

A Human Over the Loop Forecast Paradigm at The Weather Company (TWC)

Bruce L. Rose, Dale Eck, Leon Brown, Andrew Tupaj,
James Lidrbauch, Paul Bayer and Lauren Merritt – The
Weather Company, Atlanta, Georgia.

Corresponding author address: brose@weather.com

Despite the emergence of full forecast automation for a few weather providers, there would appear to be some consensus in the weather prediction community that a fully automated end-to-end forecast system remains approximately 10-20 years away. Despite this need, the crucial lag of time inherent in the human editing of objective guidance prior to publishing of the weather forecasts is increasingly problematic. In the case of TWC, the legacy human in the loop (HITL) man-machine process added at least 2 hours delay in the numerical weather prediction, human forecast editing, and forecast publishing assembly line [and this delay notwithstanding the versioning or provenance of the underlying guidance]. In late 2012 it became clear to TWC that a Human-Over-the-Loop (HOTL) or lag-less forecast process was desirable to bridge the “Age” of the traditional human-directed forecast process and the coming “Age” of a true end-to-end automated forecast process. As summarized in Neilley et. al. (2015), The Weather Company 24 x 7 global forecast operations shifted from a HITL process to a HOTL process in Summer of 2014. The human and machine migration is described in this extended abstract. The HOTL software tool is built atop a powerful graphics engine that allows for rich visualization of most current and forecast weather information that exists in the public cloud and within TWC’s weather infrastructure.

Concurrently, forecasters draw polygonal areas and assign a rich set of properties to these objects that either [a] constrain parts of the digital hourly forecast in a predefined but relative manner, [b] add value to the computer worded forecast via direct human oversight, [c] expound on local and hyper-local weather conditions of high impact or extraordinary nature that would be impossible for the machine to circumscribe or identify, and [d] alter or manage weights that control the blend of precipitation forecasts in the upstream guidance system. Forecaster responsibility [a] are termed Weather Filters, [b] are a legacy function called Qualifiers, [c] is a functionality called Blurbs, and [d] is a background weight-changing paradigm in the Digital Forecast system that is called Short-term Weight Management or SWiM. In all cases, the application of these human-directed forecast changes and forecast amplifications are done post hoc or just in time – eliminating all human-forecaster time lag from the end-to-end process. This paper will discuss the philosophy of the new forecast paradigm, and the responsibilities of the new HOTL forecaster, including the experiences of transition from the HITL to the HOTL environment. An overview of the HOTL tool and results from the first year of HOTL usage will be detailed,

and we will document highlights of a survey of TWC forecasters on HOTL usage and experience.

I. Introduction

A Human or Meteorologist in the Loop (HOTL) process was implemented at The Weather Company in Atlanta, GA. The HOTL software is built atop the Weather Services International weather graphics system called TruView Max. HOTL workstations interact with a centralized datastore of forecaster-initiated changes in Atlanta, Birmingham England, and Andover, Massachusetts. The forecaster change transactions or objects only impact the downstream forecast products when they are requested via a forecast API. By partitioning weather forecast guidance and forecaster alterations until post-production – one can eliminate delays that are typically associated with Human in the Loop (HITL) forecast intervention.

The transition to a HOTL process was part of a larger migration of the end to end forecast system in place at TWC (Neilley et. al., 2015). There was discussion of moving to a fully automated end to forecast that included smart ensembling of model and statistical forecasts, and then fully automated post-processing of that objective guidance into consumer and business-ready products. The science and meteorology team at TWC was able to convince stakeholders that an optimal solution from the standpoint of cost, resources, and forecast skill needed to include a human component in the forecast loop – and that this component was likely viable for a 10-20 year window into the future.

We agreed on a compromise that would eliminate any delay in forecast publishing as a result of forecaster intervention. This led to a HOTL-based system that managed the continuous flow of forecasts from upstream guidance to products without the interruption of that flow.

This migration worked surprisingly well and we are now approaching the end of our first year of HOTL usage. We have a digital forecast at 4km world-wide at one hour resolution through 384 hours. This includes a 6 hour forecast at 15 minute resolution that is also subject to HOTL oversight and intervention. Forecast alterations routinely occur in CONUS and Europe domains (as before) especially within the first 48 hours. Pre and post-HOTL forecast verifications show improved temperature and precipitation forecasts (year over comparisons) but much of this improvement comes via the conversion of the core forecast system (Neilley et. al. 2015).

In this paper, we explore the conceptual evolution of HOTL, provide some details on the software realization of HOTL, and also provide some forecaster feedback and insight from the first year of usage.

II. The HOTL Concept

The Forecaster or Meteorologist Over the Loop concept is commonly “in the air” at recent professional meetings and fora. Conceptually, there is a potential MOTL (Meteorologist over the Loop) or HOTL component to the NWS/FAA NextGen strategy and it is a common subtext of on-going discussions of the future role of the human forecaster (see for example, Sills, 2009 for a good literature review and Stuart et. al. 2006 in the context of the man-machine mix).

In the case of TWC, an elaborate compute system for forecast editing and intervention was in place. And this system called the Digital Intervention Tool (Digit), was part of a full end-end system of guidance production to forecast editing to post-process to publication and products.

The end to end system served TWC well over 12 years, but was in need of updating in the following areas:

1. More explicit points from the native system to reflect the higher resolution of modern NWP
2. Reduce the delay inherent in human forecaster intervention
3. Merge the two distinct short-term (0-6 hours) and long-term (1-15 day) forecast processes or systems
4. Provide more flexibility and universality with respect to downstream product creation

TWC did consider full end-end automation to overcome the delay of human intervention. However, it was believed that the human forecaster was adding value to both short and long term forecast within about 60 hours, especially with precipitation forecasts and forecasts of unusual or high-impact weather. Full automation was considered a bridge too far, thus we pursued a HOTL solution with the following characteristics:

1. The smart ensemble or composite forecast guidance (we call this gDicast) would experience no time delay or buttoning up period. The budget for forecasters to validate and tune the guidance was 0 seconds.
2. We rejected the paradigm of forecasters adjusting upstream initial conditions or ensemble inputs as a primary intervention tool. We did allow that forecasters had control over the parameterization of the ensemble or composite forecast, but we did not expect this to take place in any near real-time sense.
3. We wanted to retain some very good aspects of the legacy system with respect to the forecaster adding qualifications and context to the digital forecasts (called Qualifiers, these allow a long list of refinements to forecasts without changing their digital make-up.) An example of a

Qualifier would be “scattered frost” added by a human to a cold, crisp autumn forecast in Missouri.

4. We wished to add the ability for the senior shift forecaster to “speak” more directly to consumers. A geographic point and shoot method for text or other information to flow from forecaster to user.

Figure 2.1 shows an example of the completed HOTL client. There are weather “underlays” to help guide the forecaster in placing filters, qualifiers, and blurbs (FBQ’s) that constrain, alter, and amplify the consumer forecasts that ultimately are produced from the high resolution gridded sensible weather.

HOTL software was launched in July of 2014. Thus, the HOTL forecasters have worked through 4 seasons of operations on the new paradigm and the new software.

We provide two appendices in this extended abstract. The first is feedback from a recent survey of the operational HOTL forecasters in the US office

in Atlanta, GA. There are some annotations provided in brackets that may clarify or define special nomenclature or references to specific tools or processes with respect to TWC.

The second appendix is a document that was developed as a concept of operations “day in the life” of the HOTL forecaster. His name is Marco, and his fictitious work shift takes place in late fall when interesting or high impact weather is expected in different parts of the CONUS.

it is hoped that these project documents from our experience with a changing HITL to HOTL paradigm may help others in the weather enterprise facing the questions surrounding the future role of the human forecaster in our rapidly changing discipline.

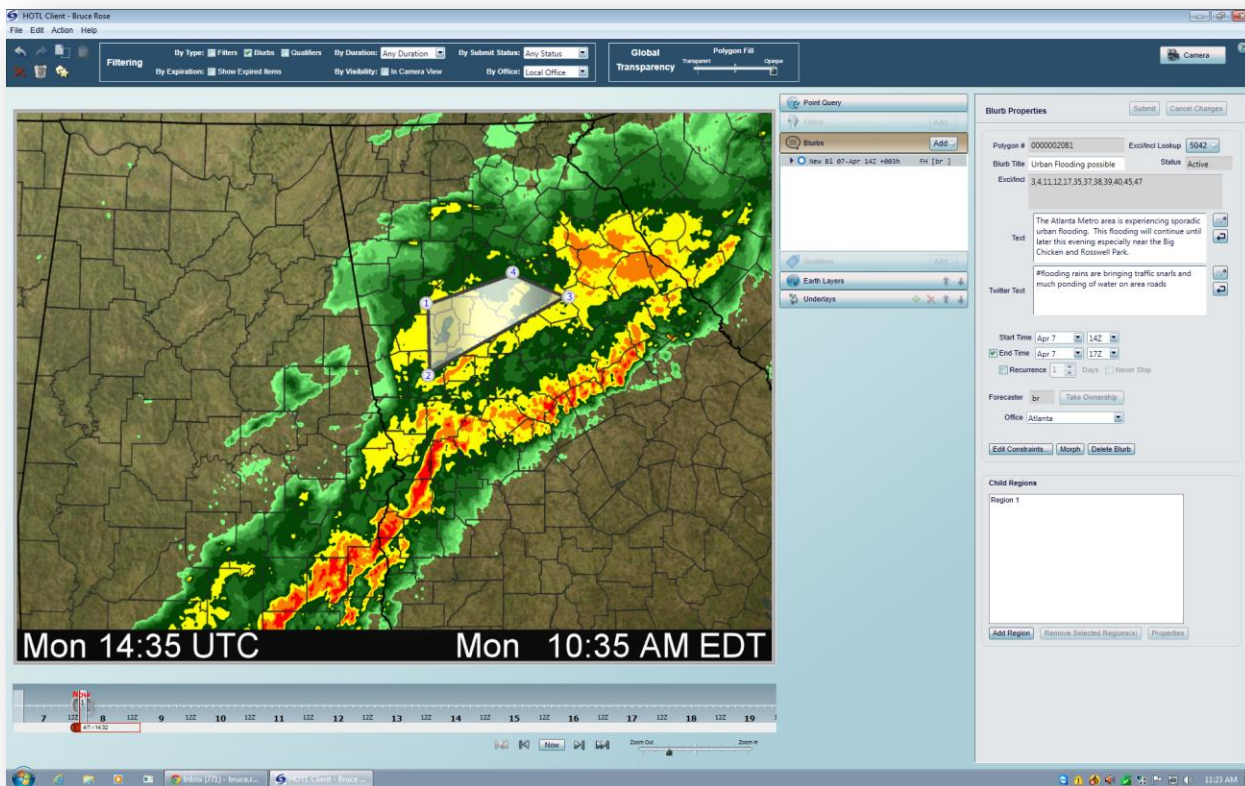


Figure 2.1 – The HOTL graphical user interface runs on Windows machines and is composed of a graphics work area showing some portion of the earth, GIS layers, and weather “underlays”. Forecasters draw polygons or clusters of polygons and then assign properties to those polygons that filter, qualify or further amplify the underlying gridded forecasts. The forecaster instructions embodied in the polygonal objects are only applied when a forecast is requested from the API thus parallelizing the flow of forecast guidance and forecast intervention from the human.

III. References

Mass, C. F. 2003: IFPS and the Future of the National Weather Service. *Wea. and Forecasting*, **18**, 75-79.

Neilley, P.P. and Coauthors. 2015: Overview of The Weather Company’s Principal Forecasting Methodologies. *AMS 27th Conference on Weather and Forecasting*. Chicago, Illinois.

Sills, D. M. 2009: On the MSC Forecasters Forum and the future role of the human forecaster. *Bull. Amer. Meteor. Soc.* **90**, 619-627.

Stuart, N. A. and Coauthors. 2006: The Future of the human forecaster in an increasingly automated forecast process. *Bull. Amer. Meteor. Soc.* **87**, 1497-1502.

Appendix I – Survey of Operational Forecasters in June 2015.

1. What was the most difficult aspects of transitioning from Human in the Loop (HITL) to Human over the Loop (HOTL)?

- Not being able to see what the result of filter/qualifier placement graphically in real time.
- Also the attitude change of less intervention was difficult at first.
- Probably giving up some of our "control" of the POP, (especially at specific time intervals) and more "trusting" the GDicast model solution of wet versus dry. Also, being a more detail-oriented person, it was a bit difficult to let some of that go and be more accepting of the overall model solution. This was especially challenging when dealing with precip-type forecasting and trying to be precise with snow QPFs. Plus, it took a while to get used to "flying blind", especially concerning forecast cloud cover and temps which have no representation graphically.
- In HITL [*Legacy editing system*], we were used to being presented with 3-hourly forecasts in an easy to see and use interface. The HOTL interface did not initially present 2-D forecasts to the METs and now that it does, it often an incorrect and untimely view of the 12-hrly roll ups (the 2-D shields lag the latest FOD forecasts and in some cases, never reflect the ST forecasts.) It is difficult to identify and correct "stupid" forecasts if they are not presented to the MET in a plain and timely fashion.....clicking around on a map to see point forecasts, though they are presented in near-real time fashion...is not an efficient way to QC FOD [*published forecast on demand forecast*]. On the plus side, it was very easy to begin to trust FOD with routine forecasts, thus allowing the met to concentrate on "trickier" forecasts (areas the models are struggling with.)
- The most difficult transition for me was giving up control. In Human in the Loop I had total control over the forecast. For Human over the Loop we are encouraged to only change "dumb" forecasts while leaving most of it alone. Thus giving up some control. It was also hard to believe that I was not improving the forecast as much as I thought I was in the Human in the Loop system.
- Not being able to know for sure what the latest forecast is, especially with the ST [*Short Term nowcast or 6 hour forecast*] changing so frequently and no real way to find out graphically in real time, the "hunt and peck" method with the FOD tool is not very efficient. We need the GE [*Google Earth layers with key forecast versus observation information*] tool to be more up to date. The fact that GDicast does not handle the inherent uncertainty with convection well and there is sometimes no way to remedy that without using the "sledgehammer" approach.

- Probably one of the most difficult things I have found in dealing with HOTL versus DIGIT [*Legacy edit tool*] is that I don't know right away what I'm changing in certain areas. The FOD browser updates very quickly so we can see what the 'new' forecast is, however other points nearby might be moved in the wrong direction. That is something that took some time to get accustomed to and is now just another part of the daily routine and isn't an issue for me.

2. How valid is the HOTL assumptions and philosophies? Do you feel it is working well?

- The assumption the models [*smart guidance or gDicast*] are generally good enough to allow for minimal intervention is optimistic.
- In theory it all seems plausible, and that having the latest model data constantly flow through the system and produce forecasts on-demand across the planet. In practice, there are often some very undesirable results in my opinion, as models flip and flop and we as forecasters are not always aware in a timely manner to intercede. It appears that the system is working, and probably well enough in a lot of cases. The HOTL graphical interface is fairly easy to use and manipulate the forecast data.
- Overall, I feel it's working very well. With convection, the assumption that high-resolution models should be taken at face value is flawed. As mentioned above, we absolutely need better tools to identify forecasts that are missing the mark.
- It is working better than I thought it would. I was convinced that we would need to be doing a lot of work correcting bad forecasts. I am still concerned models will not handle a landfalling tropical cyclone well and without a form of the hurricane wind tool [*a legacy tool that may or may not be ported to HOTL*] it could lead to a lot of work and frustration correcting bad wind/rain forecasts. While I did not experience it myself, I understand GDicast came up with 10-20 mph winds in Myrtle Beach when Tropical Storm Ana was landfalling. If that continues we could be in for big busts in high profile events this tropical season.
- The best thing is that in cases where the models do pick up on new trends or situations it will get into the forecast faster than a human can. The overall "models blend" philosophy of GDicast is good, especially in the LT [*medium term or traditional forecast*] - it is worst in the ST.
- In very mesoscale or regional type situations (Lake Effect snow, Santa Ana winds as two examples) the model base is basically totally clueless, which really renders the philosophy of letting the models be the main base of the forecast incorrect.

- In general it seems to be working well. With the higher resolution and the fact that GDicast is a superior product to our old Dicast, we are not making as many changes to the going forecast and our verification scores have been very strong. We continue to learn (passively and by necessity) new ways of tweaking and adjusting things as new seasons and new challenges come along.

3. What do you see is the biggest benefit of the change?

- Intervention time is definitely reduced compared to Digit allowing more time for details and nowcasting. We are able to make changes much more quickly. New system allows for more detail for timing events. Much less time consuming to make winter weather forecasts.
- The ability to at least do a reasonable job tuning forecasts over a wide geographic area with minimal man-power. Also, much easier to adjust the forecast over the western US with the large changes in elevation than was possible with the old dicast/digit system [*legacy guidance and edit tools*].
- Forecasters do not waste time fixing simple forecasting issues and have learned to trust FOD with those. The system is also showing increased skill with winter weather. The fact that new model data is ingested continually has eliminated the HITL step of deciding where and for what valid times new data should be loaded...I see this as a big plus. In most cases, the constraining filters methodology does a good job of putting bounds on newly ingested model data, when needed.
- The biggest benefit so far is the time savings. It used to take up to 2 hours to have a buttoned up forecast for the first 60 hours in a third of the country. Now I feel confident of my forecast in probably 30 minutes.
- The forecasts that we have to work with are better than in our previous system. There are always things that need to be adjusted, especially in the shorter term, but with the HOTL system we are able to make those changes very quickly and easily. The changes we make update very quickly so we're able to poke around on the FOD browser [*a tool capable of point queries of consumer forecasts as they will appear on the outward facing API*] to do a sanity check on what we have done and, if it isn't quite what we want, we can change and check things once again very quickly.

4. How long did it take for you to feel "comfortable" working in the new environment? Was there a light bulb that went off for you at some point?

- Still working on that. Still occasionally find a result from a qualifier(s) or filter(s) that is perplexing.

- Probably a few months overall; at least until we got into the winter season and started to use those tools and applications. It was more of a gradual thing for me rather than a light bulb moment. The more I worked with it and got experience, the more natural the process became.
- I was comfortable with HOTL very early on, though the lack of viewing real time forecasts was and still is a big concern.
- The only "light bulb" that went off for me was when I realized it was okay for the graphic maps to look differently than the ones produced via HITL (this took just a day or two!)
- I am a fairly quick learner, so I would say I was comfortable after two weeks. The 6 month practice time we had to play with the system probably helped with that. Although, the practicing did not fully translate when we went live. When practicing we had certain ideas about what we would be doing. Reality proved our assumptions to be off base. I'm not sure I had a light bulb moment. Rather it was through repetition that I grew more comfortable.
- I was able to adjust pretty quickly, except for a few quirks.
- I still feel like I'm learning to some extent every day. Each season is its own animal and as such has its own intricacies in terms of how we might attack a forecast. I felt comfortable within a few days of starting live with HOTL although there was never a 'lightbulb' that actually flickered at any point. I would say that compared to what I was expecting HOTL has been much easier to use.

5. What are your greatest frustrations/challenges with the HOTL concept?

- Not being able to see the effects of filters/qualifiers in real time graphically. Waiting an hour plus for maps to build is a huge downside.
- Pressure being put on us to be more concerned with the graphical output as opposed to the actual forecast. We foresaw this and expressed it and we were told it would not be a problem.

- Granted this is intrinsically tied to FOD, but the fact of the matter is, the models often really are not that good. This is especially true concerning the forecasting of convection. The high res RPM (for example) [*our in-house WRF instance*] may have a reasonable-looking depiction of pm convection, but if it is off by 50 miles with its individual cells and/or lines and their associated high pop, the actual forecast at any particular point can be quite bad/unrealistic. Then another model run arrives and the convection is in somewhat different locations and the process is repeated. Also frustrating is the constant changing of forecast values (like snow QPF, temperature and precip type) during the winter season. Instead of having a consistent forecast for a given market (ex. snow mixing with and changing to freezing rain, high 31), it could waffle each time a new model input gets ingested, so you end up with a mish mash of weather, temps, perhaps on the wrong side of freezing, and both disgruntled internal and external customers asking why it constantly changes.
- Again, forecasters need to have ways to easily identify forecasts that have gone awry in order to take corrective action.
- Constraining snow qpf can sometimes be a challenge due to underlying forecast changes.
- Also, a new and very large concern is the fact that graphic aesthetics are of more importance than point forecast data. Many corrective actions taken by the met in HOTL are being discouraged (even forbidden) because those actions are deemed to adversely affect graphic products, even though those actions most often produce a vastly superior forecast as compared to unaltered FOD. This concern was raised early on during the development of HOTL and forecasters were assured it would not be an issue because the number of consumers using point data forecasts far outnumber those using graphic forecasts.
- The greatest frustration for me is not knowing what new models are doing to the existing forecast visually for 20 to 30 minutes after the data is ingested. Theoretically customers can get the new forecast right after the model is ingested and the next GDicast run has completed. So, there are changes made to the forecast based on the new model that I may not be able to correct (if needed) for a half hour or more. This can lead to some flip-flopping of forecasts.
- As stated in # 1, not being able to see very well where HOTL may be not doing so well in a timely manner. Not being able to handle convection effectively. Having the model drive most of the forecast can create somewhat of a case of "meteorological cancer" where the forecast can become lazy and not really do the type of in-depth analysis

- As mentioned above making changes without knowing what will truly come out can be frustrating. I miss the DIGIT table with the values for each parameter at each polygon being right there on the screen. Being able to use FOD within seconds does make that easier to wrestle with. As well, we seem to have some issues in the shorter term with PoP running hot and cold during the convective season. It can be frustrating (and difficult to find unless you hunt and peck over an area), but we are coming up with ways to temper this type of thing.

6. Talk about the HOTL tool. Is it an effective tool for working in the "Over the Loop" world? What are its strengths, and where does it fall short?

- Can be frustrating. Moving control points, filters/qualifiers can be troublesome. The way the Station ID's show is frustrating having to zoom in and out with the ID's popping on and off.
- Per the above, it is relatively easy to use and the concepts of using filters and qualifiers to adjust the forecast work for the most part. The filters are logically written overall and make sense. The underlays are quite useful overall. It works well in limiting pop/temperatures with the clipping concept. And I think we have a fix in test to improve the snow QPF range "accuracy". Where it falls short? The weather shield depictions are rather crude and at times not totally representative of what is the actual hourly or daily weather shield for a particular point. Also, as stated above, we are not able to graphically see the cloud cover or temperature forecasts. Except for hunting and pecking specific points, we have no way of knowing what these fields look like spatially. Makes it quite difficult, for example, when trying to adjust temps for a wedge situation.
- Overall, I think it's effective. The ability to zoom in geographically on the topographic base map is very nice. Setting time limits for filters is easy and straightforward. Mouse actions for drawing new filters or altering existing ones are simple and fairly easy. Having radar and lightning data, as well as various model or FOD parameters overlayed on the maps is very helpful. Implementation of new, frequently used FBQ's [*predefined Filters and Qualifiers*] to the set lists seems to be a relatively simple process. Morphing can be a valuable tool [*ability to produce a series of polygons based on a starting position and ending position, and possible intermediate positions*].
- Not being able to see hourly temperature forecasts spatially is a big negative...even max/min temperatures are presented in a rather crude manner. Cloud cover forecasts are "invisible" to the forecaster (in the HITL world, different shades of white and gray on the base map identified cloud cover forecasts.) And, again, the ability to view real time spatial sensible weather forecasts is sorely needed.

- For the most part the tool is highly effective. Strengths are time savings, it handles most situations fairly well, it is relatively easy to make changes through the application of filters. It has fallen short lately in the temperature department. It has been too warm in the Southeast and Mid-Atlantic. With our old Dicast there was an artificial intelligence of sorts built into the system. If models had been too warm/cold Dicast would make an appropriate correction. That feature is not in GDicast so we are now [missing] that feature. One other big potential short-fall is the lack of the hurricane wind tool. We are hoping models will accurately depict landfalling tropical cyclones. If they are not there will be some major corrections needed by the forecasters.
- The tool itself is about as good as you can make it, given the other limitations that I have stated elsewhere that are more related to the concept itself and the accuracy of the base forecast.
- I believe it is a more than adequate tool to work with for us and I don't think I could come up with a better tool. Being able to constrain what is being changed helps a bunch. As well, being able to access the changing forecasts relatively quickly is very helpful in making sure our changes are what we truly want. As for where it falls short...I tend to work with what I have and not dwell on what isn't so helpful. I think the number of filters that can pile up over an area, especially during nowcasting times, can make things terribly cumbersome. Clicking down through the layers of filters...moving a polygon when you only want to move a control point...the number of control points on one FBQ can all be things that frustrate someone working with HOTL. But in the grand scheme of things these issues aren't more frustrating than fixing a typo in a paper or having to reenter your username and password several times to get into a website. They come with the territory.

7. How does the HOTL concept work as it relates to nowcasting and the Short Term system?

- I think it works pretty well for nowcasting. The emphasis on the "look" of the maps as opposed to the result of the actual forecast can be frustrating.

- This is perhaps the greatest area of frustration currently, as our hands are tied on how we can interact and improve the forecast because of aesthetic issues. We are operational forecasters, and as such, should be able to focus on making the forecast the best it can be without having to worry about "the look" and if the "fix" flows naturally. What happened to the concept we were told from the outset that our internal customers would just have to get used to the new somewhat ragged "look" to graphics? Also, having the forecast flip flop violently in the nowcast realm is another sore point. It can be very difficult to keep up with pop at a point changing from 80 to 20 to 60 which often happens in the close-in world driven off the RPM. Again, there is no way to know for sure what is coming out except to hunt and peck at various points. Very time-intensive and not very efficient in my opinion.
- Currently, we use a hunt and peck method to try to identify nowcasting and ST forecast issues....this is very inefficient. We have no way to easily see how the ST system is altering forecasts. Otherwise, the process of filtering the ST forecasts is as easy and straightforward as filtering LT forecasts. As mentioned before, having real time radar on the map is extremely helpful.
- While I do not have a great deal of experience in the short term or nowcasting I feel this may be the system's weakest point. What I am referring to is how long it takes us to see a visual representation of the forecast. Instead of that being instantaneous or at least in terms of a few minutes, it takes up to 30 minutes to "see" the forecast. We can click and get an forecast right away, but to find all of the potential problems in the nowcasting time frame it could require a lot of clicking. It could resemble finding a needle in a haystack. If the updating of the graphical representation of the forecast were quicker we could correct wrong forecasts more quickly.
- As said above, the ST changes too quickly for us to be able to keep up. Also, just trying to "work with the underlying forecast" by scaling POP is sometimes not effective when the underlying forecast is way off, which does happen especially with convection.
- As mentioned above, we are finding some issues with shorter term forecasts running hot and cold within a very short difference...this is likely related to the convective season. It happens in other seasons but convection is by far the most difficult season for this. There are some fixes coming in that will help us to smooth these types of things out in the near future. As well, the 'piling on' of filters and qualifiers in the short range can also make for a cumbersome nowcast experience.

8. Any other insights into this new way of working?

- I like that our forecasts go out basically immediately in real time as opposed to waiting an hour plus for updates to take affect...at least in the FOD system. Would be nice if the Web got the same consideration. In many ways HOTL can be easier to work with than the old DiCast system. Ptype [*precipitation type*] and qpf (especially snow) forecasting was extremely time consuming in the past, not it is much easier once you decide what you really want to do. We have more time now to look more closely at details than we did before.
- I do appreciate that we are putting time and effort into continuing to tweak the system/methodology to improve processes and outputs. We learned a lot this past winter, and I believe we will learn more over the next 6-9 months as well. Hopefully we can continue to work together to increase product accuracy for our customers as well as METs' sense of accomplishment and satisfaction.
- The transition to HOTL has been well-received by the met staff as an easy, efficient way to produce quality forecasts. The met staff is committed to the FOD/HOTL concept and understands that forecast accuracy has been increased in many instances from the HITL method. However, a few important issues should be addressed (as mentioned above) in order to provide even better forecasts for our customers.

We really need to continue to:

Improve the base forecast.

Give the forecasters a better way to see what the latest forecast is.

Appendix II – A day in the Life of a Shift Forecaster at The Weather Company in Atlanta, GA

Early morning before work

A senior forecaster named Marco awakes early on the morning of October 29, 2014. He has a 10-hour shift scheduled in the GFC (Global Forecast Center at The Weather Company) from 8:00am until 6pm. There will be two other METs with him. He'll likely do a couple of briefings for Programming (TV), write weather impact reports, interact with the tropical expert, the hurricane expert, the winter weather expert, and the Lead MET responsible for placing reporters or OCM's (On-Camera-METs) in the field.

However, the main focus of his shift will be heads-down on a system called HOTL (Human over the Loop). Marco

will spend about 60% of his shift time interacting with this system.

Marco's previous day at the office was busy. There were two tropical systems in the Atlantic and Gulf. One may threaten Florida in about 3 days. There is an early season winter storm that will probably bring heavy snow to the Arrowhead and parts of the UP of Michigan and western Canada. This same system exhibits a strengthening cold front that may be the focus for isolated severe weather in the southern plains on October 30th.

Marco noticed that the ECMWF and RPM seemed to be handling the mid-latitude system the best (including the snowfall and potential severe weather), but their tropical solutions were not that good and were not very consistent with NHC strength and position forecasts.

Marco knows from the GFC forecast discussions written overnight, that no Load Science weightings were performed and he is thinking that early in his shift he may have time to take on some weight changes for the domain in and near the CONUS. There are a number of CONUS filters to address the “strong” storm potential still active in the Plains. There is a single snowfall range filter in the far north, but the number and complexity of these filters will probably become more widespread as the storm draws near. The snowstorm may need to be limited by elevation in some areas. As well, there could be cases where points near the warmer Greater Lakes will have difficulty changing to all snow or difficulty experiencing any meaningful accumulation of snow.

All of these thoughts occur to Marco well before his shift begins. In a number of ways, Marco is anticipating the day’s big weather stories and some of the issues they will present to him and the shift of workers in the GFC during the day. In Turn, Marco is also thinking about how some of these issues will be addressed or overcome and how he can help tell the weather story more completely and how he can contribute to skillful or smart forecasts – while avoiding or filtering out unskillful or dumb forecasts.

The Beginning of the Shift

Marco arrives at work and logs into a HOTL workstation – one of five systems in the GFC. He sees a main menu and continues to the core system. There is a map of the earth that he can pan and zoom. There are lists of layers on the screen and lists of text plots on the screen. He turns on the current radar imagery and puts it into a slow animation with the map centered on Chattanooga, TN and the map or world showing about ½ of the CONUS. The radar continuously loops with a dwell or hesitation on the last image of the series. Marco can see through the radar because it is somewhat transparent. Below he sees the U.S. with state borders, ICAO’s and other landmarks displayed. He can also make out elevation differences since the map has a subtle coloring that denotes topography. The ICAO’s relate to observational METAR or SYNOP locations around the globe. There is also some number of PWS (WeatherUnderground Personal Weather Stations) stations that can also be seen. The text display is

decluttered, so the number of ICAO’s, SYNOPS, and PWS’s that appear on the geography are nicely spaced and do not overwrite one another. Each of the OBS points display a thumbtack or circle near their name/abbreviation and represents the precise latitude and location of that observation point.

Marco can double click on one of the observation points and a window pops up in the display. This window shows the current observation, the short-term now-cast for the next 6 hours, the short-term GFC forecast (worded) for the next 2 days. There are some additional links inside of the window. One of the links will expand the window display to include the 16 day forecast. Another link will show observation history of this point, and another link will give 15 minute and hourly forecast data for the point. There are about 7,000 of these points world-wide. Some 2,500 are within CONUS and include about 1,400 ICAO’s and 1,100 PWS stations.

Marco can also simply click on the map at any place that is not a thumbtack or the line-edge of an existing polygon that is being displayed. If Marco does this, he will see a very similar window and similar content, but the observation may not be available. Marco knows in both cases, that these query windows are showing the current active FOD (Forecast on Demand) forecast interpolated from a 4-km grid. In this way, Marco can quickly review key forecasts of either short or long-term for specific points on the earth.

If Marco does click on a thumbtack, then the forecast window (or information window) will appear for that specific ICAO. Only a predetermined set of ICAO’s (About 2,500) can be viewed in this manner. These plotting points are checkpoints, specially tailored to the HOTL system through an interest list. The information for an ICAO window is somewhat different. This window includes official METAR reports that cover the past several hours of METAR and SPECIs along with Hired Currents information (Synthetic surface observations). The window also includes meta-data for the point itself including Latitude, Longitude, Elevation, and Place Name (typically an airfield or population center).

Working with Polygons

Besides the Earth and thumbtacks or live objects, there will likely be a number of polygons visible on Marco's display. These polygons have different meanings and purposes depending on their color, outline, and internal attributes. In general, the polygons come into existence when forecasters explicitly draw a closed set of control points or vertices on the Earth, complete an associated form of data fields including a time duration, and at some point submit or save the polygon and its attributes to a permanent database that sits outside of the HOTL environment.

Polygons can be created or added, edited, and logically deleted. That is, a polygon can be marked as deleted but may still appear on the display and may still be seen in the database of polygons. Polygons appear on the display dependent on their position (i.e. they coincide or partially coincide with the extent of the geographic display) and their valid times.

The display has a sense of hourly time, and there is an hourly forecast valid time that is presented on the display as a number or timeline below the geographic viewer. If the polygon valid time is coincident with the exact valid time indicated on the timeline, then the polygon will display. Marco can move the timeline to any valid time within the 384 hour time horizon of the forecast. Marco can also summarize a series of contiguous hours by using his mouse to drag or broaden the timebar. This action has effect on polygons only in that it will produce a display of all polygons that fall within the geographic extent of the display that are valid or intersect in time with the range of time now indicated on the timeline. This time interval broadening can only work in the forward direction, but otherwise is not limited to any maximum time duration of that interval. However, Marco must keep in mind that a very broad time interval may bring forth a large number of polygons on the display creating a good deal of clutter and – even – polygons that overlap precisely on the display.

If Marco clicks on the edge of a polygon or the polygon outline, an editing window will open with all of the attributes of the polygon. The display changes slightly to indicate this editing mode, and the cursor or pointer related to the mouse position will point to the first

editable field. An escape key simply abandons the editing. As long as nothing is changed the display will return to the query mode as before and the editing window will disappear.

The attribute window of an existing polygon has a number of fields, most of which can be altered by the forecaster. The most important field is the polygon type. These are Filters, Qualifiers, or Blurbs (FBQ). The composition of the editing window is very much dependent on the choice of FBQ. Many fields will appear or not appear depending on the FBQ attribute. Once FBQ type is determined, it cannot be changed. The forecaster would need to remove or delete the polygon in such a case and redraw it – then choose the new FBQ type.

If Marco left-clicks and holds his mouse on the edge or line of a polygon, the entire polygon can be moved. If the polygon is moved to the edge of the display the earth will pan in the direction of movement to accommodate this relocation.

if Marco left-clicks and holds his mouse on any control point on the edge of the polygon he can move the control point only to change the shape of the polygon object. If Marco right clicks on the polygon line he can add a control point at the location of his mouse, he can delete the polygon, can change appearance of his polygon or can change the order of his polygon in relation to any polygon that lies above or below the layer of his polygon. Marco realizes that he cannot simply right-click anywhere within the polygonal area, because it would be unclear which polygon he was referring to in cases where multiple polygons overlapped one another. Thus, to access right-mouse-click menu functionality for a polygon he must right-click on the polygon outline itself.

Managing an early season snow event near the Arrowhead of Michigan and Lake Superior

After coming up to speed on the going FOD forecast and the guidance and HOTL FBQ's that are in place in the CONUS and nearby regions, Marco wants to introduce a Filter right near the shoreline of Lake Superior and Duluth, MN. The purpose of the Filter is to inhibit snowfall or accumulating snow during the daytime

period in this area. There is an existing Filter called No Accumulation that will be put into place that will ensure that no snowfall accumulation is allowed to take place during the valid time interval of that Filter. The way it does this is to simply set the snow ratio (SR) variable to a very low amount irrespective of any other weather variable. The weather forecast itself will not change with respect to temperatures, weather type, or QPF. However, the Filter will ensure very low SR's and will inhibit most or all snowfall within the user-drawn polygon and user-specified time interval.

Marco spins the Earth and pans to the general region; he then zooms into the area running from the Northern Duluth metropolitan area out over Lake Superior to areas south of Hancock, MI and Sawyer and Gwinn, MI. In this general area Marco selects the draw polygon tool and begins to introduce control points. Since Marco is in a spline-mode, a smooth outline begins to connect the points. When Marco is completed with specifying control points, he double-clicks and the spliner smoothly closes the loop and the polygonal object is now completed on the Earth.

The double-click also signals an Attribute window to appear on the r.h.s of the screen. The First field is a pull-down that has Filter, Qualifer, or Blurb as choices. Marco selects "Filter". The second field is a pull-down menu that has categories of existing Filters. Marco selects "Snowfall" as a category and a sub-menu lists about 8 choices that include the "No Snow Accumulation Filter". This is selected.

Marco wants the Filter to only act close the lakeshore. He can – therefore – draw his polygon to precisely follow the shoreline. However, he decides that it will be easier to include a conditional in the Filter so that it only acts on points that are below 800'. In this way, higher elevation points that assure a safe remove from the coast will be excluded.

Marco sets the time duration from the current time through 00Z on 30-OCT. He wants to allow snow accumulation to occur during the overnight, so he will have the Filter expire as evening approaches in this region. Marco could have the Filter recur (in time) for tomorrow afternoon as well, but he decides against

this, and will probably make mention of this decision to the second shift.

By completing these few fields in the FBQ attribute form, and saving that form he has defined and placed the Filter. The Filter will not act on FOD forecasts until Marco's entire work session is submitted to the HOTL server database. Marco will do this later. However, the act of creating this Filter during Marco's work session (i.e. the time beginning with his login to HOTL) causes the screen background to change, and also sets reminder tag in the upper r.h.s. of the screen. Moreover, if Marco were to logout of the session – he would be reminded that there are pending changes that should be submitted to the HOTL server database. Only if Marco abandons the session (via inactivity and timeout) or if he chooses to NOT submit actions to the database will those Filters go unrecognized by downstream FOD forecasts.

Now that Marco has completed the Filter creation process, he sees the "NOACCUM" polygon on the screen without activated control points. It has a distinctive whitish-gray semi-transparent color-fill and a dark blue border with dashed lines. The color-fill is unique to NOACCUM. The outline color is solid (it has no transparency) to enhance visibility on the screen. The line is dashed since the Filter has not been submitted. Dark blue is a line color for Filters. Bright red is used for Qualifiers and a vibrant green is used for the outline color of Blurbs.

In the area of snowfall or mixed rain and snow Marco thinks the nighttime temperatures will probably settle very close to freezing. He suspects that gDicast will want these temperatures to fall lower. But given the time of year, closeness to the lakes and the amount of near-surface moisture Marco thinks this is unlikely. He therefore sets out to place a Filter for the overnight hours in this region to prevent temperatures from falling too low.

He selects draws Filter form the forms menu, and he then establishes that the Filter will be hand-drawn instead of selected from preselected or predefined point sets. Marco draws a polygon that falls just south of a KMSP to KANJ line and extends this southward to

the vicinity of KORD to KCLE and over the KBFD closing off the polygon.

He selects a Filter type of DELTA. This is an open-ended or blank Filter where Marco can define both Conditionals and Actions and there will be no other logic transacted. In the Conditionals section he chooses all points where temperatures are below 31F by entering “-80 – 31” into the range field. He also enters “75-100” into cloud cover to limit the selection criteria to cases where cloud cover is nearly total. Marco then enters “32” into the temperature Actions fields. He completes the top level or master form by entering a time range using the increment controls to establish a start and end time. Again, there is no recurrence.

Marco continues to work on the early season snow event. He selects a Qualifier type and draws a polygon that circumscribes an area away from the lakes and in the sweet spot of the cold air and moisture. He selects a Qualifier number 7653 which will mention some wet snow accumulation likely occurring on grassy surfaces. The timeline for this is between about 10pm and 6am local time, so Marco uses the timebar to highlight this local time period. The results show up as UTC times in the Master polygon form.

Marco feels strongly that there will be accumulating snow in a narrow band from about KFSD to near KGRB, but away from the lake. Marco wants some snow to accumulate overnight, but not too much. He knows there are Filters to control the amount of snowfall over some time period and he selects the SNOW_2_4 Filter to accomplish this.

SNOW_2_4 works behind the scene by amplifying or deamplifying existing periods of precipitation to match the constraints of the Filter. Marco knows that temperatures must fall within a certain range and the POP needs to be above a wet threshold of 30%. If Marco is unsure of the other conditions he can actually see these in the Conditional form. He is comfortable with the change so he draws a Filter polygon, selects the Filter Name/Number and allows the Filter to simply act on all points that fall within the polygonal outline.

Marco does not establish any Blurbs for the cold part of this event. He considers writing one for the general

area where the wet snow may fall, but decides against it. He will confer with the Winter expert later in the evening to see if this situation changes.

Severe Weather Potential in the Southern Plains and Southeast

As the cold frontal system moves into the center of the country, a late season moderate threat of severe weather is taking shape from about eastern Kansas southeastward to Birmingham, Alabama. Its now about 10:00am EDT and Marco is starting to think about the afternoon. He uses HOTL to see the current radar in the southern plains and SE. He uses a version that shows precipitation type, polygonal warnings and any watch boxes that might be in effect. Marco underlays the visible satellite imagery below the radar. Marco uses a second screen to read about SPC outlooks for the area. He left clicks in the general area and inspects some of the going forecasts. He peels the forecast timebar back and forth through the afternoon hours to see the FBQ’s in effect for that area.

Marco wants to take a closer look at the 2-D forecast. He pulls up the 15-minute TruPoint forecast available for CONUS (a proprietary Nowcast) as an underlay, and he displays this current 6-hour Nowcast using the hourly precipitation amount variable as a proxy for future radar. He notices two things immediately. First, the Nowcast tends to advect existing cells too far northward in the vicinity of the existing mesoscale convective complex taking shape in Arkansas. So, instead of the MCS propagating the mass of thunderstorms east and southward it shows the system heading NE and almost NNEward. This needs to be addressed because it will impact the timing and position of the MCS over the next 6 hour period.

Marco uses a Filter to bump PoP about 20% over the polygonal area where the MCS is expected to move, knowing that some forecast on the southern edge of the precipitation region will move from dry to wet forecasts. Marco also uses a Severe Filter over a smaller area near the gust front and leading edge of the MCS to assure that if wet forecasts are triggered they will mention the possibility for strong thunderstorms.

Marco makes both of these filters of short duration. If they cannot be revisited in the next few hours, they will expire before they are overcome by events.

Finally, Marco writes an e-mail to his supervisor and the weather team pointing out the tendency for TruPoint or radar advection to undergo southward propagation near the leading edge of MCS's. This will not help the forecast center in this particular event, but may pay dividends down the road and is another responsibility of the human over the loop. That is, to record and document useful case studies of impactful weather and shortcomings, biases, or errors that occur in the weather and forecasting systems he uses in his day-to-day work.

The Look Ahead

Marco uses his second screen (non-HOTL) to inspect current medium range models, on-line analyses of those models, and the current internal discussion on the extended forecast for both the CONUS and Tropics. He plans to inspect and revise any extended FBQ's in CONUS, then spend time with the current Tropical system that is expected to become Tropical Storm Vicky by week's end.

Before he begins this phase of work he reviews the current actions in his domain of interest and submits all changes to the FBQ server database. The system acknowledges the commit of forecast changes. An icon on the upper r.h.s. of the screen changes appearance, to indicate that there are no pending FBQ adds, deletes, or edits related to the current session. The wallpaper of Marco's HOTL application also changes background color to indicate that the session status has changed to one of no pending or open changes.

Marco uses the Underlays capability to look at the currently published weather shields and maximum and minimum temperature maps for the 3 – 10 day period. He agrees with the overall progression of weather elements. Once the early season storm exits the country, a period of quiet Indian summer weather is likely to take hold with ridging in the middle of the country.

Marco sees two areas of concern or possible work. One is the near freezing in the Intermountain West at the same time as low-PoP showers arrive with an upper-level disturbance on days 4-7. He addresses this with a NOICE Filter that extends throughout the day 4-7 period. This assures that any precipitation that may be predicted for mountain points will be rain, rain and snow, or snow. Weather types such as freezing rain and sleet will not be allowed for this weather regime.

Secondly, Marco thinks the warming will be quite unusual for some parts of the Plains and Ohio Valley. Marco brings up underlays with text plots that show the maximum temperature anomaly for each day of the extended period. He sees that the anomalies peak at around 25F on day 6 near Evansville, Indiana. He draws a large area polygon in this region and uses an extended Qualifier for both days 5 and 6. The Qualifier suggests unusual warm temperatures or near record warm temperatures as possible in this timeframe.

A final change to the extended's will be made in the west for Santa Ana conditions setting up on day 3. Marco loads predicted surface pressure from the ECMWF and then an analysis of surface wind vectors. This confirms that the model is showing strong off-shore flow in the LA Basin beginning on the morning of day 3.

Marco pans and zooms to SW California. He leaves the wind vector analysis in place and scrubs the time-bar to the start time of the Filter which is Nov 2 at 15Z. Marco selects Filter Add. Instead of drawing a polygon he selects a pre-existing polygon called SANTA_ANA. A cluster of polygon shapes now fill the screen in and around the LA Basin; these are the predefined SANTA_ANA shapes. Marco assigns a time range of 15Z through 04Z and *recurs* this time range into the next day (day 4). He selects a Filter name of WINDY. This Filter assures that all points selected in the Conditionals have Wind Speeds that are equal to or above the "windy" threshold in the genome. This ensures Windy forecasts or Windy "modifiers" will attend all forecasts on Day 3 and Day 4.

Marco repeats the polygon predefined Select for SANTA_ANA and then adds an extended Qualifier that indicates the expectation of Santa Ana conditions on

Day 3. He does not *recur* the Qualifier to Day 4 for now since his confidence is a bit too low to warrant the extended Qualifier. For now, having a Windy modifier in the extended forecast is enough.

Marco is complete with over the loop attention to the Look-Ahead. He performs a Submit operation so these new and modified Filters and Qualifiers will be posted to the FBQ database. He sees the screen background change color and the system returns a message that all changes were posted to the FBQ database.

Static Weight Management (SWiM)

Marco selects SWIM from the HOTL main menu. This application lies outside of the HOTL graphical user interface and is not governed by HOTL GUI rules and behavior. HOTL simply facilitates loading SWIM and returns Marco to the HOTL main menu when Marco concludes his SWIM session(!).

During that hour window Marco manipulates static weights in the short term forecast period for the CONUS region of the world. Marco cannot directly see the potential integrated results from these weight changes. However, he is confident that he can push the gDicast guidance towards a more skillful integration because he inspects individual NWP outputs on a second screen and also reviews publically available analyses on the model solution trends.

The SWIM modifications can only be performed by Lead or senior staff. The actual static weight changes are saved into a trial or temporary area where numerous QC checks are performed. If these checks pass, the proposed weight changes are moved to production and staged for the next run of the gDicast Forecast Integrator. Marco is updated on the QC checks and staging as they occur, but this messaging takes place in the SWIM environment and is not a capability of the HOTL GUI/application.

Once weights are staged in production, Marco will write a forecast discussion describing the changes and his motivation for altering the static weights.

The Hurricane Wind Tool

Marco returns back to the HOTL GUI and visits the tropical forecasts for the two active systems. He is able to confirm that for the Atlantic Basin system that is likely to become tropical storm Vicky in about 24 hours does reflect the NHC strength and position estimates. A reasonable vortex is seen through the 120 hour NHC prediction timeframe arriving about 60 SM due south of Key West, FL by the morning of Nov 2.

The second system is in the West Pacific Basin and is not subject to the Hurricane Wind Tool intervention since bulletins are from the JTWC and not interpreted by the system. This hurricane Tetras looks to be recurving and is not threatening any international land masses in the foreseeable future so Marco allows the gDicast solution to flow unabated. He returns his attention to Vicky.

The NHC strength and position forecast is fairly consistent with gDicast fields of precipitation, temperature, and clouds. However, Marco wants to sharpen up the PoP fields on the leading edge of the storm particularly in the western Bahama's. He uses Filter deltas to clean out leading edge precipitation from arriving too quickly and bolstering PoP's near the predicted center of the storm. Marco also places Rain Squall Filters throughout the Bahama's and southern Florida coincident with the predicted storm arrival. Marco initiates an extended Qualifier for the Bahamas warning of tropical storm conditions, and a slightly different Q for southern Florida.

Marco writes a brief freestyle Blurb for Southern Florida describing storm timing and likely impacts. He lists surge potential for the coast from Pompano Beach to Key Largo as well as further south to Key West. He gives this Blurb an immediate activation (through time inputs) and allows it to expire noon the following day. It will be important to highlight this Blurb and its time parameters to the next shift so they can update the text (by simply editing the existing Blurb) and extend the expiration time further into the future.

End of Shift

Marco does his final commit of the day and turns his attention to documenting his recent shift via forecast discussions and turnover logs.

He prepares himself for an upcoming OCM and Producer briefing by talking to Experts, fellow shift forecasters and digital graphic designers.

Just before the briefing, Marco inspects 30-40 FOD forecasts using the left-click capability in HOTL. He is satisfied from these resulting or published forecasts that his weight changes and FBQ's have hit the weather forecast server and are reflected in the outward facing forecast content and products.

In all cases the forecasts fall within acceptable ranges and the HOTL actions have properly constrained the flowing gDicast forecasts.

Marco thinks the snowfall in the western Lakes is now probably overdone compared to his thinking and the trend of the forecasts some 8 hours ago. But his shift is completed. The evening forecasters will no doubt visit these accumulations and tune them appropriately. For now the published forecasts lie in an acceptable range of forecast skill.

Marco grabs his North Face lunch cooler and heads for the parking lot.