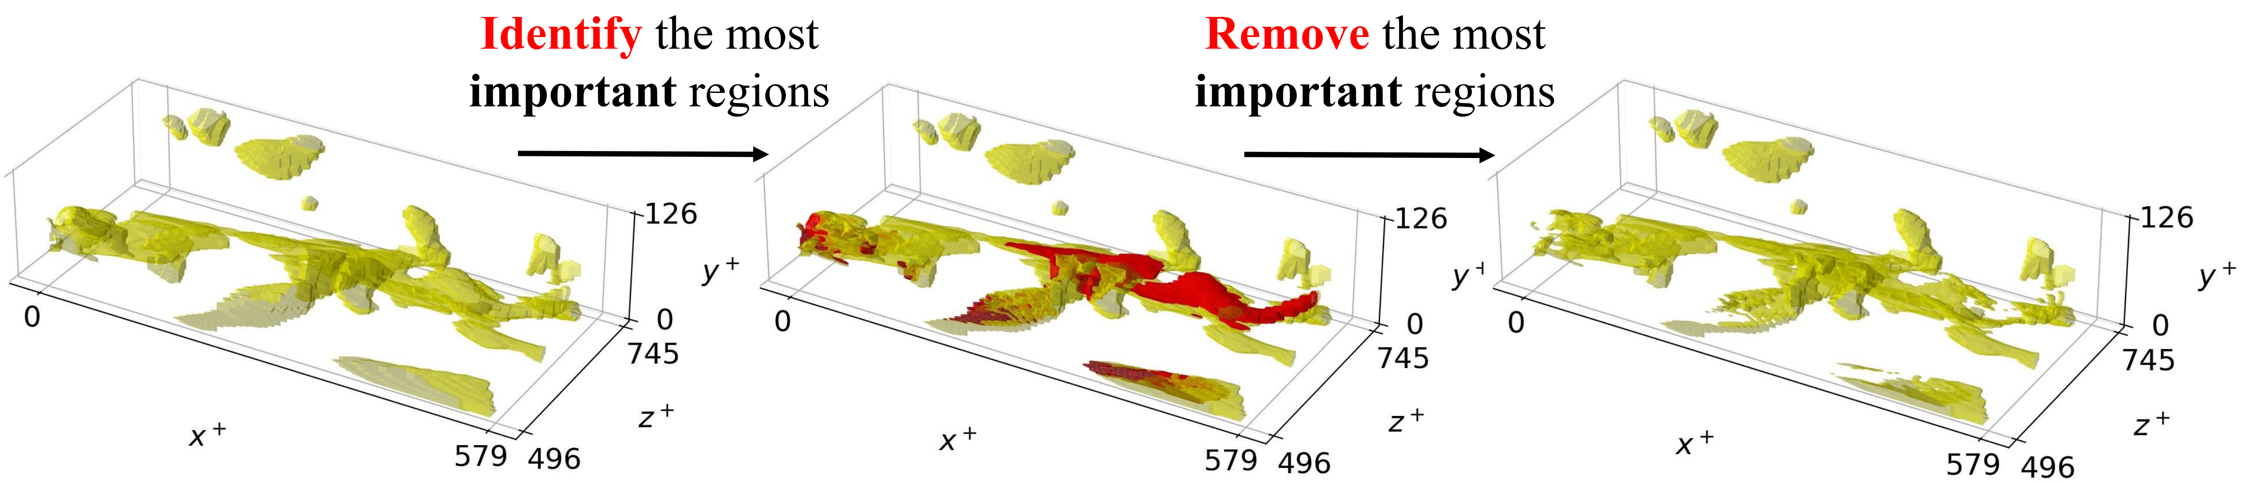

Outline

- Identify the most important regions with **explainable deep learning (XDL).**
- Remove the most important regions with deep reinforcement learning (DRL).

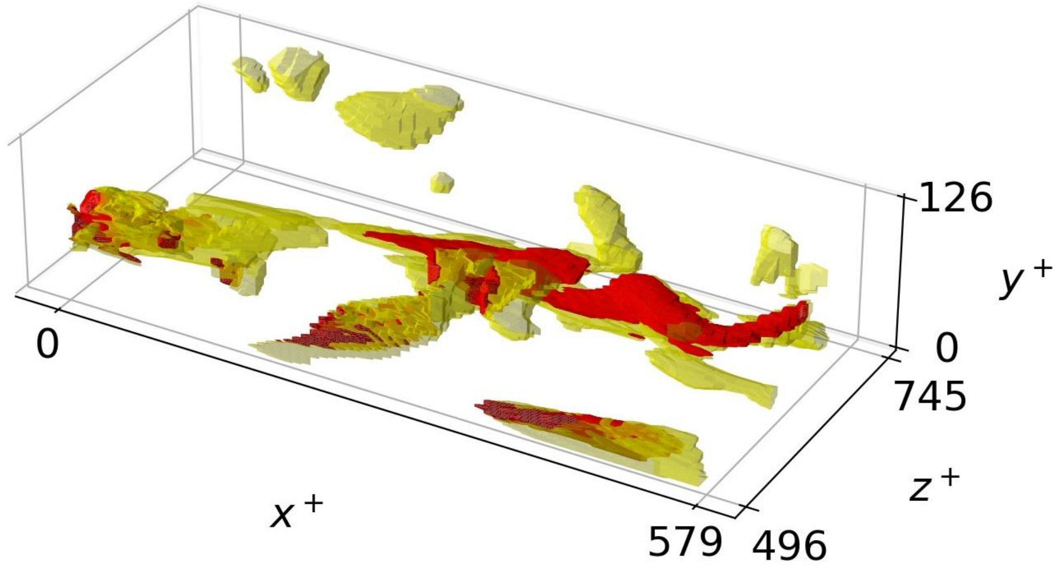

Which are the most important coherent structures in the flow? XAI framework

1. **Prediction** of a future instantaneous flow field through a **3D U-net.**
2. **Segmentation** of the domain **point by point.**
3. **The importance of each point** for the prediction is assessed by removing it from the field and re-calculating the error. XAI method based on SHapley Additive exPlanations **(SHAP).** ⇒ **By Lundberg and Lee (2017), very well established in ML (>30k citations).**

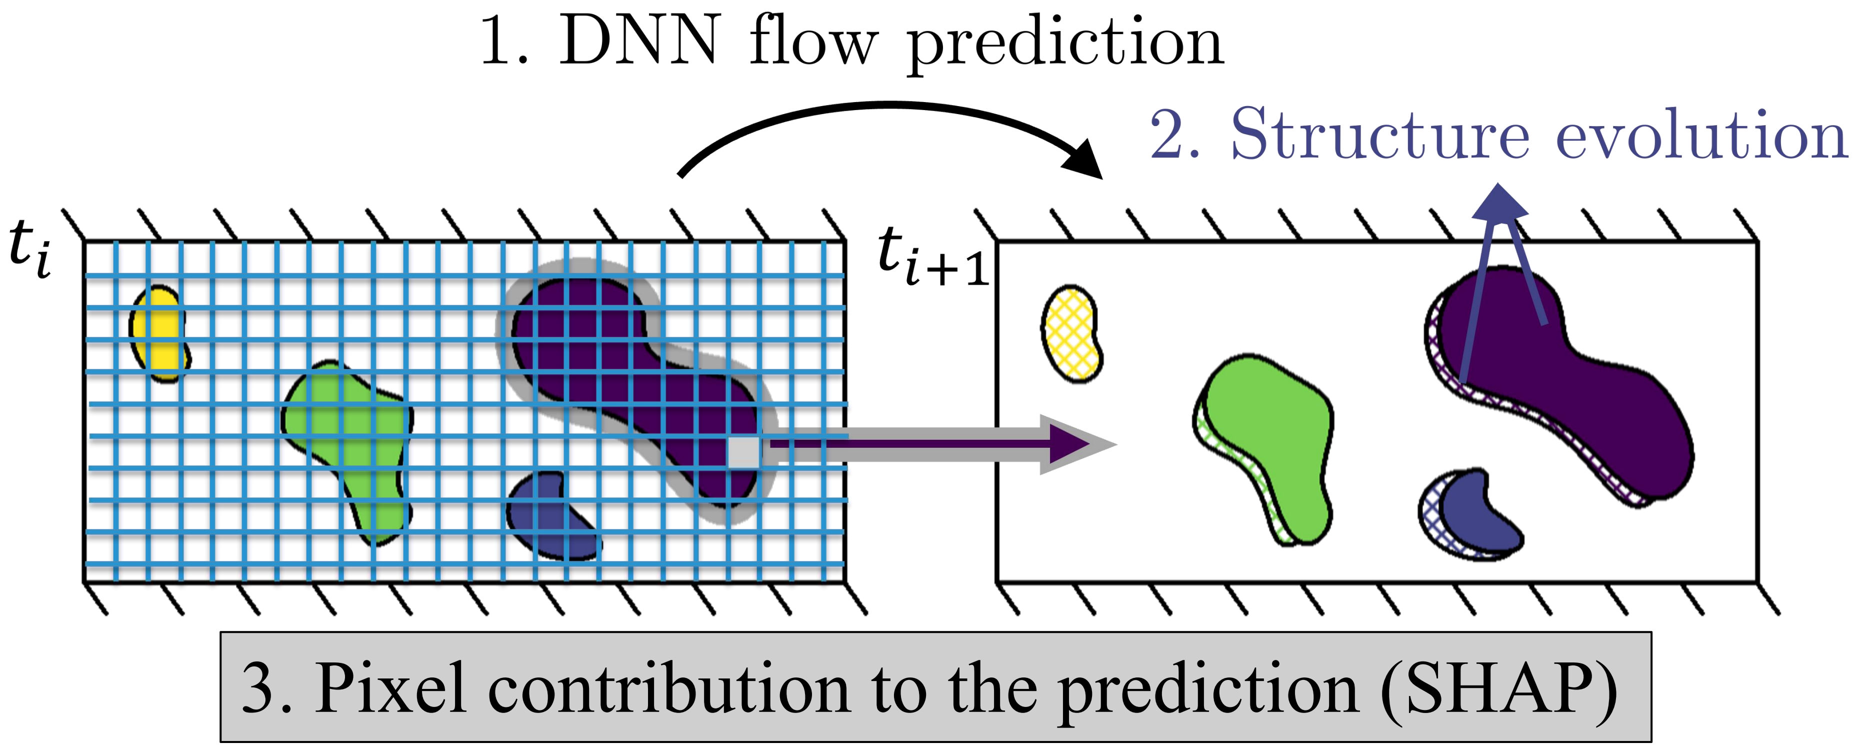

Cremades et al., Nature Communications 15, 3864 (2024)

Explainability algorithm

For each Q structure, we substitute its volume by zero fluctuations

Then we use the U-net to predict the next flow field

Predicted field at t_{i+1} (missing one Q)

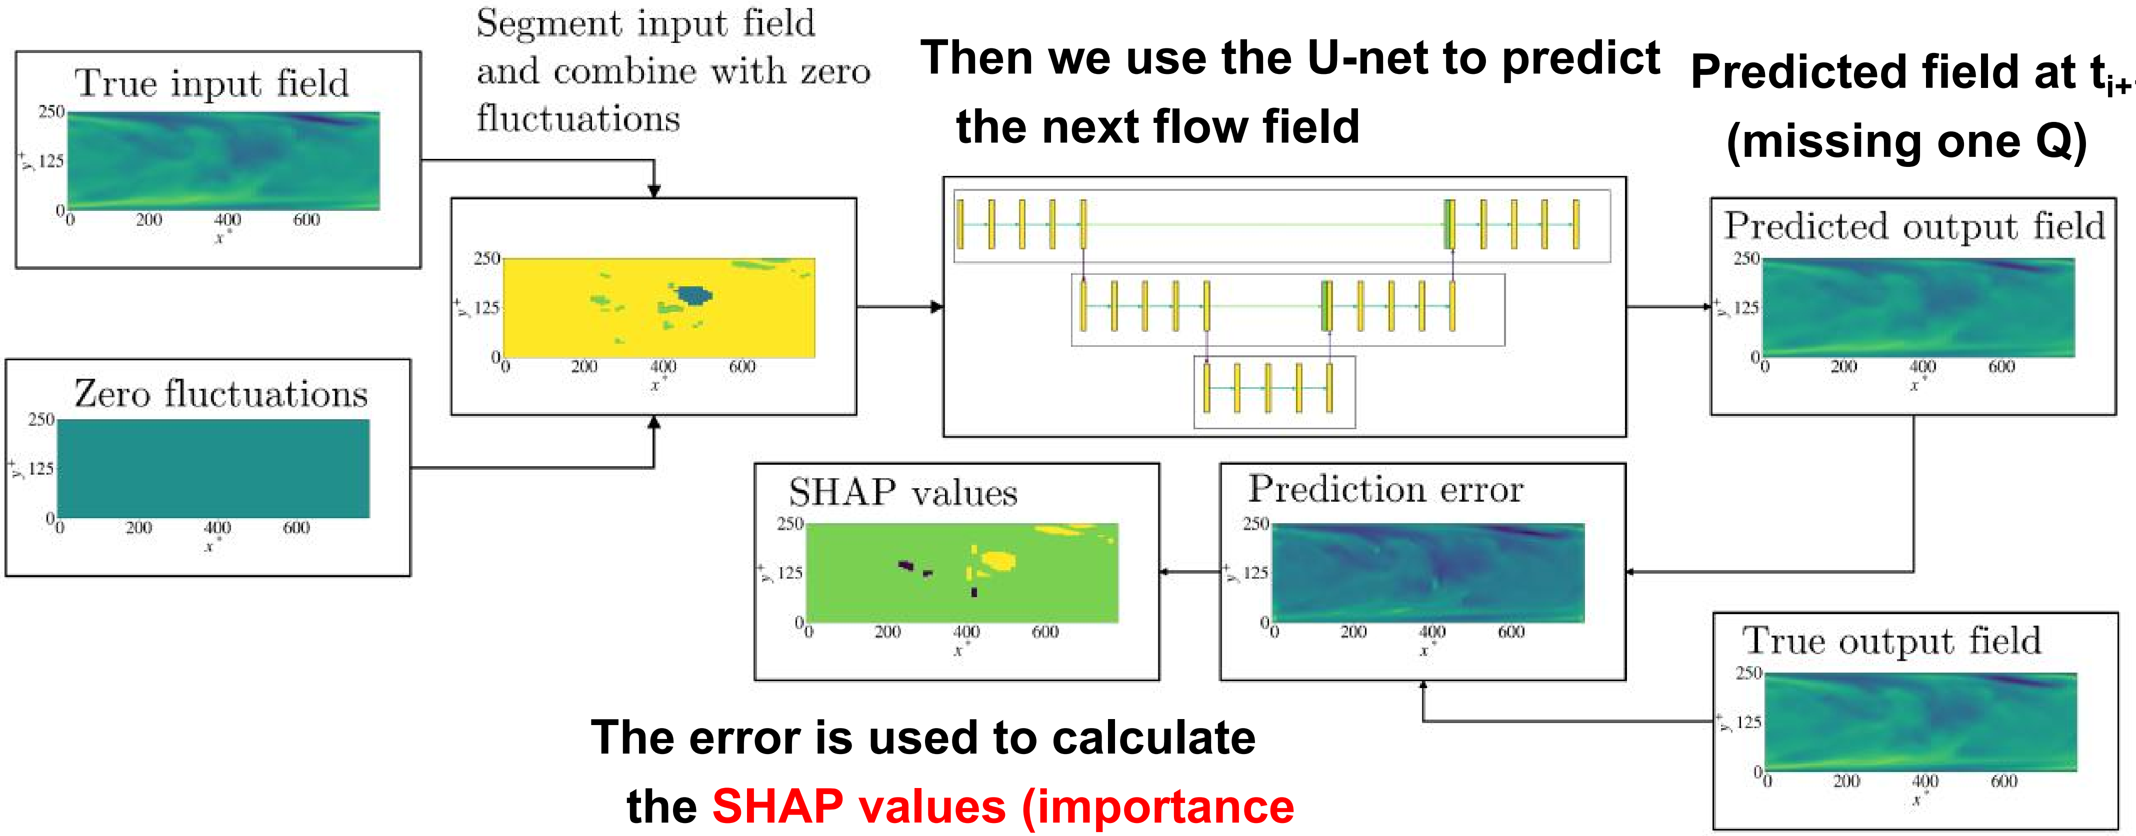

The error is used to calculate the **SHAP values (importance scores)**

Cremades et al., Nature Communications 15, 3864 (2024)

Applying SHAP point by point to 3D DNS data. Comparison with other structures

Re_τ=550

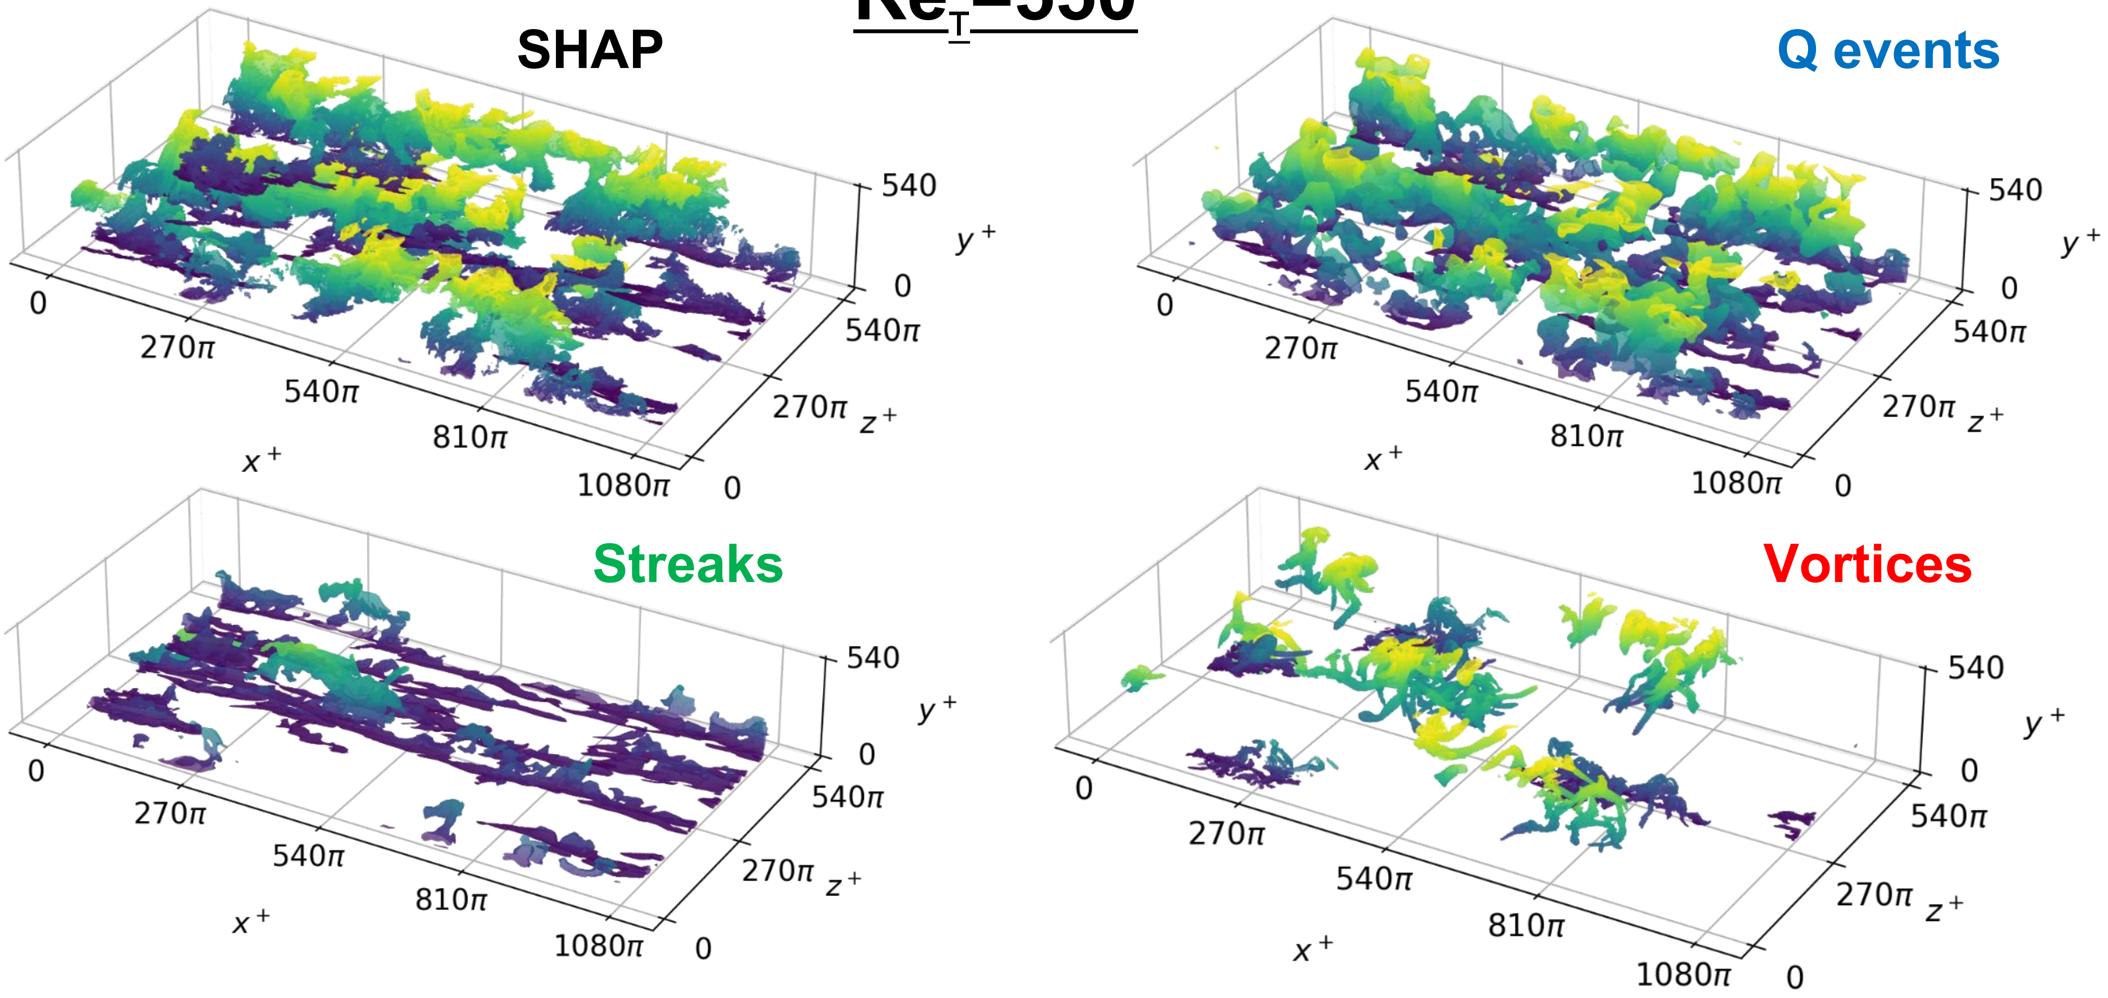

Cremades, Hoyas, Vinuesa (2024), Preprint arXiv:2410.23189v1
(Under Review in Nat. Commun.)

SHAP point by point in the 3D DNS data: u'. Similarities with other structures

Re$_\tau$=550

SHAP

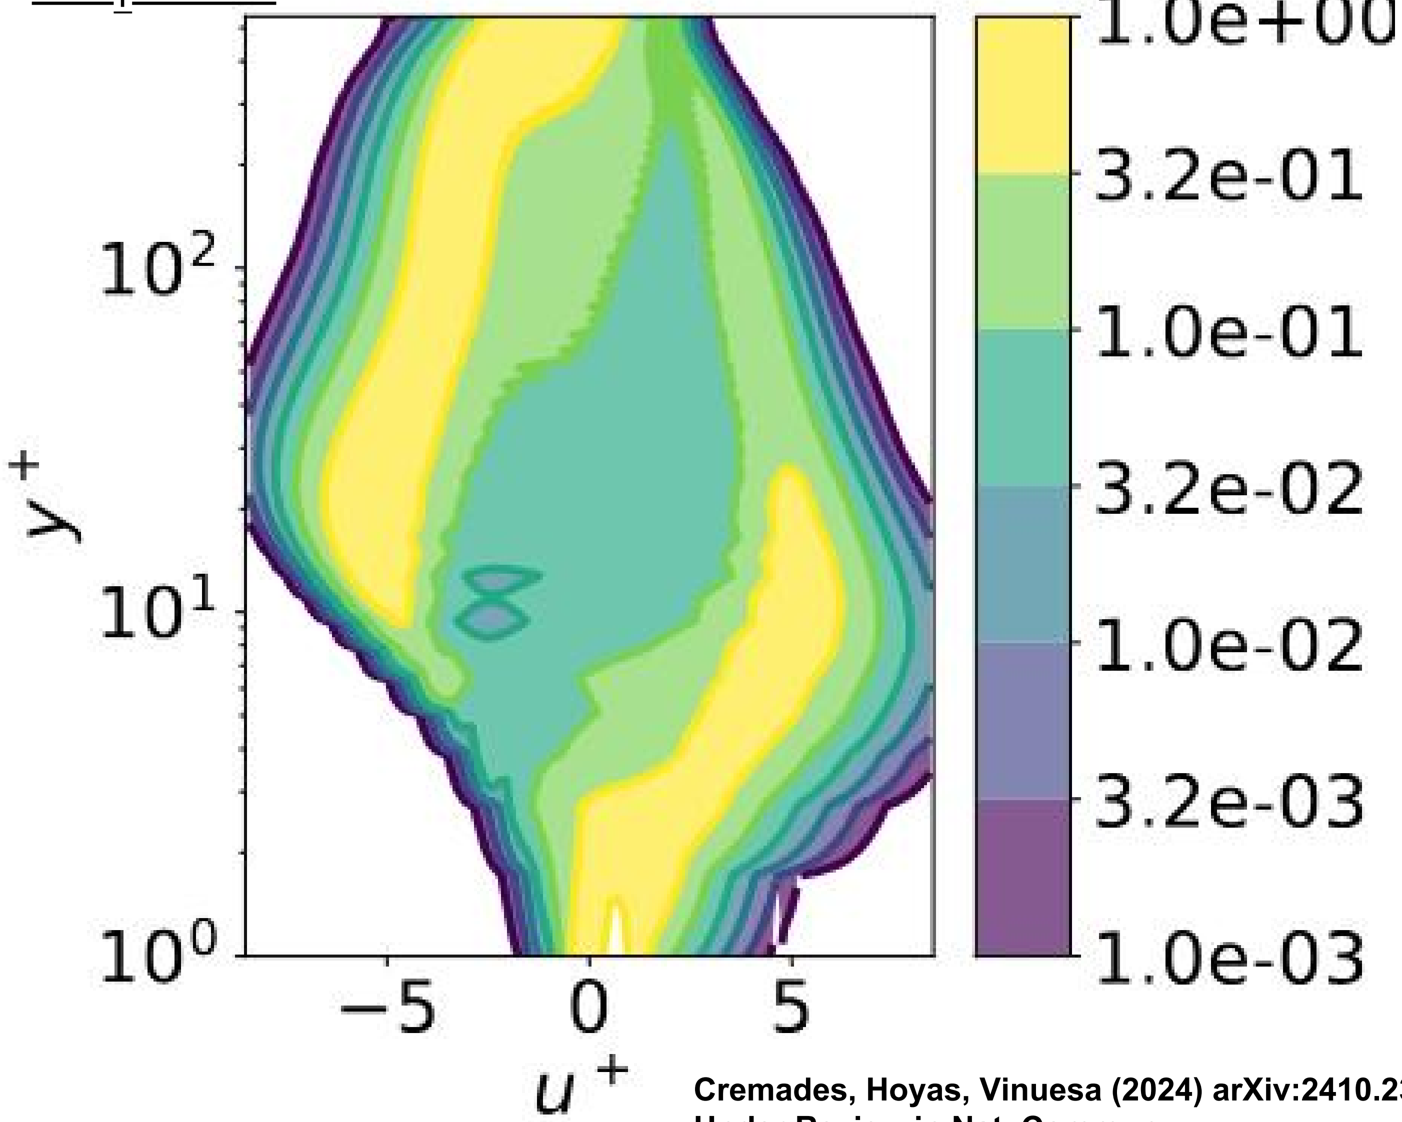

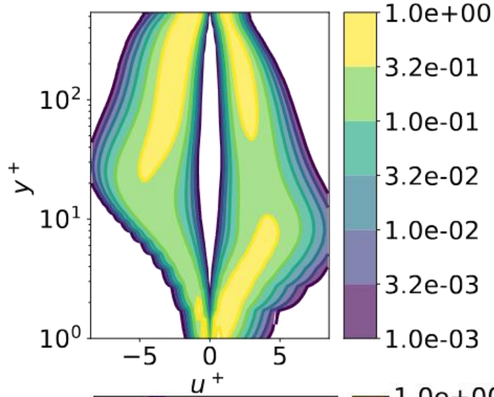

Q events

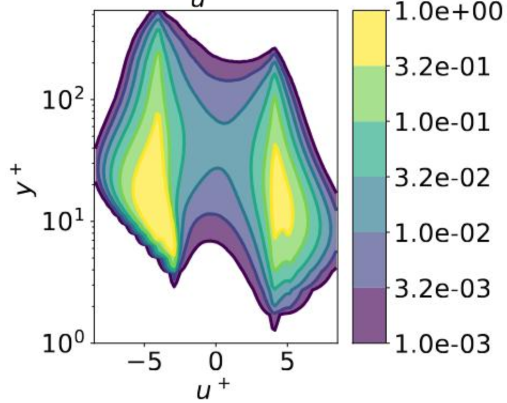

Streaks

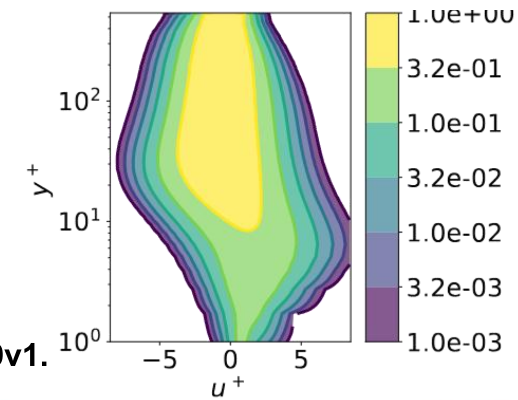

Vortices

Cremades, Hoyas, Vinuesa (2024) arXiv:2410.23189v1. Under Review in Nat. Commun.

SHAP point by point in the 3D DNS data: u'. Similarities with other structures

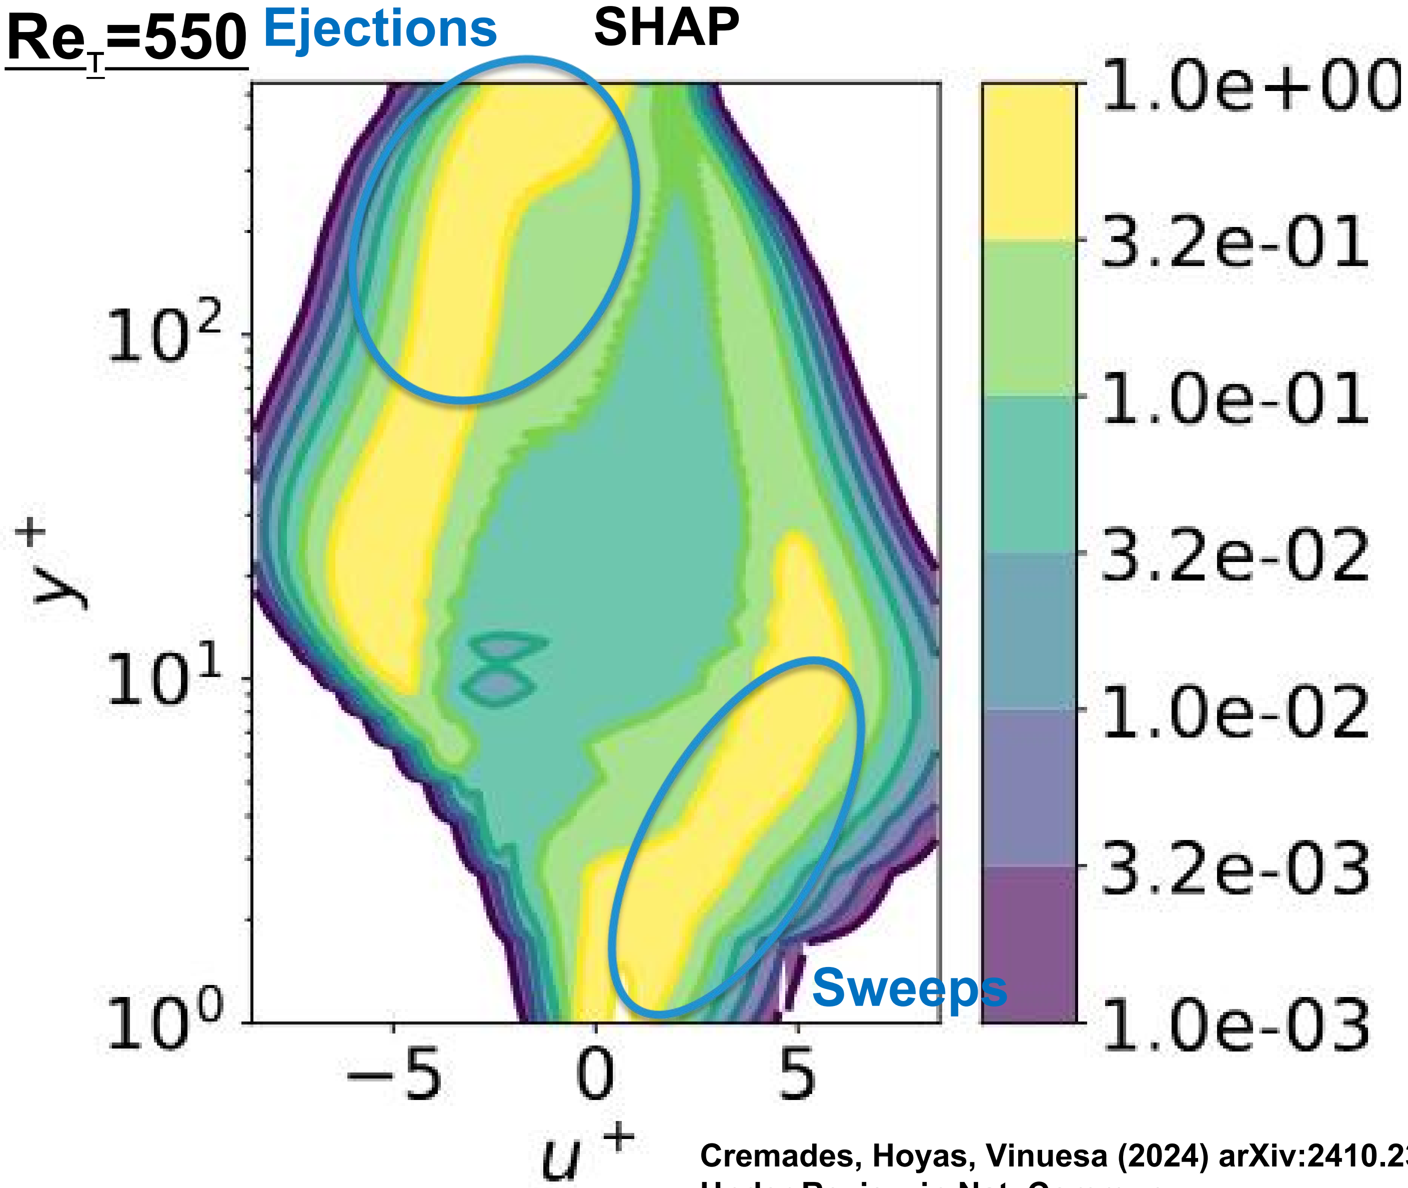

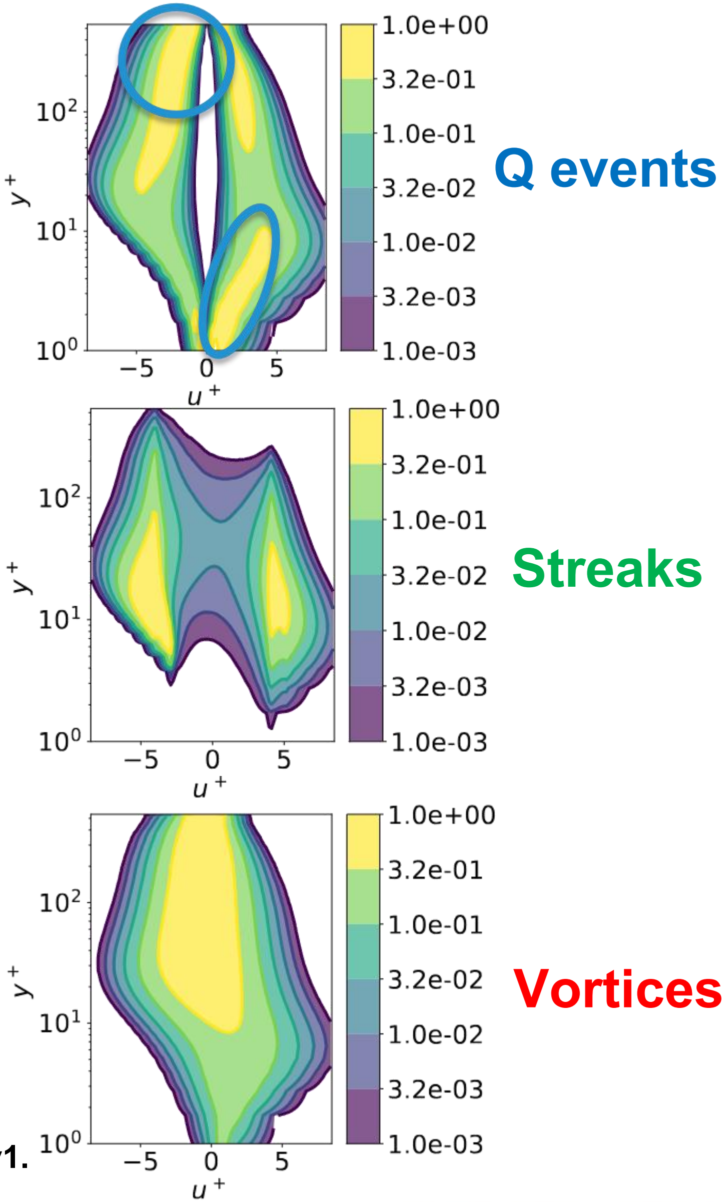

Cremades, Hoyas, Vinuesa (2024) arXiv:2410.23189v1.
Under Review in Nat. Commun.

SHAP point by point in the 3D DNS data: u'. Similarities with other structures

Re_τ=550

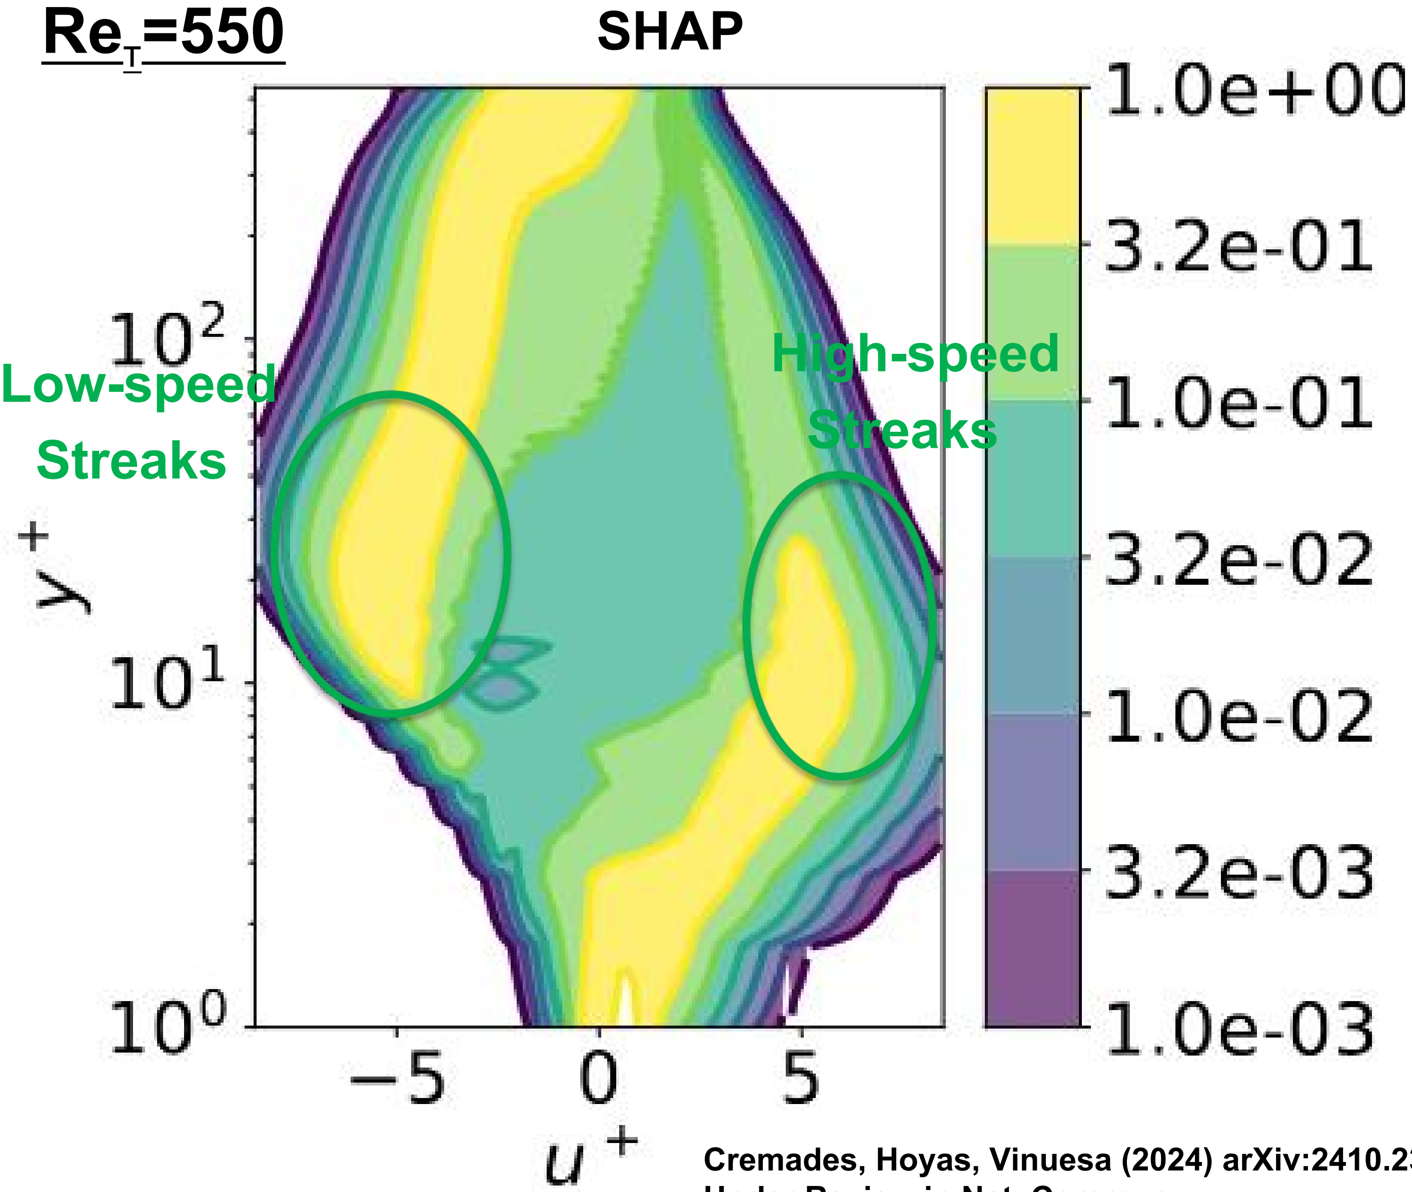

Cremades, Hoyas, Vinuesa (2024) arXiv:2410.23189v1. Under Review in Nat. Commun.

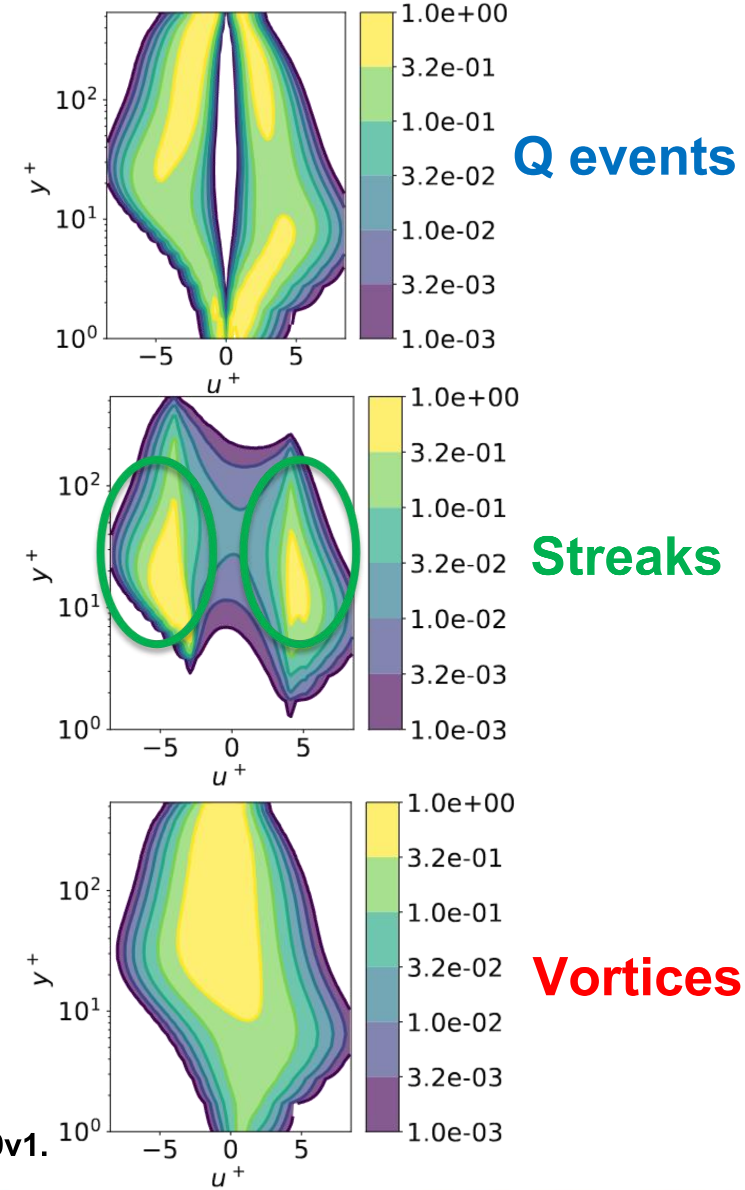

SHAP point by point in the 3D DNS data: u'. Similarities with other structures

Re_τ=550

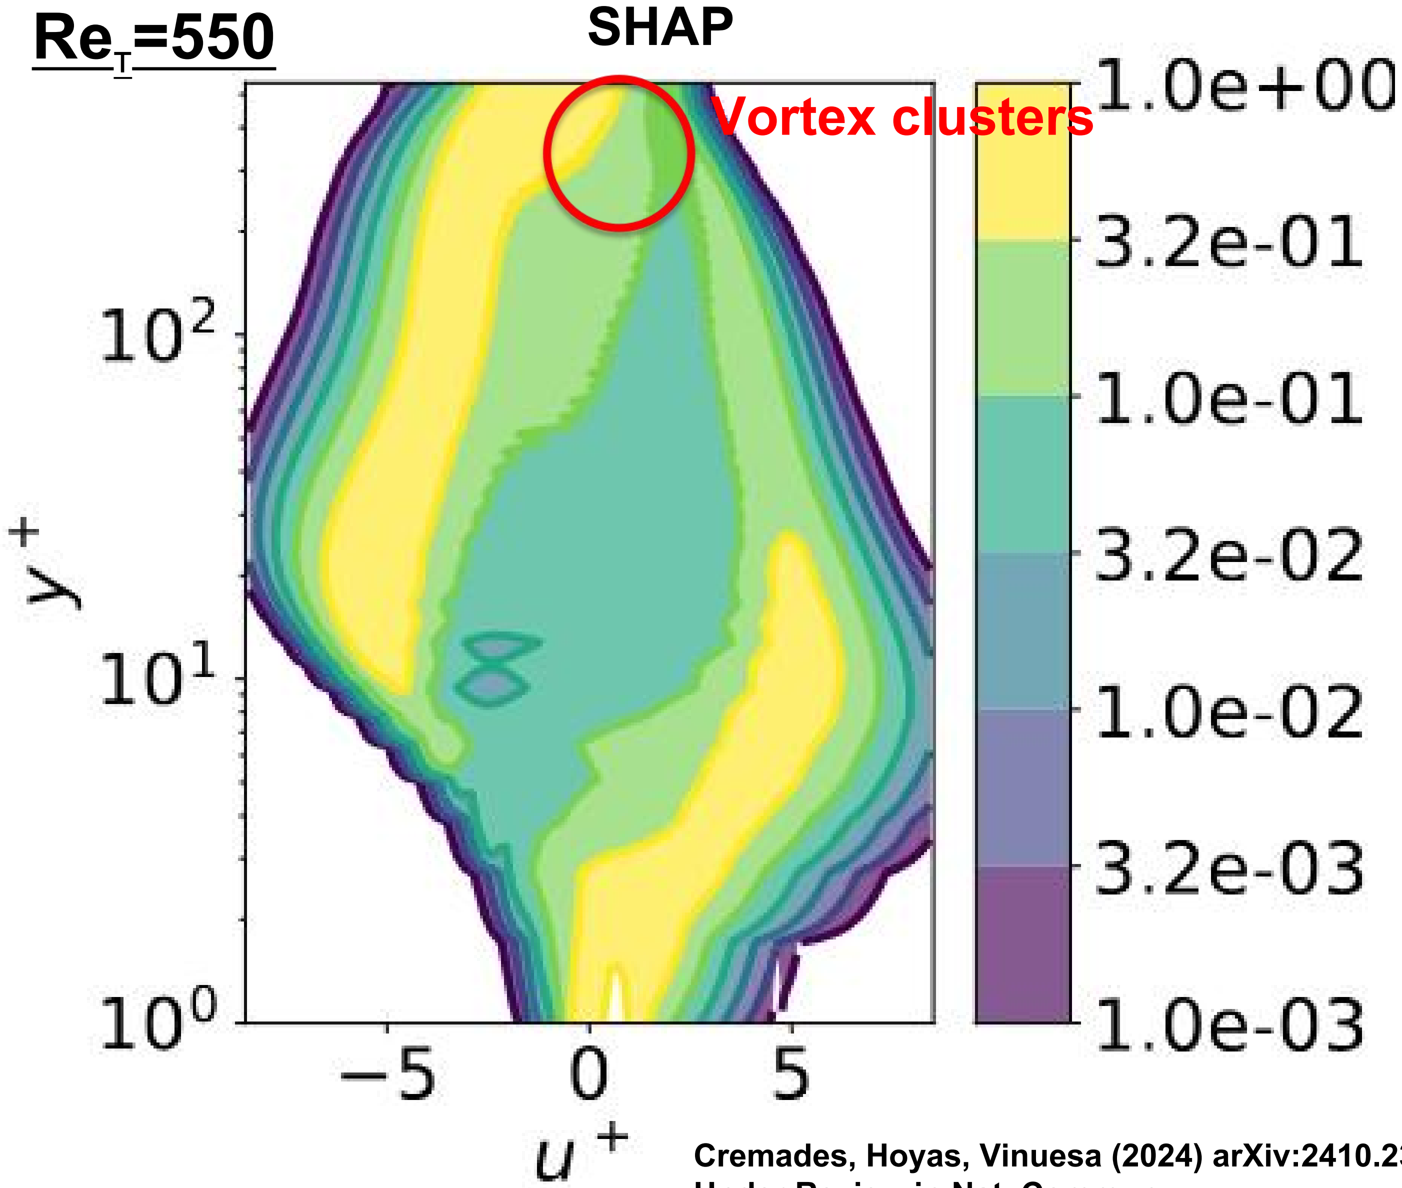

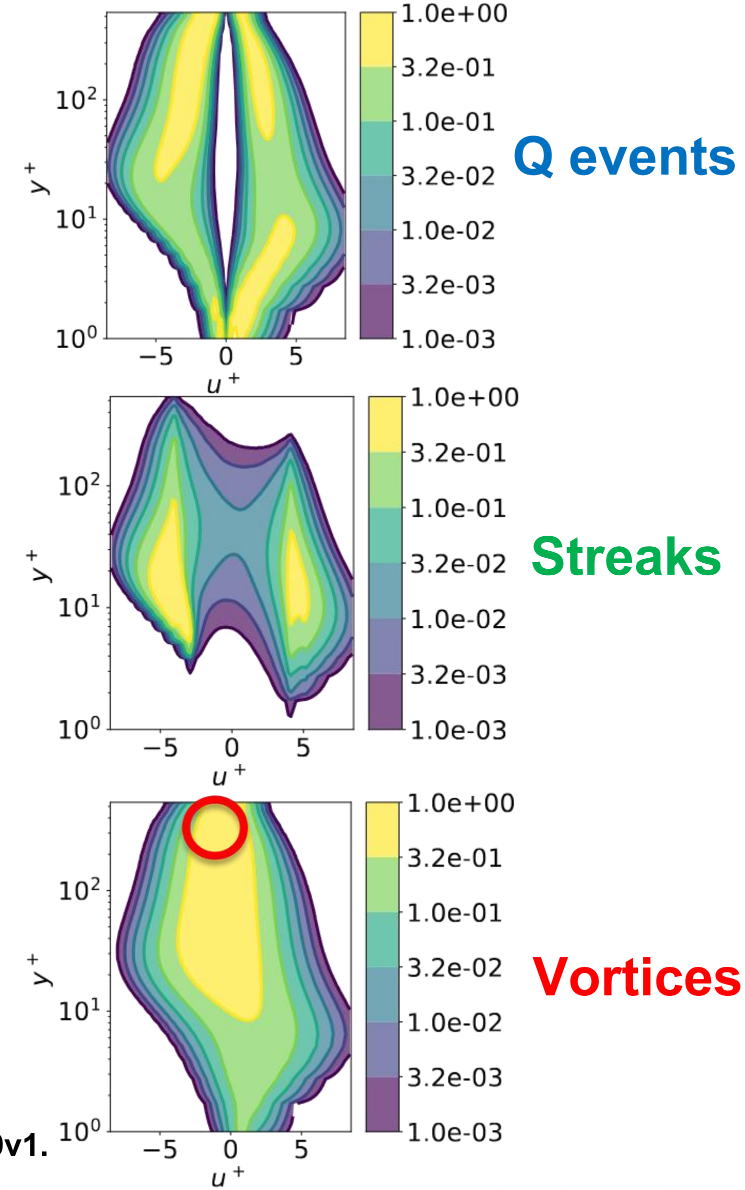

Cremades, Hoyas, Vinuesa (2024) arXiv:2410.23189v1.
Under Review in Nat. Commun.

One-to-one comparison between SHAP and other structures

<u>Re_τ=550</u>

- At **y^+=15**, the **SHAP structures are basically streaks (around 90% coincidence).**
- Close to the **wall and channel center**, the **agreement between SHAP and Q events is around 50% (ejections and sweeps).**
- At the **channel center**, the SHAP structures exhibit a **modest agremeent (around 15%) with the vortices.**

–The classically studied coherent structures only paint a partial picture of wall-bounded turbulence!

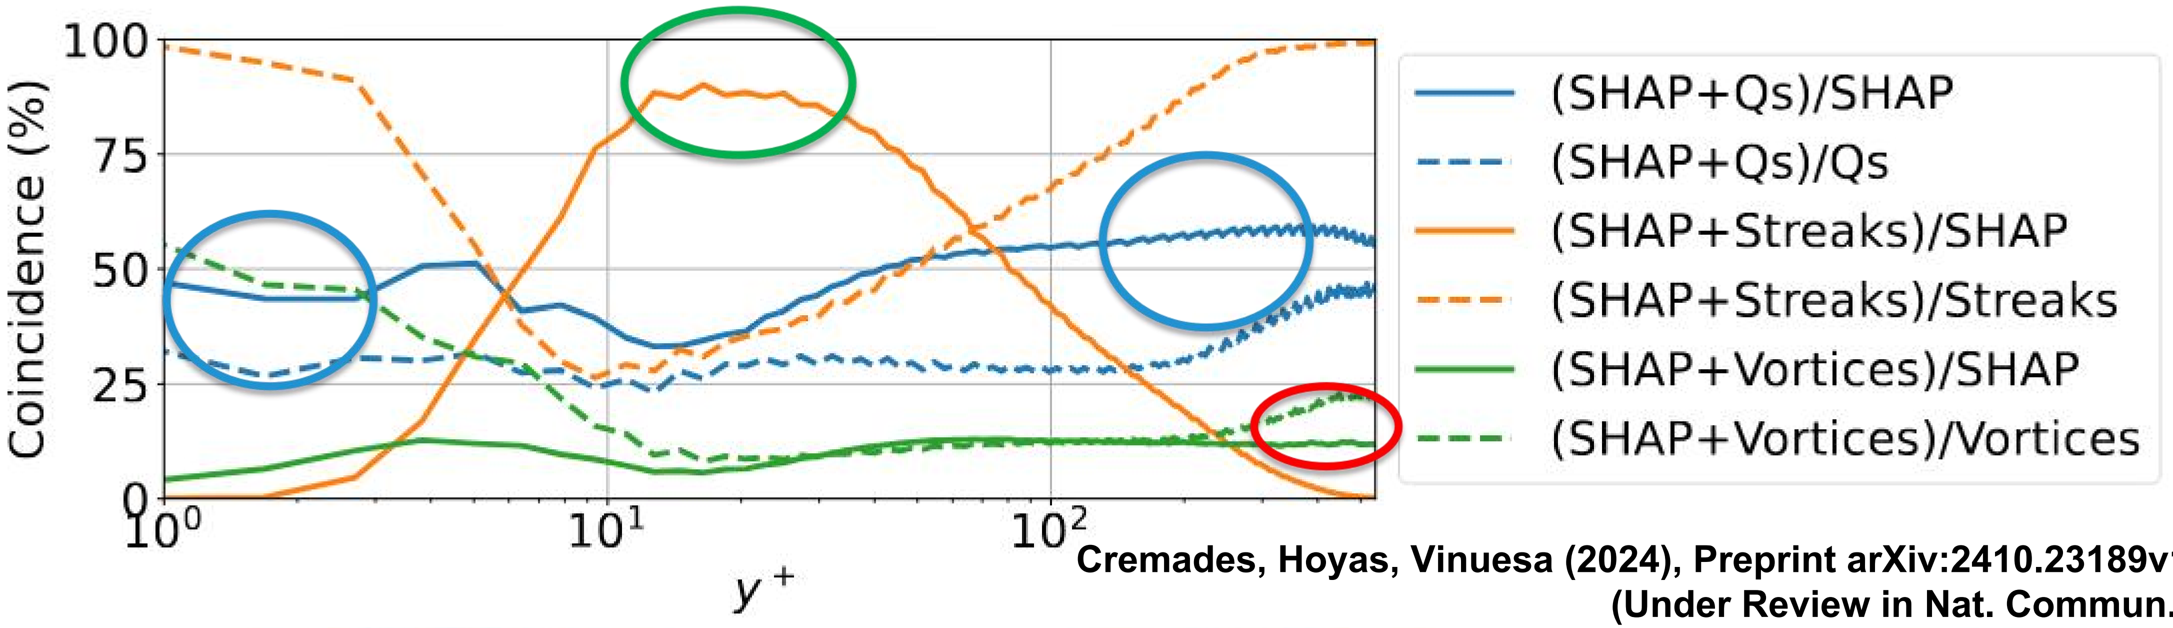

Cremades, Hoyas, Vinuesa (2024), Preprint arXiv:2410.23189v1 (Under Review in Nat. Commun.)

Outline

- Identify the most important regions with explainable deep learning (XDL).
- Remove the most important regions with **deep reinforcement learning (DRL).**

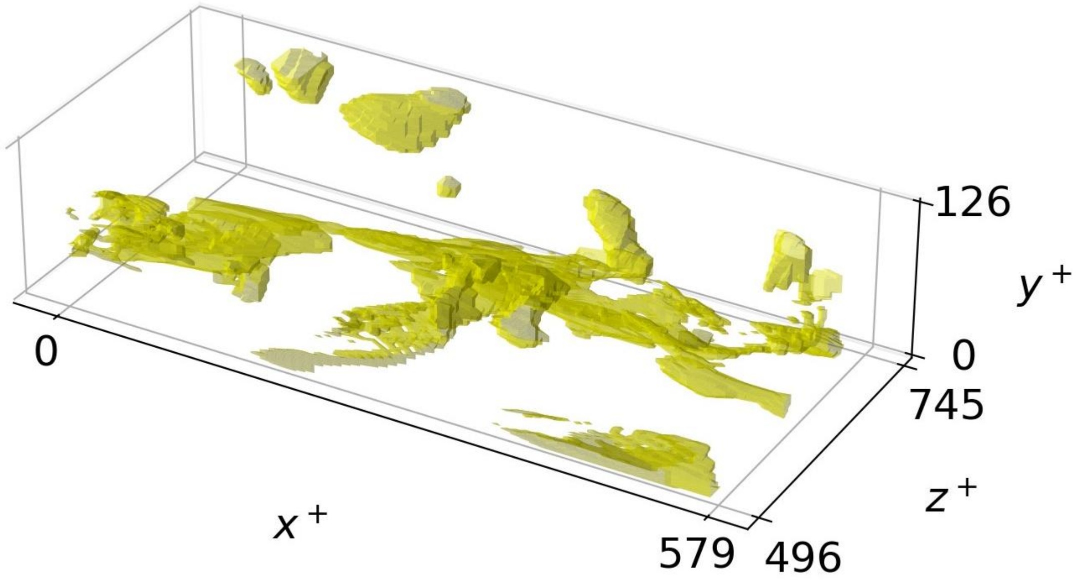

Reinforcement learning for flow control
Introduction

- Constituting elements:
  - Agent
  - Environment
  - State space $S = \{s_1, s_2, \cdots, s_N\}$
  - Action space $A = \{a_1, a_2, \cdots, a_M\}$
  - Transition function $\mathbb{P}(s_{t+1}|s_t, a_t)$
  - Reward function $r_t = R(s_t, a_t, s_{t+1})$
- Goal:

  "Define a policy $\pi(a_t|s_t)$ that maximizes the reward"

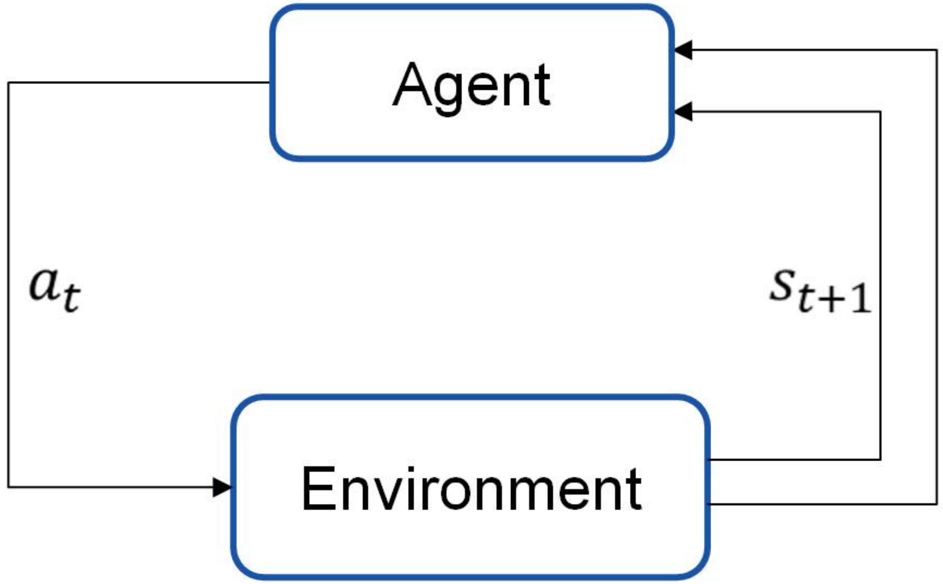

(1)

1. Sutton, R.S. Learning to predict by the methods of temporal differences. Mach Learn **3**, 9–44 (1988)

Guastoni et al., Eur. Phys. J. E, 46, 27 (2023)
Vignon et al., Phys. Fluids, 35, 031301 (2023)

Baseline: opposition control

- Introduced by Choi *et al.*[1]
- Reactive control to reduce the wall-normal fluctuations

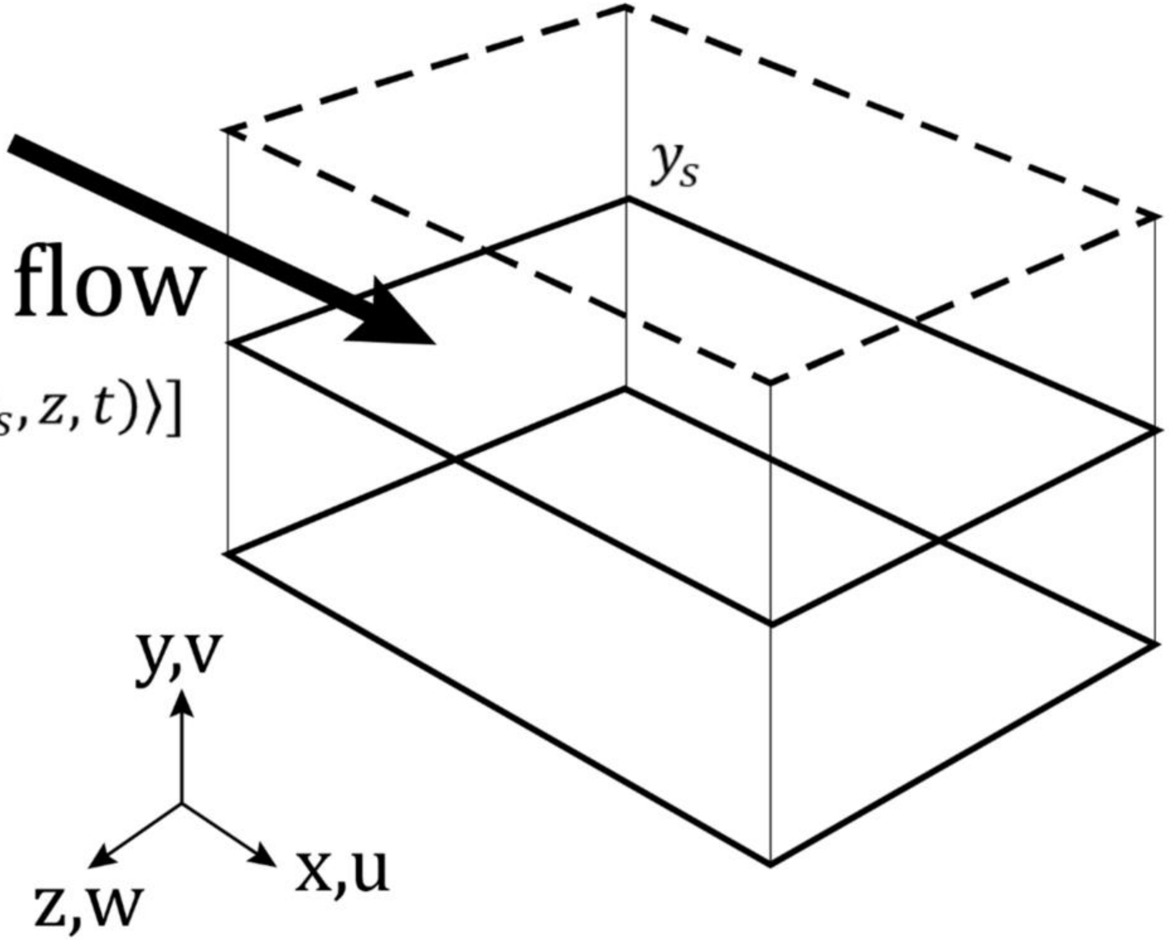

$$v_{wall}(x,z,t) = -\alpha[v(x,y_s,z,t) - \langle v(x,y_s,z,t)\rangle]$$

- y_s wall-normal location of the sensing-plane
  α positive scaling parameter

1. Choi, H., Moin, P., & Kim, J. (1994). Active turbulence control for drag reduction in wall-bounded flows. *Journal of Fluid Mechanics, 262*, 75-110

Guastoni et al., Eur. Phys. J. E, 46, 27 (2023)

DRL control targeting SHAP structures

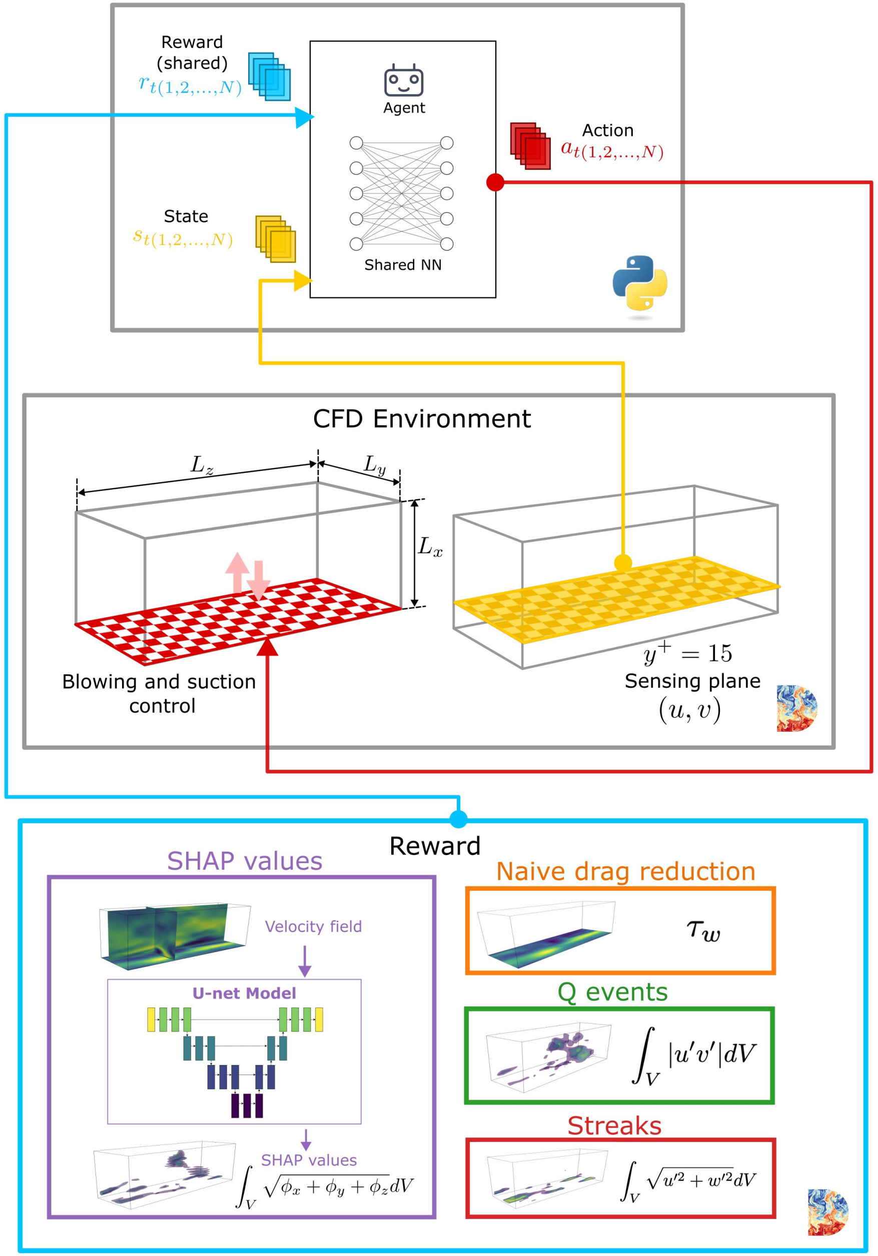

Comparison of control cases

- Opposition control (baseline)
- DRL reward: wall-shear-stress reduction
- DRL reward: Q-event reduction
- DRL reward: streak reduction
- DRL reward: SHAP reduction

Beneitez et al., Preprint arXiv:2504.02354 (2025)

DRL control targeting SHAP structures
Drag reduction

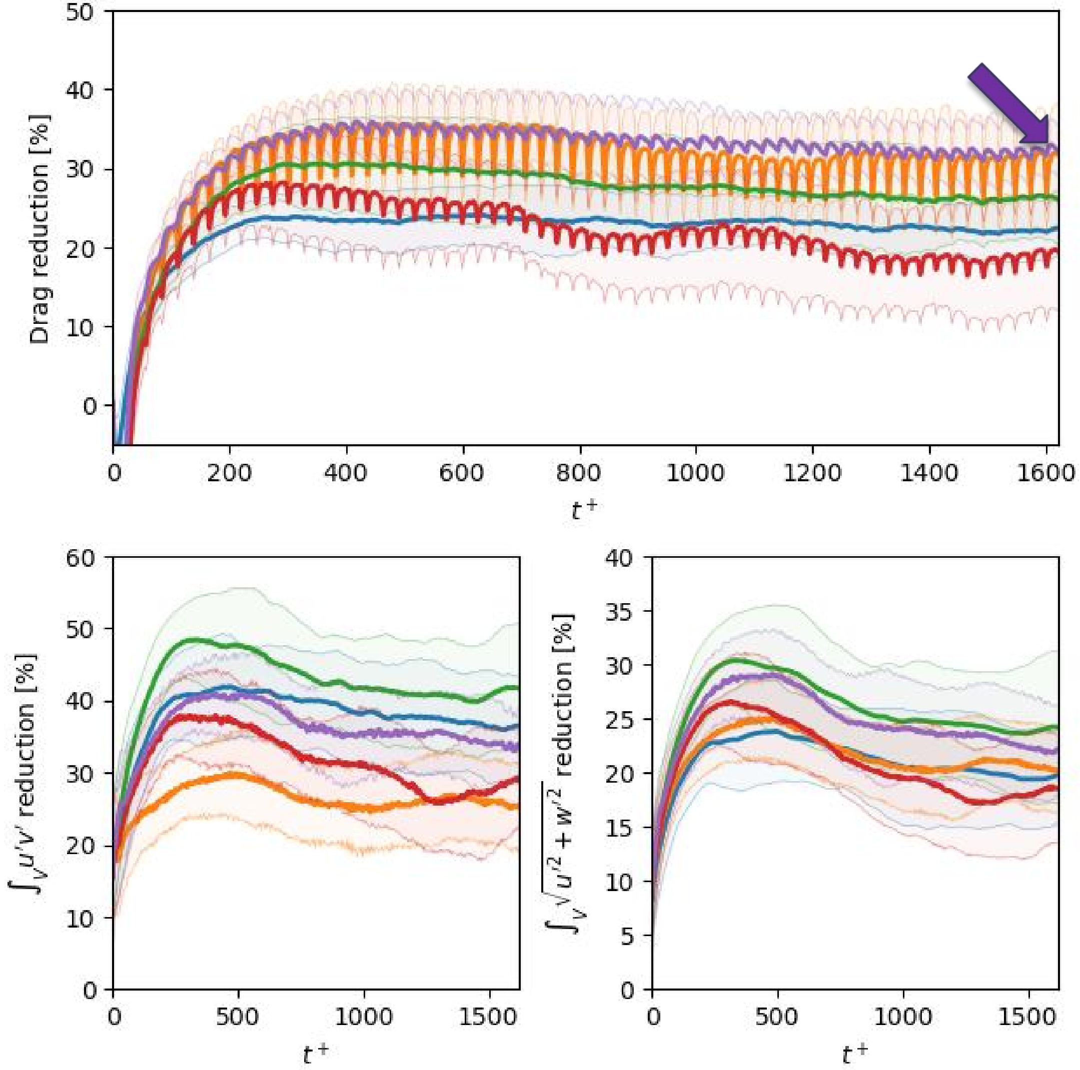

Comparison of control cases

- Opposition control (baseline): 23%
- DRL reward, wall-shear-stress reduction: 31.7%
- DRL reward, Q-event reduction: 27.5%
- DRL reward, streak reduction: 21.2%
- **DRL reward, SHAP reduction: 33.3%** (more stable than just DRL, 1.6 pp and 5% better).

Beneitez et al., Preprint arXiv:2504.02354 (2025)

DRL control targeting SHAP structures
Net energy saving

$$S = \frac{c_{f,\text{uncontrolled}} - (c_f + w_{\text{in}})}{c_{f,\text{uncontrolled}}}$$

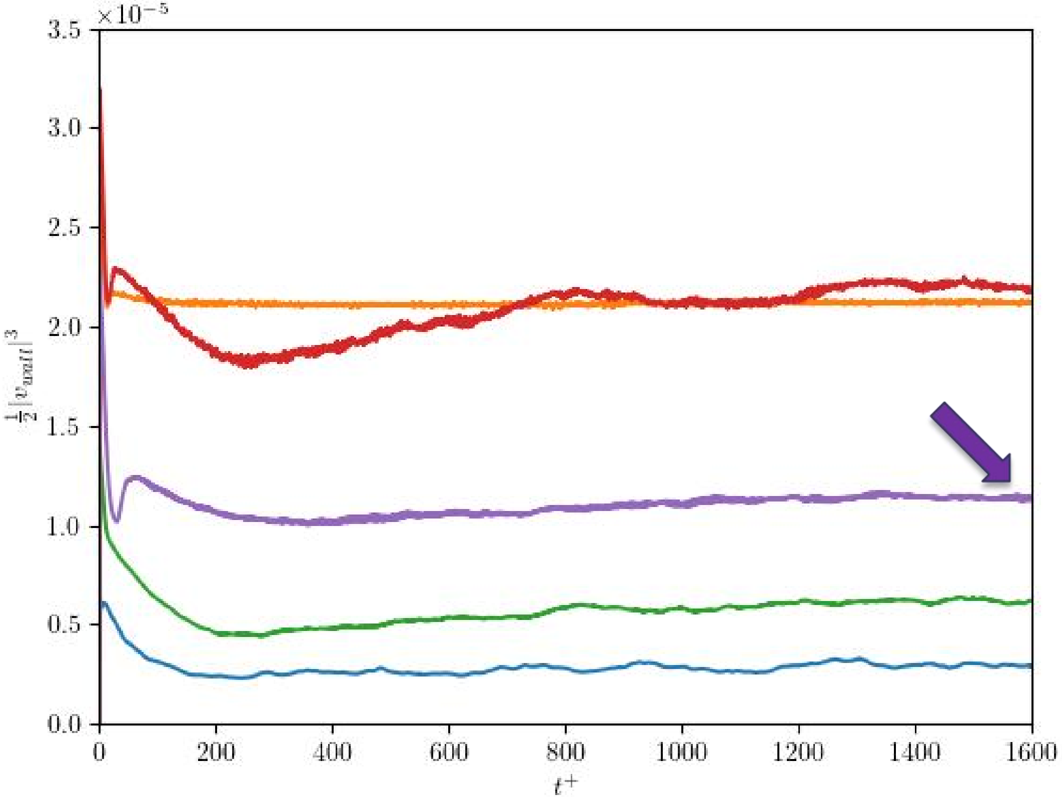

Comparison of control cases

- Opposition control (baseline): 22.96%
- DRL reward, wall-shear-stress reduction: 31.4%
- DRL reward, Q-event reduction: 27.47%
- DRL reward, streak reduction: 21%
- **DRL reward, SHAP reduction: 33.1%** (half the power than just DRL, 1.7 pp and 5.4% better).

FLOW **Beneitez et al., Preprint arXiv:2504.02354 (2025)**

Summary and Conclusions

- AI can help to achieve 79% of the SDG targets, but can be an inhibitor to 35%. **Very much needed global debate.**

- The **SHAP method** is intrusive for the surrogate, and can be applicable to data-scarce environments (e.g. experiments). **Identification of new coherent structures based on SHAP.**

- Using the **SHAP-based structures as the reward** yields the highest drag reduction through DRL!!

GitHub

Ricardo Vinuesa

@ricardovinuesa

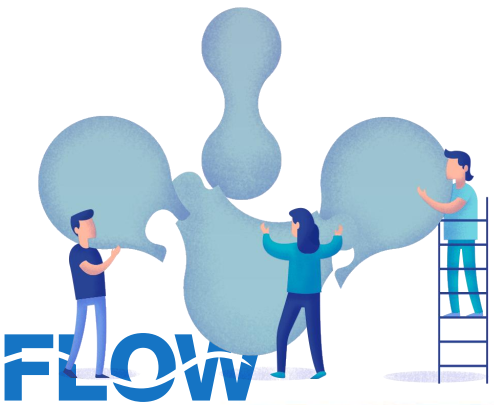

DRL control of skin friction in turbulent wings NACA4412 wing at Re_c=200,000, AoA=5 deg.

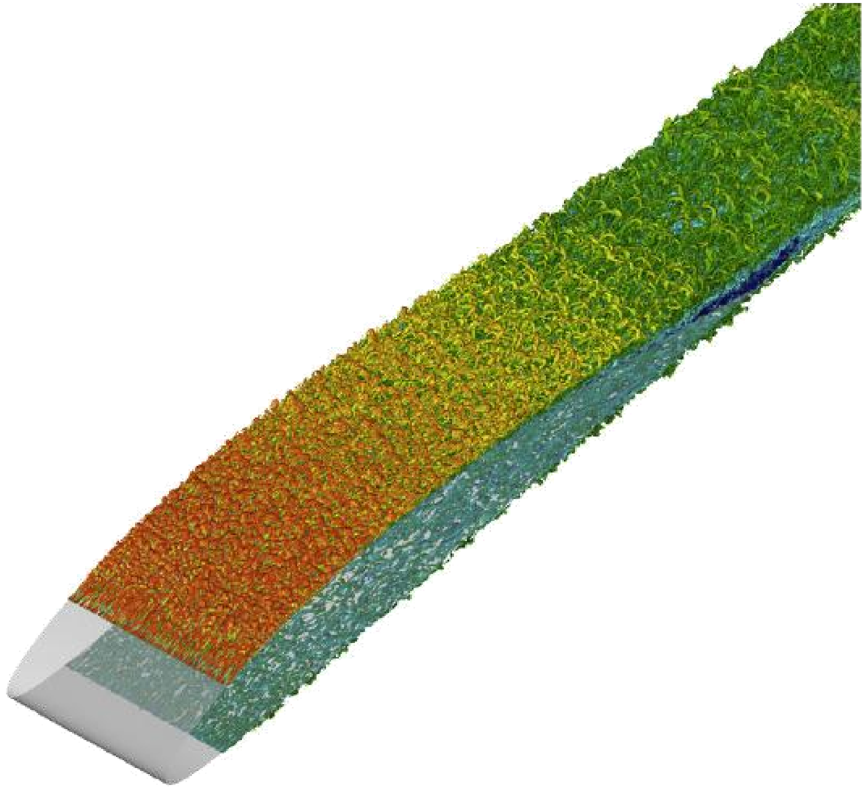

Vinuesa et al., IJHFF 72, 86 (2018)

Problems applying MARL to a turbulent wing section

- Only one periodic direction, very small number of pseudo environments for training (162).
- We need to simulate small scales in the leading edge with limit the time step.
- High-fidelity simulation of a wing is computationally expensive (65 million grid points).

Wang, Suárez and Vinuesa (2025)

DRL control of skin friction in turbulent wings NACA4412 wing at Re_c=200,000, AoA=5 deg.

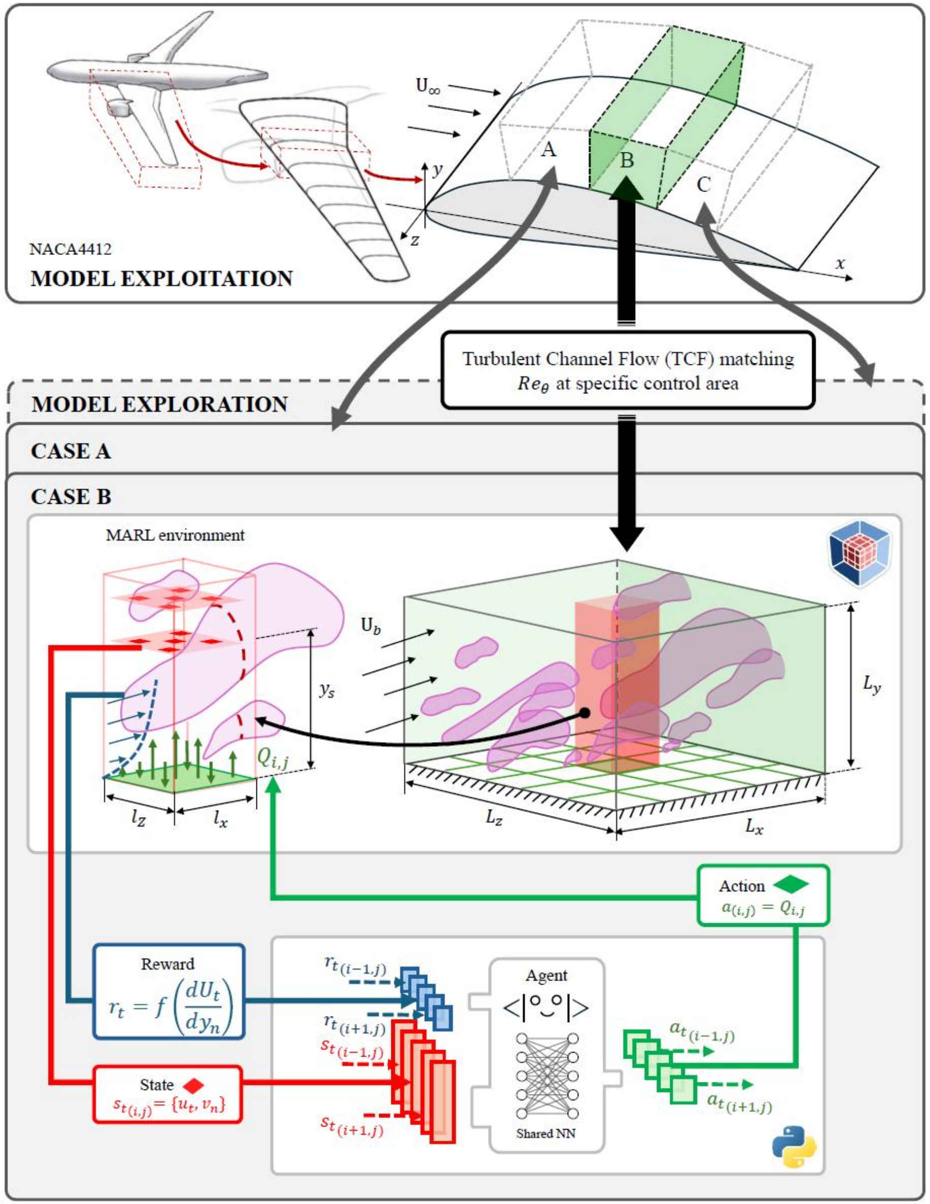

Use <u>channel flow</u> to train different wing areas (<u>matched Re_θ</u>)

- Increase number of **pseudo environments** from 162 (wing) to 3,240 (channel): **<u>Factor of 20</u>** <u>(two periodic directions and big domain</u>).
- Increase **time step** by a **<u>factor of 3</u>** (<u>we do not simulate small scales in the leading edge</u>).
- Wing has **<u>176 times more points</u>.**
- **Training in the channel: <u>60 times faster</u> and <u>~10,500 times cheaper</u>.**

Wang, Suárez and Vinuesa (2025)

DRL control of skin friction in turbulent wings NACA4412 wing at Re_c=200,000, AoA=5 deg.

- Implementation of control in two regions: x/c from 0.25 to 0.4 and from 0.4 to 0.5.
- Currently implementing several **simultaneous blocks. DRL 50% better than OC!**

DRL higher drag reduction (48.7%) than **opposition control (29.8%)**

DRL higher drag reduction (32.3%) than **opposition control (22.2%)**

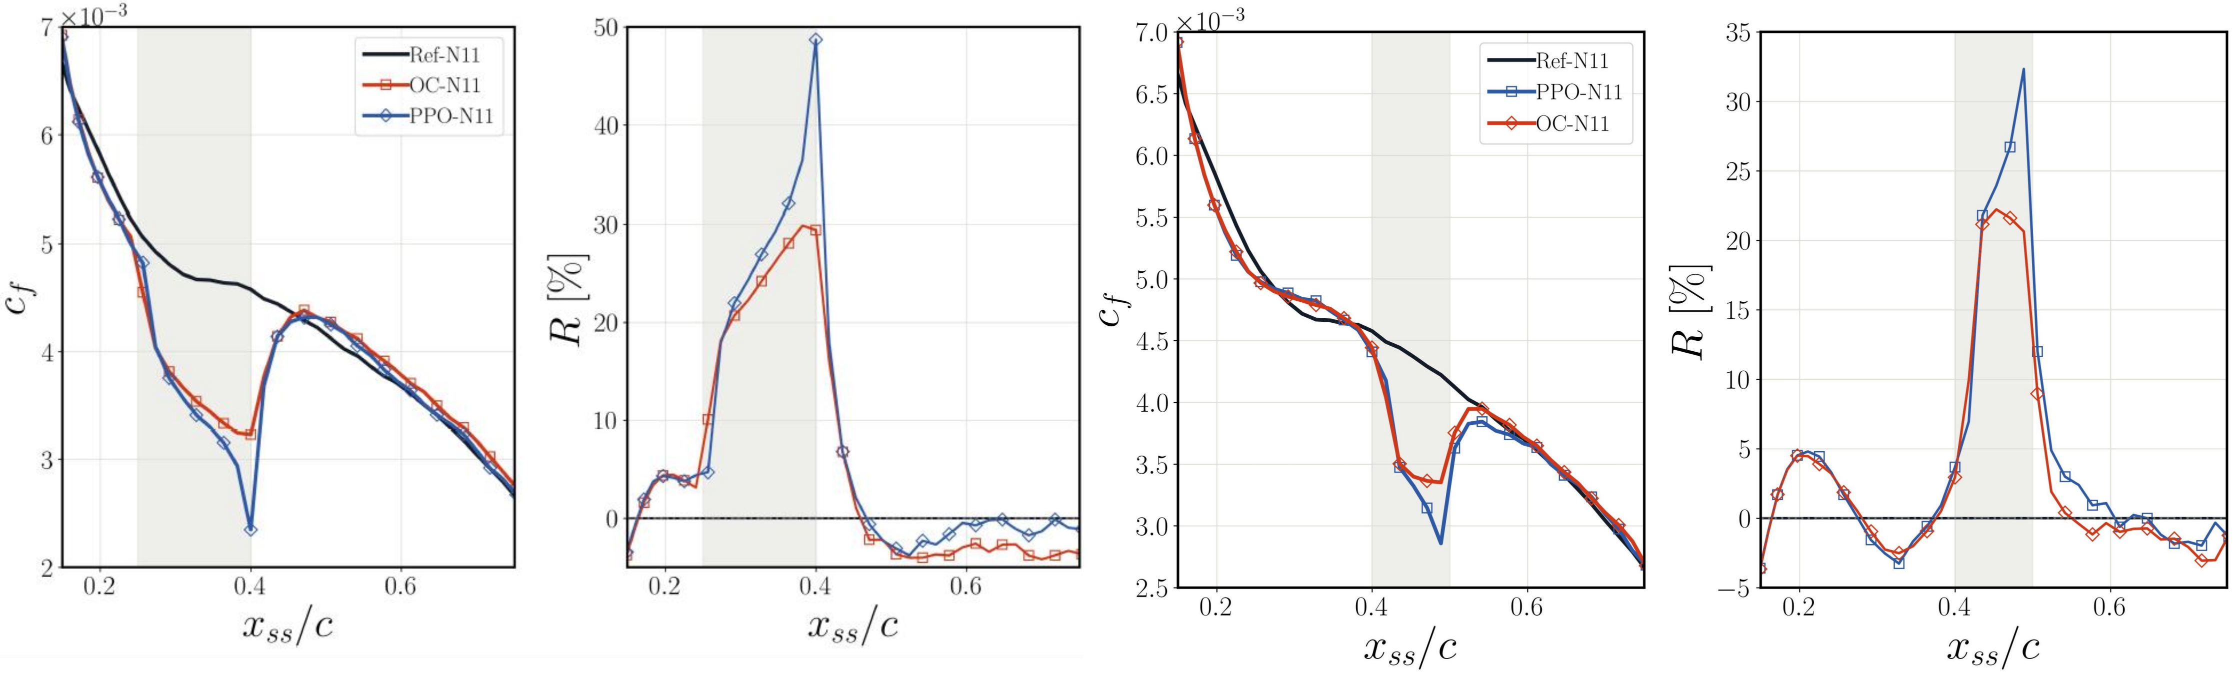

Wang, Suárez and Vinuesa (2025)

Prediction results

- Very accurate predictions of the instantaneous velocity fluctuations (periodic padding).
  - **Mean relative error is around 2%. Example u':**

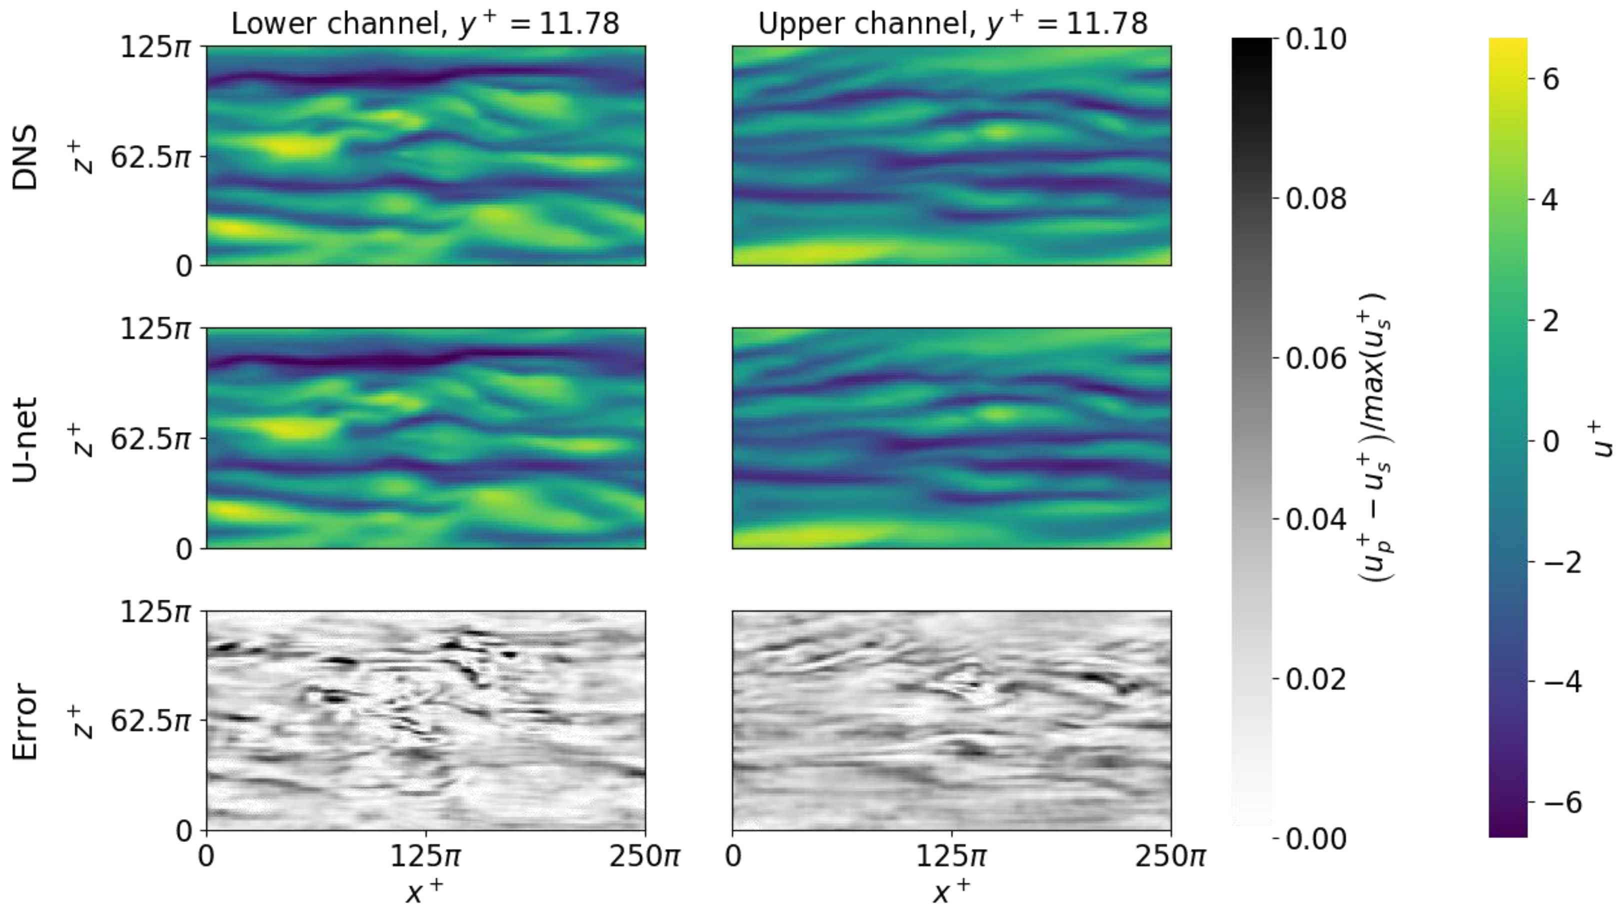

Cremades et al., Nature Communications 15, 3864 (2024)

The two main assumptions in the kernel-SHAP method

1. **<u>The model g is linear.</u>** Linear function to represent the error between the predicted and the true fields. In practice, there are so many structures (features), that the **difference** between the true error (f) and the linear model (g) is **very small: $(f\text{-}g)^2 \sim 10^{-7}$**.
2. **<u>Contribution of structures to error.</u>** Although the contribution to the error is calculated for 1 structure, and turbulence is highly chaotic, in the kernel-SHAP method the contribution is taken for a number of **coalitions**, which are groups of structures. This accounts for different **inter-structure interactions** when computing the contribution of the structure to the error. **Weighted average of the contribution to the error.**

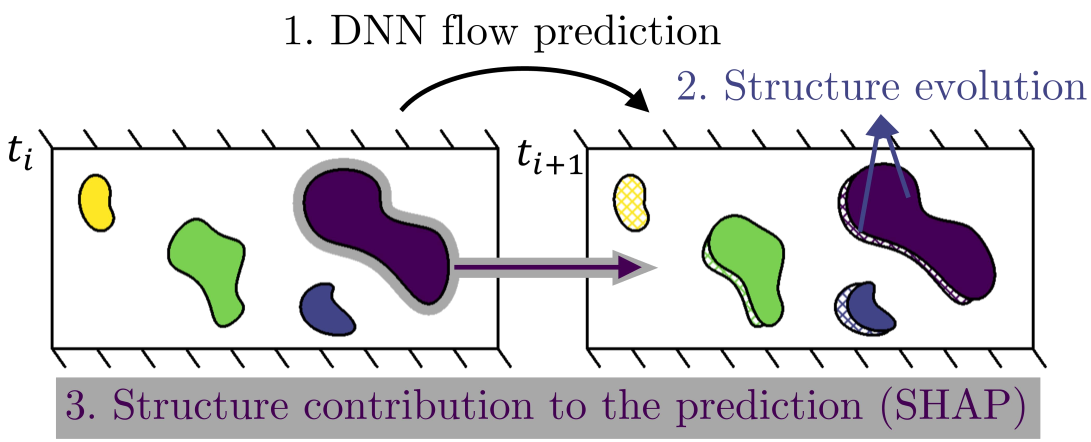

Cremades et al., Nature Communications 15, 3864 (2024)

MARL: Control and learning parameters

Control

- Actuation interval $\Delta t^+ \approx 0.6$
- Linear transition from the old value to the new one
- Limited actuation intensity: $[-u_\tau, u_\tau]$
- All actuators use the same policy:

$$\pi(s) = \pi(s|\boldsymbol{\theta})$$

Policy-learning

- Deep deterministic policy gradient (DDPG)[1] model-free off-policy actor–critic
- Policy gradient update $\Delta t^+ = 6$
- 64 minibatch gradient-based updates
- 1000 actuations per episode

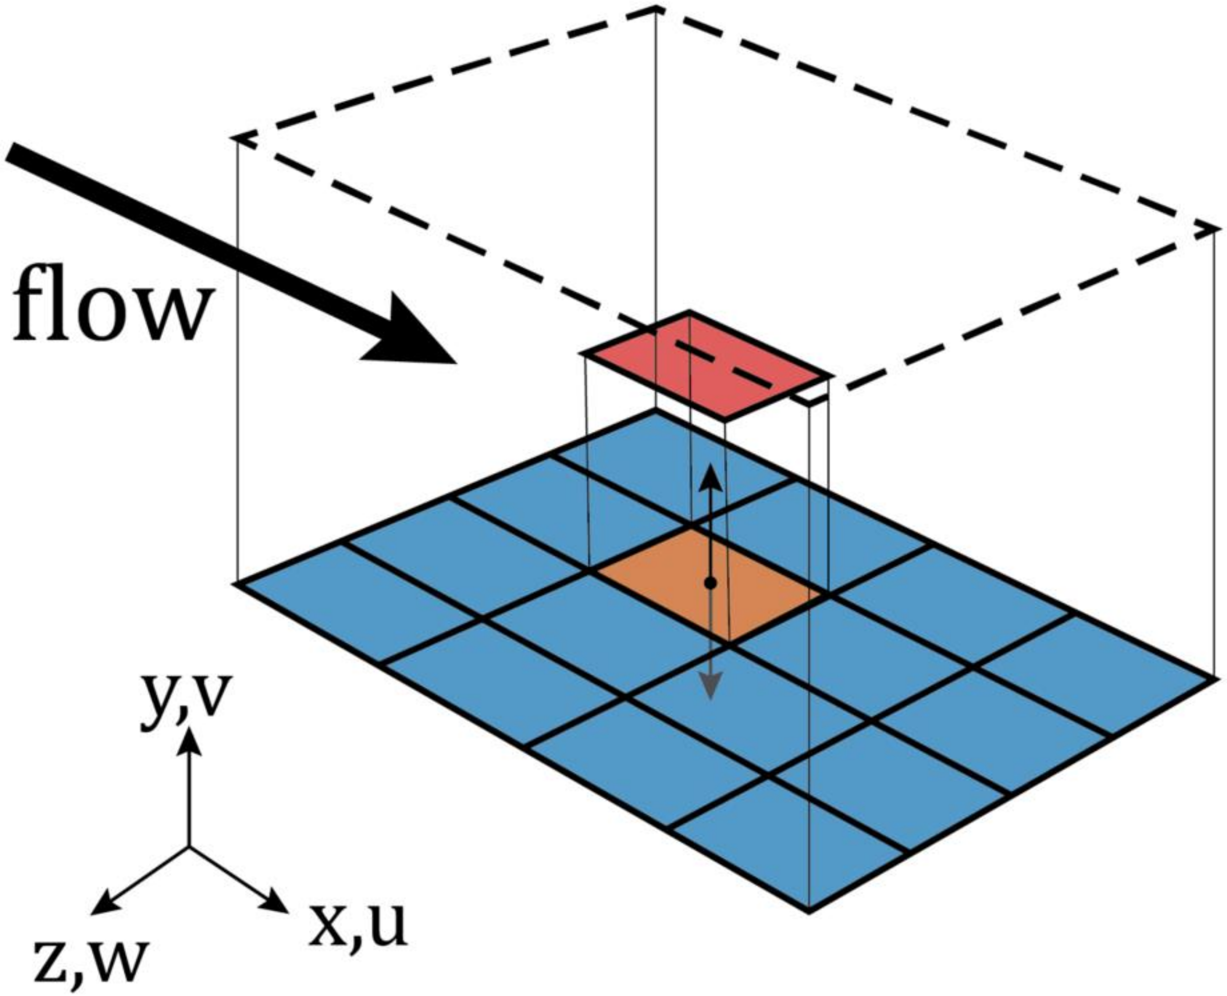

1. *T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra (2015). Continuous control with deep reinforcement learning , arXiv:1509.02971*

Guastoni et al., Eur. Phys. J. E, 46, 27 (2023)

Estimation the cost of the actuation

- How to calculate the net-energy saving?
  Referring to Kametani *et al.*[1], the input power is:

$$w_{in} = \frac{1}{2} \| v_{wall} \|^3$$

$\approx 10^{-2}$

- The net-energy saving will be:

$\approx 10^{-6}$

$$S = \frac{c_{f,\mathrm{uncontrolled}} - (c_f + w_{in})}{c_{f,\mathrm{uncontrolled}}}$$

- Note that:

$\approx 10^{-3}$

$$R = 1 - {}^{c_f}\!/_{c_{f,\mathrm{uncontrolled}}} = 1 - {}^{\tau_w}\!/_{\tau_{w,\mathrm{uncontrolled}}}$$

- Even if $w_{in,\mathrm{DRL}} > w_{in,\mathrm{Opp}}$, the actuation cost is negligible

Kametani et al., IJHFF, 55, 143-142 (2015)

Other applications: Rayleigh–Bénard convection

Objective: Reduce Nusselt number, i.e. ratio of convective to conductive heat transfer between two plates (bottom hot, top cold).

- Multi-agent DRL has been successfully applied to Rayleigh-Bénard convection[1]

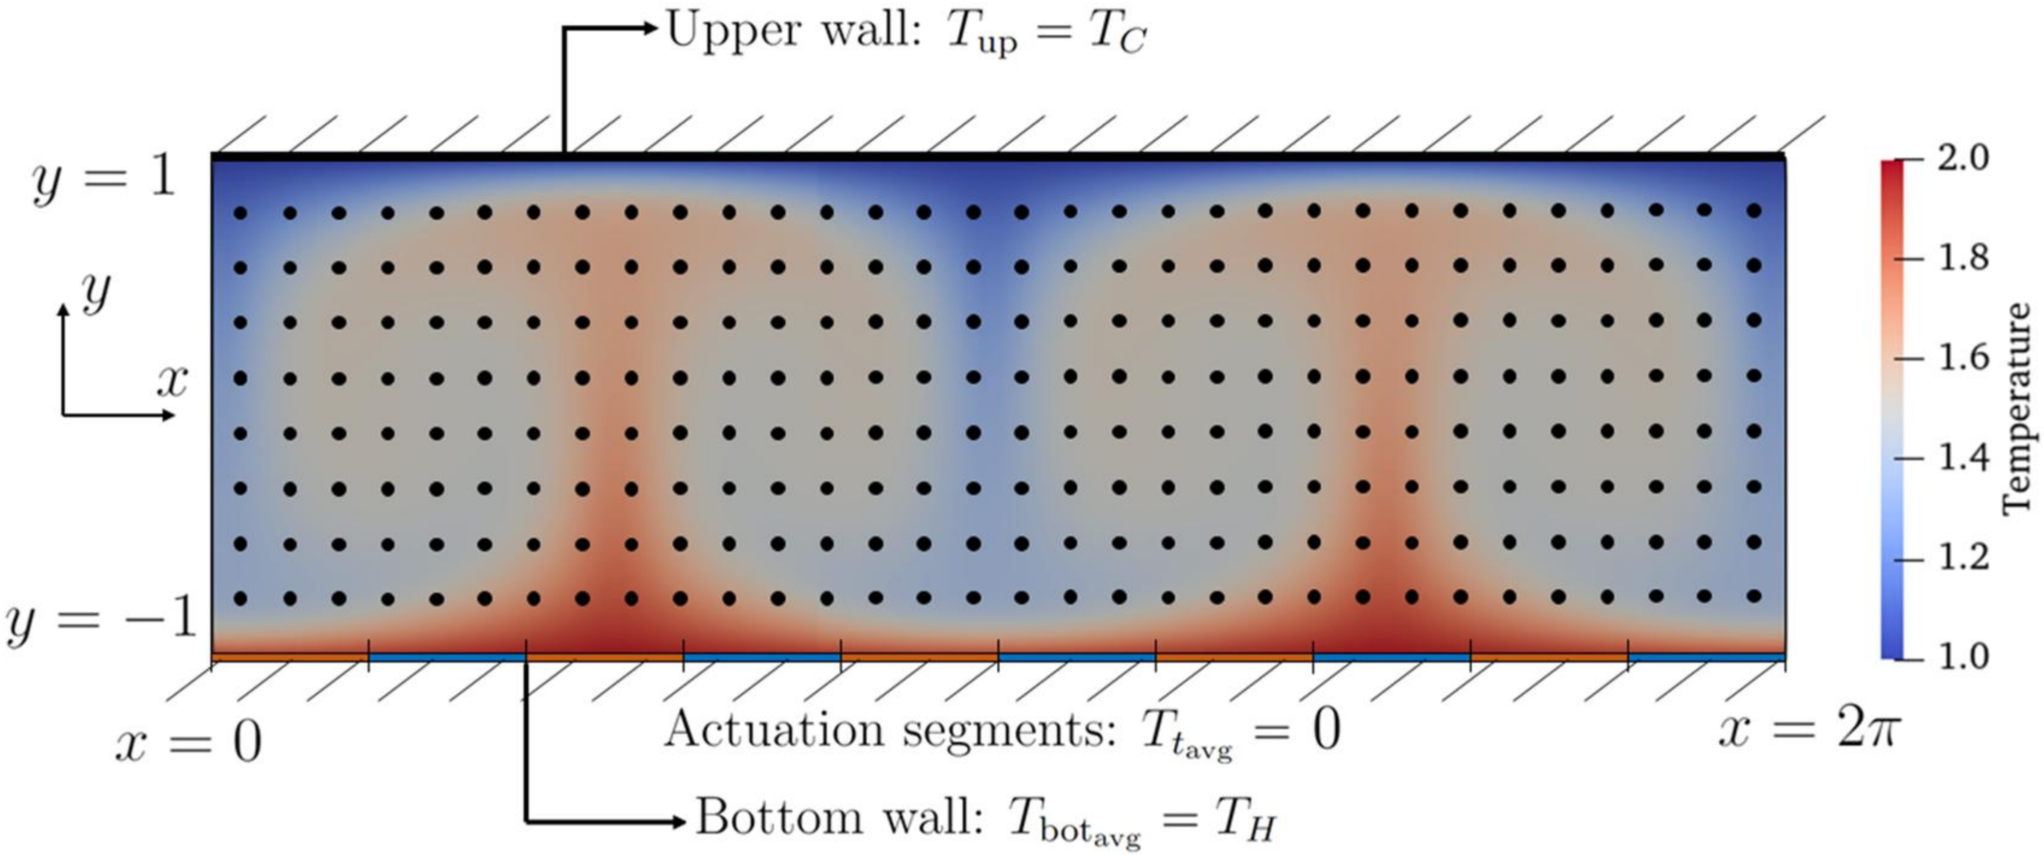

Vignon et al., Phys. Fluids 35, 065146 (2023)

Single- versus multi-agent reinforcement learning

Single-agent reinforcement learning

Multi-agent reinforcement learning

Vignon et al., Phys. Fluids 35, 065146 (2023)

Other applications: Rayleigh–Bénard convection

- Multi-agent DRL has been successfully applied to Rayleigh-Bénard convection[1]
- The objective is to reduce the Nusselt number Nu, starting from a given initial condition
- Single-agent RL only offers a limited reduction
- Multi-agent RL provides a more consistent reduction

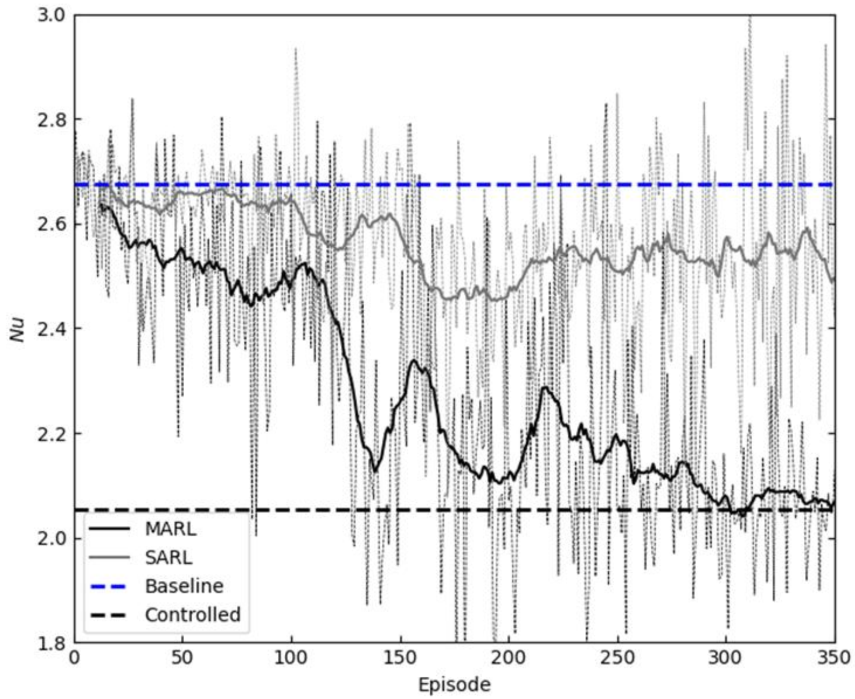

Vignon et al., Phys. Fluids 35, 065146 (2023)

The two main assumptions in the kernel-SHAP method

1. **<u>The model g is linear.</u>** Linear function to represent the error between the predicted and the true fields. In practice, there are so many structures (features), that the **difference** between the true error (f) and the linear model (g) is **very small: $(f\text{-}g)^2 \sim 10^{-7}$**.
2. **<u>Contribution of structures to error.</u>** Although the contribution to the error is calculated for 1 structure, and turbulence is highly chaotic, in the kernel-SHAP method the contribution is taken for a number of **coalitions**, which are groups of structures. This accounts for different **inter-structure interactions** when computing the contribution of the structure to the error. **Weighted average of the contribution to the error.**
   - These simplifications are needed to make the problem **computationally feasible**.
   - Typically we have around $|\mathbf{Q}|$**=150 structures per field, yielding 150! Coalitions.**
   - Kernel SHAP randomly selects **2$|\mathbf{Q}|$+2048 coalitions**, and the importance of each structure in the coalitions where it is present is **weighted to provide the final SHAP value**. Results shown to <u>converge to the values using more coalitions.</u>

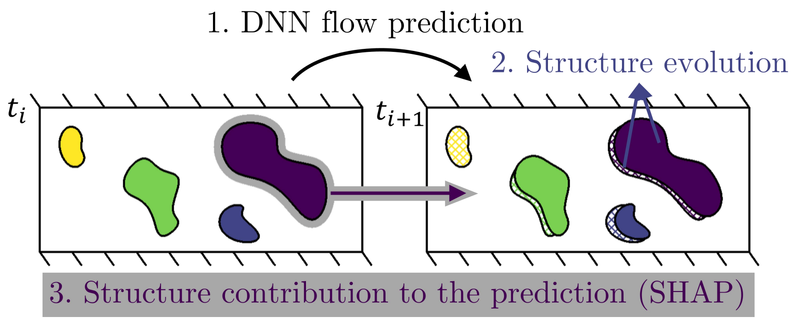

Cremades et al. (2023), Preprint arXiv:2302.01250. Nat. Commun. (To Appear)

More details on the kernel-SHAP method

– It is common to use a **linear model** for the error. Note that **$(f\text{-}g)^2$ ~10^{-7}** in this work.

Reference output of *g* (all structures are absent).

Total number of structures

Presence (1) / Absence (0) of structure 1

$$g(q') = \phi_0 + \sum_{i=1}^{N} \phi_i q'_i = \phi_0 + \boxed{\phi_1} q'_1 + \phi_2 q'_2 + \cdots + \phi_N q'_N$$

Linear model of error between predicted and true fields

Marginal contribution to error (SHAP value) of structure 1 (**negative, structure presence reduces the error**)

– We use **kernel SHAP**, which relies on two techniques: **LIME** and **Shapley values**.

– In LIME we formulate the linear optimization problem as follows:

Local kernel

$$\xi = \arg\min_{g \in \mathcal{G}} \mathcal{L}(f, g, \pi_x) + \Omega(g).$$

Loss function

Error between original and ground truth

Penalization of model complexity

More details on the kernel-SHAP method

- To ensure a unique solution, LIME needs to satisfy several **properties** (local accuracy, consistency, etc.) which are satisfied by the classical **Shapley values**.
- Shapley value: marginal contribution of a particular **structure i (feature)** to the **error f** when included in a particular **group of structures (coalition):**

Coallition without i

$$\phi_i(Q, f) = \sum_{s \subseteq Q \setminus \{i\}} \frac{|s|!(|Q|-|s|-1)!}{|Q|!} \left(f(s \cup \{i\}) - f(s)\right),$$

Set containing all the structures in a field

Total number of structures in the field

Total number of possible coalitions not containing i

- Basically evaluates all possible coalitions, identifies all where i is present, and gets a **weighted average of the contribution of i to the error f.** Emulate inter-scale interactions in turbulence.
- Extremely **challenging from a computational point of view.**

More details on the kernel-SHAP method

- Kernel SHAP is an approximation to the Shapley values:

$$\Omega(g) = 0,$$

Mapping from binary space to input space

$$\mathcal{L}(f, g, \pi_x) = \sum_{q' \in Q} \left[f(h_x(q')) - g(q') \right]^2 \pi_x(q'),$$

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

Number of nonzero structures

- LIME equation solved with linear regression, obtaining **$(f-g)^2 \sim 10^{-7}$** in this work.
- Typically we have around $|Q|=150$ structures per field, yielding 150! Coalitions.
- Kernel SHAP randomly selects $2|Q|+2048$ coalitions, and the importance of each structure in the coalitions where it is present is weighted to provide the final SHAP value.

Importance of the various structures

- **Higher SHAP absolute value** (+ o –) implies **higher importance.**

- **~60% of the structures are sweeps or ejections.**

- **97% of the SHAP** is from **ejections (75%)** and **sweeps (22%).**

- Importance **per unit volume**: **ejections 42%** and **sweeps 44%.**

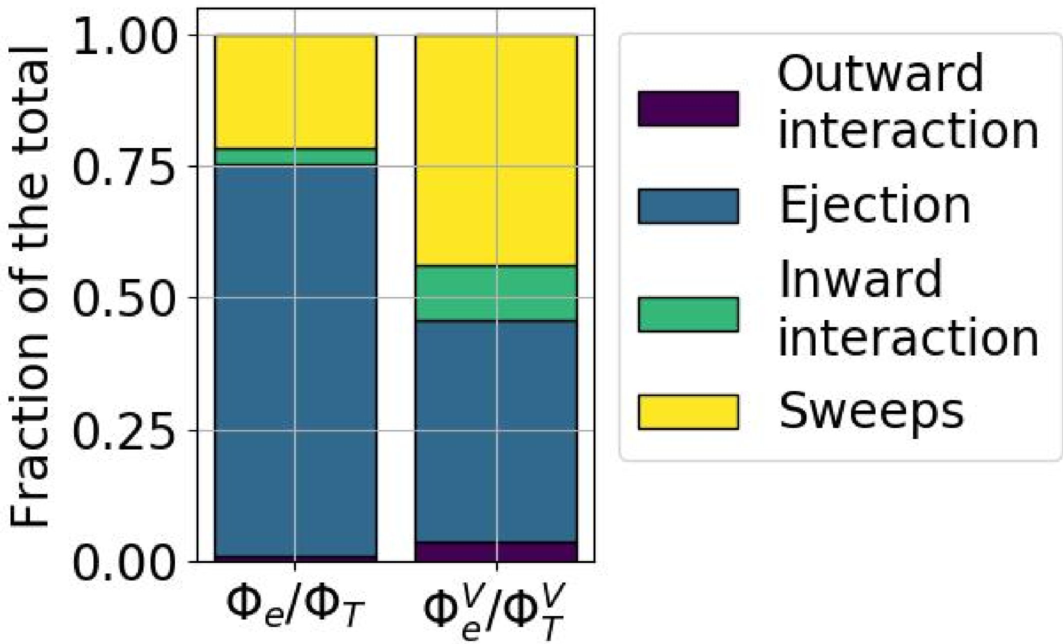

Cremades et al. (2023), Preprint arXiv:2302.01250. Nat. Commun. (To Appear)

Reinforcement learning terminology

- **Model-based RL:** Use transition probability distribution and reward function (**model of the environment**) from the **Markov decision process**.
- **Model-free RL:** Does **not use them, explicit trial and error.**
  - **Q learning:** We try to **learn everything about the system**, and how the reward changes for given actions. Q estimates the reward, assigning a value of reward to any given action for a particular state.
  - **Policy-gradient methods:** We **do not learn everything** about the system, but rather how to **maximize the reward**. Relies on a stochastic policy π (distribution for each state input) parametrized by a NN.
  - **Deep deterministic policy gradient (DDPG):** Combines aspects **of Q learning and policy gradient**. Does not differentiate positive and negative actions, use **actor-critic**:
    - **Actor approximates policy π**, which is considered deterministically.
    - **Critic computes Q** function to assess the **goodness** of those actions.
    - **Action space** is **explored** via a perturbative method through **noisy processes**.

Flat-plate APG TBLs: (MTL) wind-tunnel experiments and simulations

Both to match the LES conditions in *Re* and β as well as further extend the *Re* range. Mainly HWA measurements, incl. planar PIV with different resolutions and OFI.

<u>Styrofoam roof inserts for PG.</u>

Sanmiguel Vila et al. Phys. Rev. Fluids 5 (2020)

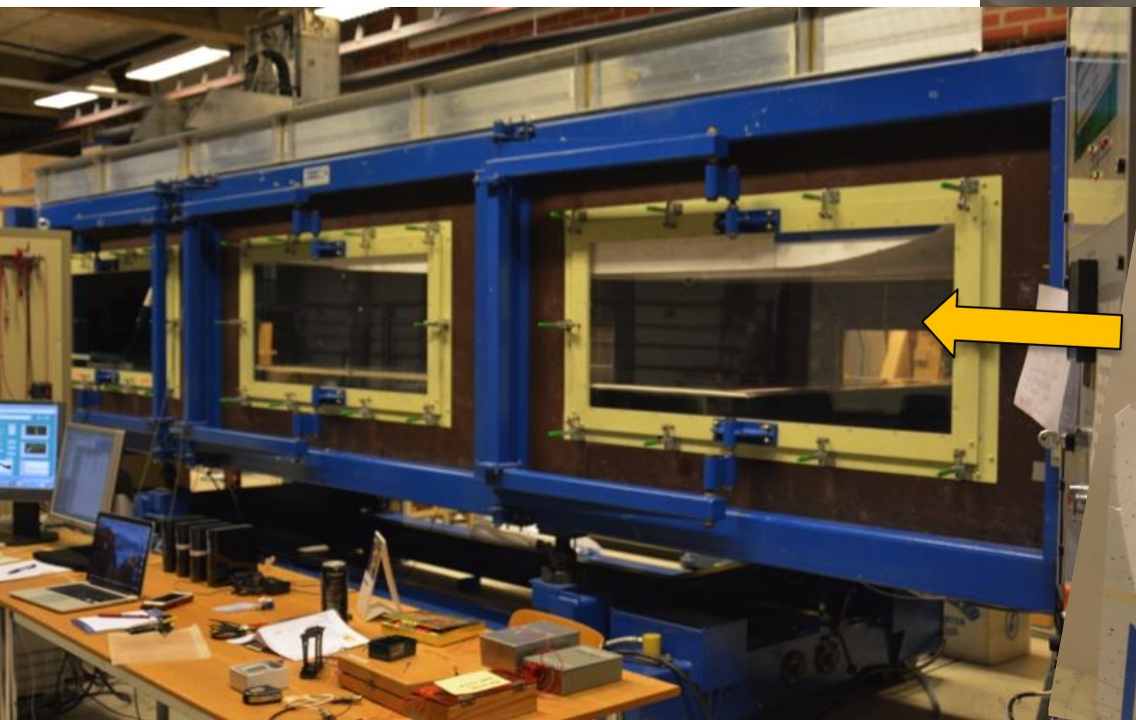

Comparisons at matched parameters: $Re_\tau \approx 4500$ & $\beta \approx 1.2$

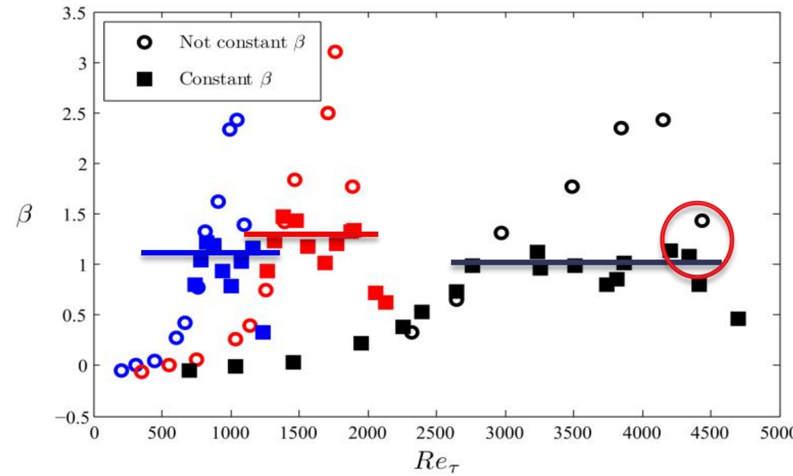

Sanmiguel Vila et al. Phys. Rev. Fluids 5 (2020)