

GENERATIVE MODELING OF SPATIO-TEMPORAL WEATHER PATTERNS WITH EXTREME EVENT CONDITIONING

Konstantin Klemmer
University of Warwick & TUM
k.klemmer@warwick.ac.uk

Sudipan Saha
TUM

Matthias Kahl
TUM

Tianlin Xu
LSE

Xiao Xiang Zhu
TUM & DLR

ABSTRACT

Deep generative models are increasingly used to gain insights in the geospatial data domain, e.g., for climate data. However, most existing approaches work with temporal snapshots or assume 1D time-series; few are able to capture spatio-temporal processes simultaneously. Beyond this, Earth-systems data often exhibit highly irregular and complex patterns, for example caused by extreme weather events. Because of climate change, these phenomena are only increasing in frequency. Here, we proposed a novel GAN-based approach for generating spatio-temporal weather patterns conditioned on detected extreme events. Our approach augments GAN generator and discriminator with an encoded extreme weather event segmentation mask. These segmentation masks can be created from raw input using existing event detection frameworks. As such, our approach is highly modular and can be combined with custom GAN architectures. We highlight the applicability of our proposed approach in experiments with real-world surface radiation and zonal wind data.

1 INTRODUCTION

Spatio-temporal modeling is important to many fields, including remote sensing (Saha et al., 2020) and climate sciences (Racah et al., 2017). With the emergence of the big-data era, large scale spatio-temporal datasets have increasingly become available across various domains: from satellite images (Sumbul et al., 2019) to large-scale mapping efforts (Weber & Haklay, 2008), and also weather data (Racah et al., 2017), the application area we are looking at in this study. The topic of climate change has gained special attention of the machine learning community in the last few years (Rolnick et al., 2019; Scher, 2018), while it has been acknowledged that deep learning may provide powerful tools for Earth-systems modeling (Reichstein et al., 2019; Camps-Valls et al., 2021). Simulation can be a useful approach for forecasting the future impacts of climate change, allowing decision makers to design data-driven intervention strategies and build a better understanding of how climate change is affecting the planet.

Apart from the great importance of the applications related to spatio-temporal climate and weather data, such data is also very interesting for machine learning researchers as it poses a number of unique challenges: the data is generally multi-modal with distribution shifts and spatio-temporal dependencies. The interactions between different climate and weather indicators is not fully understood (Racah et al., 2017). Lastly, the data exhibits various anomalies—extreme weather events—a phenomenon only accelerated by climate change.

Although spatio-temporal climate and weather modeling has been studied in past (Racah et al., 2017), most of the existing work relies on some form of feature engineering. Moreover, studies including temporal analysis often treat the data as 1-D time series and ignore the spatial complexity of the data. Inspired by the success of different generative temporal modeling methods (Xiong et al., 2018; Xu et al., 2020), in this work we train implicit generative adversarial net (GAN) mod-

els optimized for producing spatio-temporal weather patterns. In particular, we examine how prior knowledge on extreme weather events can be incorporated in these generative models. Extreme events can be detected in existing weather data using toolkits like TECA (Prabhat et al., 2012), allowing for the creation of segmentation mask from any given weather data input. As such, the contribution of this study lies in a novel approach for embedding multi-class, spatio-temporal segmentation masks as context vector for conditional GANs. Our proposed method uses a flexible mask encoder, allowing for an easy integration into existing video GANs. We highlight this applicability in experiments with real-world weather data.

The rest of this paper is organized as follows: we briefly discuss the existing literature in Section 2. Our proposed method is introduced in Section 3. We present the dataset and our experimental results in Section 4. We conclude with a brief discussion on potential future work in Section 5.

2 RELATED WORK

Considering the relevance to our work, here we briefly discuss recent developments related to: *i*) machine learning (ML) based climate and weather modeling and *ii*) spatio-temporal generative adversarial networks.

2.1 ML BASED CLIMATE AND WEATHER MODELING

Over the last years, various studies have applied machine learning, especially deep learning, to devise better methods for climate modeling and related challenges. Racah et al. (2017) presented a multi-channel spatio-temporal CNN architecture that can exploit temporal information and unlabeled climate data to improve the localization of extreme weather events. Scher (2018) presented a method for forecasting the temporal evolution of weather data. Different deep learning frameworks, e.g., autoencoder (Hossain et al., 2015), CNN (Qiu et al., 2017), and long short-term memory (LSTM) networks (Karevan & Suykens, 2018) have been used for weather prediction. In (Hossain et al., 2015), an approach for estimating air temperature from historical pressure, humidity, and temperature data is proposed using stacked denoising autoencoder. Qiu et al. (2017) leverage the correlation between multiple sites for weather prediction via a multi-task CNN. Karevan & Suykens (2018), use a spatio-temporal LSTM model for temperature prediction of 5 cities in west Europe. Kadow et al. (2020) performed monthly reconstructions of the missing climate data via image inpainting and transfer learning. Our proposed model can generate synthetic weather patterns conditioned on extreme event masks and can be useful for supplementing climate data in some of the above-mentioned tasks, e.g. localization of extreme weather events Racah et al. (2017) and forecasting temporal evolution Scher (2018).

2.2 SPATIO-TEMPORAL GANS

Aligned with its popularity in image modeling, GANs (Goodfellow et al., 2014) have been applied to spatial and spatio-temporal modeling, especially video synthesis and analysis: Xiong et al. (2018) proposed an approach to generate long-term future frames given first frame of a video. Chen et al. (2018) devise a model to generate videos from the starting frame and a corresponding audio file. Clark et al. (2019) focus on realistic temporal sequence generation through a dual strategy, employing a spatial discriminator and a temporal discriminator in combination. Several studies have investigated video to video generation / translation, e.g. for video-realistic animations of real humans Liu et al. (2019). Beyond video related applications, GANs have also been applied to other spatio-temporal data, e.g. geostatistical data (Zheng et al., 2020) and multi-temporal remote sensing data (Saha et al., 2019; Ebel et al., 2020). GANs are further becoming increasingly popular in the GIS community (Klemmer & B Neill, 2020; Zhu et al., 2019).

We propose an approach to condition a given GAN with an embedding of the extreme event mask corresponding to the input data. As such, our method can be used in combination with any of the spatio-temporal GANs mentioned here.




Figure 1: Proposed conditional GAN (agnostic of the customizable GAN backbone) with segmentation mask context, captured through the proposed mask encoder.

3 EXTREME EVENT CONDITIONAL GAN

Typically, GANs consists of a generator G and a discriminator D . G aims to generate realistic synthetic data $G(z)$ from a given random noise input z , while D learns to distinguish real data x from fake data $G(z)$. Both networks are then trained by playing a min-max game against each other.

GANs can be conditioned on some context c relevant to the data, such as class labels (Mirza & Osindero, 2014). This is facilitated by incorporating the context into both networks, so that the generator output is given as $G(z, c)$ and the discriminator output is $D([x, G(z)], c)$. However, while integrating c is trivial for binary or multi-class labels which can easily presented in vector form, this task is trickier for (multi-class) segmentation masks. We propose a flexible mask encoder to embed the segmentation mask in the desired shape, customizable with any GAN backbone. Extreme event segmentation masks come as binary values in the same shape as the input, comprising one channel per class of event.

As we require our encoder to capture both the spatial and temporal dynamics of the extreme weather events, we chose an architecture of three $2d$ -convolutional layers, each followed by a batch normalization layer and a leaky ReLU activation, in combination with two LSTM layers, again performing batch normalization between the first and second LSTM layer. The output dimension of the last LSTM layer can then be adapted according to the spatio-temporal GAN backbone. This approach follows common practice in dealing with spatio-temporal and video data, first using convolutional layers to encode spatial dynamics, while using LSTM layers to capture the sequential nature of the data. As such, the learnt embedding is able to propagate this information to both generator and discriminator.

To test out our proposed mask encoder, and the general applicability of our approach, we chose a spatio-temporal GAN recently developed by Xu et al. (2020): COT-GAN utilizes the sequential nature of temporal data to deploy optimization based on causal optimal transport (COT). This extension of optimal transport ensures that the transformation from source to target probability mass can always only depend on sequential input up to the current time step t . The authors find COT-GAN to be efficient, stable and to perform better than comparable approaches. While recent approaches (Clark et al., 2019) in video generation rely on a dual discriminator strategy to disentangle spatial and temporal complexities, COT-GAN on the other hand actively integrates the sequential nature of the data into the loss function. To compute the causal loss, COT-GAN also employs a dual discriminator approach (referred to as h and M in the original paper).

Due to these advantages, we chose to apply COT-GAN as the GAN backbone in our experiments, augmenting its discriminators and generator with our aforementioned mask encoder. Figure 1 outlines the proposed GAN framework.

4 EXPERIMENTAL RESULTS

Data: We test our conditional GAN using two different weather measures and corresponding extreme event segmentation masks, provided by the *ExtremeWeather* dataset (Racah et al., 2017). We




Figure 2: Comparison between real and generated data, including the corresponding extreme event segmentation mask. Row 1 shows a real extreme event segmentation mask for 10 consecutive time steps. Rows 2-3 show the real and generated radiative surface temperature corresponding to the extreme event pattern. Rows 4-5 show real and generated zonal winds at 850mbar corresponding to the extreme event pattern.

chose to model radiative surface temperature and zonal wind at 850mbar pressure surface. Our training data comprises the whole year of 2004 with four observations per day. At each time step, we extract the variables of interest as well as the extreme event mask, all provided as a 128×196 pixel image. The segmentation mask denotes the presence of an extreme event with pixel values of 1, opposed to 0 values. The segmentation mask comprises four channels, one for each category of extreme event: tropical depression, tropical cyclone, extratropical cyclone and atmospheric river.

Experimental setting: We run experiments generating synthetic samples of the measures described above, conditioned on extreme events. Our GANs are trained for 300 epochs using Tesla K80, Tesla T4, and Tesla P100 GPUs. We use the Adam algorithm with decoupled weight decay (Loshchilov & Hutter, 2019) to optimize our model.

Results: Our results, highlighted in Figure 2, show that training our model for a relatively short time of 300 epochs already provides impressive results, given the complexity of the data input. These initial findings point to a larger potential of implicit generative models for learning climate and weather processes. We believe that coupled with a more powerful and specialised GAN backbone, our proposed approach might become useful too for simulating and forecasting weather patterns in the presence of extreme and adverse events.

5 CONCLUSIONS

In this paper, we proposed a conditional GAN based method for generating synthetic spatio-temporal weather patterns. The proposed method effectively ingests extreme event segmentation masks to generate corresponding weather patterns. The generated data shows strong visual resemblance to real data. The proposed model can be useful for data augmentation in training robust extreme event predictors and weather nowcasting. Our work is an attempt towards better understanding the complex weather patterns and their implication in climate change.

In future studies, we want to build on these preliminary findings. In particular, we want to explore more weather and climate indicators, as well as different extreme events. We further want to devise a more specialized GAN architecture for dealing with this kind of data. Lastly, we seek to explore different options to integrate extreme events into GANs, for example through auxiliary segmentation tasks.

REFERENCES

Gustau Camps-Valls, Devis Tuia, Xiao Xiang Zhu, and Markus Reichstein (Editors). *Deep learning for the Earth Sciences: A comprehensive approach to remote sensing, climate science and*

- geosciences*. Wiley & Sons, 2021. ISBN 978-1-119-64614-3.
- Lele Chen, Zhiheng Li, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Lip movements generation at a glance. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 520–535, 2018.
- Aidan Clark, Jeff Donahue, and Karen Simonyan. Adversarial video generation on complex datasets. *arXiv preprint arXiv:1907.06571*, 2019.
- Patrick Ebel, Andrea Meraner, Michael Schmitt, and Xiao Xiang Zhu. Multisensor data fusion for cloud removal in global and all-season sentinel-2 imagery. *IEEE Transactions on Geoscience and Remote Sensing*, pp. 1–13, 2020. doi: 10.1109/TGRS.2020.3024744.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 2672–2680, 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets>.
- Moinul Hossain, Banafsheh Rekabdar, Sushil J Louis, and Sergiu Dascalu. Forecasting the weather of nevada: A deep learning approach. In *2015 international joint conference on neural networks (IJCNN)*, pp. 1–6. IEEE, 2015.
- Christopher Kadow, David Matthew Hall, and Uwe Ulbrich. Artificial intelligence reconstructs missing climate information. *Nature Geoscience*, pp. 1–6, 2020.
- Zahra Karevan and Johan AK Suykens. Spatio-temporal stacked lstm for temperature prediction in weather forecasting. *arXiv preprint arXiv:1811.06341*, 2018.
- Konstantin Klemmer and Daniel B Neill. Sxl: Spatially explicit learning of geographic processes with auxiliary tasks. *arXiv preprint arXiv:2006.08571*, 2020.
- Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeonwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. *ACM Transactions on Graphics (TOG)*, 38(5):1–14, 2019.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets, 2014.
- Prabhat, Oliver Rübel, Surendra Byna, Kesheng Wu, Fuyu Li, Michael Wehner, and Wes Bethel. TECA: A parallel toolkit for extreme climate analysis. In *Procedia Computer Science*, 2012. doi: 10.1016/j.procs.2012.04.093.
- Minghui Qiu, Peilin Zhao, Ke Zhang, Jun Huang, Xing Shi, Xiaoguang Wang, and Wei Chu. A short-term rainfall prediction model using multi-task convolutional neural networks. In *2017 IEEE International Conference on Data Mining (ICDM)*, pp. 395–404. IEEE, 2017.
- Evan Racah, Christopher Beckham, Tegan Maharaj, Samira Ebrahimi Kahou, Mr Prabhat, and Chris Pal. Extremeweather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events. In *Advances in Neural Information Processing Systems*, pp. 3402–3413, 2017.
- Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, and Prabhat. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743):195–204, feb 2019. ISSN 14764687. doi: 10.1038/s41586-019-0912-1. URL <http://www.nature.com/articles/s41586-019-0912-1>.
- David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavov, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *arXiv preprint arXiv:1906.05433*, 2019.
- Sudipan Saha, Francesca Bovolo, and Lorenzo Bruzzone. Unsupervised multiple-change detection in vhr multisensor images via deep-learning based adaptation. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pp. 5033–5036. IEEE, 2019.

- Sudipan Saha, Francesca Bovolo, and Lorenzo Bruzzone. Change detection in image time-series using unsupervised lstm. *IEEE Geoscience and Remote Sensing Letters*, 2020.
- Sebastian Scher. Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with deep learning. *Geophysical Research Letters*, 45(22):12–616, 2018.
- Gencer Sumbul, Marcela Charfuelan, Begüm Demir, and Volker Markl. Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pp. 5901–5904. IEEE, 2019.
- Patrick Weber and Muki Haklay. OpenStreetMap: user-generated street maps. *IEEE Pervasive Computing*, 2008.
- Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, and Jiebo Luo. Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2364–2373, 2018.
- Tianlin Xu, Li K Wenliang, Michael Munn, and Beatrice Acciaio. Cot-gan: Generating sequential data via causal optimal transport. In *Advances in Neural Information Processing Systems*, pp. 8798–8809, 2020.
- Qiang Zheng, Lingzao Zeng, and George Em Karniadakis. Physics-informed semantic inpainting: Application to geostatistical modeling. *Journal of Computational Physics*, 419:109676, 2020.
- Di Zhu, Ximeng Cheng, Fan Zhang, Xin Yao, Yong Gao, and Yu Liu. Spatial interpolation using conditional generative adversarial neural networks. *International Journal of Geographical Information Science*, pp. 1–24, apr 2019. ISSN 1365-8816. doi: 10.1080/1365881YYxxxxxxx. URL <https://www.tandfonline.com/doi/full/10.1080/13658816.2019.1599122>.