



Libraries and Learning Services

University of Auckland Research Repository, ResearchSpace

Version

This is the Accepted Manuscript version. This version is defined in the NISO recommended practice RP-8-2008 <http://www.niso.org/publications/rp/>

Suggested Reference

Swaminathan, R., Sridharan, M., Dobbie, G., & Hayhoe, K. (2015). Modeling Ice Storm Climatology. In B. Pfahringer, & J. Renz (Eds.), *Lecture Notes in Computer Science: AI 2015: Advances in Artificial Intelligence* Vol. 9457 (pp. 539-553). Canberra, Australia: Springer. doi: [10.1007/978-3-319-26350-2_48](https://doi.org/10.1007/978-3-319-26350-2_48)

Copyright

Items in ResearchSpace are protected by copyright, with all rights reserved, unless otherwise indicated. Previously published items are made available in accordance with the copyright policy of the publisher.

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-26350-2_48

For more information, see [General copyright](#), [Publisher copyright](#), [SHERPA/RoMEO](#).

Modeling Ice Storm Climatology

Ranjini Swaminathan^{1,3}, Mohan Sridharan², Gillian Dobbie¹, and Katharine Hayhoe³

¹ Department of Computer Science, The University of Auckland, NZ

² Department of Electrical and Computer Engineering, The University of Auckland, NZ

³ Climate Science Center, Texas Tech University, USA

ranjinis@gmail.com; m.sridharan@auckland.ac.nz; gill@cs.auckland.ac.nz;
katharine.hayhoe@ttu.edu

Abstract. Extreme weather events such as ice storms cause significant damage to life and property. Accurately forecasting ice storms sufficiently in advance to offset their impacts is very challenging because they are driven by atmospheric processes that are complex and not completely defined. Furthermore, such forecasting has to consider the influence of a changing climate on relevant atmospheric variables, but it is difficult to generalise existing expertise in the absence of observed data, making the underlying computational challenge all the more formidable. This paper describes a novel computational framework to model ice storm climatology. The framework is based on an objective identification of ice storm events by key variables derived from vertical profiles of temperature, humidity, and geopotential height (a measure of pressure). Historical ice storm records are used to identify days with synoptic-scale upper air and surface conditions consistent with an ice storm. Sophisticated classification algorithms and feature selection algorithms provide a computational representation of the behavior of the relevant physical climate variables during ice storms. We evaluate the proposed framework using reanalysis data of climate variables and historical ice storm records corresponding to the north eastern USA, demonstrating the effectiveness of the climatology models and providing insights into the relationships between the relevant climate variables.

1 Introduction

Extreme winter weather events such as snow, sleet and ice storms can cause significant damage to property and life. Perhaps the most dangerous of these are ice storms that result from freezing rain. Ice storms are a globally occurring phenomena that can account for billions of dollars of damage to urban and natural systems.

Ice storms account for roughly 60% of winter storm losses within the United States, and they have caused billions of dollars of loss in the continental United States and Canada [11,16]. The catastrophic effects of ice storms are being seen all over the world, with the 2008 ice storm in China resulting in 129 human fatalities, \$22.3 billion in direct economic losses and structural damage, and the displacement of ≈ 1.7 million people [31]. The increasing concentration of human populations in areas vulnerable to ice storms, such as the north eastern United States, will only exacerbate the impact of such ice storms in the future.

It is very challenging to accurately forecast ice storms sufficiently in advance to offset the extent of the associated losses, primarily because the complex atmospheric

processes driving ice storms are difficult to model. Traditional weather models forecast ice storms using very high resolution observed data for climate variables, and extensive interpretation by humans with domain expertise. Besides overcoming the difficulty of fully understanding and modeling the complex physical processes in the atmosphere that cause ice storms, the influence of a changing climate is a crucial factor to be considered when detecting ice storms in the future. Global Climate Models (GCMs) provide projections for climate variables decades and even a century in advance under changing climate scenarios, thereby providing the scientific community with a framework to study and model future climate events.

Existing expertise, however, is difficult to generalise and apply to future projections in order to quantify the impact of human-induced climate change on the occurrence and severity of ice storms. Such analysis is challenged by the lack of “observed” data and by the necessity to consider the influence of a changing climate on the atmospheric variables that drive ice storms. The work described in this paper is thus motivated by the need to represent complex ice storm climatologies in order to project ice storms in the future under a changing climate. We present a computational framework to model the complex phenomena of ice storms, and make the following key contributions:

- we build on classification algorithms to model ice storm climatology, and to identify fundamental climate patterns that are indicative of this climatology;
- we exploit feature selection algorithms to improve both the computational efficiency and the effectiveness of the climatology models; and
- we provide a computational representation of the relevant physical climate variables and their behaviour during ice storms, affording atmospheric scientists better insights into the relationships between these variables.

The proposed computational framework is illustrated and experimentally evaluated using reanalysis data, comprising relevant physical climate variables, and historical ice storm records, corresponding to the north eastern USA.

2 Related Work

There are no published studies in the literature that apply sophisticated AI or machine learning algorithms to model ice storm climatology. This section motivates our work by describing how ice storm climatology has been modeled historically, and by summarizing the broader applicability of computational models for climate science.

Ice Storm Forecasting: Ice storms develop due to many complex atmospheric processes that have traditionally been studied at different spatial (micro, meso and synoptic or cyclone scale) and temporal (from hourly to a few hours) resolutions. Researchers have studied the effect of topology on freezing rain, and investigated the influence of cold-air damming due to the Appalachian mountains in the north eastern USA and the Rocky Mountain areas [2,26]. Researchers have also attributed the spatial variation of freezing rain events around the Great Lakes region in the USA, to the frequency of surface cyclone tracks, availability of moisture in the Atlantic ocean, and regional topography [8]. Others have used composites of relevant climate variables over a period of five days to describe how large scale circulation patterns of these variables fostered freezing rain events in the Great Lakes region [16]. The sensitivity of the Atlantic ocean’s sea

surface temperatures to ice events in the north eastern USA has been investigated [9], and the atmospheric conditions that typically coincide with north eastern ice storms, including synoptic scale movement of moisture and temperature, have been identified [5].

The literature on studying and modeling ice storms shows that the complex factors driving ice storms are yet to be completely defined. While it is established that synoptic scale events play a key role in their occurrence, existing methods continue to use composites of high resolution observational data for climate variables. The data is further processed by computationally intensive numerical weather models, which limits these studies to single or selected features of freezing events in geographically localized areas. Furthermore, to forecast ice storms in the present or near future by observing salient weather patterns, meteorologists still rely on human expertise to make the final decision regarding intensity and occurrence. The requirements of high resolution data and extensive human input highlight the need for alternative models for ice storm climatology. This need is amplified by the question of how human-induced climate change may be affecting the atmospheric circulation patterns that lead to ice storms. Global Climate Models (GCMs) use the laws of physics to simulate many processes and interactions in the climate system. GCMs are the main tool used to study future climate; they can generate three-dimensional projections for different atmospheric variables under future climate change scenarios. GCM simulations provide data at coarser spatial granularities (250km vs 32km) and lower temporal resolutions (daily averages instead of a few hours) than weather stations. Researchers have generated future ice storm projections for the north eastern USA and eastern Canada by using statistical downscaling techniques [6,7]. These techniques establish statistical relationships between large-scale GCM projections and high-resolution observations, and use these relationships to obtain “observation-like” local data with future GCM projections [29]. However, statistical downscaling makes the assumption that the relationships between large-scale and local climate variables is stationary over time, and does not exploit the complex inter-variable relationships within and across scales. These limitations make it challenging to use statistically downscaled data for modeling complex interactions between climate variables, such as those driving ice storms.

Computational models for Climate Science: Research in climate and atmospheric sciences poses a plethora of scientific questions that could greatly benefit from advances in AI. A recent tutorial at a premier machine learning conference highlighted the complex nature of many problems of interest, and briefly summarized some of the computational solutions that have been developed for a subset of these problems [1]. A recent article describes that we are yet to explore advanced computational techniques in the context of climate and large-scale meteorological processes such as ice storms [14]. Our research study is motivated by our strong belief that computational solutions for problems in climate science (and other research domains) should go beyond just a superficial application of techniques—they should actively improve our understanding of the problem and the domain [28]. We describe a framework that addresses key computational challenges to address a hard problem in climate science. This framework also serves as an integral first step towards improving the climate science community’s understanding and ability to detect the complex climate phenomenon of ice storms.




Fig. 1. Freezing rain forms when a snowflake falls through alternating cold and warm layers before it gets supercooled and freezes on contact with the ground. The figure also shows how sleet and snow form in comparison to freezing rain. Figure adapted and modified from [27].

3 Materials and Methods

In this section, we first summarize what is currently known about the climatology of ice storms, identifying the variables that we consider in the proposed computational framework. Next, we identify and describe the sources of data corresponding to the study area in which we illustrate and evaluate our framework. Finally, we describe the computational models for classification and feature selection that are included in the proposed framework.

3.1 Ice Storm Climatology

Ice storms are caused by freezing rain and start out in a cold layer of moist air up in the atmosphere quite the same way as snow and sleet, i.e., as a snowflake. In the case of snowfall, the snowflake continues to drop down through cold air and settles as snow when it hits the surface of the earth. Sleet forms when the snowflake first melts through a layer of warm air and freezes while moving through a cold layer of air again before it hits the ground. In the case of freezing rain, the snowflake moves down through a layer of warm air and melts into rain first, then into another layer of cold air where it turns into a supercooled droplet. If the surface temperature of the earth is below freezing, then this supercooled droplet freezes as soon as it touches any surface. We see that *certain specific and relatively rare atmospheric conditions must be in place for freezing rain to occur*. This freezing rain then creates a film of ice on the ground and the surface of every object it falls on. Figure 1 shows a schematic representation of how freezing rain forms and how it differs from other forms of precipitation.

To consider the atmospheric circulation patterns that are consistent with ice storm events, we analyze the vertical profiles of three climate variables: temperature, humid-



Fig. 2. Synoptic scale plot showing the geographic range of the climate variables considered for modeling ice storm climatology in north eastern USA. The range extends from 24° N to 55° N and from 50° W to 94° W.

ity, and geopotential height. While humidity and temperature at different atmospheric pressure levels account for the moisture and alternating warm and cold layers, the geopotential height approximates the actual height of a pressure surface above mean sea level. In other words, it is the height of a constant pressure surface (e.g., 750 millibars) above the surface of the earth measured in meters. We also consider these atmospheric patterns at *synoptic scales*, as seen in Figure 2. For the chosen study area of north eastern USA, this means that we capture upper atmosphere dynamics that are characteristic of ice storms in the large area that comprises the north eastern USA, and do not attempt to localize our projections to county-level geographic locations within the area. Considering the patterns at synoptic scale helps develop more reliable models of the ice storm climatology without human feedback and downscaling errors.

3.2 Data Selection

Although we rely on climate variable projections from GCMs for estimating ice storm occurrences in the future, GCM simulations are themselves prone to uncertainties that are difficult to quantify. Uncertainty in GCM output can broadly be attributed to: (1) natural variability in the climate system; (2) scientific limitations in understanding and modeling climate; and (3) human choices that will determine the extent of future emissions [19,20]. Our initial model development is therefore done with data generated by a process called “reanalysis”, which is a systematic approach to generate consistent data for climate monitoring and research. Reanalysis is a consistent reprocessing of archived weather observations with modern forecasting systems producing gridded datasets that estimate a large variety of atmospheric, sea-state and land surface parameters [10]. Reanalysis data is also well-suited for our task because there is no physical way to obtain

comprehensive weather observations for climate variables at different heights above the earth's surface. Furthermore, reanalysis data is available at different spatial and temporal resolutions. Reanalysis data thus enables us to build more robust models than with just the GCM output. Since GCM simulations are independent of observations, the models developed also provide a framework to evaluate different GCMs and identify those that would be best suited for future ice storm projections. For our work, we used reanalysis data from the National Centers for Environmental Prediction (NCEP) provided by the National Oceanic and Atmospheric Administration, USA [24].

We model ice storm climatology in the north eastern USA by studying the atmospheric patterns that characterized historical ice storm events in the region. We chose the same 175 significant ice storm events selected in [5], all occurring in winter periods (October 1st to April 30th) between the years 1993 and 2010. These events were selected from the National Climatic Data Center's (NCDC, USA) Storm Database and a given event qualified as an ice storm if it met one of three established criteria [5]. Any ice storm events occurring in the northeastern USA is also influenced by atmospheric processes at much larger geographic scales (synoptic), which could potentially include larger masses of land and even the ocean. The synoptic scale considered sufficiently large (by climate scientists) for studying these ice storms is shown in Figure 2, spanning the area between 24° N and 55°N from 50° W to 94° W. This translates to a 145×149 grid of climate variable estimates from the reanalysis data. Data for temperature (2m, 850hPa, 925hPa, 100hPa), specific humidity (2m, 750hPa, 850hPa, 925hPa, 1000hPa) and geopotential height(200hPa, 500hPa, 700hPa, 850hPa) was used for each cell in this grid since these values were considered relevant to ice storm formation by domain experts. The temporal component of the data consists of estimates of climate variables at three-hour intervals each day.

3.3 Experimental Framework

Ice storms are characterized by signature patterns defined by climate variables at various atmospheric pressure levels. The task of learning computational models that capture this signature from historical data, and use the models to identify similar patterns in the future, can be posed as a supervised learning problem. The *Learning Task* is to learn models for climate variable patterns during ice storm and non-ice storm winter periods in the past. The *Performance Task* is to apply the learned models to identify ice storms in future GCM projections.

Classification Models for Ice Storm Climatology: We considered some popular algorithms such as decision trees, rule-based system, neural networks, and Support Vector Machines (SVMs) to classify ice storm and non-ice storm patterns defined by climate variables. The SVM-based classifier was (experimentally) observed to provide the best classification accuracy, and was hence used in all subsequent experiments. Specifically, we adapted an SVM with a linear kernel to use the climate variables as input features, providing class labels ("positive" for ice storms and "negative" for non-ice storms) as the corresponding outputs. SVMs transform the task of determining boundaries around patterns in feature spaces into the dual convex quadratic problem. Features and patterns are projected to high(er) dimensions and loss functions are used to penalize errors, resulting in sparse representations and robust decision boundaries [18,22,25].

SVMs have a simple geometric interpretation, provide robustness to overfitting, and use structural risk minimization methods [3]. SVMs are well-suited to model the complex relationships that characterize the behavior of ice storms. Since the reanalysis data was available at very high resolutions, the larger 145×149 grid was subsampled by averaging values (of each variable) across 5×5 grid cells—this size of the subsampling mask provides a granularity consistent with GCM projections. Furthermore, the three-hourly readings for the different climate variables were averaged to provide the equivalent of daily average readings, which is (again) the resolution at which GCM projections will be available. To evaluate the sensitivity of classification to different spatio-temporal scales, we considered other grid sizes and temporal resolutions—see Section 4.

We considered the relevant climate variables, i.e., temperature, humidity and geopotential height, at pressure levels ranging from the surface (2 meters) to 1000 millibars. Traditional ice storm forecasting methods consider five-day composite maps for climate variables. Instead, we build a classification model for each day in a five-day window (three days before, the day of the storm, and one day after). A total of 175 vectors of the historical record of ice storms, each with 10092 features, were used to build the classification model for each day in the window. We can now apply a sliding window on GCM projections for winter periods during future decades (or centuries) to identify five-day sequences that signal the potential for an ice storm. Section 4 describes the outcome of 10-fold cross validation on this dataset for the baseline (SVM) classifier,

Feature Selection for Ice Storms: Given the complexity of the underlying atmospheric processes, and the relatively small number of ice storm events under consideration, learning the classification model for each day can be challenging and computationally expensive with input vectors of ≈ 10000 features. Also, the learned models will be used to identify ice storms in time periods spread over several decades. One objective of this work is to identify features that contribute most significantly towards the ice storm signature so that climate scientists can gain better insights into the physical characteristics of the atmospheric processes driving ice storms. We would also like to see how the features selected change over the five-day window to better understand the change in relevance of the climate variables as the ice storm develops. Dimensionality reduction methods such as Principal Component Analysis [21] are not applicable since they transform the original features into a different dimensional space, and we lose the correspondence information needed to identify the contribution of individual climate variables. Guidelines laid out by [15] recommend first applying domain knowledge to perform feature selection and reduce the dimensionality of features for classification tasks. However, we are unable to use domain knowledge to reduce the set of features—the complexity of the relationships between climate variable driving ice storms makes it difficult for even a domain expert to identify the most relevant features.

We investigated two different measures for ranking features: (1) Pearson’s correlation coefficient; and (2) a mutual information-based information-theoretic measure. The Pearson’s correlation coefficient is computed between individual features and the class labels:

$$\rho_{Class,Attribute} = \frac{cov(Class,Attribute)}{\sigma_{Class}\sigma_{Attribute}} \quad (1)$$

where $\rho_{Class,Attribute}$ is the Pearson’s coefficient, $cov(Class,Attribute)$ is the covariance between random variables $Class$ and $Attribute$, and σ is the standard deviation. The

In the 28th Australasian Joint Conference on Artificial Intelligence (AusAI), Canberra, Australia, November 30-December 4, 2015.

Readings per day	Day 1	Day 2	Day 3	Day 4	Day 5
Four	0.594	0.640	0.717	0.780	0.780
Two	0.606	0.623	0.708	0.754	0.814
One	0.528	0.614	0.634	0.740	0.814
Daily avg.	0.571	0.606	0.683	0.769	0.751

Table 1. Classification performance of ice storm vs. non-ice storm dates in winter periods. Values represent the F-measure for a linear kernel support vector machine classifier. Every 5×5 grid in the input data becomes a single supercell for the classification.

Readings per day	Day 1	Day 2	Day 3	Day 4	Day 5
Four	0.594	0.640	0.711	0.797	0.786
Two	0.608	0.629	0.717	0.760	0.806
One	0.531	0.609	0.648	0.749	0.820
Daily avg.	0.557	0.600	0.668	0.774	0.754

Table 2. Classification performance of ice storm vs. non-ice storm dates in winter periods. Values represent the F-measure for a linear kernel support vector machine classifier. Every 3×3 grid in the input data becomes a single supercell for the classification.

information gain measure, on the other hand, evaluates the worth of each feature as the information gained about each class given the individual feature:

$$InfoGain(Class, Attribute) = H(Class) - H(Class|Attribute) \quad (2)$$

where $H(Class)$ is the marginal entropy of the class and $H(Class|Attribute)$ is the conditional entropy of the class given the feature. We experimentally evaluated how classification performance changed with different feature set sizes, to deduce the optimal set of features for the problem of modeling ice storm climatology. As described in Section 4, the feature selection algorithms reduced the number of features considered, and improved the (overall) accuracy of the classifier built using the reduced number of features. As before, we used 10-fold cross validation to evaluate the classification performance of the combination of feature selection algorithms and the SVM classifier.

4 Experimental Results

We present results from various experiments conducted in relation to ice storm climatology modeling. All algorithms were implemented and evaluated on the datasets using the WEKA data mining software [17]. As stated in Section 3, the results reported below are obtained by performing 10-fold cross validation. We primarily used the F-measure to reflect classification performance in all the experimental trials—a higher value of this measure represents better classification accuracy. We also used some other measures (see below) to evaluate the significance of the performance of different algorithms.

Classification performance at varying spatio-temporal scales: As stated in Section 3, we considered classification performances with two different spatial resolutions, one where the gridded data is averaged over a mask of 3×3 cells and another with 5×5 cells. Reanalysis data is also available at different temporal resolutions and the values of the climate variables can be considered multiple times through the day. The daily average reading is obtained by averaging all available readings for an entire day. GCM simulation data for estimating ice storms in the future will be available to us in




Fig. 3. Classification performance for different sized feature sets, using the Pearson’s coefficient for feature selection on the days before, during, and after the ice storm.




Fig. 4. Classification performance for different sized feature sets, using the mutual information gain for feature selection on the days before, during, and after the ice storm.

the form of daily average readings and hence we consider the daily average f-measure to be our baseline performance. In order to provide a more comprehensive overview, we also conducted experiments with different combinations of spatial scales and temporal scales, with the corresponding classification performance (i.e., F-measure scores) summarized in Table 1 and Table 2. We observe that classification performance peaks on Day 4, which is the day of the storm. In some cases, the classifier for the day after the storm does better and can be attributed to the fact that many of the storm events typically occur over a period of two days. In general, this tells us that the ice storm signal is strongest on the day of the storm, and that our model is able to capture this signature.

As expected, we also observe that the classification performance improves for both spatial coverages as the temporal resolution increases (i.e., if we consider multiple readings of the climate variables in a day). However, we notice that the classification performance does not drop significantly when the spatial and temporal resolutions are set to match those of the GCM projections. In other words, we are able to demonstrate that our classification framework’s performance is reasonably robust to changes in spatial and temporal resolutions. For the remaining experiments, we thus used only the GCM-scale resolution data (5×5 smoothing and daily averages).

In the 28th Australasian Joint Conference on Artificial Intelligence (AusAI), Canberra, Australia, November 30-December 4, 2015.

Feature Subsets	F-measure	Kappa Statistic	RMSE
All	0.769	0.537	0.481
Top 500 (Pearson)	0.809	0.617	0.438
Top 500 (Info-gain)	0.786	0.571	0.463

Table 3. Classification performance and statistics for the day of the storm (Day 4). The Kappa measure increases with feature selection and the RMSE decreases. The best score for each column is highlighted in bold—feature selection using the Pearson coefficient provides the best results.

Feature Subsets	All	Top 500 (Pearson)	Top 500 (Info-gain)
All		↑ 1.625	↑ 0.5893
Top 500 (Pearson)			← -1.4289
Top 500 (Info-gain)			

Table 4. The z scores for the McNemar’s test comparing classification performance with and without feature selection. The arrow points in the direction of the better performing algorithm—using the Pearson coefficient for feature selection provides the best results.

Classification performance with Feature Selection: Next, we describe the classification performance when feature selection is included with the baseline classifier evaluated above. As described in Section 3.3, we used the correlation coefficient measure and the mutual information measure to rank features. Figure 3 shows the F-measure scores when the features were ranked using the correlation between individual features and the class label. Figure 4 shows the corresponding results for when the features were ranked based on a mutual information theoretic gain measure. In both cases, the number of features was varied to find an optimal set smaller than the entire set of features (10092) that best captured the signatures of ice storms. From Figures 3, 4, we observe that for both methods of feature selection, we get the best classification performance for Day 4, which is the day of the ice storm. We also observe that using the top 500 features, as ranked by the feature ranking measures, results (on average) in the best classification performance—this *much smaller* set of features can be considered as being representative of ice storm patterns.

Many methods have been developed to evaluate the significance of classification performance [4,12,13]. We used three measures to evaluate whether the ability to detect ice storms improves significantly with feature selection. First, the *Kappa statistic* [30] compares the observed accuracy with the expected accuracy or random chance. It is a measure of the agreement between the predicted and the actual classifications. A higher Kappa statistic value indicates better classification accuracy. Second, the *Root Mean Square Error* (RMSE) [30] is a measure of the difference between the values predicted by the classifier and the ground truth (i.e., known) values. A lower RMSE indicates better classification accuracy. Finally, we used the *McNemar test* [23], a non-parametric test used to analyze matched pairs of data. It compares the classification accuracy of two algorithms on a per-instance basis and computes a z score that can be translated into confidence levels. A *One-tailed Prediction* is used to determine when one algorithm is better than the other. Tables 3, 4 summarize the result of using these measures. Table 3 shows that selecting and using the top 500 features raises the Kappa value and lowers

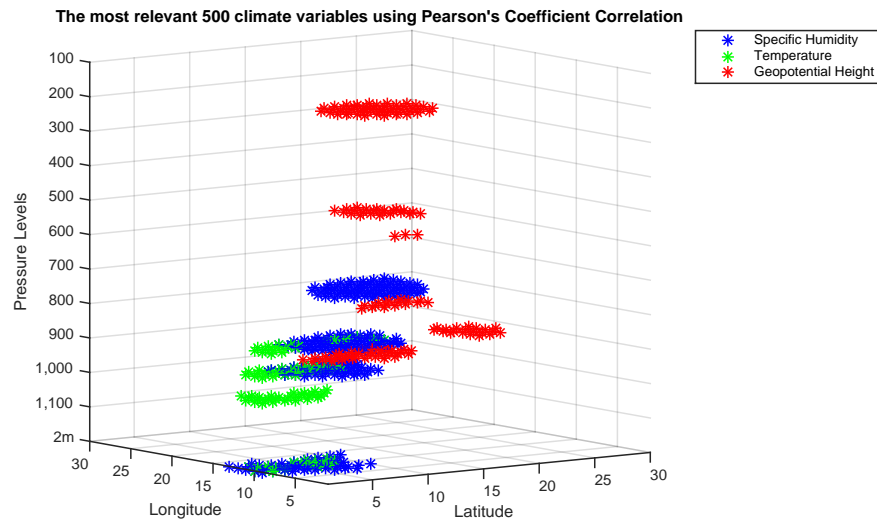


Fig. 5. The computational representation of the most relevant climate variables selected using Pearson’s Coefficient for ice storm modeling is shown for the day of the storm.

the RMSE. Table 4 shows that the classification performance improvement obtained using the Pearson’s coefficient for feature selection corresponds to a 94.8% confidence level. Due to space limitations, we only show the results for the day of the storm.

Computational representation of atmospheric processes: In addition to modeling the ice storm climatology, the results summarized above can also be used to computationally represent the climate variables that are most relevant to characterize ice storms. For instance, Figures5 shows the distribution of the climate variables at different spatial locations (specific latitude and longitude cells after 5x5 subsampling) and at different pressure levels on the day of the storm (i.e., Day 4), using the Pearson correlation coefficient for feature selection—both methods of feature selection show similarity in the atmospheric patterns (i.e., the spatial distribution of climate variables) on the day of the storm. The geopotential height variables selected indicate that a confluence of temperature and moisture variables occurs at $\approx 700mb$. This observation is consistent with the formation of snowflakes at this height. We also observe three different layers of temperature where the snowflake alternatively cools, melts and then freezes just as it hits the surface of the earth. This is clearly indicated in the representations by both methods of feature selection. Similar figures that depict the computational representations for the days before and after the storm (not shown in this paper) show the ice storm conditions slowly start to form during Day 2, peak on Day 4 and disappear on Day 5—these observations are (once again) completely consistent with the observed behavior of the ice storms. *The key finding is that we are able to computationally confirm these patterns using a much smaller subset of features typically considered while modeling ice storms.*

5 Conclusions and Future Work

Ice storms are potentially catastrophic weather events that can cause significant damage to life and property. The ability to project ice storm prevalence and intensity in the

future can help us minimize their negative impact. This paper presents a novel computational framework to model ice storm climatology, a complex problem in climate science. This framework learns atmospheric patterns that characterize historical ice storm occurrences using reanalysis data for the relevant climate variables. We show that the classification models in our framework are robust (according to climatologists) across spatial and temporal scale changes in climate data. In addition, the framework includes feature selection algorithms, which identify a smaller subset of features that increase the classification accuracy and significantly reduce the computational cost of building (and testing) the classification models on decadal data. Finally, our computational models help confirm and improve the existing understanding of what causes ice storms, in ways that will allow climate scientists to develop new insights into this complex atmospheric phenomenon.

This study opens up many directions for inter-disciplinary research. For instance, we would like to develop and use more sophisticated feature selection algorithms that can better ground the relationships between the climate variables driving ice storms. We would also like to use our models constructed with reanalysis data to evaluate Global Climate Model output for historical time periods, which will allow us to choose the Global Climate Models that are best suited for generating future ice storm projections.

Acknowledgments

This work was supported in part by the US National Science Foundation under Grant No. DEB-1457875.

References

1. A Banerjee and C Monteleoni. Tutorial: Climate Change: Challenges for Machine Learning. In *Neural Information Processing Systems*, 2014.
2. BC Bernstein. Regional and Local Influences on Freezing Drizzle, Freezing Rain, and Ice Pellet Events. *Weather and forecasting*, 15(5):485–508, 2000.
3. CM Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
4. Betul Bostanci and Erkan Bostanci. An Evaluation of Classification Algorithms Using Mc Nemars Test. In *Proceedings of Seventh International Conference on Bio-Inspired Computing: Theories and Applications (BIC-TA 2012)*, pages 15–26. Springer, 2013.
5. CM Castellano. *Synoptic and Mesoscale Aspects of Ice Storms in the Northeastern US*. Master’s Thesis in Atmospheric Science, University at Albany, SUNY, USA, 2012.
6. CS Cheng, H Auld, G Li, J Klaassen, and Q Li. Possible Impacts of Climate Change on Freezing Rain in South-Central Canada using Downscaled Future Climate Scenarios. *Natural Hazards and Earth System Science*, 7(1):71–87, 2007.
7. CS Cheng, G Li, and H Auld. Possible Impacts of Climate Change on Freezing Rain using Downscaled Future Climate Scenarios: Updated for Eastern Canada. *Atmosphere-Ocean*, 49(1):8–21, 2011.
8. J Cortinas Jr. A Climatology of Freezing Rain in the Great Lakes Region of North America. *Monthly weather review*, 128(10):3574–3588, 2000.
9. RR da Silva, G Bohrer, D Werth, MJ Otte, and R Avissar. Sensitivity of Ice Storms in the Southeastern United States to Atlantic SST-Insights from a Case Study of the December 2002 Storm. *Monthly weather review*, 134(5):1454–1464, 2006.

In the 28th Australasian Joint Conference on Artificial Intelligence (AusAI), Canberra, Australia, November 30-December 4, 2015.

10. DP Dee, M Balmaseda, G Balsamo, R Engelen, AJ Simmons, and J-N Thépaut. Toward a Consistent Reanalysis of the Climate System. *Bulletin of the American Meteorological Society*, 95(8):1235–1248, 2014.
11. AT DeGaetano. Climatic Perspective and Impacts of the 1998 Northern New York and New England Ice Storm. *Bulletin of the American Meteorological Society*, 81(2):237–254, 2000.
12. Janez Demšar. Statistical Comparisons of Classifiers Over Multiple Data Sets. *The Journal of Machine Learning Research*, 7:1–30, 2006.
13. Thomas G Dietterich. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural computation*, 10(7):1895–1923, 1998.
14. R Grotjahn, R Black, R Leung, MF Wehner, M Barlow, M Bosilovich, A Gershunov, WJ Gutowski Jr, JR Gyakum, RW Katz, et al. North American Extreme Temperature Events and Related Large Scale Meteorological Patterns: A Review of Statistical Methods, Dynamics, Modeling, and Trends. *Climate Dynamics*, pages 1–34, 2015.
15. I Guyon and A Elisseeff. An Introduction to Variable and Feature Selection. *The Journal of Machine Learning Research*, 3:1157–1182, 2003.
16. JR Gyakum and PJ Roebber. The 1998 Ice Storm-Analysis of a Planetary-scale Event. *Monthly weather review*, 129(12):2983–2997, 2001.
17. M Hall, E Frank, G Holmes, B Pfahringer, P Reutemann, and IH Witten. The WEKA Data Mining Software: an Update. *ACM SIGKDD explorations newsletter*, 11(1):10–18, 2009.
18. T Hastie and R Tibshirani. Classification by Pairwise Coupling. In *Advances in Neural Information Processing Systems*, volume 10. MIT Press, 1998.
19. E Hawkins and R Sutton. The Potential to Narrow Uncertainty in Regional Climate Predictions. *Bulletin of the American Meteorological Society*, 90(8):1095–1107, 2009.
20. E Hawkins and R Sutton. The Potential to Narrow Uncertainty in Projections of Regional Precipitation Change. *Climate Dynamics*, 37(1-2):407–418, 2011.
21. I Jolliffe. *Principal Component Analysis*. Wiley Online Library, 2002.
22. SS Keerthi, SK Shevade, C Bhattacharyya, and KRK Murthy. Improvements to Platt’s SMO Algorithm for SVM Classifier Design. *Neural Computation*, 13(3):637–649, 2001.
23. Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947.
24. F Mesinger, G DiMego, E Kalnay, K Mitchell, PC Shafran, W Ebisuzaki, D Jovic, J Woollen, E Rogers, EH Berbery, et al. North American Regional Reanalysis. *Bulletin of the American Meteorological Society*, 87(3):343–360, 2006.
25. J. Platt. Fast training of support vector machines using sequential minimal optimization. In B Schoelkopf, B Burges, and A Smola, editors, *Advances in Kernel Methods - Support Vector Learning*. MIT Press, 1998.
26. RM Rauber, LS Olthoff, MK Ramamurthy, D Miller, and KE Kunkel. A Synoptic Weather Pattern and Sounding-based Climatology of Freezing Precipitation in the United States East of the Rocky Mountains. *Journal of Applied Meteorology*, 40(10):1724–1747, 2001.
27. D Spector. Why Icestorms are so Dangerous, Business Insider Australia, 12th February 2014.
28. KL Wagstaff. Machine Learning that Matters. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 529–536, 2012.
29. RL Wilby, TML Wigley, D Conway, PD Jones, BC Hewitson, J Main, and DS Wilks. Statistical Downscaling of General Circulation Model Output: A Comparison of Methods. *Water Resources Research*, 34(11):2995–3008, 1998.
30. Ian H Witten and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2005.
31. B Zhou, L Gu, Y Ding, L Shao, Z Wu, X Yang, C Li, Z Li, X Wang, Y Cao, et al. The Great 2008 Chinese Ice Storm: Its Socioeconomic-ecological Impact and Sustainability Lessons Learned. *Bulletin of the American Meteorological Society*, 92(1):47–60, 2011.