Deep spatial transformers for autoregressive data-driven forecasting of geophysical turbulence

Ashesh Chattopadhyay[1,2]*, Mustafa Mustafa[2], Pedram Hassanzadeh[1], Karthik Kashinath[2]

Abstract

A deep spatial transformer based encoder-decoder model has been developed to autoregressively predict the time evolution of the upper layer's stream function of a two-layered quasi-geostrophic (QG) system without any information about the lower layer's stream function. The spatio-temporal complexity of QG flow is comparable to the complexity of 500hPa Geopotential Height (Z500) of fully coupled climate models or even the Z500 which is observed in the atmosphere, based on the instantaneous attractor dimension metric. The ability to predict autoregressively, the turbulent dynamics of QG is the first step towards building data-driven surrogates for more complex climate models. We show that the equivariance preserving properties of modern spatial transformers incorporated within a convolutional encoder-decoder module can predict up to 9 days in a QG system (outperforming a baseline persistence model and a standard convolutional encoder decoder with a custom loss function). The proposed data-driven model remains stable for multiple years thus promising us of a stable and physical data-driven climate model.

1   Motivation

Numerical weather prediction (NWP) models remain the state-of-the-art to predict daily weather and usually require an enormous amount of computational resources to resolve the atmospheric fields. These models employ highly accurate and expensive numerical methods that solve the discretized fluid dynamics equations in the atmosphere along with the equations of other physical processes such as radiation, phase change thermodynamics, chemistry and heat transfer. Resolving the fully turbulent flow in the atmosphere demands solving the Navier-Stokes equations on a fine grid, thus making it very expensive and challenging for even the most advanced NWPs to remain accurate for long periods of forecasting time. For that, NWPs often employ parameterization schemes that approximate the physics of processes that are unresolved (sub-grid scale processes). Data-driven models possess an advantage over these numerical models in that, if trained on samples from very high resolution simulations of the flow dynamics, or observed flow dynamics, they can more accurately represent the true tendencies of these otherwise unresolved or under-resolved processes. Such work has been done in the past for atmospheric dynamics and, more generally, for chaotic dynamical systems [1, 2, 3, 4], including ocean dynamics [5]. Another approach involves using the resolved field of a complex chaotic and turbulent flow directly in a data-driven model to be trained on and subsequently evolved forward in time as has been shown in [6, 7, 8, 9]. The assumption here, is that the data-driven model learns the effects of the unresolved or under-resolved sub-grid scale processes directly on the resolved field from a long time history of the fields in the training dataset. In this paper, we build on the problem set-up similar to Chattopadhyay et al. [8], where only part of the system's dynamics is available for training. Instead of the Lorenz 96 system, as used in Chattopadhyay et al. [8] we consider a fully turbulent flow model represented by the two-layered quasi-geostrophic equations (QG) with a baroclinically unstable jet [10, 11]. The model complexity of this QG system based on the instantaneous attractor dimension [12] of the upper layer's stream function (Ψ_1) is about 20.9 and comparable to the instantaneous attractor dimension of Z500 in the observed atmosphere [6].

In this paper, in an effort to perform fully data-driven forecasting from partial observations, we:

- Develop a deep learning model that can accurately predict the short term dynamics of Ψ_1 without any information about the lower layer's stream function Ψ_2 (even during training) although the dynamics of Ψ_1 and Ψ_2 are strongly coupled in the governing equations (see section 2).

- Implement an encoding block inside the deep learning model that can embed (within feature maps) relative translation, scale, rotation, and generic warping of local features of turbulent vortices. The advantages of

*Corresponding author: A Chattopadhyay, akc6@rice.edu [1] Rice University [2] Lawrence Berkeley National Laboratory

incorporating equivariance in convolutional networks has been demonstrated in previous studies, especially in the application of deep learning to physical systems [13].

- Compare two custom loss functions that improve on the performance and long-term climate stability of the models compared to standard root-mean-squared error loss functions.

We perform prediction of $\Psi_1(t)$ autoregressively; i.e., $\Psi_1(t+1)$ is predicted from previously predicted (or initially given) $\Psi_1(t)$ where t is in days. The encoding block inside the deep learning model implements a spatial transformer [14] that is equivariant in feature space. Similar capabilities can also be observed in capsule networks which has been used successfully to improve on the performance of predicting extreme weather events in the atmosphere over traditional convolutional neural networks [15].

2   Two-layered Quasi-Geostrophic Equations

The non-dimensional dynamical equations of QG has been developed following Lutsko et al. [10], and Nabizadeh et al. [11].

$$\frac{\partial q_k}{\partial t} + J(\Psi_k, q_k) = -\frac{1}{\tau_d}(-1)^k \left(\Psi_1 - \Psi_2 - \Psi_R\right)$$
$$-\frac{1}{\tau_f}\delta_{k2}\nabla^2\Psi_k - \nu\nabla^4 q_k \tag{1}$$

$$q_k = \nabla^2\Psi_k + (-1)^k \left(\Psi_1 - \Psi_2\right) + \beta y \tag{2}$$

In Eq. (1) and (2), q is the potential vorticity, k denotes the upper ($k = 1$) and lower ($k = 2$) level. τ_d is the Newtonian relaxation time scale while τ_f is the Rayleigh friction time scale, which only acts on the lower level represented by the Kronecker δ function (δ_{k2}). J denotes the Jacobian, β is the Rossby parameter, and ν denotes the hyperdiffusion coefficient. We have induced a baroclinically unstable jet at the center of a horizontally (zonally) periodic channel by setting $\Psi_1 - \Psi_2$ to be equal to a hyperbolic secant centered at $y = 0$. When eddy fluxes are absent, Ψ_2 is identically zero, making zonal velocity in the upper layer $U_1(y) = -\frac{\partial \Psi_1}{\partial y} = -\frac{\partial \Psi_R}{\partial y}$ where

$$-\frac{\partial \Psi_R}{\partial y} = sech^2(\frac{y}{\sigma}) \tag{3}$$

σ being the width of the jet. Eq. (1) is solved using a pseudo-spectral method with a uniform Fourier grid along x (128 grid points between $x = [0, 46]$) and Chebyshev grid along y (192 grid points between $y = [-1, 1]$). For training (see section 3) some of the upper and lower grid points along y are ignored (retaining $y = [-0.6, 0.6]$) in order to remove standing waves that are formed there due to the absence of sponge layers. For more details on the solution of Eq. (1) see Nabizadeh et al. [11]. Eq. (1) has been integrated for 120000 days with $\Delta t = 0.005$ days using a second order Adams-Bashforth method. Flow obtained from this simulation is fully turbulent, has an inverse energy cascade, and closely mimics the key features of the mid-latitude circulation variability. It presents a challenging test case to develop and test new data-driven models aimed at forecasting geophysical turbulence. The objective of the deep learning framework is to predict $\Psi_1(t)$ as a function of time autoregressively *without* any information available on $\Psi_2(t)$, although, physically, both Ψ_1 and Ψ_2 are coupled and they affect each other's dynamics, as is evident from Eq. (1).

3   Deep learning Architecture

In this paper, we explore the capability of a state-less convolutional encoder-decoder style architecture with an embedded spatial-transformer module to track rotational variances in the feature space of the input. Building on the superior performance of capsule networks on extreme weather prediction [15], we believe that, spatial transformers preserving the key equivariance property in the feature maps of subsequent convolutions [14] (similar to capsules), allow us to track the dynamics of vortices in turbulent flow. Compared to capsules, spatial transformers are more memory efficient and easier to implement as a fully differentiable layer inside any deep learning architecture. The details of the architecture is shown in the schematic given in Fig 1.

3.1   Localization Network or Encoding Block

The architecture takes in an input snapshot $\Psi_1(t)^{88 \times 128}$ and projects it onto a low dimensional encoding space via a standard convolutional encoding block. This encoding block performs three convolutions (without changing the spatial dimensions) and max-pooling which connects to a dense network with three layers. The encoded feature map connects into the spatial transformer module (see section 3.2). The convolutions inside the encoder block ensure periodic boundary conditions zonally by performing circular convolutions on each feature map inside the encoder block [16].

3.2   Spatial Transformer Module

The spatial transformer takes the feature map as the input and then applies an affine transformation $\Omega(\theta)$, where the parameters θ are learnt through backpropagation. One could use any other transformation but since the vortices rotate on the $x - y$ plane in our system, we have chosen an affine transformation so that that relative rotational and stretching features in the flow can be captured. The affine transformation results into transforming the original co-ordinate space x_i^o and y_i^o into x_i^s and y_i^s following

$$\begin{bmatrix} x_i^s \\ y_i^s \end{bmatrix} = \Omega(\theta) \begin{bmatrix} x_i^o \\ y_i^o \end{bmatrix} \tag{4}$$

The transformer module then applies a differentiable sampling kernel (a bi-linear interpolation kernel in this case) to the input feature map (denoted by F) at the locations given by x_i^s and y_i^s to produce the output feature map G. This can be written as

$$G_i = \sum_{n=0}^{n=88} \sum_{m=0}^{m=128} F_{nm} k(x_i^s - m; \Phi_x) k(y_i^s - n; \Phi_y) \tag{5}$$

Here $i \in [1, 2, \cdots 88 \times 128]$ and k is a generic sampling kernel with sampling parameters Φ_x and Φ_y. In our implementation k is a bi-linear interpolation kernel which is differentiable.

3.3   Decoding Block

The decoding block is a standard series of deconvolution layers (convolution with zero-padded upsampling) that bring the features encoded by the transformer module back into its original dimension.

3.4   Training

The architecture is trained on pairs of $\Psi_1^{train}(t)$ and $\Psi_1^{train}(t+1)$ for 100000 days (out of the generated 120000 days) with 20% left for validation. The remaining 20000 days have been used for testing. A total of 10 initial conditions have been chosen from those remaining 20000 days to test the performance of the architecture in terms of short term forecasting and sub-seasonal to seasonal scale (S2S) forecasting. Multi-timestep prediction is done autoregressively, i.e. during testing, a randomly chosen initial condition with $\Psi_1^{test}(t)$ is inputted to the architecture to obtain $\Psi_1^{test}(t+1)$. This predicted $\Psi_1^{test}(t+1)$ is fed back in to the architecture to obtain $\Psi_1^{test}(t+2)$ and so on and so forth in an autoregressive manner. We have conducted such prediction up to 1000 days and found the architecture to yield a stable and non-drifting climate.

3.5   Loss Function

Along with the standard root-mean-squared error loss function, we have explored the possibility of using custom loss function aimed at restricting drift in long term climate. For that purpose we introduce loss functions

$$\mathcal{L}_1 = (1 - \gamma) \frac{1}{N} \sum_N \left(\Psi_1^{pred} - \Psi_1^{true} \right)^2 + \frac{1}{N} \gamma \sum_N \Psi_1^{pred} \tag{6}$$

and

$$\mathcal{L}_2 = (1 - \gamma) \frac{1}{N} \sum_N \left(\Psi_1^{pred} - \Psi_1^{true} \right)^2$$
$$+ \gamma \frac{1}{N} \sum_N \left(\nabla_{xy} \Psi_1^{pred} - \nabla_{xy} \Psi_1^{true} \right)^2 \tag{7}$$

where N is the number of samples per batch during training and γ is a hyper-parameter. Ψ_1^{pred} and Ψ_1^{true} are the predicted and true values of the upper layer's stream function respectively. These changes to the loss function enable longer, stable, and non-drifting climate comparable to a baseline persistence model while partially improving short-term forecasting performance. In our experiments we found the hyper-parameter γ to be 0.1 for $\mathcal{L}_1$ and 0.25 for $\mathcal{L}_2$ to yield the best performance. $\mathcal{L}_1$ enables the jet-stream from drifting in any direction by pulling it back towards the mid-latitude while $\mathcal{L}_2$ is meant for capturing sharp features through the gradient which may otherwise contribute to the error that drifts the jet-stream.

4   Results

In this section we compare the performance of the data-driven deep spatial transformer based convolutional encoder-deocoder against a baseline persistence model and a standard convolutional encoder-decoder model. We investigate the performance with loss functions $\mathcal{L}_1$, $\mathcal{L}_2$, and a standard root-mean-squared loss function.

4.1   Short-term performance

The short-term performance is measured as the correlation coefficient (R) between $\Psi_1^{pred}(t)$ and $\Psi_1^{true}(t)$ at each day t starting from a random initial condition on an unseen dataset. Furthermore, we compute this coefficient on each snapshot between $y = -0.25$ and $y = 0.25$ so that only regions of maximum variance in the turbulent jet are considered, thus making R a fair metric. Regions further away from the jet show low variability and thus would always be predicted well. We also report the algebraic value of R instead of the absolute value so as to determine the drift in long-term data-driven climate, e.g. a negative R would suggest that the data-driven model has un-physically predicted large low-pressure anomalies near the jet that are not supposed to be there in the true flow.

Figure 2 shows time snapshots for upto 5 days prediction by the deep learning model with $\mathcal{L}_1$ loss compared against the baseline persistence model. The deep learning model shows $R = 70\%$ at 5 days which outperforms persistence by 10%. While certain high wave number features in individual snapshots are not properly captured, the overall pattern is well captured along with the amplitude near the turbulent jet. This could be attributed to the equivariance embedded into the neural network that allows it to track, and thus predict rotational features and stretching in turbulent eddies.

4.2   Sub-seasonal to seasonal scale prediction

As seen in Weyn et al. [7], persistence performs well further away in time from the initial condition where forecasting began (sub-seasonal lead days show more correlation with persistence). This is not surprising, since in a chaotic system, small errors quickly accumulate and the predicted trajectory diverges away from the true trajectory of the dynamical system. Hence, it is important to see how well the data-driven model performs as we increase the prediction timescale from 5 days to several months (seasonal scale). We can see in Fig 3 that the data-driven model with $\mathcal{L}_1$ and $\mathcal{L}_2$ losses outperform persistence and the baseline encoder-decoder without a spatial transformer even at ≈ 10 days (sub-seasonal scale). Moreover, at longer time scales ($\approx 20 - 90$ days) the model remains physical and the predicted jet-stream does not drift. This allows us to compute a physical climate from the data-driven model. The deep learning model with the standard root-mean-squared loss does not simulate a physical climate and shows a drift in the jet where low pressure anomalies dominate.

4.3   Climatology

Figure 4 shows long-term averaged quantities of the dynamics and how they behave in the data-driven models proposed as compared to the true long-term dynamics. The zonally averaged Ψ_1 has been obtained by, first, averaging over 1000 days of predicted Ψ_1 to obtain the time mean of Ψ_1 from the data driven model (and compared against true Ψ_1) and then zonally averaging this obtained time mean. A similar analysis has been done for the meridionally averaged Ψ_1. We can see that the long-term mean of the meridionally averaged Ψ_1 resembles the true long-term mean of the system more closely than persistence while in the case of zonally averaged Ψ_1, both persistence and data-driven models resemble the truth quite well.

4.4   Energetics of the predicted flow

The first left singular vector and 2nd left singular vector corresponds to the largest and the second-largest singular values obtained from a singular value decomposition of the long-term mean of Ψ_1. These would denote the most

energetically dominant modes of the flow. The data-driven models, both with $\mathcal{L}_1$ and $\mathcal{L}_2$ loss functions capture the true singular vectors and slightly outperforms persistence. Long-term averaged quantities could not be calculated for the data-driven model with a standard root-mean-squared loss function because, at sub-seasonal scales the turbulent jet drifts and produces un-physical climate as is evident from the negative R values in Fig 3.

5 Discussion

In this paper we have developed an equivariant convolutional encoder-decoder with a spatial transformer for predicting turbulent flow in a two-layered QG system. The rotational and scale invariance inside the spatial transformer [14] allows the deep learning architecture to track the rotational variances in the features of turbulent vortices in the QG flow with limited information of the dynamics during training and testing. We show excellent short-term performance upto 5 days outperforming a persistence model. While short-term forecasting of turbulent flow is a very challenging problem in geophysical and engineering applications, a fully data-driven model such as the one proposed in this paper shows great promise. We intend to extend these rotationally invariant or equivariant frameworks [13, 15] to more complicated datasets, e.g. re-analyses data, which strongly resemble the observed atmosphere. Beyond fully data-driven models, equivariant neural networks can be incorporated in building data-driven surrogates for sub-grid scale closure models or parameterization schemes that are used in complicated fully-coupled climate models [1].

The loss functions introduced in this paper help the model in avoiding long-term drift and producing a physically meaningful climate. While the gradient based loss function enables the architecture to capture sharper features near the turbulent jet, the exact mechanism that enables the model to deliver a physically consistent climate at the seasonal scale should be investigated further. The deep learning architecture without $\mathcal{L}_1$ and $\mathcal{L}_2$ produces an unphysical climate with strong drifts in the jet. Although we expect prediction errors in chaotic systems errors to accumulate quickly based on the Lyapunov exponent, the reason behind producing an unphysical climate needs to be investigated more closely. Carefully constructed stateful neural architectures [8] should be explored where predicted Lyapunov exponent can be compared and corrected compared to the true Lyapunov exponent of the system.

In summary, the key contributions of this paper are:

- a novel deep neural network architecture that incorporates equivariance to track turbulent flow dynamics better than existing data-driven methods.

- significantly improved short-term data-driven forecasting of QG turbulence.

- stable long-term predictions with a physically meaningful climate and no drift.

While the QG problem is a challenging test case, scaling this architecture for more realistic atmospheric data (e.g. re-analysis data) is the next step. Careful ablation studies on the effect of the spatial transformer, number of inputs into the architecture, and the limits of spatial resolution (on how big a domain can we predict?) are being undertaken. We believe that incorporating symmetry [13] and equivariance [15] in neural architectures would significantly improve fully data-driven prediction in atmospheric dynamics.

Acknowledgments

References

[1] S. Rasp, M. S. Pritchard, and P. Gentine, "Deep learning to represent subgrid processes in climate models," *Proceedings of the National Academy of Sciences*, vol. 115, no. 39, pp. 9684–9689, 2018.

[2] N. D. Brenowitz and C. S. Bretherton, "Prognostic validation of a neural network unified physics parameterization," *Geophysical Research Letters*, vol. 45, no. 12, pp. 6289–6298, 2018.

[3] P. A. O'Gorman and J. G. Dwyer, "Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change, and extreme events," *Journal of Advances in Modeling Earth Systems*, vol. 10, no. 10, pp. 2548–2563, 2018.
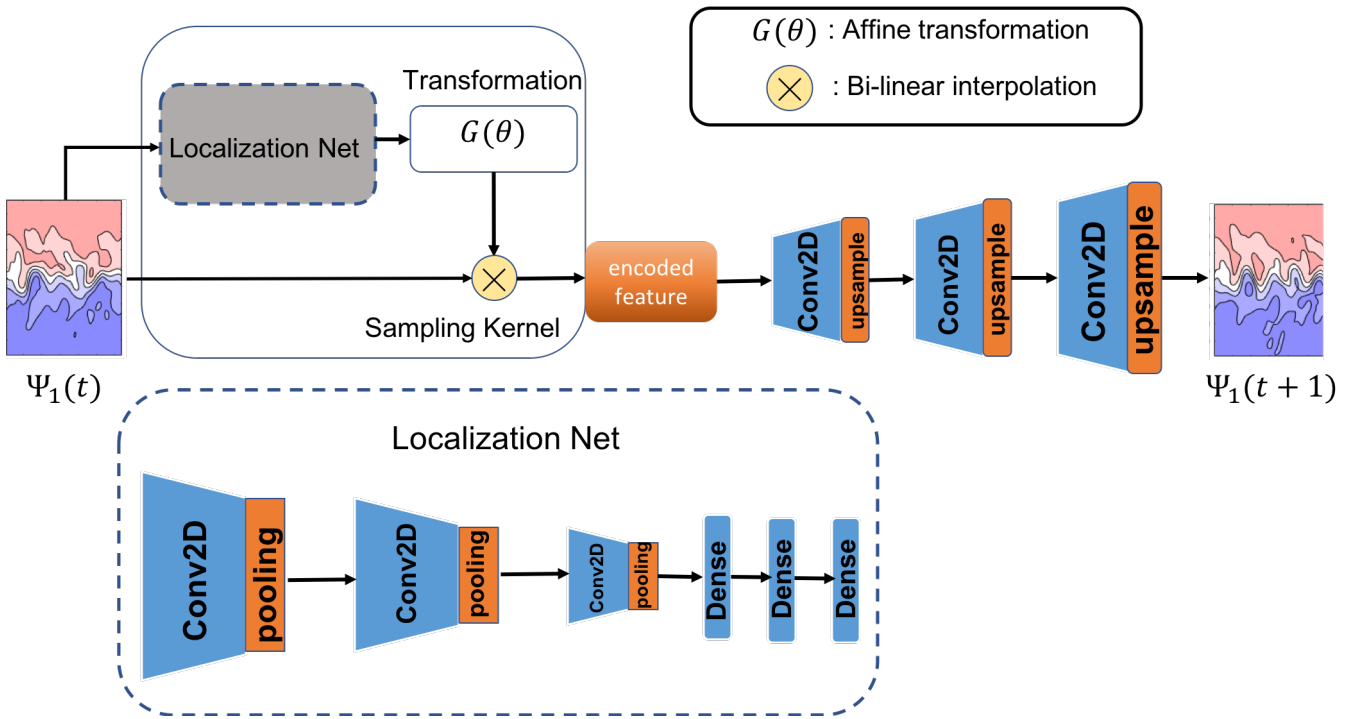
Figure 1: Schematic of the deep learning architecture with the spatial transformer module. The architecture is trained with pairs of $\Psi_1(t)$ and $\Psi_1(t+1)$.
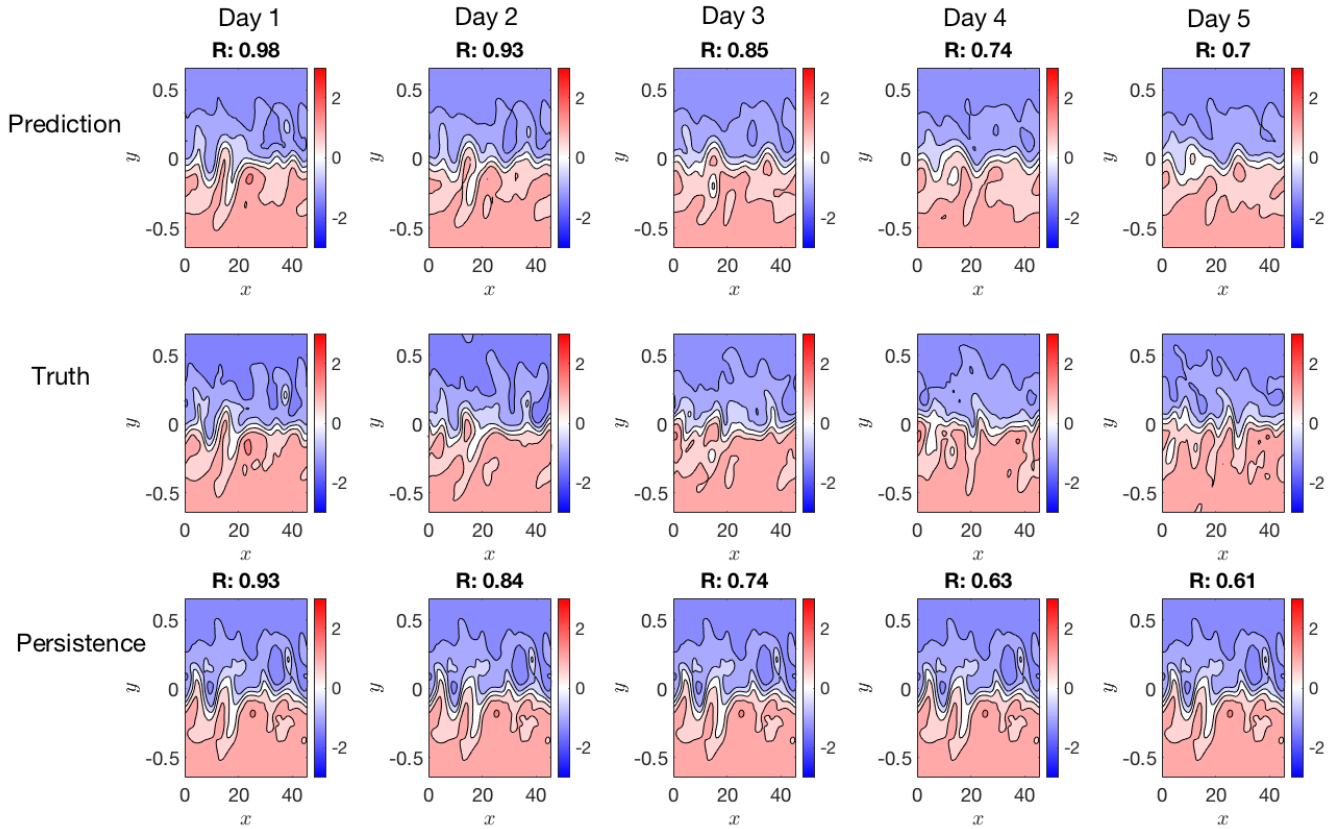

Figure 2: Short term prediction performance on Ψ_1(upto 5 days) of the deep learning framework with $\mathcal{L}_1$ as the loss function starting from a single random initial condition. The algebraic value of R (correlation coefficient between two matrices) calculated between $y = -0.25$ and $y = 0.25$ (where the maximum variability of the turbulent jet occurs) outperforms baseline persistence model by $\approx 10\%$ at 5 days.
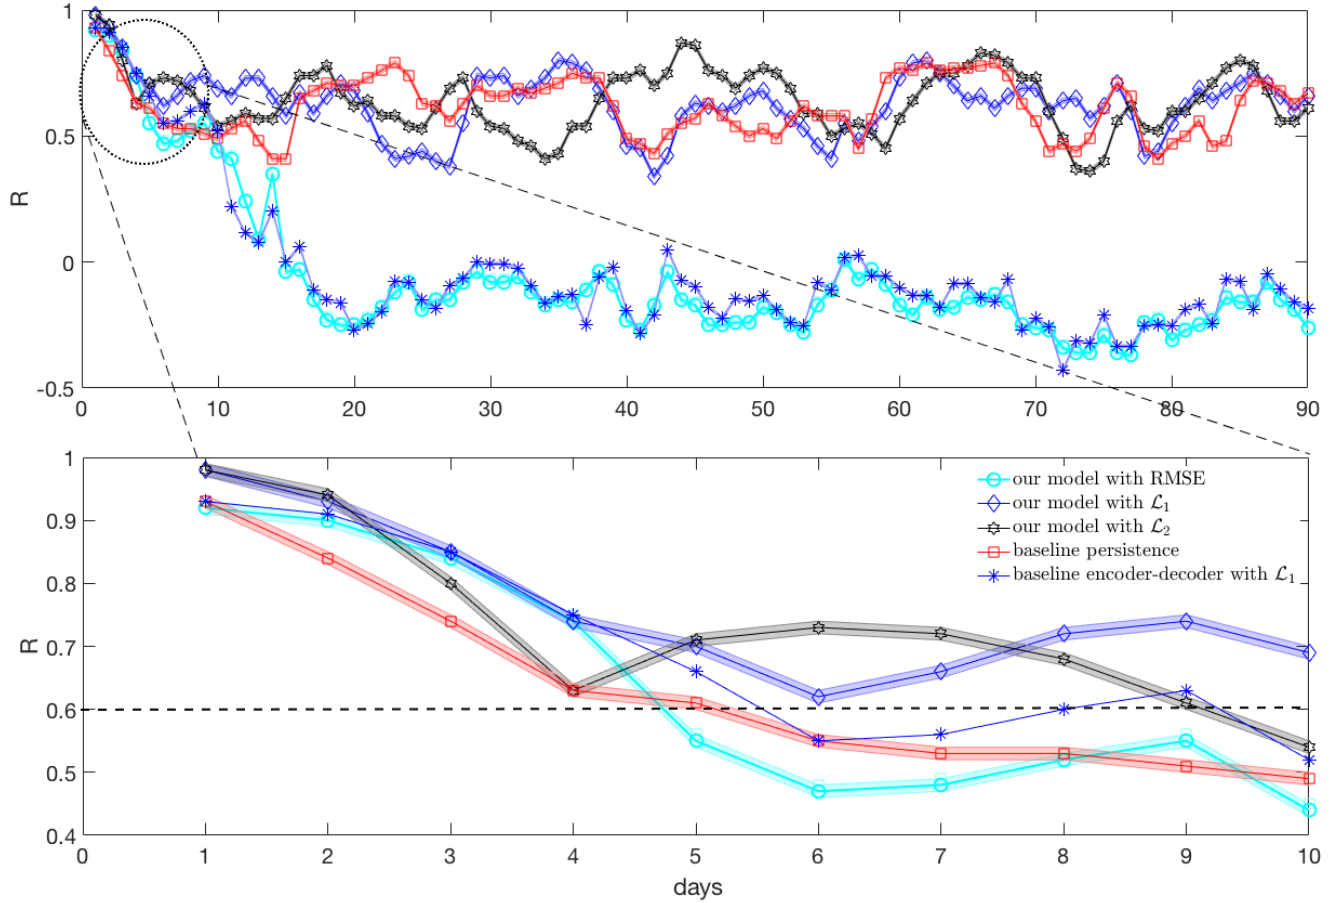
Figure 3: R, the algebraic value of the correlation coefficient between two matrices for short term and sub-seasonal to seasonal prediction with 3 different network architectures. $\mathcal{L}_1$ and $\mathcal{L}_2$ losses are defined in section 3.5 and compared against standard root mean squared error loss. While in the very short term, the performance of all three architectures are similar, at the seasonal scale (≈ 3 months), $\mathcal{L}_1$ and $\mathcal{L}_2$ improve stability and reduce drift resulting in a more physical climate with R remaining around 60% even after 3 months. All deep learning architectures outperform persistence for short term prediction upto 4 days. Beyond 4 days the architectures with $\mathcal{L}_1$ and $\mathcal{L}_2$ losses outperform persistence for upto 10 days (sub-seasonal scale) and remain comparable with persistence for upto 90 days, with a stable physical climate. All of the analyses were repeated for 10 different initial conditions chosen from the test set and separated by at least 1000 days. The mean (symbols) and standard deviation (shading) are reported in the figure. The top plot shows R upto 90 days; the bottom plot zooms into the first 10 days.
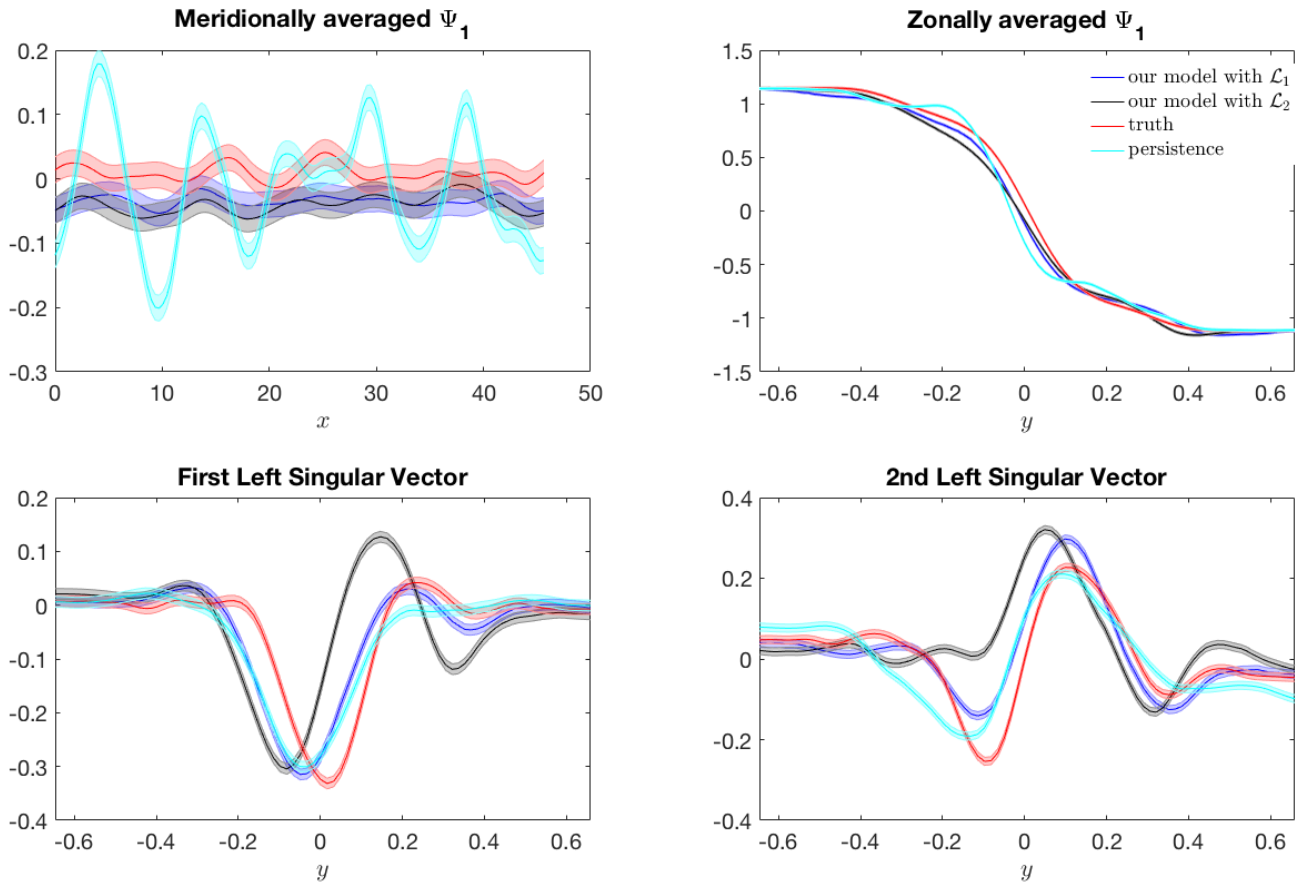
Figure 4: Figure shows long-term averaged dynamical quantities predicted by the deep learning models and compared against persistence and truth. For both meridionally and zonally averaged Ψ_1, Ψ_1 has been averaged over 1000 days. Both the left singular vectors are computed by taking the singular value decomposition of the time averaged Ψ_1. The architecture that uses a standard root-mean-squared loss function (or the baseline convolutional encoder-decoder) has not been shown because it does not achieve a stable physical climate for 1000 days and has a drifting turbulent jet with unphysical low pressure anomalies beyond 10 days.

[4] A. Chattopadhyay, A. Subel, and P. Hassanzadeh, "Data-driven super-parameterization using deep learning: Experimentation with multi-scale lorenz 96 systems and transfer-learning," *arXiv preprint arXiv:2002.11167*, 2020.

[5] T. Bolton and L. Zanna, "Applications of deep learning to ocean data inference and subgrid parameterization," *Journal of Advances in Modeling Earth Systems*, vol. 11, no. 1, pp. 376–399, 2019.

[6] S. Scher and G. Messori, "Weather and climate forecasting with neural networks: using general circulation models (gcms) with different complexity as a study ground," *Geoscientific Model Development*, vol. 12, no. 7, pp. 2797–2809, 2019.

[7] A. Weyn, Jonathan, R. Durran, Dale, and R. Caruana, "Can machines learn to predict weather? using deep learning to predict gridded 500-hPa geopotential height from historical weather data," *Journal of Advances in Modeling Earth Systems*, vol. 10, no. 8, pp. 2680–2693, 2019.

[8] A. Chattopadhyay, P. Hassanzadeh, D. Subramanian, and K. Palem, "Data-driven prediction of a multi-scale lorenz 96 chaotic system using a hierarchy of deep learning methods: Reservoir computing, ann, and rnn-lstm," 2019.

[9] P. D. Dueben and P. Bauer, "Challenges and design choices for global weather and climate models based on machine learning," *Geoscientific Model Development*, vol. 11, no. 10, pp. 3999–4009, 2018.

[10] N. J. Lutsko, I. M. Held, and P. Zurita-Gotor, "Applying the fluctuation–dissipation theorem to a two-layer model of quasigeostrophic turbulence," *Journal of the Atmospheric Sciences*, vol. 72, no. 8, pp. 3161–3177, 2015.

[11] E. Nabizadeh, P. Hassanzadeh, D. Yang, and E. A. Barnes, "Size of the atmospheric blocking events: Scaling law and response to climate change," *Geophysical Research Letters*, vol. 46, no. 22, pp. 13488–13499, 2019.

[12] D. Faranda, G. Messori, and S. Vannitsem, "Attractor dimension of time-averaged climate observables: insights from a low-order ocean-atmosphere model," *Tellus A: Dynamic Meteorology and Oceanography*, vol. 71, no. 1, p. 1554413, 2019.

[13] R. Wang, R. Walters, and R. Yu, "Incorporating symmetry into deep dynamics models for improved generalization," *arXiv preprint arXiv:2002.03061*, 2020.

[14] M. Jaderberg, K. Simonyan, A. Zisserman, *et al.*, "Spatial transformer networks," in *Advances in neural information processing systems*, pp. 2017–2025, 2015.

[15] A. Chattopadhyay, E. Nabizadeh, and P. Hassanzadeh, "Analog forecasting of extreme-causing weather patterns using deep learning," *Journal of Advances in Modeling Earth Systems*, vol. 12, no. 2, p. e2019MS001958, 2020.

[16] S. Schubert, P. Neubert, J. Pöschmann, and P. Pretzel, "Circular convolutional neural networks for panoramic images and laser data," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 653–660, IEEE, 2019.