DANISH METEOROLOGICAL INSTITUTE

ÄÄÄÄ TECHNICAL REPORT (ONLINE) ÄÄÄÄ

98-16

Forecast verification report

September 1998

Anna Hilden, Ove Kjær, Rasmus Feldberg
and Jakob Hansen

DMI
COPENHAGEN 1998

ISSN: 1399-1388,
1397-839X

Preface

The present report is the third in a series of verification reports from the DMI weather service. These reports are given out twice a year and contain central results from the verification of the Institute's forecast products. They are distributed within the DMI, as well as to collaborators and other interested parties outside the Institute. Further copies can be obtained from the Institute.

September, 1998

Anna Hilden

List of contents

1. Conventions

1.1 Seasons

All verification results presented in this report are for seasons of three months each: Winter is December through February, spring is March through May, etc. Winter seasons are numbered according to the calendar year in which they end; winter 1995 is thus December 1994 to February 1995.

1.2 Abbreviations

ME	Mean error
MAE	Mean absolute error
Hit2	Hit rate (percent correct), tolerance 2 degrees
Hit4	Hit rate (percent correct), tolerance 4 degrees
HR	Hit rate (percent correct)
HK	Hanssen-Kuypers' skill index (times 100)
FAR	False alarm rate
FOD	Frequency of detection
HI	Human intervention

For definition of HK, see [3].

2. Point temperature forecasts

2.1 The product

For several years the DMI has been producing 5-7 day point forecasts in number code of main weather parameters for five locations (airports) in Denmark: 06030 Aalborg, 06110 Vojens/Skrydstrup, 06120 Odense/ Beldringe, 06170 København/Roskilde, and 06190 Bornholm/Rønne. The forecasts are used in several end products, including the DMI tourist weather telephone service and the fax product 'Havnevejr'. In most applications the forecasts are presented as area forecasts rather than point forecasts.

The parameters forecast are the following: Maximum day temerature, minimum night temperature, amount and type of precipitation, 24-hour maximum mean wind and gust speed, direction of maximum mean wind, daytime mean cloudiness, and 24-hour accumulated potential evaporation. Except for the last parameter, which is predicted by an automatic system, the forecasts are subjective, produced by the duty meteorologist by editing a table on the screen. The day 0 subjective forecasts are entered at about 5 o'clock in the morning, the rest around noon.

2.2 Verification

2.2.1 General procedure

Only the predictions of maximum and minimum temperature are routinely evaluated by the weather service. For each month and station we calculate the verification measures mean error, mean absolute error and percent correct within 2 and 4 degrees. The numbers are then averaged over the five stations and over each season to obtain the results presented here.

2.2.2 Maximum temperature

Figures 1 a-c show results for maximum temperature for forecast days 0, 3 and 5 for the period spring 1993 through spring 1998. Signatures are explained in the figure captions and legends.

Maximum temperature is generally well predicted for day 0, with Hit2 above 80 percent for most seasons and MAE always below 2 degrees. The bias (ME) is everywhere small. Results for cold seasons are better than those for warm seasons.

The quality of the forecasts decreases with forecast length, with growing differences between the results for cold and warm seasons. For day 5 the hit rates (Hit2) are seldom above 70 percent and for some warm seasons barely exceed 50 percent. Three seasons have more than 20 percent 'misses', cases with forecast error greater than 4 degrees. The MAE is generally between 2 and 3 degrees.

No overall trend in quality is apparent.

2.2.3 Minimum temperature

Figures 2 a-c show the corresponding results for minimum temperature.

The minimum temperature forecasts are generally poorer than those of maximum temperature. This might be taken as a confirmation of the general conception that minimum temperature is more difficult to predict than maximum temperature, but is most probably in part due to the fact that lead time for minimum temperature is about 12 hours longer than for maximum temperature of the same forecast day.

For day 0 the forecasts are correct within 2 degrees in more than 70 percent of cases for most seasons. The MAE rarely exceeds 2 degrees. Most of the seasons up to autumn 1996 show a bias of around minus 0.5 degrees, i.e. a general under-estimation of the minimum temperature; for the later seasons the bias is small and positive. There is no clear seasonal variation in the verification measures.

As with maximum temperature, the quality goes down with increasing forecast length. The day 3 forecasts have Hit2 values of 60-70 percent and MAEs of 2-3 degrees; the day 5 forecasts are a little poorer. Up to autumn 1996 the day 3 forecasts are negatively biased, but not the day 5 forecasts, indicating that the forecasters' pessimism was most pronounced for the beginning of the 5-day forecast period. After autumn 1996 the bias is mostly positive, especially for day 5.

Except for the bias the figures show no trend in quality over the years.

2.3 Comments on the results

Generally speaking, the point temperature forecasts are good for day 0 and acceptable for days 3 and 5. The forecasts of maximum temperature for day 5 have in some spring and summer seasons not been quite satisfactory.

Figure 1: Verification of point forecasts of maximum temperature. Percent correct within 2 and 4 degrees (light coloured and total columns, resp.; left scale), mean error and mean absolute error (triangles and squares, resp.; right scale).

Figure 2: Same as Figure 1, but for minimum temperature.

Maximum temperature

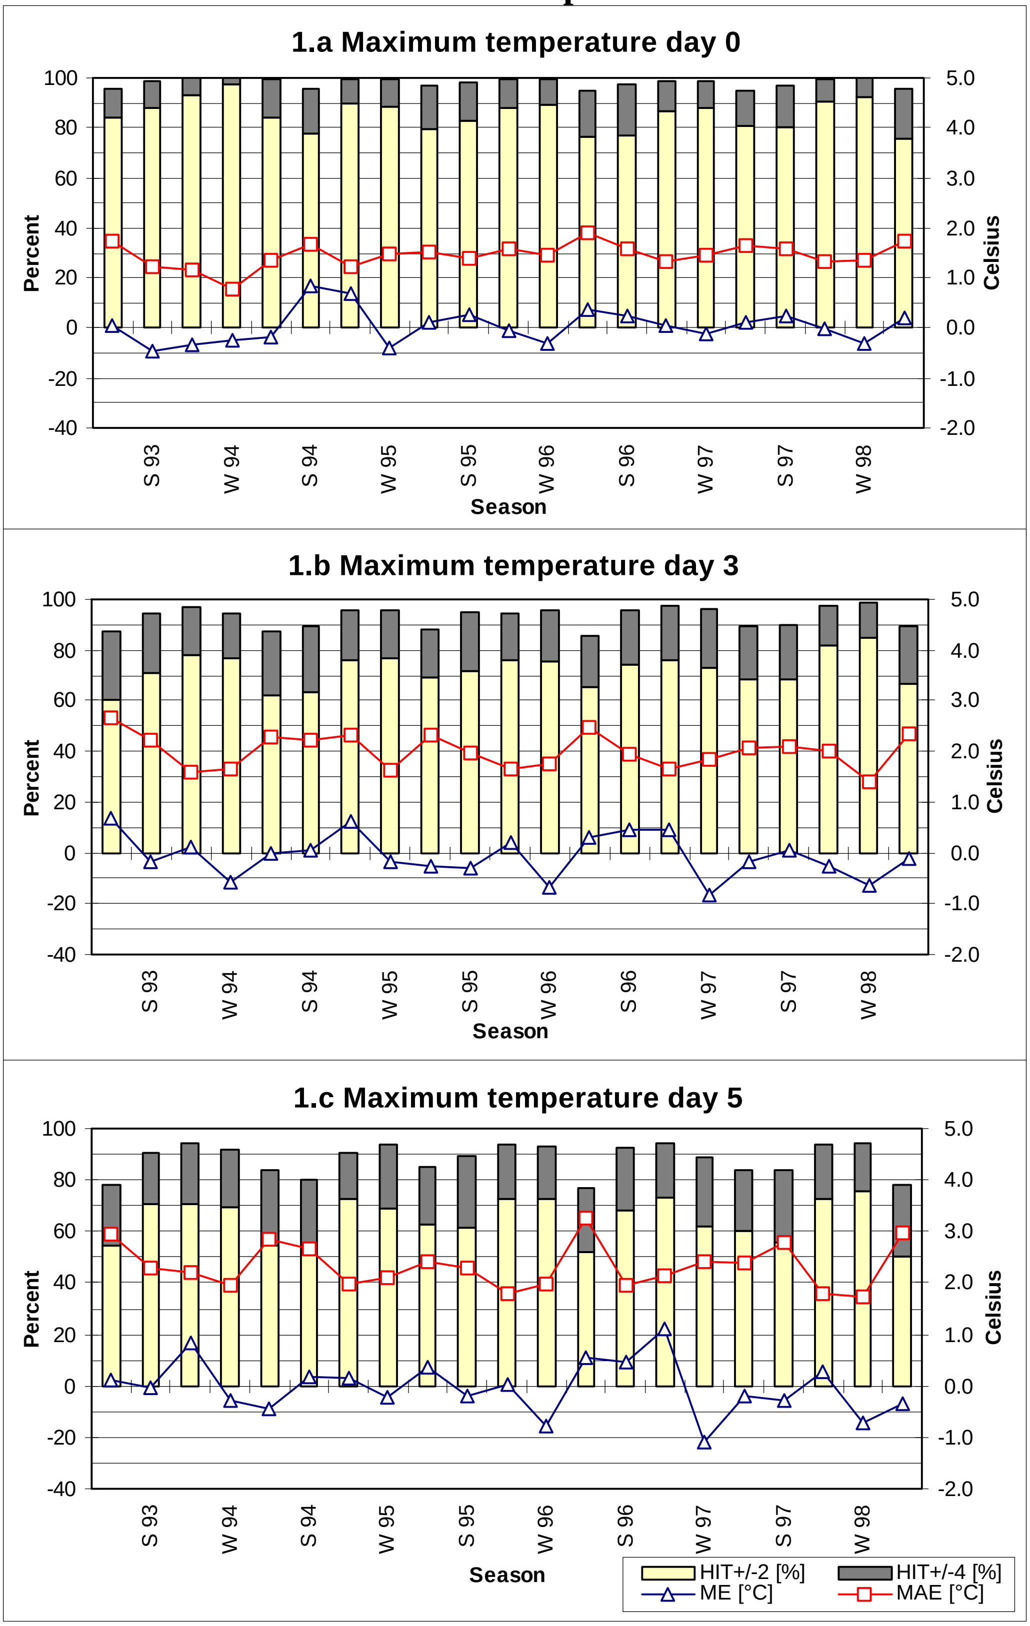

Minimum temperature

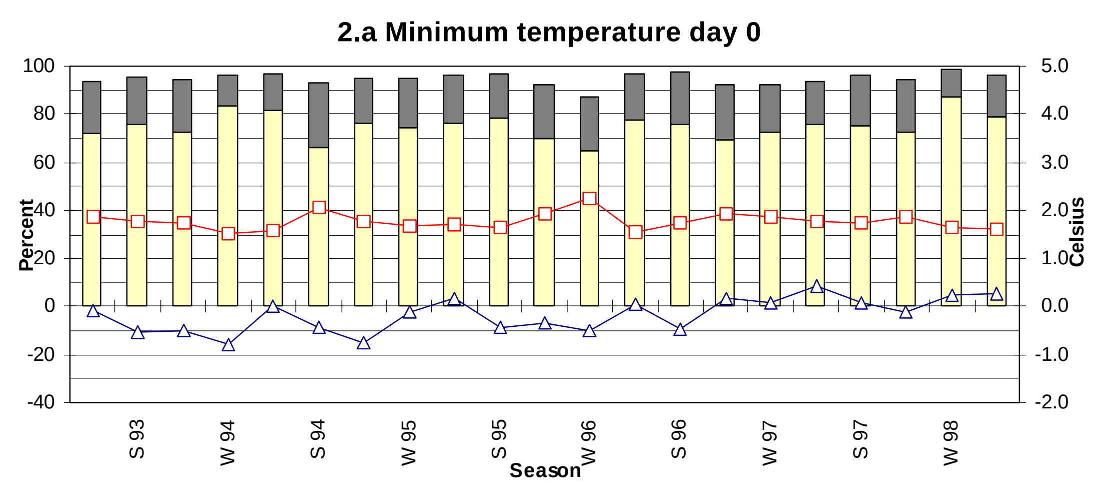

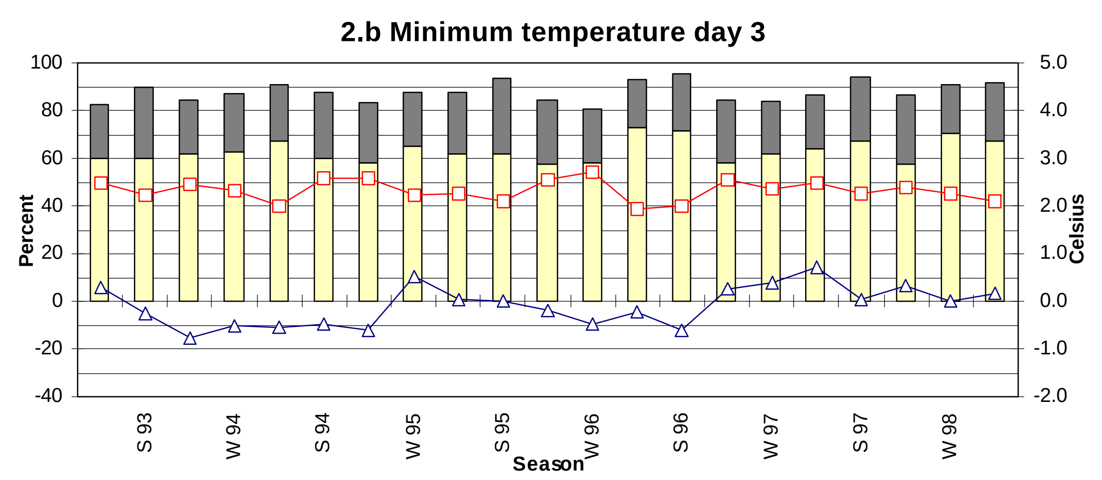

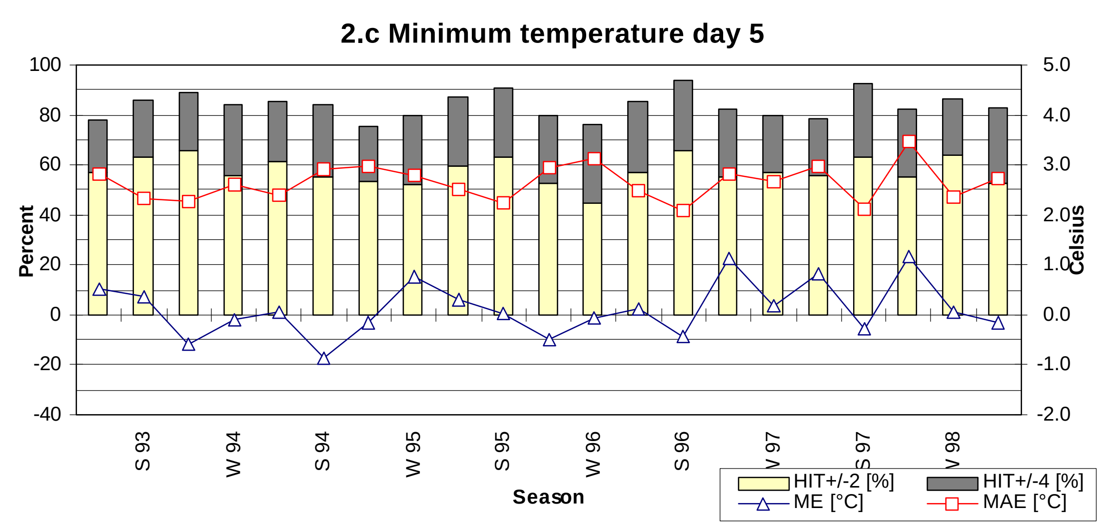

3. Marine forecasts

3.1 The product

The DMI weather service operates a semi-automatic system for generating coded 24-hour forecasts of wind, visibility and significant weather for the Danish waters. The system produces forecasts for the DMI tourist weather telephone service and for 'Havnevejr' and has been operational since August 1994. A description of the system can be found in [4].

3.2 Verification

3.2.1 General procedure

Verification of the marine forecasts is performed by comparing the coded forecasts to observations from coastal stations surrounding each district. Based on the number of suitable observing stations we have selected two districts, Kattegat and the German Bight, for detailed study; for these we use observations from 7 and 9 stations, respectively.

All verification measures are computed for lead times +6 and +18 hours, as well as for all lead times collected. For each forecast we verify also the corresponding first-guess wind forecast, produced automatically from HIRLAM model output winds. The model used is DKV up to and including Autumn 1996 and E15 thereafter.

The valid period of the forecasts verified here begins at 06 UTC. Only selected measures for the whole 24-hour period are presented. Due to upgradings, the system has been running in a non-operational mode from November 26, 1996 to May 1, 1997; thus, there are no results for winter and spring 1997.

3.2.2 Gale warnings

Predictions of the occurrence of near gale over 24 hours are verified by means of two-by-two contingency tables. Gale is considered as forecast if winds of 15 m/s or more were forecast for some part of the period; it is considered as observed if any station at any time of the period reported winds of at least 15 m/s. Observations from every whole hour are used when available. Trivial cases with gale force at the beginning of the forecast period have not been excluded from the data set.

Figure 3 a shows the false alarm rate and frequency of detection for gale in Kattegat for the seasons considered, together with the number of gale occurrences in each season. Figure 3 b shows the same for automatically generated forecasts.

The quality of the subjective (S) forecasts varies quite a lot from season to season, probably due to the small number of gale instances in Kattegat. The detection rates of the subjective forecasts range from 50 to over 90 percent, while the false alarm rates are in the interval 20 to 50 percent. No overall bias is apparent (FARs and FODs sum to around 100 percent).

Not so with the automatically produced (A) forecasts, however. Up to the break in 1997 there is a marked bias towards predicting gale too infrequently; this bias is reversed for later seasons.

Figures 4 a-b show the corresponding results for the German Bight. The gale frequency shows a marked annual cycle with a maximum in the cold months and very low numbers in summer. FARs and FODs for both types of forecast follow this cycle, the results being very good in winter, but poor in summer. The S forecasts are somewhat biased towards over-forecasting of gale, while the A forecasts are essentially unbiased.

Our verification procedure also adresses the more general problem of predicting maximum wind speed over the forecast period. The results for Kattegat show that the S forecasts are positively biased for most seasons, while the A forecasts are most often negatively biased up to 1997. For the German Bight both the S and the A forecasts for most seasons have a small positive bias up to 1997. As measured by hit rate and HK skill index, the S forecasts outperform the A forecasts in 5 of 13 seasons for Kattegat and in only 2 of 13 seasons for the German Bight.

3.2.3 Wind forecasts

Wind forecasts for individual 3-hour projections are verified using contingency tables for wind speed and direction for each of the interval widths which may occur in the forecasts (2.5 and 5 m/s for wind speed, 45, 90 and 180 degrees for direction). Each single observation is counted. Verification for the whole valid period is computed by adding up the numbers for the 3-hour projections. For wind direction, cases are stratified according to whether or not the maximum predicted wind speed exceeds 8 m/s. Only results for normal interval widths (5 m/s and 45 degrees, respectively) and - for wind direction - for stronger winds will be shown here.

Figures 3 and 4 c-d show hit rates for wind speed and direction for subjective and automatically generated forecasts for the two districts.

In general the results are not impressive. Hit rates for either parameter only rarely exceed 70 percent. Both S and A forecasts are better for Kattegat than for the German Bight; a possible explanation is that the wind field over Kattegat at any one instant is typically more homogeneous than that of the German Bight, which is a much larger area. All the figures show an annual cycle in the quality: In summer, wind speed is easier and wind direction more difficult to predict than in winter where the winds are typically stronger.

For both areas and all but a few seasons the A forecasts have higher hit rates than the S forecasts. Other results not shown here are in agreement with this, showing that the A forecasts are in most cases better than the S forecasts in predicting individual wind speeds and directions.

3.2.4 Visibility

Visibility is forecast in 4 categories: Good, moderate, poor and very poor. The limits are 10, 2 and 0.5 km. More than one category may be forecast for any given time, like in 'Moderate to good visibility' or - implicitly - in 'Rain showers, otherwise good visibility'. The verification is performed by means of contingency tables with overlapping forecast intervals. To cope with cases like the one just mentioned, an appropriate visibility category has been assigned to every weather type that may occur in the forecasts.

The results presented here are for the minimum predicted value over the 24-hour period, verified against the minimum value reported; hourly observations are used when available.

Figures 3 e and 4 e show hit rate and HK skill index for minium visibility for the two districts. The German Bight forecasts are poor, with hit rates below 50 percent and HK values mostly below 20 percent. The results for Kattegat are better, although not impressive; the best hit rate values are around 60 percent.

An investigation of other verification results for weather and visibility show that many weather types that are followed by a reduction in visibility, especially fog, tend to be under-predicted for both areas.

3.2.5 Precipitation

A simple verification of precipitation is performed using contingency tables with the four categories 'Dry', 'Liquid', 'Frozen' and 'Mix'. For the 24-hour verification considered in this report the categories are defined in the following way:

The forecast category is 'Dry' if no precipitation was forecast, 'Liquid' if for some part of the period either rain or drizzle was forecast, but no frozen precipitation, 'Frozen' if the converse was the case, and 'Mix' if both liquid and frozen precipitation were mentioned in the forecast.

The observed category is 'Dry' if no precipitation was reported at any station and time in the period, 'Liquid', if some station at some time reported liquid precipitation, but no stations reported frozen precipitation at any time, 'Frozen' in the opposite situation, and 'Mix' if both liquid and solid precipitation occurred at the stations in the period. Hourly observations are used whenever possible.

Hit rate and HK index for 4-by-4 precipitation are shown for the two districts in figures 3-4 f. The results are quite good for warm seasons for Kattegat, but too poor for Kattegat cold seasons and generally for the German Bight. Inspection of the contingency tables, not shown here, reveals a general under-prediction of precipitation and too many cases where either liquid or solid precipitation was forecast, while precipitation of both types was reported.

3.3 Comments on the results

Taken as a whole, the verification statistics for Kattegat and the German Bight are not impressive. The moderately high hit rates for wind speed and direction at individual times may be simply a reflection of the fact that nature's variability exceeds what can be expressed in the forecasts. For gale warnings the quality seems closely tied to gale frequency; if the number of gale occurences is high, so is the quality. The results for precipitation and especially for visibility are quite poor; here the main problem seems to be that the forecasts are too optimistic.

The results for wind demonstrate that human intervention may have a positive impact on gale forecasts, whereas the effect on forecasts of maximum wind over 24 hours in general, as well as on forecasts of individual wind speed and direction, is largely negative. For Kattegat the forecasters were able to compensate for the negative bias found in the automatically generated forecasts up to 1997.

Figures 3 a-b: Verification of gale warnings for Kattegat. False alarm ratio and frequency of detection (light and dark columns, resp., left scale), number of gale occurrences (rhombs, right scale).

Figures 3 c-d: Verification of wind speed and direction for Kattegat. Hit rate for wind direction and speed (light and dark columns, resp.).

Figure 3 e: Verification of visibility for Kattegat. Hit rate and HK index (light and dark columns, resp.).

Figure 3 e: Verification of precipitation for Kattegat. Hit rate and HK index (light and dark columns, resp.).

Figures 4 a-f: Same as Figures 3 a-f, but for the German Bight.

Kattegat

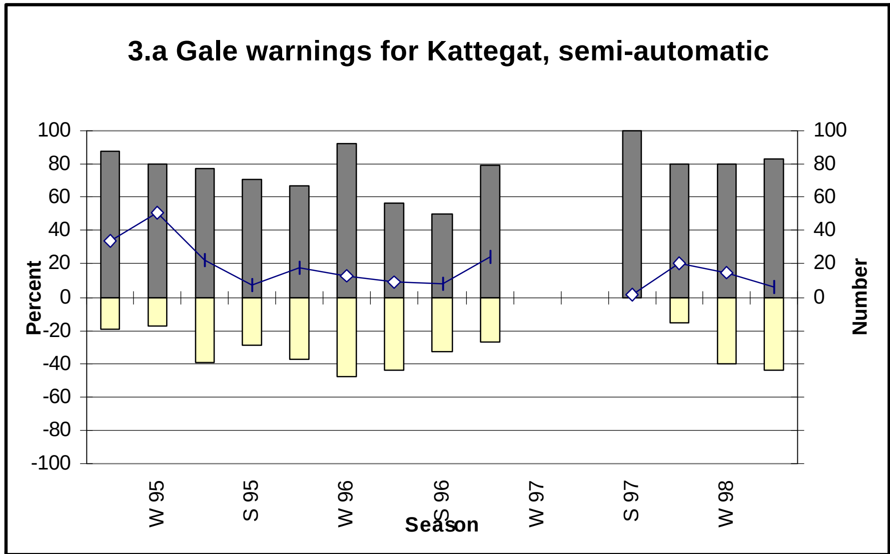
3.a Gale warnings for Kattegat, semi-automatic
Percent
Number
Season
W 95
S 95
W 96
S 96
W 97
S 97
W 98

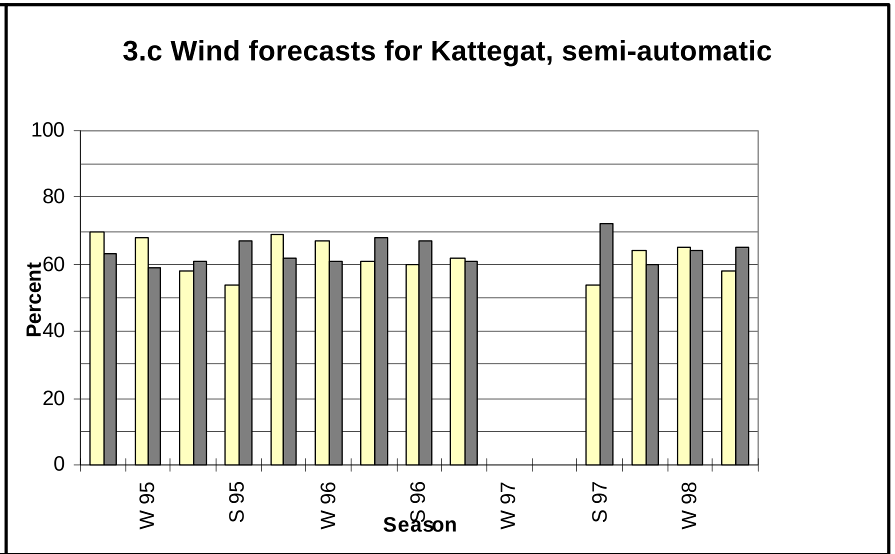
3.c Wind forecasts for Kattegat, semi-automatic
Percent
Season
W 95
S 95
W 96
S 96
W 97
S 97
W 98

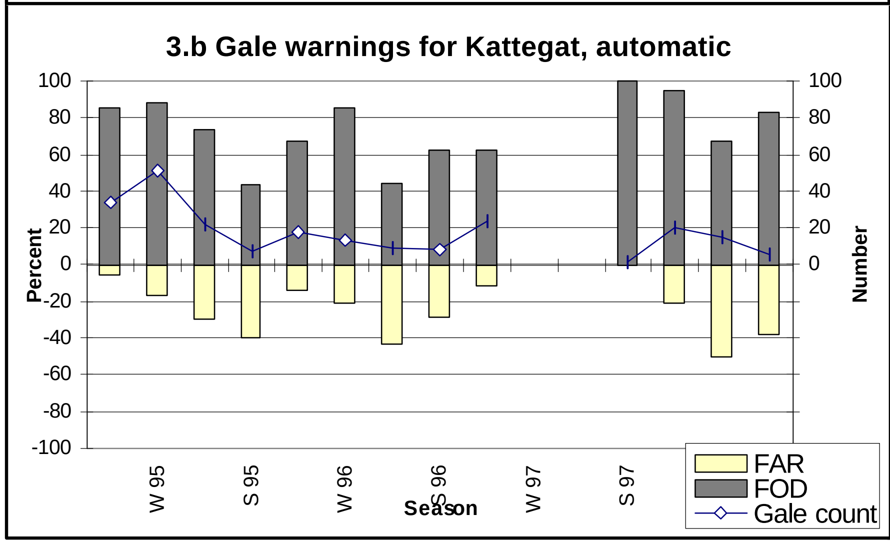
3.b Gale warnings for Kattegat, automatic
Percent
Number
Season
W 95
S 95
W 96
S 96
W 97
S 97
FAR
FOD
Gale count

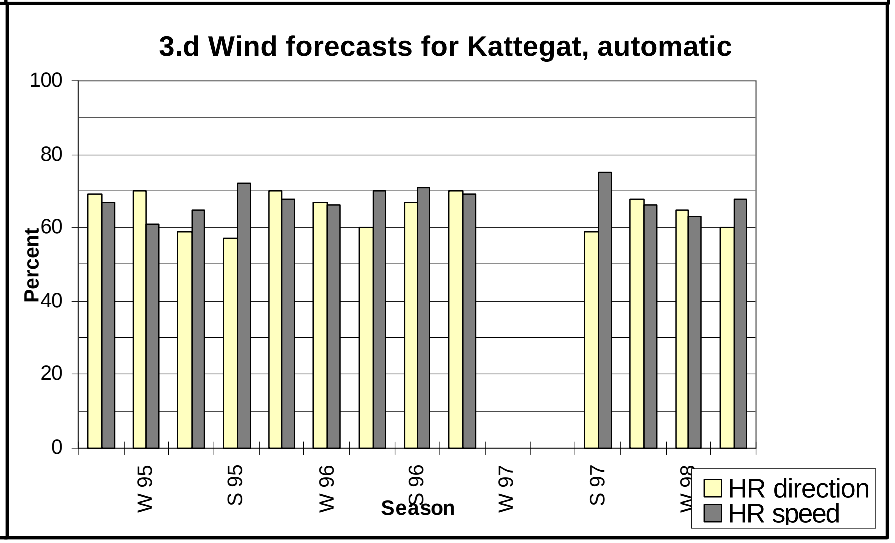
3.d Wind forecasts for Kattegat, automatic
Percent
Season
W 95
S 95
W 96
S 96
W 97
S 97
HR direction
HR speed

Kattegat

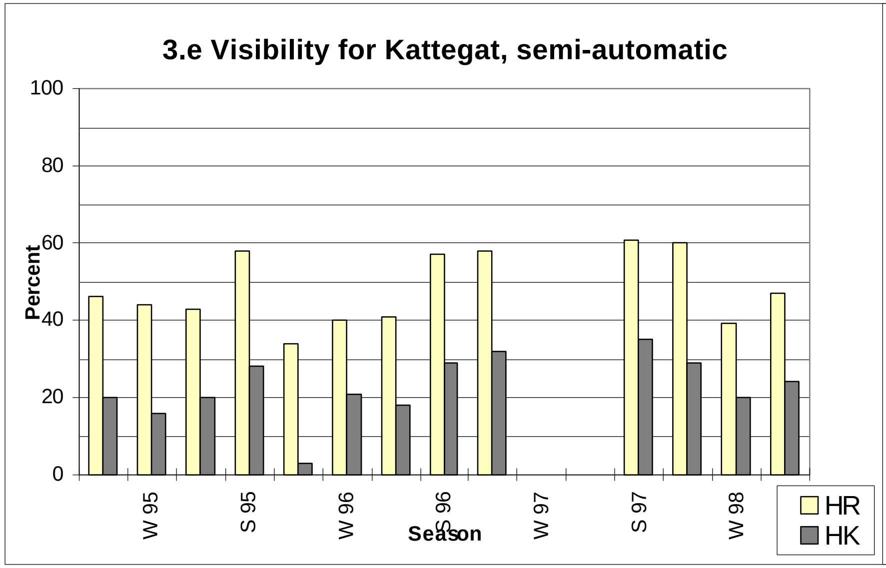
3.e Visibility for Kattegat, semi-automatic
100
80
60
40
20
0
Percent
W 95
S 95
W 96
S 96
W 97
S 97
W 98
Season
HR
HK

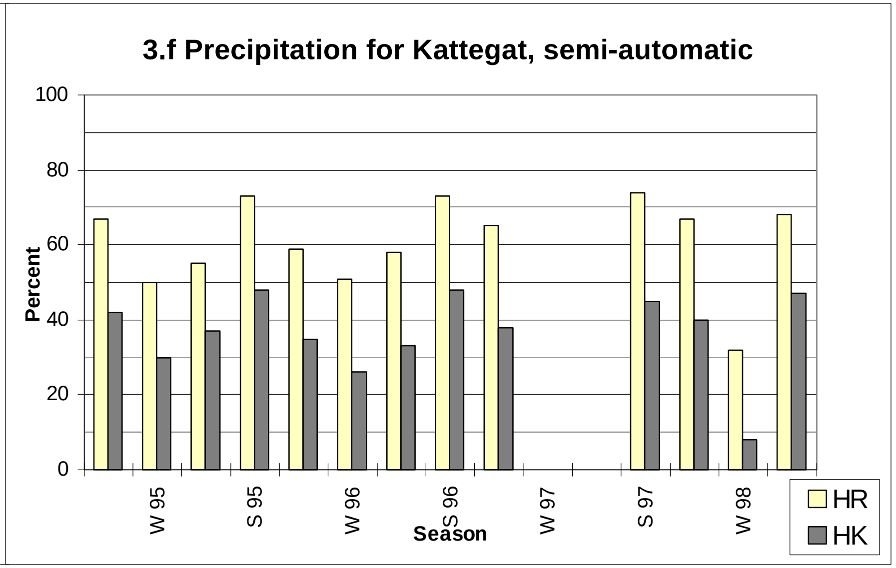
3.f Precipitation for Kattegat, semi-automatic
100
80
60
40
20
0
Percent
W 95
S 95
W 96
S 96
W 97
S 97
W 98
Season
HR
HK

German Bight

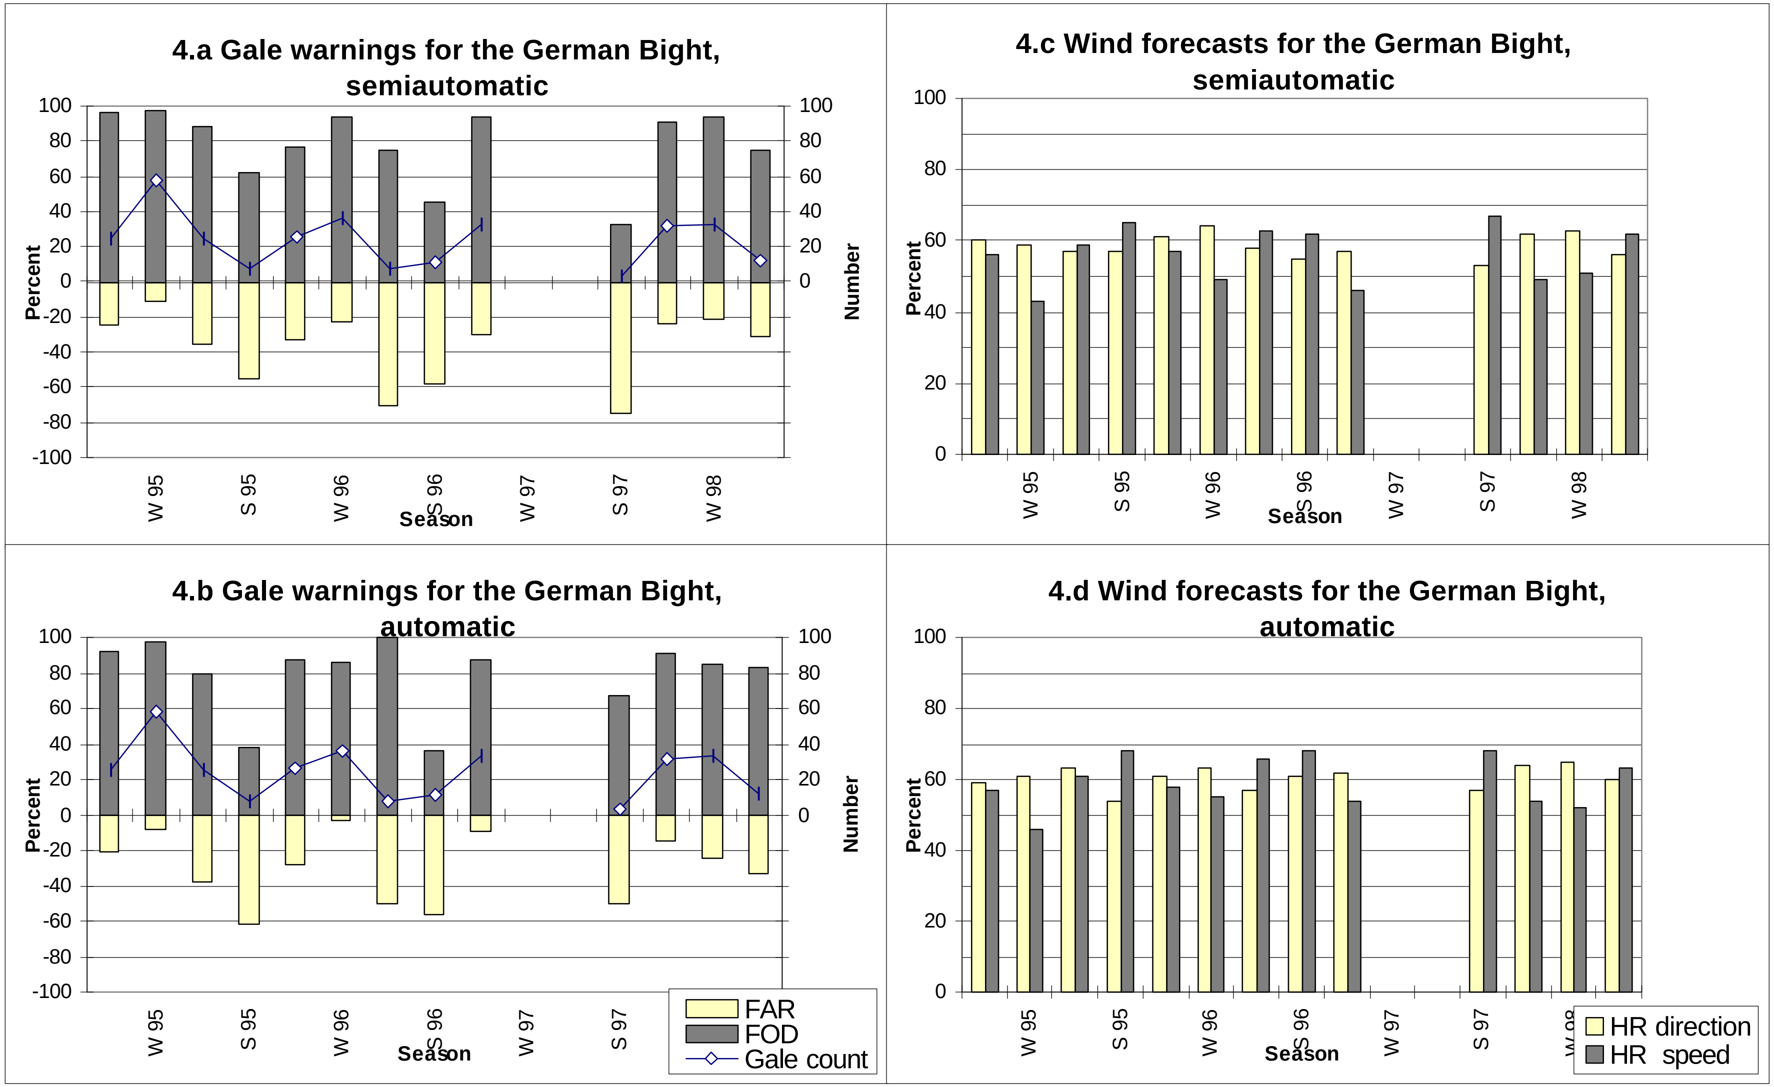

German Bight

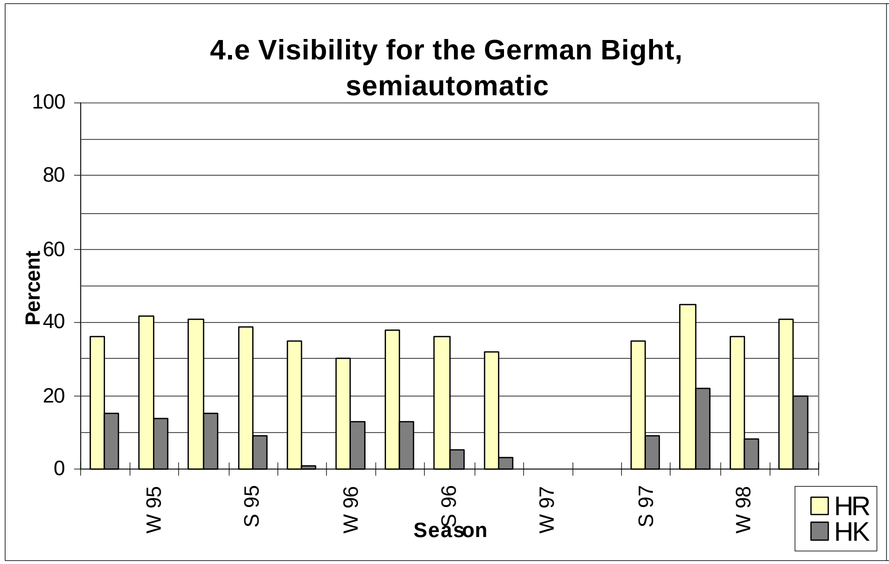

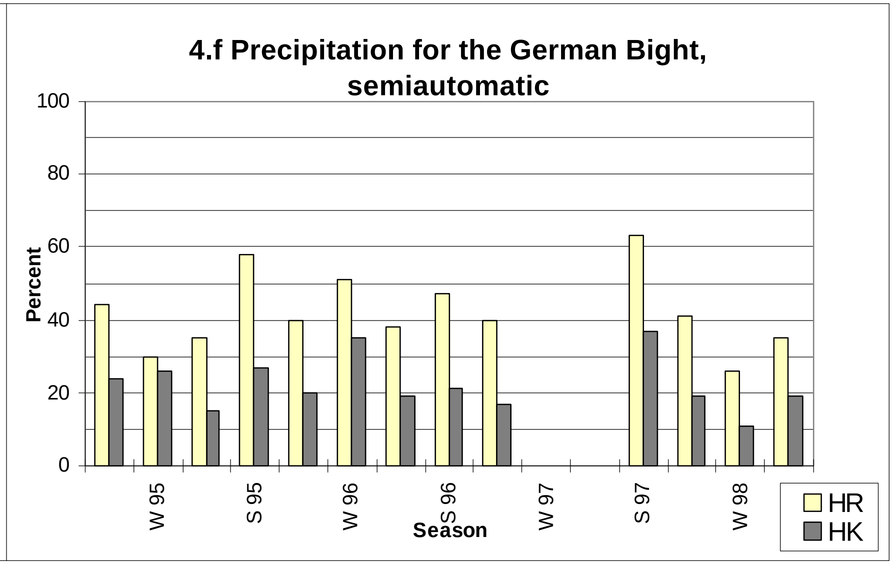

4. TAF

4.1 The product

TAFs for three airports are verified for these reports: Copenhagen/ Kastrup, Billund and Kangerlussuaq/Søndre Strømfjord. The TAFs for Copenhagen and Kangerlussuaq are produced at the local DMI offices, those for Billund at the central weather service in Copenhagen.

4.2 Verification

4.2.1 General procedure

The NORTAF scheme ([5]) is used for the verification. This scheme was developed jointly by the Nordic meteorological institutes and is an extension of the scheme devised by Gordon ([1], [2]). Here we will present some central results for the most important weather parameters. Some additional results will be mentioned in the text.

In the NORTAF scheme, as in Gordon's original work, all verification is performed for three-hour segments, so-called time blocks, of the valid period of the TAFs. We will show results for time blocks 1 and 3.

The TAFs for the two Danish airports are valid from 06 to 15 UTC, those for Kangerlussuaq from 06 UTC to 06 UTC the next day. Time blocks 1 and 3 are thus the periods 06-09 UTC and 12-15 UTC. Persistence forecasts are based on observations from 03 to 06 UTC.

Included in the scheme are also some statistics on syntax errors in TAFs and observational reports. We will present some of the results.

4.2.2 Visibility

Visibility is verified in Gordon's manner: The minimum value forecast for a given time block is compared to the minimum observed value by means of a contingency table. The minimum forecast value is defined as the smallest of all values forecast for the block, either as basic values or in change groups. The minimum observed value is the smallest value reported in METAR or SPECI.

We use a contingency table with five categories defined on the basis of the amendment criteria: 0-350 m, 350-800 m, 800-1500 m, 1500-5000 m and 5000 m or more (low-end values are included in the intervals). From the table we compute hit rate and HK skill index.

Figures 5 a-b show HR and HK for Copenhagen, blocks 1 and 3, for 8 seasons. HR is quite high, 70-80 percent, for both blocks and all seasons, except the winter seasons where values are between 50 and 70 percent. The HK is high, 40-80 percent, for block 1 and lower, 25-40 percent, for block 3. For summer 1997 the block 3 forecasts have negative skill. The reason for this is evident from the contingency table (not shown here): In all cases but one the observed minimum visibility belonged to the top category; a reference forecast based on the climatology of the data would be near perfect, thus beating the TAF forecasts which were somewhat biased towards lower visibility values.

The corresponding results for Billund, shown in figures 6 a-b, show many of the same features, but both HR and HK are generally lower than for Copenhagen.

The contingency tables for the Danish airports reveal for all seasons and time blocks a bias in the forecasts towards too low visibility values. Notably, the category 1500-5000 m was forecast too often when the reported minimum visibility was 5000 m or more. Persistence forecasts tend to outperform the TAFs at least for block 1 and often beyond.

The visibility forecasts for Kangerlussuaq (figures 7 a-b) have very high hit rates, over 80 percent, for all seasons and both time blocks. As also confirmed by the contingency tables, this is simply because visibility at Kangerlussuaq is almost always very good. That is also the reason why the HK values are sometimes high, sometimes low (even negative) and for certain seasons and time blocks undefined: The index is very sensitive to individual sets of observations and forecasts if nearly all observations belong to the same category, and becomes singular if they all do.

Even for Kangerlussuaq there is for some seasons a bias towards lower visibility values. For most seasons the persistence forecasts are better than the TAF forecasts in terms of HR throughout the 24 hours' forecast period.

The bias problems apparent in the contingency tables for visibility are confirmed by results computed with other NORTAF verification (reliability tables, contingency tables for the occurrence of alternative values).

4.2.3 Ceiling

Ceiling is defined as the base height of the lowest cloud layer covering more than 4 oktas. Ceiling forecasts are verified using Gordon's method. We use five categories, defined by the following low-end limits (included in the intervals): 0, 200, 500, 1000, and 1500 ft.

Figures 5-6 c-d show HR and HK for blocks 1-3 for Copenhagen and Billund. The results as measured by HR are quite good in the summer months, but poor, below 60 percent, in winter. HK skill is quite high for block 1, but lower for block 3; here, it tends to vary inversely to HR over the seasons, corresponding to the fact that it is difficult for the meteorologist to beat climatology when the weather is most often fine.

The contingency tables for Copenhagen show a clear bias towards forecasting too low ceiling values. This gives persistence a good case; for nearly all seasons persistence is better than the TAF forecasts for blocks 1 and 2. Billund shows less bias, but still as a rule persistence is better for block 1.

For Kangerlussuaq (figures 7 c-d) the HR results are very good. HK shows a more confused picture: For several seasons the value is negative or undefined. As confirmed by the contingency tables, this pattern appears because ceiling values are in the top category most of the time at Kangerlussuaq. For most seasons the tables show some bias towards forecasting too low ceiling heights, and persistence hits as often as the TAFs, or more often, out to block 4 or even further.

As with visibility, the pessimism in the forecasts shows itself also in results computed with other NORTAF verification methods.

4.2.4 Maximum wind speed

In the NORTAF scheme, Gordon's method is also applied to maximum wind or gust speed over a time block. Intervals of 10 kt are used in the contingency tables.

Figures 5 e-f show HR and HK for the two Danish airports, blocks 1 and 3. HR values are not outstanding: 60-75 percent for block 1 and a little lower for block 3, especially for Copenhagen. Skill values are quite high for block 1, above 40 percent, but lower for block 3, especially for Copenhagen and the spring and summer seasons.

The contingency tables for Copenhagen show little bias for most seasons and time blocks; persistence forecasts are better for some seasons and time blocks, for others not. In Billund, autumn and winter forecasts tend to be biased towards too strong wind; with the spring and summer forecasts the opposite is most often the case. Persistence is rarely better than the TAF forecasts.

Kangerlussuaq (figures 7 e-f) shows HR values at about the Copenhagen level, but the skill scores are low, especially for block 3 and the warmer seasons. This goes well with the fact that strong winds are rare at Kangerlussuaq compared to Copenhagen. For all seasons the forecasts for at least the first three time blocks are biased towards too strong winds, and persistence outperforms the TAFs out to block 3 to 4 for most seasons.

4.2.5 Precipitation

In the NORTAF scheme precipitation is verified by means of 4-by-4 contingency tables in the same way as it is done with the marine forecasts (cf. 3.2.5), except that the verification period is one time block.

According to the regulations, light precipitation must not be forecast in a TAF, so forecasts of light precipitation are considered as forecasts of dry weather. In order to avoid injustice to the forecaster, if light precipitation is reported it is counted as precipitation if precipitation was forecast, but discarded otherwise.

Figures 5-7 g-h show HR and HK for the three airports for blocks 1 and 3. The results are good, with HR values generally above 80 percent and HK values most often exceeding 60 percent. Inspection of the contingency tables reveals a general tendency towards over-forecasting of precipitation in cases of dry weather.

4.2.6 Syntax

Tables 1 a-c show the percentage of syntactically incorrect TAFs, METARs and SPECIs from each airport and season.

For all three code types, the error percentages for Copenhagen and Billund are quite low, a few percent, except for a few seasons where they exceed 5 percent. Winter 1996 was not a good season; this is probably because of the changes in the aeronautical codes introduced on January 1, 1996.

At Kangerlussuaq the error percentages are higher for all code types. The differences may be explained by the fact that the semi-automatic observation systems and TAF syntax checking software used in Denmark are not installed at Kangerlussuaq. The very high error percentages for SPECI are due to a software problem which causes the SPECI to get an incorrect time stamp.

4.3 Comments on the results

For Copenhagen and Billund the important TAF parameters of visibility and ceiling are forecast with varying success, depending on the time of year and the forecast length. For Kangerlussuaq the forecasts have very high hit rates due to the favourable climatic conditions there. The forecasts of maximum wind are of more even, although not very high, quality. Common precipitation types are well forecast for all stations.

For the parameters ceiling, visibility and precipitation there is a marked tendency towards over-forecasting of poor weather conditions, especially for the two Danish airports. This bias is presumably the primary reason why forecasts based on persistence tend to fare better than the TAFs for visibility and ceiling for the first few time blocks.

Figures 5 a-b: Verification of visibility for Copenhagen airport. Hit rate and HK index (light and dark columns, respectively).

Figures 5 c-d: Same as Figures 5 a-b, but for ceiling.

Figures 5 e-f: Same as Figures 5 a-b, but for maximum wind speed.

Figures 5 g-h: Same as Figures 5 a-b, but for precipitation.

Figures 6 a-h: Same as Figures 5 a-h, but for Billund airport.

Figures 7 a-h: Same as Figures 5 a-h, but for Kangerlussuaq airport.

Tables 1 a-c: Syntax errors in aeronautical messages for Copenhagen, Billund and Kangerlussuaq airports.

Copenhagen airport

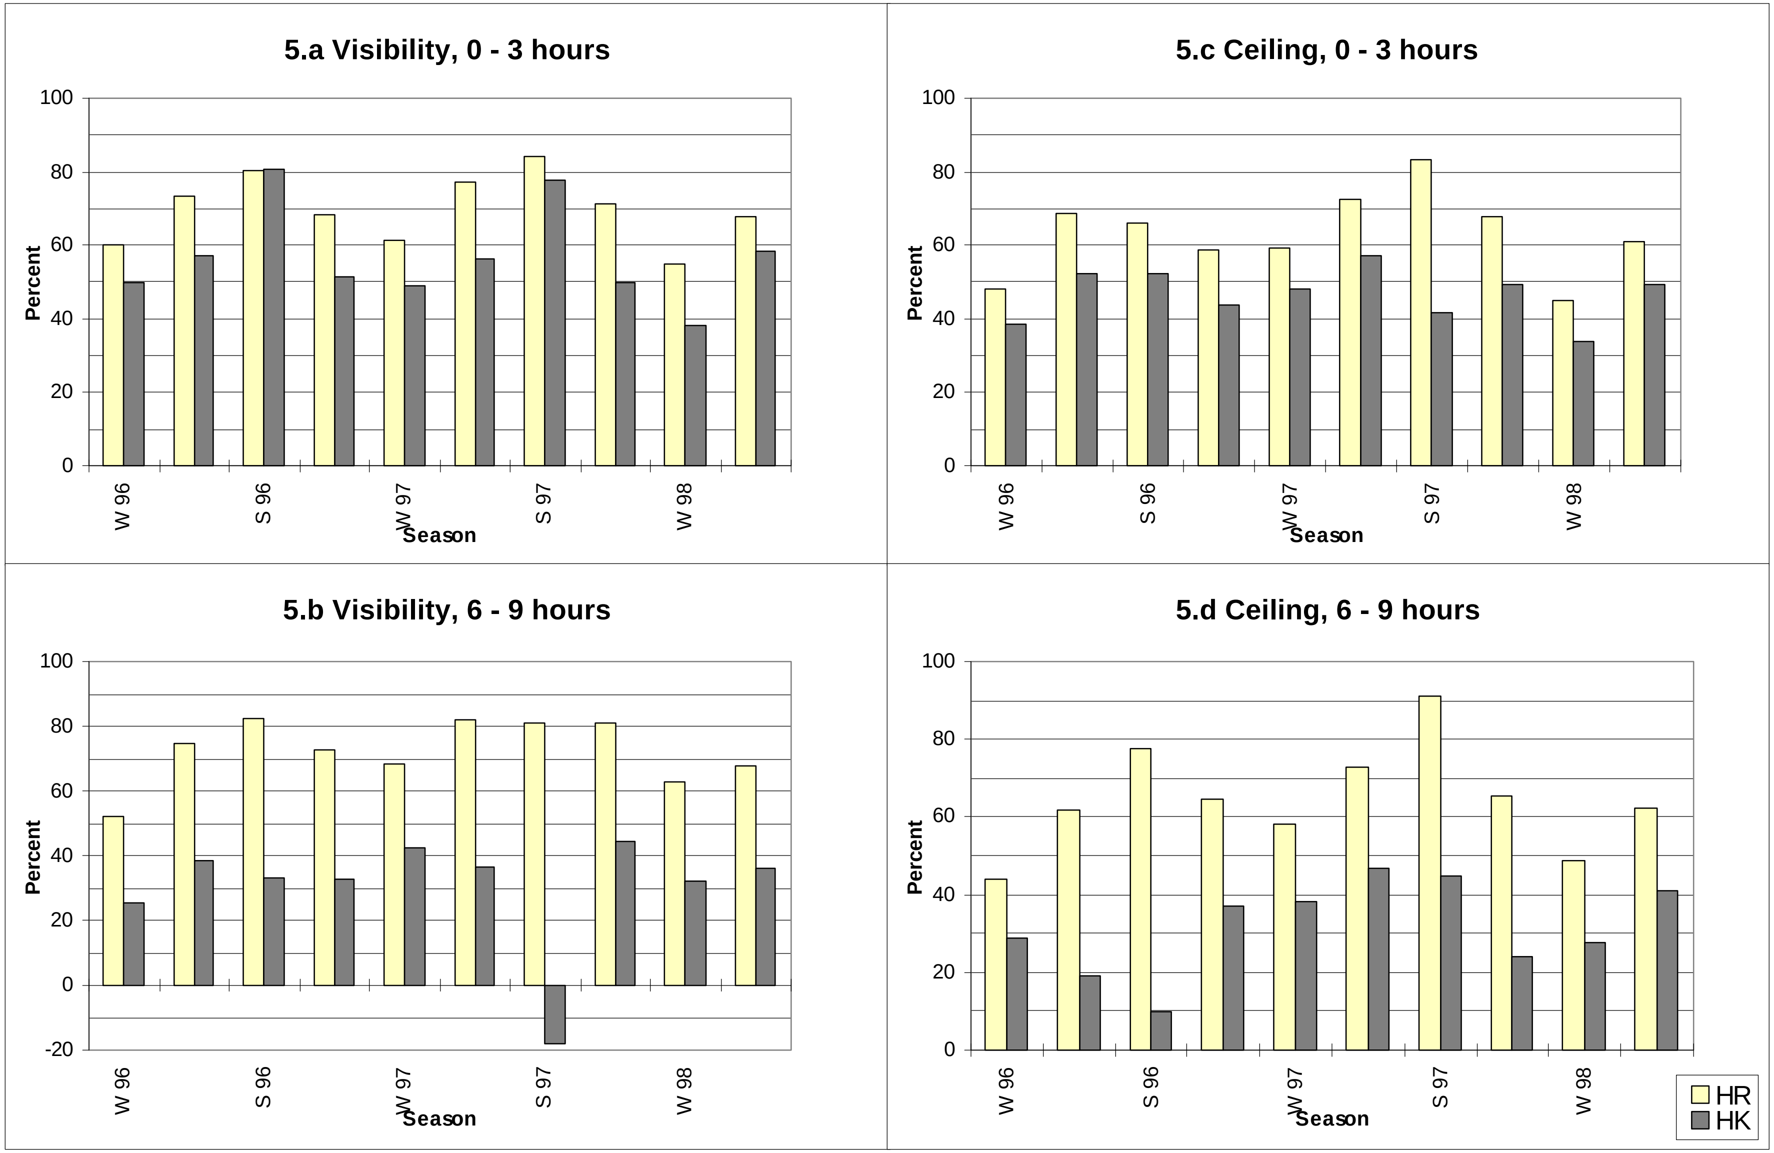

Copenhagen airport

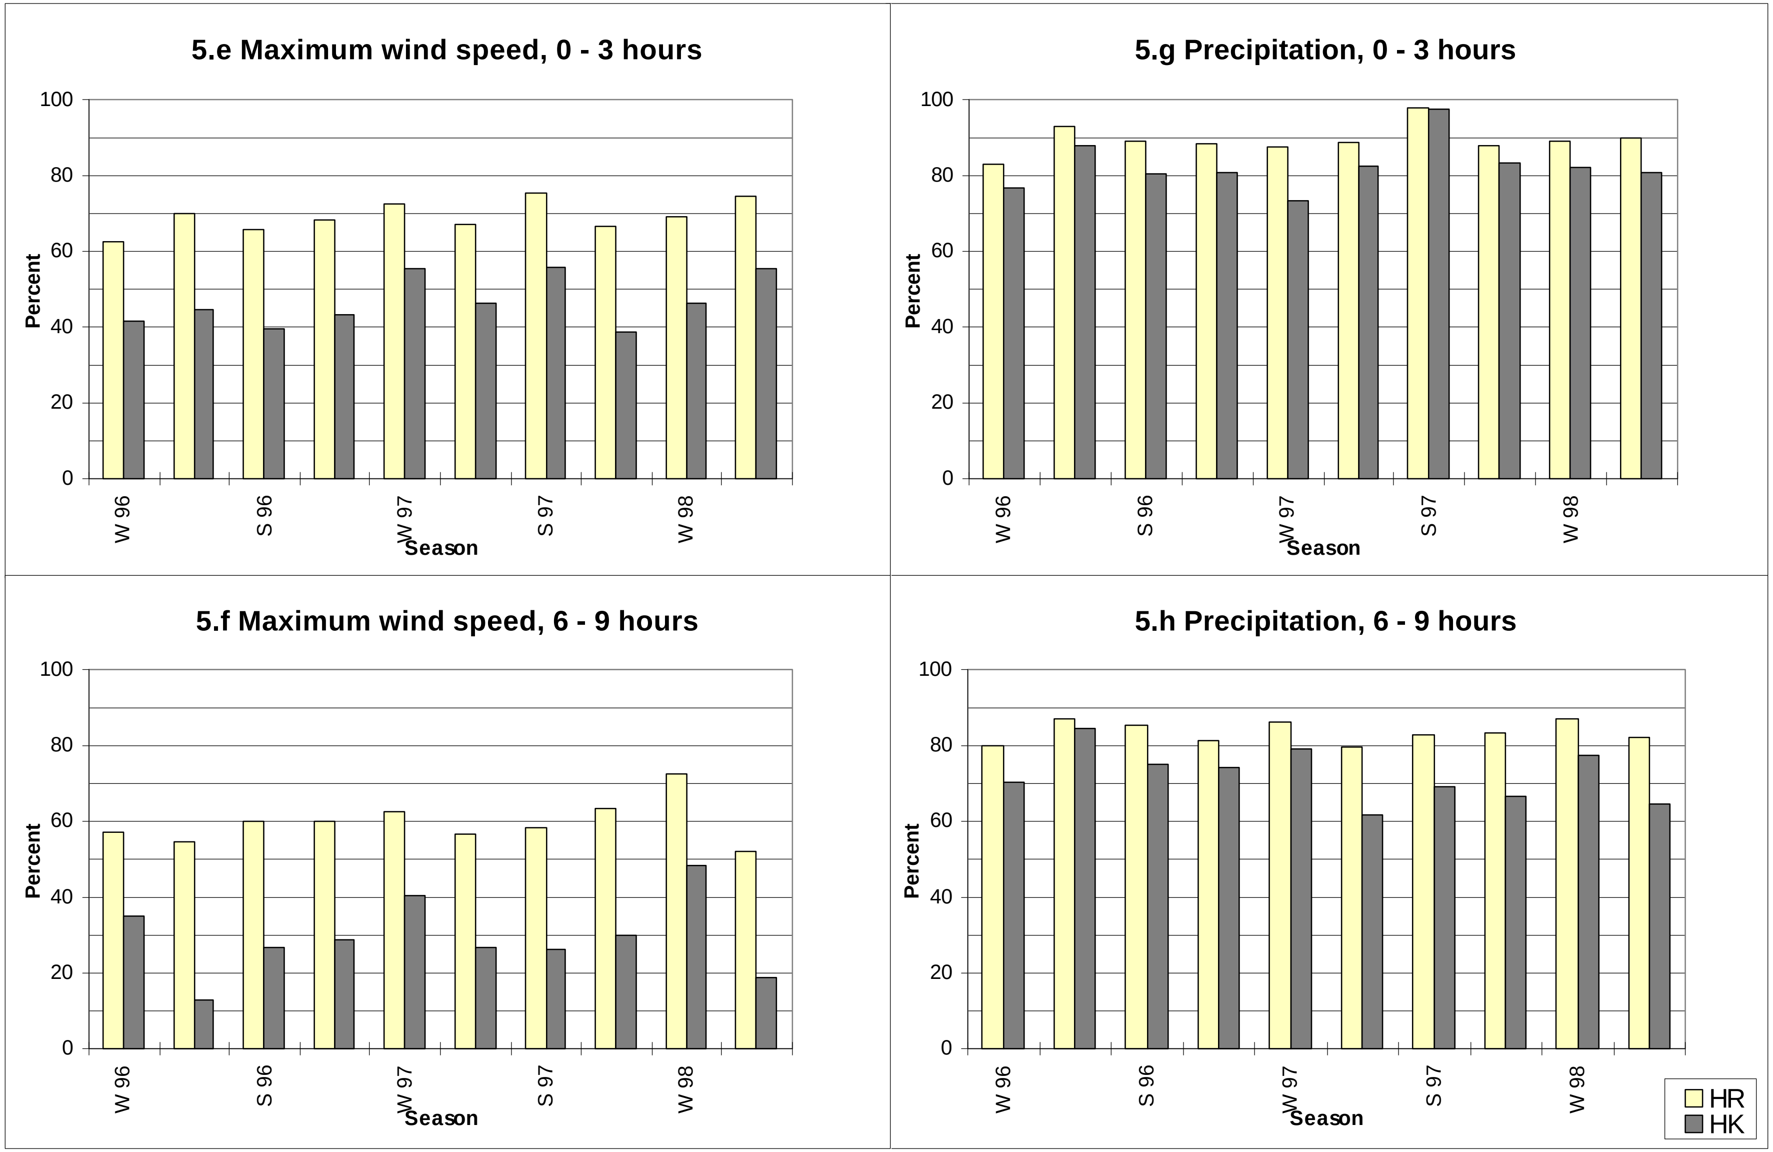

Billund airport

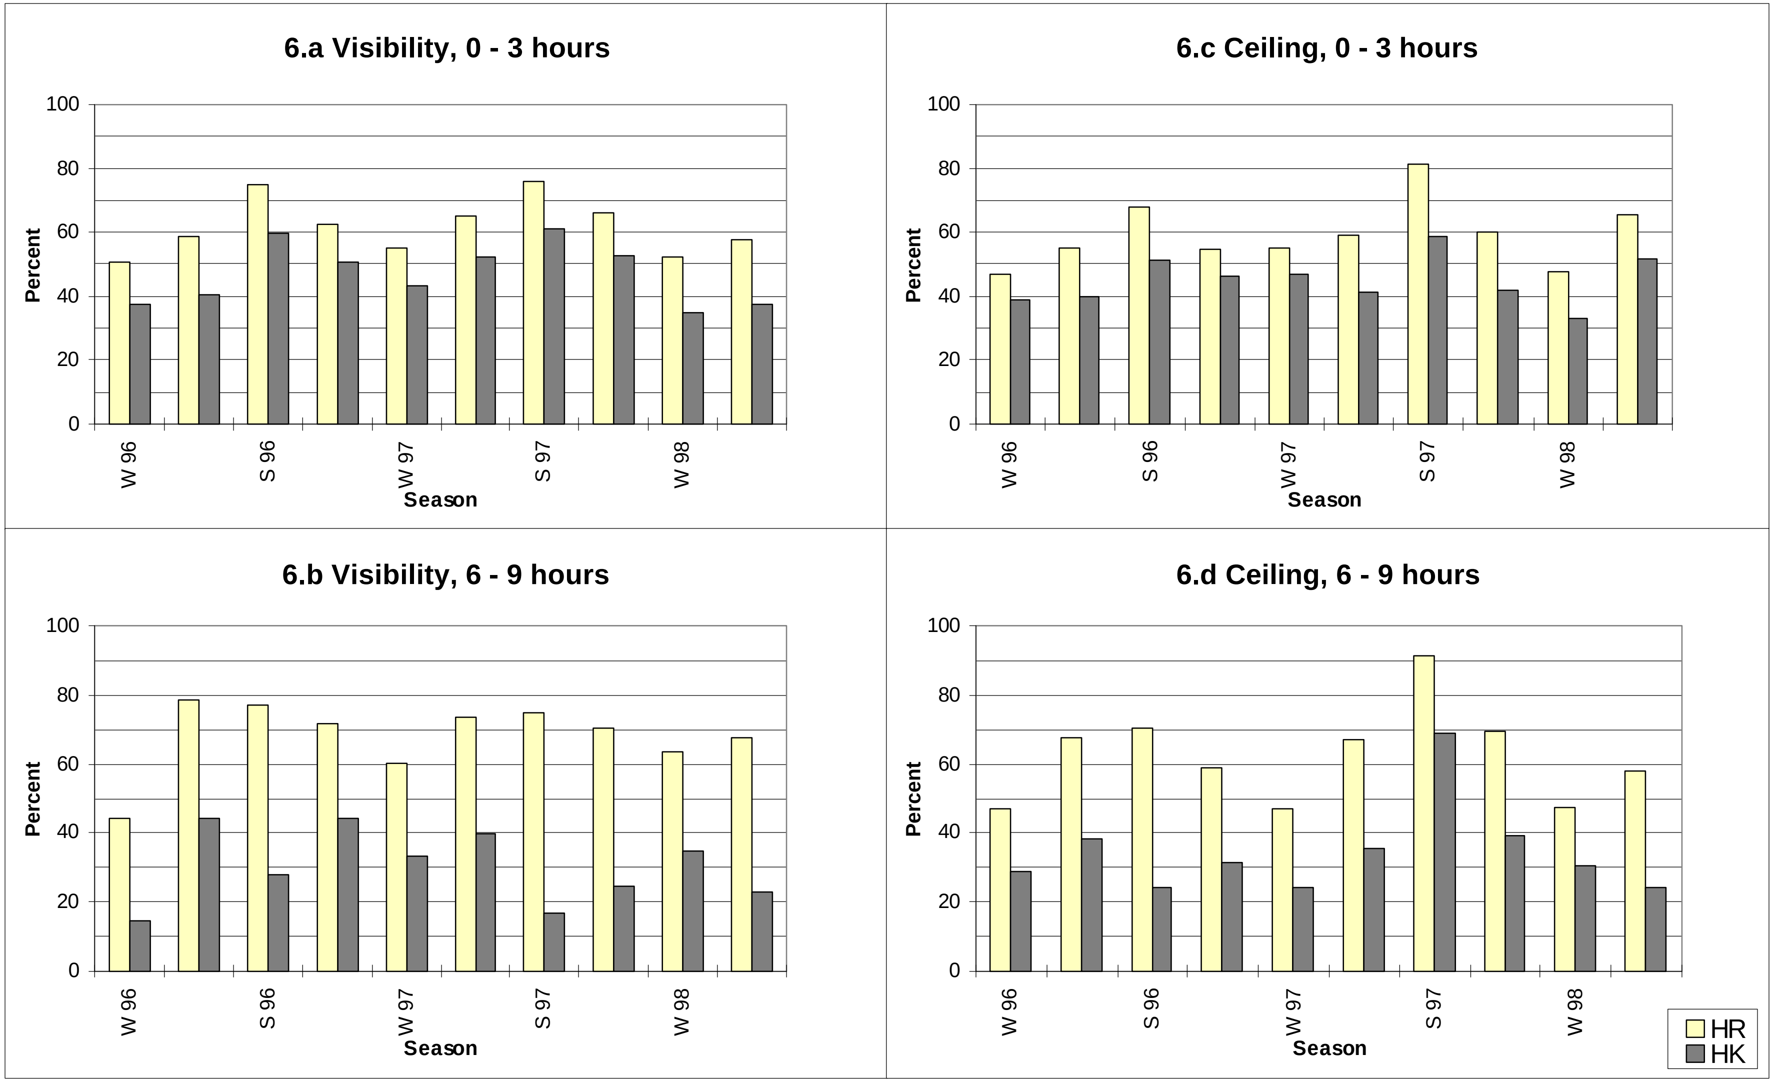

Billund airport

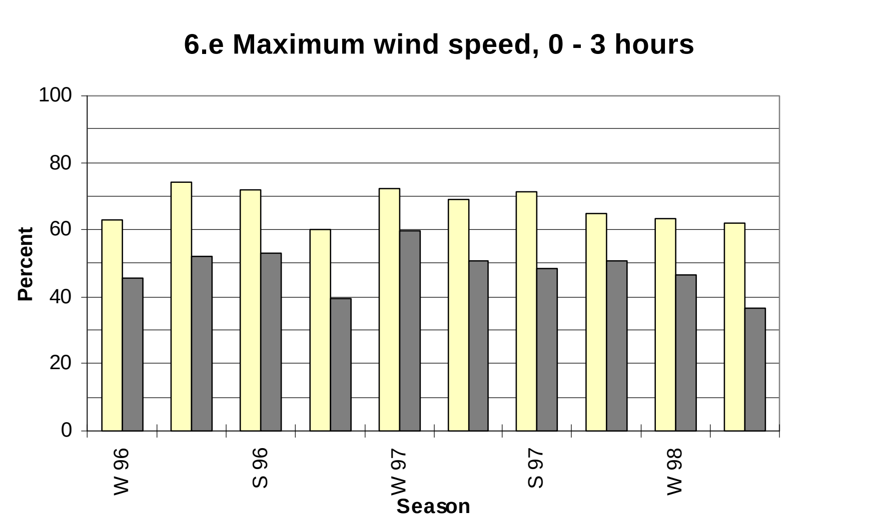

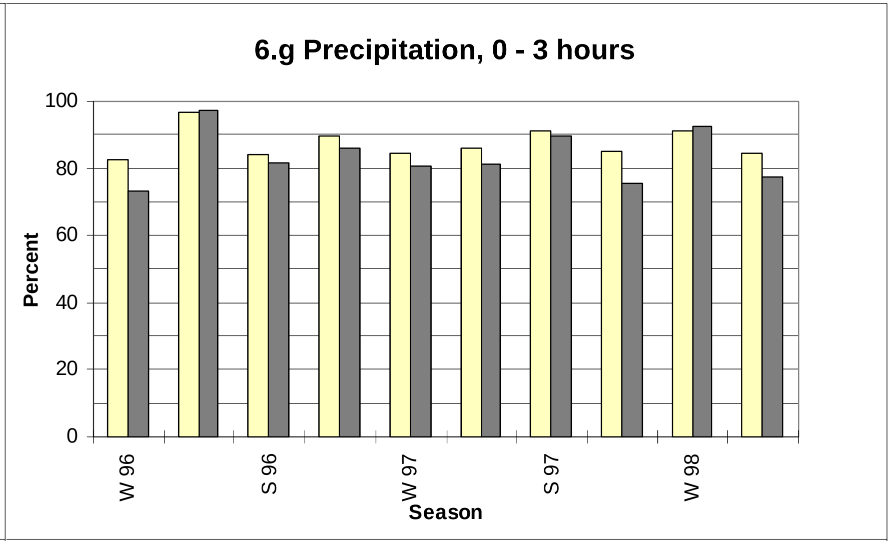

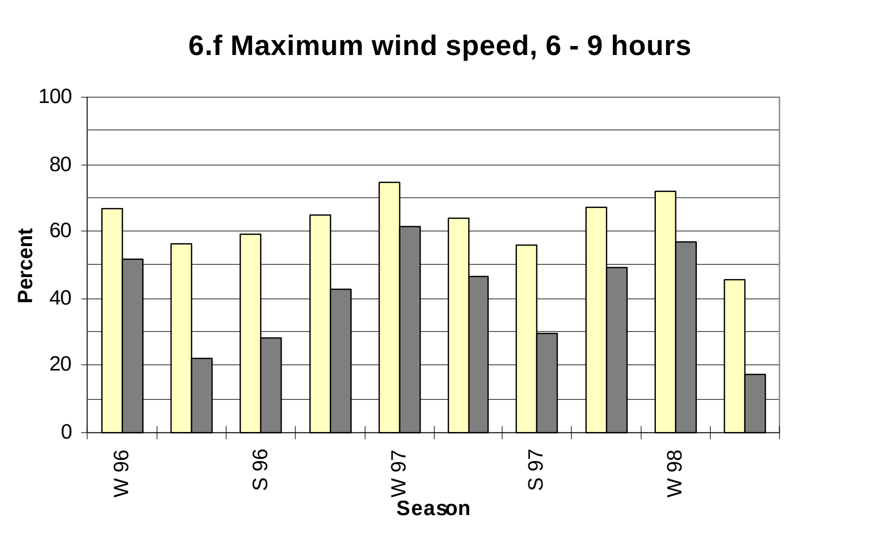

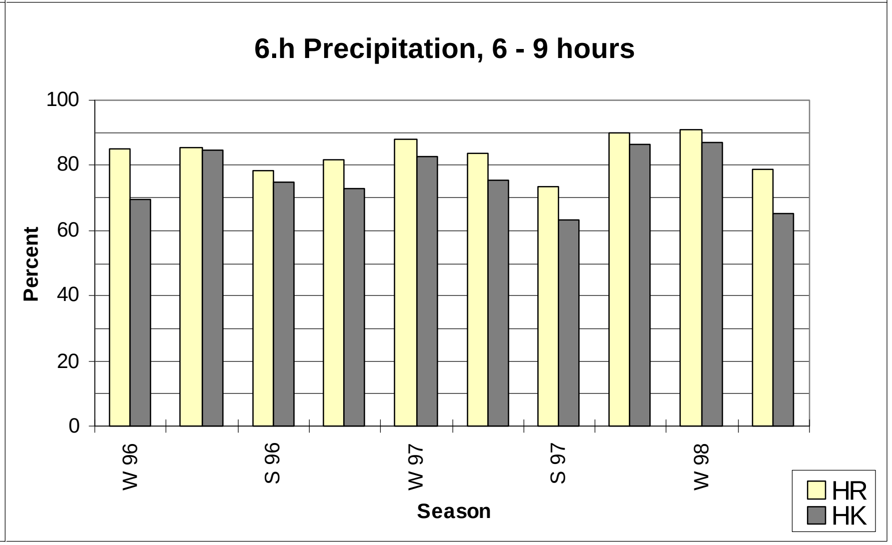

Kangerlussuaq airport

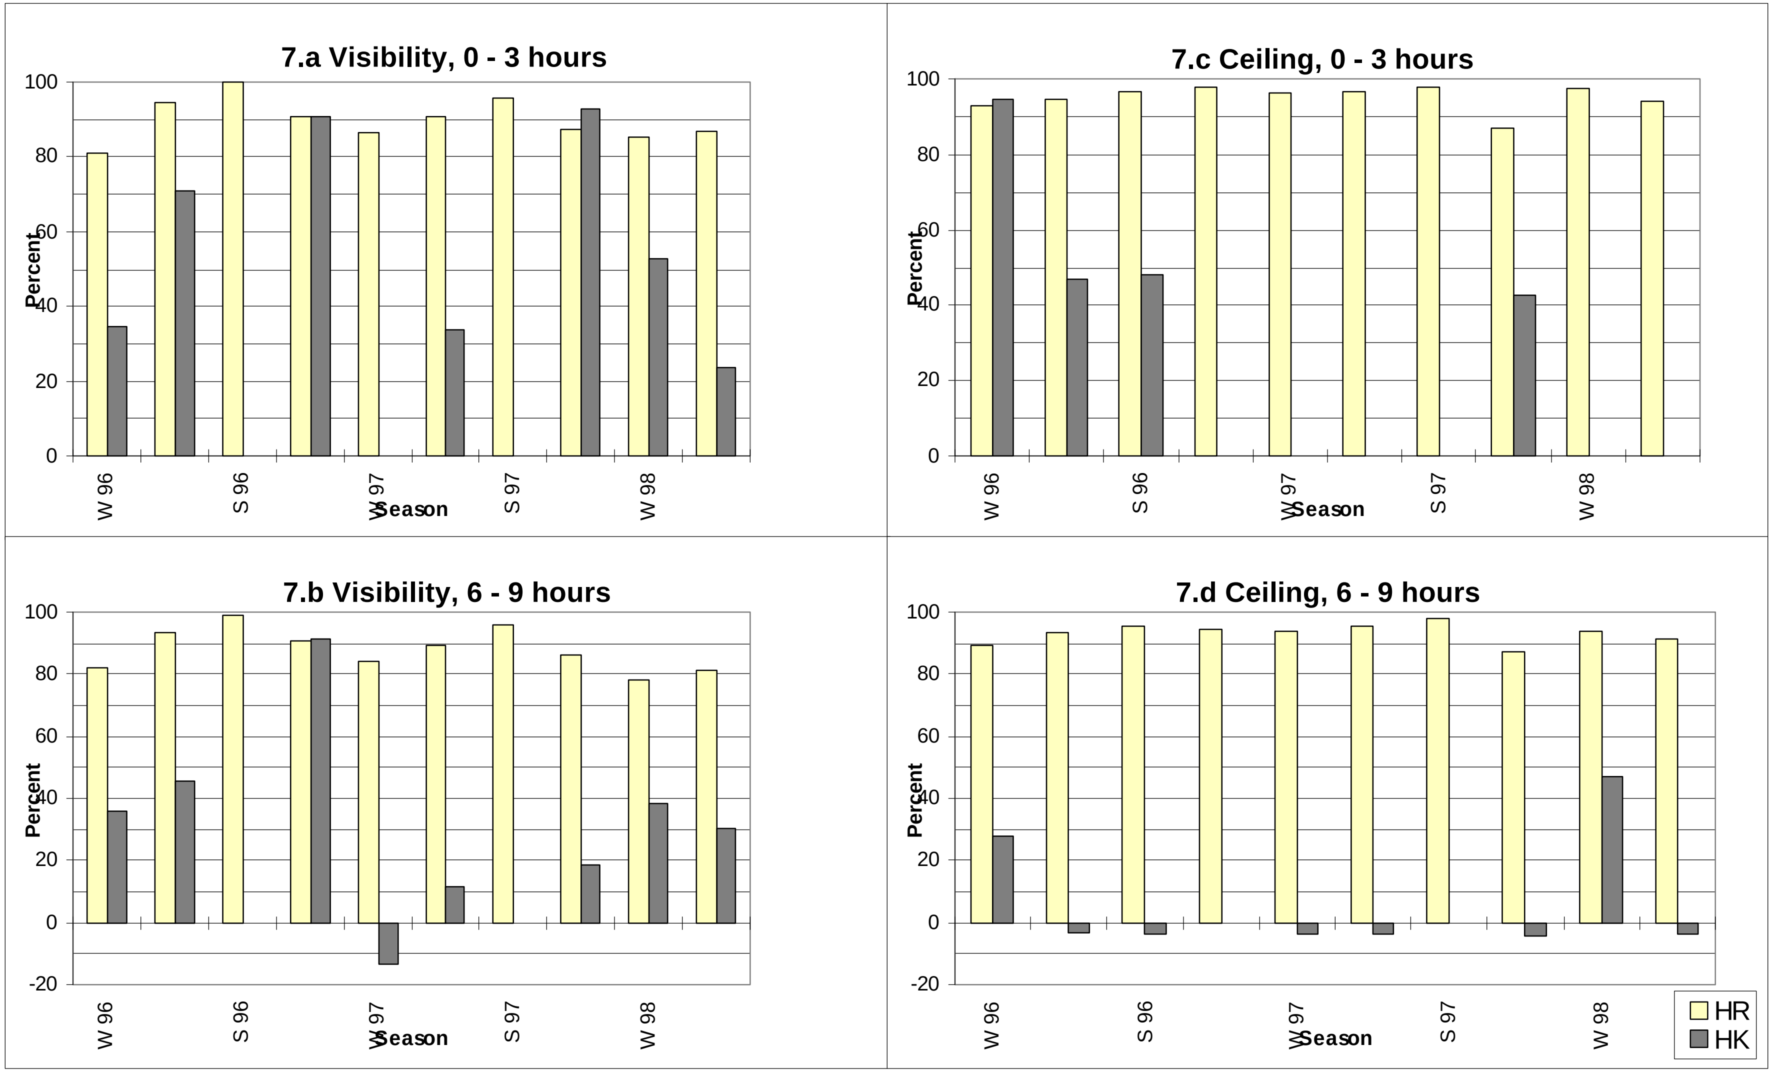

Kangerlussuaq airport

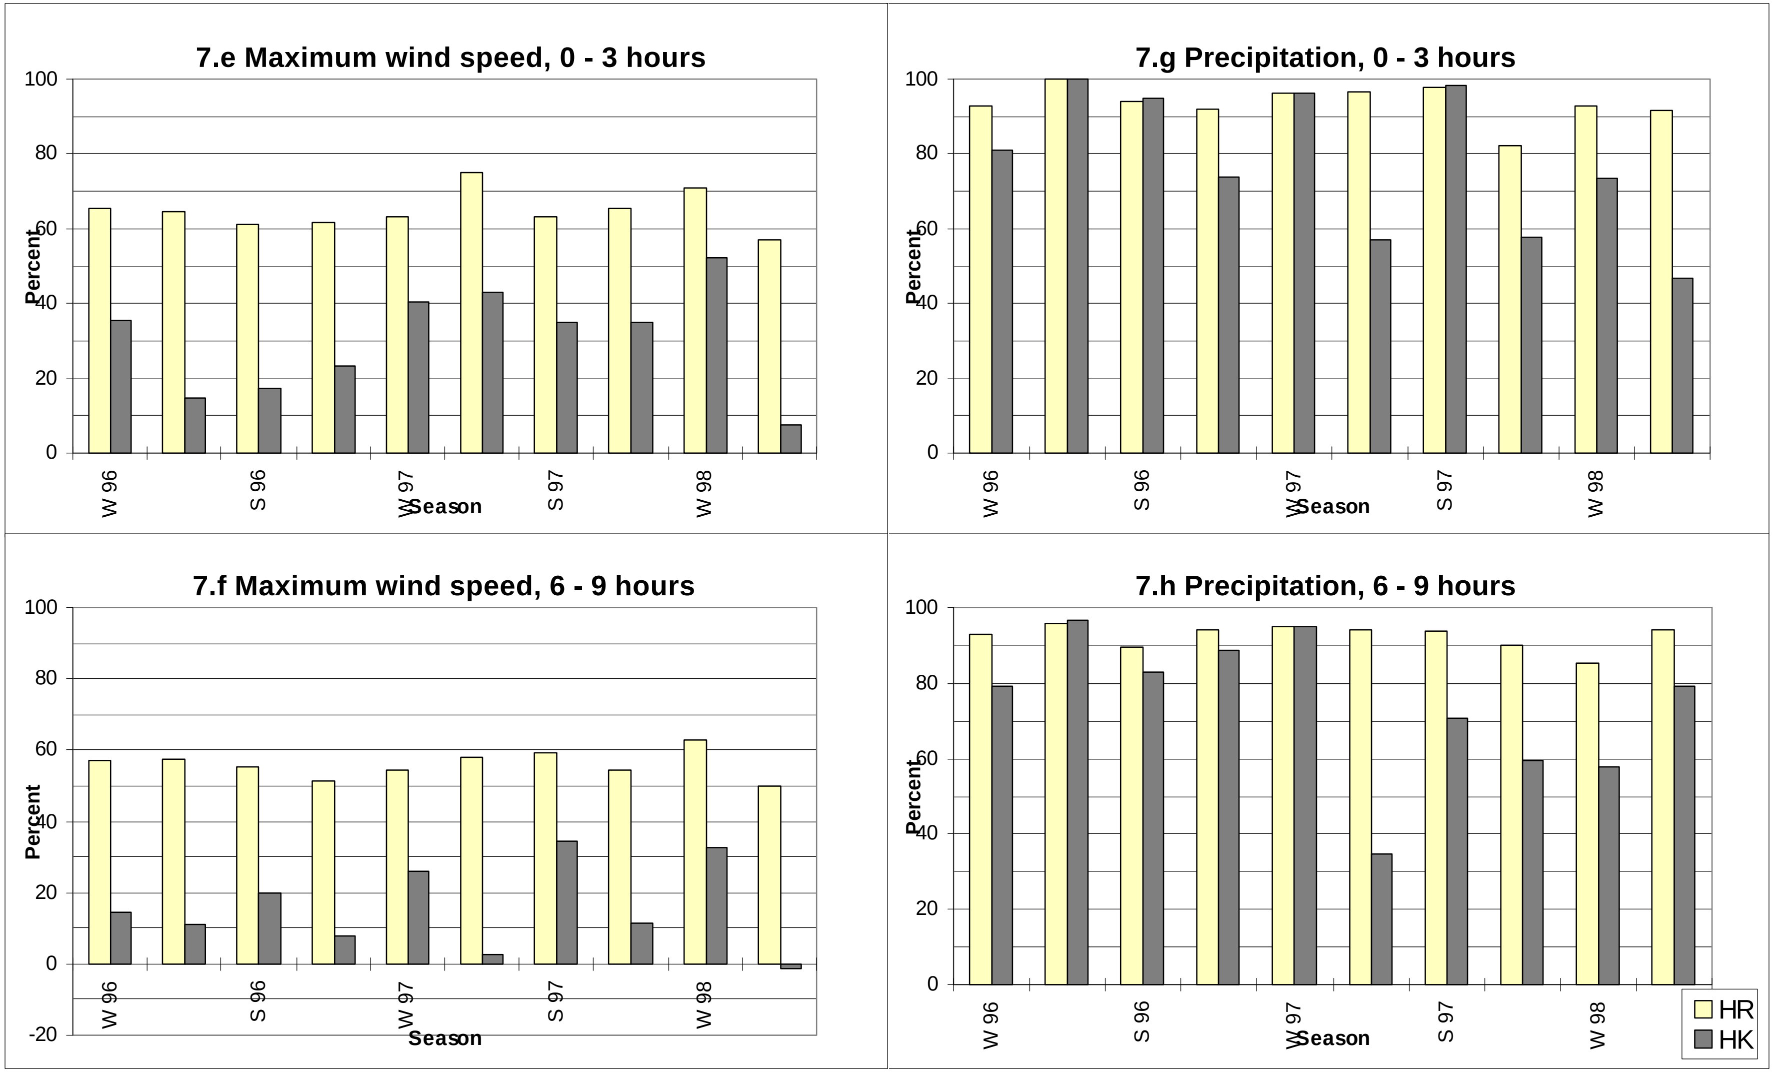

Aeronautical code error percent

Table 1.a Copenhagen			
Season	TAF	METAR	SPECI
Winter 1996	7.7%	2.7%	7.1%
Spring 1996	3.4%	1.4%	1.5%
Summer 1996	1.1%	1.5%	5.6%
Fall 1996	5.6%	1.0%	1.5%
Winter 1997	2.2%	1.9%	2.4%
Spring 1997	4.3%	1.0%	7.4%
Summer 1997	3.3%	0.7%	5.8%
Fall 1997	1.1%	0.7%	2.4%
Winter 1998	6.70%	1.20%	5.40%
Spring 1998	2.20%	0.80%	6.50%

Table 1.b Billund			
Season	TAF	METAR	SPECI
Winter 1996	10.0%	2.3%	4.3%
Spring 1996	2.2%	0.6%	1.6%
Summer 1996	3.3%	0.6%	0.0%
Fall 1996	3.3%	0.9%	3.3%
Winter 1997	6.7%	0.8%	3.9%
Spring 1997	1.1%	0.4%	0.0%
Summer 1997	1.1%	0.9%	0.0%
Fall 1997	3.3%	0.9%	2.5%
Winter 1998	2.20%	0.70%	2.80%
Spring 1998	1.10%	0.80%	2.20%

Table 1.c Kangerlussuaq			
Season	TAF	METAR	SPECI
Winter 1996	6.7%	3.5%	43.2%
Spring 1996	16.5%	2.3%	51.9%
Summer 1996	6.6%	2.5%	53.8%
Fall 1996	5.5%	3.9%	12.5%
Winter 1997	5.7%	4.7%	6.0%
Spring 1997	6.6%	3.9%	10.2%
Summer 1997	44.9%	2.9%	18.8%
Fall 1997	10.2%	2.8%	9.9%
Winter 1998	6.70%	1.40%	11.10%
Spring 1998	23.10%	2.30%	12.00%

5. References

[1] Gordon, N.D.: *Verification of terminal forecasts*. Proceedings Fifth International Conference on the Aviation Weather System, 2-6 August 1993, Vienna, Virginia.

[2] Gordon, N.D.: *Why and how to verify TAFs*. CaEM-X / Doc. 45 part G, October 1994, Geneva, Switzerland.

[3] Hanssen, A.W., and W.J.A. Kuipers: *On the relationship between the frequency of rain and various meteorological parameters*. De Bilt, KNMI, Mededelingen en Verhandelingen, 81, 1965, 65 pp.

[4] Hilden, A.: *Semi-automatic generation of sea area forecasts*. Preprints, Det 19. nordiske meteorologmøte, June 1994, Kristiansand, Norge.

[5] Hilden, A., A. Heidegård, K.H. Midtbø, J. Santakari and T. Vävargård: *The NORTAF verification system. Status report*. May 1996, DMI, Copenhagen, Denmark.