A Technique For Spatial Evaluation of Severe Thunderstorm Warnings Issued by the National Weather Service

NICOLE R. RAMSEY[1,2], KIEL L. ORTEGA[2], VALLIAPPA LAKSHMANAN[1,2]

[1] School of Meteorology, University of Oklahoma, Norman, OK
[2] NOAA/National Severe Storms Laboratory, Norman, OK

ABSTRACT

In 2007, the National Weather Service (NWS) changed severe weather warnings from county-based warnings to storm-based warnings. The primary goal of this change was to reduce the area under warning in order to reduce the perception of a false alarm by people unaffected by the storm in the warned counties. While the NWS changed the way in which warnings were issued, the method of verifying warnings has remained the same. A Severe Thunderstorm Warning (SVR) is considered correct if there is at least one report of hail 25.4mm in diameter, winds 25.93 ms^{-1} or greater, reports of wind damage, such as broken tree limbs, or a tornado within the temporal and spatial bounds of the warning. This type of warning verification might assist in determining the percentage of severe weather events which were warned, but it offers no information on how much of the warning area was justified or how much of the area outside of the warning experienced severe weather. The National Severe Storms Laboratory (NSSL) has developed a number of multi-radar, multi-sensor severe weather products which could serve as a proxy in determining the accuracy and geographical specificity of severe storm warnings. The geographical specificity of NWS warnings was studied through the use of time-accumulated swaths of hail size fields and assessing the location of warning polygons with respect to storm motion and intensity.

1. Introduction

On 1 October 2007, the method used by the National Weather Service (NWS) for issuing Severe Thunderstorm Warnings (SVR) changed from a county-based warning system to a storm-based warning system. Although the warning technique changed, the method for verification remained fundamentally the same (National Weather Service Instruction 10-1601). One report of winds in excess of 25.93 ms^{-1} or hail greater than 25.4mm in diameter within the spatial and temporal bounds of the SVR verified the warning (National Weather Service Instruction 10-511). While a single report is enough for the NWS to consider a warning correct, that single report does not divulge any information about the geographical specificity, hereafter "efficiency," of the warning as a whole. Since the purpose of switching to storm-based warnings was to make the warnings more efficient, using a single report to validate a warning is a poor choice of verification method. It also lends itself to gaming – overly large polygons, though not efficient, will lead to a higher likelihood of reports.

When the NWS switched from county-based warning to storm-based warning, a new method for evaluating the placement of polygons was created. County Area Ratio takes the ratio of the area of a polygon to the area of the warned counties. This allows the NWS to study the efficiency of reducing the size of warnings through the use of polygons rather than county-based warnings (Waters 2006).

County Area Ratio provided the NWS with a way to analyze the improvements associated with shrinking the warnings from county-based warnings to polygons, but it offered no assistance when determining the accuracy of warnings. The method of point verification remains as the NWS warning verification method. One problem with point verification concerns the accuracy and volume of reports being used by NWS offices in their verification procedure. Reports are collected by NWS offices in three ways: reports from "probing phone calls" by NWS officials to citizens affected by the storm, reports sent into the NWS by local officials or storm spotters, and reports gathered from local newspapers or media outlets (Witt et al. 1998). Once a warning is verified by a single report, data collection is often halted to save time and manpower (Amburn and Wolf 1997). Another common problem with single report verification is the lack of reports in unpopulated areas. If a severe thunderstorm passes over such an area, a lack of reports will cause it to remain unverified (Amburn and Wolf 1997).

We investigated SVR efficiency through the visual study of waning placement with respect to the associated hail swaths. The National Severe Storms Laboratory's (NSSL) Warning Decision Support System – Integrated Information (WDSSII) (Lakshmanan et al. 2007) was used to display a composite of SVRs and Maximum Expected Size of Hail (MESH) swaths (Witt et al. 1998). The overlay of these two products will allow for a visual investigation of the accuracy of SVRs. Eventually, this research can be used to determine whether an automated SVR evaluation grid system, as opposed to a point verification system, can be used for SVR verification.

2. Method

An evaluation of NWS warnings was conducted by analyzing fifteen case studies from November 2008 – May 2010. This time constraint was decided based upon data available in the NSSL data archive. NSSL has a nearly complete archive of MESH swaths spanning five minute intervals for this time period. The SVR text files used to create the warning grids were downloaded from the National Climatic Data Center (NCDC).

Once data acquisition was complete, five minute MESH swaths were overlaid with one minute warning bounds to visually examine the accuracy of SVRs. If the majority of warnings on a day were efficient, the day was marked as having been warned well. An example of this occurring can be seen in Figure 1. The reflectivity (Figure 1a) shows three distinct supercells, each with a SVR. The associated MESH (Figure 1b) is well covered by warnings, with the strongest MESH in the middle of each warning. Conversely, if the MESH was unwarned or uncontained by the warning, the storm was marked as poorly warned. This can be seen in Figure 2 as the main body of low level reflectivity (Figure 2a) is located well outside the SVR and TOR, and the associated MESH (Figure 2b) is also uncontained. In this way, fifteen case studies were chosen: five decently warned days, five poorly warned days, and five intermediately warned days (Tables 1).

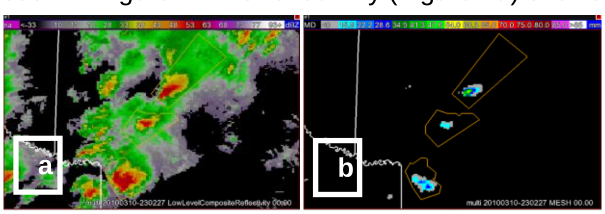

FIG 1: Correctly warned storms in AR on 10 March 2010
(a) Low level reflectivity showing three distinct supercells, each properly covered by SVRs
(b) Associated MESH and SVRs show the hail to be located within the warnings

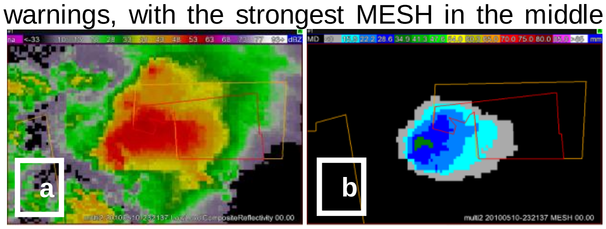

FIG 2: An incorrectly warned storm in OK on 10 May 2010
(a) Low level reflectivity of a supercell shows the SVR and TOR to the northeast of the strongest reflectivity
(b) Associated MESH and SVRs show the strongest hail located outside of the warnings

Table 1: 15 Case Studies

Starting Date	Start Time	Ending Date	End Time	Location
5 Decently Warned Days				
25-Mar-09	18Z	26-Mar-09	00Z	TX
10-Mar-10	15Z	11-Mar-10	09Z	TX AR
25-Apr-10	20Z	26-Apr-10	02Z	WV VA MD
2-May-10	21Z	3-May-10	03Z	MS AL
7-May-10	00Z	7-May-10	09Z	NE KS
5 Poorly Warned Days				
15-May-09	11Z	16-May-09	10Z	KS MO IL
2-Jun-09	17Z	3-Jun-09	05Z	TX
26-Jun-09	16Z	17-Jun-09	05Z	PA NJ NY CT
6-Apr-10	12Z	7-Apr-10	07Z	OK KS MO IA
10-May-10	18Z	10-May-10	23Z	KS OK
Intermediately Warned Days				
10-Feb-09	19Z	11-Feb-09	07Z	TX OK
5-Jun-09	20Z	6-Jun-09	08Z	WY NE TX OK
15-Jun-09	15Z	16-Jun-09	11Z	SD NE KS OK
17-Jun-09	17Z	18-Jun-09	07Z	SD NE KS MO
28-May-10	17Z	29-May-10	03Z	AL TN NC VA

Since this study focused solely on the efficiency of SVR issuance for hail, another factor in determining which case studies to use was the extent of tornadic activity and wind damage that occurred during a given day. Warnings issued solely for wind could appear as false alarm areas in a MESH-based warning evaluation grid. Tornado warnings (TOR) were also not considered, causing any hail warned by TORs to appear as a miss. The storm reports for each day can be seen in Figures A1 – A3 in the appendix. Two examples of inopportune days are 10 May 2010 when a tornado outbreak occurred simultaneously with the hail reports and 28 May 2010 where the hail producing storms had associated wind damage.

Each day was then examined for warning efficiency. MESH less than 15 mm in size was discarded from the database during data processing. This threshold was chosen because hail reports tend to be associated with MESH values above this threshold (Ortega et al. 2009). During data processing, a composite of MESH swaths and warning bounds was created for each individual warning. Next, a composite for all MESH swaths (Figure 3a) and warning bounds (Figure 3b) was produced for each day. The MESH was then overlaid with the warning bounds and each grid point (pixel) in the final verification grid was colored based upon whether the value of the MESH grid at that pixel was above 15mm and whether or not the pixel was within a warning polygon (Figure 3c). Pixels that were warned and experienced hail were colored green to indicate a hit. Pixels that were unwarned but experienced hail were colored red to indicate a miss. Lastly, pixels that were warned but had no hail were colored yellow to indicate a false alarm. These files were displayed using WDSSII.

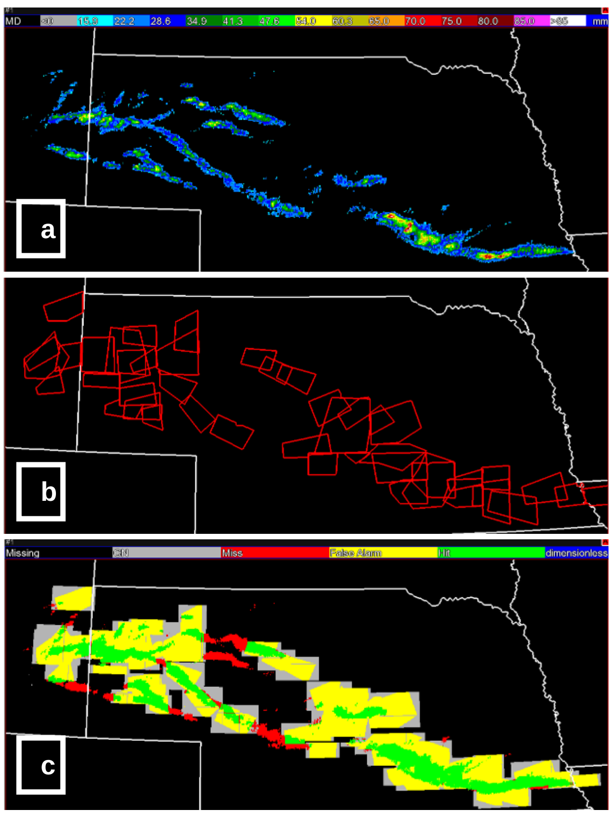

FIG 3: Daily composites for NE on 5 June 2009
(a) Composite of all MESH present that day
(b) Composite of all SVRs issued that day
(c) Verification grid showing the daily composite of hit areas where MESH was properly warned (green), miss areas where MESH was unwarned (red), and false alarm areas where a warning existed but no MESH occurred (yellow)

SVR efficiency was initially studied by comparing individual warning bounds to the associated MESH tracks. The direction, size, orientation, and overlap of warnings were the characteristics examined in the hopes of determining SVR

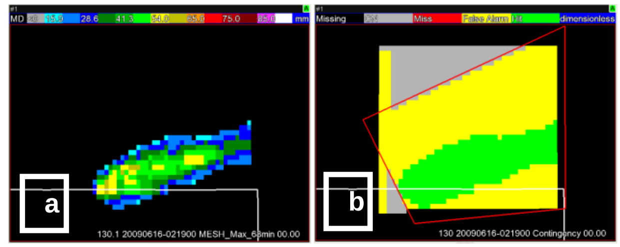

FIG 4: A case example of a proper warning in OK on 15 June 2009
(a) MESH swath from storm during the life of the warning
(b) SVR warning bound (red) surrounding the warning composite. The MESH crosses directly through the warning with a boundary of false alarm on either side of the MESH swath and no areas of miss

efficiency. Figure 4 shows a decently warned storm in Oklahoma on 15 June 2009. In this example, the MESH predicts hail sizes up to 60.3 mm, which warrants a SVR. The MESH swath (Figure 4a) is shown to cross directly through the SVR warning, leading to zero miss areas and a marginal false alarm area on either side of the MESH swath (Figure 4b). Although zero miss area is preferable, a marginal miss area at the start of a warning is still acceptable since the NWS currently uses a warn-on-detection forecast method (Stensrud et al. 2009).

After individual SVRs were investigated for warning efficiency, the daily composite of MESH, warnings, and hits, misses, and false alarm areas were studied. Just like with the individual warnings, a storm in Missouri on 17 June 2009 was chosen to act as a proxy for differentiating properly warned storms from poorly warned storms. Figure 5a shows the associated

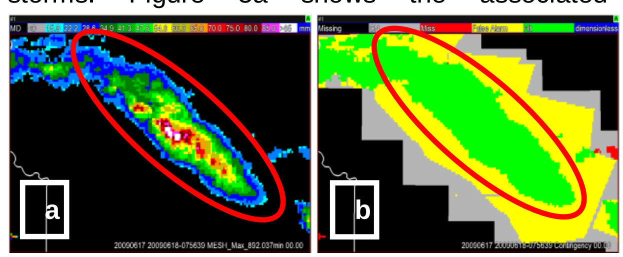

FIG 5: A case example of a properly warned storm in MO on 17 June 2009
(a) The MESH swath from the storm is circled in red. All other MESH was caused by other storms.
(b) The hit, miss, and false alarm area caused by the MESH. The MESH crosses directly through the warnings with a boundary of false alarm on either side of the MESH swath and no areas of miss.

MESH swath, with MESH values well above the minimum threshold for SVRs. Figure 5b shows the daily composite image for that storm. With no area of miss and a small area of false alarm surrounding the hit area, or the MESH track, it is very evident that the storm passed through the center of the SVRs and was very well warned.

Finally, the differences evident between the individual warning composites and the daily composites were investigated. This comparison looked to differentiate between the information that could be gained from individual warning composites and the information that could be gained through the study of daily or storm composites. Inevitably, this last step assisted in determining whether the evaluation of NWS warnings would be more informative when conducted on daily basis as opposed to a point verification procedure.

3. Individual Warning Evaluation

One of the leading problems in warning issuance is storm growth. As a storm begins to develop and produce MESH, a warning is issued. If the storm's growth becomes stunted and the hail dissipates, much of the warning remains unverified. An instance of this occurring is on 6 April 2010 in Missouri where Figure 6a shows the end of the storm's hail producing life. The single warning contingency (Figure 6b) shows a warning being issued on the remaining MESH from that storm. The beginning of the warning verifies, due to the initial MESH present within its bounds, but the remainder of the warning is marked as a large false alarm area since the MESH diminishes and fails to cross through to the other side.

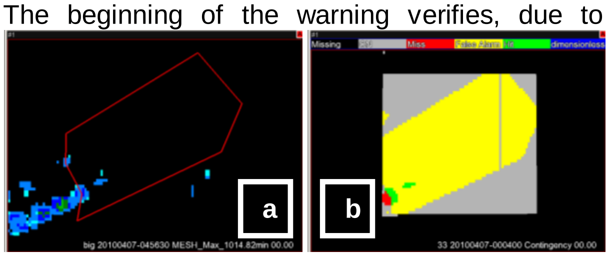

FIG 6: A dissipating storm led to a large area of false alarm in MO on 6 April 2010
(a) Dissipating MESH swath with SVR bound shown in red
(b) Warning contingency showing a small area of miss at the beginning of the warning followed by a large area of false alarm due to the dissipating storm

Studying individual SVRs not only displays the effects of dissipating storms on warning efficiency, but it also amplifies inaccurate warning placement due to storm motion problems. On 2 June 2009, storms in Texas were moving on a northwest to southeast axis, but the associated SVR for one particular cell was oriented on a west to east axis (Figure 7). The hail quickly exited the warning creating a large false alarm area in the eastern part of the warning as well as a miss area to the south of the warning.

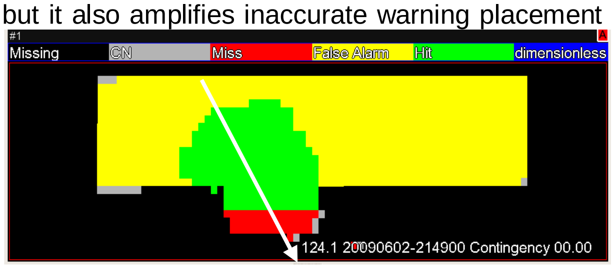

FIG 7: Incorrect storm motion affecting warning placement in TX on 2 June 2009

A third example of a misplaced warning can be seen in Pennsylvania on 26 June 2009. MESH values ranging from 22.2 mm to 34.9 mm (Figure 8a) are large enough to warrant a SVR, but the issued warning misses the main body of MESH to the north. This misplacement of the SVR can very clearly be seen in the individual warning composite (Figure 8b). Because of this error, a substantial area of false alarm is present to the north as well as a large area of miss to the south. In order to ensure that this warning was issued for hail, not wind, official storm reports were searched for wind damage occurring within the warning bounds. As expected, all of the surrounding wind reports were located to the south of the warning, along the path of the hail (National Climatic Data Center).

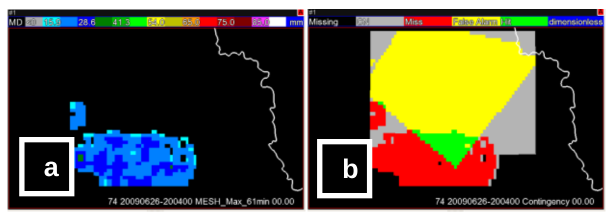

FIG 8: Misplaced warning in PA on 26 June 2009 results in a large miss and false alarm.
(a) MESH from storm during the duration of the warning
(b) Areas of miss and false alarm associated with the misplaced warning

4. Daily Composite Evaluation

On 2 May 2010, two storms crossed Mississippi, eventually dissipating in Alabama. Figure 9a shows the composite image for both storms. This storm composite contains no areas of miss and a limited area of false alarm, all indicators that the storms were well warned. Figure 9b shows the MESH swaths from each storm while Figure 9c shows the warnings issued for the storms. As one can clearly see, the SVR orientation matches the motion of the storms, rarely overlap, and closely fit the MESH. While the MESH is well covered, a bias in the warning placement of the southerly storm with respect to the MESH is evident. The comparison example for a properly warned storm (Figure 5b) shows an equal area of false alarm to the right and to the left of the MESH track. An equal amount of false alarm on either side of the MESH swaths is not apparent on the Mississippi storms, as the MESH frequents the left most boundary of the SVRs. This leaves a large false alarm area on the right hand side of the warning.

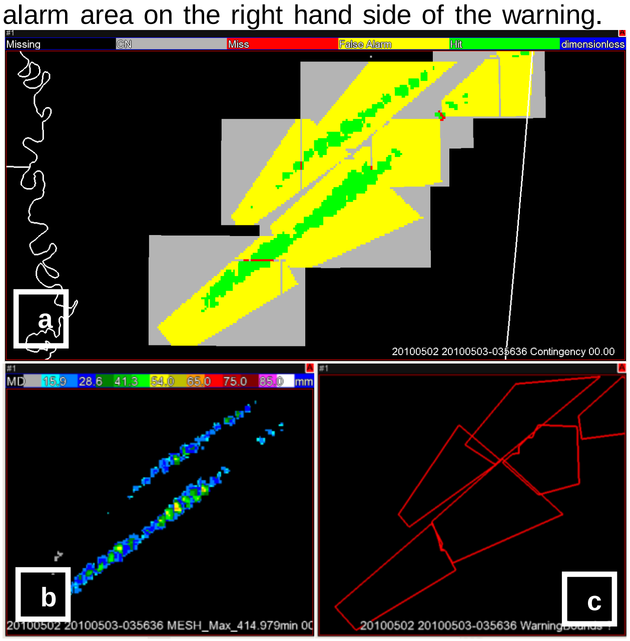

FIG 9: Hail producing storm in MS on 2 May 2010 shows warning placement bias with respect to the MESH track.
(a) Daily composite shows no areas of miss and a relatively small false alarm area, but a bias in the SVRs with respect to the MESH swath
(b) MESH swaths and (c) warnings bounds for both storms

A similar example of a properly warned storm is a storm that crossed central Texas on 25 March 2009. The main body of the storm is well warned (Figure 10a), and the extraneous warnings to the north are explained by storms that passed through the area later in the day. While there are minimal areas of miss to the south of the storm, only a small area of false alarm is located along the storm path itself. The warning bounds for this storm (Figure 10b) follow the correct path of the MESH (Figure 10c), but is also beneficial to note that the storm motion was most likely slower than NWS forecasters originally expected since the start of each new SVRs frequently overlapped with the end of previous warnings.

The warning circled in red (Figure 10a) indicates an extreme example of a warning issued to perfectly fit the storm path, as indicated by the lack of false alarm surrounding the MESH. Many times a large false alarm area is indicative of a poor warning, but a negligible false alarm area is also undesirable. SVRs are issued to alert the public of impending severe weather (National Weather Service Instruction 10-511). Giving no distance between the severe hail and the warning bounds will not adequately warn the public of the potential

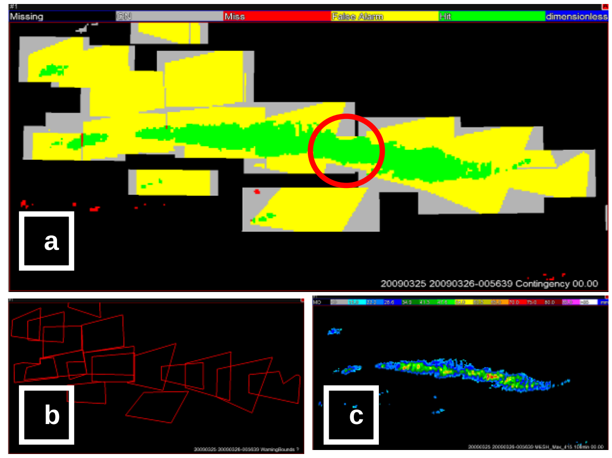

FIG 10: Hail producing supercell in TX on 25 March 2009 that showcases a SVR issued to close to the MESH swath.
(a) Composite of MESH and SVRs indicate minor miss areas and minimal false alarm areas
(b) Composite of SVRs
(c) MESH swath present for this storm

threat to their life and property, especially if the path of the storm changes slightly. If a person's location is just outside of a warning bound, they are likely to take different precautionary measures then they would if they were located within the warning. Thus, if the MESH had deviated ever so slightly in either direction, people's safety may have been compromised.

On 25 April 2010, a cluster of hail producing cells moved across northern Virginia and Maryland. As with the previous two daily composites, the daily composite (Figure 11a) shows multiple areas of MESH fully covered by SVRs. The few miss areas present in eastern Maryland are an error of the program caused by a large number of vertices in a small area. In actuality there were no areas of miss on this day. Examining the associated MESH swaths (Figure 11b), it becomes clear that these storms were not excessively strong and were incapable of producing a substantial and unbroken MESH swath. Although the MESH swath supported the existence of relatively weak storms, there is a number of overlapping warnings (Figure 11c) across West Virginia and Virginia. Due to this discontinuity between the strength of the MESH and the abundance of the warnings, official storm reports were investigated for severe wind reports within the warning bounds (National Climatic Data Center). Unfortunately, there were no reports of thunderstorm winds, and therefore the abundance of overlapping warnings was deemed unnecessary given the weakness and brevity of the storms.

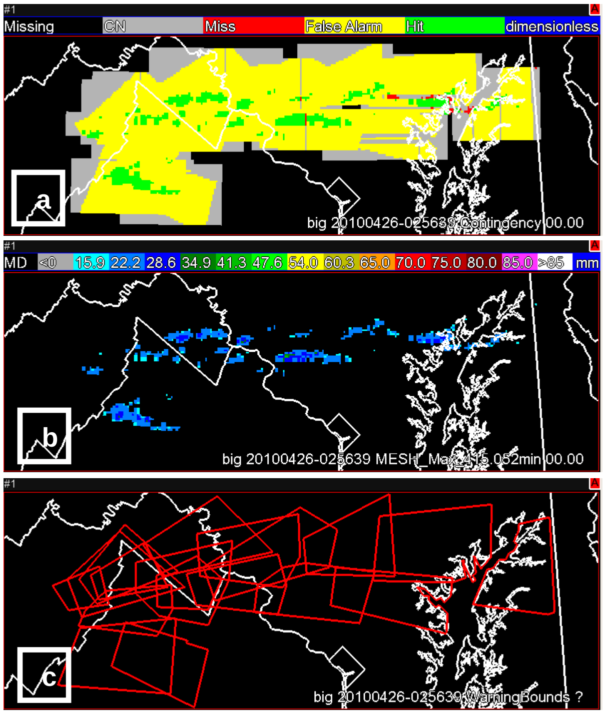

FIG 11: Overabundance of warnings on a storm system in WV, VA, and MD on 25 April 2010
(a) Daily composite of MESH and SVRs shows decent coverage of storms
(b) MESH swath from storms shows marginal MESH and broken MESH swaths
(c) SVR contingency showing an abundance of warnings

15 May 2009 was the worst day out of the 15 cases studied. Misplaced warnings plagued the central United States, and the day ended with large areas of miss and false alarm across much of eastern Kansas and Missouri. Figure 12a shows the overall hit, miss, and false alarm areas with Figure 12b indicating the numerous hail swaths produced by the storms. Official storm reports confirmed numerous reports of hail and thunderstorm winds confirming that these storms were poorly warned.

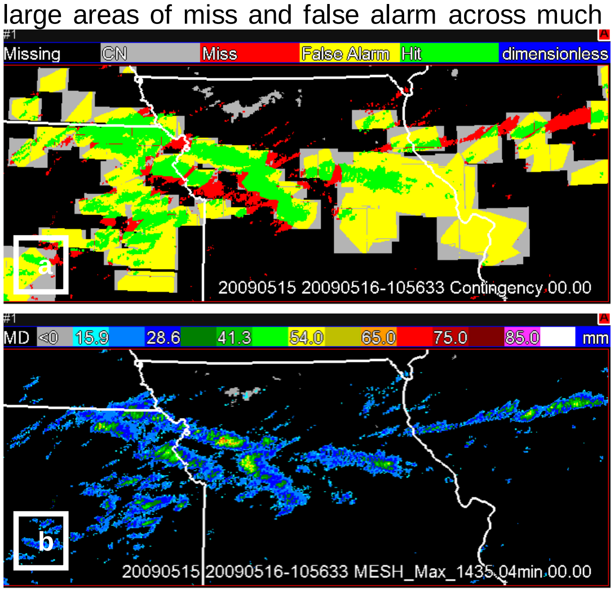

FIG 12: Poorly warned storms in KS and MO on 15 May 2009 indicated by an large area of miss
(a) Daily composite showing a large area of miss
(b) MESH swaths from associated storms

5. Limitations of SVR Grid Evaluation

The leading limitation of warning evaluation through the use of daily composites is the inability to recognize the underlying errors within individual warning bounds. Multiple days, as indicated in Figures 6–8, present an error with individual warnings, including inaccurate storm motion or overall misplacement. If the daily composite indicates an accurately warned day, but the individual warnings show underlying problems, then the NWS would receive credit for warnings that were incorrect. A prime example of this occurring can be seen on 10 February 2009 in central Oklahoma. In the daily MESH composite (Figure 13a), it is evident that the storms are moving in a southwest to northeast direction. This indicates that warnings should be oriented in a similar fashion. The daily composite (Figure 13b)

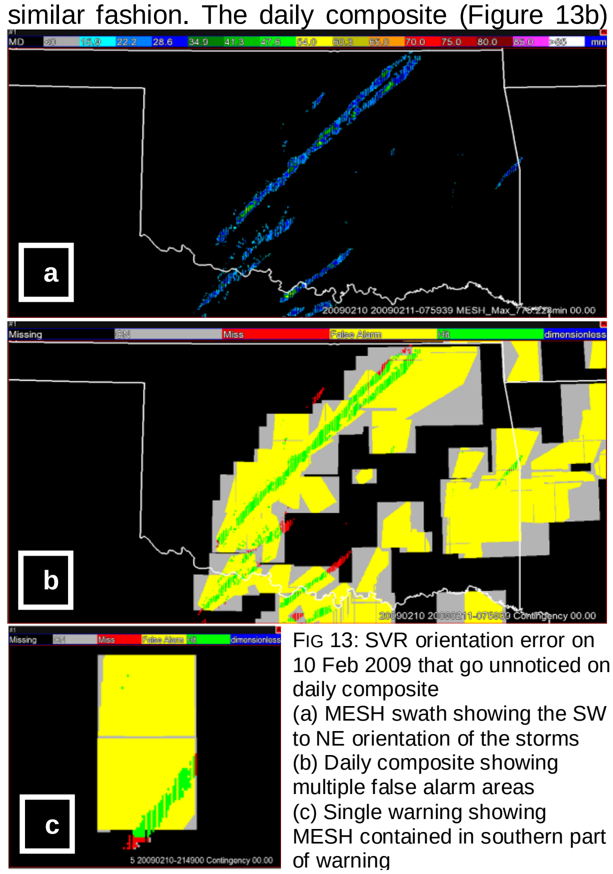

FIG 13: SVR orientation error on 10 Feb 2009 that go unnoticed on daily composite
(a) MESH swath showing the SW to NE orientation of the storms
(b) Daily composite showing multiple false alarm areas
(c) Single warning showing MESH contained in southern part of warning

indicates that the MESH was well covered with warnings. Upon further investigation, Figure 13c shows a north to south oriented warning. Strangely, the MESH swath is still contained within this warning, albeit in the southeastern most corner, even though the warning orientation is incorrect. Indications of these poorly oriented warnings can be seen across much of central Oklahoma, as a larger area of false alarm to the north of the MESH than to the south indicated a track bias, or in this case, orientation bias. This, along with the storms in Texas on 2 June 2009 (Figure 7), is a key example of the daily composite obscuring more subtle problems with NWS SVRs. Had the individual warnings never been studied, the warning orientation problems would never have been discovered, and the warnings would have been verified without further investigation. This limitation presents an encumbering issue when using an automated warning evaluation grid based solely upon the daily composite.

Similarly, if the daily composite denotes a substantial area of miss or false alarm, the cause for these errors will go unknown unless further investigation of individual warnings is conducted. Analyzing the daily composites of hits, misses, and false alarms for each case was found to be a generalized indicator of properly or poorly warned days, but unlike the study of individual warnings, the specifications of warning errors cannot be determined in the daily composite. As Figure 12a showed, large areas of miss were present in Missouri on 15 May 2009. Very little was actually known about these areas of miss by studying only the daily composite. These areas of miss could very easily have been a case example on limitation of the program caused by the lack of TOR warnings in the program, as discussed earlier. It was not until further investigation of individual warnings that the cause for these errors was blamed on inaccurate warning issuance.

Figure 14 displays five cases of inaccurate storm warnings that represent the warning errors present that day. Figures 14a–14b depict MESH forming before the warning was issued, leading to a miss on the southern edge of the warning. The MESH continued to extend past the northern edge of the warning boundary, leading to a secondary area of miss. Figures 14c–14d show MESH exceeding the warning bounds, leading to areas of miss along the sides of the warning. If the MESH present outside of the warning was beneath SVR criterion, then a warning was not necessary and the miss area would be written off as a limitation of the technique. The MESH swaths quickly confirm that MESH greater then 25.4mm were located outside

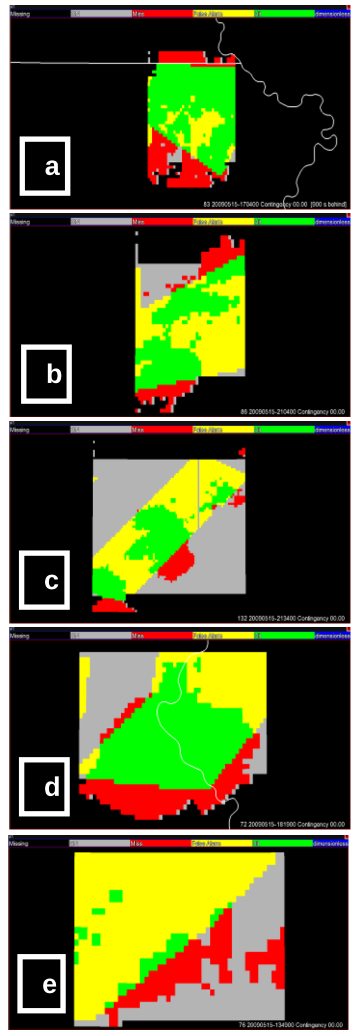

FIG 14: Incorrect SVRs
(a) and (b) MESH forms before the SVR is issued and extends past the northern boundary
(c) and (d) MESH outgrows SVR
(e) MESH forms to the SE of the warning

of the warning, warranting a larger SVR. Finally, Figure 14e shows a warning placed to the north of the MESH, leading to a miss area on the south side of the warning. The daily composite provided a generalized image of the day's warning accuracy, but it was unable to discern what went wrong and why.

Another limitation with the daily composite is its inability to account for both the excessive overlap of SVRs and the reason for a SVR's issuance. The program used to create the daily composite creates one cell of either false alarm, hit, or miss regardless of how many warnings were overlapping in that area. This can been seen in Figure 11 where the daily composite (Figure 11a) shows properly warned storms but the warning composite (Figure 11c) shows an excessive overlap of warning bounds. Examining only the daily composite masked this overlap and led to an initial indication of warning accuracy.

Details currently in warning texts are unable to differentiate between SVRs issued for wind and those issued for hail. Thus, when warnings issued only for wind are ingesting into the daily composite, they appear as false alarm areas due to the lack of MESH. Even if the warnings were verified by reports of severe wind, it still appears as a false alarm area since the program is unable to account for the wind damage. A permanent fix for this problem requires further investigation since the complete removal of all warnings issued solely for wind would eliminate a substantial portion of SVRs, causing this method of SVR evaluation to be insufficient.

6. Benefits for SVR Grid Evaluation

Many factors influence the accuracy of SVRs, causing a majority of warnings to share limited similarities with Figure 2. Studying individual warning composites allows the associated error to be investigated. Figures 6-8 all display individual cases of SVR placement inaccuracy. It can easily be discerned from each warning composite what went wrong. Studying individual warning composites allows for this in depth analysis of SVRs. Figure 14 also showcases the benefits of studying individual warning composites. It was evident from the daily composite in Figure 12 that imprecise warnings were issued throughout the day. The cause for such inaccuracies was not known until individual warnings were studied (Figure 14). These individual warning composites showed precise reasons for the large area of miss that was present that day.

While individual warning composites assist in the location and diagnosis of poor SVR warnings, daily composites provide useful information concerning the day's overall warning accuracy. Figure 9, showing the tack biases in the SVRs, proved that information can be lost when studying only individual warning composites. If the daily composite had not been studied, this bias would have gone unnoticed on a warning to warning basis. Studying the daily composite also clearly displays cases of overlapping warnings. Figure 11c shows a plethora of warnings overlapping in Virginia and West Virginia. Individual warning composites would be unable to account for overlapping warnings, and thus the daily composite provides more reliable data concerning the efficiency of SVR warning placement during a particular day.

7. Summary and Conclusion

Point verification is an inappropriate method for determining the accuracy of SVRs. It says very little about the accuracy of warnings, and it can treat, as correct, SVRs that were poorly located or oriented. Studying the efficiency of SVRs on a daily scale offers a more detailed analysis concerning the placement of warnings with respect to MESH swaths. Although it offers limited information concerning the cause for inefficient warnings, it presents a big picture look at the warning biases for that day. If further investigation is required, a closer look at each individual warning is also possible.

8. Bibliography

Amburn, S.A., and Wolf, P.L., 1997: VIL Density as a Hail Indicator. *Wea. Forecasting,* **12,** 473-478.

Lakshmanan, V., T. Smith, G. J. Stumpf, and K. Hondl, 2007: The Warning Decision Support System - Integrated Information (WDSS-II). Weather and Forecasting 22, 592-608.

National Climatic Data Center, cited 2010: SRRS Text Products/Bulletin Selection. [Available online at http://.has.ncdc.noaa.gov.]

Ortega, K.L., T. M. Smith, K. L. Manross, K. A. Scharfenberg, A. Witt, A. G. Kolodziej, J. J. Gourley, 2009: The Severe Hazards Analysis and Verification Experiment. *BAMS,* 1519-1530.

Stensrud, D. J., M. Xue, L. J. Wicker, K. E. Kelleher, M. P. Foster, J. T. Schaefer, R. S. Schneider, S. G. Benjamin, S. S. Weygandt, J. T. Ferree, J. P. Tuell, 2009: Convective-Scale Warn-On-Forecast System – A Vision for 2020. *BAMS,* 1487-1499.

Storm Prediction Center, cited 2010: NWS Local Storm Reports. [Available online at http://www.spc.noaa.gov.]

The National Weather Service, cited 24 February 2009: NATIONAL WEATHER SERVICE INSTRUCTION 10-1601 – Verification Procedures. [Available online at http://www.weather.gov/directives/sym/pd01016001curr.df .]

The National Weather Service, cited 19 April 2010: NATIONAL WEATHER SERVICE INSTRUCTION 10-511 – WFO Severe Weather Products Specifications. [Available online at http://www.weather.gov/directives/sym/pd01005011curr .pdf.]

Waters, K., cited 2006: Definition of Severe Weather Metric "County Area Ratio". [Available online at http://www.weather.gov/regsci/gis/CAR.pdf.]

Witt, A., Eilts, M.D., Stumpf, G.J., Mitchell, D., 1998: Evaluating the Performance of WSR-88D Severe Storm Detection Algorithms. *Wea. Forecasting,* **13,** 513-518.

9. Appendix

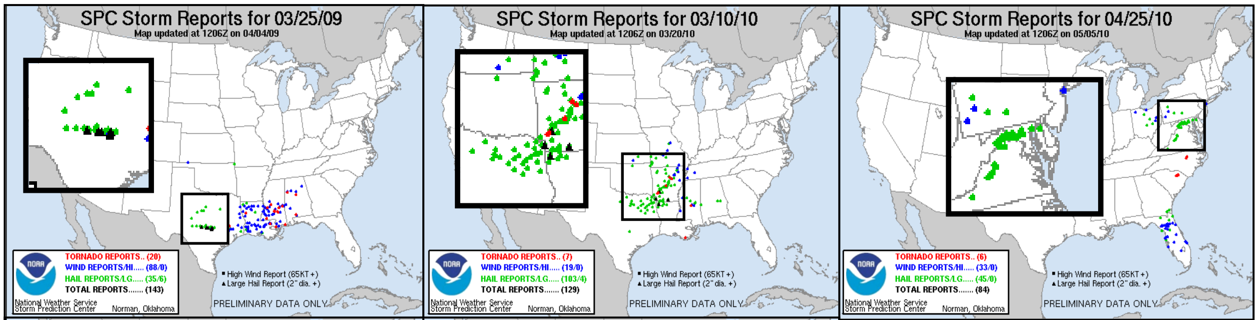

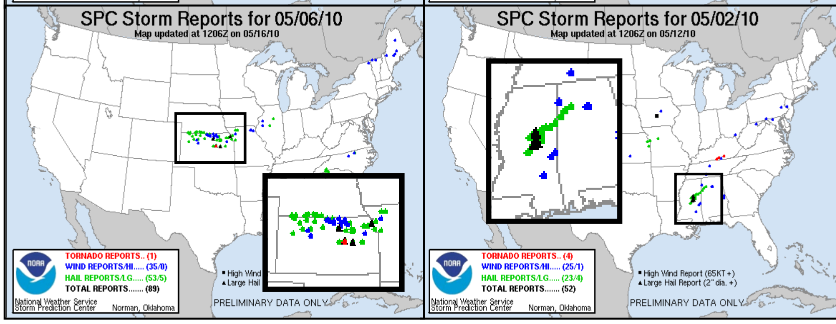

FIG 1A: 5 Cases for Decently Warned Storms. Each image contains the days overall reports and an outlet displaying the specific storm or set of storm studied in the data set. Images from spc.noaa.gov

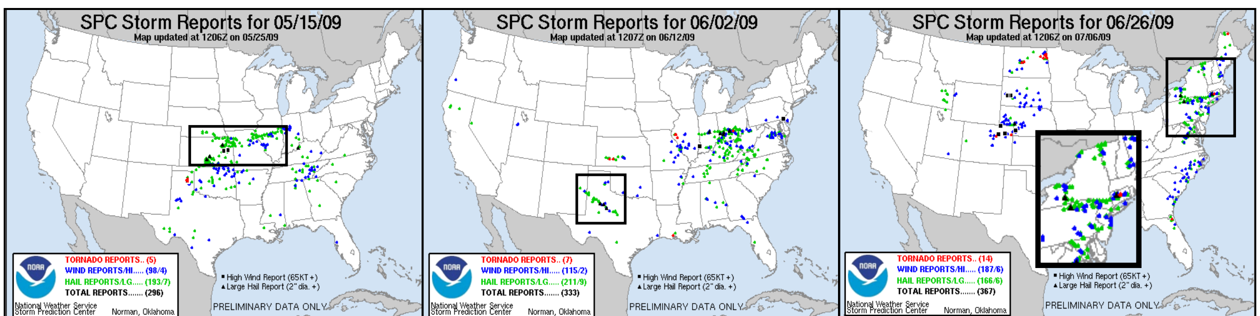

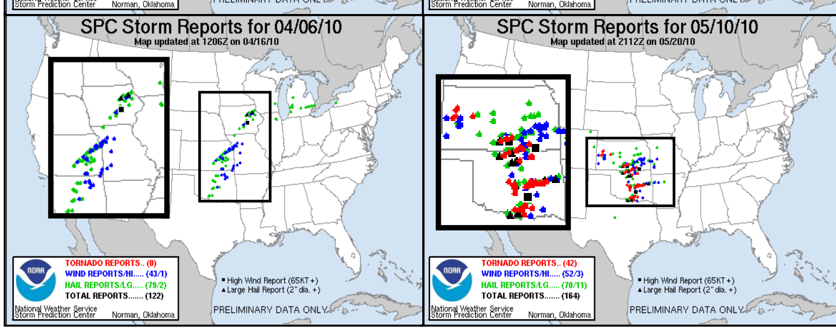

FIG 2A: 5 Cases for Poorly Warned Storms. Each image contains the days overall reports and some contain an outlet displaying the specific storm or set of storm studied in the data set. Images from spc.noaa.gov

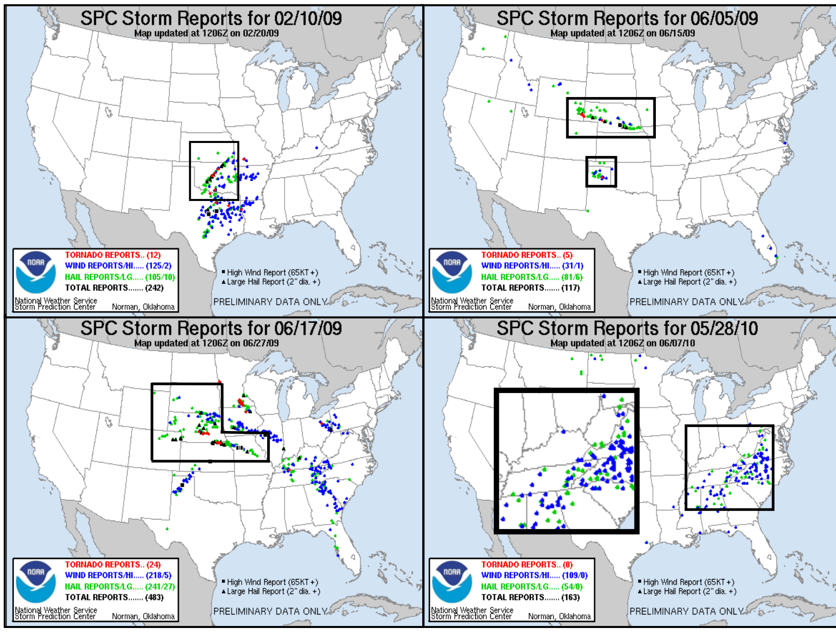

FIG 3A: 5 Cases for Intermediately Warned Storms. Each image contains the days overall reports and some contain an outlet displaying the specific storm or set of storm studied in the data set. Images from spc.noaa.gov