

Alexandra K. Anderson-Frey* and Yvette P. Richardson
 Pennsylvania State University, University Park, PA

Andrew R. Dean, Richard L. Thompson, and Bryan T. Smith
 NWS Storm Prediction Center, Norman, OK

1. INTRODUCTION

Our understanding of and ability to forecast tornadoes has improved markedly (Brooks 2004) over the past several decades with the deployment of new technology (e.g., Doppler radar), as well as the development of improved conceptual models for tornadoes and severe storms (e.g., Markowski and Richardson 2014, Davies-Jones 2015). Despite these advances, there is considerable room for improvement in warning verification statistics: over a period of interest encompassing Jan 01 2003 through June 30 2013, the probability of detection (POD; the percentage of tornadoes for which a warning was issued ahead of time) was 68%, and the false-alarm ratio (FAR; the percentage of tornado warnings for which no confirming report was ever received) was 77%.

Tornado climatologies often focus exclusively on tornadoes originating from traditional right-moving supercells (e.g., Alexander 2010, Rasmussen and Blanchard 1998, Rasmussen 2003), which are responsible for producing over 95% of all tornadoes rated EF3+ (Smith et al. 2012), but are responsible for only about three-quarters of all tornadoes reported during the aforementioned time period. Since many non-supercellular storm modes have poorer tornado warning statistics than traditional right-moving supercells, we take a holistic look at the entire spectrum of tornadic storms, which includes storm modes such as quasi-linear convective systems and more disorganized clusters.

We make use of a large dataset consisting of 12,090 tornado reports and 40,357 tornado warnings issued over the period of interest, the associated warning verification information, model mesoanalysis data as proxies for the near-storm environment, and manually compiled storm mode information based on radar observations (Smith et al. 2012). This work is a proof-of-concept for the use of self-organizing maps (SOMs) in clustering and analyzing the two-dimensional environments surrounding tornadic events [see Nowotarski and Jensen (2013) for an example of the application of SOMs to one-dimensional atmospheric soundings].

2. METHODS AND DATA

*Corresponding author address: Alexandra K. Anderson-Frey, Department of Meteorology, The Pennsylvania State University, University Park, PA 16801. aka145@psu.edu

2.1 The Near-Storm Environmental Dataset

The archived mesoanalysis data corresponding to each tornado event originates from the Storm Prediction Center (SPC) and consists of proximity sounding data obtained from the Rapid Update Cycle model (Benjamin et al. 2002) until April 2012, and thereafter originates from the Rapid Refresh model (RAP). Smith et al. (2012) have, in addition, manually classified each of these tornado events in terms of their parent storm mode based on visual inspection of the radar data; the two storm modes we discuss here are right-moving supercells (RMS) and quasi-linear convective systems (QLCS). The geographical distributions of these two storm modes are illustrated in Fig. 1.

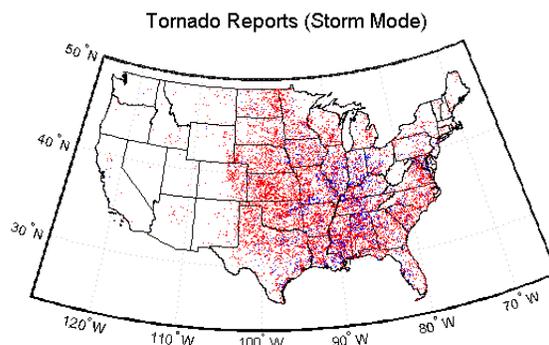


Figure 1: Geographical distribution of right-moving supercell (red) and QLCS (blue) tornado reports received between January 2003 and June 2013.

In order to gain a baseline understanding of the near-storm environments that characterize these tornado reports and warnings, we plot their distributions in two parameter spaces. The first is mixed-layer convective available potential energy (MLCAPE) versus 0-6 km vector shear magnitude (SHR6, the magnitude of the vector difference between the wind at the surface and the wind at 6 km); this combination of parameters has been shown to discriminate fairly well between supercellular and non-supercellular convection (Brooks et al. 2003). The second parameter space, mixed-layer lifting condensation level (MLLCL) versus 0-1 km storm-relative helicity (SRH1), has been shown to discriminate fairly well between nontornadic and significantly tornadic (i.e., EF2+) supercells.

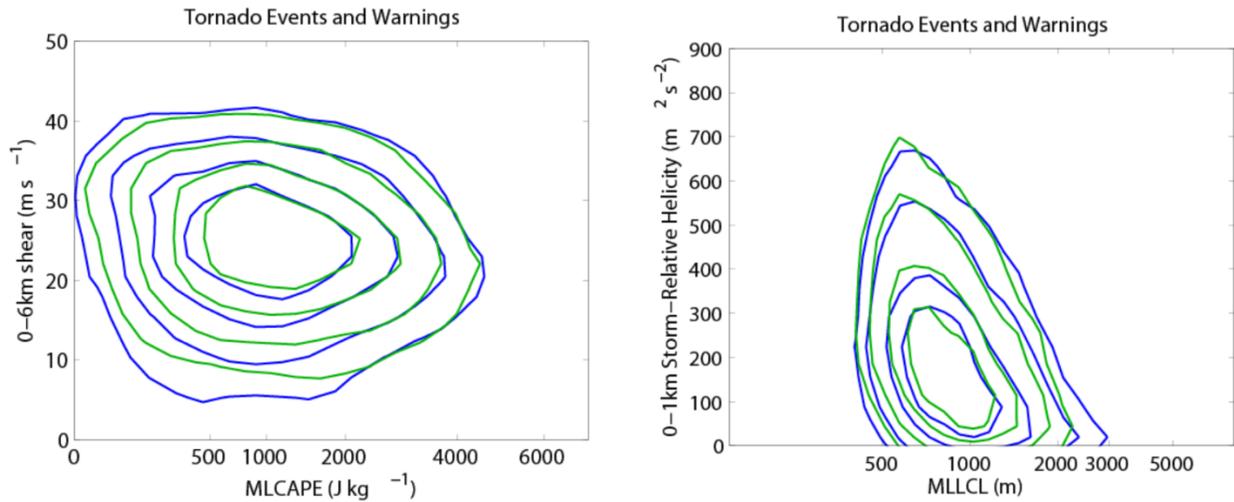


Figure 2: Distribution of tornado events (blue) and warnings (green) for the MLCAPE--SHR6 (left) and MLLCL--SRH1 (right) parameter spaces. Data are smoothed using kernel density estimation, where each plot is centered on the point of highest data density. The innermost contour contains 25% of the data, followed by 50%, 75%, and 90% moving outward.

Figure 2 plots the tornado reports and warnings in our two parameter spaces. There is no readily apparent “smoking gun” here: no part of the parameter space is very obviously being warned with no corroborating reports, and no part of the parameter space features tornadoes reported without any warnings issued; in short, the two distributions overlap through most of the parameter space. Looking at the entire set of tornado reports and warnings, overall the two distributions match well; peak report and warning densities are both centered near 1250 J kg⁻¹ of MLCAPE, 25 m s⁻¹ of SHR6, 800 m MLLCL heights, and 200 m² s⁻² of SRH1.

Given that these four environmental parameters change according to factors such as time of day (Fig. 3), time of year, and geographical location, we would

expect the distributions of tornado reports and warnings within these parameter spaces to shift accordingly. The effect of these shifts on tornado warning skill is outside the scope of this study, and is discussed in additional literature (Anderson-Frey et al. 2014, 2016).

To help tease out any signal that may be buried when looking at the dataset as a whole, we can also create several clusters of statistically distinct near-storm environments through the use of self-organizing maps. This approach will enable us to analyze several distinct tornadic near-storm environments and discuss the warning skill of each one individually.

2.2 Self-Organizing Maps

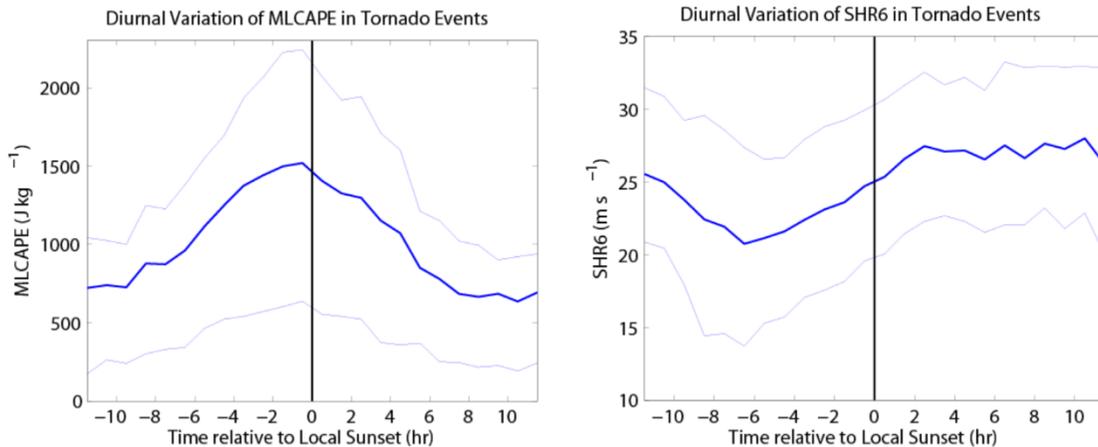


Figure 3: Diurnal variation of MLCAPE and SHR6. The bold blue line is the mean value within each bin, while the thin blue lines represent the 25th and 75th percentiles of all reports for each hour.

A self-organizing map (SOM; Kohonen 1982) is a type of artificial neural network that facilitates the visualization (and hence the analysis) of large datasets. It can be used as a clustering method, making use of pattern recognition to cluster statistically similar groups of data according to the extraction of distinct features, which then enables the data-mining of those groups as separate entities. The technique has been widely applied to a plethora of atmospheric and oceanic observations [see Liu and Weisberg (2011) for a review of many such applications], but has thus far been relatively underused in the field of tornadic near-storm environments, with the focus remaining solely on sounding data (Nowotarski and Jensen 2013) rather than two-dimensional storm environments.

SOMs can receive a variety of inputs; in our case, the inputs consist of $n \times n$ maps of particular environmental parameters. The algorithm the SOM follows is loosely summarized as follows (for more information, see Vesanto et al. 2000):

1. Create a user-specified number M of “nodes”, which are $n \times n$ maps of randomly generated parameter values. The number of nodes will also be the final number of clusters, and is the only user-specified aspect of the entire unsupervised learning process.
2. The first input map is selected randomly from the list of input maps. This map is compared with each of the M nodes via point-by-point analysis of Euclidean distance.

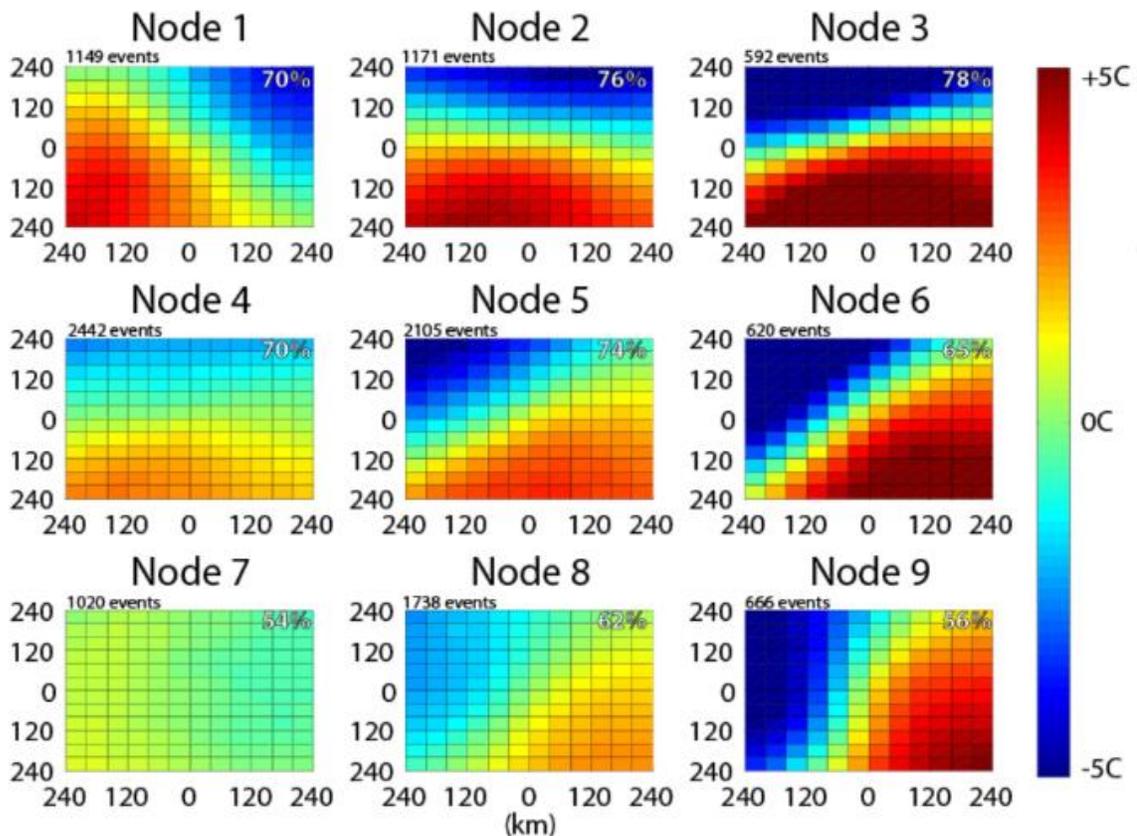


Figure 4: Self-organizing map results for 3x3 nodes of temperature anomalies. The tornado is located closest to the gridpoint (0,0) within each node; the plot shows the difference between the temperature value at a given gridpoint and the temperature value at (0,0). For example, Node 2 shows relatively cold air north of the tornado and relatively warm air south of the tornado. The number of events sorted into each cluster are listed above the top-left corner of each node. Probability of detection values for each cluster are in the top-right corner of each node.

3. Each node is “nudged” slightly by the input map; nodes that are more similar to the input map (i.e., those with smaller Euclidean distance) are more strongly nudged toward the input map values. Nodes that are less similar to the input map are not as strongly nudged.
4. A second input map is selected randomly from the list of input maps. It is then compared with each of the new nodes point-by-point, resulting in the same nudging process.
5. This process is repeated across all input maps, and then iterated several times until the nodes stabilize into M statistically distinct maps.
6. This time, when each input map is compared with the nodes, it is assigned to the cluster corresponding to the node it most closely

resembles. Each cluster can then be analyzed separately.

Thus, the input for a SOM is a large number of $n \times n$ plots of an environmental variable, and the output is a much smaller number of clusters of $n \times n$ plots of that environmental variable, each of which is summarized by its statistically distinct characteristic node.

2.3 Methodology

To create the input for the self-organizing map, we make use of 480 km \times 480 km grids of an environmental parameter at 40 km² resolution, centered on the position of each tornado report. For this simple proof-of-concept experiment, we use plots of surface temperature anomaly in K (i.e., surface temperature at a given position minus the surface temperature at the gridpoint nearest the tornado report).

The fact that users can select the number of nodes in a SOM is both a blessing and a curse: on the one hand, it can be extremely valuable to specify the complexity of the result, i.e., the final number of clusters. On the other hand, the specification process is often fairly arbitrary: a SOM that chooses too few nodes will result in the maps contained in each of those nodes having very similar statistics to the dataset as a whole, since each cluster encompasses a massive chunk of the dataset. A SOM that chooses too many nodes, on the other hand, will result in redundant nodes with functionally identical appearance (i.e., two nodes will show essentially the same map).

Through sensitivity testing, we determined that selecting $M = 9$ nodes seemed to strike the best balance between capturing important features in the data and reducing redundancy in the final nodes. We iterated the SOM process 200 times (sensitivity tests did not show appreciable differences in the nodes with higher numbers of iterations) before grouping each map into its component cluster.

3. RESULTS

Figure 4 depicts the results of our proof-of-concept SOM. Each image is a depiction of a node created by the self-organizing map, which can be thought of as a two-dimensional temperature map that is characteristic of the maps contained in that cluster. There are considerable differences in the magnitude and distributions of temperature anomalies in each node, ranging from an extremely tight temperature gradient in Node 3 to a virtually nonexistent temperature gradient in Node 7.

The orientation of the temperature gradient is also different from node to node; while Node 9 has a similar temperature gradient to Node 3, the orientation of the dividing line between relatively warm and cool air is more north-south than east-west oriented.

Probability of detection values, in outlined white font at the upper right of each node, show that the statistics of these clusters differ from the 68% average POD nationwide. In the case of Node 3, with its strong temperature gradient, we see 78% POD; Node 7's weaker temperature gradient results in a POD of only 54%. This distinction could be a function of the fact that Node 7 contains a higher percentage of EF0 tornadoes than does Node 3; these marginal cases have a lower POD.

A less immediately apparent difference is between the POD for Node 3 (78%) and Node 9 (56%); despite the fact that both nodes have similar temperature gradients, the orientation appears to make a difference in terms of forecast skill. Node 9 contains a disproportionately high percentage of tornadoes originating from QLCS storms, which could help account for the lower POD; QLCS tornadoes have an average POD of only 50%, versus 79% for RMS tornadoes (Anderson-Frey et al. 2016).

Despite the simplicity of the proof-of-concept experiment (nobody in their right mind would use only temperature anomalies to forecast tornadoes!), the SOM is still able to extract meaningful information from the data, namely the effect of both strength and orientation of local temperature gradients on tornado forecast skill.

4. FUTURE WORK

A lot of work remains to be done in the wake of this proof-of-concept experiment. Given that we have access to full mesoanalyses corresponding to each tornado report, we have barely scratched the surface of these data. Similar analyses are underway for environmental parameters of more immediate tornadic relevance, such as MLCAPE, SHR6, MLLCL, SRH1, surface dewpoint temperature, and even composite indices such as the Significant Tornado Parameter and the Supercell Composite Parameter.

We will also be able to delve more deeply into the statistics associated with each of the SOM-created clusters: what percentage of this cluster occurs at night, and how does that percentage compare with other clusters? What percentage of the events in this cluster were associated with fatalities or injuries? What can we learn about storm morphology by examining these SOM nodes? How does warning skill change as a function of the heterogeneity of the near-storm environment?

By exploring these and other questions, we will investigate the effects of the tornadic environment on tornado warning skill, and from there we will focus our attention on the parts of the parameter space that would most benefit from enhanced forecaster training and a deeper conceptual understanding.

5. ACKNOWLEDGEMENTS

The authors are grateful for the assistance from Brenton MacAloney for obtaining the verification data. This work has benefited tremendously from helpful discussions with Martin Tingley, Russ Schneider, Richard Grumm, the forecasters at the NWS State College weather office, Steven Weiss, Bill Bunting, Roger Edwards, and Paul Markowski, as well as the mesoscale research group at Penn State University. Anderson-Frey is partially supported through NSERC Postgraduate Scholarship PGSD3-462554-2012, and Richardson and Anderson-Frey's time is supported by a NOAA CSTAR Program Award NA14NWS4680015.

6. REFERENCES

- Alexander, C.R., 2010: A mobile radar based climatology of supercell tornado structures and dynamics. Ph.D. thesis, School of Meteorology, University of Oklahoma, 251 pp.
- Anderson-Frey, A., Y. Richardson, A. Dean, R. Thompson, and B. Smith, 2014: Tornado environments, metrics, and warnings: Lessons from a ten-year climatology. Preprints, *27th Conf. on Severe Local Storms*, Madison, WI, AMS, 4B.6.
- Anderson-Frey, A., Y. Richardson, A. Dean, R. Thompson, and B. Smith, 2016: *Investigation of near-storm environments for tornado events and warnings*. Manuscript in preparation.
- Benjamin, S., and Coauthors, 2002: RUC20—The 20-km version of the Rapid Update Cycle, NWS tech. proc. bull. 490. 30 pp.
- Brooks, H., J. Lee, and J. Craven, 2003: The spatial distribution of severe thunderstorm and tornado environments from global reanalysis data. *Atmos. Res.*, **67–68**, 73–94.
- Brooks, H., 2004: Tornado warning performance in the past and future: A perspective from signal detection theory. *Bull. Amer. Meteor. Soc.*, **85**, 837–843.
- Davies-Jones, R., 2015: A review of supercell and tornado dynamics. *Atmos. Res.*, **158–159**, 274–291.
- Kohonen, T., 1982: Self-organized formation of topologically correct feature maps. *Biol. Cybern.*, **43**, 59–69.
- Liu, Y., and R. Weisberg, 2011: A review of self-organizing map applications in meteorology and oceanography. *Self-organizing maps – Applications and novel algorithm design*. J. Mwasiagi, Ed., InTech, 253–272.
- Markowski, P., and Y. Richardson, 2014: The influence of environmental low-level shear and cold pools on tornadogenesis: Insights from idealized simulations. *Atmos. Sci.*, **71**, 243–275.
- Nowotarski, C., and A. Jensen, 2013: Classifying proximity soundings with self-organizing maps toward improving supercell and tornado forecasting. *Wea. Forecasting*, **28**, 783–801.
- Rasmussen, E., 2003: Refined supercell and tornado forecast parameters. *Wea. Forecasting*, **18**, 530–535.
- Rasmussen, E., and D. Blanchard, 1998: A baseline climatology of sounding-derived supercell and tornado forecast parameters. *Wea. Forecasting*, **13**, 1148–1164.
- Smith, B., R. Thompson, J. Grams, C. Broyles, and H. Brooks, 2012: Convective modes for significant severe thunderstorms in the contiguous United States. Part I: Storm classification and climatology. *Wea. Forecasting*, **27**, 1114–1135.
- Vesanto, J., J. Himberg, E. Alhoniemi, and J. Parkhankangas, 2000: SOM toolbox for MATLAB 5. Helsinki University of Technology Tech. Rep. A57, 8 pp.