


METHODS PAPER

A precipitation downscaling method using a super-resolution deconvolution neural network with step orography

P. Jyoteeshkumar Reddy¹ , Richard Matear¹, John Taylor², Marcus Thatcher³ and Michael Grose¹

¹CSIRO, Environment, Hobart, TAS, Australia

²CSIRO, Data61, Canberra, ACT, Australia

³CSIRO, Environment, Melbourne, VIC, Australia

Corresponding author: P. Jyoteeshkumar Reddy; Email: jyoteesh.papari@csiro.au

Received: 01 February 2023; **Revised:** 31 March 2023; **Accepted:** 31 May 2023

Keywords: Australia; deconvolution neural network; deep learning; precipitation downscaling; super-resolution

Abstract

Coarse spatial resolution in gridded precipitation datasets, reanalysis, and climate model outputs restricts their ability to characterize the localized extreme rain events and limits the use of the coarse resolution information for local to regional scale climate management strategies. Deep learning models have recently been developed to rapidly downscale the coarse resolution precipitation to the high local scale resolution at a much lower cost than dynamic downscaling. However, these existing super-resolution deep learning modeling studies have not rigorously evaluated the model's skill in producing fine-scale spatial variability, particularly over topographic features. These current deep-learning models also have difficulty predicting the complex spatial structure of extreme events. Here, we develop a model based on super-resolution deconvolution neural network (SRDN) to downscale the hourly precipitation and evaluate the predictions. We apply three versions of the SRDN model: (a) SRDN (no orography), (b) SRDN-O (orography only at final resolution enhancement), and (c) SRDN-SO (orography at each step of resolution enhancement). We assess the ability of SRDN-based models to reproduce the fine-scale spatial variability and compare it with the previously used deep learning model (DeepSD). All the models are trained and tested using the Conformal Cubic Atmospheric Model (CCAM) data to perform a 100 to 12.5 km of hourly precipitation downscaling over the Australian region. We found that SRDN-based models, including orography, deliver better fine-scale spatial structures of both climatology and extremes, and significantly improved the deep-learning downscaling. The SRDN-SO model performs well both qualitatively and quantitatively in reconstructing the fine-scale spatial variability of climatology and rainfall extremes over complex orographic regions.

Impact Statement

Extreme rainfall events significantly impact society, ecosystems, and the economy. Developing reliable, locally-relevant information about extremes is critical for building actionable adaptation and mitigation strategies. Statistical downscaling using machine learning algorithms has attracted the attention of researchers for its ability to provide downscaled high-resolution precipitation data at a much lower computational cost than dynamical downscaling. The present study applies a set of deep learning models to rapidly perform 8× downscaling of hourly precipitation data. Our best model outperforms existing deep learning models in reproducing the localised spatial variability of mean and extreme precipitation. Our model can rapidly produce multiple ensemble members of high-resolution precipitation data with reasonable accuracy and with less computational cost, allowing us to better characterise precipitation extremes.

1. Introduction

Spatially downscaling precipitation from coarse resolution datasets and model outputs (~100's of kms) to local scale resolution (~10's of kms or less) allows for a more accurate representation of the spatial variability at fine scales. For climate projections, this will be useful for developing region-specific climate adaptation strategies. For these purposes, recent deep learning methods, especially, the super-resolution approaches, are being widely used for empirical precipitation downscaling (Vandal et al., 2017; Harilal et al., 2021; Kumar et al., 2021; Chandra et al., 2022; Wang et al., 2023). In precipitation downscaling, super-resolution models are trained to learn the relationship between the coarse-scale precipitation data and the fine-scale precipitation data, and then apply this relationship to generate high-resolution precipitation maps. These super-resolution deep learning models are computationally efficient and relatively fast compared to traditional dynamical downscaling.

Previous studies have developed super-resolution-based deep learning models for precipitation downscaling using a convolutional neural network (CNN) (Vandal et al., 2017; Chandra et al., 2022; Wang et al., 2023). For example, Vandal et al. (2017) used the super-resolution CNN (SRCNN) model architecture (Dong et al., 2014) and applied it to the precipitation downscaling by introducing orography as an auxiliary input to the model and named it DeepSD. They showed that DeepSD outperformed the commonly used empirical downscaling techniques like bias correction spatial disaggregation (BCSD) for precipitation downscaling (for a downscale factor of 8) across the United States for grid point-based metrics such as mean bias, root mean square error, and correlation coefficient. Work with DeepSD showed that providing orography as an additional input enhances the accuracy of model output based on grid point metrics (Vandal et al., 2017; Kumar et al., 2021). However, it is not clearly shown that adding orography improves the fine-scale spatial variability of downscaled precipitation products from coarse to a fine grid. Recently, Wang et al. (2023) developed super-resolution deep residual network models with customised loss functions to downscale the hourly precipitation data (for a downscale factor of 12) over the northern region of the Gulf of Mexico and showed their model performs better than the quantile mapping based empirical downscaling technique. In these previous studies, super-resolution deep learning models are mostly based on CNNs. CNNs are mainly used for feature extraction and in these CNN-based models, upscaling is performed using interpolation techniques (with fixed parameters). Unlike CNNs, deconvolution (or transposed convolution) neural networks (DNs) are used for reconstruction purposes and upscaling is done with learnable parameters. In this study, we develop DN-based model architecture for the downscaling of hourly precipitation data over the Australian region. Further, we provided the orographic information at various steps (2, 4, and 8 \times) of resolution improvement in the DN model to systematically assess the influence of the additional orographic information on the precipitation downscaling. We compare our DN models with the baseline DeepSD model and evaluate the ability of deep learning models to reconstruct the fine spatial scale structure of hourly precipitation from its coarse resolution product.

2. Data

For this study, we use high-quality modern reanalysis output dynamically downscaled to fine resolution using the Conformal Cubic Atmospheric Model (CCAM) (McGregor and Dix, 2008). This creates a "model as truth" or "perfect model" exercise to rigorously assess the deep learning models. Output is high-resolution (at 12.5 km) hourly precipitation data and static orography data over the Australian region from 1980 to 2020. The ERA5 reanalysis data (Hersbach et al., 2020) drives the CCAM model. These high-resolution data are conservatively averaged by a scale factor of 8 to 100 km coarse resolution. Deep learning models are trained using data from 1980 to 2012 (until September 15, 2012) and tested on data from 2012 (from September 16, 2012) to 2020. This data split between training and test sets is based on the commonly used 80–20 split rule and allows us to effectively scale over multiple nodes with proper load balance. We have normalized precipitation and orography input data using the training period's maximum and minimum values. For the SRDN models, the normalization is scaled to 100 so that all the values are in

the range of 0 to 100. We have adopted this normalization scheme because most of the hourly precipitation values are small and normalizing them to 0–1 produces a training dataset of nearly all zero values. For the DeepSD model, precipitation is scaled from 0 to 100 and orography from 0 to 1 because our solutions with scaling orography from 0 to 100 produce noisy results (not shown).

3. Methods

3.1. DeepSD

DeepSD is built by stacking the three separate SRCNN models (SRCNN1, SRCNN2, and SRCNN3) with orography as an additional input at every step of the resolution enhancement (Vandal et al., 2017). Each SRCNN model performs $2\times$ resolution enhancement. For example, the SRCNN1 model takes in low-resolution precipitation input at 100 km, interpolates (bicubic) to 50 km, and uses orography at 50 km as a secondary input and then performs convolution operations and outputs the precipitation at 50 km. Both the 100 km input and 50 km target precipitation data for training the SRCNN1 model are produced by conservatively averaging the 12.5 km precipitation data. SRCNN1 has three convolution layers, layer-1 has $64 \times 9 \times 9$ filters with a rectified linear unit (ReLU) non-linear activation, layer-2 has $32 \times 1 \times 1$ filters with ReLU activation, and the last convolution layer has a 5×5 filter that linearly maps the previous layers 32 feature maps to the output.

The model architecture is the same for all three SRCNN models. Only the resolution of the input and target data differs between the three SRCNN models. SRCNN2 model is trained with 50 km precipitation and 25 km orography as inputs to output 25 km precipitation data. Similarly, SRCNN3 model is trained with 25 km precipitation and 12.5 km orography as inputs and outputs 12.5 km. For training purposes of SRCNN2 and SRCNN3, both the input and target precipitation data used for each model come from conservatively averaging the 12.5 km high-resolution CCAM data to the appropriate resolution.

In DeepSD, each SRCNN model is trained separately for 100 epochs with mean square error (MSE) as a loss function using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 1×10^{-4} . The learning rate is further reduced by a factor of 0.1 to a minimum of 1×10^{-5} when the model stopped improving for at least 10 epochs. During prediction, all three SRCNN trained models are stacked together, that is, the SRCNN1 model output precipitation and associated high-resolution orography flow in as an input for the SRCNN2 and similarly SRCNN2 output with respective high-resolution orography is input for SRCNN3. We monitored the loss curve of training and testing data to ensure the model is not overfitting/underfitting the data. The loss curve of these models represents a good fit with no evidence we are overfitting the data (Figure 1).

3.2. SRDN

Our super-resolution DN (SRDN) model is primarily built on deconvolution layers. It has three main deconvolution layers; each consists of $64 \times 7 \times 7$ filters with ReLU non-linear activation and a convolution layer at the end to linearly map the last deconvolution layer feature maps to the output. Each deconvolution layer does a $2\times$ resolution enhancement. This model does a straight $8\times$ resolution enhancement, which means it takes 100 km of precipitation as an input and outputs 12.5 km of precipitation. Guided by studies with DeepSD, we explored including orography as an additional input to the model in several ways.

We develop two models by providing orography at different stages of the model architecture. Version 1 adds to the SRDN by concatenating high-resolution (12.5 km) orography as a channel to the SRDN output high-resolution precipitation (12.5 km) and passing it through a convolution layer of $64 \times 7 \times 7$ filters with ReLU activation and one other convolution layer at the end to linearly map previous layer 64 feature maps to the output. This model is named SRDN-O. Version two, called SRDN-SO, is built on the SRDN-O model by providing orography at multiple steps as feature maps (i.e., 50 km resolution orography is appended to the first deconvolution layer and 25 km resolution orography to the second deconvolution layer, respectively). All three models are trained with MSE as a loss function for




Figure 1. The DeepSD models' (SRCNN1, SRCNN2, and SRCNN3) MSE computed from training (solid line) and test (dashed line) normalized datasets during the model training process over 100 epochs. The MSE values shown are calculated using the normalized data. To estimate the MSE of DeepSD, we need to stack the three models together. Hence the errors in panel (a) flow to panel (b), and the errors in panel (b) flow to panel (c), which leads to model prediction with greater MSE than what is shown in panel (c). The MSE of DeepSD is summarized in Table 1.

30 epochs using the Adam optimizer with a learning rate of 3×10^{-3} . The learning rate is further reduced by a factor of 0.1 to a minimum of 1×10^{-5} when the model stopped improving for at least 10 epochs. The loss curve of these models represents a good fit with no evidence of overfitting the data (see Figure 2).

All the models considered in this study are developed using the TensorFlow (Abadi et al., 2015) and Keras (Chollet et al., 2015) application programming interfaces and trained in a distributed training framework using the Horovod package (Sergeev and Del Balso, 2018). These models are trained on 80 Nvidia V100 GPUs with 32 GB memory at the National Computational Infrastructure in Canberra on the Gadi supercomputer.

4. Results

First, we compare the model's performance based on standard metrics (peak signal-to-noise ratio [PSNR] and mean structural similarity index measure [MSSIM]) of image resolution enhancement models in computer vision applications. The PSNR and MSSIM are calculated for each predicted precipitation map

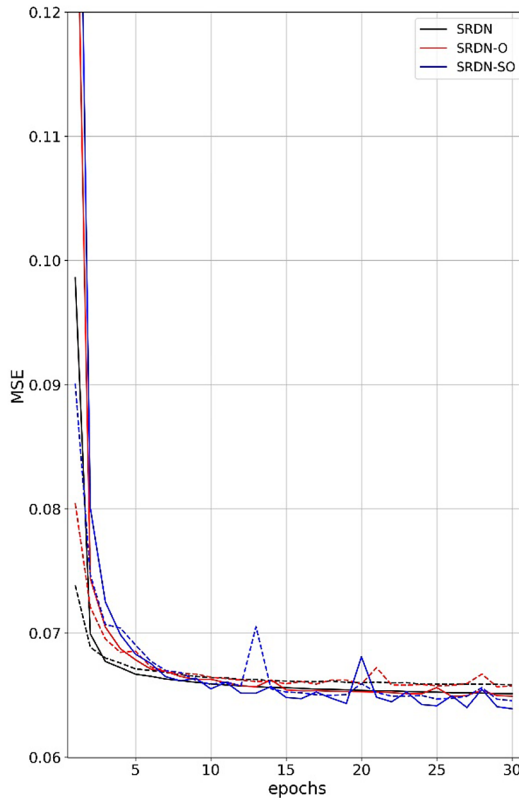


Figure 2. The SRDN models’ MSE error computed from training (solid line) and test (dashed line) normalized datasets during the model training process over 30 epochs. The MSE values shown are calculated using the normalized data.

Table 1. Comparison of mean peak signal-to-noise ratio (PSNR), mean structural similarity index measure (MSSIM), and mean square error (MSE) of predicted precipitation during test period between different models

Model	MSSIM	PSNR	MSE
DeepSD	0.969	40.924	0.1825
SRDN	0.970	41.031	0.1768
SRDN-O	0.970	41.053	0.1760
SRDN-SO	0.970	41.138	0.1729

and are averaged over all test data samples to compute their mean (shown in Table 1). The MSSIM results suggest that all the models perform similarly. The mean PSNR suggests that the SRDN-SO model performs slightly better than the other models in this study in enhancing the resolution of the hourly precipitation field at a downscale factor of 8 (100 to 12.5 km). However, these standard metrics are similar across the models used and are insufficient to identify the best deep-learning model. To quantitatively assess the ability of the deep learning models to reconstruct the 12.5 km from the 100 km precipitation data, we used 2D spatial Fourier Transforms analysis. We compute the predicted precipitation map’s power spectral density (PSD) and compare it with PSD of the target precipitation (Manepalli et al., 2020; Kashinath et al., 2021). Looking at only PSD may provide misleading conclusions when the predicted




Figure 3. The diagonal of one of the quadrants (all quadrants are symmetrical) of mean spatial PSD of the hourly precipitation field versus wavelength (λ) using the target data and the models' predicted output during the test period.

data are noisy. Hence, we also examine the spatial maps qualitatively to gauge how well the model predictions are representing the precipitation data.

First, the PSD is calculated for each of the predicted precipitation maps and averaged over all time samples to get the mean PSD. The diagonal of one quadrant (symmetrical across quadrants) of the mean PSD is shown in Figure 3. All the considered models show less power than the target data at short wavelengths. The DeepSD model's PSD drops substantially more at the short wavelength (around 50 km) than the SRDN models. SRDN-SO model shows slightly more power in both mid-range and short wavelengths than the other models. The additional use of the orography at each step of resolution enhancement appears to improve the model representation of the spatial variability. However, the effect is likely to be even more notable over mountain ranges than minor orographic regions, and this effect may be washed out in this domain statistic, so next, we examine spatial maps.

Figure 4 shows the climatology rainfall with zoomed-in plots in two regions with complex orography (i.e., Papua New Guinea [PNG] and Southeast Australia [SEA]). On average, SRDN-SO captures the mean well (Figure 4a vs. b). At large scale, all deep learning models are quite effective in representing the mean rainfall (figure not shown). If we zoom in on the two regions (PNG and SEA), we start to see the differences between the models and the benefit of orography. It can be visually seen that the SRDN-SO model performs well in reconstructing the spatial structure over complex orographic regions of PNG (Figure 4c–g) and SEA (Figure 4h–l) compared to other models. To quantify the subtle spatial differences between the deep learning models, we show the PSD of climatology rainfall (Figure 4m). The PSDs of SRDN-O and SRDN-SO models closely match the target PSD at mid-range wavelengths; however, they slightly overestimated at short wavelength space.

Figure 5 shows the 95th percentile plot and the respective PSD of the models and target data. On a domain scale, SRDN-SO predicts the 95th percentile well (Figure 5a vs. b). Further, SRDN-SO model extracts the fine-scale spatial structure and magnitude of the 95th percentile precipitation field much better than other models, particularly over the PNG (Figure 5c–g) and SEA (Figure 5h–l) regions. However, the SRDN-SO model underestimates the PSD of the 95th percentile precipitation compared to the target data (Figure 5m). An important note on these results is that all the deep learning models underestimate the target 99th percentile extreme value (not shown). This may be due to the lower sampling size at the extremes of the distribution and structural model features of CCAM that cannot be replicated by deep learning models. We plan to investigate this further.

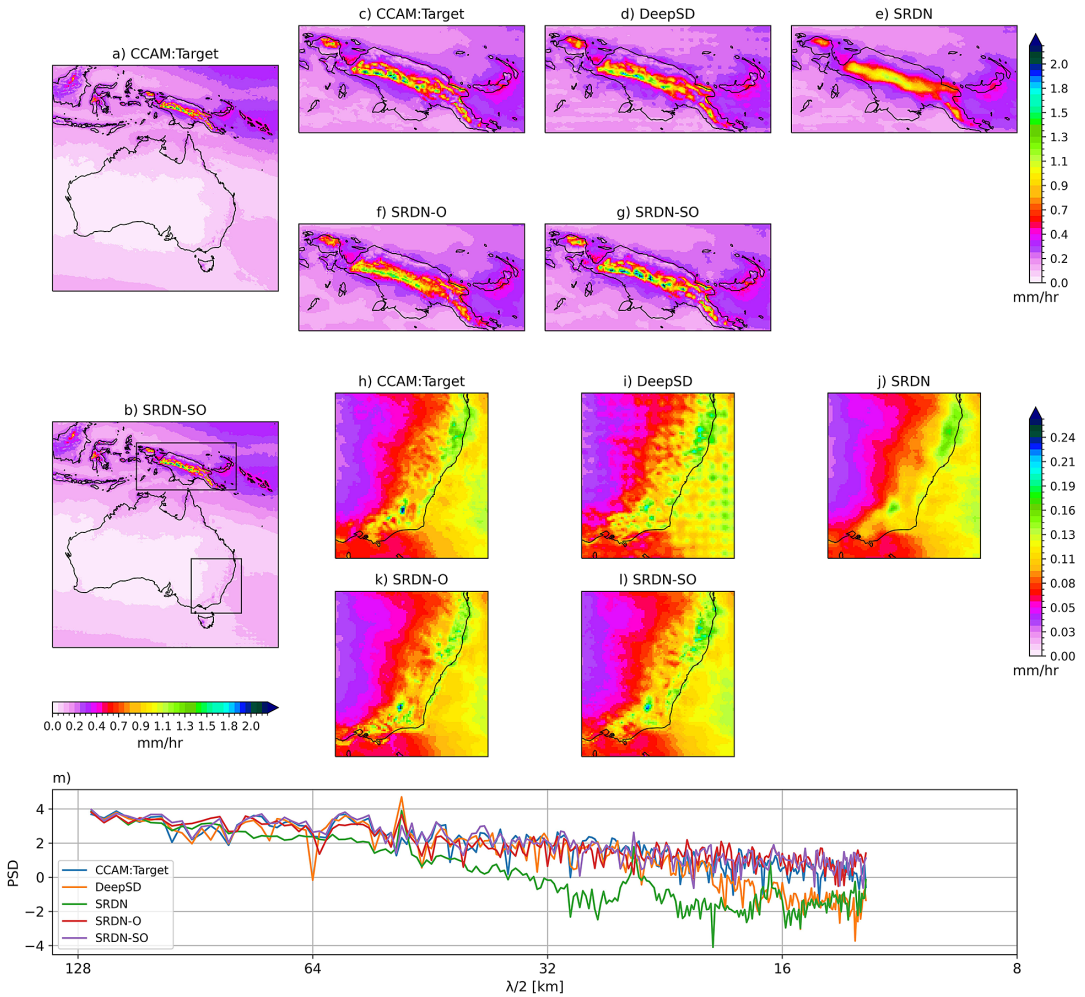


Figure 4. Hourly precipitation climatology over the study region of target (a) and SRDN-SO model (b) predictions during the test period. The top panels (c–g) and middle panels (h–l) show climatology over zoomed-in regions of Papua New Guinea and southeast Australia, respectively. The bottom panel (m) line plot shows the PSD (shown only for mid-range and short wavelengths) of the climatology of the whole domain for different considered models and the target.

5. Conclusion

In this study, we have developed an SRDN model architecture for rapidly downscaling the hourly precipitation by a scale factor of 8 (100 to 12.5 km) over the Australian region that successfully downscales precipitation to accurately reproduces the mean and extreme (95th percentile). We systematically evaluated the use of orography to generate the downscaled precipitation. By utilizing orography at multiple resolution steps, we improved the model’s ability to reconstruct the fine-scale spatial variability of precipitation of both the mean and extreme (95th percentile). Our SRDN-SO model outperforms a previously used deep-learning model (DeepSD) with better skill in reproducing the fine-scale spatial structure of both the climatology and the extremes. This model architecture can be applied to different regions for downscaling hourly precipitation.

The proposed SRDN-SO model has huge potential to rapidly downscale hourly precipitation products from various global climate models ensemble members dynamically downscaled by a coarse resolution

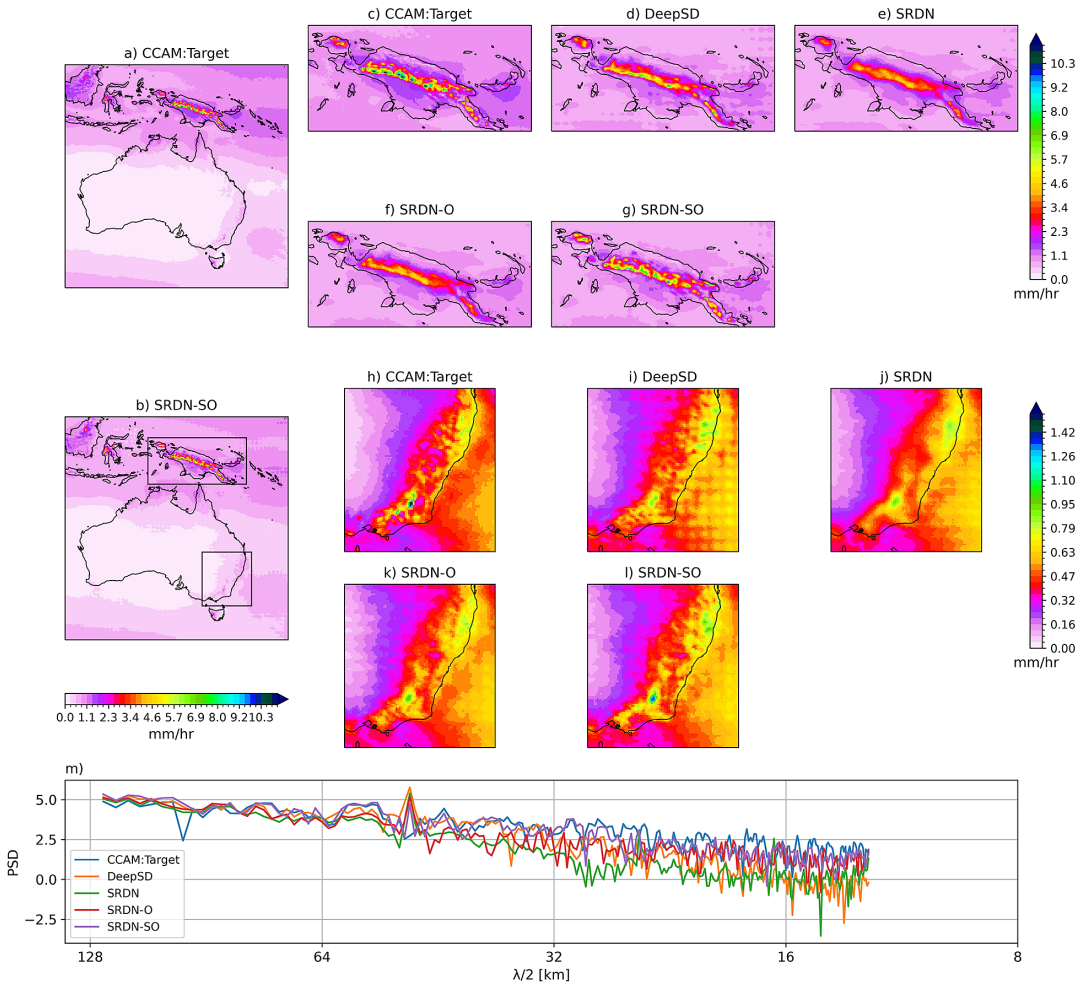


Figure 5. The 95th percentile hourly precipitation for the target (a) and SRDN-SO model (b) predictions during the test period. The top panels (c–g) and middle panels (h–l) show Papua New Guinea and southeast Australia, respectively. The bottom panel (m) line plot shows the corresponding PSD of the target and the deep-learning models.

CCAM to high resolution with less computational cost than straight dynamical downscaling to the desired resolution. This can help produce large ensembles of high-resolution precipitation data needed to characterizing extreme rainfall events. The results here also illustrate the benefits of systematic assessment and evaluation of the results of deep learning models, and we suggest that this should become a more routine part of analysis using deep learning models for downscaling. However, there are a few limitations to our proposed SRDN-SO model. The model cannot accurately represent the PSD at short wavelengths compared to the target data. Several recent studies show models based on Generative Adversarial Networks (GANs), particularly Enhanced Super-Resolution GAN (ESRGAN; Wang et al., 2019) models can accurately represent the PSD even at short wavelengths in wind field for downscaling by a factor of 2 and 4 (Singh et al., 2019; Manepalli et al., 2020; Kashinath et al., 2021). However, these previous studies are based on downscaling winds by a factor of 2 and 4, and not precipitation downscaling by a factor 8. The power spectrum of wind maps is expected to differ from precipitation maps because of the more sporadic and chaotic nature of the precipitation (particularly, convective precipitation), making precipitation downscaling to finer spatial scales more challenging. Hence, based on the previous studies, we

cannot be definite that ESRGAN could accurately represent the PSD for precipitation downscaling of a downscale factor of 8 (100 to 12.5 km). This needs to be further examined in future work. Following the RainNet initiative of providing the precipitation dataset for deep learning-based downscaling (Chen et al., 2022), we will make the data set available as a precipitation downscaling benchmark for the Oceania region and encourage downscaling practitioners to apply their models to this data. In future precipitation downscaling we will assess the benefit of additional input variables such as geopotential height, wind speed, and direction, and explore the inclusion of temporal dependency in the model based on the Chen et al. (2022) approach.

Acknowledgments. We would like to thank National Computing Infrastructure (NCI) Australia for providing computational resources. We also thank TensorFlow, Keras, and Horovod teams for openly providing their APIs.

Author contribution. Conceptualization: P.J.R., R.M., J.T.; Formal analysis: P.J.R.; Methodology: P.J.R., R.M., J.T., M.T., M.G.; Writing–original draft: P.J.R.; Writing–review & editing: R.M., J.T., M.T., M.G.

Competing interest. The authors declare no competing interests exist.

Data availability statement. ERA5 data used to drive the CCAM downscaling is freely available at: <https://doi.org/10.24381/cds.bd0915c6>. The high-resolution CCAM data can be made available upon request to the corresponding author.

Funding statement. The authors would like to acknowledge the funding support of the CSIRO and Australian Climate Service.

Provenance statement. This article is part of the Climate Informatics 2023 proceedings and was accepted in *Environmental Data Science* on the basis of the Climate Informatics peer review process.

References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mane D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viegas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y and Zheng X (2015) *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*. Available at www.tensorflow.org (accessed 3 August 2022).
- Chandra S, Sharma M and Mitra A (2022) ResDeepD: A residual super-resolution network for deep downscaling of daily precipitation over India. *Environmental Data Science 1*, e19. <https://doi.org/10.1017/EDS.2022.23>
- Chen X, Feng K, Liu N, Ni B, Lu Y, Tong Z and Liu Z (2022) RainNet: A large-scale imagery dataset and benchmark for spatial precipitation downscaling. In *36th Conference on Neural Information Processing Systems (NeurIPS 2022)*. Available at https://proceedings.neurips.cc/paper_files/paper/2022/file/3fbf0c1ea0716c03dea93bb6be78dd6f-Paper-Conference.pdf.
- Chollet F, et al. (2015) *Keras*. Available at <https://keras.io> (accessed 3 August 2022).
- Dong C, Loy CC, He K and Tang X (2014) Learning a deep convolutional network for image super-resolution. In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8692 LNCS (PART 4). Cham: Springer, pp. 184–199. https://doi.org/10.1007/978-3-319-10593-2_13
- Harilal N, Singh M and Bhatia U (2021) Augmented convolutional LSTMs for generation of high-resolution climate change projections. *IEEE Access 9*, 25208–25218. <https://doi.org/10.1109/ACCESS.2021.3057500>
- Hersbach H, Bell B, Berrisford P, Hirahara S, Horányi A, Muñoz-Sabater J, Nicolas J, Peubey C, Radu R, Schepers D, Simmons A, Soci C, Abdalla S, Abellan X, Balsamo G, Bechtold P, Biavati G, Bidlot J, Bonavita M, de Chiara G, Dahlgren P, Dee D, Diamantakis M, Dragani R, Flemming J, Forbes R, Fuentes M, Geer A, Haimberger L, Healy S, Hogan RJ, Hólm E, Janisková M, Keeley S, Laloyaux P, Lopez P, Lupu C, Radnoti G, de Rosnay P, Rozum I, Vamborg F, Villaume S and Thépaut JN (2020) The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society 146*(730), 1999–2049. <https://doi.org/10.1002/QJ.3803>
- Kashinath K, Mustafa M, Albert A, Wu JL, Jiang C, Esmacilzadeh S, Azizzadenesheli K, Wang R, Chattopadhyay A, Singh A, Manepalli A, Chirila D, Yu R, Walters R, White B, Xiao H, Tchelepi HA, Marcus P, Anandkumar A, Hassanzadeh P and Prabhat (2021) Physics-informed machine learning: Case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A 379*(2194), 20200093. <https://doi.org/10.1098/RSTA.2020.0093>
- Kingma DP and Ba JL (2015) Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 – Conference Track Proceedings*. Available at <https://arxiv.org/pdf/1412.6980.pdf>
- Kumar B, Chattopadhyay R, Singh M, Chaudhari N, Kodari K and Barve A (2021) Deep-learning based down-scaling of summer monsoon rainfall data over Indian region. *Theoretical and Applied Climatology 143*(3–4), 1145–1156. <https://doi.org/10.1007/s00704-020-03489-6>

- Manepalli A, Singh A, Mudigonda M, White B and Albert A** (2020) Generalization properties of machine learning based weather model downscaling. In *International Conference on Learning Representations (ICLR 2020)*. Available at <https://ai4earthscience.github.io/iclr-2020-workshop/papers/ai4earth25.pdf>
- McGregor JL and Dix MR** (2008) An updated description of the conformal-cubic atmospheric model. In *High Resolution Numerical Modelling of the Atmosphere and Ocean*. New York, NY: Springer, pp. 51–75. https://doi.org/10.1007/978-0-387-49791-4_4/COVER
- Sergeev A and Del Balso M** (2018) *Horovod: Fast and easy distributed deep learning in TensorFlow*. ArXiv./abs/1802.05799; <https://doi.org/10.48550/arxiv.1802.05799>
- Singh A, White B and Albert A** (2019) Numerical weather model super-resolution. In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*. Available at https://ml4physicalsciences.github.io/2019/files/NeurIPS_ML4PS_2019_75.pdf
- Vandal T, Kodra E, Ganguly S, Michaelis A, Nemani R and Ganguly AR** (2017) DeepSD: Generating high resolution climate change projections through single image super-resolution. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Part F129685*. Association for Computing Machinery, pp. 1663–1672. <https://doi.org/10.1145/3097983.3098004>
- Wang F, Tian D and Carroll M** (2023) Customized deep learning for precipitation bias correction and downscaling. *Geoscientific Model Development* 16(2), 535–556. <https://doi.org/10.5194/GMD-16-535-2023>
- Wang X, Yu K, Wu S, Gu J, Liu Y, Dong C, Loy CC, Qiao Y and Tang X** (2019) ESRGAN: Enhanced super-resolution generative adversarial networks. In: Leal-Taixé L and Roth S (eds), *Computer Vision – ECCV 2018 Workshops. ECCV 2018. Lecture Notes in Computer Science*, vol. 11133. Cham: Springer. https://doi.org/10.1007/978-3-030-11021-5_5; Available at <https://github.com/xinntao/ESRGAN> (accessed 2 March 2023).

Cite this article: Reddy P.J, Matear R, Taylor J, Thatcher M and Grose M (2023). A precipitation downscaling method using a super-resolution deconvolution neural network with step orography. *Environmental Data Science*, 2: e17. doi:10.1017/eds.2023.18