

Physics-Constrained Deep Learning for Climate Downscaling

Paula Harder

*Fraunhofer ITWM, Kaiserslautern, Germany
Mila Quebec AI Institute, Montreal, Canada
TU Kaiserslautern, Kaiserslautern, Germany*

PAULA.HARDER@FRAUNHOFER.ITWM.DE

Venkatesh Ramesh

*Mila Quebec AI Institute, Montreal, Canada
University of Montreal, Montreal, Canada*

Alex Hernandez-Garcia

*Mila Quebec AI Institute, Montreal, Canada
University of Montreal, Montreal, Canada*

Qidong Yang

*Mila Quebec AI Institute, Montreal, Canada
New York University, New York, USA*

Prasanna Sattigeri

IBM Research, New York, USA

Daniela Szwarcman

IBM Research, Brazil

Campbell D. Watson

IBM Research, New York, USA

David Rolnick

*Mila Quebec AI Institute, Montreal, Canada
McGill University, Montreal, Canada*

Abstract

The availability of reliable, high-resolution climate and weather data is important to inform long-term decisions on climate adaptation and mitigation and to guide rapid responses to extreme events. Forecasting models are limited by computational costs and, therefore, often generate coarse-resolution predictions. Statistical downscaling, including super-resolution methods from deep learning, can provide an efficient method of upsampling low-resolution data. However, despite achieving visually compelling results in some cases, such models frequently violate conservation laws when predicting physical variables. In order to conserve physical quantities, we develop methods that guarantee physical constraints are satisfied by a deep learning downscaling model while also improving their performance according to traditional metrics. We compare different constraining approaches and demonstrate their applicability across different neural architectures as well as a variety of climate and weather datasets. Besides enabling faster and more accurate climate predictions, we also show that our novel methodologies can improve super-resolution for satellite data and standard datasets.

1. Introduction

Accurate modeling of weather and climate is critical for taking effective action to combat climate change. In addition to shaping global understanding of climate change, local and regional predictions guide adaptation decisions and provide impetus for action to reduce greenhouse gas emissions (Gutowski et al., 2020). Predicted and observed quantities such as precipitation, wind speed, and temperature impact decisions in sectors such as agriculture, energy, and transportation. While these quantities are often required at a fine geographical and temporal scale to ensure informed decision-making, most climate and weather models are extremely computationally expensive to run (sometimes taking months even on super-computers), resulting in coarse-resolution predictions. Thus, there is a need for fast methods that can generate high-resolution data based on the low-resolution models that are commonly available.

The terms *downscaling* in climate science and *super-resolution* (SR) in machine learning (ML) refer to a map from low-resolution (LR) input data to high-resolution (HR) versions of that same data; the high-resolution output is referred to as the super-resolved (SR) data. Downscaling via established statistical methods—*statistical downscaling*—has been long used by the climate science community to increase the resolution of climate data (Maraun and Widmann, 2018). In parallel, computer vision SR has evolved rapidly using various deep learning architectures, with such methods now including super-resolution convolutional neural networks (CNNs) (Dong et al., 2016), generative adversarial models (GANs) (Wang et al., 2018a), vision transformers (Yang et al., 2020), and normalizing flows (Lugmayr et al., 2020). Increasing the temporal resolution via frame interpolation is also an active area of research for video enhancement (Liu et al., 2017) that can be transferred to spatiotemporal climate data. Recently, deep learning approaches have been applied to a variety of climate and weather datasets, covering both model output data and observations. Climate super-resolution has mostly focused on CNNs (Vandal et al., 2017), recently shifting towards GANs (Stengel et al., 2020; Wang et al., 2021).

Generating high-resolution data with machine learning can produce realistic-looking images and good predictive accuracy. However, a major obstacle often encountered when applying ML to a physical system such as the Earth’s atmosphere is that the predicted output values can violate physical laws such as conservation of energy, momentum, and mass. Even slight violations of constraints may be problematic - leading to errors that potentially compound as climate models are run iteratively on their own output. Moreover, the inability to respect physical laws can lead to a lack of trust and lower adoption among domain scientists. Besides climate science, there are numerous domains of ML for societal benefit in which satisfaction of physical constraints is fundamentally important. Examples include the discovery of new materials for energy and healthcare, aerodynamics simulations for efficient vehicles, and optimal control in industrial settings. There are certain tasks that are more suited for hard-constraining than others. One important point is that there exists a relationship between low-resolution and high-resolution samples given by an equation. This can be the case when modeling physical quantities, with for example mass or energy conservation that exists between LR and HR pairs. If we are looking at compressed or blurry images and trying to remove the effects of compression or blur, there may be no known constraint between low and high resolution, and so constraining methodologies would not be

applicable. On the other hand, for some data from e.g. satellites or telescopes, images are created by summing photons across a given field of view, so the value at a given pixel can be interpreted as the sum of values at unobserved subpixels; in such cases, hard constraints could potentially be useful.

In this work, we introduce novel methods to strictly enforce physical constraints between low-resolution (input) and high-resolution (output) images. We do this via a constraint layer at the end of a neural architecture, which renormalizes the prediction either additively, multiplicatively, or with an adaptation of the softmax layer. We use datasets based on ERA5, WRF, and NorESM data, spanning different quantities such as water content, temperature, water vapor, and liquid water content. For ERA5 data we also look increasing the resolution by different factors, we create datasets with a enhancement of factors ranging from 2 over 4 and 8 to 16. We show the utility of our methods across architectures including CNNs, GANs, CNN-RNNs, and a novel architecture that we introduce to apply super-resolution in both spatial and temporal dimensions. Besides climate datasets we show that our methods are able to improve predictive accuracies for lunar satellite imagery super-resolution as well as on standard image super-resolution benchmark datasets, like Set5, Set14, Urban100 and BSD100. Our code is available at <https://github.com/RoInickLab/constrained-downscaling> and our main dataset can be found at <https://drive.google.com/file/d/1IENhP1-aTYyqOkRcnmCIvxXkvUW2Qbdx/view>.

Contributions Our main contributions can be summarized as follows:

- We introduce a novel constraining methodology for deep learning-based downscaling methods, which guarantees that physical constraints such as mass and energy conservation are satisfied in the prediction.
- We show that our method improves predictive performance across different deep learning architectures on a variety of climate datasets.
- Additionally, we show that our method increases the accuracy of super-resolution in other domains, such as standard images and satellite imagery, and we introduce a new deep learning architecture for downscaling along both spatial and temporal dimensions.

2. Related Work

Deep Learning for Climate Downscaling There exists extensive work on ML methods for climate and weather observations and predictions, from CNN architectures (Vandal et al., 2017) to GANs (Stengel et al., 2020) and normalizing flows (Groenke et al., 2020). Recently, especially GANs have become a very popular architecture choice, including many works on precipitation model downscaling (Wang et al., 2021; Watson et al., 2020; Chaudhuri and Robertson, 2020) as well as other quantities such as wind and solar data (Stengel et al., 2020). Unified frameworks comparing methods and benchmarks were introduced by Baño Medina et al. (2020) to assess different SR-CNN setups and by Kurinchi-Vendhan et al. (2021) with the introduction of a new dataset for wind and solar SR. To date, there has been limited work on spatiotemporal SR with climate data. Some authors have looked at super-resolving multiple time steps at once, but not increasing the temporal resolution (Harilal et al., 2021;

Leinonen et al., 2021), whereas Serifi et al. (2021) increases the temporal resolution by just treating the time steps as different channels and using a standard SR-CNN.

Constrained Learning for Climate Various works on ML for climate science have attempted to enforce certain physical constraints via soft penalties in the loss (Beucler et al., 2019), linearly constrained neural networks for convection (Beucler et al., 2021), or aerosol microphysics emulation (Harder et al., 2022). A different line of work incorporates constraints into machine learning based on flux balances (Sturm and Wexler, 2020, 2022; Yuval et al., 2021). These strategies use domain knowledge of how properties flow to ensure conservation of different quantities, instead of predicting tendencies directly, fluxes are predicted. Outside climate science, recent works have emerged on enforcing hard constraints on the output of neural networks (e.g. Donti et al. (2021)).

Constrained Learning for Downscaling In super-resolution for turbulent flows, Mesh-freeFlowNet (Jiang et al., 2020) employs a physics-informed model which adds PDEs as regularization terms to the loss function. In parallel to our work, the first approaches employing hard constraints for climate-related downscaling were introduced: Geiss and Hardin (2020) introduces an enforcement operator applied to multiple CNN architectures for scientific datasets. A CNN with a multiplicative renormalization layer is used for atmospheric chemistry model downscaling in (Geiss et al., 2022). Hess et al. (2022) introduces one global constraint to be applied to post-process the precipitation prediction generated by a GAN. We are the first to compare a variety different hard-constraining approaches and also apply them to multiple deep learning architectures.

3. Enforcing Constraints

When modeling physical quantities such as precipitation or water mass, principled relationships such as mass conservation can naturally be established between low-resolution and high-resolution samples. Here, we introduce a new methodology to incorporate these constraints within a neural network architecture. We choose hard constraints enforced through the architecture over soft constraints that use an additional loss term. Hard constraints guarantee certain constraints even at inference time, whereas soft constraining encourages the network to output values that are close to satisfying constraints, while minimizing a penalty during training, but do not provide any guarantees. Additionally, for our case hard constraining increases the predictive ability, and soft constraining can lead to unstable training and an accuracy-constraints trade-off (Harder et al., 2022).

3.1 Setup

Consider the case of downscaling by a factor of N in each linear dimension, and let $n := N^2$. Let $y_i, i = 1, \dots, n$ be the values in the predicted high-resolution patch that correspond to low-resolution pixel x . The set $\{y_i\}$ for $i = 1, \dots, n$ is also referred to as a super-pixel. Then, a conservation law takes the form of the following constraint:

$$\frac{1}{n} \sum_{i=1}^n y_i = x. \quad (1)$$

Depending on the predicted quantity, there may additionally be an inequality constraint associated with the data. One example is when we consider physical variables that can only assume positive values, such as water mass. This would give us:

$$y_i \geq 0 \quad \forall i = 1, \dots, n \quad (2)$$

3.2 Constraint layer

We introduce three different constraint layers: additive constraining, multiplicative constraining, and softmax-based constraining. These are all added at the end of any neural architecture, as shown in Figure 2, and all satisfy (1) by construction. The constraints are applied for each pair of input pixel x and the corresponding SR $N \times N$ patch. An illustration is shown in Figure 1. We will use $\tilde{y}_i, i = 1, \dots, n$ to denote the intermediate outputs of the neural network before the constraint layer and $y_i, i = 1, \dots, n$ to be the final outputs after applying the constraints.

Additive constraining For our Additive Constraint Layer (AddCL), we take the intermediate outputs and reset them using the following operation:

$$y_j = \tilde{y}_j + x - \frac{1}{n} \sum_{i=1}^n \tilde{y}_i. \quad (3)$$

A similar approach is used in Geiss et al. (2022). Note that this approach in general violates nonnegativity constraints, so is sometimes inapplicable.

We also consider a more complex additive approach, the Scaled Additive Constraint Layer (ScAddCL), which was first introduced by Geiss and Hardin (2020):

$$y_j = \tilde{y}_j + \left(x - \frac{1}{n} \sum_{i=1}^n \tilde{y}_i\right) \cdot \frac{\sigma + \tilde{y}_j}{\sigma + \frac{1}{n} \sum_{i=1}^n \tilde{y}_i}, \quad (4)$$

with $\sigma := \text{sign}\left(\frac{1}{n} \sum_{i=1}^n \tilde{y}_i - x\right)$.

Multiplicative constraining For the Multiplicative Constraint Layer (MultCL) approach, we rescale the intermediate output using the corresponding input value x :

$$y_j = \tilde{y}_j \cdot \frac{x}{\frac{1}{n} \sum_{i=1}^n \tilde{y}_i}. \quad (5)$$

Softmax constraining For predicting quantities like atmospheric water content, we want to enforce the output to be nonnegative and, therefore, physically valid. Here, we use a softmax multiplied by the corresponding input pixel value x :

$$y_j = \exp(\tilde{y}_j) \cdot \frac{x}{\sum_{i=1}^n \exp(\tilde{y}_i)}. \quad (6)$$

This Softmax Constraint Layer (SmCL) enforces $y_i \geq 0, i = 1, \dots, n$.

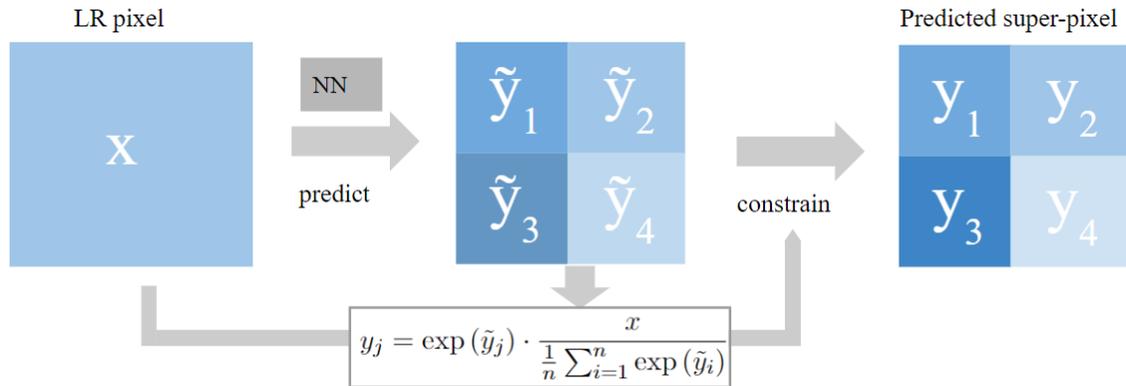


Figure 1: Our Softmax Constraining Layer (SmCL) is shown for one input pixel x and the corresponding predicted 2×2 super-pixel for the case of 2 times upsampling.

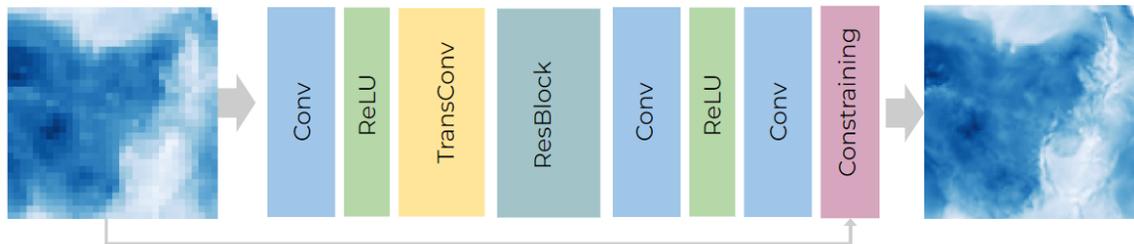


Figure 2: The CNN architecture for 2 times upsampling including the constraint layer (in red).

4. Data

To test and evaluate our proposed method, we create a variety of datasets as well as use existing and established ones. We generate multiple datasets based on the ERA5 reanalysis product using average pooling to create the LR inputs, which has been the standard methodology in climate downscaling studies (see e.g. Serifi et al. (2021); Leinonen et al. (2021)). We also use datasets based on the outputs of models such as the Weather and Research Forecasting (WRF) Model and the Norwegian Earth System Model (NorESM) that contain real low-resolution simulation data matched to high-resolution data. Finally, we test our methods on non-climate datasets: lunar satellite imagery and natural images. An overview of all the different datasets used can be found in Table 1.

4.1 ERA5 dataset

The ERA5 dataset (Hersbach et al., 2020) is a so-called *reanalysis* product from the European Center for Medium-Range Weather Forecast (ECMWF) that combines model data with worldwide observations. The observations are used as boundary conditions for numerical

Table 1: The different datasets we use to test our constraint layers.

SOURCE	QUANTITY	DIMENSIONS	SIZE
		LR/HR	TRAIN/VAL/TEST
ERA5	WATER CONTENT	(1,32,32)/(1,128,128)	40k/10k/10k
ERA5	WATER CONTENT	(1,32,32)/(1,128,128)	40k/10k/10k
ERA5	WATER CONTENT	(1,32,32)/(1,128,128)	40k/10k/10k
ERA5	WATER CONTENT	(1,32,32)/(1,128,128)	40k/10k/10k
ERA5	WATER CONTENT	(3,32,32)/(3,128,128)	40k/10k/10k
ERA5	WATER CONTENT	(2,32,32)/(3,128,128)	40k/10k/10k
ERA5	WATER VAPOR	(3,32,32)/(3,128,128)	40k/10k/10k
	LIQUID WATER		
	TEMPERATURE		
WRF	TEMPERATURE	(1,45,45)/(1,135,135)	20k/4k/4k
NORESM	TEMPERATURE	(1,32,32)/(1,64,64)	24k/12k/12k
LUNAR IMAGERY	PHOTON COUNT	(1,32,32)/(1,128,128)	132k/16k/16k
NATURAL IMAGES	RGB	(1,128,128)/(1,512,512)	4300/100

models that then predict various atmospheric variables. ERA5 is available as global, hourly data with a $0.25^\circ \times 0.25^\circ$ resolution, which is roughly 25 km per pixel in the mid-latitudes. It covers all years starting from 1950.

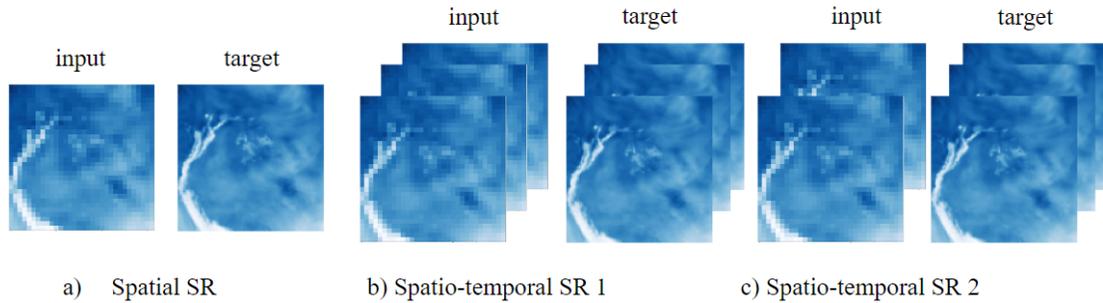


Figure 3: Samples of the three different dataset types used in this work. a) A data pair we use for our standard spatial super-resolution task. The input is an LR image and the target is the HR version of that. b) A data pair for performing SR for multiple time steps simultaneously. The input is a time series of LR images and the output is the same time series in HR. c) A data pair where SR is performed both temporally and spatially, with two LR time steps as input and 3 HR time steps as a target.

Total water content dataset For this work, one quantity we focus on is the total column water (tcw) that is given in kg/m^2 and describes the vertical integral of the total amount of atmospheric water content, including water vapour, cloud water, and cloud ice but not precipitation.

Spatial SR data To obtain our high-resolution data points we extract a random 128×128 pixel image from each available time step (each time step is 721×1440 and there are roughly 60,000 time steps available). We randomly sample 40,000 data points for training and 10,000 for each validation and testing. The low-resolution counterparts are created by taking the mean over $N \times N$ patches, where N is our upsampling factor. A sample pair is shown in Figure 3 a). This operation is physically sound, considering that conservation of water content means that the water content (density per squared meter) described in an LR pixel should be equal to the average of the corresponding HR pixels.

Spatio-Temporal datasets Including the temporal evolution of our data, we create two additional datasets. For the first dataset, one sample consists of 3 successive time steps, the same time steps for both input and target, but at different resolutions. This is done to perform spatial SR for multiple time steps simultaneously, see Figure 3 b). We select three random 128×128 pixel areas per global image, resulting in the same number of examples as the procedure described above. We split the data randomly as before, and each time step is downsampled by taking the spatial mean. To increase both spatial and temporal dimensions, we again crop three images out of a series of three successive time steps. To create the low-resolution input, we take every other time step and compute the mean spatially, resulting in two LR inputs, see Figure 3 c).

OOD dataset For the datasets described above the train-val-test split is done randomly. To understand how our constraining influences out-of-distribution generalization, we create a dataset with a split in time. We train on older data and then test on more recent years: for training we use the years 1950-2000, for validation 2001-2010, and for final testing 2011-2020.

Energy dataset Also originating from the ERA5 data, we create a second dataset including different physical variables coming with different constraints as well. This dataset is constructed to preserve moist static energy and water masses while predicting water vapor, liquid water content, and air temperature. The variables are taken from the pressure level at 850hPa.

4.2 WRF data

In Watson et al. (2020), a dataset using the Advanced Research version of the Weather Research and Forecasting (WRF) Model is introduced. It comprises hourly operational weather forecast data for Lake George in New York, USA from 2017-01-01 to 2020-03-20. More details about the model and its configuration can be found in Watson et al. (2020). The variable we consider for this work is the temperature at 2m above the ground. Unlike the previous datasets, this one does not involve synthetic downsampling but includes two forecasts run at different resolutions with different physics-based parameterizations: one at 9 km horizontal resolution and one at 3 km. Our goal is to predict the 3 km resolution temperature field given the 9 km one and builds on work by Auger et al. (2021), which used the same dataset.

4.3 Lunar data

Recent work (Delgado-Centeno et al., 2021) on super-resolution for lunar satellite imagery has shown how deep learning can be used to enhance the captured data to help future missions

to the moon. To increase the resolution of images from regions like the south pole, where there is no high-resolution data available, a machine learning-ready dataset has been created. It consists of 220,000 images cropped out of the Narrow-Angle Camera (NAC) imagery from NASA’s Lunar Reconnaissance Orbiter (LRO); for more details see Delgado-Centeno et al. (2021). Here we use a 4x upsampling version of the dataset to verify if our constraining methodologies can increase the performance of super-resolution outside of climate science.

4.4 Natural images

The standard benchmark datasets for super-resolution deep learning architectures applied to natural images include the OutdoorSceneTRaining (OST), DIV2K, and Flickr2k datasets for training and Set5, Set14, Urban100 and BSD100 for testing, as for example in Wang et al. (2018b). Here, we use a version resized to 512×512 pixels for HR and apply average pooling to downsample them.

4.5 Constraints in our datasets

In predicting distinct physical quantities, there are different constraints we need to consider. Most of our datasets include the downscaling constraints given by (1), which are satisfied by the LR-HR pairs either approximately (for simulations that are run at LR and HR with quantities respecting physical conservation laws) or exactly (in the case of average pooling for creating the LR version). We detail the constraints in the following subsections.

Water mass conservation For predicting the total water mass, we are given the low-resolution water mass $Q^{(LR)}$ and must obtain the super-resolved version $Q^{(SR)}$. The downscaling constraint or mass conservation constraint (1) for each LR pixel $q^{(LR)}$ and the corresponding super-pixel $(q_i^{(SR)})_{i=1,\dots,n}$ is then given by

$$\frac{1}{n} \sum_{i=1}^n q_i^{(SR)} = q^{(LR)}. \quad (7)$$

Moist static energy conservation One of our tasks includes predicting water vapor, liquid water, and temperature while conserving both water mass and moist static energy. As described above, water mass conservation is straightforward, directly applying our constraining methodology. On the other hand, the moist static energy S is given by:

$$S = ((1 - Q_v) \cdot c_{pd} + Q_L \cdot c_l) \cdot T + L_v \cdot Q_v, \quad (8)$$

where

$$L_v = 2.5008 \cdot 10^6 + (c_{pw} - c_L) \cdot (T - 273.16)$$

is the latent heat of vaporization in (Jkg^{-1}) . The water vapor $Q_v[kg \cdot kg^{-1}]$, the liquid water $Q_L[kg \cdot kg^{-1}]$, and the temperature $T[K]$ are being predicted, whereas c_{pd} , c_{pw} and $c_L[J \cdot K^{-1} \cdot kg^{-1}]$ are heat capacity constants.

We use the following procedure to predict these quantities while conserving moist static energy:

1. Given LR T, Q_V, Q_L

2. Calculate LR S with (8)
3. Predict SR S, Q_v, Q_L while enforcing (1) using one of our constraint layers
4. Calculate SR T using (8) and SR S, Q_v, Q_L .

This means we predict T not directly, but by predicting S . We are then able to predict the temperature T while ensuring energy conservation by applying our constraint layer to the prediction of S .

Different simulations If the LR-HR pairs are not created by taking the local mean of the HR but by using two simulations run at different resolutions, the downscaling constraint is not automatically satisfied in the data. This is the case for our WRF and NorESM datasets (NorESM data is discussed in the appendix; here, we focus on WRF). Even though the downscaling constraint is not exactly obeyed (see Figure 4), it is approximately, and we can still apply our constraining in the same way as before. If the real low-resolution data and the downsampled high-resolution data are not significantly dissimilar, constraining can still benefit the predictive ability.

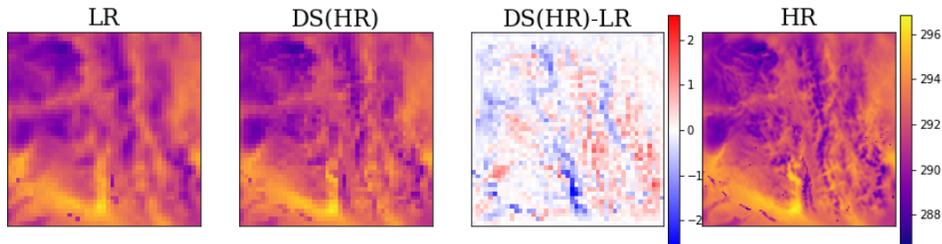


Figure 4: A LR-HR pair from the WRF temperature data. Here we compare the real LR with the low-resolution data created by average pooling of the HR, written as DS(HR).

5. Experimental Setup

We conduct three sets of experiments:

1. Compare different constraining methods and different upsampling factors on the ERA5 water content data.
2. Show the applicability of our constraining method to other neural network architectures.
3. Show the applicability of our constraining method to different datasets and different constraints types.

In most of our experiments, we use synthetic low-resolution data created by applying average pooling to the original high-res samples, as usually done in this research area.

Additionally, we consider cases with pairs of real low-res and high-res simulations to show that our methods work in the intended final application.

5.1 Architectures

We test our constraints methods throughout a variety of standard deep learning SR architectures including an SR CNN, conditional GAN, a combination of an RNN and CNN for spatio-temporal SR, and a new architecture combining optical flow with CNNs/RNNs to increase the resolution of the temporal dimension. The original, unconstrained versions of these architectures then also serves as a comparison for our constraining methodologies.

SR-CNNs Our SR CNN network, similar to Lim et al. (2017), consists of convolutional layers using 3×3 kernels and ReLU activations. The upsampling is performed by a transpose convolution followed by residual blocks (convolution, ReLU, convolution, adding the input, ReLU). The architecture for 2 times downscaling is shown in Figure 2.

SR-GAN A conditional GAN architecture (Mirza and Osindero, 2014) is a common choice for super-resolution (Ledig et al., 2016). Our version uses the above-introduced CNN architecture as the generator network. The discriminator is used from (Ledig et al., 2016), it consists of convolutional layers with a stride of 2 to decrease the dimensionality in each step, with ReLU activation. It is trained as a classifier to distinguish SR images from real HR images using a binary cross-entropy loss. The generator takes as input both Gaussian noise as well as the LR data and then generates an SR output. It is trained with a combination of an MSE loss and the adversarial loss given by the discriminator, like a standard SR GAN, e.g. Ledig et al. (2017).

SR-ConvGRU We apply an SR architecture based on the GAN presented by Leinonen et al. (2021), which uses ConvGRU layers to address the spatio-temporal nature of super-resolving a time series of climate data. Here we only use the generator trained without the discriminator, providing a deterministic approach.

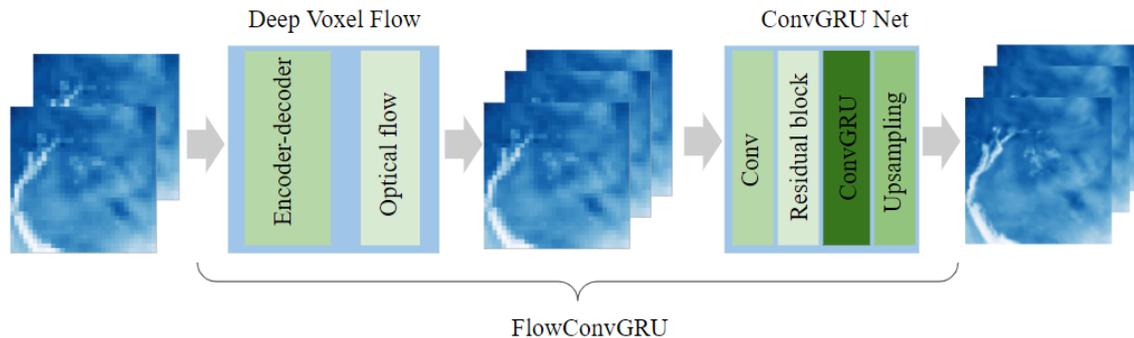


Figure 5: Our novel spatio-temporal architecture, combining Deep Voxel Flow and a ConvGRU.

SR-FlowConvGRU To increase the temporal resolution of our data we employ the Deep Flow Flow method (Liu et al., 2017), a deep learning architecture for video frame interpolation combining optical flow methods with neural networks. We introduce a new architecture combining the Deep Flow model and the ConvGRU network (FlowConvGRU): First, we increase the temporal resolution resulting in a higher-frequency time-series of LR images on which we then apply the ConvGRU architecture to increase the spatial resolution. The combined neural networks are then trained end-to-end. The architecture is shown in Figure 5.

5.2 Training

Our models were trained with the Adam optimizer, a learning rate of 0.001, and a batch size of 256. We trained for 200 epochs, which took about 3–6 hours on a single NVIDIA A100 Tensor Core GPU, depending on the architecture. All models use the MSE as their criterion, the GAN additionally uses its discriminator loss term. All the data is normalized between 0 and 1 for training, except for the cases where the ScAddCL is applied. In the case of this constraint layer we scale the data between -1 and 1 as proposed in Geiss and Hardin (2020), only for our time-dependent model, ConvGRU and FlowConvGRU, we use scaling between 0 and 1, because -1 and 1 let to NaN-values during training.

5.3 Baselines

Pixel enlargement This baseline consists of scaling the LR input to the same size as the HR by duplicating the pixels. We include this to have reference metrics that reflect how close the LR is to the HR data. This baseline conserves mass by construction.

Bicubic upsampling As a simple non-ML baseline, we use bicubic interpolation for spatial SR and take the mean of two frames for temporal SR.

Soft constraining Soft-constraining has been successfully applied before to a variety of physics-informed deep-learning tasks. Here we use it to see how it compares to hard constraints. Soft-constraining is done by adding a regularization term to the loss function. Our MSE loss is then changed to the following:

$$\text{Loss} = (1 - \alpha) \cdot \text{MSE} + \alpha \cdot \text{Constraint violation}, \quad (9)$$

where the constraint violation is the mean overall constraint violations between an input pixel x and the corresponding super-pixel $y_i, i = 1, \dots, n$:

$$\text{Constraint violation} = \text{MSE} \left(\frac{1}{n} \sum_{i=1}^n y_i, x \right).$$

We conducted an experiment to investigate the impact of α values on final model performance; the results are reported in the appendix. For our main paper we choose $\alpha = 0.99$.

Unconstrained counterparts Furthermore, we always compare against an unconstrained version of the above-introduced standard SR NN architectures (SR-CNN, SR-GAN, SR-ConvGRU, SR-FlowConvGRU).

Clipping We also run the standard CNN, but with clipping applied at inference. This is a common practice to remove negative values. Results can be found in the appendix, see Table 6. This method does not guarantee mass conservation or significantly improves performance.

6. Results and discussion

For evaluating our results, we use typical metrics for weather and climate super-resolution: root-mean-square error (RMSE), mean absolute error (MAE) and mean bias as well as typical metrics for super-resolution: peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), multi-scale SSIM (MS-SSIM) and Pearson correlation. We show RMSE and MS-SSIM in the main paper, while the others can be found in the appendix. Most metrics are highly correlated in our case. For the GAN giving a probabilistic prediction, we also use continuous ranked probability score (CRPS). Because we are interested in the violation of conservation laws and predicting non-physical values, we also look at the average constraint violation, the number of (unwanted) negative pixels, and the average magnitude of negative values.

6.1 Different constraining methods

Whereas hard-constraining shows exact conservation and appears to enhance performance, the application of soft-constraining on the other hand does decrease constraint violation, but still maintains a significant magnitude of it, which can be seen in Figure 6 for example. Also, soft-constraining seems to suffer from an accuracy-constraints trade-off, where depending on the regularization factor α , either the constraint violation is reduced, or the accuracy increases, but it struggles to do both simultaneously. A table for different α is shown in the appendix. Among the hard-constraining methodologies, the multiplicative renormalization layer, MultCL, performs the weakest in terms of predictive skills (e.g. see Figures 6, 7, 8 and 9). The three other methods, ScAddCL, AddCL, and SmCL, often have very similar measurements. SmCL shows the advantage of also enforcing positivity when necessary (see Figure 6).

6.2 Different architectures

As shown in Figures 6, 7, 8, and 9 for all architectures (CNN, GAN, ConvGRU, FlowConvGRU), adding the constraint layers enforces the constraint and improves the evaluation metrics compared to the CNN case. Constraining the GAN leads to less of a performance boost, but AddCL and SmCL still enhance the predictions compared to the unconstrained GAN. Including the temporal dimensions, the constraining improves the prediction quality much more significantly than in the case with just a single time step (see Figures 8 and 9).

6.3 Different datasets and constraints

The success of our constraining methodology does not depend on the upsampling factor: in Figure 6, we can see that the constraining methods work well and improve all metrics for upsampling factors of 2, 4, 8, and 16. When applied to our out-of-distribution dataset, the improvement achieved by adding constraints is even more pronounced than for the

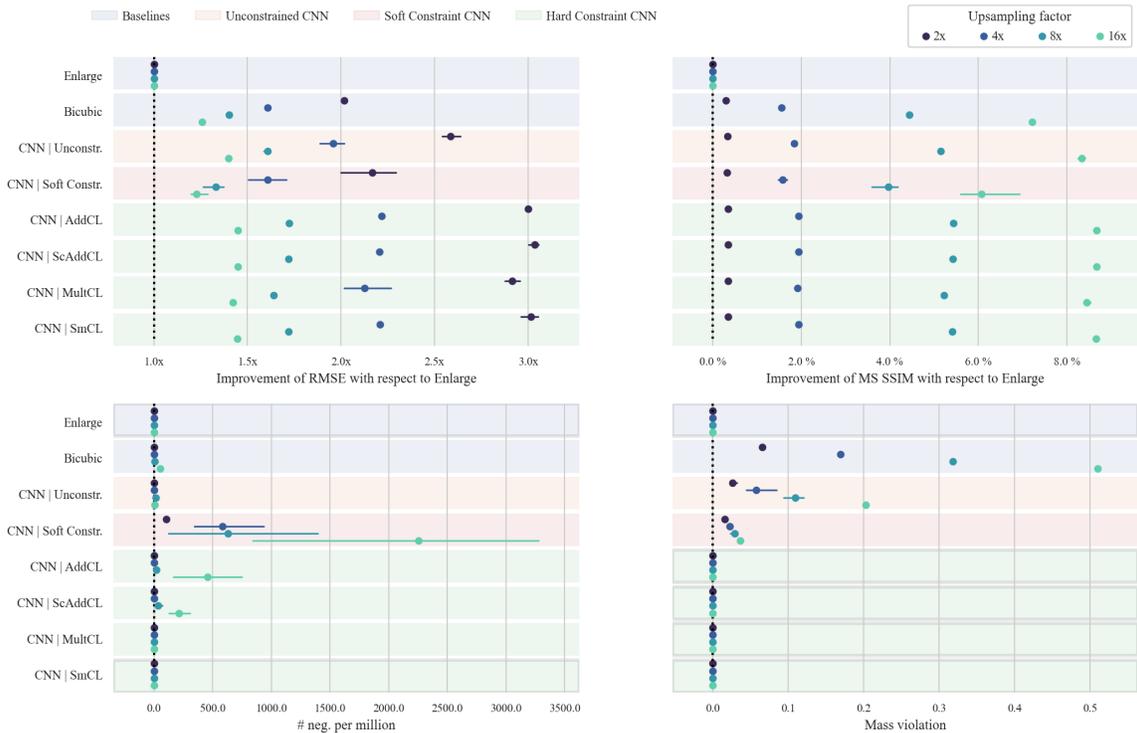


Figure 6: Metrics for different constraining methods applied to an SR CNN, calculated over 10,000 test samples of the water content data. The mean and confidence interval from 3 runs is shown relatively to the Enlarge baseline. The framed box indicates a method that achieves zero violation of the physics, no negative pixels or mass conservation up to numerical precision. A table with more metrics can be found in the appendix

randomly split data (see results in Figure 10). The constraints can help architectures with their generalization ability.

Not only mass can be conserved, but other quantities such as moist static energy. We show that moving on to different quantities of the ERA5 dataset, temperature, water vapor, and liquid water. Looking at Table 12 (see appendix), one can observe similar results for liquid water Q_L and water vapor Q_v as for the total water content: ScAddCL, AddCL, and SmCL significantly improve results in all measures over the unconstrained CNN, while enforcing energy and mass conservation. For temperature, on the other hand, MultCL performs the strongest, followed by SmCL, whereas AddCL and ScAddCL achieve smaller improvements in the scores.

Our WRF temperature dataset includes low-resolution data points drawn from a separate simulation, rather than downsampling, and therefore it results in much harder tasks. Table 2 shows that the scores are improved slightly with our constraint layer.

Finally, we also show that applying our constraint methodology can improve results in other domains, even in cases where there is no physics involved. We see that both for the

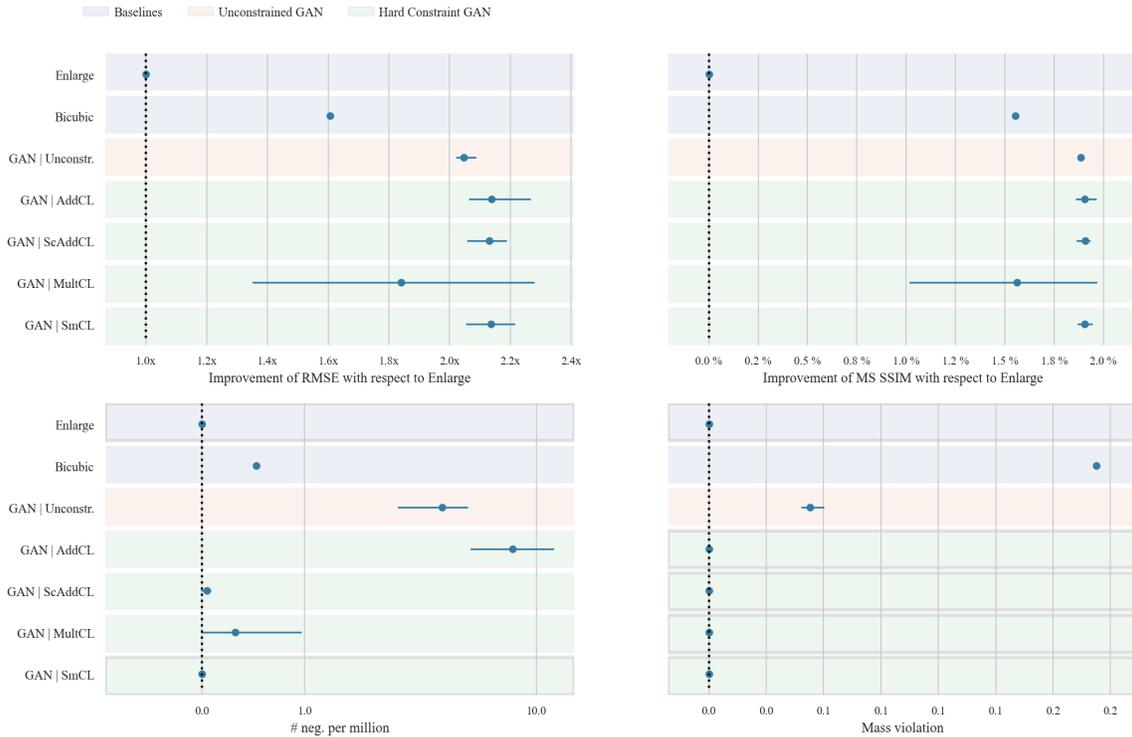


Figure 7: Metrics for different constraining methods applied to an SR GAN, calculated over 10,000 test samples of the 4x upsampling water content data. The mean and confidence interval from 3 runs are shown, for RMSE and MS-SSIM relatively to the Enlarge baseline for number of negative pixels (per mil.) and mass conservation violation the absolute values are shown. The framed box indicates a method that achieves zero violation of the physics, no negative pixels or mass conservation up to numerical precision. A table with more metrics can be found in the appendix, see Table 8.

lunar satellite imagery and the natural images benchmark datasets, the application of our SmCL improves the traditional metrics, as shown in Tables 3 and 4.

6.4 Perceptual quality of predictions

Additionally to an enhancement quantitatively, we can see an improved visual quality for some examples, as shown in Figure 13 and 14 for the water content data. For the WRF temperature forecast data, we see a very significant improvement in the perceptual quality of the prediction. Looking at an example, such as shown in Figure 11, we can see how much more detail is added to the prediction when adding our constraining. For the lunar satellite imagery, Figure 12 shows that applying constraints can make the image slightly less blurry.

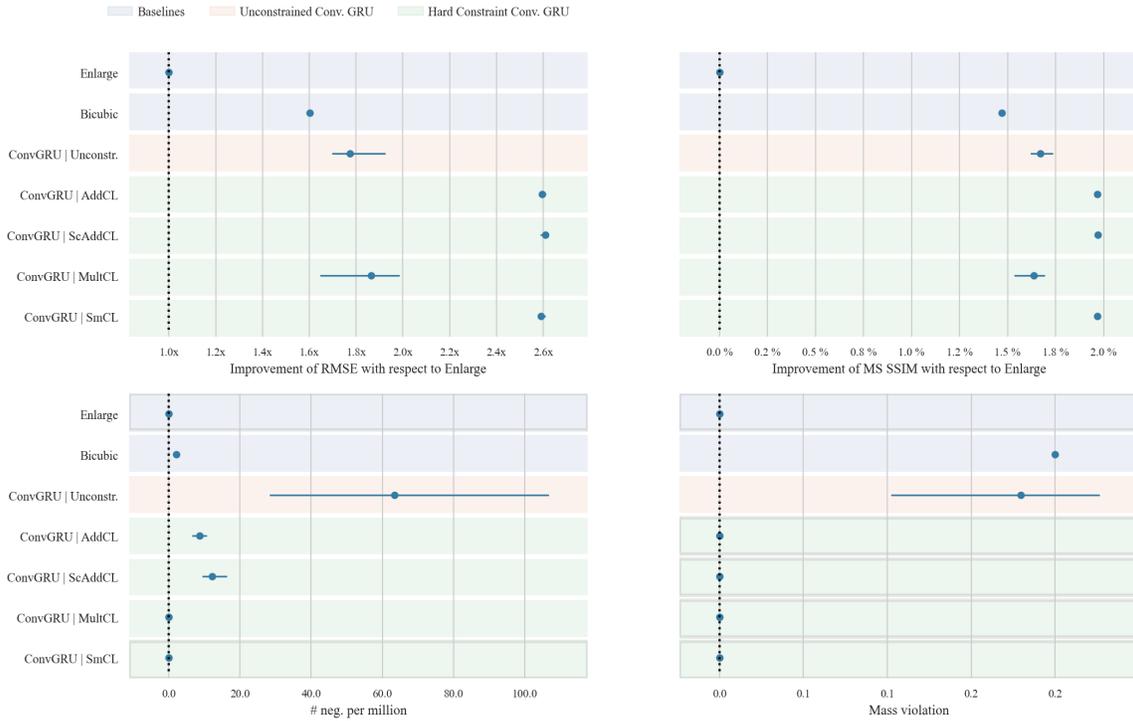


Figure 8: Metrics for different constraining methods applied to an SR ConvGRU, calculated over 10,000 test samples of the 4x upsampling water content data. The mean and confidence interval from 3 runs are shown, for RMSE and MS-SSIM relatively to the Enlarge baseline for number of negative pixels (per mil.) and mass conservation violation the absolute values are shown. The framed box indicates a method that achieves zero violation of the physics, no negative pixels or mass conservation up to numerical precision. A table with more metrics can be found in the appendix, see Table 9.

6.5 Development of error during training

Observing how the MSE develops during training (see Figure 15), we can see that the curve of the constrained network is generally lower than the unconstrained one. Additionally, it can be seen that constraining helps smooth both the training and validation curves.

6.6 Spatial distribution of errors

A known issue in downscaling methods is the so-called coastal effect, where errors of predictions tend to be more pronounced in coastal regions. Besides coastal region areas, mountain ridges can also be critical. In Figure 16, we show the error of the unconstrained prediction for water content and the softmax-constrained prediction. We can see that both predictions show more errors in coastal and mountainous regions. However, if we analyze the

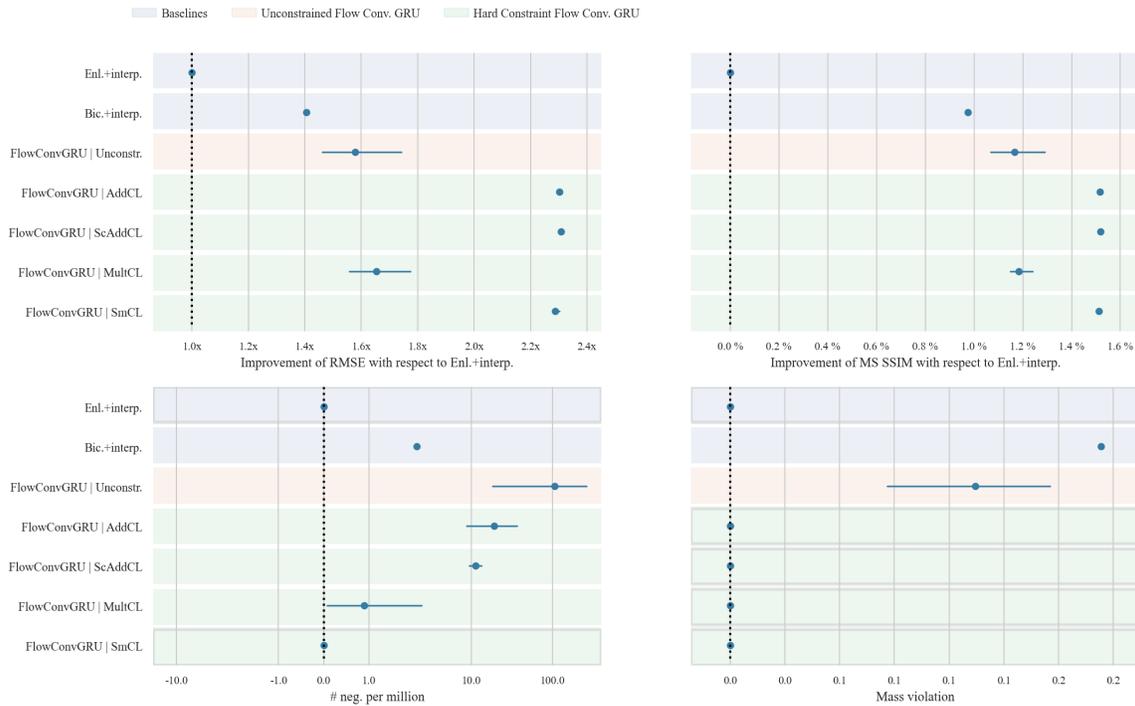


Figure 9: Metrics for different constraining methods applied to an SR FlowConvGRU, calculated over 10,000 test samples of the 4x upsampling water content data. The mean and confidence interval from 3 runs are shown, for RMSE and MS-SSIM relatively to the Enlarge baseline for number of negative pixels (per mil.) and mass conservation violation the absolute values are shown. The framed box indicates a method that achieves zero violation of the physics, no negative pixels or mass conservation up to numerical precision. A table with more metrics can be found in the appendix, see Table 10

difference in errors between the unconstrained and constrained versions, we can see in Figure 17 that constraining leads to lower errors in those areas.

6.7 Limitations

In the case of our WRF dataset, we have seen that the constraining methodology can improve predictive performance even if the underlying constraints are slightly violated by the original data. In cases where low-resolution and its high-resolution counterpart are too far apart, our model is not always able to increase the predictive skill. We built a dataset from two different resolutions of the Norwegian Earth System Model (NorESM) (Seland et al., 2020), and applying our constraining methods improved the visual similarity of the predictions, but decreased the predictive ability. We provide scores and plots in the appendix.

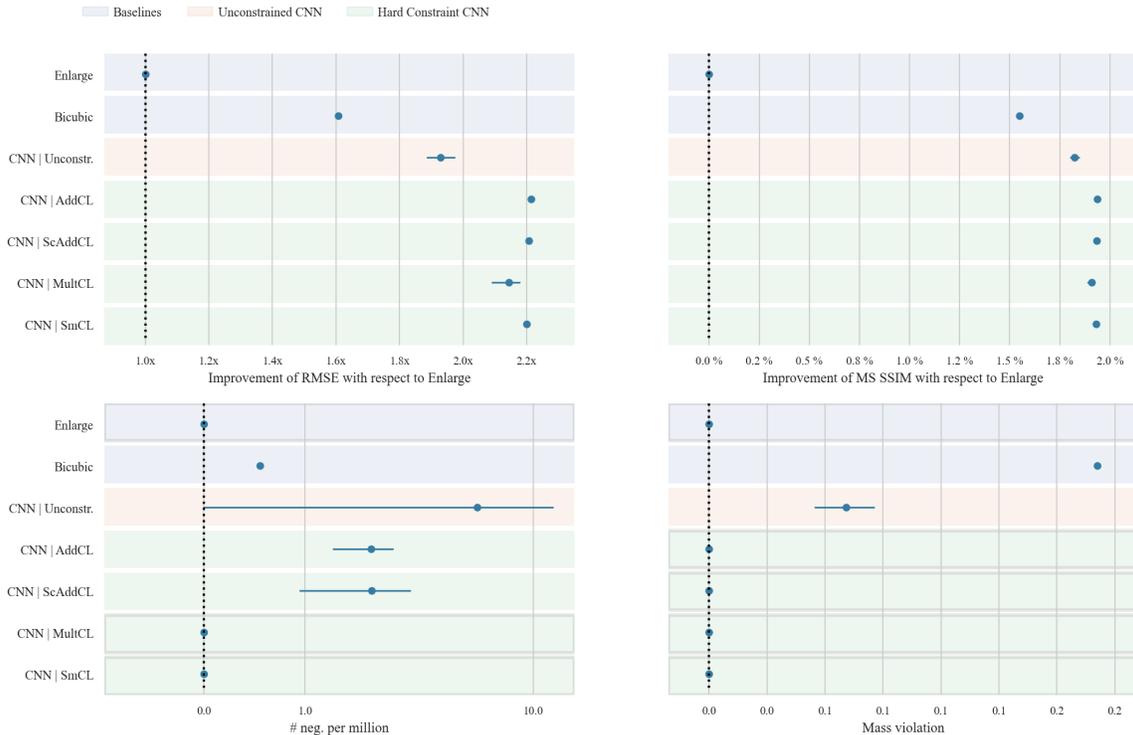


Figure 10: Metrics for different constraining methods applied to the SR CNN applied on the OOD water content dataset, calculated over 10,000 test samples. The mean and confidence interval from 3 runs are shown, for RMSE and MS-SSIM relatively to the Enlarge baseline for number of negative pixels (per mil.) and mass conservation violation the absolute values are shown. The framed box indicates a method that achieves zero violation of the physics, no negative pixels or mass conservation up to numerical precision. A table with more metrics can be found in the appendix, see Table 11.

7. Conclusion and future work

This work presents a novel methodology to incorporate physics-based constraints into neural network architectures for climate downscaling. We show that this method performs well across different deep learning architectures, upsampling factors, predicted quantities, and datasets. We demonstrate its skill both on standard downscaling datasets and on data created by independent simulations. Our constrained models are not only guaranteed to satisfy conservation laws such as mass conservation, but also increase predictive performance across metrics and use cases. Compared to soft-constraining through the loss function, our methodology does not suffer from the common accuracy-constraints enforcement trade-off. Our hard-constraining performance enhancement is not only limited to climate super-resolution but also noticeable in satellite imagery of the lunar surface as well as standard benchmark datasets for natural images. Within the climate context, our constraint layer can

Table 2: Metrics for different constraining methods applied to the SR CNN applied on the WRF temperature data, calculated over 10,000 test samples. The mean is taken over 3 runs. The best scores are highlighted in bold blue.

DATA	MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	CONSTR. VIOL.
T2 WRF	ENLARGE	NONE	1.015	0.648	94.51	0.000
T2 WRF	CNN	NONE	0.952	0.618	94.92	0.181
T2 WRF	CNN	SOFT	1.020	0.660	94.57	0.032
T2 WRF	CNN	SMCL	0.950	0.592	95.25	0.000

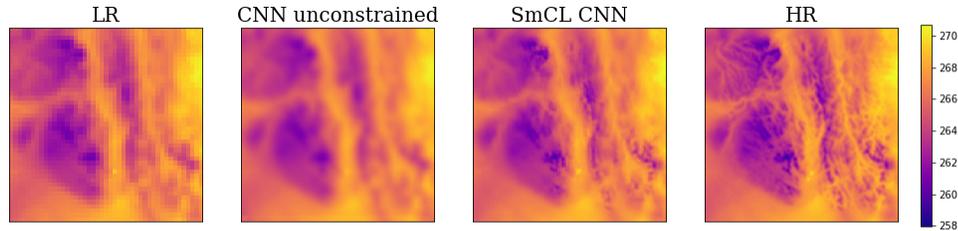


Figure 11: A random prediction for the WRF temperature dataset. We compare unconstrained and softmax-constrained predictions.

help with common issues connected to deep learning applied to downscaling: it dampens the coastal effect, errors get lower in critical regions, out-of-distribution generalization is improved and training can be more stable.

Future work could extend the application of our constraint layer to other tasks than downscaling. Climate model emulation (e.g. Beucler et al. (2021) and Harder et al. (2021)) for example could strongly benefit from a reliable and performance-enhancing method to enforce physical laws.

Table 3: Metrics for different constraining methods applied to the SR-CNN, calculated over the test samples of the lunar dataset. The mean is taken over 3 runs. The best scores are highlighted in bold blue.

DATA	MODEL	CONSTRAINT	RMSE	MAE	SSIM	PSNR
LUNAR	CNN	NONE	0.00217	0.00146	90.08	37.57
LUNAR	CNN	SMCL	0.00213	0.00144	90.40	37.74

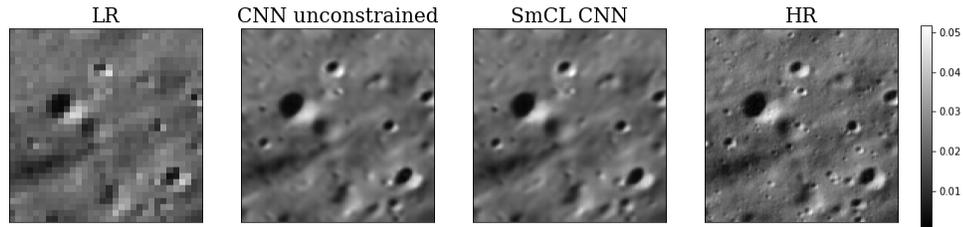


Figure 12: A random sample prediction from the lunar dataset is shown. We compare the unconstrained with the constrained prediction.

Table 4: Metrics of the SR-GAN with and without SmCL calculated over the test datasets Set5, Set14, Urban100, BSD100. The better scores are highlighted in bold blue.

DATA	MODEL	CONSTRAINT	RMSE	MAE	SSIM	PSNR
SET5	SR-GAN	NONE	8.57	4.80	92.48	29.47
SET5	SR-GAN	SMCL	6.61	4.01	93.95	31.73
SET14	SR-GAN	NONE	15.75	8.82	86.06	24.28
SET14	SR-GAN	SMCL	14.07	8.12	87.37	25.16
URBAN100	SR-GAN	NONE	25.00	14.57	81.40	20.17
URBAN100	SR-GAN	SMCL	23.25	13.60	83.19	20.80
BSD100	SR-GAN	NONE	14.38	8.28	85.95	24.97
BSD100	SR-GAN	SMCL	13.52	7.82	87.09	25.50

Acknowledgement and Disclosure of Funding

PH acknowledges the funding received by the Fraunhofer Institute for Industrial Mathematics. DR was funded in part by the Canada CIFAR AI Chairs Program. The authors also are grateful for support from the NSERC Discovery Grants program, material support from NVIDIA in the form of computational resources, and technical support from the Mila IDT team in maintaining the Mila Compute Cluster.

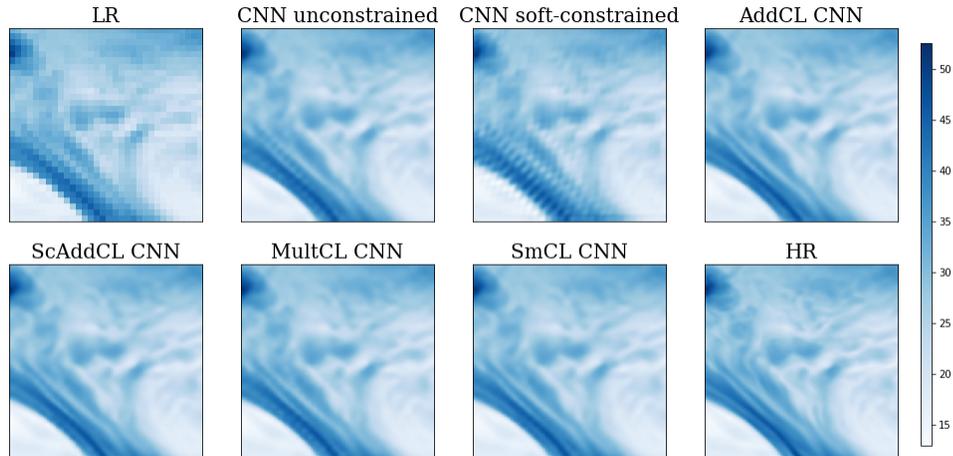


Figure 13: One example image from the test set. Shown here are the LR input, different constrained and unconstrained predictions, and the HR image as a reference.

Appendix A.

Here we investigate the influence of the factor α on the soft-constraining method in more detail. Table 5 shows how the increase of α improves the mass conservation but only up to a value between 0.014 and 0.017. At the same time, it shows that the predictive skill decreases with the increase of α significantly.

Table 5: Metrics calculated over 10,000 validation samples. The best scores are highlighted in bold blue, second best in bold black.

ALPHA	RMSE	MAE	MS-SSIM	MASS VIOL.	#NEG
0.0001	0.241	0.102	99.95	0.021	199
0.001	0.237	0.100	99.96	0.022	19
0.01	0.247	0.103	99.96	0.022	228
0.1	0.252	0.104	99.95	0.023	68
0.9	0.268	0.110	99.95	0.020	2761
0.99	0.297	0.133	99.94	0.014	5085
0.999	0.477	0.261	99.84	0.016	98557
0.9999	0.706	0.433	99.71	0.017	634336
1	2.618	1.814	94.22	0.017	157510

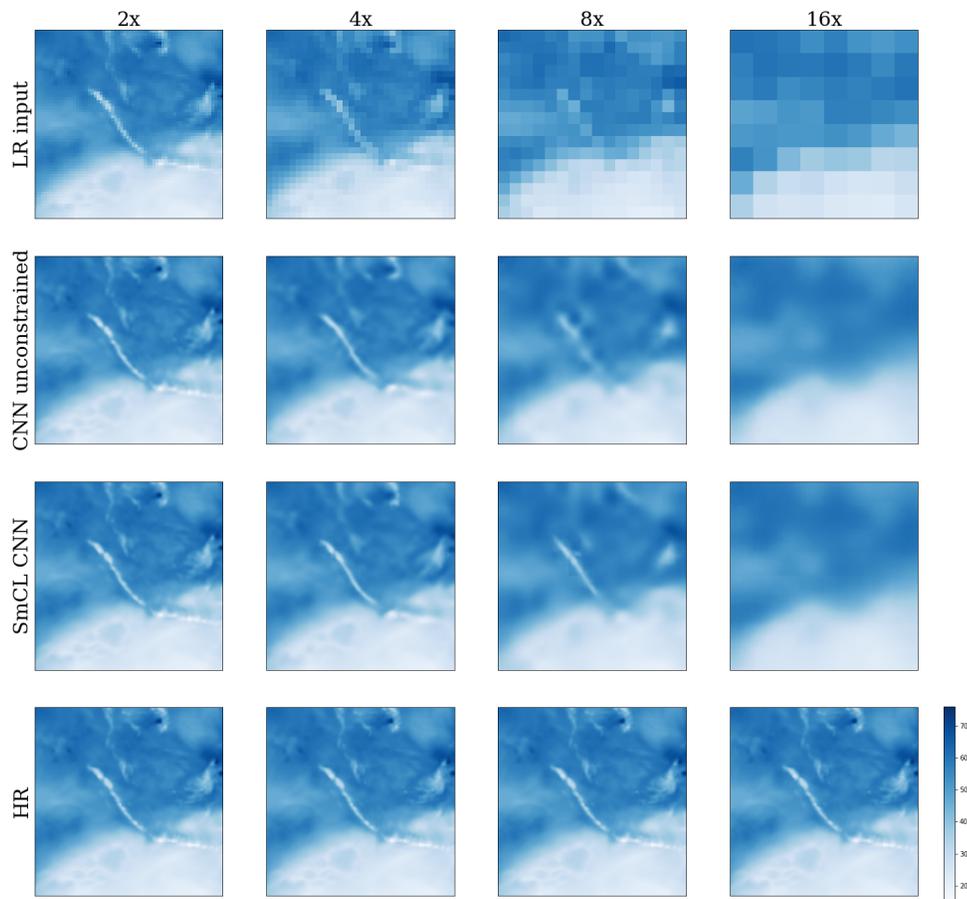


Figure 14: One example image is chosen randomly from the test set. Each model was trained for the same target resolution but with a different upsampling factor. The first row shows the LR inputs for each resolution and the last row the corresponding HR ground truth. The second and third rows show the prediction of an unconstrained CNN and with the SmCL, respectively.

Appendix B.

As natural RGB images have a well-defined range, it is common in CNN and GAN implementations to clip the pixels at inference time to the desired range, removing negative values, for

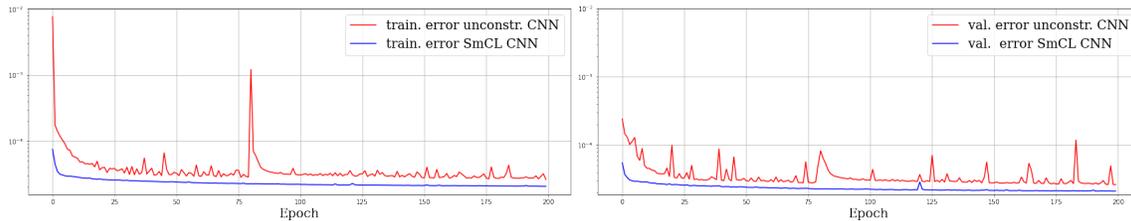


Figure 15: The development of training and validation errors with increasing iterations. Shown for an unconstrained CNN and CNN+SmCL applied to the water content data.

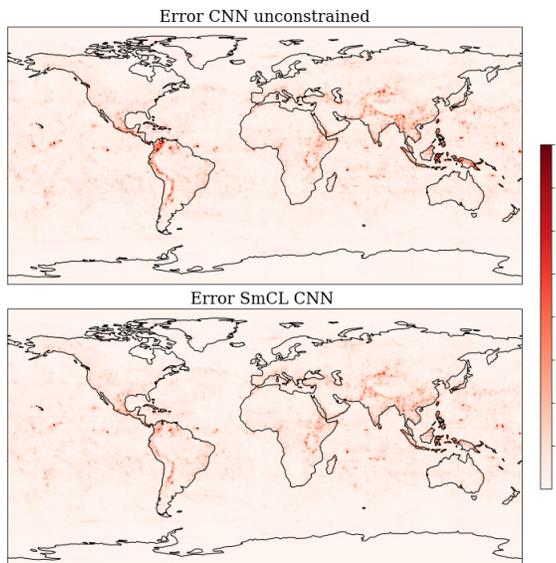


Figure 16: The errors of the global predictions for unconstrained and constrained (SmCL) CNNs, when compared to the ground truth.

example. Here, in Table 6 we show that doing that gives a very small increase in performance, but still performs significantly worse than SmCL, which achieves also zero negative values.

Appendix C.

We show the tables with the mean scores that are displayed as Figures in the main paper and additionally include the MAE.

Appendix D.

We look at additional scores for our water content dataset. We investigate the mean bias (mean over the difference for each pixel value of prediction and truth), the peak signal-to-noise ratio (PSNR), the structural similarity index measure, the Pearson correlation (Corr),

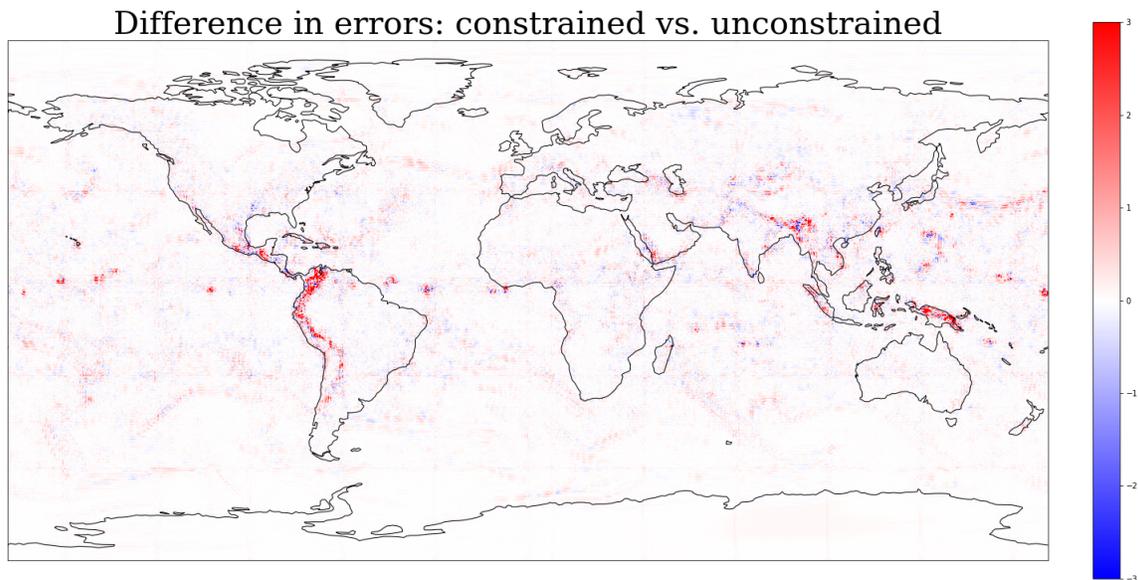


Figure 17: The difference in the errors of constrained and unconstrained predictions from Figure 16. Positive values (red) mean a higher error in the unconstrained version. We trim values at 3, so everything that has a difference greater than 3 is shown as full red for better visibility.

Table 6: Metrics for different constraining methods applied to the SR CNN + clipping applied on the water content dataset, calculated over 10,000 test samples. The mean is taken over 3 runs. The best scores are highlighted in bold blue.

DATA	MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	MASS VIOL.	#NEG
WC	CNN	NONE	0.661	0.327	99.39	0.059	396
WC	CNN	CLIP	0.657	0.326	99.440	0.058	0
WC	CNN	SMCL	0.582	0.291	99.49	0.000	0

and the negative mean (the average magnitude of predicted negative values, the average is calculated over all predicted values, including positive, that are set to zero to calculate the negative mean). These metrics show a similar trend then the metrics shown in the main paper: all of them are improved by adding constraints in our architecture. Without or with soft constraining there are small biases appearing in the predictions, but hard constraining removes those biases. PSNR is a function of the MSE and therefore shows the same trend as it. SSIM and correlation give very similar results, with ScAddCL, AddCL,

Table 7: Metrics for different constraining methods applied to an SR CNN, calculated over 10,000 test samples of the water content data. The mean is taken over 3 runs. The best scores are highlighted in bold blue, second best in bold.

FACT.	MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	MASS VIOL.	#NEG
2X	ENLARGE	NONE	0.422	0.361	99.61	0.000	0
2X	BICUBIC	NONE	0.322	0.137	99.90	0.066	42
2X	CNN	NONE	0.251	0.105	99.95	0.026	230
2X	CNN	SOFT	0.301	0.137	99.23	0.016	17,163
2X	CNN	ADDCL	0.216	0.092	99.96	0.000	215
2X	CNN	ScADDCL	0.199	0.0876	99.96	0.000	3
2X	CNN	MULTCL	0.223	0.094	99.96	0.000	0
2X	CNN	SmCL	0.215	0.094	99.96	0.000	0
4X	ENLARGE	NONE	1.286	0.717	97.60	0.000	0
4X	BICUBIC	NONE	0.800	0.401	99.12	0.169	87
4X	CNN	NONE	0.657	0.326	99.40	0.058	396
4X	CNN	SOFT	0.801	0.410	99.15	0.023	95,373
4X	CNN	ADDCL	0.580	0.290	99.50	0.000	234
4X	CNN	ScADDCL	0.575	0.289	99.50	0.000	11
4X	CNN	MULTCL	0.606	0.300	99.47	0.000	0
4X	CNN	SmCL	0.582	0.291	99.49	0.000	0
8X	ENLARGE	NONE	2.181	1.294	92.39	0.000	0
8X	BICUBIC	NONE	1.557	0.900	96.49	0.318	1,076
8X	CNN	NONE	1.358	0.782	97.15	0.109	2,539
8X	CNN	SOFT	1.640	0.965	96.06	0.029	103,702
8X	CNN	ADDCL	1.267	0.733	97.41	0.000	3,196
8X	CNN	ScADDCL	1.264	0.734	97.41	0.000	26
8X	CNN	MULTCL	1.331	0.733	97.22	0.000	18
8X	CNN	SmCL	1.268	0.734	97.40	0.000	0
16X	ENLARGE	NONE	3.425	2.159	85.55	0.000	0
16X	BICUBIC	NONE	2.723	1.730	91.72	0.510	8,803
16X	CNN	NONE	2.450	1.545	92.68	0.203	682
16X	CNN	SOFT	2.794	1.776	90.74	0.036	369,127
16X	CNN	ADDCL	2.364	1.491	92.96	0.000	75,004
16X	CNN	ScADDCL	2.368	1.495	92.94	0.000	348
16X	CNN	MULTCL	2.409	1.518	92.77	0.000	29
16X	CNN	SmCL	2.368	1.492	92.95	0.000	0

and SmCL showing the best scores. Overall we can see that soft-constraining leads to the most significantly negative predictions, which would cause issues in the context of climate models and predictions.

Table 8: Metrics for different constraining methods applied to an SR GAN, calculated over 10,000 test samples of the 4x upsampling water content data. The mean is taken over 3 runs. The best scores are highlighted in bold blue, and the second best in bold.

MODEL	CONSTRAINT	RMSE	MAE	CRPS	MS-SSIM	MASS VIOL.	#NEG
GAN	NONE	0.628	0.313	0.1522	99.44	0.0453	569
GAN	ADDCL	0.602	0.306	0.1519	99.46	0.000	1211
GAN	SCADDCL	0.604	0.305	0.1508	99.46	0.000	8
GAN	MULTCL	0.732	0.406	0.1978	99.13	0.000	0
GAN	SMCL	0.603	0.310	0.1520	99.46	0.000	0

Table 9: Metrics for different constraining methods applied to an SR ConvGRU, calculated over 10,000 test samples of the water content data. The best scores are highlighted in bold blue, second best in bold.

MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	MASS VIOL.	#NEG
ENLARGE	NONE	1.292	0.718	97.72	0.000	0
BICUBIC	NONE	0.807	0.402	99.16	0.169	355
CONVGRU	NONE	0.672	0.340	99.42	0.102	9094
CONVGRU	ADDCL	0.499	0.260	99.64	0.000	222,793
CONVGRU	SCADDCL	0.499	0.260	99.64	0.000	1736
CONVGRU	MULTCL	0.903	0.472	98.98	0.000	42
CONVGRU	SMCL	0.500	0.260	99.64	0.000	0

Appendix E.

Here we present some visualizations, a prediction by the GAN (Figure 18), the FlowConvGRU (Figure 19, unconstrained and constrained example prediction from BSD100 and Urban100 (Figure 20), and a global prediction for water content (Figure 21).

Appendix F.

Our NorESM dataset is based on the second version of the Norwegian Earth System Model (NorESM2), which is a coupled Earth System Model developed by the NorESM Climate modeling Consortium (NCC), based on the Community Earth System Model, CESM2. We build our dataset on two different runs: NorESM-MM which has a 1-degree resolution for model components and NorESM2-LM which has a 2-degree resolution for atmosphere and land components. We use the temperature at the surface (tas) and a time period from 2015 to 2100. The scenarios ssp126 and ssp585 are used for training ssp370 for validation and ssp245 for testing. By cropping into 64×64 and 32×32 pixels, each scenario contains

Table 10: Metrics for different constraining methods applied to our FlowConvGRU, calculated over 10,000 test samples of the water content dataset. The best scores are highlighted in bold blue, second best in bold.

MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	MASS VIOL.	#NEG
INTERPOLATION	NONE	0.834	0.428	99.10	0.169	352
FLOWCONVGRU	NONE	0.673	0.352	99.40	0.072	2997
FLOWCONVGRU	ADDCL	0.509	0.275	99.63	0.000	6085
FLOWCONVGRU	ScADDCL	0.509	0.274	99.63	0.000	2199
FLOWCONVGRU	MULTCL	0.719	0.383	99.27	0.000	0
FLOWCONVGRU	SMCL	0.514	0.276	99.62	0.000	0

Table 11: Metrics for different constraining methods applied to the SR CNN applied on the OOD water content dataset, calculated over 10,000 test samples. The mean is taken over 3 runs. The best scores are highlighted in bold blue.

DATA	MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	MASS VIOL.	# NEG
OOD	ENLARGE	NONE	1.274	0.711	97.60	0.000	0
OOD	BICUBIC	NONE	0.792	0.397	98.63	0.167	91
OOD	CNN	NONE	0.661	0.327	99.39	0.059	810
OOD	CNN	ADDCL	0.575	0.287	99.50	0.000	272
OOD	CNN	ScADDCL	0.573	0.288	99.50	0.000	35
OOD	CNN	MULTCL	0.591	0.294	99.47	0.000	0
OOD	CNN	SMCL	0.579	0.289	99.49	0.000	0

12k data points. The results for the NorESM data are shown in Table 14: the best scores are in all cases achieved by the unconstrained CNN. This is probably due to the stronger violation of the downscaling constraints between low-resolution and high-resolution samples. We can see a significant difference between the real LR and the HR downsampled, as shown in Figure 23. The visual quality of the prediction, on the other hand, seems to be improved by constraining, an example is shown in Figure 22.

Table 12: Metrics for different constraining methods applied to the SR CNN, calculated over the test set for water vapor, liquid water, and temperature. The mean is taken over 3 runs. For Q_L , RMSE, MAE, and Constr. violation are scaled by a factor of 10^3 for readability. The best scores are highlighted in bold blue, second best in bold.

DATA	MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	CONSTR. VIOL.
Q_v	ENLARGE	NONE	0.474	0.262	94.74	0.000
Q_v	BICUBIC	NONE	0.326	0.182	97.12	0.07
Q_v	CNN	NONE	0.260	0.141	98.14	0.02
Q_v	CNN	ADDCL	0.250	0.133	98.28	0.00
Q_v	CNN	ScADDCL	0.250	0.133	98.28	0.00
Q_v	CNN	MULTCL	0.250	0.133	98.28	0.00
Q_v	CNN	SMCL	0.248	0.132	98.30	BF0.00
Q_L	ENLARGE	NONE	0.0217	0.00862	98.34	0.00000
Q_L	BICUBIC	NONE	0.0186	0.00765	98.96	0.00236
Q_L	CNN	NONE	0.0157	0.00617	99.15	0.00067
Q_L	CNN	ADDCL	0.0155	0.00588	99.18	0.00000
Q_L	CNN	ScADDCL	0.0155	0.00588	99.17	0.00000
Q_L	CNN	MULTCL	0.0166	0.00647	99.06	0.00000
Q_L	CNN	SMCL	0.0155	0.00585	99.17	0.00000
T	ENLARGE	NONE	0.470	0.288	99.03	0.0
T	BICUBIC	NONE	0.281	0.156	99.67	159.1
T	CNN	NONE	0.459	0.287	99.03	139.7
T	CNN	ADDCL	0.276	0.160	99.67	0.0
T	CNN	ScADDCL	0.280	0.163	99.67	0.0
T	CNN	MULTCL	0.270	0.155	99.69	0.0
T	CNN	SMCL	0.272	0.155	99.68	0.0

Table 13: More metrics for different constraining methods applied to an SR CNN, calculated over 10,000 test samples. The best scores are highlighted in bold blue, second best in bold.

FACT.	MODEL	CONSTRAINT	MEAN BIAS	PSNR	SSIM	CORR	NEG MEAN
2X	ENLARGE	NONE	0.000	45.36	98.65	99.75	0.000
2X	BICUBIC	NONE	0.000	51.46	99.71	99.95	0.000
2X	CNN	NONE	-0.003	53.62	99.82	99.97	0.002
2X	CNN	SOFT	-0.002	52.07	99.74	99.94	0.192
2X	CNN	ADDCL	0.000	54.91	99.85	99.98	0.002
2X	CNN	SCADDCL	0.000	55.66	99.87	99.98	0.000
2X	CNN	MULTCL	0.000	54.65	99.84	99.97	0.000
2X	CNN	SMCL	0.000	54.95	99.85	99.98	0.000
4X	ENLARGE	NONE	0.000	39.43	94.91	98.98	0.000
4X	BICUBIC	NONE	0.000	43.55	98.29	99.63	0.000
4X	CNN	NONE	-0.015	45.26	98.70	99.74	0.001
4X	CNN	SOFT	-0.001	43.55	98.15	99.59	0.546
4X	CNN	ADDCL	0.000	46.35	98.89	99.80	0.001
4X	CNN	SCADDCL	0.000	46.42	98.90	99.79	0.000
4X	CNN	MULTCL	0.000	45.98	98.83	99.78	0.000
4X	CNN	SMCL	0.000	46.31	98.88	99.79	0.000
8X	ENLARGE	NONE	0.000	34.84	89.08	96.95	0.000
8X	BICUBIC	NONE	+0.0001	37.77	95.40	98.50	0.006
8X	CNN	NONE	-0.0148	38.96	95.93	98.82	0.012
8X	CNN	SOFT	-0.0071	37.32	94.37	98.22	0.656
8X	CNN	ADDCL	0.000	39.56	96.23	98.96	0.011
8X	CNN	SCADDCL	0.000	39.58	96.24	98.97	0.000
8X	CNN	MULTCL	0.000	39.13	95.99	98.87	0.000
8X	CNN	SMCL	0.000	39.55	96.21	98.96	0.000
16X	ENLARGE	NONE	0.000	30.92	85.20	92.19	0.000
16X	BICUBIC	NONE	+0.0090	32.91	91.99	95.15	0.063
16X	CNN	NONE	-0.0091	33.83	92.48	95.94	0.006
16X	CNN	SOFT	+0.0115	32.70	90.45	94.63	4.233
16X	CNN	ADDCL	0.000	34.14	92.67	96.20	0.581
16X	CNN	SCADDCL	0.000	34.13	92.67	96.18	0.007
16X	CNN	MULTCL	0.000	33.98	92.54	96.07	0.000
16X	CNN	SMCL	0.000	34.13	92.68	96.19	0.000

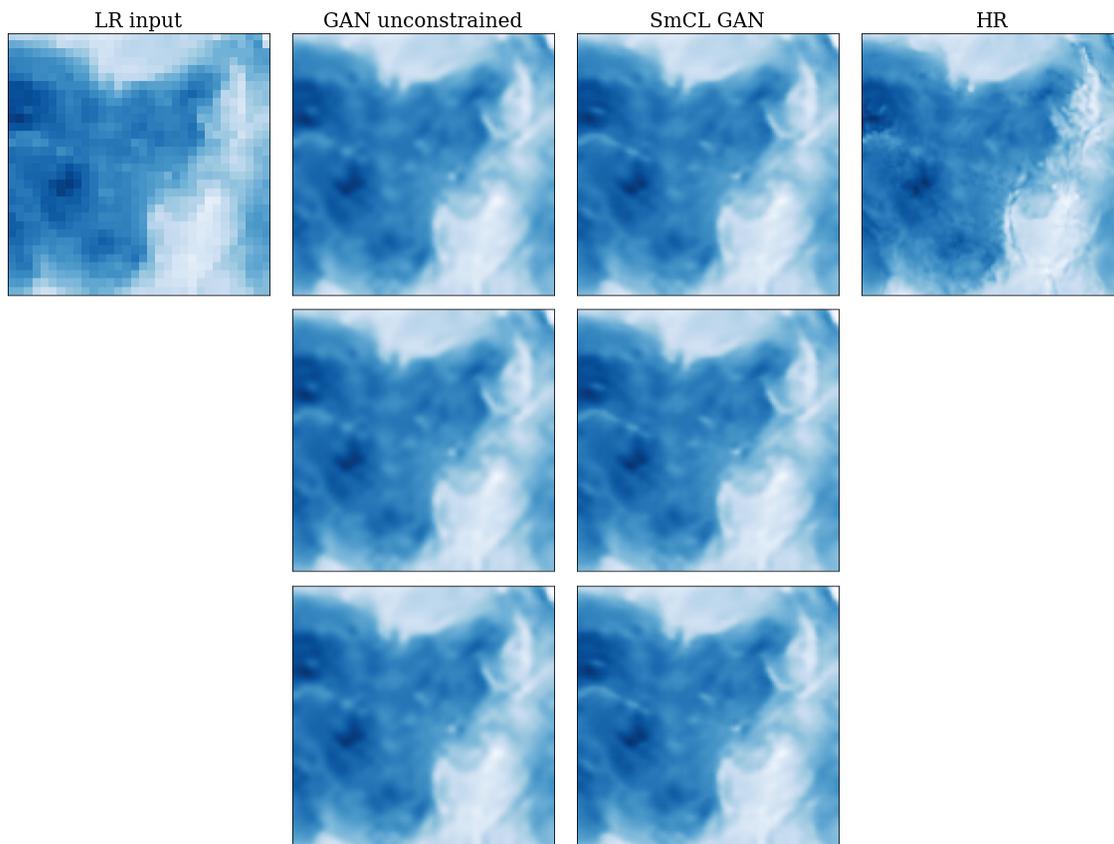


Figure 18: A random sample for the GAN predictions, showing 3 different outputs from the ensemble, constrained and unconstrained.

Table 14: Metrics for different constraining methods applied to the SR CNN, calculated over the test samples of the NorESM dataset. The mean is taken over 3 runs. Best scores are highlighted in bold.

DATA	MODEL	CONSTRAINT	RMSE	MAE	MS-SSIM	CONSTR. VIOL.
TAS NORESM	ENLARGE	NONE	2.987	1.915	95.96	0.000
TAS NORESM	BICUBIC	NONE	2.910	1.864	96.36	0.073
TAS NORESM	CNN	NONE	2.348	1.559	96.93	1.034
TAS NORESM	CNN	SOFT	2.928	1.874	96.28	0.041
TAS NORESMF	CNN	ADDCL	2.885	1.847	96.45	0.000
TAS NORESM	CNN	ScADDCL	2.884	1.846	96.46	0.000
TAS NORESM	CNN	MULTCL	2.888	1.859	96.43	0.000
TAS NORESM	CNN	SmCL	2.885	1.847	96.45	0.000

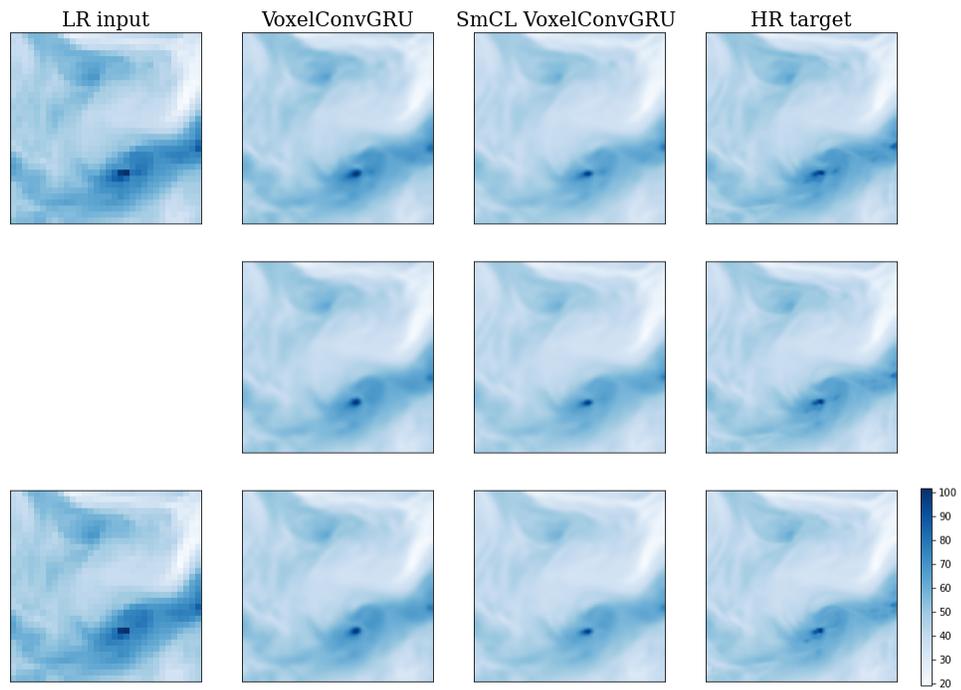


Figure 19: One random test sample and its prediction. Shown here are the two LR input time steps, predictions by both a constrained and unconstrained version of the FlowConvGRU, and the HR sequence as a reference.



Figure 20: Two random images from both the BSD100 and the Urban100 datasets. The first row shows the unconstrained prediction, the second row the constrained prediction using softmax constraining.

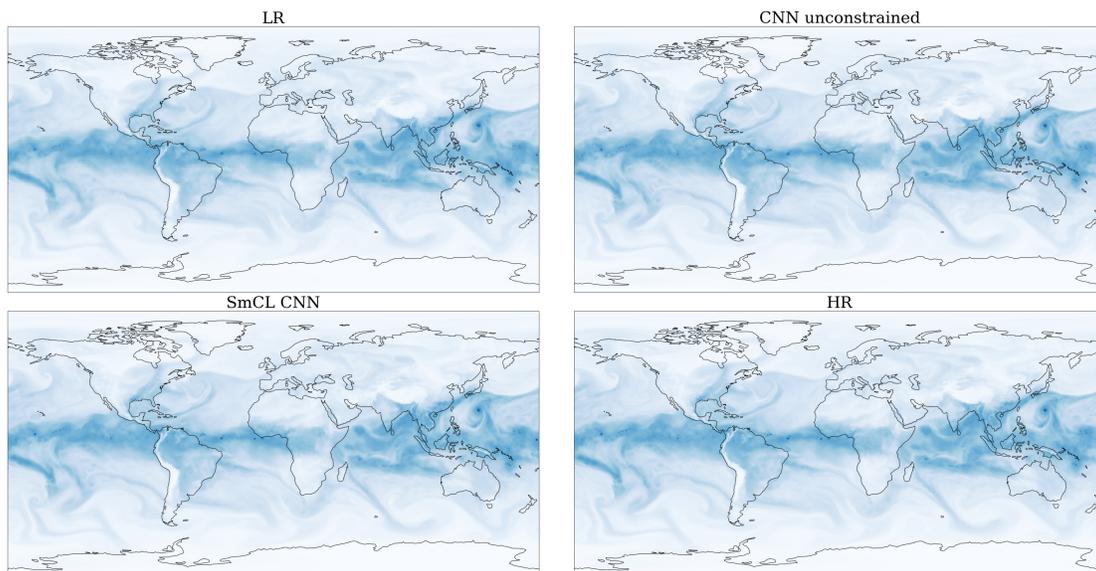


Figure 21: Global data: LR, unconstrained prediction, constrained prediction, and HR.

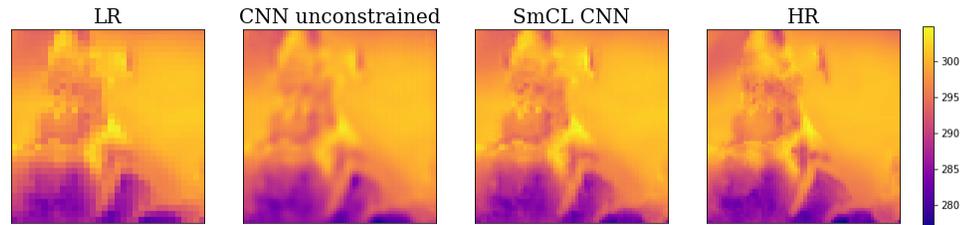


Figure 22: A sample prediction for the NorESM temperature, unconstrained and constrained.

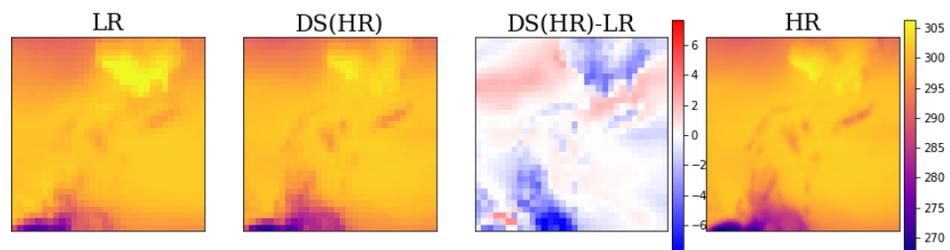


Figure 23: A sample from the NorESM temperature dataset. We compare the low-resolution simulation to the downsampled high-resolution counterpart.

References

- G. A. R. Auger, C. D. Watson, and H. R. Kolar. The influence of weather forecast resolution on the circulation of lake george, ny. *Water Resources Research*, 57(10):e2020WR029552, 2021. doi: <https://doi.org/10.1029/2020WR029552>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020WR029552>. e2020WR029552 2020WR029552.
- J. Baño Medina, R. Manzananas, and J. M. Gutiérrez. Configuration and intercomparison of deep learning neural models for statistical downscaling. *Geoscientific Model Development*, 13(4):2109–2124, 2020. doi: 10.5194/gmd-13-2109-2020. URL <https://gmd.copernicus.org/articles/13/2109/2020/>.
- T. Beucler, S. Rasp, M. Pritchard, and P. Gentine. Achieving conservation of energy in neural network emulators for climate modeling, 2019. URL <https://arxiv.org/abs/1906.06622>.
- T. Beucler, M. Pritchard, S. Rasp, J. Ott, P. Baldi, and P. Gentine. Enforcing analytic constraints in neural networks emulating physical systems. *Phys. Rev. Lett.*, 126:098302, Mar 2021. doi: 10.1103/PhysRevLett.126.098302. URL <https://link.aps.org/doi/10.1103/PhysRevLett.126.098302>.
- C. Chaudhuri and C. Robertson. Cligan: A structurally sensitive convolutional neural network model for statistical downscaling of precipitation from multi-model ensembles. *Water*, 2020.
- J. Delgado-Centeno, P. Harder, B. Moseley, V. Bickel, S. Ganju, F. Kalaitzis, and M. Olivares-Mendez. Single image super-resolution with uncertainty estimation for lunar satellite images. *NeruIPS Workshop ML for Physical Sciences*, 2021.
- C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016. doi: 10.1109/TPAMI.2015.2439281.
- P. Donti, D. Rolnick, and J. Z. Kolter. Dc3: A learning method for optimization with hard constraints. In *International Conference on Learning Representations*, 2021.
- A. Geiss and J. C. Hardin. Strict enforcement of conservation laws and invertibility in cnn-based super resolution for scientific datasets, 2020. URL <https://arxiv.org/abs/2011.05586>.
- A. Geiss, S. Silva, and J. Hardin. Downscaling atmospheric chemistry simulations with physically consistent deep learning. *Geoscientific Model Development Discussions*, 2022: 1–26, 2022. doi: 10.5194/gmd-2022-76. URL <https://gmd.copernicus.org/preprints/gmd-2022-76/>.
- B. Groenke, L. Madaus, and C. Monteleoni. Climalign: Unsupervised statistical downscaling of climate variables via normalizing flows. In *Proceedings of the 10th International Conference on Climate Informatics*, CI2020, page 60–66, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450388481. doi: 10.1145/3429309.3429318. URL <https://doi.org/10.1145/3429309.3429318>.

- W. J. Gutowski, P. A. Ullrich, A. Hall, L. R. Leung, T. A. O'Brien, C. M. Patricola, R. W. Arritt, M. S. Bukovsky, K. V. Calvin, Z. Feng, A. D. Jones, G. J. Kooperman, E. Monier, M. S. Pritchard, S. C. Pryor, Y. Qian, A. M. Rhoades, A. F. Roberts, K. Sakaguchi, N. Urban, and C. Zarzycki. The ongoing need for high-resolution regional climate models: Process understanding and stakeholder information. *Bulletin of the American Meteorological Society*, 101(5):E664 – E683, 2020. doi: 10.1175/BAMS-D-19-0113.1. URL <https://journals.ametsoc.org/view/journals/bams/101/5/bams-d-19-0113.1.xml>.
- P. Harder, D. Watson-Parris, D. Strassel, N. Gauger, P. Stier, and J. Keuper. Physics-informed learning of aerosol microphysics. *arXiv preprint arXiv:2109.10593*, 2021.
- P. Harder, D. Watson-Parris, P. Stier, D. Strassel, N. R. Gauger, and J. Keuper. Physics-informed learning of aerosol microphysics, 2022. URL <https://arxiv.org/abs/2207.11786>.
- N. Harilal, M. Singh, and U. Bhatia. Augmented convolutional lstms for generation of high-resolution climate change projections. *IEEE Access*, 9:25208–25218, 2021. doi: 10.1109/ACCESS.2021.3057500.
- H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut. The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. doi: <https://doi.org/10.1002/qj.3803>. URL <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803>.
- P. Hess, M. Druke, S. Petri, F. M. Strnad, and N. Boers. Physically constrained generative adversarial networks for improving precipitation fields from earth system models. *Nature Machine Intelligence*, 4, 2022.
- C. M. Jiang, S. Esmaeilzadeh, K. Azizzadenesheli, K. Kashinath, M. Mustafa, H. A. Tchelepi, P. Marcus, Prabhat, and A. Anandkumar. Meshfreeflownet: A physics-constrained deep continuous space-time super-resolution framework. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '20*. IEEE Press, 2020. ISBN 9781728199986.
- R. Kurinchi-Vendhan, B. Lütjens, R. Gupta, L. Werner, and D. Newman. Wisosuper: Benchmarking super-resolution methods on wind and solar data, 2021. URL <https://arxiv.org/abs/2109.08770>.
- C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network, 2016. URL <https://arxiv.org/abs/1609.04802>.
- C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative

- adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- J. Leinonen, D. Nerini, and A. Berne. Stochastic super-resolution for downscaling time-evolving atmospheric fields with a generative adversarial network. *IEEE Transactions on Geoscience and Remote Sensing*, 59(9):7211–7223, 2021. doi: 10.1109/TGRS.2020.3032790.
- B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee. Enhanced deep residual networks for single image super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1132–1140, 2017. doi: 10.1109/CVPRW.2017.151.
- Z. Liu, R. A. Yeh, X. Tang, Y. Liu, and A. Agarwala. Video frame synthesis using deep voxel flow. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4473–4481, 2017. doi: 10.1109/ICCV.2017.478.
- A. Lugmayr, M. Danelljan, L. Van Gool, and R. Timofte. Srflow: Learning the super-resolution space with normalizing flow. In *ECCV*, 2020.
- D. Maraun and M. Widmann. *Statistical Downscaling and Bias Correction for Climate Research*. Cambridge University Press, 2018. doi: 10.1017/9781107588783.
- M. Mirza and S. Osindero. Conditional generative adversarial nets, 2014. URL <https://arxiv.org/abs/1411.1784>.
- Ø. Seland, M. Bentsen, D. Olivié, T. Toniazzo, A. Gjermundsen, L. S. Graff, J. B. Debernard, A. K. Gupta, Y.-C. He, A. Kirkevåg, J. Schwinger, J. Tjiputra, K. S. Aas, I. Bethke, Y. Fan, J. Griesfeller, A. Grini, C. Guo, M. Ilicak, I. H. H. Karset, O. Landgren, J. Liakka, K. O. Moseid, A. Nummelin, C. Spensberger, H. Tang, Z. Zhang, C. Heinze, T. Iversen, and M. Schulz. Overview of the norwegian earth system model (noresm2) and key climate response of cmip6 deck, historical, and scenario simulations. *Geoscientific Model Development*, 13(12):6165–6200, 2020. doi: 10.5194/gmd-13-6165-2020. URL <https://gmd.copernicus.org/articles/13/6165/2020/>.
- A. Serifi, T. Günther, and N. Ban. Spatio-temporal downscaling of climate data using convolutional and error-predicting neural networks. *Frontiers in Climate*, 3, 2021. ISSN 2624-9553. doi: 10.3389/fclim.2021.656479. URL <https://www.frontiersin.org/articles/10.3389/fclim.2021.656479>.
- K. Stengel, A. Glaws, D. Hettlinger, and R. N. King. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences*, 117(29):16805–16815, 2020. doi: 10.1073/pnas.1918964117. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1918964117>.
- P. O. Sturm and A. S. Wexler. A mass- and energy-conserving framework for using machine learning to speed computations: a photochemistry example. *Geoscientific Model Development*, 13(9):4435–4442, 2020. doi: 10.5194/gmd-13-4435-2020. URL <https://gmd.copernicus.org/articles/13/4435/2020/>.

- P. O. Sturm and A. S. Wexler. Conservation laws in a neural network architecture: enforcing the atom balance of a julia-based photochemical model (v0.2.0). *Geoscientific Model Development*, 15(8):3417–3431, 2022. doi: 10.5194/gmd-15-3417-2022. URL <https://gmd.copernicus.org/articles/15/3417/2022/>.
- T. Vandal, E. Kodra, S. Ganguly, A. Michaelis, R. Nemani, and A. R. Ganguly. DeepSD: Generating high resolution climate change projections through single image super-resolution. *Association for Computing Machinery*, page 1663–1672, 2017. doi: 10.1145/3097983.3098004. URL <https://doi.org/10.1145/3097983.3098004>.
- J. Wang, Z. Liu, I. Foster, W. Chang, R. Kettimuthu, and V. R. Kotamarthi. Fast and accurate learned multiresolution dynamical downscaling for precipitation. *Geoscientific Model Development*, 14(10):6355–6372, 2021. doi: 10.5194/gmd-14-6355-2021. URL <https://gmd.copernicus.org/articles/14/6355/2021/>.
- X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, C. C. Loy, Y. Qiao, and X. Tang. Esrgan: Enhanced super-resolution generative adversarial networks, 2018a. URL <https://arxiv.org/abs/1809.00219>.
- X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. C. Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *The European Conference on Computer Vision Workshops (ECCVW)*, September 2018b.
- C. D. Watson, C. Wang, T. Lynar, and K. Weldemariam. Investigating two super-resolution methods for downscaling precipitation: ESRGAN and CAR, 2020. URL <https://arxiv.org/abs/2012.01233>.
- F. Yang, H. Yang, J. Fu, H. Lu, and B. Guo. Learning texture transformer network for image super-resolution. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5790–5799, 2020. doi: 10.1109/CVPR42600.2020.00583.
- J. Yuval, P. A. O’Gorman, and C. N. Hill. Use of neural networks for stable, accurate and physically consistent parameterization of subgrid atmospheric processes with good performance at reduced precision. *Geophysical Research Letters*, 48(6):e2020GL091363, 2021. doi: <https://doi.org/10.1029/2020GL091363>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020GL091363>. e2020GL091363 2020GL091363.