Verification of WRF Experiments Over the Southeast United States

Peter Bogenschutz
University of Utah, Department of Meteorology*,* Salt Lake City, UT

1) Introduction

The purpose of the Coastal Storms Initiative (CSI) project is to lessen the impacts of storms on coastal communities. In 2003 the pilot project of the local modeling component involved the first known installation of the Weather Research and Forecast (WRF) modeling system at a National Weather Service (NWS) Forecast Office (WFO) (Shaw et al 2003, Welsh et al 2003). It was shown that the WRF framework, utilizing the Eulerian based Advanced Research Weather (ARW) core, improved twenty-four hour forecasts of precipitation, sea breeze transition detection and propagation, and visibility compared to the traditionally used Eta 12 km model (Bogenschutz et. al. 2005)

This study furthers verification work for the CSI local modeling component by assessing a variety of experiment WRF configurations. The first experiment tests the value of running a WRF configuration utilizing a larger domain compared to one with substantially smaller boundaries. The second experiment tests the difference in forecast accuracy between a WRF configuration initialized through local data versus one initialized using forecasts from the Eta model. Statistical evaluation is performed using Forecast System Laboratories (FSL) Real Time Verification System (RTVS), while object oriented verification techniques are performed using the Ebert & McBride Technique (EMT) (Ebert and McBride 2000) for precipitation and a modified version of the Contour Error Mapping (CEM) (Case et al. 2002) method for sea breeze verification.

2) Experiment Setup

a) Big / Small Domain Experiment

The purpose of the Big/Small Domain (BSD) experiment is to assess any potential advantages in running a WRF configuration with a larger domain as opposed to a much smaller domain. Both configurations are identical with the only difference being the extension of the lateral boundaries. The larger domain (hereafter referred to as the 'WRF-CONUS') covers all of the continental US with portions of Mexico and Canada. The smaller domain (hereafter referred to as the 'WRF-CSI') covers north Florida as well as portions of southern Georgia and South Carolina. The WRF-CSI configuration covers the forecast responsibility area for the JAX WFO.

Each configuration utilizes a 5 km grid resolution in each horizontal direction and 38 full vertical levels. It is important to clarify that the WRF-CONUS configuration does not incorporate nesting and that the horizontal grid spacing is 5 km throughout the entire domain. Likewise, the domain for the WRF-CSI, is the mother domain and not a product of a nest from a larger domain. Therefore, given the high resolution and much larger boundaries for the WRF-CONUS, a substantial improvement in forecast capability for precipitation and sea breeze transition, amongst other variables, would have to be witnessed to constitute the practical worth for possible operational implementation.

b) Hot Vs. Cold Start Experiment

The purpose of the Hot Vs. Cold Start (HCS) experiment is to test the difference in forecast capabilities when two exact configurations are initialized differently. This experiment utilizes the operational WRF forecasts provided by the JAX WFO. The horizontal domain for this experiment is very similar to the domain used in the BSD experiment, although capturing a larger portion of the Florida panhandle and eastern Alabama. The "hot-start" simulation is initialized using FSL's Local Analysis and Prediction System (LAPS) (Albers et al 1996) 10 km grids (this simulation hereafter referred to as 'WRF-LAPS'). The "cold-start" simulation is initialized using the six hour forecast from the 00 UTC run of the Eta 218 grids (this simulation hereafter referred to as 'WRF-Eta'). Bogenschutz et al (2005) found minimal differences in forecast skill between the WRF-LAPS and WRF-Eta simulations during a brief examination of the 2003/2004 verification of the local modeling CSI project, and this experiment simply expands that research by including more forecast cases.

3)Methodology

The verification period for the BSD experiment runs from April 1 – July 30, 2005 while the verification period for the HCS experiment is from May 1 – July 30, 2005. Verification periods are segregated by 24 hours. Therefore, at times, the BSD experiment verification will compare skill of the 1-24 vs. 25-48 forecast hours for the WRF-CONUS and WRF-CSI. Since the WRF configurations for the HCS experiment only forecast out to 24 hours, obviously there is only one period of verification for this experiment.

a) Real-Time Verification System

A statistical verification approach is performed utilizing FSL's Real-Time Verification System (RTVS) for surface and upper-air sensible parameters (Mahoney 1997). The RTVS configurations for the BSD and HCS experiments both differ in terms of the observational data set used\ because the RTVS for the BSD experiment utilizes the WRF Verification System developed by NCEP and FSL. The RTVS for the BSD utilizes metar sites to serve as surface observations while the RTVS for the HCS uses madis observations. The WRF Verification system uses radiosonde observations for upper air verification whereas the HCS RTVS does not perform upper-air verification. The RTVS for both experiments is primarily used to verify temperature and wind speed. Neither RTVS configuration is used to verify precipitation.

b) Precipitation Verification

Rather than focusing on the traditional statistics, the Ebert & McBride Object Oriented Technique (EMT) is used to verify precipitation. The EMT introduces the idea of a Contiguous Rain Area (CRA), which acts as the union of the forecast and observation rain field, set by a user defined threshold. A 0.25" threshold is used for the purposes of this study. Therefore, instead of verifying precipitation forecasts by individual grid points, the EMT allows us to verify precipitation forecasts by the actual entities. The EMT verifies entities assuming there is no displacement error, which is ideal for assessing warm season Florida forecasts. After the best fit between forecast and observations has been determined, systematic error decomposition due to volume, displacement, and pattern can then be found for the forecast entity.

For a CRA to count as a 'hit' for the model, the shifted forecast entity must be correlated at the 95% confidence interval and meet a size and mass criteria set relative to the observed entity. This is to ensure that CRAs which bear little resemblance, or forecast by pure chance, are not counted as a hit. Should a CRA be forecast and observed but does not meet the aforementioned criteria it is either classified as an 'overforecast', 'underforecast', or 'significant pattern error' CRA depending on the bias score and ratio of forecast to observed grid points. The EMT is not applied to these CRAs, however they are segregated from the typical 'false alarm' and 'miss' CRAs to exhibit that there was limited skill in forecasting those rain entities. The data set used as observations is Stage IV NCEP precipitation 6 hour accumulations.

c) **Sea Breeze Verification**

Verification of the WRF-ARW forecasts in JAX during the 2003 and 2004 warm seasons utilized Contour Error Mapping (CEM) (Case et al. 2002) to verify sea breeze transition and propagation. Due to the relatively coarse temporal and spatial resolution of the model and observational data, a methodology based off the CEM is used to verify only sea breeze transition and not propagation. Model forecasts are compared to mesonet and metar observations along the coastlines of JAX, Tampa, and Cape Canaveral. If a sea breeze is detected through observations and forecasts, it is counted as a hit for the model. Each hit, miss, and false alarm is subjectively looked at to ensure fair results and if there is any question relating to the category of a modeled or observed sea breeze the case is thrown out. In addition, synoptic weather maps and temperature gradients are observed on days when a sea breeze is detected to ensure that a true detection is being assessed. For each sea breeze hit, statistics for temperature, wind speed, and wind direction are also examined to gain a better understanding of model behavior for these sea breezes.

4) Results

a) Big / Small Domain Experiment

i) Statistical Evaluation

An examination of the RMSE of temperature for each forecast hour for both set-ups in the BSD experiment yields similar results for the entire period (April 1st – July 30th). While the WRF-CONUS appears to hold a slight advantage over the WRF-CSI for most of the early forecast hours, the differences are so small that they may not be statistically significant. A diurnal cycle in the errors is witnessed, with both simulations experiencing maximum RMSE during the daytime hours and errors in the second day being more considerable. Both configurations exhibit the same general pattern in terms of temperature bias during the first 36 forecast hours. Thereafter the trend for the WRF-CONUS is to exhibit a cool temperature bias during the time of maximum heating of the 24-48 hour forecast cycle whereas the WRF-CSI does not experience this bias.

ii) Ebert McBride Verification

Through the Ebert and McBride Technique (EMT), for precipitation verification, it was determined that 452 CRAs were detected through observations for the periods examined. It should be noted that only days in which both the WRF-CONUS and WRF-CSI grids are available are included in this tally. In addition, it should be kept in mind that two accumulation periods for the models are being examined, so each 24 hour accumulated precipitation grid is assessed twice for CRAs. Therefore, there were actually 452/2 CRAs detected in true observations.

Both simulations have very comparable overall detection rates, with the WRF-CSI having a slight advantage with 46.2% correlated detections compared to the WRF-CONUS detection rate of 45.4%. In terms of systematic errors for each simulation, the difference lies in the displacement category, in which the WRF-CSI exhibits a 10% higher score compared to the WRF-CONUS. While the WRF-CSI may exhibit a slightly better detection score, an examination of CRA statistics shows that the WRF-CONUS generally has a higher average correlation coefficient (for both unshifted and shifted forecasts), lower RMSE, more comparable average rain rates compared to observations, and better lat/lon displacement. This says that while the WRF-CSI is performing slightly better at forecasting the number of CRAs, the WRF-CONUS has an advantage at more accurately representing the placement, pattern, and volume of the precipitation entities. One might expect that with a larger domain the WRF-CONUS would perform better in terms of detection, especially with synoptic forced precipitation. A more detailed analysis of the results, including examinations of the 1-24 hr and 25-48 hr accumulations for both simulations and detection rate by the size of the entities will be performed. In addition, any interesting cases will be examined in which will help characterize the tendencies in the model.

First an examination of the two different accumulation periods is assessed for each simulation (hereafter, the 1-24 hr accumulation period will be referred to as "period one" and the 25-48 hr period will be referred to as "period two"). The CRA statistics for period one yields similar results to that of the entire period study, with the WRF-CONUS having a slight advantage over the WRF-CSI in every category. Interestingly enough, the WRF-CONUS also outperforms the WRF-CSI in terms of detection, although the differences are very small. This leads to the conclusion that the WRF-CSI detects more CRAs than the WRF-CONUS during period two. Some interesting features here is the relatively high false alarm rate from the WRF-CONUS and, while not shown for the sake of brevity, the statistics which are still very comparable between the two simulations.

The fact that the WRF-CSI actually detects more CRAs than the WRF-CONUS during period two is puzzling due to the boundary limitations for this smaller domain. If a frontal structure moved through the domain during the second period, logically one would think the WRF-CONUS would be the better model for the case, since most likely the feature was initialized within the model domain at some location. These surprising results then require further investigation of the CRA statistics in order to determine model discrepancies.

Although both models have roughly the same number of overall detections during the second period, a simple correlation is run for both periods to see if the same number of detections are occurring on the same day. The correlation coefficient for period one yields a score of 0.97, indicating that the number of detections for each day are very similar for each simulation. This is also mirrored in the average correlation coefficient for 24 h accumulated precipitation grids between both simulations, with a score of 0.70. These correlation scores for the second period are quite different. The correlation testing the number of detections for each day yields a score of 0.80, while the correlation between the simulation accumulation grids for period two is 0.45. This decrease in correlation between the two periods for these parameters suggests that while the accumulated precipitation grids are similar for the first period, they differ somewhat in the second period.

A comparison on the skill of each simulation on the detection of different size CRAs is conducted. In this analysis the CRAs are divided into three categories; pop-up (very small CRAs, less than 100 grid points in size), mesoscale (less than 1000 but greater than 100 grid points), and synoptic (1000 or greater grid points) convection. One of the most notable features in this analysis was the high miss rate for all simulations and time periods for the pop-up convection (accounting for nearly 70% of the misses for both simulations). This is to be expected as it is difficult for even high resolution models to forecast these extremely small precipitation features accurately. These CRAs are left out of the tables to reduce clutter and focus on the mesoscale and synoptic precipitation features.

The first period, as suspected, yields similar results between the two simulations for detections of both mesoscale and synoptic entities. For the second period there is a noticeable difference in the hit rate for mesoscale precipitation entities, where the WRF-CSI experiences 8 hits over the WRF-CONUS. When we view synoptic entities we find that the WRF-CONUS holds an advantage of 7 CRAs over the WRF-CSI. However, the WRF-CONUS also has a false alarm rate of synoptic precipitation entities, higher than the WRF-CSI by four. These results suggest that the WRF-CONUS does have an apparent advantage in forecasting synoptic CRAs in period two, however the extension of the boundaries may also be hurting the simulation in a few cases as it is suspected that false synoptic influences may be advected into the JAX domain and setting off synoptic CRAs and missing mesoscale entities.

iii) Sea Breeze Verification

For sea breeze verification, the same time periods are used as in the EMT. As already stated, here we examine sea breeze detection skill and also look at wind speed, wind direction, and temperature statistics for each sea breeze event. For the 91 days of complete data and for the three sub sea breeze domains studied, there were a total of 87 sea breeze transitions detected through observations. Of the 87 observed sea breezes both simulations perform quite comparably when looking at the entire data set, the WRF-CSI with a 93% detection rate and the WRF-CONUS with a 92% detection rate. However, the WRF-CSI does have a slightly higher false alarm rate when compared to the WRF-CONUS. When that is factored in, both simulations have a near equal Critical Success Rate (CSR) of 0.89 and 0.885 for the WRF-CSI and WRF-CONUS simulations, respectively.

In terms of the average statistical parameters for the sea breeze detection cases, the WRF-CONUS holds a slight advantage over the CSI simulation. Both experience a cool temperature bias averaging around 0.5 K for the coast on these events, along with a slight tendency to underforecast the wind speed. As with the analysis of the EMT, the forecasts will be assessed for sea breezes during the 1-24 forecast hours and the 25-48 forecast hours. Similar to the results of the EMT, the first period of forecasts yields very similar skill scores for the detection of sea breezes for both simulations. However, period two shows differences between the two simulations, the WRF-CONUS exhibiting four additional missed cases compared to the WRF-CSI, which has three more false alarms over the WRF-CONUS. Statistical parameters for period two shows a more aggressive cold bias of -0.7 K for the WRF-CONUS and overall slightly more errors for each simulation when compared to the overall statistics for the entire period. Once again, these interesting results warrant a more in detail evaluation of some cases and of the synoptic influences which could be inhibiting or helping model performance.

iv) Case Verification

A subjective examination of synoptic weather maps and each 24 hour accumulation period for both simulations shows that when synoptic patterns are not influencing the model domain (and if synoptic influence is more than two days away from the JAX area at time of initialization), the precipitation and sea breeze forecasts are nearly identical between the WRF-CONUS and WRF-CSI. Every case where there is a substantial difference in forecast ability between the simulations are on days in which the synoptic influences were around 18-36 hours away from the JAX domain at the time of initialization. If these influences are not present, the high resolution of both simulations allows for the typical surface heating on the land which leads to a modeling of the sea breeze circulation and hence, identical detection rates. However, different detection rates between the two simulations are observed, mostly during the second period, on days with synoptic influence. In some cases the WRF-CONUS performs better, while the WRF-CSI is the winner in the others.

Obviously, one of the precipitation cases in which the WRF-CSI does not perform well is hurricane Dennis. While both simulations underforecast the spatial features of the observed precipitation field for all time steps in which the cyclone affected the JAX domain, the WRF-CONUS more accurately predicts the placement, volume, and size of the features. Since Dennis was never initialized in the WRF-CSI model domain, this result is not surprising. Figure 1 illustrates the 25-48 hour accumulation period and exhibits superior skill of the WRF-CONUS over the WRF-CSI. These results also hold true for tropical storm Arlene, which didn't provide as much precipitation, yet was another case where the storm was not initialized in the WRF-CSI domain.

The June 02, 2005 case exhibits an example of a situation where a cold front moved through the JAX domain during the second period. While the WRF-CONUS detects the precipitation as result of the frontal boundary, the WRF-CSI misses this structure altogether and forecasts a sea breeze CRA along the east coast of Florida (Figure 2). Obviously, the WRF-CONUS has the advantage in this situation because of the larger boundaries, as with the previous tropical cyclone case examined. While the WRF-CONUS detects the frontal precipitation in this case, it is important to note the bias of the simulation to overforecast the size of the precipitation field by a substantial amount. This type of situation occurred four times during the experiment and led to the relatively high miss rate of synoptic scale CRAs and a false alarm rate of mesoscale CRAs for the WRF-CSI. However, this does not explain the somewhat high false alarm rate for synoptic CRAs and the miss rates for mesoscale CRAs for the WRF-CONUS.

The June 06, 2005 case, among others, helps to explain the aforementioned detection rates for the WRF-CONUS setup. Here we have a situation where a frontal system from the north makes a slow propagation towards the southeast during the first period and becomes nearly stationary in southern Tennessee and northern Georgia during the second period. The JAX WRF domain never observed direct precipitation effects of this front during this period and normal sea breeze circulations develop throughout the peninsula (figure 3). The WRF-CSI correctly forecasts the precipitation from the resulting sea breeze along the west coast of Florida, however the WRF-CONUS misses this feature. While the WRF-CONUS correctly stalls the front in the correct location, albeit slightly farther south than observed, it also overforecasts the size of the precipitation and cloud fields, similar to the previous case examined. As a result, the precipitation field is observed in the JAX domain and helps to cut off sufficient surface heating necessary for sea breeze development. While this situation occurred a few other times throughout the experiment, they were not as drastic as this case.

b) Hot / Cold Start Experiment

i) Statistical Verification

For the average of all forecast hours from the period of May 1st through July 31, 2005 for temperature, the WRF-LAPS simulation appears to hold a very slight advantage over the WRF-Eta, with both simulations apparently exhibiting very similar forecasts on each day. Overall, both simulations have a slight warm bias. This warm tendency is the effect of both simulations to overforecast the temperature during the early morning hours, or the first six forecast hours. These biases reach their maximum during the 3rd forecast hour, with average scores of 1.42 K and 1.7 K for the hot and cold starts respectively. The

Eta simulation also has higher errors for this forecast hour. The afternoon cold bias is not as strong in magnitude as the morning bias, nor is it as strong as the initial installation tests of the version 1.3 WRF framework (Bogenschutz et al 2005). In this case we see the WRF-LAPS with a slightly stronger bias and nearly equal error magnitude when compared to the WRF-Eta. For the most part, the temperature forecasts differ very little for each individual day for all forecast hours, with one exception. On June 15th, the error in temperature at the 15th forecast hour for the grid differs by nearly 1 K, with errors of 2.7 and 3.6 for the hot and cold starts respectively. On this day there is a considerable cold bias exhibited by the WRF-Eta, which is simply a result of this simulation forecasting too much stratus cloud cover during the early morning and afternoon hours. This cloud cover is most likely a result of the initialization, as it develops during the early forecast hours.

ii) Ebert McBride Verification

Unlike the BSD experiment, EMT verification for this setup only focuses on one accumulation period, since the simulations in this experiment are only set up to be run for 24 hours. Precipitation results for this experiment yield very similar findings between the WRF-LAPS and the WRF-Eta simulations. For the sixty days of available forecasts, a total of 249 CRAs were detected in the JAX domain. Of these, the WRF-LAPS detects 72% CRAs which satisfy the minimum correlation requirements while the WRF-Eta detects slightly less with 71.2%. Both simulations experience twenty-three false alarms and thus have roughly the same Critical Success Index scores. The results between these two comparisons are even more similar than the results between the WRF-CONUS and WRF-CSI from the BSD experiment from period one. The correlation coefficient for the number of hits for each particular day is 0.975, which is slightly higher than period one of the BSD experiment. However, the correlation between the WRF-LAPS and WRF-Eta for the 24 accumulation grids is 0.84, compared to 0.70 for the BSD, suggesting that the effects of a different size domain may have more impact on the forecasts than initialization procedures for precipitation.

While both simulations exhibit comparable statistics, with the WRF-LAPS having a slight advantage, it is interesting to note that the trend for each simulation is to underforecast the size of the CRAs. This differs from the BSD experiment in which both the NMM simulations overforecast the CRAs by a fairly substantial amount. In addition, the average rain rates are more comparable to observations than the NMM runs. Although, it should be restated that a direct comparison between the ARW and NMM runs cannot be made. An analysis of results focusing on different precipitation sizes shows that there is no substantial difference between each simulations, with both the WRF-LAPS and WRF-Eta forecasting nearly the same number of CRAs for mesoscale and synoptic scale precipitation entities.

iii) Sea Breeze Verification

In terms of sea breeze verification, both the WRF-LAPS and WRF-Eta have detection rates comparable to the NMM runs, with detections rates of 92.5 and 91.8% respectively for the 63 observed sea breezes. As with precipitation verification, the statistics for sea breezes are similar between each simulation and comparable to those results from the NMM runs, with a tendency to overforecast the wind speed and temperatures near the coast which are slightly underforecast. The WRF-LAPS holds a slight advantage over the WRF-Eta for each statistical parameter.

5) Conclusions

It was shown that a larger domain configuration of the WRF framework over the southeast United States shows no overall advantage compared to the small domain run. While both experiment configurations experience roughly the same verification scores, CRA detections, and sea breeze detections during the first 24 hour period, it was shown that the second 24 hour forecast period yields different results in a few cases. Namely, while the WRF-CONUS has an advantage in the detection and positioning of synoptic influences because of the larger domain, it can also suffer due to the tendency to overforecast the precipitation field, which sometimes helped to reduce heating of the surface and hence the sea breeze circulations. While the WRF-CSI detects slightly more sea breezes and CRAs, the WRF-CONUS experiences better statistics overall for these features. However, any improvements seen in the

WRF-CONUS over the smaller domain is most likely not significant enough to warrant the excess computational cost and time to run the much larger domain version, especially when the WRF-CSI proves to be just as skillful for at least the first 24 hour forecasts.

Results from the Hot vs. Cold Start experiment shows a slight improvement in forecast ability of the WRF-LAPS compared to the WRF-Eta for every parameter tested. However, the differences are very small. The precipitation and wind forecasts have a strong correlation between the two simulations for each day included in the study except for one, in which the WRF-Eta experiences a severe cold temperature bias. The forecasts of the two configurations in this experiment were more similar than the first 24 hour forecasts of the two simulations in the Big Small Domain experiment. Both simulations perform quite well in terms of precipitation and sea breeze detection and neither suffers from any fundamental errors to disrupt these features (such as the consistently extreme temperature bias from the 2003/2004 WRF-ARW version 1.3 experiments).

Acknowledgements

The work by Ligia Bernardet, John Snook, and the JAX WFO for providing the model runs is greatly appreciated. In addition, this work could not be completed without the help of John McGinley, Andrew Loughe, Mike Kay, Christopher McAloon, and many others. The work was funded by NOAA and the NCAR WRF DTC.

References

Albers, S.J., J. McGinley, D. Birkenheuer, and J. Smart, 1996: The Local Analysis and Prediction System (LAPS): Analysis of clouds, precipitation, and temperature. *Wea. Forecasting,* **11**, 273-287.

Bogenschutz, Peter A., Ruscher, P., and P. Welsh, 2005: Skill of the new weather research and forecast model on warm season precipitation in the southeast united states. *Mon. Wea. Rev.* (submitted).

Case, J. L., Manobianco, J. Manobianco, J.E. Lane, C.D. Immer, and F.J. Merceret, 2004: An objective technique for verifying sea breezes in high-resolution numerical weather prediction models. *Wea. Forecasting*, **19,** 690-705.

Davis, Christopher A., Manning, K., Carbone R., Trier, S., and J. Tuttle, 2003: Coherence of warm-season continental rainfall in numerical weather prediction models. *Mon. Wea. Rev.,* **131**, 2667-2679.

Ebert, E.E. and J.L. McBride, 2000. Verification of precipitation in weather systems: determination of systematic errors. *J. Hydrol.,* **239**, 179-202.

Fritsch, Michael J., and Carbone R.E., 2004: Improving quantitative precipitation forecasts in the warm season: a USWRP research and development strategy. *Bull Amer. Meteo. Soc.,* **85**, 955-965.

Janjic, Z. I., 2003. A nonhydrostatic model based on a new approach. *Meteorol. Atmos. Phys*. **82**, 271—285.

Jascourt, S., 2002. An in detail look at precipitation in the Eta model, http://meted.ucar.edu/nwp/pcu3/cases/etapbomb/cover.htm

Kain, J.S., Stensrun, D.J., Baldwin, M.E., Manikin, G.S., 1998. A comparison of two convective parameterization schemes in NCEP's MesoEta model and the implications for quantitative precipitation forecasting. *Preprints,* 11th *Conference on Weather Anal. Forecasting*, Phoenix, AZ. Am Meterol. Soc.

Mahoney, J.L., J.K. Henderson, B.G. Brown, J.E. Hart, A. Loughe, C. Fischer, and B. Sigren, 2002: The real-time verification system (RTVS) and it's application to aviation weather forecasts. *Preprints, 10th Conference on Aviation, Range, and Aerospace Meteorology,* Portland, OR, Amer. Meteor. Soc.

Nance, L., L. Bernardet, H. Chuang, G. DiMego, M. Demirtas, R. Gall, S. Koch, Y. Lin, A. Loughe, J. Mahoney and M. Pyle, 2005. The WRF Developmental Testbed Center: A status report. The Joint WRF/MM5 Users Workshop, Boulder, CO.

Shaw, B., Welsh P.T., A. Wildman, J. Smart, P. Ruscher, J. Savadel, J. McGinley, M. Kay, J. Mahoney, J. McQueen, 2004. Applying the weather research and forecast (WRF) model to national weather service forecast office operations. *Preprints 20th Conf. on Wea. and Forecasting/16th Conf. on NWP*, Seattle, WA, Amer. Meteor. Soc. *11-15 January 2004*.

Welsh, P.T., A. Wildman, B. Shaw, J. Smart, P. Ruscher, J. Savadel, and J. McGinley, 2004: Implementing the weather research and forecast (WRF) model in a national weather service forecast office (WFO). *Preprints, 20th Conf. on Wea. Anal. and Forecasting/16th Conf. on NWP*, Seattle, WA, Amer. Meteor. Soc. *11-15 January 2004*.

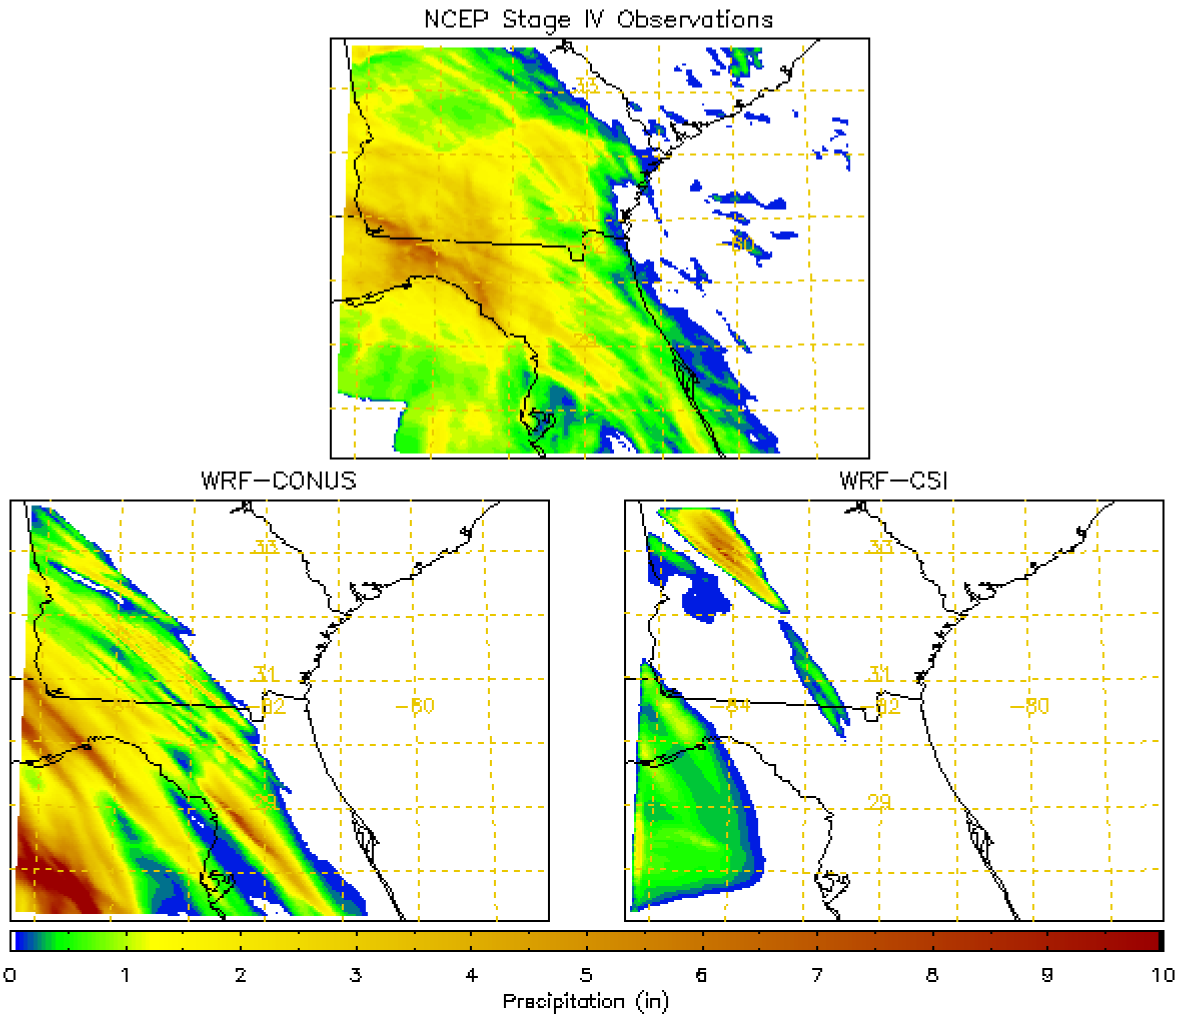

Figure 1. Twenty-four precipitation accumulation of 25-48 hour forecasts from June 9th, 2005 model runs for hurricane Dennis.

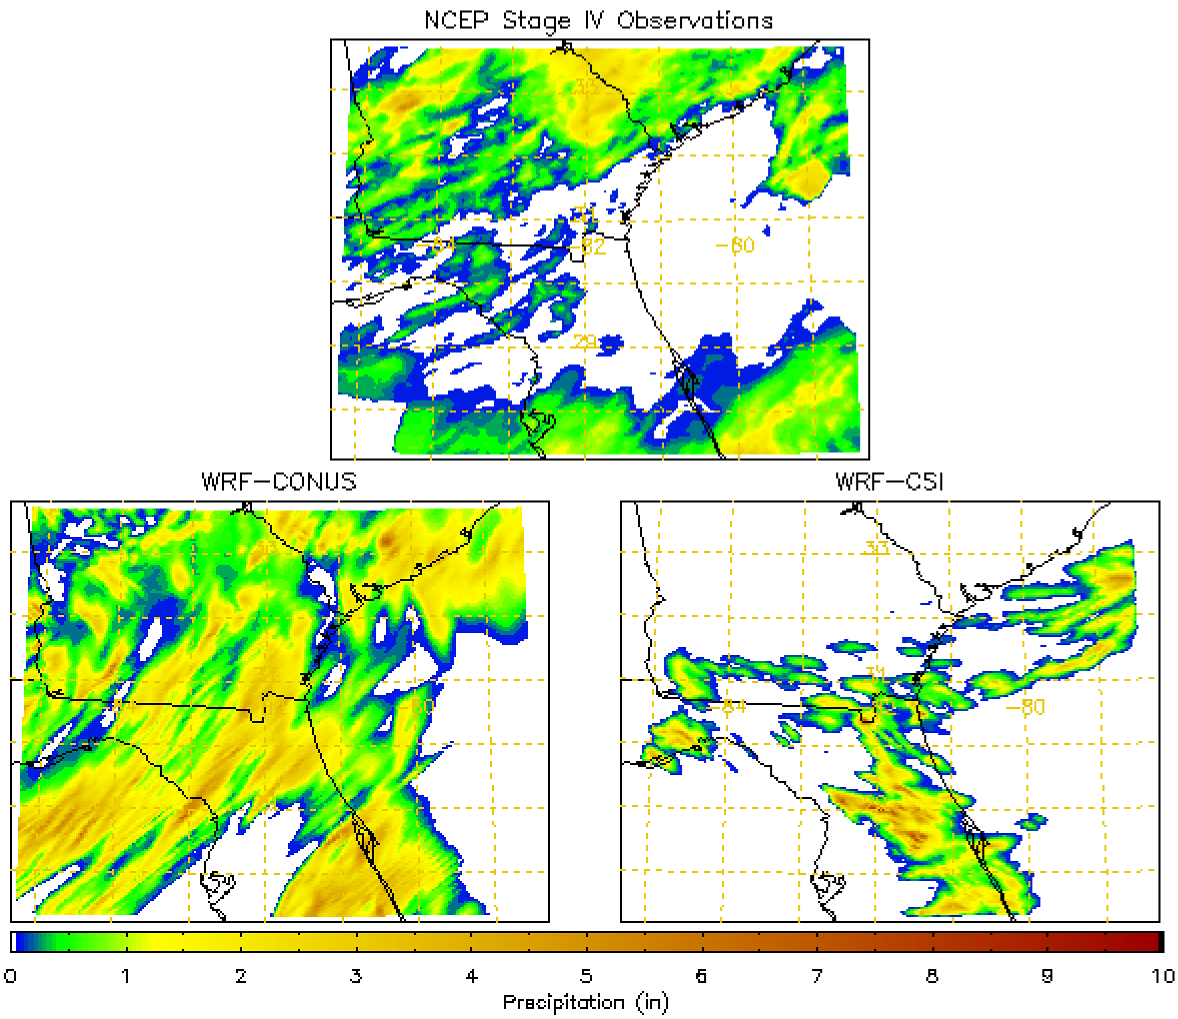

Figure 2. Twenty-four precipitation accumulation of 25-48 hour forecasts from June 2th, 2005 model runs.

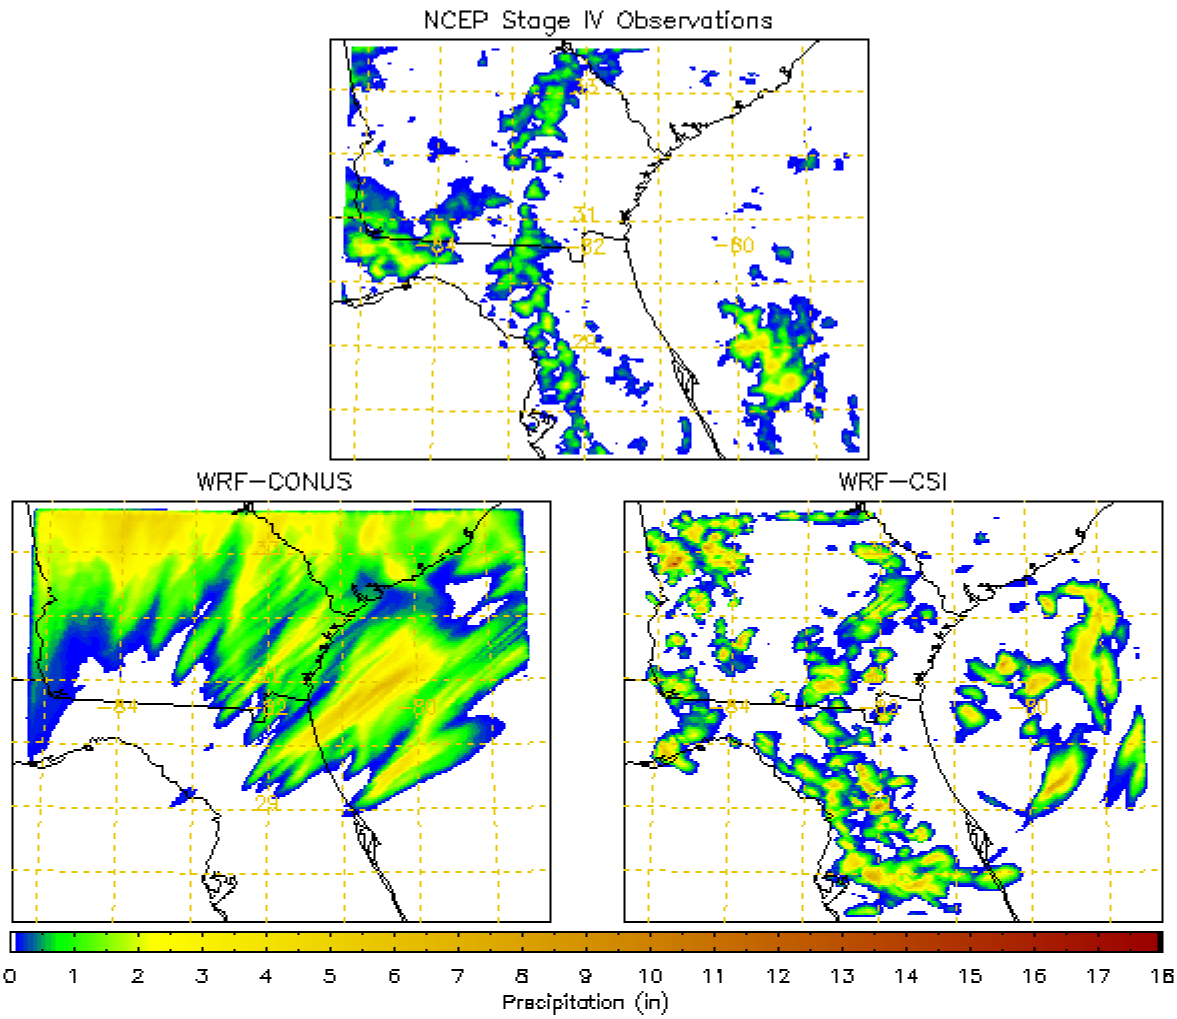

Figure 3. Twenty-four precipitation accumulation of 25-48 hour forecasts from June 6th, 2005 model runs.