A test case for application of convolutional neural networks to spatio-temporal climate data: Re-identifying clustered weather patterns

Ashesh Chattopadhyay[1], **Pedram Hassanzadeh**[1,*]**, and Saba Pasha**[1,2]

[1]Rice University, Houston, 77005, USA
[2]University of Pennsylvania, Philadelphia, 19104, USA
[*]pedram@rice.edu

ABSTRACT

Convolutional neural networks (CNNs) can potentially provide powerful tools for classifying and identifying patterns in climate and environmental data. However, because of the inherent complexities of such data, which are often spatio-temporal, chaotic, and non-stationary, the CNN algorithms must be designed/evaluated for each specific dataset and application. Yet to start, CNN, a supervised technique, requires a large labeled dataset. Labeling demands (human) expert time, which combined with the limited number of relevant examples in this area, can discourage using CNNs for new problems. To address these challenges, here we 1) Propose an effective auto-labeling strategy based on using an unsupervised clustering algorithm and evaluating the performance of CNNs in re-identifying these clusters; 2) Use this approach to label thousands of daily large-scale weather patterns over North America in the outputs of a fully-coupled climate model and show the capabilities of CNNs in re-identifying the 4 clustered regimes. The deep CNN trained with 1000 samples or more per cluster has an accuracy of 90% or better. Accuracy scales monotonically but nonlinearly with the size of the training set, e.g. reaching 94% with 3000 training samples per cluster. Effects of architecture and hyperparameters on the performance of CNNs are examined and discussed.

Introduction

Classifying and identifying specific patterns in the spatio-temporal climate and environmental data are of great interest for various purposes such as finding circulation regimes and teleconnection patterns[1-5], identifying extreme-causing weather patterns[6-8], studying the effects of climate change[9-12], understanding ocean-atmosphere interaction[13-15], and investigating air pollution transport[16,17], just to name a few. Such classifications are often performed by employing empirical orthogonal function (EOF) analysis, clustering algorithms (e.g., K-means, hierarchical, self-organizing maps), linear regression, or specifically designed indices, such as those used to identify atmospheric blocking events. Each approach suffers from some major shortcomings (see the reviews by Grotjahn *et al.*[6] and Monahan *et al.*[18]); for example, there are dozens of blocking indices which frequently disagree and produce conflicting statistics on how these high-impact extreme-causing weather patterns will change with climate change[10,19].

In recent years, applications of machine learning methods for accelerating and facilitating scientific discovery have increased rapidly in various research areas. For example, in climate science, neural networks have produced promising results for parameterization of convection and simulation of clouds[20-24], weather forecasting[25,26], and predicting El Niño[27]. A class of supervised deep learning architectures, called convolutional neural network (CNN), has transformed pattern recognition and image processing in various domains of business and science[28,29] and can potentially become a powerful tool for classifying and identifying patterns in the climate and environmental data. In fact, in their pioneering work, Liu *et al.*[30] and Racah *et al.*[31] have shown the promising capabilities of CNNs in identifying tropical cyclones, weather fronts, and atmospheric rivers in large, labeled climate datasets.

Despite the success in applying CNNs in these few studies, to further expand the applications and usefulness of CNNs in climate and environmental sciences, there are some challenges that should be addressed[32]. One major challenge is that unlike the data traditionally used to develop and assess CNN algorithms such as the static images in ImageNet[33], the climate and environmental data, from model simulations or observations, are often spatio-temporal and highly nonlinear, chaotic, high-dimensional, non-stationary, multi-scale, and correlated. For example, the large-scale atmospheric circulation, whose variability strongly affects day-to-day weather and extreme events, is a high-dimensional turbulent system with length scales of 100 km to 10000 km and time scales of minutes to decades (and beyond), with strongly coherent and correlated patterns due to various physical processes, and non-stationarity due to, e.g., atmosphere-ocean coupling and anthropogenic effects[34-36]. An additional challenge with observational datasets is that they are usually short and sparse and have measurement noise.

As a result, to fully harness the power of CNNs, the algorithms (architecture, hyperparameters etc.) have to be designed and evaluated for each specific climate or environmental data and for each specific application. However, to start, CNN, as a supervised technique, requires a large labeled dataset for training/testing. Labeling data demands (human) expert time and resources and while some labeled datasets for specific types of data and applications are now publicly available[31,37], can discourage exploring the capabilities of CNN for various problems. With these challenges in mind, the purpose of this paper is twofold:

1. To propose an effective, simple approach for labeling any spatio-temporal climate and environmental data based on using K-means clustering, which is an easy-to-implement, unsupervised classification technique,

2. To use this approach and label thousands of large-scale weather patterns over North America in the outputs of a state-of-the-art climate model and show the capabilities of CNNs in identifying the four different clusters and examine how the performance of CNNs depend on the architecture, hyperparameters, and size of the training dataset.

Methodology

The approach proposed here involves two steps: i) the spatio-temporal data is clustered into n classes using an unsupervised technique such as K-means[38], which assigns an index (1 to n) to each pattern in the dataset, and ii) the cluster indices are used to label the patterns in the dataset, 1 to n. The labeled dataset is then used to train and test the CNN. The performance of CNN in re-identifying the cluster indices of patterns in the testing phase can be used to evaluate and improve the CNN algorithms for each specific dataset. Note that while here we use K-means clustering for indexing, other algorithms such as hierarchical, expectation-maximization, or self-organizing maps[3,4,39] can be used instead. However, the K-means algorithm, which clusters the data into *a priori* specified n based on Euclidean distances, provides an effective, simple method for the objective here, which is to label the dataset for evaluating CNN, rather than finding the most meaningful (if possible[13]) number of clusters in the spatio-temporal data.

The approach proposed here can be used for any climate or environmental data such as wind, precipitation, or sea-surface temperature patterns or distributions of polluters, to name a few. For the case study presented here, we focus on the daily weather patterns over North America in summer and winter. The data and K-means clustering and CNN algorithms are presented in Data and Methods, but their key aspects are briefly discussed below. We use data from the Large Ensemble (LENS) Community Project[40], which consists of a 40-member ensemble of fully-coupled Community Earth System Model version 1 (CESM1) simulations with historical radiative forcing from 1920 to 2005. We focus on daily averaged geopotential height at 500 hPa (Z500 hereafter), whose isolines are approximately the streamlines of the large-scale circulation at mid-troposphere and are often used to represent weather patterns[41]. Daily Z500 from 1980 to 2005 provides ~ 95000 samples for summer months and for winter months over North America.

As discussed in Data and Methods, the K-means algorithm is used to classify the winter days and summer days (separately) into $n = 4$ clusters. The clustering analysis is performed on zonal-mean-removed daily Z500 anomalies projected on 22 EOFS that retain approximately 95% of the variance; however, the computed cluster index for each day is used to label that day's full Z500 pattern. The four cluster centers in terms of the full Z500 field are shown in Figure 1. Labeled full Z500 patterns are used as input to CNN for training and testing. We work with the full Z500 fields, rather than the computed anomalies, because one hopes to use CNN with minimally pre-processed data. Therefore, we focus on the more difficult task of re-identifying the clusters in the full Z500 fields, which include the complex temporal variabilities such as the seasonal cycle and non-stationarity resulting from the low-frequency coupled ocean-atmosphere modes and changes in the radiative forcing between 1980 and 2005. We further emphasize that the spatio-temporal evolution of Z500 field is governed by high-dimensional, strongly nonlinear, chaotic, multi-scale dynamics[41].

The architecture of our CNN algorithm is shown in Figure 2. In general, the main components of a CNN algorithm are: convolutional layers in which a specified number of kernels (filters) of specified sizes are applied to extract the key features in the data and produce feature maps; Rectified Linear Unit (ReLU) layers in which the ReLU activation function, $f(x) = \max(0, x)$, is applied to the feature maps to introduce nonlinearity; pooling layers that reduce the dimensions of the feature maps to increase the computational performance, control overfitting, and induce translation and scale invariance (which is highly desirable for the chaotic spatio-temporal data of interest here); and finally, fully connected layers[28,29]. The inputs to CNN are the full Z500 fields that are converted to images and down-sampled to reduce redundancies in small scales (Figure 3). During the training phase, the images and their cluster indices, from a randomly drawn training set (TR), are inputted into CNN and the kernels (i.e. their weights) are learned using backpropagation[29]. The major advantage of CNNs over traditional image processing methods is that the appropriate kernels are learned for each dataset, rather than being hand-engineered and specified *a priori*. During the testing phase, images, from a randomly drawn testing set (TS) that has no overlap with TR, are inputed into the CNN and the output is the predicted cluster index. If the CNN has learned the key features of these high-dimensional, nonlinear, chaotic, non-stationary patterns, then the predicted cluster indices should be largely correct.

Table 1. The confusion matrix for CNN4 (CNN2) applied to summer months. A TR of length $N = 12000$ (3000 samples per cluster) and a TS, consisting of 5 independent sets each with 1000 samples per cluster, are used (see Data and Methods). The TR and TS are randomly selected and have no overlap. Each number shows how many patterns from a given cluster in TS are identified by the trained CNN to belong to each cluster (the mean and standard deviation from the 5 sets of TS are reported). The results are from the best trained CNN2 and CNN4. The overall test accuracy, calculated as the sum of the diagonal numbers, i.e. all correctly identified patterns, divided by the total number of patterns, i.e. 3000, and turned to percentage is $93.3\% \pm 0.2\%$ (CCN4) and $89.0\% \pm 0.3\%$ (CNN2).

	Identified as C1	Identified as C2	Identified as C3	Identified as C4
True C1	$\mathbf{915 \pm 3}\,(\mathbf{959 \pm 8})$	$30 \pm 3\,(14 \pm 4)$	$55 \pm 3\,(27 \pm 2)$	$0 \pm 0\,(0 \pm 0)$
True C2	$17 \pm 2\,(78 \pm 3)$	$\mathbf{906 \pm 3}\,(\mathbf{819 \pm 2})$	$48 \pm 2\,(58 \pm 2)$	$29 \pm 3\,(45 \pm 4)$
True C3	$17 \pm 1\,(61 \pm 1)$	$8 \pm 1\,(30 \pm 1)$	$\mathbf{955 \pm 3}\,(\mathbf{857 \pm 3})$	$20 \pm 2\,(52 \pm 2)$
True C4	$0 \pm 0\,(0 \pm 0)$	$18 \pm 2\,(49 \pm 2)$	$26 \pm 3\,(23 \pm 2)$	$\mathbf{956 \pm 3}\,(\mathbf{928 \pm 3})$

Table 2. Same as Table 1 but for winter months. The overall test accuracy is $93.8\% \pm 0.1\%$ (CNN4) and $86.6\% \pm 0.3\%$ (CNN2).

	Identified as C1	Identified as C2	Identified as C3	Identified as C4
True C1	$\mathbf{937 \pm 2}\,(\mathbf{783 \pm 3})$	$13\,(98 \pm 1)$	$37 \pm 1\,(20 \pm 1)$	$13 \pm 1\,(99 \pm 1)$
True C2	$71 \pm 2\,(28 \pm 2)$	$\mathbf{920 \pm 2}\,(\mathbf{951 \pm 2})$	$0 \pm 0\,(0 \pm 0)$	$9 \pm 1\,(21 \pm 1)$
True C3	$49 \pm 2\,(61 \pm 1)$	$0 \pm 0\,(0 \pm 0)$	$\mathbf{984 \pm 3}\,(\mathbf{822 \pm 2})$	$3 \pm 0\,(129 \pm 0)$
True C4	$23 \pm 0\,(4 \pm 0)$	$29 \pm 2\,(82 \pm 2)$	$37 \pm 2\,(3 \pm 1)$	$\mathbf{911 \pm 2}\,(\mathbf{911 \pm 3})$

In this paper we developed two CNNs, one with two convolutional layers (CNN2) and another one with four convolutional layers (CNN4). The effects of hyperparameters and other practical issues as well as scaling of the accuracy with the size of the training set are examined and discussed.

Results

Performance of CNN
Tables 1 and 2 show the test accuracies of CNN2 and CNN4 for the summer and winter months, respectively. The CNN4 has the accuracy of $93.3\% \pm 0.2\%$ (summer) and $93.8\% \pm 0.1\%$ (winter) while CNN2 has the accuracy of $89.0\% \pm 0.3\%$ (summer) and $86.6\% \pm 0.3\%$ (winter). The reported accuracies are the mean and standard deviation of the accuracies of the 5 sets in the TS. The $4\% - 7\%$ higher accuracy of the deeper net, CNN4, comes at the price of higher computational demands (time and memory) because of the two additional convolutional layers, however, the robust test accuracy of $\sim 93\%$ is significant for the complex patterns studied here.

Deep CNNs are vulnerable to overfitting as the large number of parameters can lead to perfect accuracy on the training samples while the trained CNN fails to generalize and accurately classify the new unseen samples because the CNN has memorized, rather than learned, the classes. In order to ensure that the reported high accuracies of CNNs here are not due to overfitting, during the training phase, a randomly chosen validation set (which does not have any overlap with TS or TR) was used to tune the hyperparameters (see Data and Methods). For each case in Tables 1 and 2, the reported test accuracy is approximately equal to the training accuracy after the network converges, which along with small standard deviations among the 5 independent sets in the TS, indicate that the classes have been learned rather than overfitted. It should be mentioned that for this data with the TR of size $N \leq 12000$, we have found that overfitting occurs if more than 4 convolutional layers are used.

Scaling of the test accuracy with the size of the training set
An important practical question that many ask before investing in labeling data and developing their CNN algorithm is "how much data do I need to get reasonable accuracy with CNN?". However, a theoretical understanding of the bound or scaling of CNNs' accuracy based on the number of the training samples or number of tunable parameters of the network is currently unavailable[43]. Given the abundance of the labeled samples in our dataset, it is an interesting experiment to examine how the test accuracy of CNNs scales with the size of the TR, N. Figure 4 shows that the test accuracy of CNN2 and CNN4 scales monotonically but nonlinearly with N for summer and winter months. With $N = 500$ (125 training samples per cluster), the test accuracy of CNN4 is around 64%. The accuracy jumps above 80% with $N = 1000$ and then increases to above 90% as N is increased to 8000. Further increasing N to 12000 will slightly increase the accuracy to 93%. The accuracy of CNN2 qualitatively shows the same behavior, although it is consistently lower than the accuracy of CNN4 for the same N. While

the empirical scaling presented here is very likely problem dependent and cannot replace a theoretical scaling, it provides an example of how the test accuracy might depend on the size of the TR.

Incorrectly classified patterns

While the results presented above show an outstanding performance by CNN and test accuracy of $\sim 93\%$ with $N = 12000$ for CNN4, the cluster indices of a few hundred testing samples (out of the 4000) have been incorrectly identified. From visually comparing examples of correctly and incorrectly identified patterns, inspecting the cluster centers in Figure 1, or examining the results of Tables 1 and 2, it is not easy to understand why patterns from some clusters have been more (or less) frequently mis-classified. For example, in summer months using CNN4, patterns in cluster C2 (C4) are the most (least) frequently mis-classified. Patterns in C2 are most frequently mis-identified to belong to C3 (48 samples) while patterns in C3 are rarely mis-identified to belong to C2 (8). There are many examples of such asymmetries in mis-classification in Tables 1 and 2, although there are some symmetric examples too, most notably no mis-classification between C1 and C4 in summer or C2 and C3 in winter. It should be also noted that while CNN4 has robustly better overall test accuracy compared to CNN2 for summer/winter or as N changes, it may not improve the accuracy for every cluster (e.g. 915 C2 samples are correctly identified by CNN4 for summer months compared to 959 by CNN2). Visual inspection of cluster centers does not provide much clues on which clusters might be harder to re-identify or mix up, e.g., patterns in C2 in winter months are frequently (71 samples) mis-classified as C1 while rarely mis-classified as C3 (0) or C4 (9 samples) even though the cluster center of C2, which has a notable ridge over the eastern Pacific ocean and a low-pressure pattern over north-eastern Canada, is (visually) distinct from the cluster center of C1 or C3 but resembles that of C4.

While understanding *how* a CNN learns or *why* some patterns are identified and some are mis-identified can be of great interest for many applications, particularly those involved addressing a scientific problem, answering such questions is not straightforward with the current understanding of deep learning[44]. In the results presented here, there are two potential sources of inaccuracy: imperfect learning and improperly labeled patterns. The former can be a result of unoptimized choice of the hyperparameters or insufficient number of training samples. As discussed in Data and Methods, we have explored a range of hyperparameters, manually for some and using optimization algorithms such as ADAM[45] for others. Still there might be room for further systematic optimization and improvement of the test accuracy. The results of Figure 4 suggest that increasing N would have a small effect on the test accuracy. Training CNN4 for summer with $N = 18000$ increases the best test accuracy from 93.3% (obtained with $N = 12000$) to just 94.1%. These results suggest that the accuracy might be still further improved, although very slowly, by increasing N.

Another source of inaccuracy might be related to how the patterns are labeled using the K-means cluster indices. The K-means algorithm is deterministic and assigns each pattern to one and only one cluster index. In data that have well-defined classes, the patterns in each cluster are very similar to each other (high cohesion) and dissimilar from patterns in other clusters (well separated). However, in chaotic, correlated, spatio-temporal data, such as those studied here, some patterns might have similarities to more than one cluster, however, the K-means algorithm assigns them to just one (the closest) cluster. As a result, two patterns that are very similar might end up in two different clusters and thus assigned different labels. The presence of such borderline cases in the TR can degrade the learning process for CNN and their presence in the TS can reduce the test accuracy. The silhouette value s is a measure often used to quantify how a pattern is similar to its own cluster and separated from the patterns in other clusters[46]. Large, positive values of s indicate high cohesion and strong separation, while small and in particular negative values indicate the opposite.

To examine whether part of the inaccuracy in the testing phase is because of borderline cases, in Table 3 we show the percentage of samples correctly classified or incorrectly classified for ranges of high and negative silhouette values. The results indicate that poorly clustered (i.e. borderline) patterns, e.g. those with $s < 0$, are more frequently mis-classified compared to well-clustered patterns, e.g., those with $s > 0.4$ (11% versus 4.8%). This analysis suggests that part of the 6.7% testing error of CNN4 for summer months might be attributed to poor clustering and improper labeling (one could remove samples with low s from the TR and TS, but here we chose not to in order to have a more challenging task for the CNNs).

Note that soft clustering methods (e.g. fuzzy c-means clustering[47]) in which a pattern can be assigned to more than one cluster might be used to overcome the aforementioned problem if it becomes a significant source of inaccuracy. In any case, one has to ensure that the labels obtained from the unsupervised clustering technique form a learnable set for the CNN and be aware of the potential inaccuracies arising from poor labeling alone.

Discussion

The unsupervised auto-labeling strategy proposed here facilitates exploring the capabilities of CNNs in studying problems in climate and environmental sciences. The method can be applied to any spatio-temporal data and allows one to examine the power and limitations of different CNN architectures and scaling of their performance with the size of the training dataset for

Table 3. Percentage of samples correctly classified or incorrectly classified for different ranges of silhouette values, s. Silhouette values, by definition, are between -1 and 1 and high (low and particularly negative) values indicate high (low) cohesion and strong (weak) separation. Percentages show the fraction of patterns in a given range of silhouette values. The samples are from summer months and for CNN4.

Silhouette value	$s < 0$	$s > 0.2$	$s > 0.4$
Correctly identified	89.0%	94.2%	95.2%
Incorrectly identified	11.0%	5.8%	4.8%

each type of data before further investing in labeling the patterns to address specific scientific problems, e.g. to study patterns that cause heat waves or extreme precipitation.

The analysis conducted on daily large-scale weather patterns shows the outstanding performance of CNNs in identifying patterns in chaotic, multi-scale, non-stationary, spatio-temporal data with minimal pre-processing. Building on the promising results of previous studies[30, 31], our analysis goes beyond their binary classifications and shows over 90% test accuracy for 4-cluster classification once there are at least 2000 training samples per cluster.

The promising capabilities of CNNs in identifying complex patterns in climate data further opens frontiers for prediction of some weather, and in particular extreme weather, events using deep learning methods. Techniques such as recurrent neural networks (RNNs) with long short-term memory (LSTM) and tensor-train RNNs have shown encouraging skills in predicting time series in chaotic systems[48, 49]. Coupling CNNs with these techniques can potentially provide powerful tools for spatio-temporal prediction; e.g., a convolutional LSTM network has been recently implemented for precipitation nowcasting[50].

Data and Methods

Data from the Large Ensemble (LENS) Community Project

We use data from the publicly available Large Ensemble (LENS) Community Project[40], which consists of a 40-member ensemble of fully-coupled atmosphere-ocean-land-ice Community Earth System Model version 1 (CESM1) simulations at the horizontal resolution of $\sim 1^\circ$. The same historical radiative forcing from 1920 to 2005 is used for each member; however, small, random perturbations are added to the initial state of each member to create an ensemble. We focus on daily averaged geopotential height at 500 hPa (Z500). Z500 isolines are approximately the streamlines of the large-scale circulation at mid-troposphere and are often used to represent weather patterns[41]. We focus on Z500 from 1980 to 2005 for the summer months of June-August (92 days per summer) for all 40 ensemble members (total of 95680 days) over North America, $30^\circ - 90^\circ$ north and $200^\circ - 315^\circ$ east (resulting in 66×97 latitude–longitude grid points). Similarly, for winter we use the same 26 years of data for the months of December, January, and February (90 days per winter and a total of 95508 days).

Clustering of weather patterns

The daily Z500 patterns over North America are clustered for each season into $n = 4$ classes. Following Vigaud et al.[8], first, an EOF analysis is performed on the data matrix of zonal-mean-removed Z500 anomalies and the first 22 principal components (PCs), which explain 95% of the variance, are kept for clustering analysis. The K-means algorithm[38] is used on these 22 PCs and repeated 1000 times with new initial cluster centroid positions, and a cluster index $k = 1, 2, 3$ or 4 is assigned to each daily pattern.

It should be noted that the number of clusters $n = 4$ is not chosen as an optimal number, which might not even exist for these complex, chaotic, spatio-temporal data[13]. Rather, for the purpose of the analysis here, the chosen n should be large enough such that the cluster centers are reasonably distinct and there are several clusters to re-identify to evaluate the CNNs in a challenging multi-class classification problem, yet, small enough such that there are enough samples per cluster for training and testing.

Labeling and up/down-samplings

Once the cluster index for each daily pattern is computed, the full Z500 daily patterns are labeled using these indices. We focus on the full Z500 fields, rather than the anomalies, for several reasons: (1) The differences between the patterns from different clusters are more subtle in the full Z500 compared to the anomalous Z500 fields; (2) The full Z500 fields contain all the complex, temporal variabilities and non-stationarity resulting from ocean-atmosphere coupling and changes in the radiative forcing while some of these variabilities might be removed by computing the anomalies; As a result of (1) and (2), re-identifying the cluster indices in the full Z500 fields provide a more challenging test for CNNs; (3) One hopes to use CNNs with no or minimal pre-processing of the data, consequently, we focus on the direct output of the climate model, i.e., full Z500 field, rather than the pre-processed anomalies.

In our algorithm, the only pre-processing conducted on the data is the up-sampling/down-sampling shown in Figure 3. The down-sampling step is needed to remove the small-scale, transient features of the chaotic, multi-scale atmospheric circulation from the learning/testing process. Inspecting the cluster centers in Figure 1 shows that the main differences between the four clusters are in large-scale. If the small-scale features, which are associated with processes such as baroclinic instability, are not removed via down-sampling, the CNN will try to learn the distinction between these features in different classes, which is futile as these features are mostly random. We have found in our analysis that without the down-sampling step, we could not train the CNN using a simple random normal initialization of the kernel weights (if instead of random initialization, a selective initialization method such as Xavier[51] is used, the network can be trained for the full-sized images although the test accuracy remains low due to overfitting on small-scale features.) The need for down-sampling in applications of CNNs to multi-scale patterns has been reported previously in other areas[52]. In the applications that involve the opposite case, i.e. when the small-scale features are of interest and have to be learned, techniques such as localization can be used[53].

Note that although Z500 is a scalar field, here we have used the three channels of RGB to represent it because we are focusing on only one variable. In the future applications, when several variables are studied together, each channel can be used to represent one scalar field, e.g. temperature and/or components of the velocity vector.

Convolutional Neural Network (CNN)

The CNN is developed using the Tensorflow library[54] following the Alex Net architecture[33]. We have trained and tested two CNNs: one with two convolutional layers, named CNN2, and a deeper one with 4 layers, called CNN4.

CNN2

The shallow neural network has two convolutional layers with 16 and 32 filters, respectively. Each filter has a kernel size of 5×5. In each convolutional layer, zero padding around the borders of images is used to maintain the size before and after applying the filters. Each convolutional layer is followed with a ReLU activation function and a max-pooling layer that has a kernel size of 2×2 and stride of 1 (stride is the number of pixels the filter shifts over in the pooling layer[29]). The output feature map is $7 \times 7 \times 64$ which is fed into a fully connected neural network with 1024 neurons. The cross entropy cost function is accompanied by a L_2 regularization term with a hyperparameter λ. Furthermore, to prevent overfitting, dropout regularization with hyperparameter p has been used in the fully connected layer. An adaptive learning rate α, a hyperparameter, is implemented through the ADAM optimizer[45]. The final output is the probability of the input pattern belonging to each cluster. A softmax layer assigns the pattern to the cluster index with the highest probability.

CNN4

The deeper neural network, CNN4, is the same as CNN2, except that there are four convolutional layers, which have $8, 16, 32$ and 64 filters, respectively (Figure 2). Only the last two convolutional layers are followed by max-pooling layers.

Training, validating, and testing procedures

For the case with $N = 12000$, 3000 labeled images from each of the four clusters is selected randomly (the TR set). Separately, 4 validation datasets, each with 1000 samples per cluster are randomly selected. For the testing set (TS), 5 datasets, each with 1000 samples per cluster, are randomly selected. The TR, validation sets, and TS have no overlap. The equal number of samples from each cluster prevents class imbalance in training and testing.

In the training phase, the images and their labels, in randomly shuffled batches of size 32, are inputted into the CNN and hyperparameters α, λ, and p are varied until small loss and high accuracy are achieved. Figure 5 shows examples of how loss and accuracy vary with epochs for properly and improperly tuned CNNs. Note that only an initial value of α is specified, which is then optimized using the ADAM algorithm. Once the CNN is properly tuned, the 4 validation sets are used to check the accuracy of CNN in re-identifying the cluster indices. If the accuracy is not high, λ and p are varied manually and training/validation is repeated until they both have similarly high accuracy. We found the best test accuracy with the hyperparameters shown in Figure 5(a)-(b). Furthermore, we explored the effect of other hyperparameters such as the number of convolutional layers (from 2 to 8) and the kernel sizes (in the range of 5×5 to 11×11) in the convolutional layers on the performance of CNN for this dataset. We found that a network with more than 4 convolutional layers overfits on 12000 samples thus producing test accuracy lower than what is reported for CNN4 in Tables 1 and 2. Again, the best test accuracy is found with the architecture shown in Figure 2 and described above.

In the testing phase, the best trained CNN is applied on the 5 datasets of TS once. The mean and standard deviation of the computed accuracy among these 5 datasets are reported in Tables 1 and 2.

For the cases with $N = 500$ to 8000, conducted to study the effect of the size of the training set N on the performance of CNN, $N/4$ labeled images from each of the four clusters is selected randomly and used to train the CNN while testing is done on $N/8$ (to the nearest integer) images from each class.

Alternative approach: Applying CNN on data matrix rather than images

While CNNs are often used on images, they can be used directly to find features in the data matrices as well. For example, we can get the same accuracy as the CNN applied on images with CNN applied on a data matrix of labeled patterns. In such a data matrix X, each column contains the full Z500 over 97×66 grid points for each day. The CNN is applied to X, although the best results are obtained with a CNN whose architecture is slightly different from the one applied to images. In this case, the four convolutional layers have 8, 8, 16 and 32 filters while the fully connected layer has 200 neurons.

References

1. Mo, K. & Ghil, M. Cluster analysis of multiple planetary flow regimes. *J. Geophys. Res. Atmospheres* **93**, 10927–10952 (1988).

2. Thompson, D. W. J. & Wallace, J. M. The Arctic Oscillation signature in the wintertime geopotential height and temperature fields. *Geophys. Res. Lett.* **25**, 1297–1300 (1998).

3. Smyth, P., Ide, K. & Ghil, M. Multiple regimes in northern hemisphere height fields via mixturemodel clustering. *J. Atmospheric Sci.* **56**, 3704–3723 (1999).

4. Bao, M. & Wallace, J. M. Cluster analysis of Northern Hemisphere wintertime 500-hPa flow regimes during 1920–2014. *J. Atmospheric Sci.* **72**, 3597–3608 (2015).

5. Sheshadri, A. & Plumb, R. A. Propagating annular modes: Empirical orthogonal functions, principal oscillation patterns, and time scales. *J. Atmospheric Sci.* **74**, 1345–1361 (2017).

6. Grotjahn, R. *et al.* North American extreme temperature events and related large scale meteorological patterns: a review of statistical methods, dynamics, modeling, and trends. *Clim. Dyn.* **46**, 1151–1184 (2016).

7. Barnes, E. A., Slingo, J. & Woollings, T. A methodology for the comparison of blocking climatologies across indices, models and climate scenarios. *Clim. Dyn.* **38**, 2467–2481 (2012).

8. Vigaud, N., Ting, M., Lee, D.-E., Barnston, A. G. & Kushnir, Y. Multiscale variability in North American summer maximum temperatures and modulations from the North Atlantic simulated by an AGCM. *J. Clim.* **31**, 2549–2562 (2018).

9. Corti, S., Molteni, F. & Palmer, T. N. Signature of recent climate change in frequencies of natural atmospheric circulation regimes. *Nature* **398**, 799 (1999).

10. Barnes, E. A., Dunn-Sigouin, E., Masato, G. & Woollings, T. Exploring recent trends in Northern Hemisphere blocking. *Geophys. Res. Lett.* **41**, 638–644 (2014).

11. Horton, D. E. *et al.* Contribution of changes in atmospheric circulation patterns to extreme temperature trends. *Nature* **522**, 465 (2015).

12. Hassanzadeh, P. & Kuang, Z. Blocking variability: Arctic Amplification versus Arctic Oscillation. *Geophys. Res. Lett.* **42**, 8586–8595 (2015).

13. Fereday, D. R., Knight, J. R., Scaife, A. A., Folland, C. K. & Philipp, A. Cluster analysis of North Atlantic–European circulation types and links with tropical Pacific sea surface temperatures. *J. Clim.* **21**, 3687–3703 (2008).

14. McKinnon, K. A., Rhines, A., Tingley, M. P. & Huybers, P. Long-lead predictions of eastern United States hot days from Pacific sea surface temperatures. *Nat. Geosci.* **9**, 389 (2016).

15. Anderson, B. T., Hassanzadeh, P. & Caballero, R. Persistent anomalies of the extratropical Northern Hemisphere wintertime circulation as an initiator of El Niño/Southern Oscillation events. *Sci. Reports* **7**, 10145 (2017).

16. Zhang, J. P. *et al.* The impact of circulation patterns on regional transport pathways and air quality over Beijing and its surroundings. *Atmospheric Chem. Phys.* **12**, 5031–5053 (2012).

17. Souri, A. H., Choi, Y., Li, X., Kotsakis, A. & Jiang, X. A 15-year climatology of wind pattern impacts on surface ozone in Houston, Texas. *Atmospheric Res.* **174**, 124–134 (2016).

18. Monahan, A. H., Fyfe, J. C., Ambaum, M. H. P., Stephenson, D. B. & North, G. R. Empirical orthogonal functions: The medium is the message. *J. Clim.* **22**, 6501–6514 (2009).

19. Woollings, T. *et al.* Blocking and its response to climate change. *Curr. Clim. Chang. Reports* **4**, 287–300 (2018).

20. Schneider, T., Lan, S., Stuart, A. & Teixeira, J. Earth system modeling 2.0: A blueprint for models that learn from observations and targeted high-resolution simulations. *Geophys. Res. Lett.* **44**, 12,396–12,417 (2017).

21. Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G. & Yacalis, G. Could machine learning break the convection parameterization deadlock? *Geophys. Res. Lett.* **45**, 5742–5751 (2018).

22. Brenowitz, N. D. & Bretherton, C. S. Prognostic validation of a neural network unified physics parameterization. *Geophys. Res. Lett.* **45**, 6289–6298 (2018).

23. Rasp, S., Pritchard, M. S. & Gentine, P. Deep learning to represent subgrid processes in climate models. *Proc. Natl. Acad. Sci. United States Am.* **115**, 9684–9689 (2018).

24. O'Gorman, P. A. & Dwyer, J. G. Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change and extreme events. *J. Adv. Model. Earth Syst.* **10** (2018).

25. Rasp, S. & Lerch, S. Neural networks for post-processing ensemble weather forecasts. *arXiv preprint arXiv:1805.09091* (2018).

26. Dueben, P. D. & Bauer, P. Challenges and design choices for global weather and climate models based on machine learning. *Geosci. Model. Dev.* **11**, 3999–4009 (2018).

27. Nooteboom, P. D., Feng, Q. Y., López, C., Hernández-García, E. & Dijkstra, H. A. Using network theory and machine learning to predict El Niño. *Earth Syst. Dyn.* **9**, 969–983 (2018).

28. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436 (2015).

29. Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. *Deep learning*, vol. 1 (MIT press Cambridge, 2016).

30. Liu, Y. *et al.* Application of deep convolutional neural networks for detecting extreme weather in climate datasets. *arXiv preprint arXiv:1605.01156* (2016).

31. Racah, E., Beckham, C., Mahranj, T., Prabhat & Pal, C. Semi-supervised detection of extreme weather events in large climate datasets. *arXiv preprint arXiv:1612.02095* (2016).

32. Karpatne, A., Ebert-Uphoff, I., Ravela, S., Babaie, H. A. & Kumar, V. Machine learning for the geosciences: Challenges and opportunities. *IEEE Transactions on Knowl. Data Eng.* (2018).

33. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 1097–1105 (2012).

34. Williams, P. D. *et al.* A census of atmospheric variability from seconds to decades. *Geophys. Res. Lett.* **44**, 11,201–11,211 (2017).

35. Ma, D., Hassanzadeh, P. & Kuang, Z. Quantifying the eddy–jet feedback strength of the annular mode in an idealized gcm and reanalysis data. *J. Atmospheric Sci.* **74**, 393–407 (2017).

36. Kosaka, Y. & Xie, S.-P. Recent global-warming hiatus tied to equatorial Pacific surface cooling. *Nature* **501**, 403 (2013).

37. Prabhat, S. B. *et al.* TECA: Petascale pattern recognition for climate science. In *International Conference on Computer Analysis of Images and Patterns*, 426–436 (2015).

38. Lloyd, S. Least squares quantization in pcm. *IEEE Transactions on Inf. Theory* **28**, 129–137 (1982).

39. Cheng, X. & Wallace, J. M. Cluster analysis of the Northern Hemisphere wintertime 500-hPa height field: Spatial patterns. *J. Atmospheric Sci.* **50**, 2674–2696 (1993).

40. Kay, J. E. *et al.* The Community Earth System Model (CESM) large ensemble project: A community resource for studying climate change in the presence of internal climate variability. *Bull. Am. Meteorol. Soc.* **96**, 1333–1349 (2015).

41. Holton, J. R. & Hakim, G. J. *An introduction to dynamic meteorology*, vol. 88 (Academic press, 2012).

42. Girshick, R., Donahue, J., Darrell, T. & Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 580–587 (2014).

43. Zhang, C., Bengio, S., Hardt, M., Recht, B. & Vinyals, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530* (2016).

44. Lin, H. W., Tegmark, M. & Rolnick, D. Why does deep and cheap learning work so well? *J. Stat. Phys.* **168**, 1223–1247 (2017).

45. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

46. Rousseeuw, P. J. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **20**, 53–65 (1987).

47. Bezdek, J. C., Ehrlich, R. & Full, W. Fcm: The fuzzy c-means clustering algorithm. *Comput. & Geosci.* **10**, 191–203 (1984).

48. Vlachas, P. R., Byeon, W., Wan, Z. Y., Sapsis, T. P. & Koumoutsakos, P. Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks. *Proc. R. Soc. A* **474**, 20170844 (2018).

49. Yu, R., Zheng, S., Anandkumar, A. & Yue, Y. Long-term forecasting using tensor-train RNNs. *arXiv preprint arXiv:1711.00073* (2017).

50. Xingjian, S. *et al.* Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems*, 802–810 (2015).

51. He, K., Zhang, X., Ren, S. & Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, 1026–1034 (2015).

52. Lin, T.-Y. *et al.* Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, 4 (2017).

53. Kim, S. *et al.* Resolution reconstruction of climate data with pixel recursive model. In *IEEE International Conference on Data Mining Workshops*, 313–321 (2017).

54. Abadi, M. *et al.* Tensorflow: a system for large-scale machine learning. In *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation*, vol. 16, 265–283 (2016).

Acknowledgements
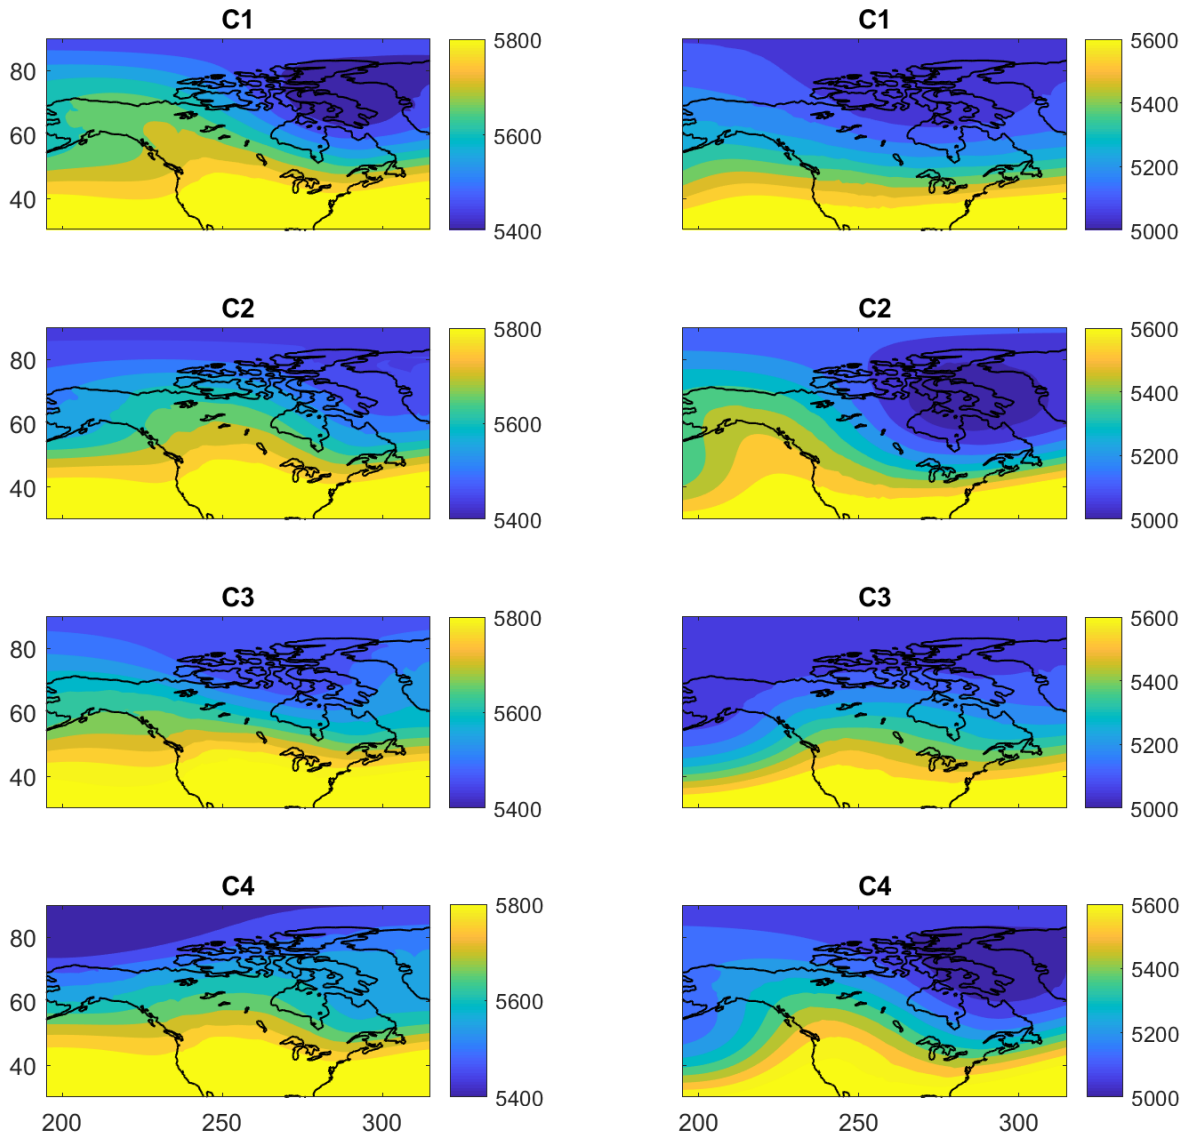
Figure 1. Centers of the four K-means clusters in terms of the full Z500 field (with unit of meters) for summer months, June-August (left column) and for winter months, December-February (right column). The K-means algorithm finds the cluster centers based on *a priori* specified number of clusters n ($= 4$ here) and assigns each daily pattern to the closest cluster center based on Euclidean distances. The assigned cluster indices are used as labels for training/testing CNNs. Note that K-means clustering is performed on daily zonal-mean-removed Z500 anomalies projected onto their first 22 EOFs, but the cluster indices are used to label the full Z500 patterns to minimize pre-processing and retain the complex temporal variabilities of the Z500 field (see Data and Methods for further discussions).
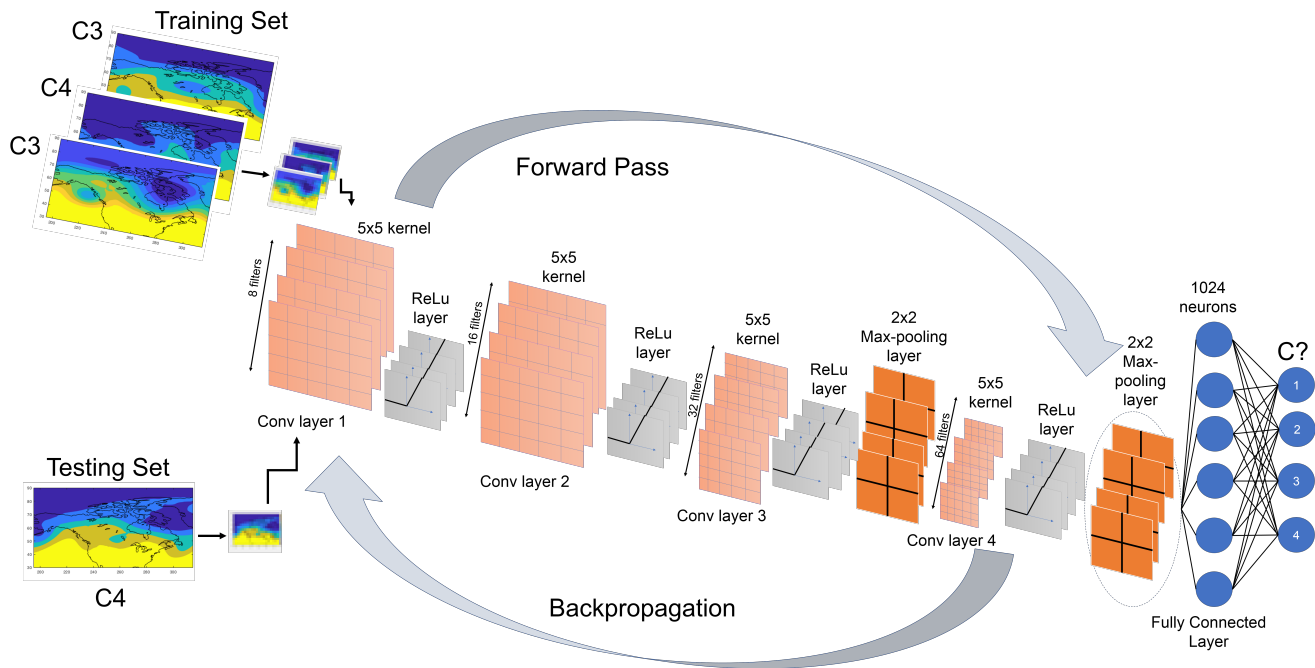
Figure 2. The architecture of CNN4, which has 4 convolutional layers that have $8, 16, 32$ and 64 filters, respectively. Each filter has a kernel size of 5×5. Filters of the max-pooling layer have a kernel size of 2×2. The convolutional layers at the beginning capture the low-level features while the latter layers would pick up the high level features[42]. Each convolution step is followed by the ReLU layer that introduces nonlinearity in the extracted features. In the last two layers, a max-pooling layer after the ReLU layer retains only the most dominant features in the extracted feature map while inducing translation and scale invariance. These extracted feature maps are then concatenated into a single vector which is connected to a fully connected neural network with 1024 neuron. The output is the probability of each class. The input images into this network have been first down-sampled using bi-cubic interpolation to only retain the large-scale features in the circulation patterns (Figure 3).
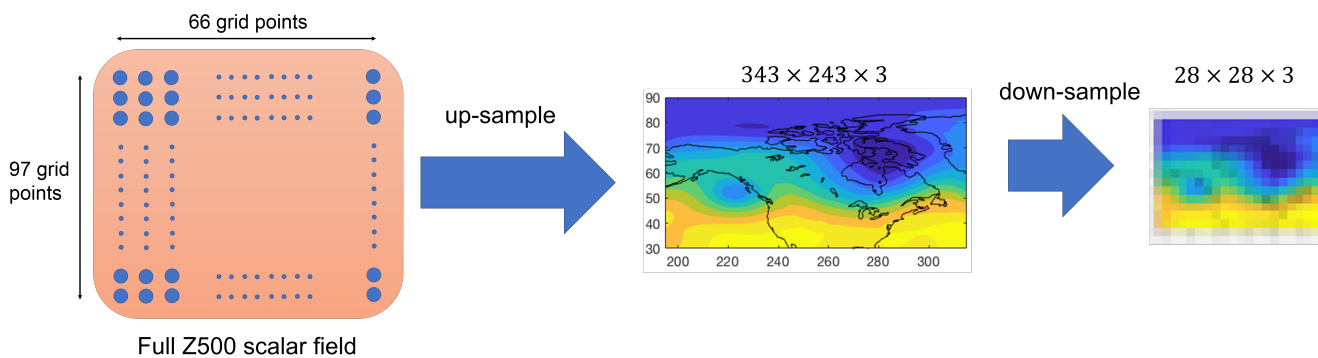
Figure 3. Schematic of the up-sampling and down-sampling steps. Each daily full Z500 pattern, which is on a 66×97 latitude-longitude grid, is converted to a contour plot represented by a RGB image of size 342×243 pixels with 3 channels representing red, green, and blue. This up-sampled image is then down-sampled to an image of size $28 \times 28 \times 3$ using bi-cubic interpolation and further standardized by subtracting the mean and dividing by the standard deviation of the pixel intensities. These images are the inputs to CNN for training or testing. The down-sampling step is used to remove redundant features at small scales from each sample. Trying to learn such small features, which are mostly random, can result in overfitting of the network (see Data and Methods for further discussions).
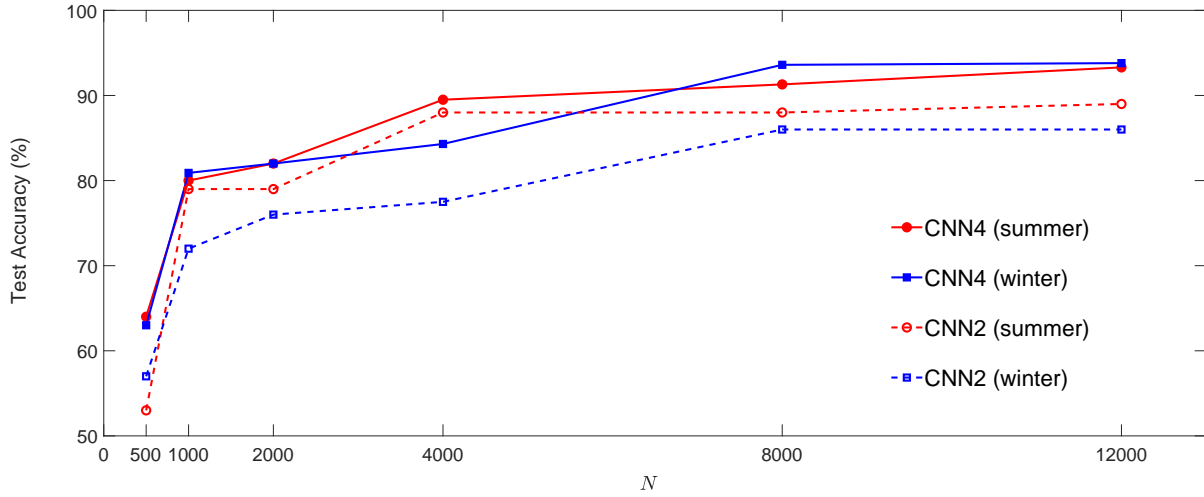
Figure 4. Test accuracy of CNN4 and CNN2 as a function of the size of the training set N. To avoid class imbalance, $N/4$ samples per cluster are used. A $3 : 1$ ratio between the number of samples per cluster in the training and testing sets are maintained.
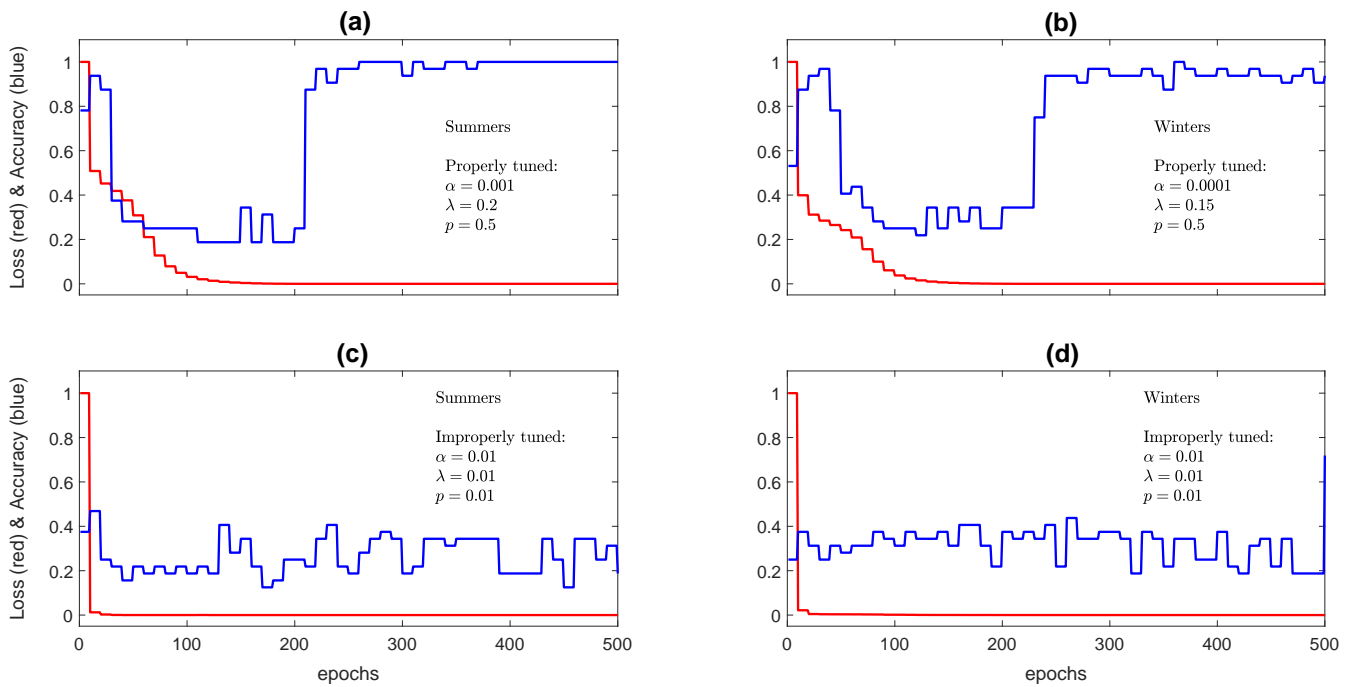
Figure 5. Examples of how loss and accuracy change with epochs during training for CNN4 for properly tuned and improperly tuned CNNs. Loss is measured as cross entropy (CE) normalized by its maximum value while the training accuracy is measured by the number of training samples correctly identified at the end of each epoch. Hyperparameters α, λ, and p are, respectively, the initial learning rate, regularization constant, and dropout probability. (**a**) $\alpha = 0.001$, $\lambda = 0.2$ and $p = 0.5$ for summers with the test accuracy of 93.3%. (**b**) $\alpha = 0.001$, $\lambda = 0.15$ and $p = 0.5$ for winters with the test accuracy of 93.8%. (**c**) $\alpha = 0.01$, $\lambda = 0.01$ and $p = 0.01$ for summers with the test accuracy of 25%. (**d**) $\alpha = 0.01$, $\lambda = 0.01$ and $p = 0.01$ for winters with the test accuracy of 60%). Several kernel sizes were tried and it was found that 5×5 kernel size gives the best validation accuracy and consequently the best test accuracy.