

Available online at www.sciencedirect.com



Procedia Computer Science 9 (2012) 887 - 896

Procedia Computer Science

International Conference on Computational Science, ICCS 2012

Sub-daily Statistical Downscaling of Meteorological Variables Using Neural Networks

Jitendra Kumar^{a,1}, Bjørn-Gustaf J. Brooks^b, Peter E. Thornton^c, Michael C. Dietze^b

^aComputer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA ^bDepartment of Plant Biology, University of Illinois, Urbana-Champaign, IL, USA ^cEnvironmental Sciences Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA

Abstract

A new open source neural network temporal downscaling model is described and tested using CRU-NCEP reanalysis and CCSM3 climate model output. We downscaled multiple meteorological variables in tandem from monthly to sub-daily time steps while also retaining consistent correlations between variables. We found that our feed forward, error backpropagation approach produced synthetic 6 hourly meteorology with biases no greater than 0.6% across all variables and variance that was accurate within 1% for all variables except atmospheric pressure, wind speed, and precipitation. Correlations between downscaled output and the expected (original) monthly means exceeded 0.99 for all variables, which indicates that this approach would work well for generating atmospheric forcing data consistent with mass and energy conserved GCM output. Our neural network approach performed well for variables that had correlations to other variables of about 0.3 and better and its skill was increased by downscaling multiple correlated variables together. Poor replication of precipitation intensity however required further post-processing in order to obtain the expected probability distribution. The concurrence of precipitation events with expected changes in subordinate variables (*e.g.*, less incident shortwave radiation during precipitation events) were nearly as consistent in the downscaled data as in the training data with probabilities that differed by no more than 6%. Our downscaling approach requires training data at the target time step and relies on a weak assumption that climate variability in the extrapolated data is similar to variability in the training data.

Keywords: statistical downscaling, artificial neural networks, CRU-NCEP, CCSM, climate

1. Introduction

Many Earth science disciplines require data at time steps that differ from the available measurements. Remotely sensed satellite products provide superb spatial coverage but typically have sampling frequencies of days to weeks. General circulation model (GCM) output often have global coverage but variables of interest are not always available at the desired time step, which can be an obstacle for application to model intercomparisons [*e.g.*, 1, MsTMIP,

Email addresses: jkumar@climatemodeling.org (Jitendra Kumar), bjorn@climatemodeling.org (Bjørn-Gustaf J. Brooks),

thorntonpe@ornl.gov (Peter E. Thornton), mdietze@life.illinois.edu (Michael C. Dietze)

¹Corresponding author

PalEON]. Two general solutions for generating climate/meteorological data at higher sampling frequencies exist: 1) Dynamical downscaling using a mesoscale model (*e.g.*, PRECIS, WRF) to simulate local effects at sub-daily time steps that are still informed by background climate [2], or 2) statistical downscaling to extrapolate data to finer sampling frequencies [3, 4]. Running a nested mesoscale model (option 1) over the entire domain at short time steps is not always practical. On the other hand statistical downscaling (option 2) may not be ideal because it relies on assumed relationships [*cf.*, 5] that may not hold true when extrapolating to new states not realized in the source data. Hybrid methods, as in statistical-dynamical downscaling [6, 7, 8], have been developed to require less computing time than dynamical downscaling while still producing three-dimensional climate state variables. Our focus in this paper however, is to describe a fast statistical downscaling procedure that does not rely on linear covariance relationships among variables [3, 9, 10], does not require additional atmospheric modeling [*e.g.*, 8], and is capable of temporally downscaling measurements based on observed patterns at the target sampling rate rather than imposing randomness to synthesize data [*e.g.*, 4]. This approach has advantages for downscaling over regions that have strong non-random short-term variability such as the North American 'monsoon' [11] and other cyclical shorter-term processes that feed back into climate.

1.1. Dynamical Downscaling

Dynamical downscaling uses process-based regional models to simulate atmospheric interactions within boundary conditions prescribed by a GCM. This permits direct modeling of the regional climate system. Dynamically down-scaled output has finer spatial and temporal resolution than GCM's and is more consistent with the observed climate at any given measurement station. The United Kingdom Hadley Centre Meteorological Office has for several years provided the software and support to downscale GCM projections to regional scales using a stand-alone desktop version of the Hadley Centre Regional Climate Model (PRESIS).

The National Centers for Environmental Prediction (NCEP) Regional Spectral Modeling (RSM) system provides publicly available downscaled climate products from a two component system that includes a low resolution global model coupled to higher resolution regional models [12, 13]. The data flow is one-way and is used to inform the RSM so that regional responses to large-scale climate can be simulated.

Combining regional and global climate models is a preferred approach when resources and time permit, because it provides locally specific forecasts that are informed by climate trends. However, this approach tends to have a higher computational cost and potentially requires many terabytes of storage space, along with computational and scientific expertise for model set-up.

1.2. Temporal Statistical Downscaling Methods

Temporal statistical downscaling methods for increasing sampling frequency most commonly work by dividing each time step into fractions of random size and distributions at the target sampling frequency. A number of studies have patterned their techniques on the early work by Richardson [3]. Weather generator approaches, such as WGEN [14] or weathergen (meteo.unican.es), operate by downscaling precipitation first to the target sampling frequency, then downscaling other variables based on their covariance with precipitation. Implementations such as LARS-WG [10] and WeaGETS are particularly flexible and can run on desktop computers. Typically the target sampling frequency is daily, although sub-daily implementations exist [5]. An neural networks approach has been applied by other researchers [15] to temporally downscale watershed drainage using GCM output with time lags.

Temporal statistical downscaling methods are generally fast but their downscaled output does not preserve the temporal and/or spatial coherence in the synoptic meteorology. For example, given a downscaling of precipitation using the methods above, the total mass of a moving front would be conserved in the downscaled product. However, that precipitation will be randomly distributed in rain events (like static) in the target time series when in fact the precipitation may have been more continuous temporally and spatially. The significance of this issue for such methods grows as the gap between the input and target sampling frequency increases. In other words, high frequency weather patterns are not present in downscaled product. Furthermore, spatial coherence at the target time step may also be randomized if each grid cell is downscaled independently and randomly.

Since statistical models rely on ancillary data for training, spatial statistical methods will project any shortcomings from the training data to the downscaled product. Furthermore training transfer models, like neural networks and multiple linear regressions, using averages can result in output that neglects important variability. Thus it is important to carefully select target data in training capable of conferring the desired spatial heterogeneity.

2. Methods

Our approach uses an artificial neural network (ANN) for pattern analysis on data sampled at the target time step. The algorithm develops and iteratively optimizes a transfer function capable of converting infrequently sampled data to a time series of synthetic data that preserves observed high frequency patterns at the target sampling rate. For example, our method can downscale monthly means to 6-hourly time steps by training the neural network on 6-hourly data and then applying it to monthly means. The target data provides 'expected' 6-hourly values, which need not necessarily cover the same time period as the 'input' monthly means to be downscaled provided that one is willing to make a weak assumption that sub-daily variability in the data requiring downscaling is similar to that observed in the expected data. In fact, in this paper we downscale monthly means can be downscaled, but the target (expected) data should be of sufficient length to allow the network to 'learn' using multiple instances of the desired output. This means that a network trained on 10 years of 6-hourly data may not contain sufficient examples of variability to result in synthetic output that captures the expected variance.

2.1. Artificial Neural Networks

Artificial Neural Network (ANN) is a mathematical model inspired by the structure and operation of biological neural networks. ANN is a densely interconnected network of independent adaptive processing units called *neurons*. Neurons in an ANN are arranged in a layered structure. While there are no connections between the neurons within a layer, they are connected to the neurons in the adjacent layer(s), Each input connection to a neuron has an associated adaptive weight to model the synaptic learning. A *multi-layer perceptron* (MLP) ANN consists of an input layer, an output layer and one or more hidden layer(s) (Figure 1). The input data are presented to the ANN at the input layer, which is processed forward through the hidden layers to calculate the output at the output layers. This process of one directional forward movement of information is called *feed forward method*.



Units Variable Short Name Downwelling longwave radiation lwdown W/m² Total precip. (stratiform+convective) precipf mm/6 hours Surface pressure Pa psurf Specific humidity kg/kg qair Downwelling shortwave radiation W/m^2 swdown 2 meter air temperature tair Κ wind speed wind m/s

Input layer Hidden layer Output layer Figure 1: Schematic of multilayer Artificial Neural Network.

 Table 1:
 Meteorological variables from CRU-NCEP

 datasets and their corresponding units

Activations of the input neurons are carried forward to the output layers through the connection weights. The output at a neuron is calculated by the passing weighted inputs through an activation function. A sigmoid activation function was used for the study presented here. Computed outputs at the output layer are compared with the expected values to calculate a sum squared error. The error is backpropagated through the network to update the connection weights using a general delta learning rule [Rumelhart *et al.*, 1986 16]. the generalized delta rule is based on gradient descent and changes the weights in proportion to the error derivative with respect to each weight. The feed forward and backpropagation training steps are repeated iteratively until convergence.

2.2. Data used in the study

The downscaling of climate data is frequently performed to generate relevant data for regional studies from Global Circulation Models (GCM). For this study GCM meteorological output was extracted from Community Climate System Model-3 (CCSM3) output. The particular simulation used here (b30.125.1150AD_2.cam2.h0.1150) was used to investigate the quasi-decadal volcanic aerosol forcing as a possible trigger for global climatic cooling for several large volcanic events near 1250 A.D. [17]. Archived monthly means from this fully coupled simulation (A.D. 1150–1700) included one hundred variables each with three to four dimensions. The size of each monthly mean file is

30 MB, and the total size for this simulation exceeds 200 GB. While both spatial and temporal downscaling of these data were performed, only temporal downscaling will be focused on in this paper.

No observed high frequency meteorological records are available for the last millennium. However, 6 hourly values covering the past 110 years are available from the CRU-NCEP dataset, which was used to develop neural network models for downscaling the monthly meteorological datasets to 6 hourly values. CRU-NCEP data were produced by [18] who combined two existing datasets, the Climate Research Unit (CRU) TS.3.1 climatology and the National Center for Environmental Prediction (NCEP) reanalysis product, which was derived from observed climatology. Table 1 shows the CRU-NCEP meteorological variables for which neural networks downscaling models were developed.

2.3. Neural Network Design for Temporal Downscaling

The design of ANN configuration is a critical step in development of the ANN model for the variable of interest. ANNs combine the nonlinear functions of the variables presented as inputs to model the output(s). Careful selection of types and number of input variables to the ANN model will have a strong bearing on the reliability of the model. Neural network models were developed using a time series of monthly mean meteorological variables from the CRU-NCEP dataset [18] as *predictors* and 6 hourly CRU-NCEP data as *predictands*. To place synthesized output from the neural network in the proper context a useful first step is to examine the relationships among variables. A cross correlation analysis was carried out for all 6 hourly and monthly time series of CRU-NCEP variables (Table 1). The upper triangular matrix in Table 2 shows the cross correlation coefficients for the monthly and 6 hourly time series of all seven meteorological variables. Although all variables used are physically interdependent and some exhibit particularly strong correlations they are not all equally related. This could be an important consideration when designing a downscaling network that combines multiple different variables as predictands in order to estimate a particular variable (*e.g.*, using 6 hourly time series of multiple variables to downscale precipitation monthly means).

To downscale each variable to 6-hourly values, monthly mean time series of all seven variables were used as inputs (predictors) to the neural network model. This approach accounts for the interdependency among meteorological variables by using the same set of mean monthly time series from all seven variables to downscale monthly mean input for each variable sequentially. In addition to strong cross correlations these meteorological variables are strongly autocorrelated over shorter windows with each value at any given time being strongly dependent on its recent previous values.

ANN models for temporal downscaling were developed to cover an entire year in order to train the network on sub-monthly to diel variability. Monthly mean values for all seven variables $(12 \times 7 = 84)$ of the year were used as input for the ANN model to predict one year of 6 hourly values for the output variable $(365 \times 4 = 1460)$. For each variable a number of ANN models with different hidden neuron configurations were trained using the error back-propagation method. Each variable dataset had different variability requiring that the ANN model be trained across a range of hidden neurons to determine the best configuration. For example, the best configuration for longwave radiation during January for 7 input predictors included 20 neurons in the hidden layer for 1460 samples in the output. Best configuration was determined across a series of hidden neuron trials with a single hidden layer. Table 3 shows the best configuration identified for the seven meteorological variables for the reported case study.

Special handling for precipitation data: Table 2 shows that precipitation is the least correlated of all seven variables. In general precipitation occurrence is difficult to forecast, but also the intensity of that precipitation obeys a different probability distribution from the other variables. We noticed that when downscaling monthly precipitation according to the method used by all other variables, precipitation occurrence was reliable but its distribution in the downscaled output was more like Poisson or a Gaussian distribution rather than the expected gamma distribution. This could be the result of the low autocorrelation on scales longer than about 12 hours and the low correspondence between precipitation intensity and intensity changes of other variables (*e.g.*, solar radiation, humidity). The day can be cloudy for example but that does not mean it will rain.

Previous studies have addressed this issue by simulating precipitation occurrence as a Markov chain random process, and by scaling the precipitation amount according to a gamma distribution [*cf.* 14, 9, 10, 19]. However, stochastic precipitation generation may not necessarily reproduce the regularity in precipitation occurrence (if any) or variation in dry-spell length that is observed at a location, which may have relevant ecosystem implications during

some seasons. The downscaling design used here is already capable of reproducing dominant patterns of precipitation regularity, however the gamma distribution of intensity was rarely achieved and required that we optimize the computed precipitation values by fitting them to the expected gamma distribution on a per month basis. Because the distribution of precipitation values for any single month was not necessarily gamma distributed (but all months obey a gamma distribution) gamma fitting was done iteratively for each moth using a random gamma number generator until a sequence matching the monthly mean precipitation value was achieved (see precipf.m in supplementary material).

3. Results

From the 110 years (1901-2010) of meteorological records available from CRU-NCEP datasets, 88 years (80% of the record) were used to train the ANN model while 22 years (20% of total record) of the data were selected for validation of the ANN downscaling model. While the downscaling was carried out for each $0.5^{\circ} \times 0.5^{\circ}$ grid cell, all results presented in this paper reflect one example grid cell ~35°N,75°W.

Table 2: Cross-correlation matrix of meteorological variables. The upper triangular matrix shows the cross correlations for Expected/Computed monthly mean time series of the variables (*predictor*). The lower triangular matrix shows the cross correlation for Expected/Computed 6 hourly time series of the variables (*predictand*).

	lwdown	precipf	psurf	qair	swdown	tair	wind	
lwdown	-	0.33/0.39	-0.42/-0.48	0.96/0.95	0.77/0.77	0.96/0.95	-0.90/-0.90	
precipf	0.23/0.07	-	-0.17/-0.17	0.33/0.40	0.20/0.25	0.28/0.33	-0.22/-0.31	
psurf	-0.34/-0.28	-0.23/-0.01	-	-0.44/-0.50	-0.53/-0.60	-0.47/-0.49	0.39/0.50	
qair	0.90/0.85	0.16/0.08	-0.24/-0.22	-	0.77/0.77	0.97/0.95	-0.88/-0.89	
swdown	0.14/0.16	0.11/0.04	-0.00/-0.03	0.22/0.24	-	0.82/0.81	-0.69/-0.70	
tair	0.89/0.83	0.12/0.06	-0.25/-0.22	0.93/0.88	0.37/0.37	-	-0.88/-0.89	
wind	-0.34/-0.36	0.00/-0.02	-0.12/-0.01	-0.36/-0.40	-0.05/-0.04	-0.38/-0.38	_	

3.1. Downscaling model validation

We began by investigating the neural network's ability to downscale output without interfering with the correlational relationships across meteorological variables. Results from the cross correlation analysis (see the lower triangle of 6 hourly correlations in Table 2) show that across six variables (precipitation excluded) the downscaled and expected correlations differed by no more than 0.11, and in most cases by 0.03 or less. Precipitation and surface pressure showed both the weakest correlations to other variables and the most difference between expected and computed correlations. The correlation for precipitation:longwave (0.23/0.07), and precipitation:pressure (-0.23/-0.01) were particularly divergent between the expected and computed series. The pressure:wind correlation was also small and noticeably different (0.12/-0.01).

Computed 6 hourly output showed greater similarity to the expected correlations for all variables other than precipitation and pressure. Long and shortwave radiation, specific humidity, air temperature and wind speed showed the strongest correlations to other variables in both expected and computed datasets, usually near r = 0.8 to 0.9. These suggest that over all variables the neural network is capable of replicating 6 hourly output meanwhile retaining correlations between all variables that have correlations larger than about 0.25. Variables with weak correlations tend to have correlations in the downscaled output that diverge from those in the original training data.

We next examined summary differences in averages and distributions between the downscaled (computed) and original (expected) 6 hourly data, which are listed in Table 3. Again paired variables exhibiting the weakest correlations tended to have larger bias values and error in standard deviation and skew. Percent bias was calculated as the percent difference between the computed and the expected mean values. Percent bias for downscaling models for all the variables was less than 0.7%. Wind and precipitation showed the largest bias with the values computed 6 hourly data both being overall slightly larger than the expected values, which is also reflected in the larger means for the computed counterparts.

Standard deviations in Table 3 show that variances between expected and computed were very consistent across variables, with the most different being surface pressure ($\sigma = 571.3/549.1$). Asymmetries between the expected and computed series are also very comparable. The largest difference is for precipitation, where both expected and computed are strongly skewed to the right of the mean, but the computed is slightly more skewed (5.5) toward large

positive values than the expected (5.1). An illustration of the divergence between individual computed values and expected values is given by the scatter plot in Figure 2.

Table 3: Summary for all variables comparing statistics between the original CRU-NCEP 6 hourly values (Expected) and the 6 hourly values computed by the ANN

Variabla	ANN Model	Expected			Computed				
variable	AININ MOULI	Mean	Std. dev	Skewness	Mean	Std. dev	Skewness	% Bias	RMSE
Long Wave Rad.	84-20-1460	344.90	46.04	-0.42	344.70	45.46	-0.36	-0.06	28.88
Precipitation	48-100-1460	0.87	2.25	5.12	0.87	2.49	5.52	0.60	3.24
Pressure	84-15-1460	100763.21	571.26	-0.01	100773.91	549.11	0.06	0.01	576.03
Humidity	84-15-1460	0.009	0.004	0.35	0.009	0.004	0.39	0.22	0.0016
Temperature	84-20-1460	289.32	8.55	-0.37	289.14	8.47	-0.74	-0.06	4.12
Short Wave Rad.	84-25-1460	194.96	213.74	0.62	195.97	211.90	0.61	0.52	45.31
Wind	84-25-1460	4.01	2.08	0.67	4.04	1.89	0.67	0.69	1.97

3.2. Case Study tests for consistency across variables

An important consideration for downscaling meteorological data is that scenarios of synthesized data should be consistent with atmospheric understanding of the effects of cloud cover and precipitation on radiation, humidity, and other variables. We first tested the validity of these relationships in our computed data against the expected reanalysis data by visually inspecting a series of case study windows centered on precipitation events. Figure 3 shows multiple precipitation events alongside corresponding time series for other variables. Our expectation for physically consistent data is that obvious precipitation events should be linked to increased longwave radiation due to cloud cover, elevated humidity due to moisture content in the air, and decreased shortwave radiation from solar interception by clouds. Note that the sampling frequency here is 6 hourly, and values represent averages over that time step. The gray bands in Figure 3 locate chosen precipitation events and illustrate that the synthetic data appear visually consistent. That is, computed data show changes in precipitation-related variables that are similar to the changes in the original data, and these changes for both sets seem consistent with our meteorological expectations.

To quantify the consistency between meteorological variables for both expected and computed data we performed two statistical tests. These tests measure the consistency of precipitation occurrence with expected changes in long and shortwave radiation, and humidity. Precipitation events were selected that represented obvious precipitation events (the 10% of precipitation events with the most precipitation). These events were evaluated against proximal events when there was no precipitation. Statistical likelihoods that the values of related variables (long, shortwave, humidity) changed as expected with precipitation events were calculated, and differences between the expected and computed series were compared.

We calculated the statistical probability of occurrence, p, that a precipitation event (e) at time t will have more longwave radiation, for example, than a proximal event with no precipitation (n) within ±3 days of t. The probability (p = e/n) was 0.88 in the expected data and 0.85 in the computed. Thus both expected and computed datasets have greater than 80% probability that precipitation will be associated with enhanced longwave radiation. Specific humidity probability for expected was 0.92 and computed was 0.86. Shortwave probability was 0.50 for expected and 0.51 for computed (Note that this probability is smaller because nearly half of the precipitation events occurred at night when shortwave radiation was zero). Over all three related variables the computed ANN output probability differed only by 1-6% from the expected, which indicates that ANN output is as physically consistent as the expected data to within about 6%.

3.3. Application to CCSM3 Downscaling

Neural network models for temporal downscaling of the seven meteorological variables were developed and validated for $0.5^{\circ} \times 0.5^{\circ}$ grid cells using a CRU-NCEP dataset for contemporary time period. The trained models were then applied to downscale monthly mean output from the Community Climate System Model-3 [CCSM3, 17]. Figure 4 shows the time series of monthly mean temperatures from CCSM3 and downscaled 6 hourly time series. Our intention in downscaling CCSM3 output to 6 hourly time steps is to develop a dataset of meteorological forcing data for terrestrial ecosystem models (*cf.* paleonproject.org) that is 1) resolved to a time step useful to most models, yet still informed by the long term climate trends and variability, and 2) informed by weather variability from a higher resolution mesoscale reanalysis approach (NCEP). It is critical that the downscaled output closely correspond to the monthly





Figure 2: Comparison of downscaled results with original 6 hourly CRU-NCEP data. These show how well the neural network output computed from CRU-NCEP monthly means replicates the 6 hourly 'observed' values they were based on. The middle column of histograms of computed neural network variables show that the synthesized data have very similar frequency distributions to the original data in the left column. The XY scatter plots show the fit of the computed output to the expected values. The red, magenta, and blue bands locate the 33%, 66%, and 95% confidence intervals respectively. This indicates that for some variables, such as shortwave radiation, most of the data are tightly packed around the 1:1 line.



Figure 3: Case study comparison of consistency between meteorological variables for several precipitation events. The expected and computed both represent data extracted from the same grid cell (\sim 35°N,75°W) covering the same time (January 1–31 from an example year). These show that for both expected and synthetic ANN output related variables respond accordingly to precipitation events. Longwave radiation and specific humidity increase marginally and shortwave radiation decreases.

means from the GCM output, but also that the consistency in the original mesoscale training data be preserved, which are discussed above.

An important caveat for our downscaling approach is that we are making an assumption that weather variability over the contemporary period (1901-2010) is an acceptable proxy for weather variability over other periods. In this case we are applying our weather variability to CCSM3 data covering the historical period A.D. 1150–1700.

4. Software Development

The neural network based downscaling approach developed in this work has been implemented as a part of "Climate Observations and Model Data Analysis and Synthesis Toolkit (COMDAST)" package. The neural network tool has been developed in C and uses Message Passing Interface (MPI) for parallel execution of multiple ANN configurations on distributed memory parallel architectures. The package also provides tool sets for pre-processing, post-processing, statistical analysis and plotting. The software package is being distributed under open source GNU General Public License (GPL). http://www.climatemodeling.org/COMDAST



Figure 4: Downscaling of the validation (CRU-NCEP) dataset from monthly temperature to 6 hourly appears on the left in blue. The trained model is then applied to CCSM3 monthly temperature data on the right to obtain synthetic 6 hourly output.

5. Discussion and Future Work

Previous studies that implemented neural networks to downscale precipitation found that ANN's performed poorly due to inaccuracy in predicting precipitation occurrence [20] or had a tendency to underestimate precipitation extremes [21]. Haylock and others [21] however introduced a mixture model approach for estimating precipitation probability shape parameters allowing subsequent resampling to better model precipitation extremes. Neural networks have also been notably applied for use in measuring boundary layer dynamics [22, 23, 24]. Although similar these were intended for gap filling moisture, heat and CO_2 fluxes and interpolating over regions of sparse eddy covariance coverage.

In this paper we developed an open source framework for using neural networks to temporally downscale a variety of meteorological variables from monthly to 6 hourly, although it could be applied to bridge other time ranges as well. We found that by downscaling multiple correlated variables simultaneously with precipitation the mismatch error of the downscaled output was reduced (data not shown). If a variable showed weak correlations to other variables downscaled output was more likely to contain large errors from expected values. Precipitation, despite being a variable with strong influence on other variables such as solar radiation and humidity, was not predictably related to other variables or to any sequence of occurrence and required a second post-processing step to fit a generalized gamma distribution in order to achieve realistic precipitation intensity values [see 21]

The application of our neural network temporal downscaling approach differs from these previous approaches in that it allows us to downscale many variables across hundreds of years of monthly means to sub-daily values in just a few minutes per location or grid cell. We conducted varied validation tests between downscaled output and the original 6 hourly values. Our tests show that the means of downscaled output are accurate to within 0.6% across all variables (Table 3). Differences between expected and computed standard deviations were also similar to within 1% for longwave and shortwave radiation, specific humidity, and air temperature. Standard deviations for surface pressure, wind, and precipitation were notably different with differences of 4%, 10%, and 11% respectively, which should be considered if the intensity of extremes in these three variables is critical.

Table 2 shows that correlations between variables are highly consistent between expected and computed CRU-NCEP data for five variables, with surface pressure and precipitation being the exceptions. Variables exhibiting the weaker correlations to other variables such as shortwave radiation and wind (Table 2) tended to have larger bias values and error in standard deviation and skew (Table 3). This suggests that the correlation across the input variables handled by the neural network is important for training the neural network to constrain bias and variance in the synthetic output.

Finally, synthetic output from neural networks is based on the patterns observed across many instances (example years) in the training data. Neural networks downscaling may not be a realistic approach if downscaling data that extends to a period with significantly different weather variability. Also potential biases in the training data (*e.g.*, exaggerated diel cycles) may also be manifest in downscaled output.

Acknowledgment

This work was funded in part by the National Science Foundation, grant #1065848. For data contributions we thank Nicolas Viovy for developing and making publicly available the CRU-NCEP data, and also the National Center for Environmental Prediction (NCEP) and the Climate Research Unit, University of East Anglia who produced the source datasets for CRU-NCEP. We also thank Yafang Zhong for providing the CCSM-3 model output.

Oak Ridge National Laboratory is managed by UT-Battelle, LLC for the U.S. DOE under contract DE-AC05-00OR22725. A contractor of the U.S. Government has authored the submitted manuscript. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

References

- [1] T. G. F. Kittel, N. A. R. J. A. Royle, C. Daly, W. P. Gibson, H. H. Fisher, P. Thornton, D. N. Yates, S. Aulenbach, C. Kaufman, R. McKeown, D. Bachelet, D. S. Schimel, V. Participants, VEMAP phase 2 bioclimatic database. I. Gridded historical (20th century) climate for modeling ecosystem dynamics across the conterminous USA, Climate Research 27 (2) (2004) 151–170, doi: 10.3354/cr027151.
- [2] R. G. Jones, M. Noguer, D. Hassell, D. Hudson, S. Wilson, G. Jenkins, J. M. J, Generating high resolution climate change scenarios using PRECIS (2004).
- [3] C. W. Richardson, Stochastic simulation of daily precipitation, temperature, and solar radiation, Water Resources Research 17 (1) (1981) 182–190, doi: 10.1029/WR017i001p00182.
- [4] M. F. Hutchinson, Stochastic space-time weather models from ground-based data, Agricultural and Forest Meteorology 73 (3–4) (1995) 237–264, doi: 10.1016/0168-1923(94)05077-J.
- [5] R. W. Katz, M. B. Parlange, Mixtures of stochastic processes: application to statistical downscaling, Climatic Research 7 (2) (1996) 185–193, doi: 10.3354/cr007185.
- [6] F. Frey-Buness, D. Heimann, R. Sausen, A statistical-dynamical downscaling procedure for global climate simulations, Theoretical and Applied Climatology 50 (1995) 117–131, doi: 10.1007/BF00866111.
- [7] U. Fuentes, D. Heimann, An improved statistical-dynamical downscaling scheme and its application to the alpine precipitation climatology, Theoretical and Applied Climatology 65 (2000) 119–135, doi: 10.1007/s007040070038.
- U. Busch, D. Heimann, Statistical-dynamical extrapolation of a nested regional climate simulation, Climate Research 19 (2001) 1–13, doi: 10.3354/cr019001.
- M. A. Semenov, E. M. Barrow, Use of a stochastic weather generator in the development of climate change scenarios, Climatic Change 35 (4) (1997) 397–414, doi: 10.1023/A:1005342632279.
- [10] M. A. Semenov, Simulation of extreme weather events by a stochastic weather generator, Climatic Research 35 (3) (2008) 203–212, doi: 10.3354/cr00731.
- [11] A. R. Desai, D. J. Moore, W. Ahue, P. Wilkes, S. de Wekker B. G. Brooks, T. Campos, B. B. Stephens, R. K. Monson, S. P. Burns, T. Quaife, S. M. Aulenbach, D. S. Schimel, Seasonal pattern of regional carbon balance in the central Rocky Mountains from surface and airborne measurements, Journal of Geophysical Research 116 (G4), doi: 10.1029/2011JG001655.
- [12] H.-M. H. Juang, M. Kanamitsu, The NMC nested regional spectral model, Monthly Weather Review 122, doi: 10.1175/1520-0493(1994)122₁0003:TNNRSM₆2.0.CO;2.
- [13] H.-M. H. Juang, S.-Y. Hong, M. Kanamitsu, The ncep regional spectral model: An update, Bulletin of the American Meteorological Society 78 (1997) 2125–2143, doi: 10.1175/1520-0477(1997)078;2125:TNRSMA¿2.0.CO;2.
- [14] C. W. Richardson, D. A. Wright, WGEN: A model for generating daily weather variables (1984).
- [15] P. Coulibaly, Y. B. Dibike, F. Anctil, Downscaling precipitation and temperature with temporal neural networks, Journal of Hydrometeorology 6 (2005) 483–496, doi: 10.1175/JHM409.1.
- [16] D. E. Rumelhart, G. E. Hinton, R. J. Williams, Learning representations by back-propagating errors, Nature 323 (6088) (1986) 533–536. URL http://dx.doi.org/10.1038/323533a0
- [17] Y. Zhong, G. Miller, B. Otto-Bliesner, M. Holland, D. Bailey, D. Schneider, A. Geirsdottir, Centennial-scale climate change from decadallypaced explosive volcanism: a coupled sea ice-ocean mechanism, Climate Dynamics 37 (11) (2011) 2373–2387, doi: 10.1007/s00382-010-0967-z.
- [18] N. Viovy, P. Ciais, A combined dataset for ecosystem modelling, http://dods.extra.cea.fr/10 data/p529viov/cruncep/readme.htm (2009).
- [19] W. Kleiber, R. W. Katz, B. Rajagopalan, Daily spatiotemporal precipitation simulation using latent and transformed Gaussian processes, Water Resources Research 48 (1) (2012) W01523, doi: 10.1029/2011WR011105.
- [20] R. L. Wilby, T. M. L. Wigley, D. Conway, P. D. Jones, B. C. Hewitson, J. Main, D. S. Wilks, Statistical downscaling of general circulation model output: A comparison of methods, Water Resources Research 34 (11) (1998) 2995–3008, doi: 10.1029/98WR02577.
- [21] M. R. Haylock, G. C. Cawley, C. Harpham, R. L. Wilby, C. M. Goodess, Downscaling heavy precipitation over the United Kingdom: a comparison of dynamical and statistical methods and their future scenarios, International Journal of Climatology 26 (10) (2006) 1397–1415, doi: 10.1002/joc.1318.
- [22] D. Papale, R. Valentini, A new assessment of European forests carbon exchanges by eddy fluxes and artificial neural network spatialization, Global Change Biology 9 (4) (2003) 525–535, doi: 10.1046/j.1365-2486.2003.00609.x.
- [23] A. Schmidt, T. Wrzesinsky, O. Klemm, Gap filling and quality assessment of co₂ and water vapour fluxes above an urban area with radial basis function neural networks, Boundary-Layer Meteorology 126 (3) (2008) 389–413, doi: 10.1007/s10546-007-9249-7.
- [24] A. L. Neal, H. V. Gupta, S. A. Kurc, P. D. Brooks, Modeling moisture fluxes using artificial neural networks: can information extraction overcome data loss?, Hydrology and Earth System Sciences 15 (1) (2011) 359–368, doi: 10.5194/hess-15-359-2011.