13th Australasian Natural Convection Workshop (13ANCW)
Canberra, Australia, 1st and 2nd December 2025

Analysis of Vertical Natural Convection using β-Variational Autoencoders

Jiyoung Lee[1*], Tony Zahtila[2], Wilson Lu[1], Wai Kit Lam[1], Leon Chan[1], Richard Sandberg[1] and Andrew Ooi[1]

[1] Department of Mechanical Engineering, University of Melbourne, Parkville, VIC 3010, Australia
[2]Center for Turbulence Research, Stanford University, Stanford, CA 94305, USA

*lee16@student.unimelb.edu.au

1. INTRODUCTION

Vertical natural convection (VNC) arises in many engineering applications, such as solar chimneys, electronics cooling, and building ventilation. Numerical simulations conducted to study this flow are typically two-dimensional (2D) due to the reduced computational cost. However, 2D VNC flows vastly differ from turbulent three-dimensional (3D) simulations. In 2D periodic flow such as the setup shown in fig. 1(a), it is found that the flow exhibits chaotic oscillations, including flow reversals and relaminarisation. Figure 1(b)−(c) shows the switching periods over the course of a long simulation time and example snapshots of the temperature field. The underlying physics of this challenging behaviour can be effectively studied using modal analysis, with proper orthogonal decomposition (POD) being a traditional approach. Recent advances in machine learning have shown that autoencoders can be an effective tool for producing reduced-order models [1]. The present study aims to validate the use of β-variational autoencoders (β-VAEs) for natural convection flows and to uncover the mechanisms behind the differing behaviours observed in 2D and 3D simulations.

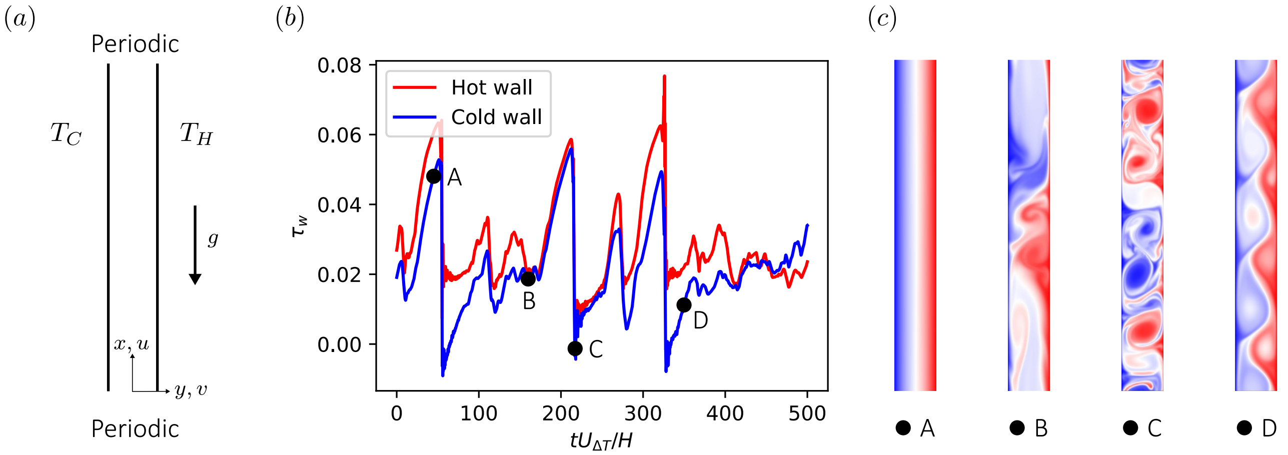

Figure 1. (a) Flow domain showing a vertical 2D periodic channel, (b) wall shear stress on both walls verses nondimensional time, and (c) instantaneous temperature contours at selected timesteps throughout the simulation.

Autoencoders (AEs) are a type of neural network that reduce data dimensionality sequentially through an encoder and a decoder. In between, compression is achieved through a bottleneck layer with a latent space vector typically of dramatically reduced dimension to the high dimensional input data. Variational autoencoders (VAEs) extend this framework by representing the latent space as a probabilistic distribution, typically parameterised by a mean (μ) and variance (σ). The corresponding loss function therefore includes an additional term to assess how well the latent distribution is approximated. Adjusting the weight (β) of this term has been shown to improve orthogonality while maintaining reconstruction accuracy (β-VAE approach) [2].

For our study, we employed a β-VAE on a 2D DNS dataset (Rayleigh number 1×10^6, Prandtl number 0.71), using a four layer encoder to compress the data into 16 latent-space variables. The input dataset comprised 5000 snapshots of normalised streamwise (u) and crossflow (v) velocities along with temperature (T) fields, interpolated onto a spatial grid of $(n_x, n_y) = (128, 64)$, sampled at Δt of 0.1. The first 80% of the data (i.e., up to 400 time units) was used for training, with the remainder reserved for validation. The model was trained using the Adam optimiser for 1200 epochs, until both training and validation losses reached stability.

2. RESULTS

Figure 2(a) shows the cumulative energy captured as the number of modes increases, where for β-VAE, the energy is computed as the L_2 norm of the reconstruction relative to the input data. The reconstruction accuracy for a given number of modes can thus be assessed by the proportion of energy captured. We observe that the β-VAE captures high energy features more rapidly than POD, particularly where the POD curve slows in capturing energy for the higher order modes. This indicates that the β-VAE architecture effectively captures the dominant reduced-order features of the flow with a compact nonlinear representation.

The other two sub-figures show the reconstruction of u velocity fields, at a time (b) from the training set and (c) from the validation set (unseen during training). In (b), both methods produce accurate reconstructions, confirming that β-VAEs can effectively capture complex flows and can identify high-energy structures. For future states in (c), prediction is particularly challenging as it involves forecasting the evolution of the simulation. Nevertheless, the β-VAE generally captures the overall flow trends, indicating that the trained model has successfully learned the complex dynamics.

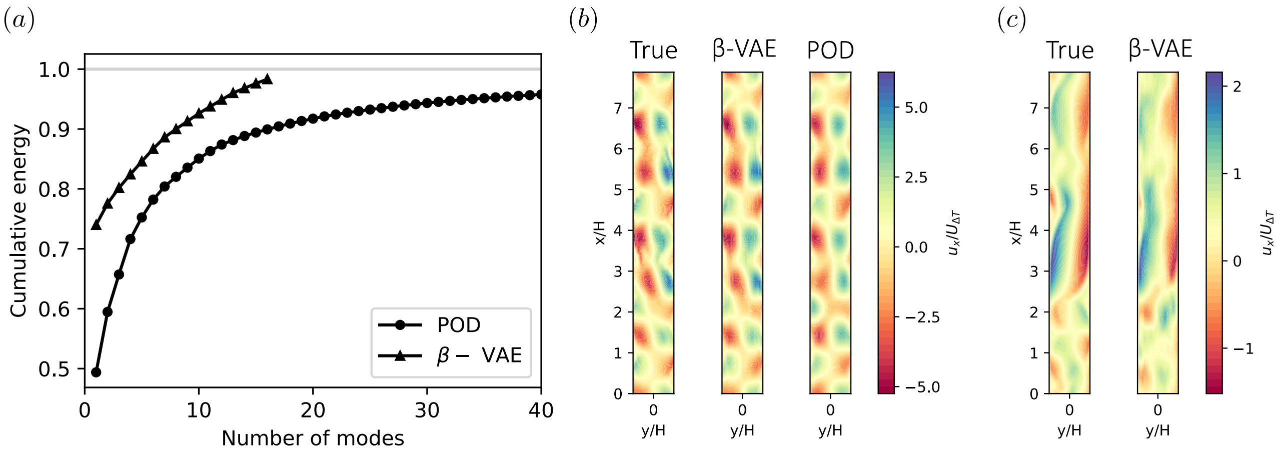

Figure 2. (a) Cumulative energy versus number of modes. Reconstruction of u velocity from (b) the training set (POD reconstructed with 16 modes), at $t = 217$, and (c) from the validation set, at $t = 410$. True represents the input field (2D DNS).

3. CONCLUSION

The present study demonstrates the modal analysis of VNC using a β-VAE. Further work will focus on enhancing the β-VAE architecture, investigating sensitivity to the input data, and ultimately analysing spanwise effects by comparing the latent-space representations of 2D and 3D flows.

REFERENCES

[1] Vinograd MY, Clark Di Leoni P. Reduced representations of Rayleigh–Bénard flows via autoencoders. J. Fluid Mech. 2025;1006:A10. doi:10.1017/jfm.2025.21

[2] Solera-Rico, A., Sanmiguel Vila, C., Gómez-López, M. et al. β-Variational autoencoders and transformers for reduced-order modelling of fluid flows. Nat. Commun. 15, 1361 (2024). https://doi.org/10.1038/s41467-024-45578-4