5.5 REGIME-BASED PERSISTENCE PROBABILITY OF CEILING AND VISIBILITY

Alister Ling* and Bjarne Hansen
National Laboratory for Aviation Meteorology
Environment Canada

Ismail Gultepe
Cloud Physics and Severe Weather Research Section
Environment Canada

1. INTRODUCTION

This paper describes some preliminary results from the Aerodrome Vicinity Guidance (AVGuide) project. The goals of the AVGuide project are to improve ceiling and visibility (C&V) guidance, to make it more relevant and easily accessible in forecast operations, review past efforts, to explore options and take first steps towards the automation of TAFs.

More than 180 TAFs (coded forecasts for airport terminals) for Canadian airspace are produced by forecasters working in the two regional offices of the CMAC (Canadian Meteorological Aviation Centre) of Environment Canada. Because each forecaster is typically responsible for 15 TAFs and graphical aviation products, it is important that guidance is available in near real-time in a high glance format, to help them focus on the hubs and where they can add value.

Despite a few decades of cumulative effort by several teams of various National Meteorological Services into developing automated TAFs, the autoTAFs tend to be used sporadically by operational forecasters as guidance or a second opinion, and are not considered “good enough” to be transmitted as official products. Some private companies do offer autoTAF services (not yet assessed in this project).

This state of affairs, combined with the TAF production model of the CMAC, prompted us to consider an alternative approach: to first identify “straightforward” TAFs with a high degree of confidence and to automate them, and flag the forecaster to the ones whose guidance is assessed to be of low confidence.

* Corresponding author’s address: Alister Ling, Room 200 Weather Office, 4999 98 Ave NW, Edmonton AB, T6R 2B1 Canada;
email: Alister.Ling@ec.gc.ca

The goal is not just a matter of increasing efficiency, but to ensure that guidance and products are of known performance levels and reliability from verification results. From this, end users gain a more appropriate sense of confidence and can make better decisions for their operations. Similarly, decision support systems that integrate or blend different input can be improved. After a literature review, discussions with researchers, operational forecasters, and end users such as airline dispatchers, it was decided to begin with an analysis of C&V occurrence and recurrence (persistence) and a simple definition of weather regime, in part based on ideas contained in Gultepe at al. (2007).

Persistence is a standard benchmark against which many algorithms are compared; a variety of authors (e.g. Rudack and Ghirardelli 2008, Black et al. 2008; and others) note that persistence is a competitor in the first hours of a forecast.

In section 3 the concept of weather regime for this study is provided. Section 4 and 5 present and discuss the results that demonstrate when persistence and duration statistics can be used in combination as a means to identify situations when persistence is expected to fail. Section 6 highlights parts of a prototype display for near real-time operational forecasting. Regime dependent behavior, specifically the character of variability is discussed in section 7. Our focus on regime-based performance has led us to question the appropriateness of using persistence as a forecast methodology, and of using recent past performance in forming a level of confidence in future results. These are raised in sections 8 and 9 respectively. In the latter the concept of estimating a measure of “future confidence” is proposed.

2. DATA AND DEFINITIONS

2.1 Dataset

The data for this work were provided by the Canadian Climate Center and consists of the hourly METARs of Canadian airports (more than 190 TAF sites) during the period from 1971 to 2005. Several of the regional airports operate on a reduced schedule or came into service part way through the period and thus have a smaller dataset.

Data used for verification are from an independent set from the year 2009.

A comment regarding SPECIs is made in the section discussing variability.

2.2 Flight categories

In this initial phase of the project, we have chosen to work with *flight category* to reduce the complexity compared to ceiling and visibility separately. Although conceptually equivalent to the U.S. (e.g. Bateman et al 2008) the most common thresholds at Canadian airports were used and are summarized in Table 1 for Visual Flight Rules (VFR) and Instrument Flight Rules (IFR), as well as the operationally relevant Marginal Flight Rules (MVFR) and Low IFR (LIFR).

Flight Rules	Ceiling (ft)	Visibility (mi)
VFR	> or = 2500	> or = 6
MVFR	< 2500	< 6
IFR	< 1000	< 3
LIFR	< 400	< 1

Table 1: Flight Rules and Associated Ceiling and Visibility Limits. Note the limits and inequalities differ slightly from those in the United States.

2.3 Persistence

A straightforward definition of persistence is used: two sequential observations are in the same flight category.

Duration: the number of hours a category persists. Given a category at 01:00Z of IFR, if the condition is IFR at 02:00Z but not at 03:00Z, then the duration was classed at 1 hour; if there was a change at 01:25Z, the duration was assigned 0 hours. From an operational perspective, the condition must "hold" until the top of the upcoming hour for it to count.

2.4 Weather regime (situation)

Section 3 contains a complete discussion on weather regime. For the purposes of this study, regime is defined by:

- surface wind direction in 45 degree bins
- occurrence and type of "significant" precipitation (larger than drizzle drops or ice crystals).

3. REGIME – CHOOSING THE RIGHT TOOL FOR THE JOB

3.1 Concept

"Dirty northwesterly flow", "moist, stagnant airmass" "numerous pulse thunderstorms" and "off-lake snowsqualls" represent a few examples of how forecasters might characterize particular weather regimes, or situations.

Through instruction and experience, forecasters typically have a good idea of how well certain algorithms/techniques perform in each regime and whether or not to use them. For instance, a radiation fog technique for forecasting the onset of low visibilities based on dewpoint spread should not be used when it is cloudy and windy. If little change is expected in crucial weather parameters, then current conditions are more likely to persist. If it looks like a cold front will cross in the next hour, then using persistence is likely a bad idea.

The key is to pick the algorithm(s) most suited to the situation. It becomes increasingly important when the current conditions are different from those that predominate overall. An algorithm that works well 80% of the time may perform very poorly in the current regime, but this result may go unnoticed when bulk performance numbers are examined.

In the quest to demonstrate the overall effectiveness of a prediction system, care should be taken not to summarily dismiss the uncommon occurrences because they contribute little to the overall score. Although a cold frontal passage crossing an individual airport in the next hour or two occurs only a small percentage of time, this is one of the events that matters.

In principle, an automated system could evaluate the regime and based on historical performance eliminate the poor algorithms while keeping the one that is best, or blend the results from the good ones. Better yet, it would also assess upcoming changes in regime, such as a cold frontal passage, and suggest the appropriate algorithms to use subsequently, both to the forecaster and to other decision support systems through numerical input.

One parameter of specific interest is the "crossover time" when the skill of persistence as a predictor drops below that of other algorithms. Fig. 1 in Golding (1998), a conceptual schematic, portrays the relative information content (or skill) of nowcasting and NWP models methodologies decreasing with lead time. Instead of using a long term average or its value in recent hours or days, the crossover time could be derived from past performance in the anticipated regime.

Most of the climatology-based guidance used in operational forecasting of C&V is presented as probabilities of specific flight categories. After initial stratification by season or month, the counts of the flight categories are typically further stratified by precipitation type and finally binned into hour of the day <u>or</u> surface wind direction. Trying to do both often resulted in the problems associated with empty bins.

This conditional climatology (CC) guidance has definite value, for example in estimating onset and dissipation of fog, and checking that forecast combinations of weather, wind direction, and C&V in a TAF are reasonable. However, it does not directly address the issue of change. Change can be a gradual transition or a rapid jump from one category to another, but it can also be a complex chain of events described as "highly variable conditions."

The research was motivated in part in trying to answer the question posed by both forecasters and end users: "What is the probability that current conditions are going to persist, in this situation?"

3.2 <u>Regime in practice - first step</u>

Most verification systems are too generic to answer the question pressing the operational forecaster: "how does this guidance perform in the current regime?" This approach appears promising; evaluating it for the persistence forecast has begun.

In order to more quickly investigate the potential of regime-based performance, the concept of regime was simplified to a basic form: surface wind direction in 45 degree bins and whether or not there was "significant" precipitation (larger than drizzle drops or ice crystals).

It is anticipated that increasing the complexity of the regime concept will improve the skill but incorporating too many parameters will result in well-known sample size issues (Wilks 2006). It is unknown how much sophistication is possible when regional aggregation of sites for synoptic scale events is performed.

4. RESULTS

It is demonstrated how regime (for this study surface wind direction and occurrence of precipitation) significantly affects the probability of persistence in the upcoming hours. Verification of the events when the probability goes to zero (i.e. that persistence is expected to fail) in the first few hours shows Frequency of Hits (FOH) near or above 0.70.

4.1 <u>Probability of Persistence (climate-based)</u>

The following sample graphs confirm the operational forecasting experience that the duration of persistence depends on the direction of flow and the initial category.

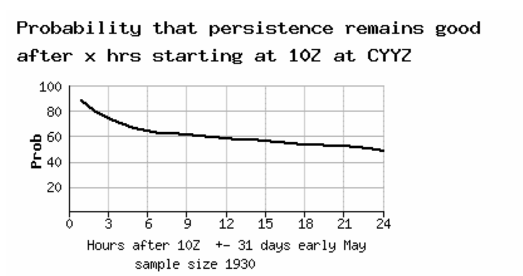

Figure 1. Probability is simply: (number of observations that have persisted) / (sample size). Not stratified by regime. Pearson International Airport, CYYZ, is Canada's principal hub near Toronto, Ontario, near the northwest shore of Lake Ontario.

Fig. 1 shows the probability that persistence will continue to hold (or not change) for each hour after 10Z in May, regardless of initial conditions or wind direction. There is a 70% chance that persistence will last for 3 hours, and a 60% chance that the current category holds for 11 hours. When a filter of LIFR is added, the story is considerably different as can be seen in Fig. 2. The probability that LIFR will persist drops to about 30% in 3 hours.

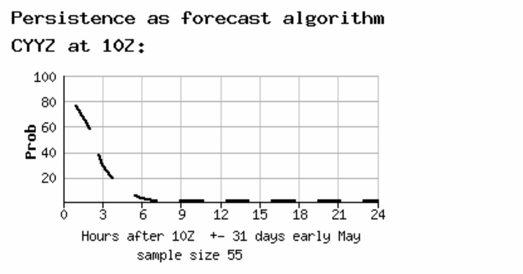

Figure 2. As is Fig. 1 except only those cases that were initially LIFR at 10Z were retained. The graph title was changed simply to convey the concept of persistence as an algorithm.

When further stratified by wind direction (Fig. 3), the climate data indicates that with a westerly surface wind, persistence of LIFR has a high probability of failing in the next two hours!

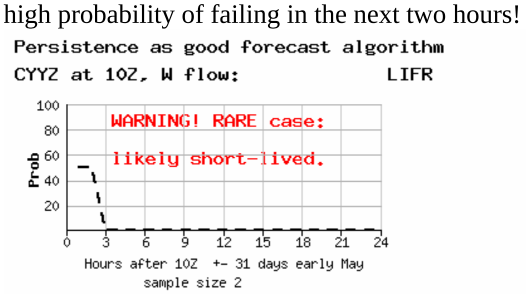

Figure 3. As in Fig 1. and 2. except a further stratification by wind direction (west in this case) has been made. This implementation can be termed conditional persistence climatology (CPC).

The system also indicates the sample size, which flags the decision support system as well as the forecaster that this is an uncommon situation and to treat results with suspicion. Instead of retreating behind larger bins to avoid small sample problems, it is desirable to highlight the likelihood that the output from most techniques will be unreliable and to place more attention on this forecast – humans should focus on where they can add value.

A probability of 0 signifies that the current condition persisting beyond a certain time has not occurred in the period of record. However, during the verification period some conditions did persist past the expected time, which implies that the probabilities are not truly calibrated, though they should be quite close.

The preliminary results, coupled with the notion of "right tool for the job" encouraged us to consider that conditional persistence probability can be a useful tool in identifying some cases when persistence fails (i.e. a change will occur), like that in Fig 5. Conversely, when conditions are highly likely to persist, it appears less useful as a tool: it would not be good at discriminating between conditions that do persist and the small percent of occasions when surprises happen.

Therefore in the following preliminary verification only the subset of cases when the conditional persistence probability went to zero in the first 4 hours are examined. Most C&V forecast systems note that persistence is often the poorest predictor beyond 4-6 hours (e.g. Black et al. 2008; Rudack and Ghirardelli 2008).

4.2 Preliminary verification

Frequency of Hits (FOH) is the characteristic that has been examined. The creation of ROC curves and a wider range of skill scores will be undertaken in subsequent work. Similarly, a more complete assessment of regime-based performance for conditional persistence probability will be performed at a later date.

From averification point of view (with a classic 2x2 contingency table), only hits and false alarms and the derived Frequency of Hits are considered. FOH = hits / (hits+false alarms). The event is "persistence shall fail," in other words, a forecast hit means a change in category, while a false alarm is that no change occurred. By definition, persistence as a method is incapable of forecasting a change.

The primary test was "when the system forecasts the probability of persistence to be zero by hour x, did persistence fail to reach hour x?" It should be noted that these results are for cases from November 25 to December 21, 2009 and precipitation (larger than drizzle or ice crystals) was present.

Secondary tests, to allow for an hour's grace were "did persistence fail to reach hour x+1" and two hours grace "did persistence fail to reach hour x+2?" Fig 4. depicts the FOH for the primary test for hour 2 is above 60 percent, and increases when the event does occur in the subsequent two hours.

The verification results for the other hours followed similar patterns, so they are listed together in Table 2.

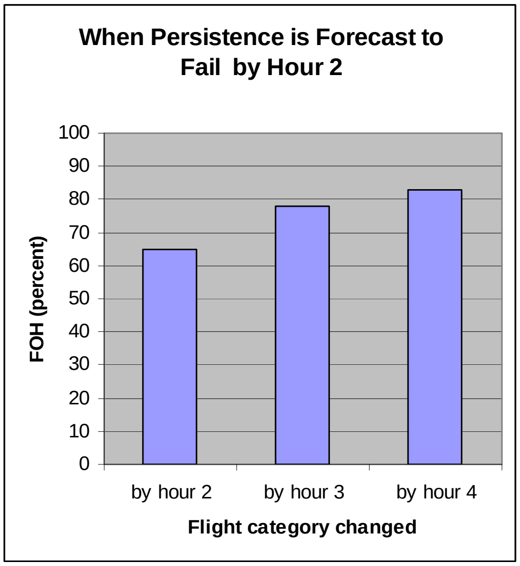

Figure 4. For the event "persistence shall fail by hour 2" (a category change in less than 2 hours), the FOH is shown for a change in conditions by hour 2, and if a more relaxed time condition is allowed, that the change arrived 1 hour late, (by hour 3), or 2 hours later than forecast (by hour 4).

	FOH	**FOH by hr+1**	**FOH by hr+2**
Pers. fails in hr 0-1	**39**	**60**	**69**
Pers. fails by hr 2	**65**	**78**	**83**
Pers. fails by hr 3	**76**	**84**	**87**
Pers. fails by hr 4	**84**	**85**	**86**

Table 2: Frequency of Hits, in percent, for the event "persistence shall fail by hour x" and when one and two hours' grace is permitted.

5. DISCUSSION

5.1 Discussion of primary results

The preliminary results demonstrate that one can use regime-based duration of persistence probability to forecast change, in particular when the probability goes to zero in the first few hours. For example, when an operational forecaster is presented with guidance that says "the chance of the current category persisting beyond 2 hours is zero", the system can also present the verification results from Table 2 that the forecast tool works more than 60% of the time, and more than 75% of the time if one allows an hour's grace.

Automated systems that use persistence as one of several inputs can benefit from this information by dynamically removing it from the integration scheme or changing its relative weight, instead of using an average weight or one based on recent performance.

The fact that the algorithm (probability of persistence is zero by hour x) can be wrong shows that a 30 year period of record is not long enough for the probabilities to be naturally calibrated.

The FOH is the lowest for the event "persistence will fail in hour 0-1", that is a positive forecast of change before the top of the next hour. This should not be surprising since these would be the most uncommon conditions hence less reliable. Furthermore, this data subset includes forecasts that are triggered from SPECIs with less than 15 minutes before the top of the hour. Conditions are much less likely to change in just 15 minutes.

The paucity of non-precipitation events (not shown here) suggests that increasing the stratification in the non-precipitation events is possible.

Also noteworthy is that allowing one or two hour's grace makes less difference with increasing hour. This suggests that the underlying statistics are more robust, though this has yet to be investigated. In addition it might be interesting to see if there is a commonality in the regimes when the conditions persist despite the forecast.

Subsequent analyses of a longer verification dataset will also set aside autostations in case their additional SPECIs distort the results.

It is planned to follow up on the investigation of lead-time metrics by Loughe et al. (2008).

5.2 <u>Confirmation of previous research</u>

We were initially puzzled that the forecast community has not made better use of this type of information. Martin (1972) discussed "Occurrence and Reoccurrence Profiles." He presented a graph (his Fig. 14) showing the probabilities that observed ceilings of <= 200 ft will reoccur 2 hours later. Although differing from our project in the details, the intent of "persistency characteristics of ceilings occurring under different synoptic situations" is similar.

We concur with several of his conclusions, and in particular that "the extension of these procedures and techniques to other weather elements and stratifying parameters is straightforward and recommended." Unfortunately there was little follow-up in the community at large. This may in part have been due to the challenges of data access and computing resources at the time.

5.3 <u>Regime and small sample size</u>

It is relatively straightforward to derive regime-based persistence duration and probabilities because of the long period of record of surface observations. A judicious selection of upper air parameters can be reconstructed from historical data available in datasets like the NARR (North American Regional Reanalysis) to help stratify the results.

It will be considerably more difficult to derive stable regime-based information from algorithms whose output cannot be historically reconstructed. For example some recent algorithms for C&V rely on upgraded NWP model variables that were not available two years ago. We will have to accept greater uncertainties and lower levels of confidence as a result of combining output regionally and broadening the date range as long as it is still appropriate.

6. IMPLEMENTATION OF A PROTOTYPE

At the CMAC (Canadian Meteorological Aviation Centre) forecasters are frequently responsible for more than 15 TAFs per person, in addition to producing graphical aviation products. It continues to be a challenge designing a dashboard display for monitoring TAFs that dynamically adjusts to the level of risk in a minimum of screen real estate. A popular utility TAFWarn (de Groot and Honch 2008) displays a warning when the TAF at some point in the near future becomes inconsistent with the current observation. That is, if conditions persist, the TAF will require amending (within 2 hours, in the case shown in Fig. 5).

TAFWarn has been adapted by adding a line to display a warning flag when current conditions are uncommon, and the probability that it will persist through each subsequent hour, using conditional persistence climatology as described above.

SITE	NOW(14Z)	1 HR (15Z)	2 HR (16Z)
CYYT VFR c/v persist	OK 42%	OK 0%	/CTGY/WX 0%

Figure 5. A portion of the TAFWarn_next prototype display. Beneath the timeline is the status of the TAF assuming conditions persist (red indicates a potential amendment). The bottom line has a flag for uncommon conditions, followed by the probability that they will persist to the top of the following hour.

The display is refreshed every two minutes, pulling information from a database that is updated with each TAF issue or amendment, SPECI, and hourly METAR for the more than 190 TAFs in Canadian airspace.

More detailed "drill down" information is available by mouse click. The portion of the content related to conditional persistence duration is shown in Fig. 6.

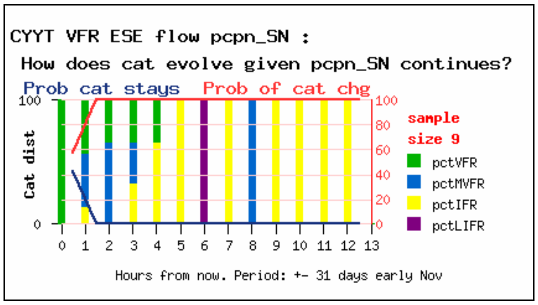

Figure 6. An example of conditional persistence duration and probability, for the case in Fig 5. A probability (that the flying category persists – blue line) dropping to zero in a couple of hours is one typical result of small sample size statistics, which we use as a flag that most algorithms, not just conditional persistence probability, are likely to be unreliable. In fact, the vast majority of uncommon events do not persist, otherwise they would be more common. This example also shows that by hour 2, all cases had changed category, but some changed back to VFR after the first hour.

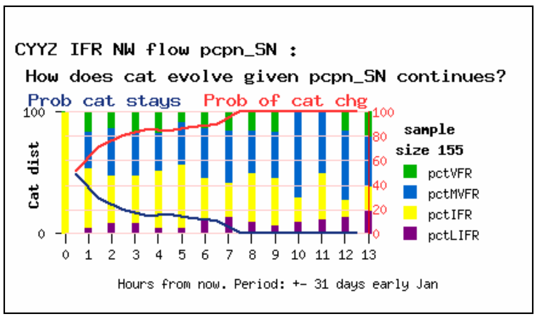

Figure 7. A typical example of conditional persistence duration and probability with a larger sample size. The vertical bars show the distribution of flight categories for each subsequent hour, but the probability line is offset because it is the likelihood the current category "makes it" to the top of the next hour.

7. VARIABILITY

One of the reasons for the interest in regimes is to improve forecasts of challenging situations, such as snowsqualls or fog from a nearby river wafting back and forth across the runway. In both of these examples, the C&V often swings wildly from practically zero to nearly wide open, two or three times an hour.

The classic approach of C&V forecasts is to use a probability distribution across the categories and choose the best category based on a set of rules and thresholds. But this does not capture the character of the highly variable conditions – in some cases a forecaster can confidently predict that the best forecast for a few hours is 1/2SM and P6SM NSW.

It may be tempting to argue that the overall percentage of times this happens is low and simply to ignore it. However, these are occasions when it matters to the client. Until meteorologists are capable of initializing and running models at sufficiently high resolution to reasonably capture the highly variable details, forecasters may be rendering a disservice to the end users if there is a policy requirement to minimize the range of categories in a part period of a TAF, resulting in numerous amendments.

Now that pilots self-brief over the internet, it is easier for them to launch on an improving condition and be surprised in tens of minutes when the poor C&V return. A future product that shows "probability of rapid change" could be quite useful for these times.

The fundamental characteristics of variability are range and frequency. A preliminary assessment can be made with hourly METARs; over the period of record, change will occasionally take place on the hour. Some of this variability is reflected in our results, when persistence fails to reach the top of the next hour. However, in order to properly tackle this challenge an analysis of SPECIs is required.

8. PERSISTENCE: AN INAPPROPRIATE FORECAST METHOD?

Fundamentally, forecasting the weather is about the ability to forecast change. Two common questions dispatchers ask our forecasters are "do you still think the ceilings will lift at 17Z?" and "could that fog form/dissipate earlier/later than your TAF currently indicates?"

Change in the weather is a primary driver for fuel loading decisions. Change however, is something that persistence as an algorithm is incapable of forecasting. At face value, persistence appears to be a reasonable benchmark against which to measure another algorithm's performance, but this is in hindsight.

Picture the case, not uncommon, of weather following a simple "step function" from P6SM down to 1/2SM FG and a few hours later a rapid improvement. Several verification methodologies will award persistence a good score overall, made up from the perfect forecast leading up to the change, a good score from a single category forecast during the fog episode, and a good score for constant conditions after the improvement. Of course there will be some penalties for the miss and false alarm at the transition points. Most TAFs, in order to account for uncertainty in the timing and depth of the fog, will contain TEMPOs which trigger numerous miss and false alarm counts. Even if the TAF did not require an amendment, the persistence forecast typically has a better score. In hindsight, it appears to have skill for "correctly" forecasting constant conditions for a few hours.

From the dispatcher's point of view, persistence as an algorithm would be obviously inappropriate and unacceptable: conditions are never expected to deteriorate, then suddenly they are bad and never expected to improve, except without notice conditions are once again good forever.

In summary: the duration of persistence in various regimes can be a useful tool for evaluating the likelihood of change, but using persistence as a forecast algorithm should be reconsidered. Lead-time verification metrics better reflect the reality that persistence is not capable of forecasting what matters: change.

9. REGIME-BASED "FUTURE CONFIDENCE" AND REAL-TIME VERIFICATION

Another common question that dispatchers ask of forecasters is "what is your confidence level in your TAF?" Because forecasters have a reasonable idea of their skill in each regime, their confidence can be used by dispatchers to modify the amount of hedging, or "discretionary" fuel loading they have, on top of the regulatory requirements. V-CMAC, a convenient situational awareness interface that encourages unobtrusive and rapid dispatcher-forecaster interaction via text, has resulted in substantially better fuel loading decisions (Kothbauer et al. 2008).

One of the benefits anticipated from regime-based algorithms and verification is an automated assessment of the confidence level and anticipated skill of the forecast in the future, based on the calculated regime for that point in time. For example, in the situation of an approaching cold front, the dispatcher will want a degree of confidence in the forecast conditions after frontal passage, and in principle a regime-based system can provide it. The performance of the forecast in the recent past and up to the current time is irrelevant and potentially misleading in this situation.

Current real time verification systems, such as CIWS (Corridor Integrated Weather System) in their web-based help file state "Accuracy for the current time is a measure of how well the forecast that was made an hour ago compares to the current weather. It is not an estimate of how good the current forecast will be at the validation time." Yet it is acknowledged that "The Forecast Accuracy scores enable traffic managers to quickly assess the performance of the CIWS RCWF product, thereby instilling a sense of confidence in the forecast products." This is okay as long as there are not impending changes.

Because weather tends to be relatively lengthy periods of persistence punctuated by short episodes of change, it is human nature to remember the times a system performed as expected and forget the times it did not. Too easily we forget that "past performance does not guarantee future results."

While regime-based "future confidence" will not be perfect, it promises to be more appropriate. One of the intermediate values calculated in Hansen's (2007) C&V forecast system (known as "WIND") is a set of similarity indices for the historical analogs, which we have been using as a simple way of estimating the degree of confidence at each hour. It is hoped that the verification process will begin soon.

10. CONCLUSIONS

A project to investigate the potential of improving ceiling and visibility (C&V) guidance and take first steps towards TAF automation has begun in Environment Canada's National Laboratory for Aviation Meteorology. The initial work in this project has focused on the concept of regimes, and the use of persistence as a forecast algorithm.

10.1 <u>Preliminary results</u>:

- there is value in forecasting the failure of persistence when the regime-based conditional persistence probability goes to zero in the first few hours. It is least valuable when the probability is high.

- persistence as a forecast algorithm should be reconsidered. Its appearance of skill comes from hindsight. In real-time, it is incapable of forecasting change, which is a primary driver in fuel loading decisions.

10.2 Future effort and possibilities

- There is a need to expand the period of verification and investigate more characteristics, including the ROC.

- Exploration of highly variable regimes could benefit end users and forecasters alike by creating an awareness of the potential for rapid changes.

- Regime-based algorithms and verification should allow an automated assessment of the confidence level and anticipated skill of the forecast in the upcoming regimes.

10. ACKNOWLEGEMENT

The authors wish to acknowledge: the contribution of discussions and ideas over the years from the operational forecasters at the CMAC (Canadian Meteorological Aviation Centre) and colleagues at the National Laboratories of the MSC (Meteorological Service of Canada); the advice and help on computers and scripts from Erik de Groot, Jack Dunnigan, Ron Goodson and Michael Skarupa.

The views expressed do not necessarily represent the official policy or position of the MSC.

11. REFERENCES

Black, J.L., Bateman, R., Herzegh, P.H., Wiener, G., Cowie, J., and Kessinger, C.J. 2008: An Automated National-Scale Ceiling and Visibility Forecast System: Development Progress. Preprints, *13th Conference on Aviation Range and Aerospace Meteorology*. New Orleans, Amer. Meteor. Soc., P7.2.

de Groot, E., and Honch, R., 2008. TAFWarn – a situational awareness tool: helping forecasters monitor TAFs in an age of information overload. Poster, *13th Conference on Aviation Range and Aerospace Meteorology*. New Orleans, Amer. Meteor. Soc., P1.28.

Golding, B.W., 1998: Nimrod: A system for generating automated very short range forecasts. Meteorol. Appl., **5**, 1-16.

Gultepe, I., Tardif, R., Michaelides, S.C., Cermak, J., Bott, A., Bendix, J., Muller, M.D., Pagowski, M., Hansen, B., Ellrod, G., Jacobs, W., Toth, G.. and Cober, S.G. 2007: Fog Research: A Review of Past Achievements and Future Perspectives. *Pure appl. Geophys*, **164**, 1121-1159.

Hansen, B., 2007: A fuzzy logic-based analog forecasting system for ceiling and visibility. *Wea. Forecasting*, **22**, 1319-1330.

Kothbauer, M., Laroche, S., and Larochelle, B. : Use of new prototype products from the Canadian Air Navigation System : 'TAFplus' and 'V-CMAC' : the dispatch perspective. *2008 Congress of the Canadian Meteorological and Oceanographic Society*. Kelowna, Canada. 1B5.2

Loughe, A., Madine, S., Mahoney, J. and Graf, M. 2008: A lead-time metric for assessing skill in forecasting the onset of IFRevents. Preprints, *13th Conference on Aviation Range and Aerospace Meteorology*. New Orleans, Amer. Meteor. Soc., P11.4.

Martin, D.E., 1972: Climatic Presentations for Short-Range Forecasting Based on Event Occurrence and Reoccurrence Profiles. *J.App. Met.*, **11**, 1212-1223.

Rudack, D.E., and Ghirardelli, J.E. 2008: Statistical guidance for terminal aerodrome forecasts from the Localized Aviation MOS Program. Preprints, *13th Conference on Aviation Range and Aerospace Meteorology*. New Orleans, Amer. Meteor. Soc., P1.1.

Wilks, D.S., 2006: *Statistical Methods in the Atmospheric Sciences*. Academic Press, 627 pp.