

VISUALIZING THE CONSEQUENCES OF CLIMATE CHANGE USING CYCLE-CONSISTENT ADVERSARIAL NETWORKS

Victor Schmidt*, Alexandra Luccioni*, S. Karthik Mukkavilli, Kris Sankaran, & Yoshua Bengio
Montreal Institute for Learning Algorithms
Montreal, Canada
{schmidtv, luccionis, mukkavis}@mila.quebec

Narmada Balasooriya
ConscientAI Labs,
Colombo, Sri Lanka

Jennifer Chayes
Microsoft Research New England
Cambridge, Massachusetts

ABSTRACT

We present a project that aims to generate images that depict accurate, vivid, and personalized outcomes of climate change using Cycle-Consistent Adversarial Networks (CycleGANs). By training our CycleGAN model on street-view images of houses before and after extreme weather events (e.g. floods, forest fires, etc.), we learn a mapping that can then be applied to images of locations that have not yet experienced these events. This visual transformation is paired with climate model predictions to assess likelihood and type of climate-related events in the long term (50 years) in order to bring the future closer in the viewer’s mind. The eventual goal of our project is to enable individuals to make more informed choices about their climate future by creating a more visceral understanding of the effects of climate change, while maintaining scientific credibility by drawing on climate model projections.

1 INTRODUCTION

It is difficult to downplay the importance of fighting climate change. A recent report from the Intergovernmental Panel on Climate Change has determined that dramatic and rapid changes to the global economy are required in order to avoid increasing climate-related risks for natural and human systems (IPCC, 2018). However, necessary system overhauls require governmental interventions, which are difficult without strong public support. In fact, recent studies have shown that political will is currently the main obstacle to keeping temperature rise within the limits proposed by the IPCC, i.e. 1.5°C (Smith et al., 2019).

Unfortunately, public awareness and concern about climate change often does not match the magnitude of its threat to humans and our environment (Pidgeon, 2012; Weber & Stern, 2011). One reason for this mismatch is that it is difficult for people to mentally simulate the complex and probabilistic effects of climate change (O’Neill & Hulme, 2009). People often discount the impact that their actions will have on the future, especially if the consequences are long-term, abstract, and at odds with current behavior and identity (Stoknes, 2016). To contribute to overcoming these challenges, an easily accessible tool is needed to help the public understand - both rationally and viscerally - the consequences of not taking sufficient action against climate change.

*equal contribution

2 OUR PROPOSAL

We propose to develop a Machine Learning (ML) based tool showing in a personalized way the probable effect that climate change will have on a specific location familiar to the viewer. Given an address, it generates an image projecting transformations which are likely to occur there, based on a formal climate model. The hope is that such visualizations would help to visceralize climate change: one might be more willing to take action when seeing the consequences of climate change on their home, their neighbourhood, or the street that they grew up on. The first prototype version of our tool simply generates images of flooded locations based on a binary random variable from a climate model of whether flooding will be present at a given place within a static time frame (in 2050). Eventually, we will extend our model to incorporate other climate-related events (fires and droughts, etc.), varying time horizons, and ‘decision knobs’ allowing the viewer to choose actions and make decisions and see their impact on the projected consequences of climate change.

In our prototype, we are able to generate images of the projected impact of flooding by training a CycleGAN network (Zhu et al., 2017) on Google Street View images of both flooded and unflooded streets and houses (Anguelov et al., 2010). The advantage of using the CycleGAN model is that paired one-to-one mapping is not necessary (i.e. we do not have to have images of the same house before and after a flood). Instead, the model uses domain-level mapping in order to learn the transformation necessary to transform a non-flooded house into a flooded one. We present our approach in more detail in Section 4.

3 RELATED WORK

Climate change requires solutions to several urgent problems facing humanity and the planet. Since climate sciences have entered the era of big data, ML - which has been widely successful in several domains - has brought forward immense potential to contribute to problems in climate sciences. However such applications introduce new challenges for ML due to unique climate physics properties encountered in each problem, requiring novel research in ML. Nonetheless, there are several cross-cutting research themes in problems such as super-resolution, classification, climate down-scaling, forecasting, emulating simulations, localization, detection and tracking of extreme events or anomalies, that are applicable across climate science and ML problems, which requires deep collaboration for synergistic advancements in both disciplines (Monteleoni et al., 2013; Joppa, 2017; Racah et al., 2017; Schneider et al., 2017; Gil et al., 2018; Hwang et al., 2018; Karpatne et al., 2018; Rasp et al., 2018).

Furthermore, ML can help bridge the gap between numerical physics and personalized predictions by improving the accuracy of the physics models. For instance, applying ML techniques along with a physics-guided understanding of meteorology and climate has been shown to significantly improve the prediction of high-impact events (Karpatne et al., 2017; Rupe et al., 2017). Also, ML techniques can extract otherwise unavailable information from climate forecasts by fusing model output with observations to provide additional decision support for forecasters and users (McGovern et al., 2017). Finally, climate science-motivated discovery could lead to advances in ML, as demonstrated in the application of deep learning methods for pixel-level segmentation of extreme events by Kurth et al. (2018).

In this work, we use CycleGANs (Zhu et al., 2017) to depict photo-realistic visuals of the potential effects of climate-change events on individual houses and streets. While other approaches to visualize climate change have used both selecting specific images that best represent climate change impacts (Sheppard, 2012; Corner & Clarke, 2016) as well as using artistic renderings of possible future landscapes (Giannachi, 2012) and even video simulations of flooded streets due to rising water levels (Gianatasio, 2014), to our knowledge, our project is the first application of generative models for the specific purpose of generating images of future climate change impact.

4 CONTRIBUTIONS

While the final version of our visualization tool will include various climate events and incorporate different types of metrics from the climate model, for the initial version of our GAN model, we

focused on generating images of houses and buildings specifically after flooding events. In this section, we present the data collection and training approach used for our model.

4.1 FLOODING IMAGE DATASET

One of the most challenging aspects of generating realistic images using GANs is collecting the training data needed in order to extract the mapping function. CycleGAN training assumes that there is some underlying relationship between the domain - for instance a change of seasons in a landscape - which is why we collected many images of streets and houses before and after flooding, with as few extraneous objects (such as vehicles or people) as possible. To collect the necessary training data, we manually searched open source photo-sharing websites for images of houses from various neighborhoods and settings, such as suburban detached houses, urban townhouses, and apartment buildings. We gathered over 500 images of non-flooded houses and the same number of flooded locations, which were all re-sized to 300x300 pixels.

In order to increase the quantity of images that we could use for training, we performed several data augmentation techniques such as: random crops of a subset of each image, horizontal flipping, small rotations, etc., which enabled us to increase our data set five-fold to over 5000 images total. However, a challenge that we encountered was the fact that flooding is not truly a one-to-one mapping such as the one assumed by the CycleGAN approach, but in fact a many-to-one mapping, i.e. roads, grass, dirt, fences are all mapped to water. For this reason, our data collection was constrained to houses surrounded by lawns, which were then mapped to water by the model.

4.2 MODEL ARCHITECTURE AND TRAINING

We use the same architecture for our generative network as that used in the original CycleGAN paper (He et al., 2016). We trained the networks using the publicly available PyTorch (Paszke et al., 2017) implementation¹. The unique aspect of the CycleGAN approach is the cycle consistency loss, which is used along with the traditional adversarial loss to reduce the space of possible domain-to-domain mapping functions by ensuring that for each image x from domain \mathcal{X} , the image translation cycle should be able to bring x back to the original image (and vice-versa for a given image y from the other domain, \mathcal{Y}). We trained our CycleGAN model for 200 epochs on the training images, using the Adam solver (Kingma & Ba, 2015) with a batch size of 1, training the model from scratch with a learning rate of 0.0002. As per the CycleGAN training procedure, the learning rate is constant for the first 100 epochs and linearly decayed to zero over the next 100 epochs. We present some of our results below.

4.3 RESULTS

As can be seen in Figure 1, our CycleGAN model was able to learn an adequate mapping between grass and water, and this mapping could be applied to generate fairly realistic images of flooded houses. The mapping works best with single-family, suburban-type houses which are surrounded by an expanse of grass. There are still improvements to be made with regards to the color scheme of the generated images and the visual artifacts that remain, as well as the coverage of more types of buildings and houses. From the 80 images in the test set, we found that about 70% were successfully mapped to realistically flooded houses (see 5 for more information about image evaluation).

The information about whether or not a house is flooded at specific locations for the CycleGAN images above is sourced from climate model flood hazard outputs, which were converted to binary global flood maps. First, for inland lakes and rivers, a binary flood hazard map was based on each of the 10, 20, 50 and 100 year return runs globally at 1km resolution. We show a 50 year return run in Figure 2 using data from Dottori et al. (2016). Secondly, probabilistic projection data of Extreme Sea Levels until the end of the 21st century along the global coastline was extracted (for the 50th quantile) from Vousdoukas et al. (2018) in year 2050 to create a second binary map. This second binary map is based on projections of the decade-window, 2050, under a representation concentration pathway of 4.5 °C global warming scenario with a greater than 20 cm sea level increase exceedance threshold compared to a baseline sea level from 1980-2014.

¹<https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix>



Figure 1: Images of flooded houses generated by our model^a

^aNote some artifacts in the sky of the second image, most likely due to trees or clouds in the flooding images used for training

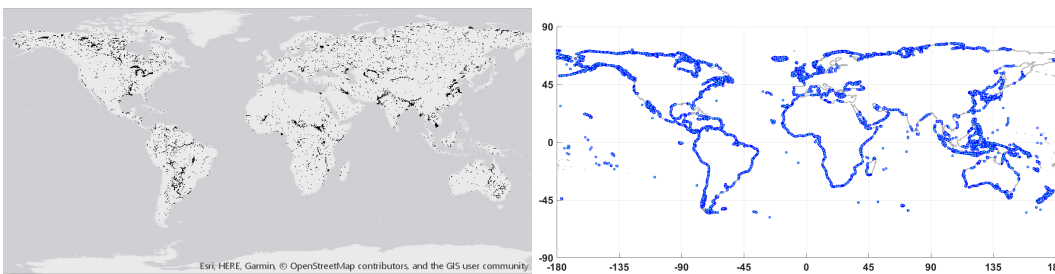


Figure 2: Binarised maps of global inland lake and river flooding 50-year projection based on Dottori et al. (2016) and coastal flood hazard maps for year 2050 based on Vousdoukas et al. (2018) ^a

^aFrom left to right: flooding of lakes and rivers inland binary map; binary map of coastal flooding for RCP 4.5 °C and higher than 20cm rise w.r.t baseline; Colors: black = inland flood hazard, blue = coast flood hazard

5 DISCUSSION AND FUTURE DIRECTIONS

The initial version of the CycleGAN model that we have developed in the present paper is a prototype to illustrate the feasibility of applying generative models to create personalized images of an extreme climate event, flooding, that is expected to increase in frequency based on climate change projections. Subsequent versions of our model will integrate more varied types of houses and surroundings, as well as different types of climate-change related extreme event phenomena (i.e. droughts, hurricanes, wildfires, air pollution etc), depending on the expected impacts at a given location, as well as forecast time horizons.

Furthermore, to channel the emotional response into behavioural change or actions, another important planned improvement to our model is the eventual addition of ‘choice knobs’, to enable users to visually see the impact of their personal choices, such as deciding to use more public transportation, as well as the impact of broader policy decisions, such as carbon tax and increasing renewable portfolio standards. The effects of an individual turning these knobs could be based on the best available climate model projections, such as the one used for our binary flood map (Dottori et al., 2016), integrated with economic and policy assessment models. Ultimately, by integrating these ‘knobs’ into our system, we aim to help the general population progress towards greater and more visible public support for climate change mitigation steps on a national level, facilitating governmental interventions and helping make the required rapid changes to a global sustainable economy.

However, there are several challenges which we are currently facing that require gaps to be bridged in research at the intersection of climate science and ML. Current climate models make projections based on the physics of fluid motion, energy transfer, mass conservation or chemical transport, not using Deep Learning approaches. Furthermore, the spatial resolutions of these physics models are at best regional, which is much coarser than individual households investigated in this problem.

Moreover, the model outputs provide physics variables that are non-trivial to translate into equivalent photo-realistic representations. We therefore believe that there a need to explore physical constraints to GAN training in order to incorporate more physical knowledge into these projections. This is important so that a GAN model will not only transform a house to its projected flooded state, but also take into account the forecast simulations of the flooding event represented by the physical variable outputs and probabilistic scenarios by a climate model for a given location.

REFERENCES

- Dragomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stéphane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, and Josh Weaver. Google street view: Capturing the world at street level. *Computer*, 43(6):32–38, 2010.
- Adam Corner and Jamie Clarke. *Talking climate: From research to practice in public engagement*. Springer, 2016.
- Francesco Dottori, Peter Salamon, Alessandra Bianchi, Lorenzo Alfieri, Feyera Aga Hirpa, and Luc Feyen. Development and evaluation of a framework for global flood hazard mapping. *Advances in water resources*, 94:87–102, 2016.
- David Gianatasio. ‘world under water uses streetview to visualize flooding from climate change. *Adweek*, 2014.
- Gabriella Giannachi. Representing, performing and mitigating climate change in contemporary art practice. *Leonardo*, 45(2):124–131, 2012.
- Yolanda Gil, Suzanne A Pierce, Hassan Babaie, Arindam Banerjee, Kirk Borne, Gary Bust, Michelle Cheatham, Imme Ebert-Uphoff, Carla Gomes, Mary Hill, et al. Intelligent systems for geosciences: an essential research agenda. *Communications of the ACM*, 62(1):76–84, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Jessica Hwang, Paulo Orenstein, Karl Pfeiffer, Judah Cohen, and Lester Mackey. Improving sub-seasonal forecasting in the western us with machine learning. *arXiv preprint arXiv:1809.07394*, 2018.
- IPCC. *Global Warming of 1.5 C: An IPCC Special Report on the Impacts of Global Warming of 1.5 C Above Pre-industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty*. Intergovernmental Panel on Climate Change, 2018.
- Lucas N Joppa. The case for technology investments in the environment, 2017.
- Anuj Karpatne, Gowtham Atluri, James H Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331, 2017.
- Anuj Karpatne, Imme Ebert-Uphoff, Sai Ravela, Hassan Ali Babaie, and Vipin Kumar. Machine learning for the geosciences: Challenges and opportunities. *IEEE Transactions on Knowledge and Data Engineering*, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- Thorsten Kurth, Sean Treichler, Joshua Romero, Mayur Mudigonda, Nathan Luehr, Everett Phillips, Ankur Mahesh, Michael Matheson, Jack Deslippe, Massimiliano Fatica, et al. Exascale deep learning for climate analytics. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, pp. 51. IEEE Press, 2018.
- Amy McGovern, Kimberly L Elmore, David John Gagne, Sue Ellen Haupt, Christopher D Karstens, Ryan Lagerquist, Travis Smith, and John K Williams. Using artificial intelligence to improve real-time decision-making for high-impact weather. *Bulletin of the American Meteorological Society*, 98(10):2073–2090, 2017.
- Claire Monteleoni, Gavin A Schmidt, and Scott McQuade. Climate informatics: accelerating discovering in climate science with machine learning. *Computing in Science & Engineering*, 15(5): 32–40, 2013.

- Saffron J O'Neill and Mike Hulme. An iconic approach for representing climate change. *Global Environmental Change*, 19(4):402–410, 2009.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. *NIPS 2017 Workshop Autodiff Submission*, 2017.
- Nick Pidgeon. Public understanding of, and attitudes to, climate change: Uk and international perspectives and policy. *Climate Policy*, 12(sup01):S85–S106, 2012.
- 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.
- Stephan Rasp, Michael S Pritchard, and Pierre Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39):9684–9689, 2018.
- Adam Rupe, James P Crutchfield, Karthik Kashinath, et al. A physics-based approach to unsupervised discovery of coherent structures in spatiotemporal systems. *arXiv preprint arXiv:1709.03184*, 2017.
- Tapio Schneider, Shiwei Lan, Andrew Stuart, and João Teixeira. Earth system modeling 2.0: A blueprint for models that learn from observations and targeted high-resolution simulations. *Geophysical Research Letters*, 44(24), 2017.
- Stephen RJ Sheppard. *Visualizing climate change: a guide to visual communication of climate change and developing local solutions*. Routledge, 2012.
- Christopher J Smith, Piers M Forster, Myles Allen, Jan Fuglestedt, Richard J Millar, Joeri Rogelj, and Kirsten Zickfeld. Current fossil fuel infrastructure does not yet commit us to 1.5 c warming. *Nature communications*, 10(1):101, 2019.
- Espen Stoknes. Why the human brain ignores climate change - and what to do about it. In *Environmental Reality: Rethinking the Options*, pp. 75–81. Swedish Royal Colloquium 2016, 2016. URL <https://files.acrobat.com/a/preview/1ef80b88-177c-4e5d-b879-d6d3a059c694>.
- Michalis I Voudoukas, Lorenzo Mentaschi, Evangelos Voukouvalas, Martin Verlaan, Svetlana Jevrejeva, Luke P Jackson, and Luc Feyen. Global probabilistic projections of extreme sea levels show intensification of coastal flood hazard. *Nature communications*, 9(1):2360, 2018.
- Elke U Weber and Paul C Stern. Public understanding of climate change in the united states. *American Psychologist*, 66(4):315, 2011.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2223–2232, 2017.