A Deep Learning Based Hail Climatology for the Contiguous United States 1979-2018

Stan Biryukov, Eric Hewitt, Jim Krzak, Anthony Marro, and Alex Kubicek
Understory, Inc.
Madison, WI
info@understoryweather.com

ABSTRACT

Hail storms are unique climatological events that are difficult to forecast yet unrelenting in their capacity to cause wide-spread agricultural and property damage. In this study, we aim to elucidate the spatial and temporal trends in hail days over the Continental United States from 1979 to 2018. We leverage the discriminatory power of radar derived hail signatures, combined with ground validated hail sizes, to distinguish daily environments that produce potentially damaging hail. In order to quantify the likelihood of damaging hail for any given day, we develop a neural network model which utilizes ERA5 based reanalysis variables. Our model reveals that annual hail days are significantly increasing since 1979. We find a statistically significant increase in severe hail environments in both frequency and coverage, with 1.4 percent more land mass of the Continental United States experiencing hail every decade.

INTRODUCTION

Hail is one of the leading causes of weather-related property damage in the United States. Over the past decade, billion dollar storm events have been increasing in frequency (1) (2) with hail making up a large climatologies have thus far either focused on modeling storm reports, reviewing trends in radar derived Maximum Estimated Hail Size (MESH), or have been confined to analysis of mesoscale parameters or derived indices. We aimed to directly model daily radar derived hail signatures with atmospheric parameters from a re-analysis dataset, thereby extending the spatial and temporal discriminatory power of radar onto a 40-year climatological scale.

Our methodology differs from previous climatological studies in that we let our model learn a transfer function from a dense matrix of numerous atmospheric features. Rather than limiting our algorithm to a select few variables or computing radar derived hail indices, we framed hail formation as a machine learning classification problem and posited that interactions across environmental variables are more informative to hail formation than any single variable.

METHODS

We pooled all filtered hail signatures from NOAA's Severe Weather Data Inventory (SWDI) from 2005 to 2017. While the radar derived hail signatures provide a rich spatiotemporal profile of potential hail environments, the raw dataset is prone to multiple sources of error. Specifically, false positives may occur when hail melts before it reaches the surface due to a high-freezing level, and overestimation may occur when the melting level is below the lowest radar beam (5). In order to remove as many potential false positives from our dataset, we calibrated the indirect measurement of radar signatures to surface validated hail measurements. We developed a discriminatory Damage Likelihood Index based on the severe probability threshold of all SWDI hail signatures, and their spatiotemporally matched surface measured hail size. We queried all available dates of hail capture from Understory's network of ground based weather sensors (6), and merged over 3,800 overlapping spatio-temporal records with the SWDI data. For each hail storm, we computed the average measured maximum size of hail across all stations reporting hail. Fig 1 shows the relationship between the SWDI severe probability filter and ground measured average maximum size. We found that the severe probability threshold of 100 within the SWDI dataset provided the optimal threshold for detecting potentially ground-level damaging hail - a surface level average measured maximum size of .75 inches.

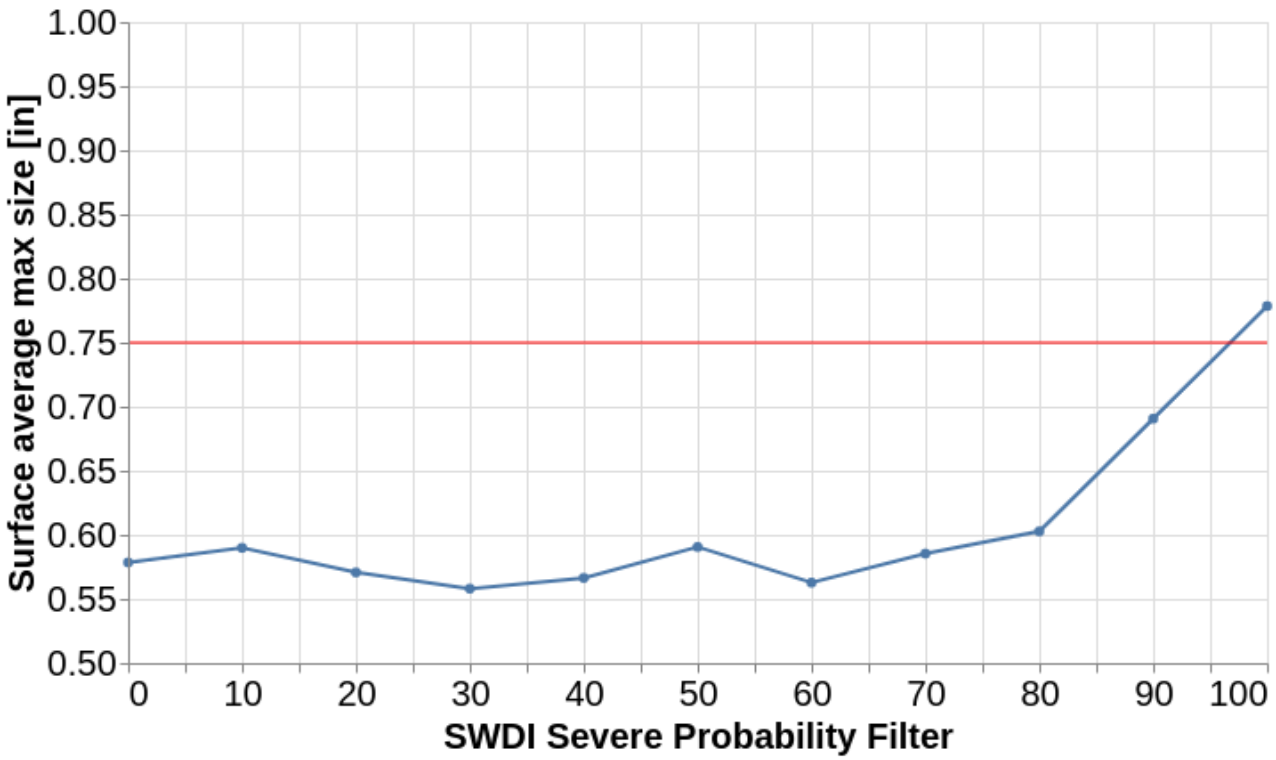

Figure 1: Severe Weather Data Inventory Damage Likelihood Index

While Storm Prediction Center reports are naturally biased towards population centers and reported intensities are quantized into categorical thresholds based on reference objects (7) (8), we validated our choice of radar filter utilizing all spatiotemporal matched storm reports. Notably, Fig 2 reveals reported hail sizes begin a steady increase above a severe probability filter of 60 and exhibit a similar peak at a severe probability filter of 100%. A one-sided t-test confirms the global maximum at 100 distinguishes storms that generate potentially damaging hail.

All resulting SWDI hail signatures with a 100% severe probability of hail were re-gridded onto a standard .25 x .25 decimal degree (~30 km) grid corresponding to the ERA-5 dataset resolution. The maximum MESH size was aggregated for each day to create a daily hail swath. Each grid cell was then transformed into a daily corresponding binomial count of 92,887 unique space-time observed hail days.

To elucidate relationships between atmospheric environments and conditions favorable to hail formation, we pooled surface-level data from ERA5, a methodologically consistent global weather and climate reanalysis dataset (9). Our covariates of interest included freezing-degree height, K-index, total column ice water content, convective available potential energy, and convective inhibition, as these are parameters often used to forecast hailstorms (10) (11). We computed the extreme values of each variable for the day and merged the data onto the corresponding CONUS hail grid.

Discovering and including interaction effects in modelling a response variable of interest is a fundamental problem across disciplines (12). Numerical Weather Prediction models have advanced largely by incorporating interactions across large scale domains (13). In addition to capturing informative interactions, avoiding inference from extraneous variables and collinear terms helps generalize a model's performance; it is just as important to avoid overfitting parameters to an outcome of interest such that spurious correlations are not drawn. In practice, computing the relevance of predictors along with their interaction pairs in algorithm development can be a computationally challenging task (14); variables are often included based on study-specific approximated thresholds. In order to scale our optimization problem, we leveraged the stochastic gradient optimization advantages of neural networks. We designed a regularized multi-layered feed-forward neural network architecture, with a

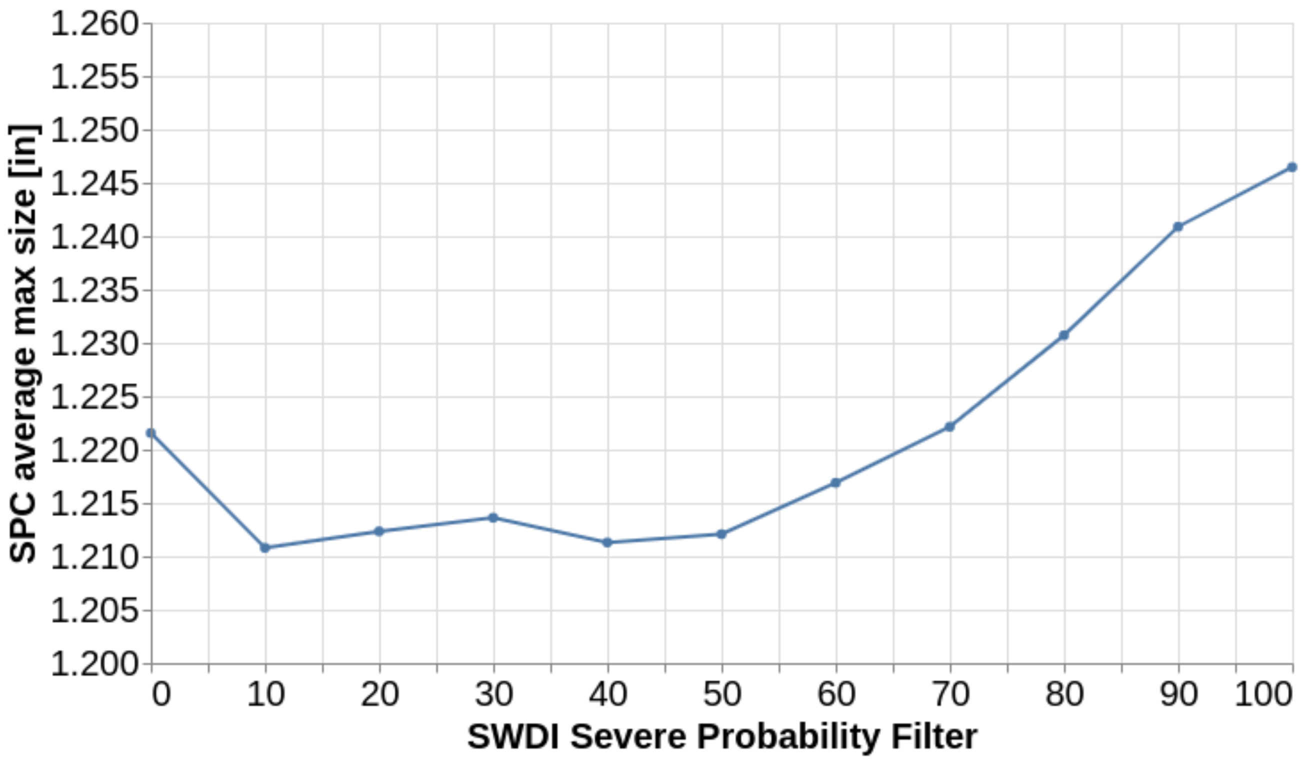

Figure 2: SPC-dervied Severe Weather Data Inventory Damage Likelihood Index

logistic output layer, allowing our model to exploit a fully dense matrix containing all covariates up to their third-degree polynomial feature interactions.

Our tuned neural network utilized stochastic gradient optimization, with L_2 regularization (15), enabling training with a total of 119 features, and over 1 million observations. A decaying learning rate helped guide the optimization procedure to a minimum while avoiding oscillation along the error surface (16). Five K-Fold hyperparameter cross-validation was performed on the optimizer, activation functions, number of hidden neuron layers and feature-interaction polynomial terms. In reviewing some of the learned relationships, we find the activation matrix of the neural network has a bimodal shape centered about zero, indicating regularization primed the 256-layer network to focus on the most relevant features. Model cross-validation reveals our neural network outperforms a standard logistic regression model, and thus infers latent effects that correlate with hail formation from daily tabulated atmospheric variables. The resulting model provides a spatially and temporally resolute posterior likelihood of hail over the Continental United States for any given day.

RESULTS

Our neural network exhibited strong discrimination between hail-days and non-hail days on a random validation set of 100,000 space-time grid points, with an area under the receiver operating curve of .94. Fig 3 shows the corresponding receiver operating curve.

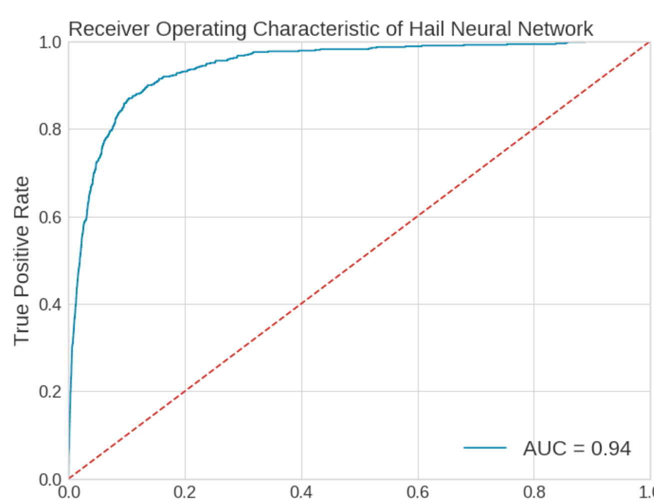

Figure 3: Receiver Operating Characteristic Curves for Hail Neural Network

We find hail-days increasing at a statistically significant rate of 1.1 percent per year over all of the Continental United States, or 11.6 percent per decade. Fig 4 displays the annual change in hail-days by State, which shows Kansas leading the annual increase of hail days with 54 new hail-days per year, followed by Nebraska, South Dakota, Oklahoma, and North Dakota at 41, 36, 31, and 23 additional hail-days, respectively.

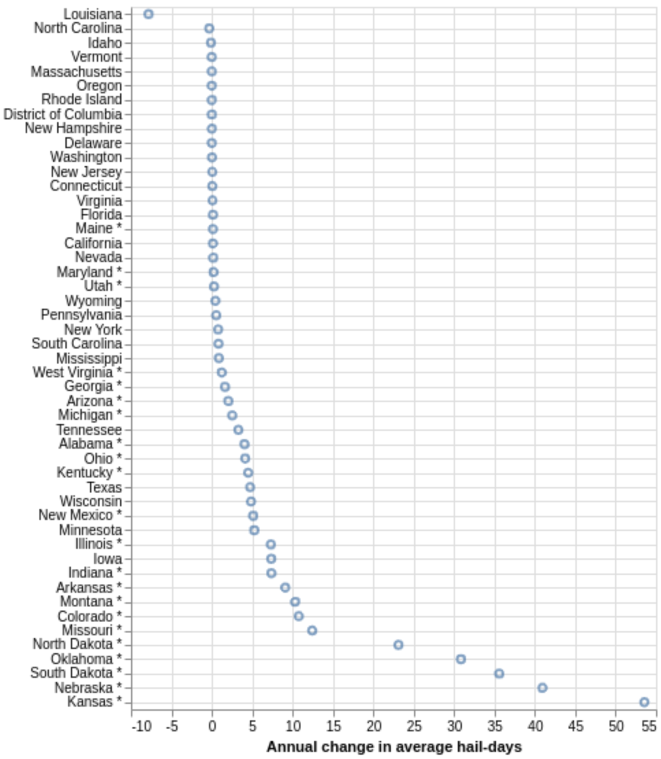

Figure 4: Annual change in average hail-days by state. Asterisk indicates significant at 95% confidence interval.

Our model captures a significant Westward and Eastern regional expansion of hail-days, as shown in Fig 6. In reviewing atmospheric and surface level trends, we find a notable increase of convective parameters that correlate with our observed upward trend of hail days. Convective inhibition appears to be decreasing in the West, while a steady annual increase in average CAPE may concomitantly explain some of the increased hail days, shown in Figs 7, and 8, respectively. Particularly noteworthy is the dramatic annual increase in CAPE in the Southwest and Northeast. This is in line with other studies that have observed an increase in environmental conditions suitable for hail formation over the United States (17).

While we find hail-days are expanding from the originally confined 'Hail-Alley' stretch from the Texas Panhandle up into the Dakotas, there is no evidence that frequency of hail is decreasing there - in fact, they remain areas with the largest annual increase in hail-days. Instead, it appears synoptic conditions are evolving to additionally favor hail formation in larger swaths of the country, with .14 percent, or 4,300 square miles, more of the Continental United States experiencing hail each year. To quantify the increased societal impact, we merged gridded global population numbers (18) and computed the proportion of the 2010 CONUS population that experienced any hail event during the year. Relative to the 2010 population, we find 1.5 percent more of the population experiences hail per decade, demonstrating how hail environments are encroaching on a greater share of the population.

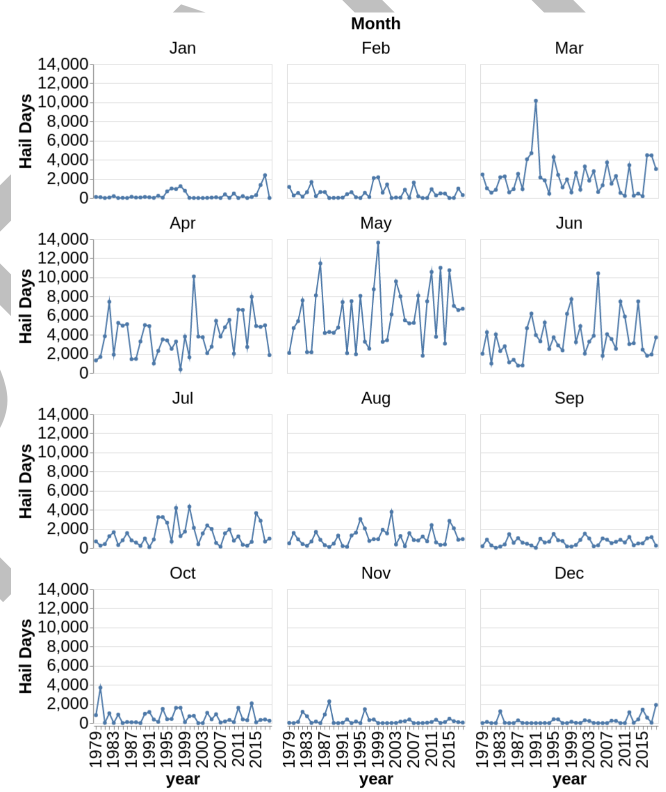

Figure 5: Annual trend in total hail-days by month. Asterisk indicates significant at 95% confidence interval.

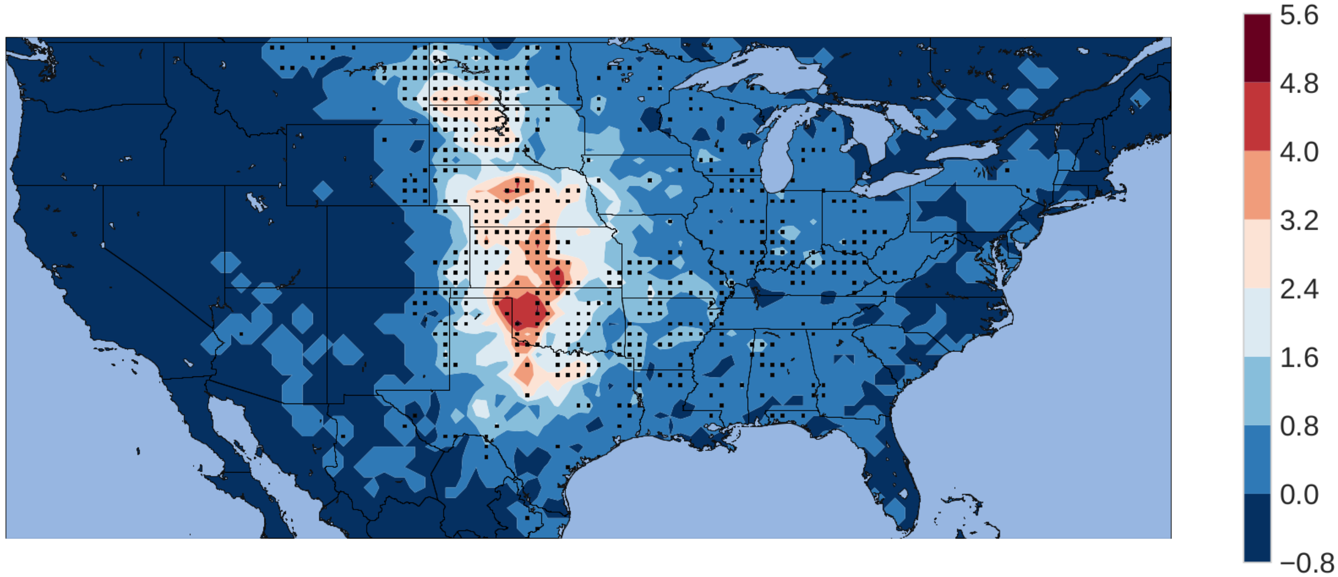

Figure 7: Change in Hail Days per decade. Shaded areas indicate significance at the 95% confidence interval.

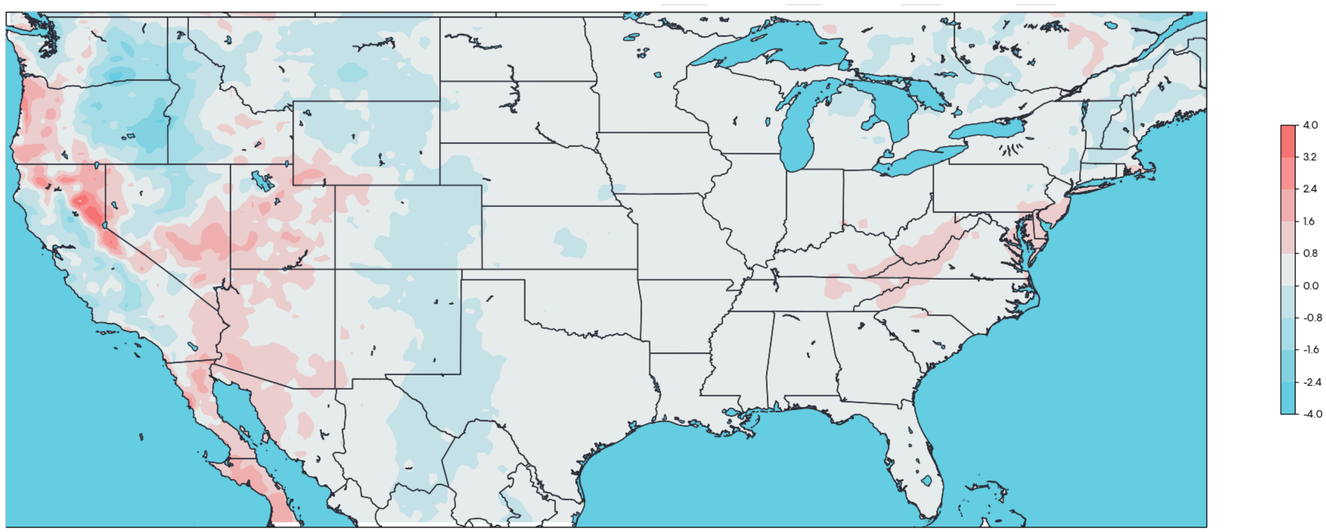

Figure 6: Annual percent change of 10 year moving average: Convective Inhibition

DISCUSSION

By directly modeling radar detected hail days, we are able to resolve a climatology that is consistent with other research and reveals a new phenomenon of hail-days spreading outside of the central US core hail region. Given our main source of inputs were NEXRAD hail signatures and ERA5 covariates, one must be mindful of the intrinsic biases of the inputs. Some radar stations may be limited to specific azimuthal scan strategies and may not pick up hail as regularly as other locations (19). Nonetheless, radar derived signatures of hail, calibrated with surface-level observations, allow for a much more spatially and temporally resolute model parameterization, as opposed to modeling conditions associated with storm reports. Likewise, utilizing daily local environments, as opposed to monthly or annual temporal averages, allows for a more realistic mapping and accumulation of hail events, as the inherent stochasticity is preserved, enabling simulation-based risk forecasting.

The background convective environments are based off of ERA5 model parameterizations and are also subject to those biases. Nonetheless, previous studies have shown strong correlation between ERA5 convective parameters and radiosonde data (11). ERA5 contains marked improvements over ERA-Interim, from integrating new input observations and expanding variational bias schemes, to drastically increasing the

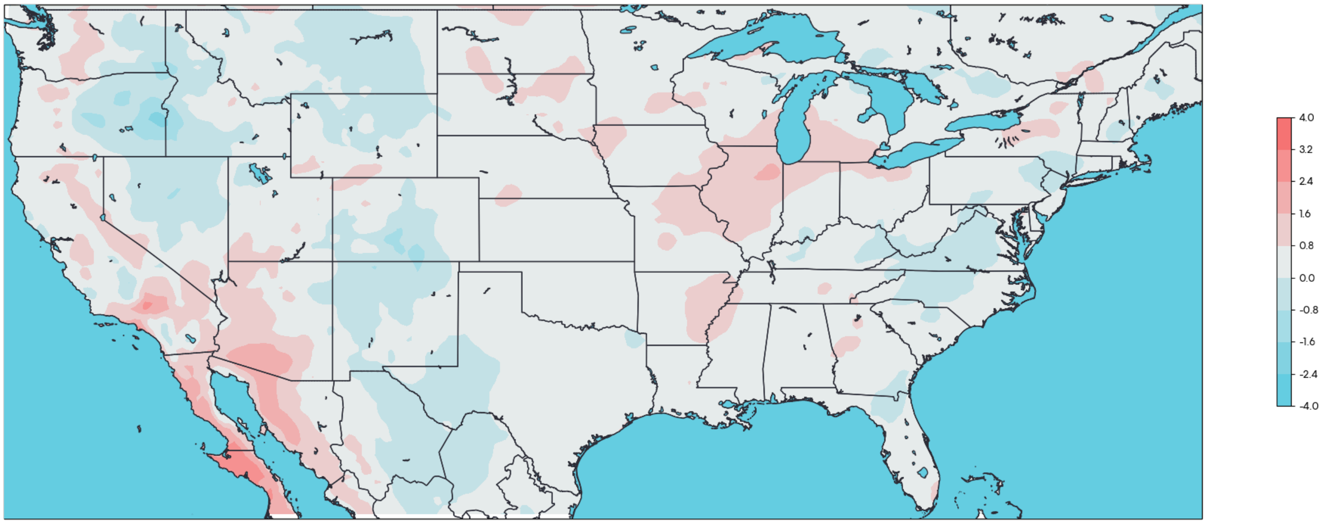

Figure 8: Annual percent change of 10 year moving average: CAPE

spatial resolution of model outputs. While the spatial resolution has improved, we find our model underestimates hail in the Colorado Front-Range. As with other studies, there appears to be a bias in re-analysis models over the Rocky Mountain Region such that orographic effects are not adequately captured right off of the mountains (17), (20). However, it appears ERA5 is able to resolve some of those effects in the South Dakota Black Hills as instability and low-level moisture convolves with the lee troughing off of the Rockies. Overall, there is strong agreement between our pooled daily modeled hail swaths and radar derived MESH, consistent with radar based climatologies (19). By utilizing ERA5 variables to model daily hail swaths, we are able to infer statistically significant time trends in hail expansion that would not be possible at the short-term scale of radar availability.

CONCLUSION

Our study reveals that radar derived severe hail days are increasing in frequency since 1979, in part, due to changing environmental conditions that favor hail formation. Hail-days are both increasingly manifesting outside of the regular 'Hail Alley' (21) and becoming more common outside of the typical March-June season. The Hail Alley of CONUS is extending eastward into the Appalachians, and southwestward into the Phoenix metro area. This expansion into large population centers, if trends hold, will almost assuredly cause property losses to increase nonlinearly. As new population centers begin to experience hail events at more regular intervals, responding to and mitigating hail damage will become ever more important.

The trends of increasing hail days in this study admittedly could have other factors contributing to it - namely due in part to natural variability of decadal oscillations of contributing meteorological factors, and in part due to a connection to climate change (22). Fully accepting that a small sample size of data lends to making presumptuous assumptions over the long term, we nonetheless find notable trends in space and time over a forty year period, in agreement with other studies (23).

Further work will include extending our model to the globe and expanding our training dataset to include additional pressure-level covariates. A deeper dive into hail growth specifics such as focusing on the -10C to -30C layer CAPE, and substituting MUCAPE values over surface based CAPE would be things to consider implementing to fine tune results (24). Involving satellite data and its relation to severe weather such as identifying overshooting top instances and their relationship to very strong updrafts could narrow the focus to severe hail cases (25). With the advent of real-time hail sensing platforms providing a rich stream of hail distribution size data (6), we aim to further extend our analysis by modeling ground-truth hail size with re-analysis variables.

ACKNOWLEDGEMENT

NOAA Severe Weather Data Inventory, a joint project of NOAA's National Climatic Data Center, UNCA's National Environmental Modeling and Analysis Center, and the Renaissance Computing Institute.

ERA5 Reanalysis, a project of ECMWF, the European Centre for Medium-Range Weather Forecasts.

REFERENCES

1. Smith AB, Matthews JL. Quantifying uncertainty and variable sensitivity within the US billion-dollar weather and climate disaster cost estimates. Nat Hazards. 2015 Jul;77(3):1829–51.
2. Smith AB, Katz RW. US billion-dollar weather and climate disasters: data sources, trends, accuracy and biases. Nat Hazards. 2013 Jun;67(2):387–410.
3. Changnon SA. Increasing major hail losses in the U.S. Clim Change. 2009 Sep 1;96(1):161–6.
4. Barthel F, Neumayer E. A Trend Analysis of Normalized Insured Damage from Natural Disasters [Internet]. Rochester, NY: Social Science Research Network; 2011 May [cited 2020 Apr 2]. Report No.: ID 1831633. Available from: https://papers.ssrn.com/abstract=1831633
5. Schlie EE-J, Wuebbles D, Stevens S, Trapp R, Jewett B. A radar-based study of severe hail outbreaks over the contiguous United States for 2000–2011. Int J Climatol. 2019;39(1):278–91.
6. Willmot E, Bussmann S, Homeier N, Kubicek A. Examples of ground-based hail, rain, and wind sensor networks operating in real-time. In: 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS). 2017. p. 5987–9.
7. Allen JT, Tippett MK, Sobel AH. An empirical model relating U.S. monthly hail occurrence to large-scale meteorological environment. J Adv Model Earth Syst. 2015;7(1):226–43.
8. Schaefer J, Levit J, Weiss S, Mccarthy D. The frequency of large hail over the contiguous United States. Bull Am Meteorol Soc. 2004 Jan 1;
9. Forecasts EC for M-RW. ERA5 Reanalysis [Internet]. UCAR/NCAR - Research Data Archive; 2017 [cited 2020 Apr 2]. Available from: https://rda.ucar.edu/datasets/ds630.0/
10. Skaggs RH. The Diuranal Distribution of Large Hail Over the United States. Prof Geogr. 1980;32(3):293–9.
11. Ukkonen P, Mäkelä A. Evaluation of Machine Learning Classifiers for Predicting Deep Convection. J Adv Model Earth Syst. 2019;11(6):1784–802.
12. Koo CL, Liew MJ, Mohamad MS, Mohamed Salleh AH. A Review for Detecting Gene-Gene Interactions Using Machine Learning Methods in Genetic Epidemiology. BioMed Res Int [Internet]. 2013 [cited 2020 Feb 21];2013. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3818807/
13. Bauer P, Thorpe A, Brunet G. The quiet revolution of numerical weather prediction. Nature. 2015 Sep;525(7567):47–55.
14. Prein AF, Holland GJ. Global estimates of damaging hail hazard. Weather Clim Extrem. 2018 Dec 1;22:10–23.
15. PedregosaFabian, VaroquauxGaël, GramfortAlexandre, MichelVincent, ThirionBertrand, GriselOlivier, et al. Scikit-learn: Machine Learning in Python. J Mach Learn Res [Internet]. 2011 Nov 1 [cited 2020 Apr 2]; Available from: https://dl.acm.org/doi/abs/10.5555/1953048.2078195
16. Kanter I, Lecun Y, Solla S. Second-order properties of error surfaces: learning time and generalization. In 1991.
17. Tang BH, Gensini VA, Homeyer CR. Trends in United States large hail environments and observations. Npj Clim Atmospheric Sci. 2019 Dec 6;2(1):1–7.
18. Gridded Population of the World (GPW), v4 | SEDAC [Internet]. [cited 2020 Feb 21]. Available from: https://sedac.ciesin.columbia.edu/data/collection/gpw-v4
19. Cintineo JL, Smith TM, Lakshmanan V, Brooks HE, Ortega KL. An Objective High-Resolution Hail Climatology of the Contiguous United States. Weather Forecast. 2012 May 16;27(5):1235–48.
20. Barratt M. Impact of orographic drag on forecast skill [Internet]. ECMWF. 2017 [cited 2020 Apr 2]. Available from: https://www.ecmwf.int/en/newsletter/150/meteorology/impact-orographic-drag-forecast-skill
21. Dewey KF. Hail-Producing Storms on the Great Plains. Weatherwise. 2019 Sep 3;72(5):30–5.
22. Kravtsov S, Grimm C, Gu S. Global-scale multidecadal variability missing in state-of-the-art climate models. Npj Clim Atmospheric Sci. 2018 Nov 20;1(1):1–10.
23. Brimelow JC, Burrows WR, Hanesiak JM. The changing hail threat over North America in response to anthropogenic climate change. Nat Clim Change. 2017 Jul;7(7):516–22.
24. Craven JP, Brooks H. Baseline climatology of sounding derived parameters associated with deep, moist convection. Natl Wea Dig. 2004 Jan 1;28:13–24.
25. Punge HJ, Bedka KM, Kunz M, Reinbold A. Hail frequency estimation across Europe based on a combination of overshooting top detections and the ERA-INTERIM reanalysis. Atmospheric Res. 2017 Dec 1;198:34–43.

SUPPLEMENT

Supplement 1: Monthly average Hail-Days: January

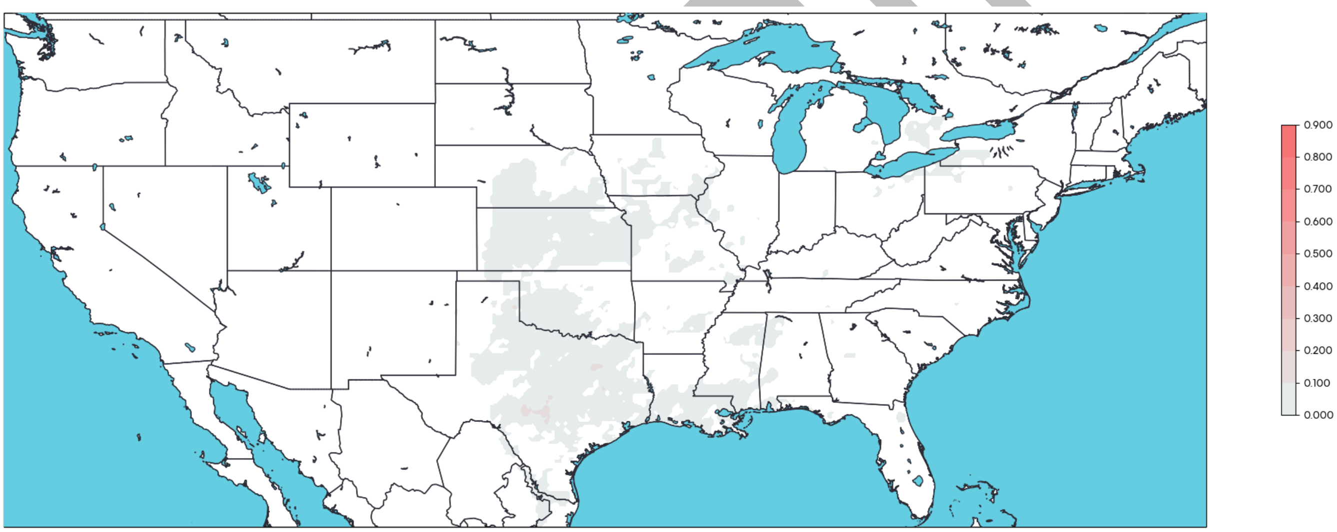

Supplement 2: Monthly average Hail-Days: February

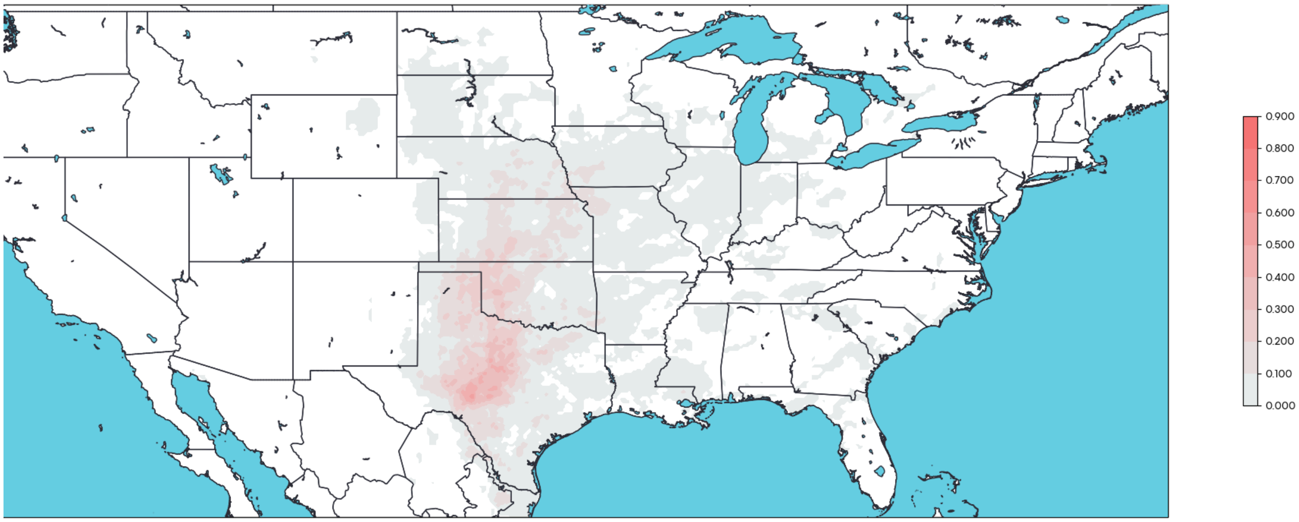

Supplement 3: Monthly average Hail-Days: March

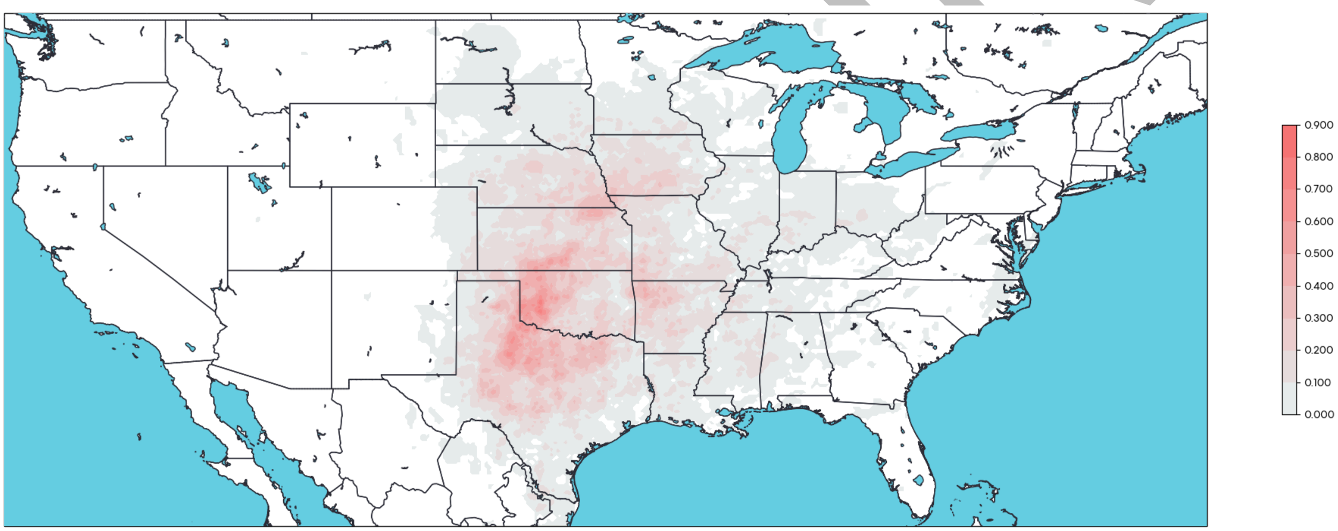

Supplement 4: Monthly average Hail-Days: April

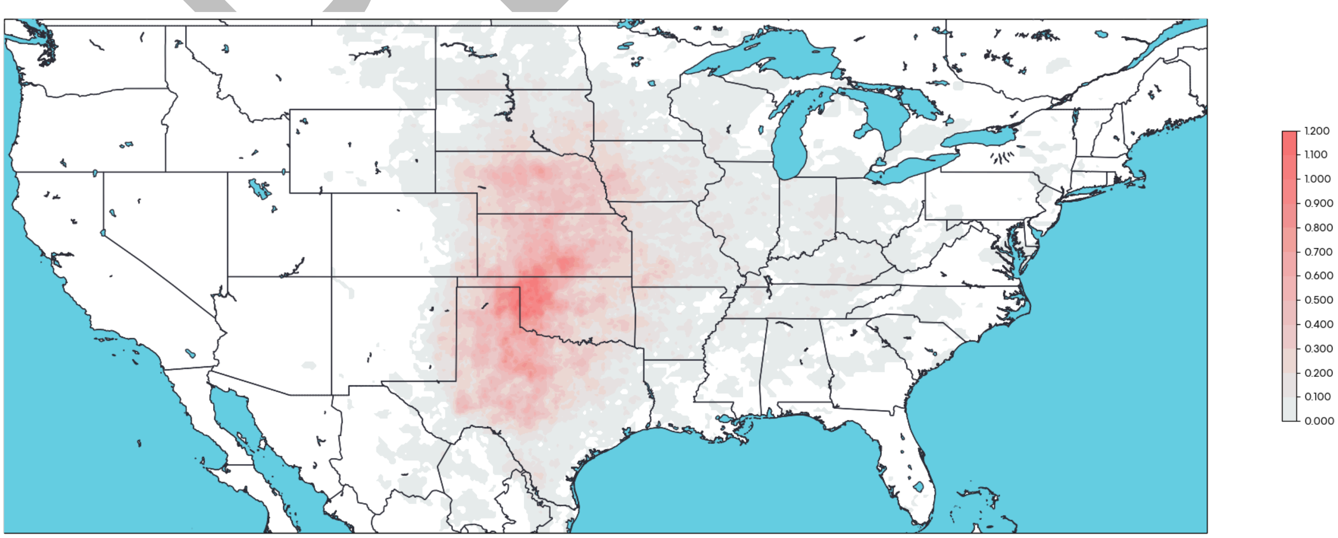

Supplement 5: Monthly average Hail-Days: May

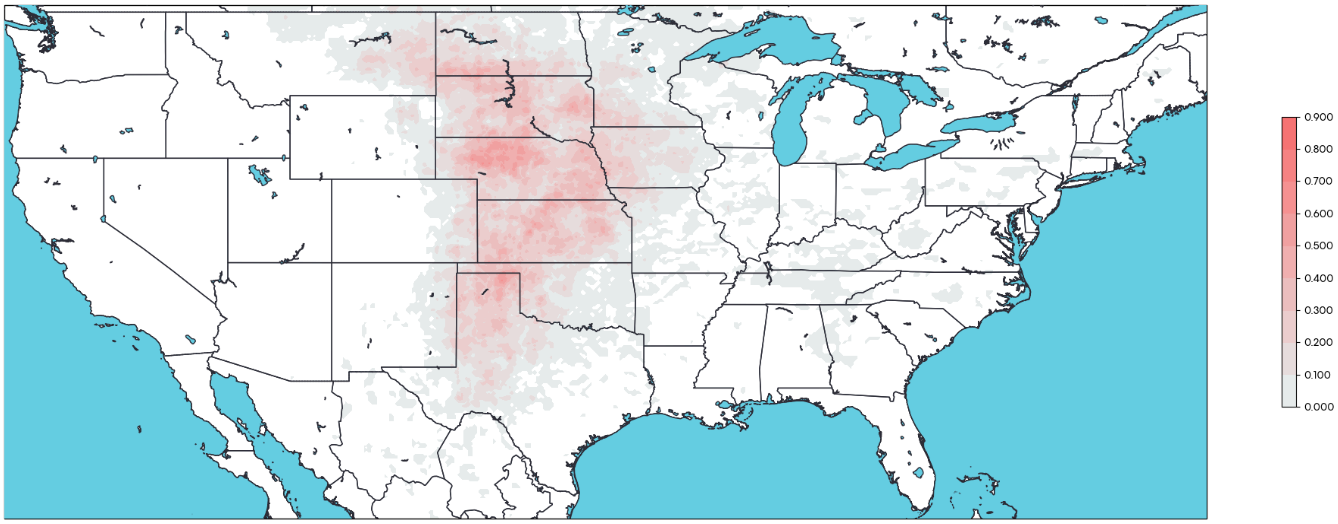

Supplement 6: Monthly average Hail-Days: June

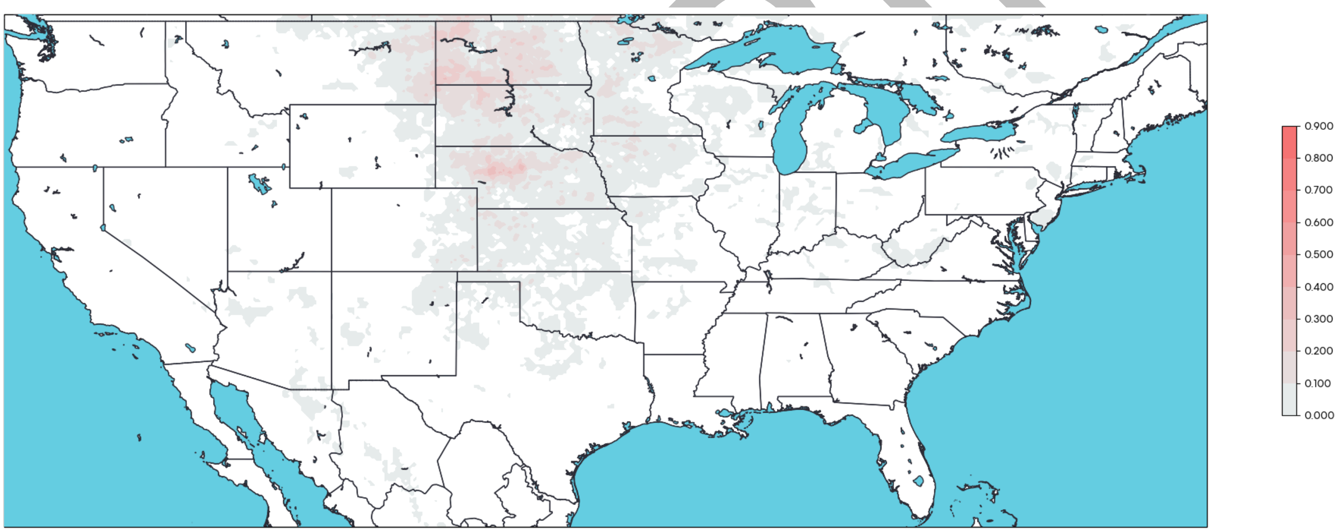

Supplement 7: Monthly average Hail-Days: July

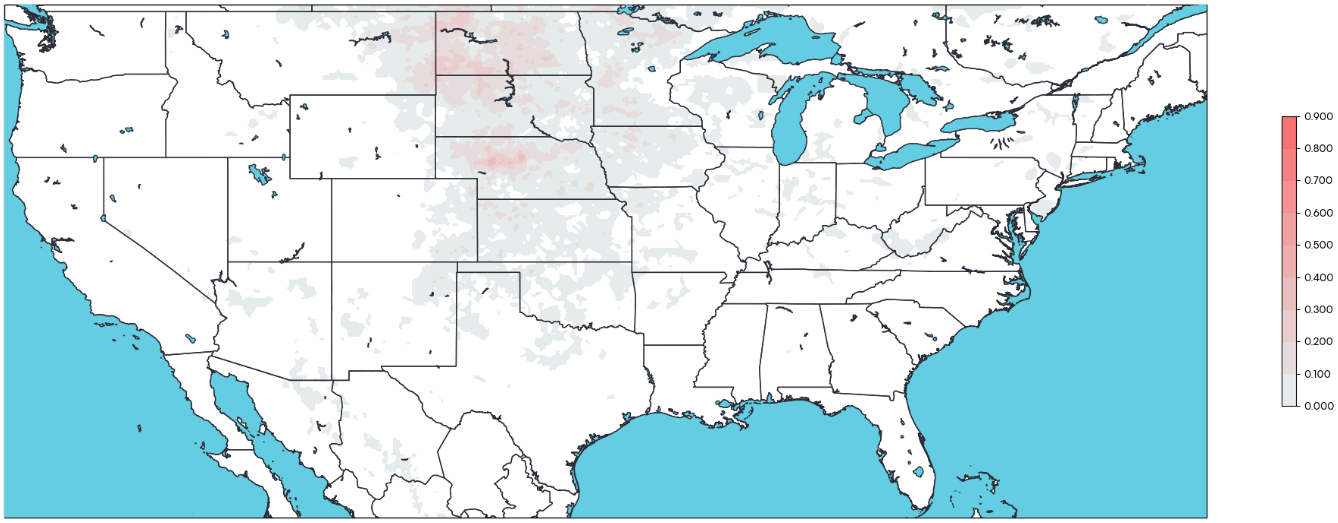

Supplement 8: Monthly average Hail-Days: August

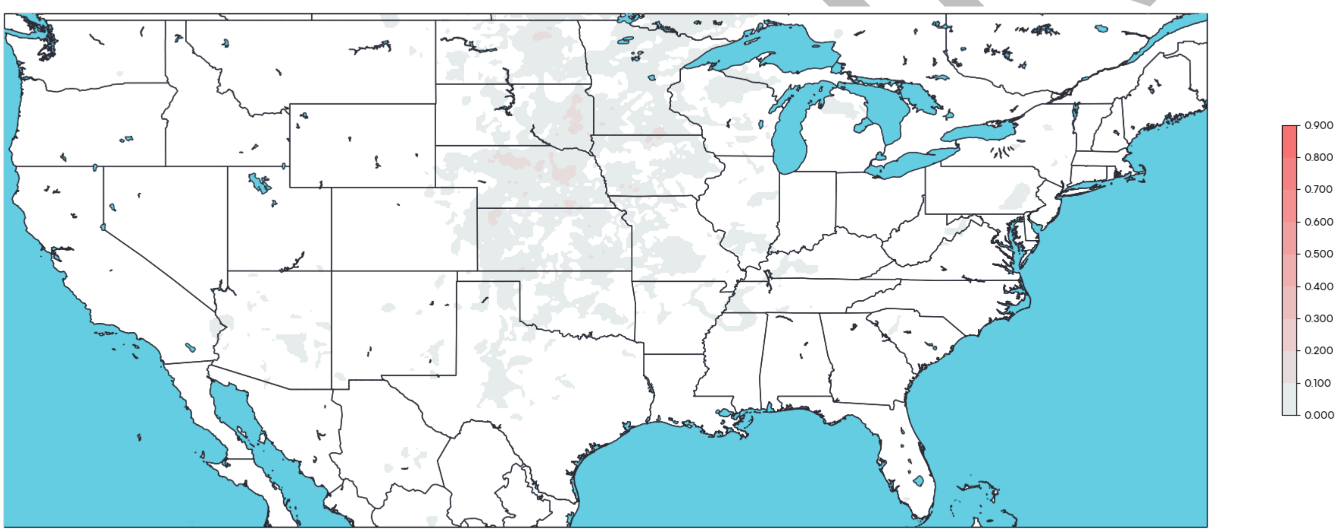

Supplement 9: Monthly average Hail-Days: September

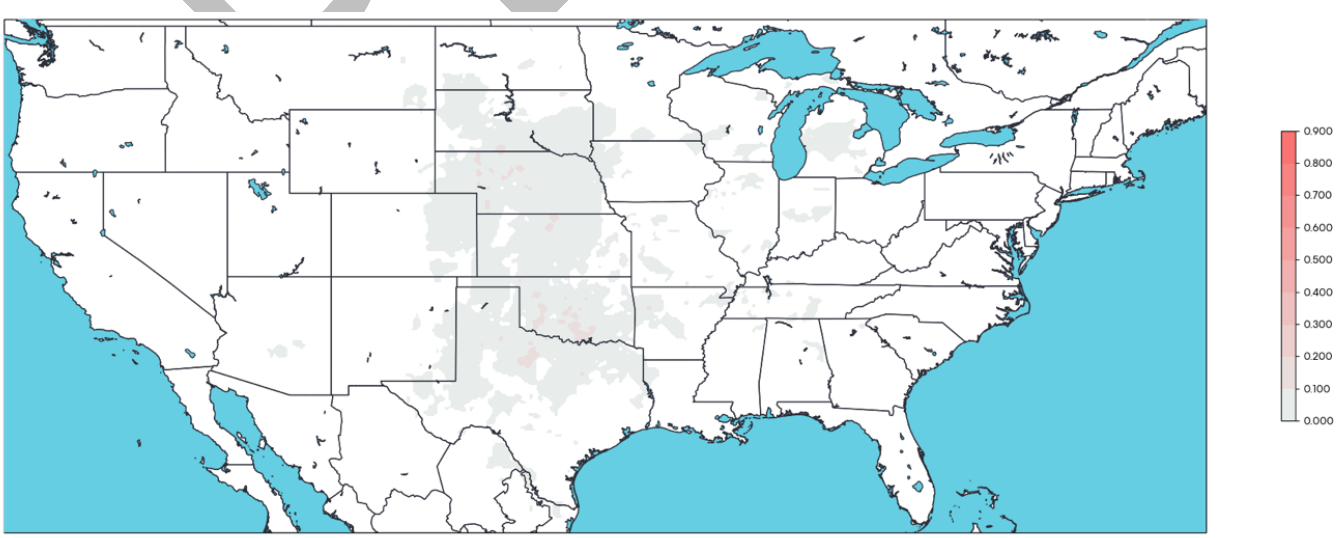

Supplement 10: Monthly average Hail-Days: October

Supplement 11: Monthly average Hail-Days: November

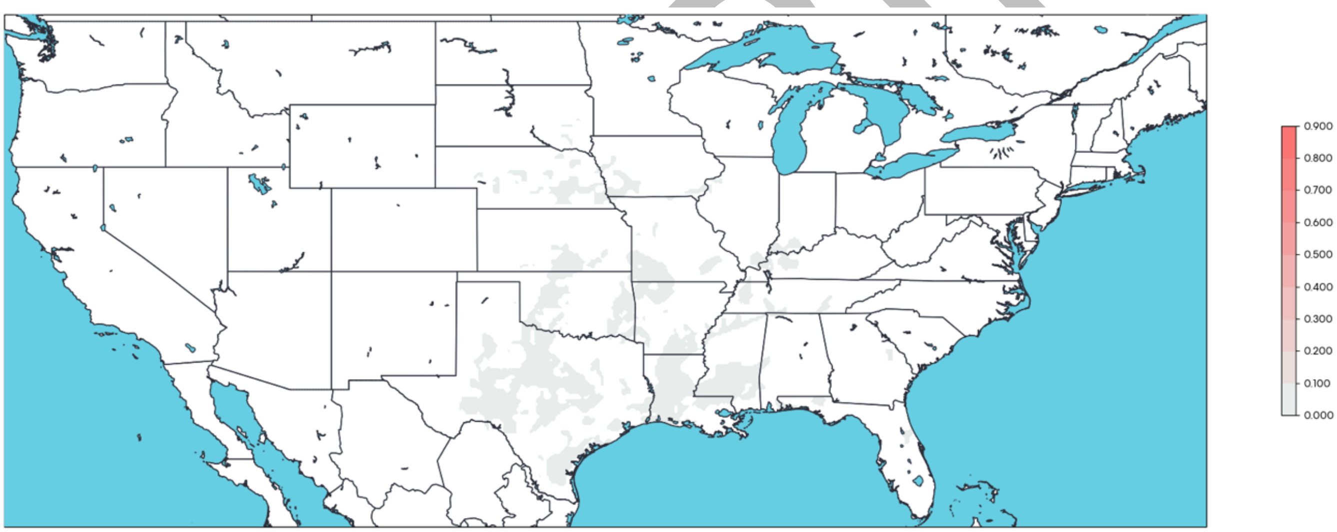

Supplement 12: Monthly average Hail-Days: December

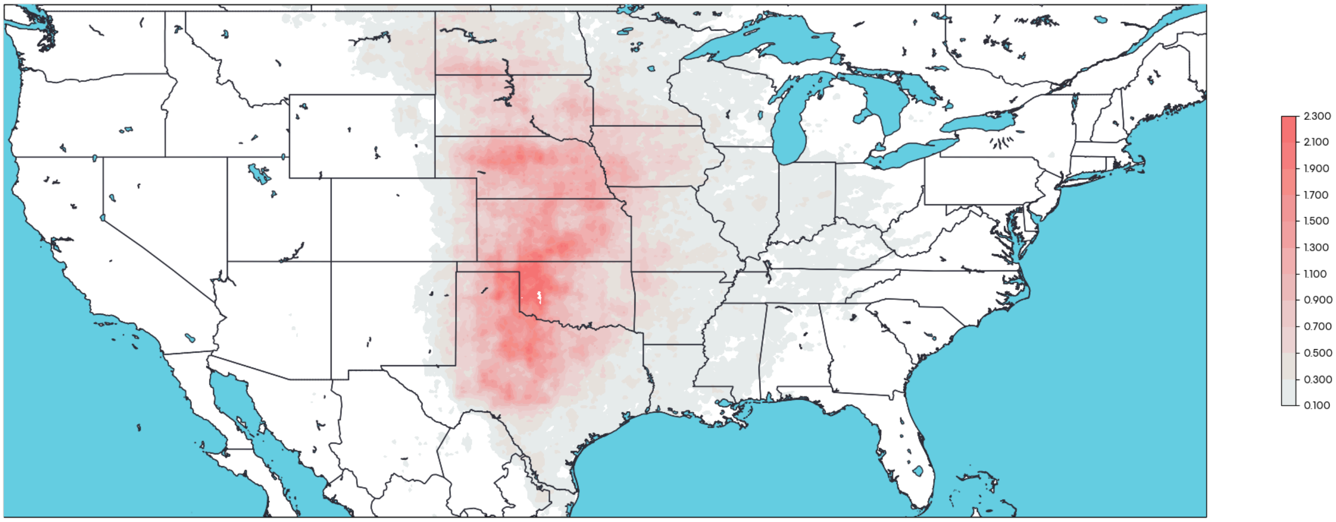

Supplement 13: Average annual hail-days

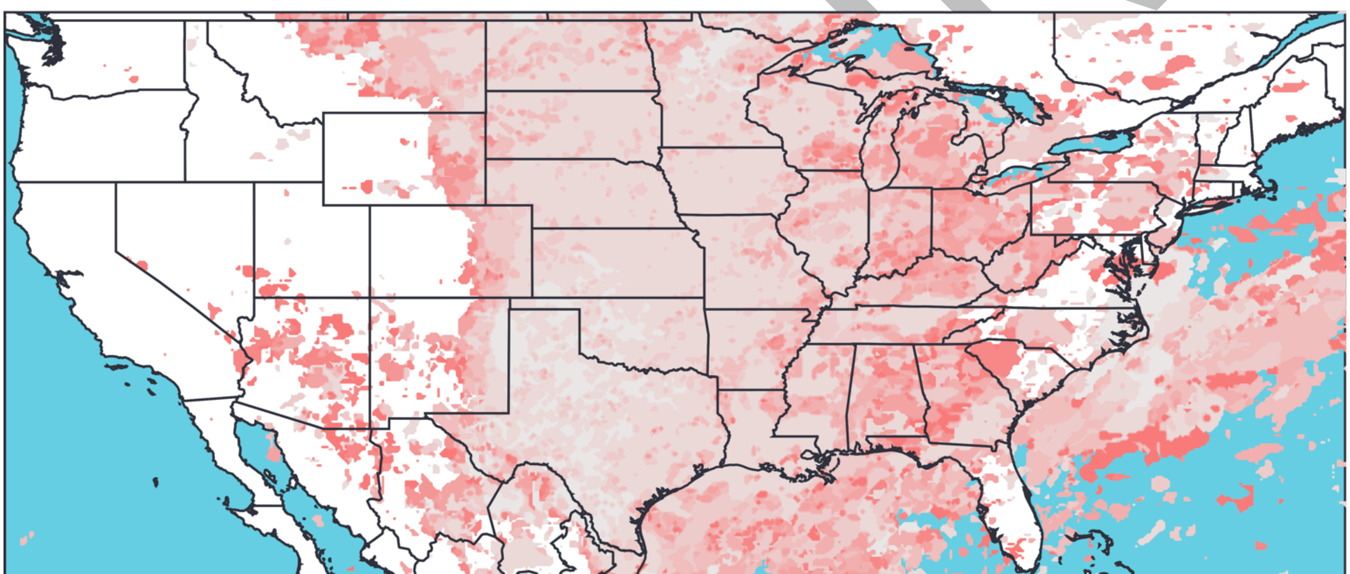

Supplement 14: First modeled occurrence of hail

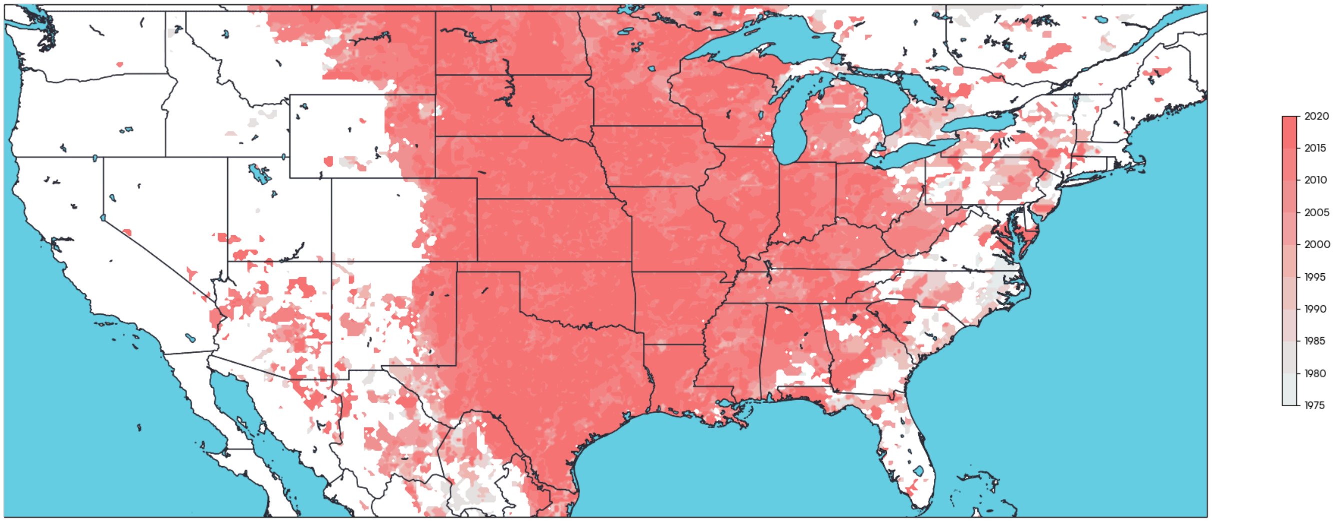

Supplement 15: Last modeled occurrence of hail

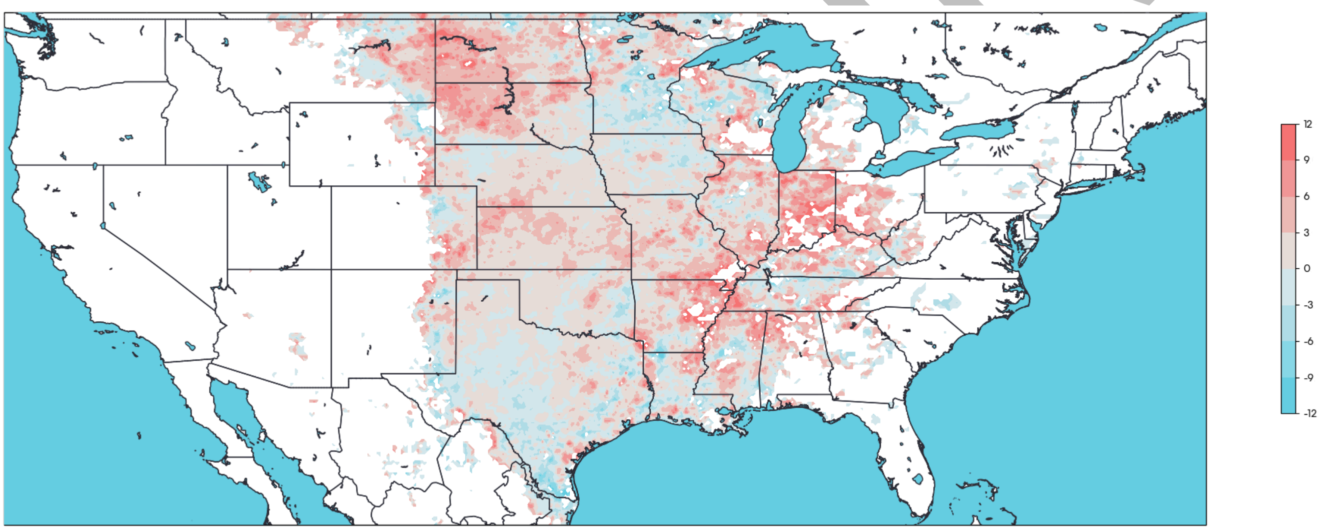

Supplement 16: Average annual percent change in hail days: 1979-1998 to 1999-2018.

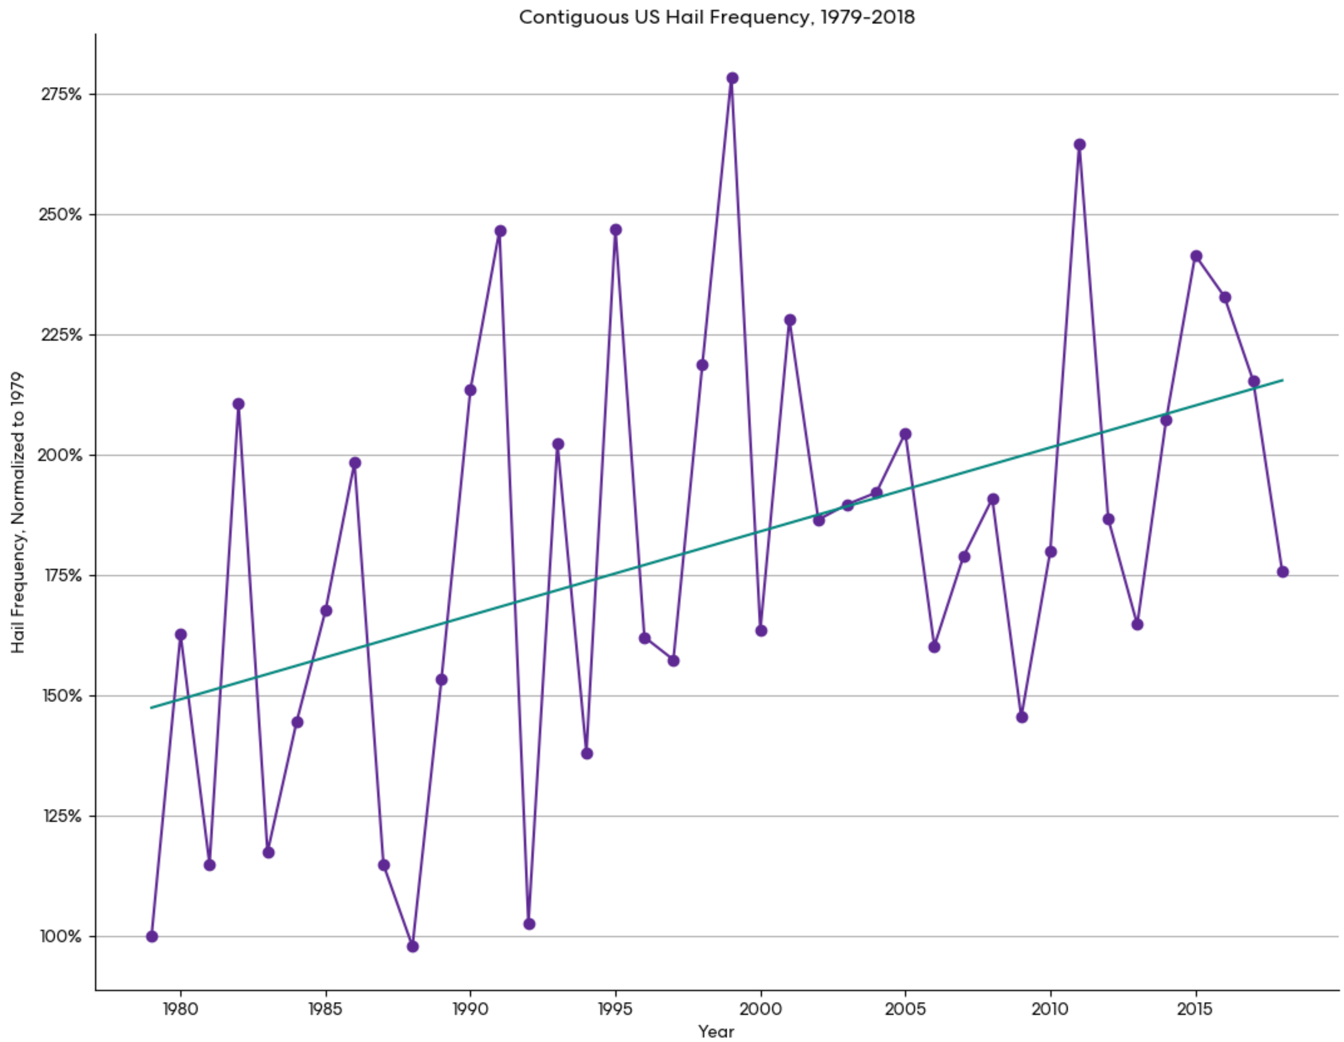

Supplement 17: Total hail-days per year