Optimization of Rainfall Interpolation

I. Soenario, R. Sluiter

De Bilt, March 2010 | Intern rapport; IR 2010-01

Koninklijk Nederlands
Meteorologisch Instituut
Ministerie van Verkeer en Waterstaat

Optimization of Rainfall Interpolation

Versie 2.0

Datum	Maart 2010
Status	Definitief

Colofon

Titel	Optimization of Rainfall Interpolation
Versienummer	2.0
Contactpersoon	Raymond Sluiter T 030 2206 446 raymond.sluiter@knmi.nl Koninklijk Nederlands Meteorologisch Instituut (KNMI) Wilhelminalaan 10 | 3732 GK De Bilt Postbus 201 | 3730 AE | De Bilt
Bijlage(n)	Geen
Auteurs	Ivan Soenario Maarten Plieger Raymond Sluiter (editor)

Contents

Voorwoord

In opdracht van de Waterdienst is een onderzoek op het KNMI uitgevoerd naar methoden voor optimale vergridding (interpolatie) van relevante parameters voor het stroomgebied van de Rijn in Nederland (deelovereenkomst 31007187.0002 onderdeel 2).

In dit rapport zijn de resultaten voor parameter neerslag weergegeven.

N.B. de onderzoeksresultaten zijn in het Engels opgesteld.

1 Introduction

There is increasing demand for gridded products of meteorological and climatological variables with high quality and spatial resolution from many different disciplines. The KNMI initiated several research projects regarding interpolation with the following general objectives (Sluiter, 2008):

- Acquire knowledge through literature review.
- Give an overview of current products and possible interpolation and data providing applications.
- Determine which interpolation method is best for interpolation of meteorological and climatological variables (a.o. rain, temperature, relative humidity).

On behalf of the last objective, research on temperature interpolation has already been executed (Salet, 2009). Because rainfall is very irregular and fickly, interpolation is hard and requires special attention. For this reason rainfall interpolation has been chosen as subject of this research.

Current interpolation of meteorological and climatological data at the KNMI is done with a spline technique, which results in annotated images (figure 2.1) available through the internet (Sluiter, 2008; Salet, 2009). Four comments can be made:

- Spline methodology has some demerits (see below).
- Due to manual intervention reproducibility is difficult.
- Only images and not the underlying data are available through the internet.
- Images are not GIS maps and thus not usable in GIS environments.

The spline technique has the following demerits: first, it can't deal with external factors that could influence meteorological data, such as distance to sea, height above sea level and prevailing wind direction. Second, the weight with which measurement points influence a prediction at an unvisited location depends only on the distance. It is a mathematical interpolation method, not a geostatistical method. Furthermore, the results are presented without error indication. Cross validation is possible, but has not been executed for the present products.

1.1 Objectives

The objectives are formulated as follows:

- Find an optimal interpolation method for rainfall, which is automated and reproducible.
- Provide the results as "real" data that can be used in further analysis.
- Provide information about the prediction errors.

The research was limited to the interpolation of daily rainfall. Additionally, monthly totals were created for external purposes by summing daily rainfall maps. To compare the outcomes of different techniques, two specific days representing two specific rainfall patterns have been chosen following Schuurmans et al. (2007). April 1, 2004 is an example of stratiform rainfall. This type of rainfall is characterised by an evenly spread spatial distribution, relatively low but continuous rain intensity which results in a fair amount of rainfall in 24 hours. May 1, 2004 presents convective rainfall, which is usually short-lived, intense and very local.

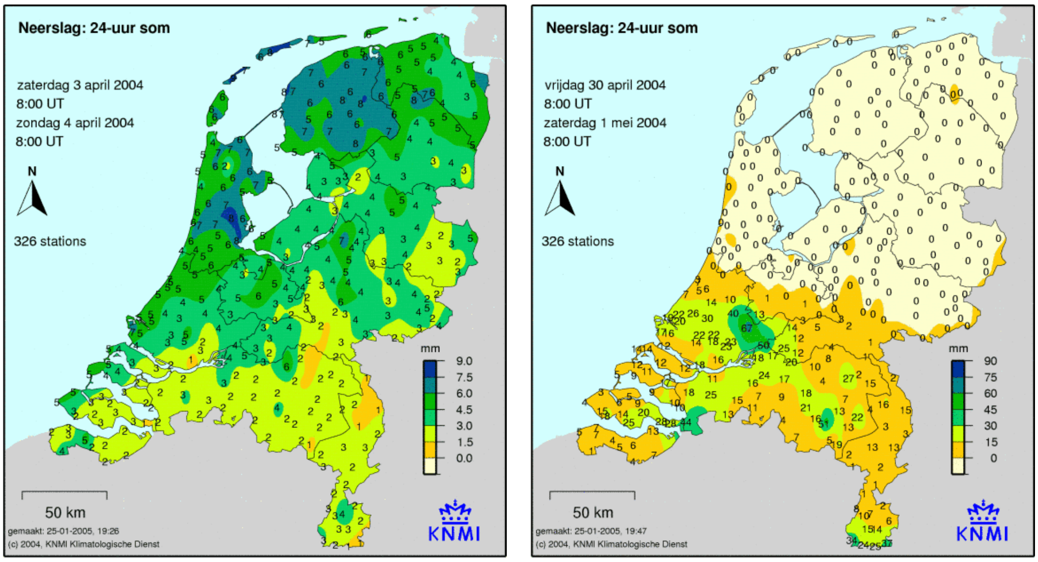

Figure 1.1 – Cumulative 24h precipitation maps available at the KNMI website, for April 4 2004 and May 1 2004.

2 Data and interpolation environment

2.1 Input data: voluntary precipitation network

Rainfall data is collected by a voluntary network, which consists of more than 300 stations. See appendix 1 for a map with the observation locations. Each 24 hour cumulative rainfall is measured at 08:00 UTC. This implies that a measurement dated May 1st includes the interval from April 30th 08:00 to May 1st 08:00. The data is stored in the Klimaat Informatie Systeem database (KIS). A recipe extracts data from KIS with a query (see section 2.2).

For each observation point the following information is obtained:
- Unique station number.
- Coordinates (rijksdriehoekstelsel).
- Actual date of measurement.
- Amount of rainfall (kg/m2).

The following restrictions apply:
- Station start time before 1988.
- Station stop time after 2007.
- It must contain daily cumulative rainfall.
- Rainfall is not null (the interpolation software needs real values).

With these restrictions, the number of observations used for an interpolation ranges from 290 stations in 1991 to 309 stations in 2006. See appendix 1 for a map with measurement locations and a graph showing the development of the number of available measurement locations.

2.2 Interpolation environment

2.2.1 Rheinblick/HYRAS interface

Most of the research was done using the Rheinblick/HYRAS interface, developed at KNMI. This interface shows maps via web mapping services and has a “recipe manager” to develop and execute recipes, queries and R-scripts for interpolation. The recipes are the main control. They define parameters such as the start and stop time of the period over which maps are calculated, output location, the legend to use and which query and R-script should be used.

2.2.2 Implementation of R scripts

The actual interpolation is done with R, software for statistical computing, which uses packages for geospatial analysis (R-Project, 2009). R is embedded in a web interface as described in section 2.2.1.

The packages used in R are:
- `sp` which allows R to deal with spatial objects.
- `gstat`, containing the geostatistical tools.
- `automap`, which automates the interpolation process, by automatically estimating a semi-variogram and performing Kriging.

The interpolation process is highly automated by using the automap package. The automap package calls functionality from the gstat module. More specifically, `autoKrige` does the interpolation and uses `autofitVariogram`, which uses `fit.variogram`, with certain assumptions made for the variogram model (Hiemstra, 2009). This will be further explained in the methods section.

2.3 Reference data

2.3.1 HYRAS-REGNIE

Within the Rheinblick/HYRAS project, Deutsche Wetterdienst (DWD) has interpolated precipitation for the entire Rhine catchment using the REGNIE (Regionalisierung der Niederschlachshöhen) method. The source data is the same as in this research. REGNIE uses background fields to calculate quotient anomalies. Through these background fields elevation is taken into account. First, the value of a weather station is assigned to a grid cell, which is then divided by the background field value at that grid cell. These quotients are then interpolated using IDW. This result is then multiplied with the background field to obtain precipitation values (Steiner, 2009). The HYRAS-REGNIE product is one of the products that have been used for comparison. As the southern part of the Netherlands drains to the Maas it is not included in the dataset as shown in figure 4.5, bottom images.

2.3.2 Precipitation radar

The KNMI operates two C-band Doppler radars, located in De Bilt and Den Helder. The resolution of the images available in this study is 2.5km x 2.5km (1 x 1 km since 2008). For both radars, the surface rainfall intensities are accumulated for each pixel from 0800 to 0800 UTC on the following day to match the observations of the voluntary network. Because both underestimation and overestimation take place, observations from the voluntary network are used to perform a daily range correction (Holleman, 2004)

3 Methods

In preceding studies (Buishand et al., 2008; Sluiter, 2008) IDW and several types of Kriging were indentified as most promising techniques. Moreover it was recommended to study the work of Haylock et al. (2008) and Schuurmans et al. (2007) in more detail. Therefore, in this study a mathematical method (IDW) and two geostatistical methods (Kriging) are used to interpolate the rainfall maps. Prior to interpolating rainfall, the data distribution is investigated following Schuurmans et al. (2007).

3.1 Data distribution

Data distribution has been researched with a normal quantile-quantile plot (q-q plot). This plot shows if data is close to being normally distributed. It is a combination of the data with a theoretical line, on which all data should lie if they were normally distributed. The effect of two transformations will be determined. Both log and square root transformation are commonly used distribution transformations.

3.2 Interpolation

3.2.1 IDW

The mathematical interpolation method applied is Inverse Distance Weighted (IDW). The assumption made for IDW is that the value of an attribute z at an unvisited point is a distance-weighted average of data points occurring within a neighborhood or window surrounding the unvisited point. It is forced to be an exact interpolator, since it produces infinities at the data points. Inverse distance interpolators commonly have a ‘bulls-eye’ pattern around solitary data points with values that differ greatly from their surroundings. IDW has no built in method of testing for the quality of predictions (Burrough and McDonnell, 1998)

3.2.2 Kriging

When spatial variation is too high to be described adequately by a simple function, the variation may be better described by a stochastic function as used in the Kriging method (Burrough and McDonnell, 1998). There are many types of Kriging, and each method has many options. Choosing the right Kriging method and the right options requires decision making. This allows to “tune” the interpolation strategy specifically for rainfall. Customisation of the approach to the problem at hand is a necessity for a good geostatistical study (Isaaks and Srivastava, 1989). Kriging uses a variogram model to characterise spatial correlation. A variogram describes in terms of variances how spatial variability changes as a function of distance and direction (Isaaks and Srivastava, 1989).

In this research, ordinary (point) Kriging and block Kriging have been used. (Burrough and McDonnell, 1998; Bivand et al., 2008). In short, ordinary Kriging is the basic form of Kriging. The prediction by ordinary Kriging is a linear combination of the measured values. The spatial correlation between the data, as described by the variogram, determines the weight. Block Kriging predicts averages of larger areas or volumes. It tends to smooth the prediction surface. Also, in contrast with ordinary kriging, the prediction errors are smaller, since much of the variability averages out.

A q-q plot is used to visualize the distribution of the measurements. Following the q-q plot, Kriging is performed on a transformed data set. To obtain back-transformed rainfall values simply taking the square of the Kriging prediction (based on the

square root-transformed rainfall data) makes the distribution positively skewed, resulting in a mean larger than the median. So percentiles are calculated, which are squared (Schuurmans et al., 2007).
As mentioned before, rainfall can be zero. This information is valuable for interpolation, therefore the predicted rainfall values were forced to contain the same percentage of zeros as in the dataset as in Schuurmans et al. (2007).

3.3 Validation

The current KNMI spline interpolation product is not validated by any measure of success. A measure of success may be the results of cross validation (CV) and or the Kriging variance. A way to implement CV is data splitting. However, the spatial variation of precipitation may be too high for a successful application of this technique (Sluiter, 2008). Another form of CV is Leave One Out Cross Validation (LOOCV). One by one, each measurement location is left out of the dataset, and at that point the value is interpolated. The difference is obtained and the process moves on to the next point (Burrough and McDonnell, 1998; Bivand et al., 2008).

3.3.1 Leave One Out Cross Validation (LOOCV)

LOOCV has been performed on IDW and ordinary Kriging. The cross validation results in residues: the difference between prediction and observation for all observations. With these residues an R^2 has been calculated, following the method described in Bivand et al. (2008). This method compares the residues with the mean. Thus, a higher R^2 means that the interpolation method performs better than the mean.

3.3.2 Quantitative map comparison

Maarten Plieger (KNMI) performed RMSE comparisons over the available data sets to investigate how they perform and relate to each other. See appendix 3 for a detailed description of the method.

3.3.3 Visual interpretation

The aim of validation is to pin down objectively which interpolation method performs optimal. Two quantitative methods are used, namely LOOCV and RMSE, but the importance of visual interpretation of the interpolation results must not be overlooked nor underestimated. Visual interpretation may seem subjective, but reviewing a lot of interpolation results gives valuable insight into the credibility of an interpolation method. Comparison with other spatial products increases this insight.
Radar images demonstrate clear spatial distribution patterns, which make it ideal to compare with the interpolation results. Radar data is available for a period of almost ten years.

4 Results and discussion

4.1 Data distribution and transformation

Due to its irregular nature, rainfall shows a non-normal distribution, especially on days with convective rainfall. See figure 4.1, left images. As mentioned in the previous chapter, for Kriging a (multivariate) Gaussian distribution is preferred. In general, square root transformation gives a good result in approximating the Gaussian model (Schuurmans et al., 2007). It doesn't heavily distort a distribution that is already quite good, like on April 4 2004, where a log transformation would exaggerate the non-normal distribution. Figure 4.2 shows variograms for the selected days. The model has been fixed to exponential and the nugget is set to zero, so the predictions are forced to coincide with the observations.

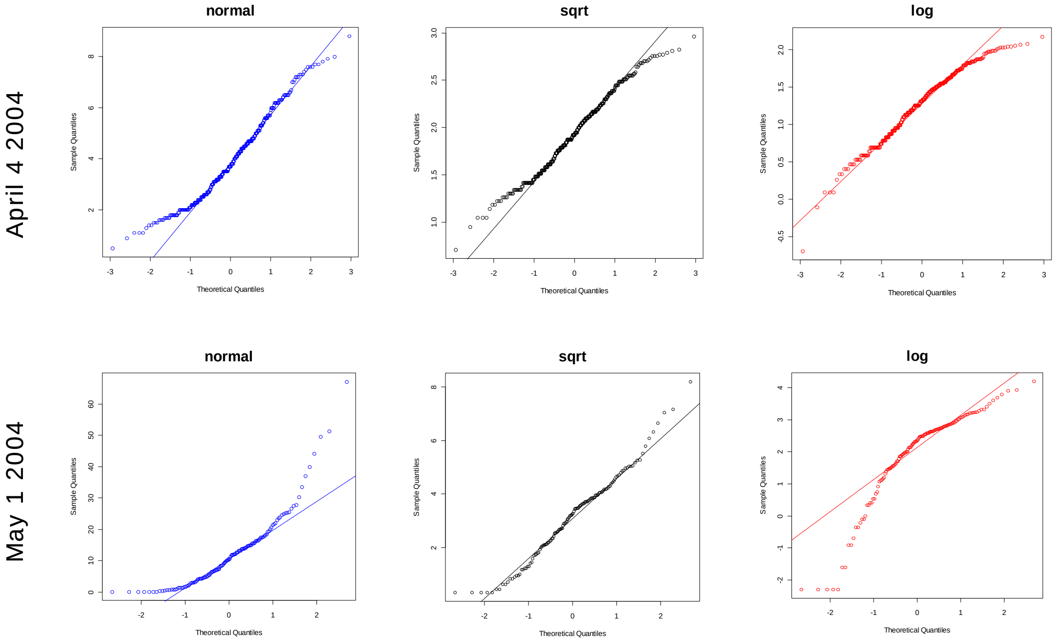

Figure 4.1 - normal qq-plots for April 4 2004 and May 1 2004: without transformation, with square root transformation and log transformation.

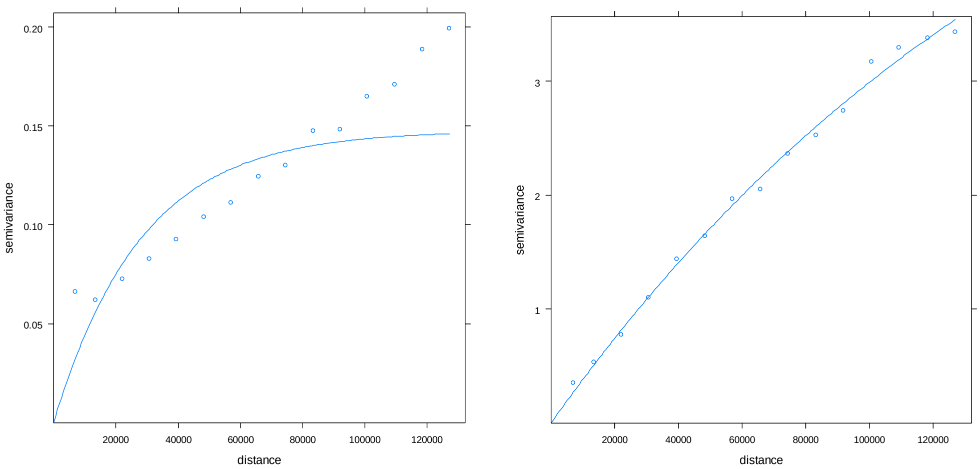

Figure 4.2 - Variograms for April 4 2004 (left) and May 1 2004 (right).

4.2 IDW interpolation

Usually, with abundant observations, any interpolation method can be expected to perform well. Because IDW is simple and fast, this method was performed initially.
Figures 4.3a and 4.3c show standard IDW interpolation for April 4 and May 1, 2004. The bulls-eye pattern is very strong, but the overall spatial variation is quite in concordance with the radar images in figure 4.5. The extent of the area with very low (<0.1 mm) precipitation is overestimated by IDW interpolation and exaggerated by this visualization; however this is a technical issue that can be easily solved.
To smooth out the bulls-eyes, a block IDW was executed as well. For block IDW the following parameters have been applied:

- Block size of 50km.
- Maximum search distance of 60km.
- Maximum number of nearest observations used for interpolation at an unknown location of 60.

The latter two parameters were set because IDW showed the best LOOCV R^2 for this combination. Moreover, calculation time is shortened using the above parameters.
The results are shown in figure 4.3b and 4.3c. Even with a block size of 50km, the bulls-eye pattern that is characteristic of IDW is still visible, especially on May 1, 2004. Also, a regular block pattern is introduced (artefact). Not surprisingly, the heavily smoothed block IDW method does not match the irregular distribution pattern as shown by the radar images in figure 4.5.

4.3 Block Kriging

Figure 4.4 shows results for block Kriging. The same parameters as for IDW with block have been applied (section 4.2). Observations are square root transformed and back-transformed after interpolation using quantiles calculation as described in the methods section. The main advantage over IDW is the absence of the bulls-eye pattern, inherently to the Kriging method. However, hardly anything is left of the spatial patterns that are visible in the radar images (figure 4.5). The small scale variation of rainfall distribution is smoothed out.

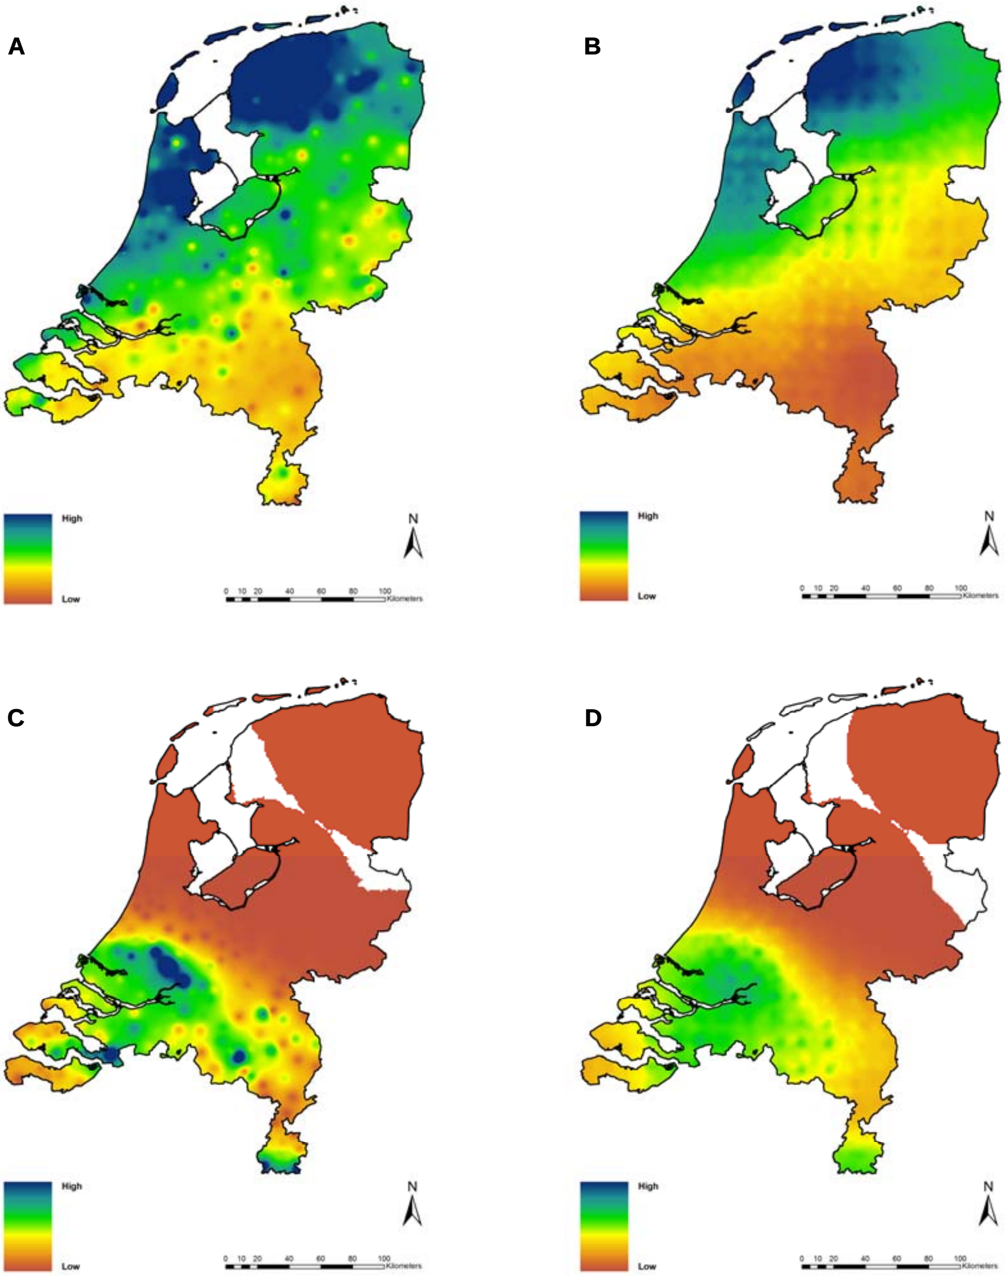

Figure 4.3 - IDW interpolation results. A) April 4, 2004 standard IDW B) April 4, 2004 IDW with block C) May 1, 2004 standard IDW D) May 1, 2004 IDW with block. Each image is individually scaled from “low” to “high” using standard deviation stretching to enhance the visualisation of the patterns. Therefore the absolute values differ in these images.

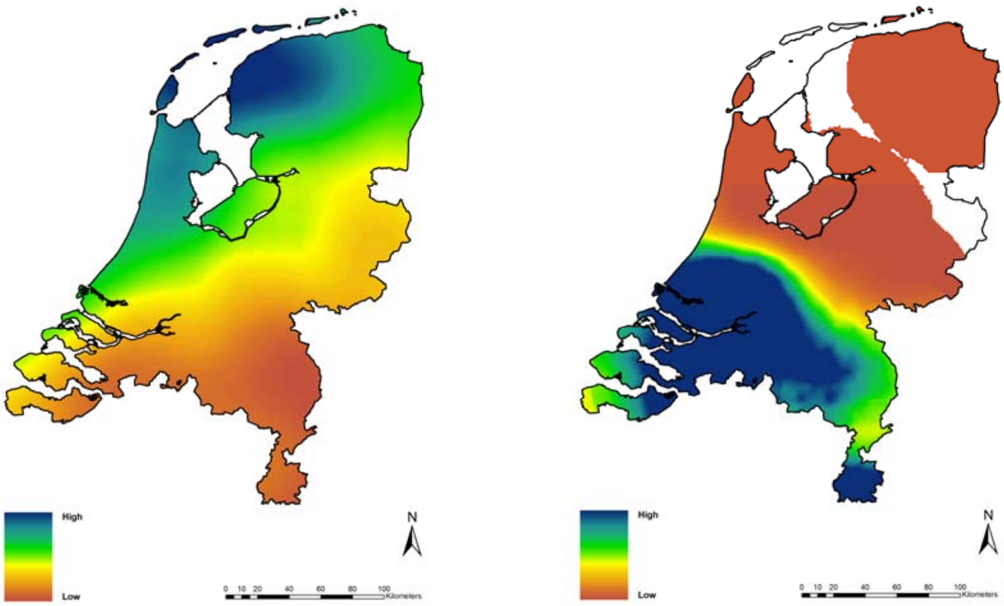

Figure 4.4 - Block kriging results for April 4, 2004 (left) and May 1, 2004 (right). Each image is individually scaled from “low” to “high” using standard deviation stretching to enhance the visualisation of the patterns. Therefore the absolute values differ in these images.

4.4 Ordinary Kriging

Figure 4.5 shows from top to bottom the ordinary kriging results, the 24h cumulative radar image and the HYRAS-REGNIE results. These three products, based on totally different methods, show great overlap for the two selected dates. Each of these products has advantages and disadvantages. Not surprisingly, radar images give the best representation of spatial patterns.

Ordinary Kriging is not as good as radar at showing spatial variability, but it is more accurate with estimating rainfall amounts. Moreover, Kriging is more “sophisticated” than IDW, since it uses the variogram to establish the distance weights used in the calculation. If we compare ordinary Kriging with IDW (figure 4.3) the main difference is the tendency of a bulls-eye pattern in the IDW results. Ordinary Kriging performs much better on this aspect. If we compare ordinary kriging with HYRAS-REGNIE we see that HYRAS-REGNIE is less smooth than ordinary kriging. The bulls-eye pattern is less pronounced in HYRAS-REGNIE than in the standard IDW interpolation (figure 4.3). However if we zoom into the HYRAS-REGNIE result a block/edge pattern appears as shown in figure 4.6. This pattern is likely caused by the background fields.

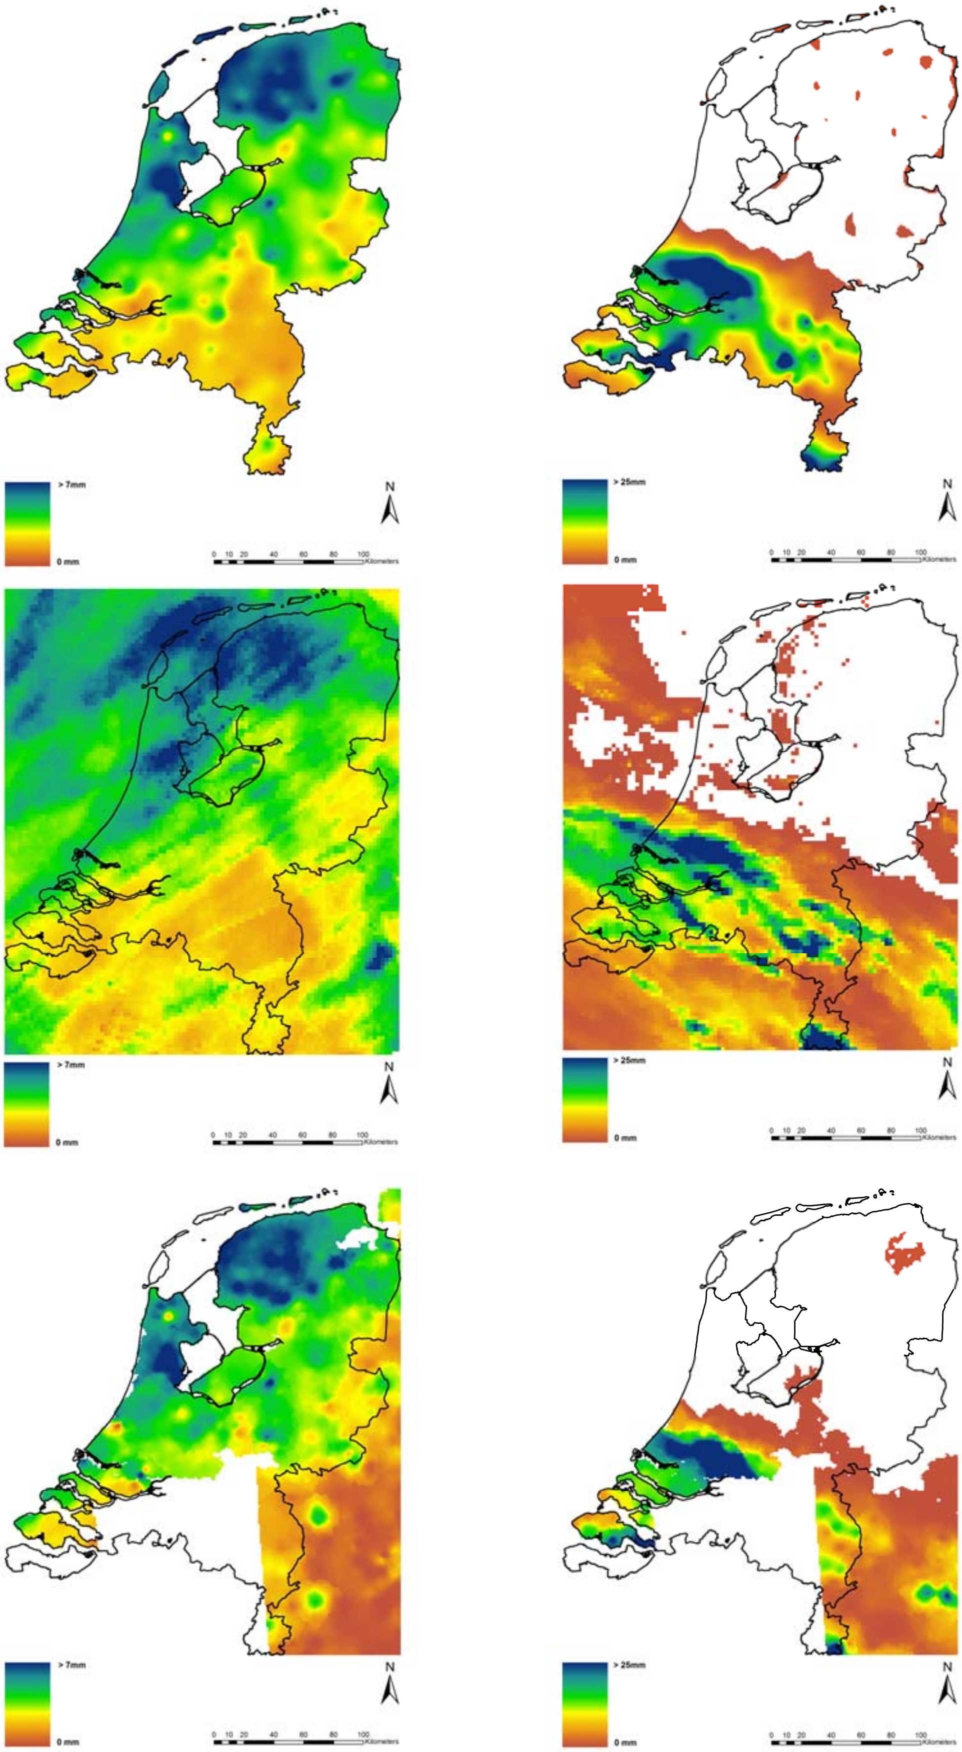

Figure 4.5 - Ordinary Kriging (top), radar (middle) and HYRAS-REGNIE (bottom). Images on the left represent April 4, 2004 (stratiform rainfall), on the right represent May 1, 2004 (convective rainfall). The images are scaled using linear minimum-maximum stretching. Maximum legend values differ between the two dates to enhance the visualisation of the patterns.

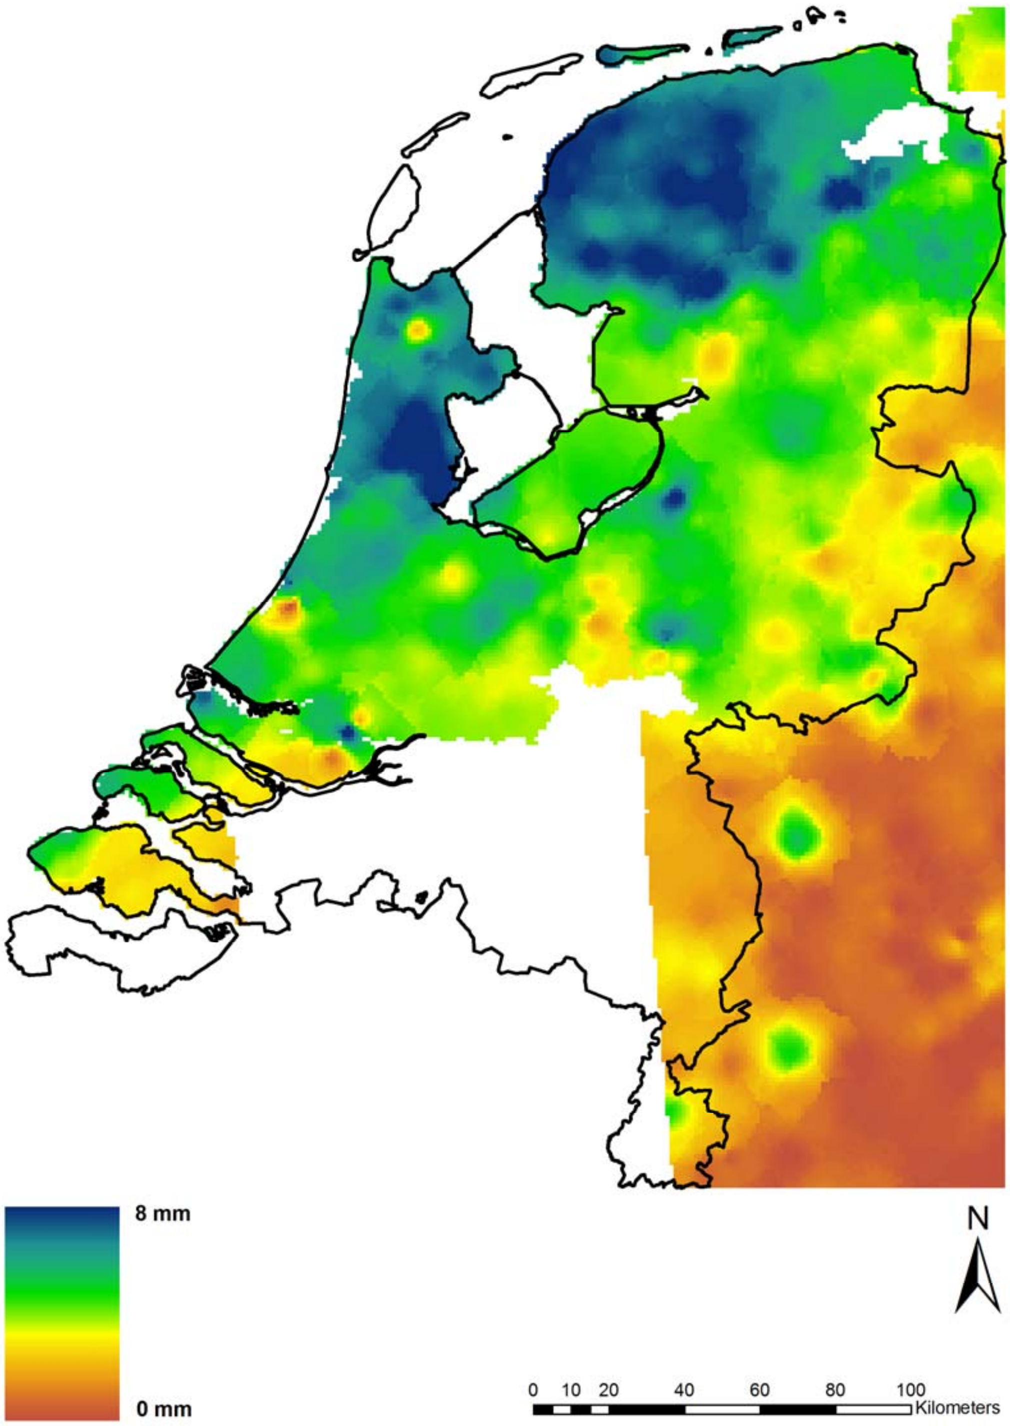

Figure 4.6 - HYRAS-REGNIE result showing the block/edge artefacts

5 Validation

5.1 Kriging variance

All Kriging methods create both a prediction map and a variance map. In figure 4.7 the Kriging variances for the two selected dates are shown. The behaviour of convective and stratiform rainfall is clearly visible, the variance on April 4 is practically zero while on May 1 it is quite large, due to unequal spreading of rain. In the April 4 image is visible that the Kriging variance relatively increases towards the Dutch borders, where unknown locations are no longer surrounded by observations.

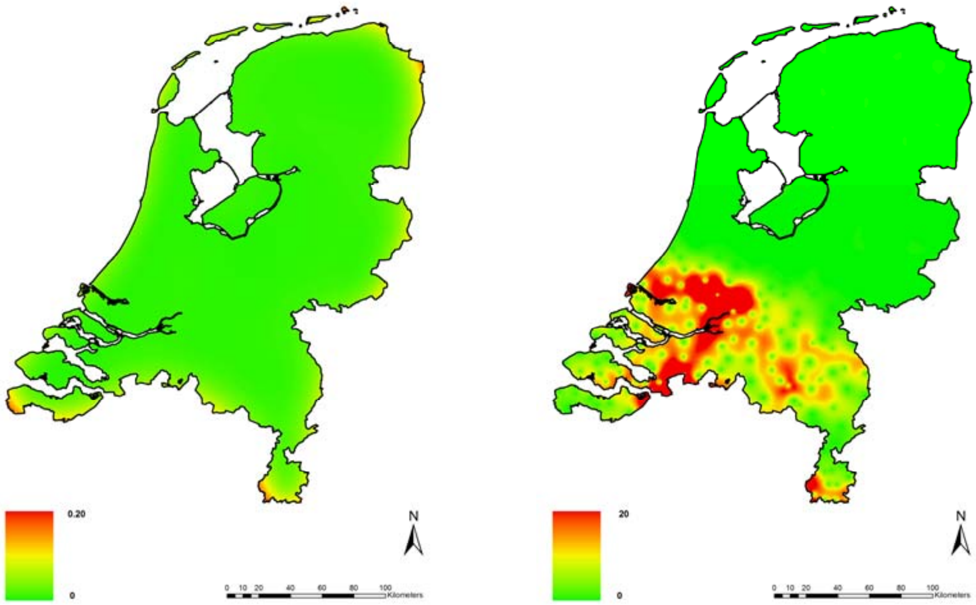

Figure 4.7 – Ordinary Kriging variance on April 4, 2004 (left) and May 1, 2004 (right). The images are scaled using linear minimum-maximum stretching. Maximum legend values differ between the two dates to enhance the visualisation of the patterns.

5.2 Leave One Out Cross Validation (LOOCV)

As can be seen in figure 4.8, ordinary Kriging has a higher LOOCV R^2 than IDW. For calculation of this boxplot, days with a total precipitation < 2 mm and a mean < 0.25 were left out. Both total and mean precipitation were calculated with all available observations on that day which have non-NA values. These limits were used because ordinary Kriging is sensitive to low precipition. Probably because it is harder to calculate a “robust” variogrammodel for low rainfall situations. Moreover IDW is overestimating low precipitation, resulting in better LOOCV R^2 values for low rain situations. More detailed information on this boxplot and R^2 calculation can be found in appendix 2.

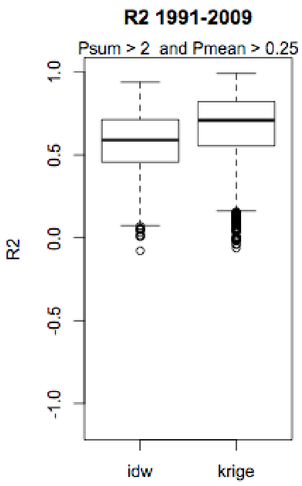

Figure 4.8 - Boxplot comparing LOOCV R^2 for standard IDW and ordinary Kriging, based on data from 1991-2009.

5.3 Root Mean Square Error (RMSE) comparison

The methods and results of the RMSE comparison are described in detail in appendix 3. Figure 4.9 shows a low RMSE for the combination KNMI (ordinary Kriging) and DWD (HYRAS-REGNIE). The RMSE is higher for the combinations KNMI-radar and DWD-radar, due to the generally smoother surfaces delivered by the interpolation methods. On average, the RMSE KNMI-radar is slightly lower than the RMSE DWD-radar.

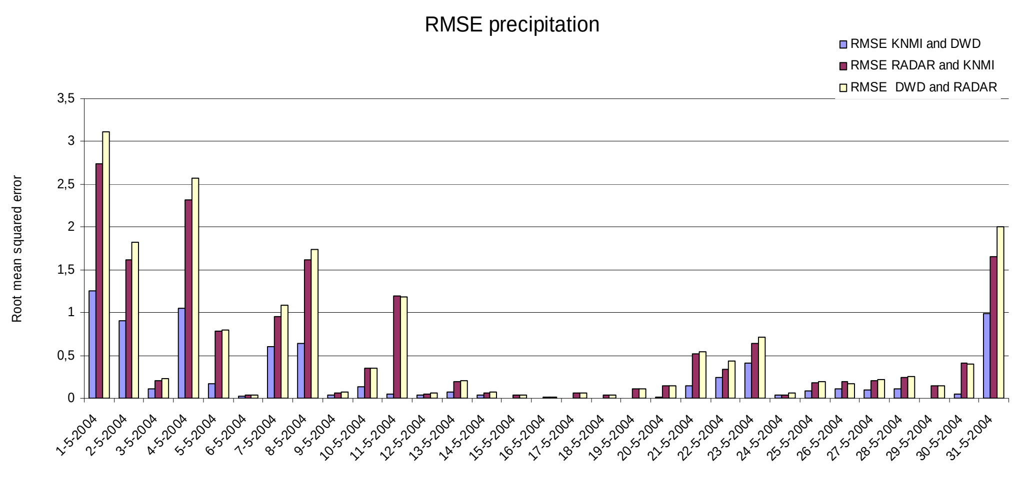

Figure 4.9 - Root mean squared error (RMSE) between the three grids: ordinary Kriging (KNMI), HYRAS-REGNIE (DWD) and radar.

6 Discussion

Ordinary Kriging outperforms standard IDW for interpolating rainfall data. LOOCV R^2 values are higher for ordinary Kriging when very low rainfall conditions are filtered out. The overall spatial patterns of standard IDW and ordinary Kriging match very well but the main drawback of standard IDW is the dominant and unrealistic bulls-eye pattern. The use of block-IDW to minimize the bulls-eye pattern produces a too smooth and unrealistic result. The same applies to block Kriging.

If we compare ordinary kriging with HYRAS-REGNIE we see that HYRAS-REGNIE is less smooth than ordinary kriging. If we consider the entire Dutch area covered by HYRAS-REGNIE as a whole, there is just little difference in the RMSE values between HYRAS-REGNIE and ordinary Kriging. The bulls-eye pattern is less pronounced in HYRAS-REGNIE than in the standard IDW interpolation but a block/edge pattern caused by the background fields is visible. Based on the RMSE study (section 5.3) Ordinary Kriging differs less from radar imagery than HYRAS-REGNIE. These factors in combination with the limited Dutch coverage of HYRAS-REGNIE make ordinary Kriging the preferred product for (hydrological) studies in the Netherlands. For international *large scale* studies ordinary Kriging and HYRAS-REGNIE can be combined without introducing too large errors: absolute values of ordinary Kriging and HYRAS-REGNIE differ not much.

A few preliminary tests with use of anomaly quotient, as in HYRAS-REGNIE and Haylock et al. (2008), show no great improvement over ordinary Kriging alone. The differences are very small and are only discernible at edges of rainfall fields: precipitation in the Netherlands is, as expected, only limited influenced by altitude compared to the rest of Europe, which is the extent of Haylocks research and the HYRAS-REGNIE research.

Cumulative radar images for days with considerable amounts of rainfall, often show distinctive patterns or trends due to prevailing winds, for both stratiform and convective precipitation. It is interesting to see if a distinction can be made between convective and stratiform precipitation, purely on data alone and to research if anisotropy can further improve spatial distribution. Also, the use of external variables like wind direction and distance to sea could improve predictions in specific cases. However, the improvements will be limited as it is impossible to reproduce the detailed patterns visible in the radar images.

As can be seen in figure 4.7, variance increases towards the border of the Netherlands. Incorporation of measurements from Belgium and Germany will likely decrease the variance. Actual implementation is complicated, since observation time differs between countries. Rainfall in the Netherlands is recorded at 8:00 UT, while the data used by HYRAS-REGNIE is recorded at 6:30 UT. Moreover access to the data is limited.

Observations with 0 mm rainfall give valuable information. So far, these 0 mm measurements are used to calculate a fraction, which is imposed on the result. It may be possible to use 0 mm information by using indicator Kriging, resulting in a boolean map: either it rains or it rains not. This map can then be used as extra input for the interpolation as in Haylock et al. (2008).

In this study all data is transformed using the same square root transformation to obtain a “more normal” distribution. A custom transformation for each day could be programmed in R.

7 Conclusions and recommendations

Based on this research we can conclude the following:

- Ordinary Kriging outperforms standard IDW.
- The overall prediction patterns of ordinary Kriging and HYRAS-REGNIE are quite similar but on the detailed level ordinary Kriging shows less artefacts. Moreover, ordinary Kriging differs less from radar imagery than HYRAS-REGNIE.
- Ordinary Kriging is the preferred product for (hydrological) studies in the Netherlands compared to HYRAS-REGNIE due to the artefacts and limited Dutch coverage of HYRAS-REGNIE.
- For international *large scale* studies the ordinary Kriging product and HYRAS-REGNIE can be combined without introducing too large errors.
- The use of an anomaly quotient, as in HYRAS-REGNIE and Haylock et al. (2008), shows no great improvement over ordinary Kriging alone because the effect of altitude is very limited in the Netherlands on daily basis.

Some aspects were left unattended in this research, but might improve the interpolation results:

- Anisotropy could lead to interpolation results that resemble cumulative radar images more closely. Especially on days where a clear directional pattern is visible.
- A custom transformation to make rainfall more normally distributed could be programmed in R instead of a fixed square root transformation for each individual day.
- Using data from weather stations from neighbouring countries has potential in reducing Kriging variance at the country's border. Actual implementation is complicated, since observation time differs between countries and access to the data is limited. This problem has also been encountered by Haylock et al. (2008).
- Interaction between factors that influence rainfall is complex. Haylock et al. (2008) suggest the use of stochastic simulations to quantify the probabilistic errors.
- An alternative to deal with 0 mm observations could be indicator Kriging, which as used by Haylock et al. (2008).

References

Bivand, R.S., E.J. Pebesma & V. Gomez-Rubio (2008). Applied Spatial Data Analysis with R: Springer.

Buishand, T.A., J.J. Beersma, R. Sluiter & T. Kroon (2008). Definitiestudie rasterdata meteorologie. De Bilt, KNMI / RWS Waterdienst, 2008

Burrough, P.A. & R.A. McDonnell (1998). Principles of Geographical Information Systems. Oxford: Oxford University Press.

Haylock, M.R., N. Hofstra, A.M.G. Klein Tank, E.J. Klok, P.D. Jones & M. New (2008). A European daily high-resolution gridded data set of surface temperature and precipitation for 1950-2006. Journal of Geophysical Research 113, pp. .

Hiemstra, P. (2009). Package ‘automap’. [23/02/2010] http://cran-mirror.cs.uu.nl/web/packages/automap/index.html.

Holleman, I. (2004). VPR adjustment using a dual CAPPI technique. Third European Conference on Radar Meteorology (ERAD), Visby, Island of Gotland, Sweden, Copernicus Gmbh 2

Isaaks, E.H. & R.M. Srivastava (1989). An Introduction to Applied Geostatistics.

R-Project (2009). The R Project for Statistical Computing. [01/03/2010] http://www.r-project.org/.

Salet, F.W.J. (2009). Het interpoleren van temperatuurgegevens. De Bilt, Royal Netherlands Meteorological Institute (KNMI). http://www.knmi.nl/bibliotheek/stageverslagen/stageverslag%20Salet.pdf

Schuurmans, J.M., M.F.P. Bierkens & E.J. Pebesma (2007). Automatic Prediction of High-Resolution Daily Rainfall Fields for Multiple Extents: The Potential of Operational Radar. Journal of Hydrometeorology 8, pp. 1204 - 1224.

Sluiter, R. (2008). Interpolation methods for climate data: literature review. De Bilt, Royal Netherlands Meteorological Institute (KNMI). http://www.knmi.nl/bibliotheek/knmipubIR/IR2009-04.pdf

Steiner, H. (2009). HYRAS report chapter 5 - Interpolationsverfarhen Deutsche Wetterdienst.

Appendix 1 - Observation locations

Observation locations of the voluntary network (2006).

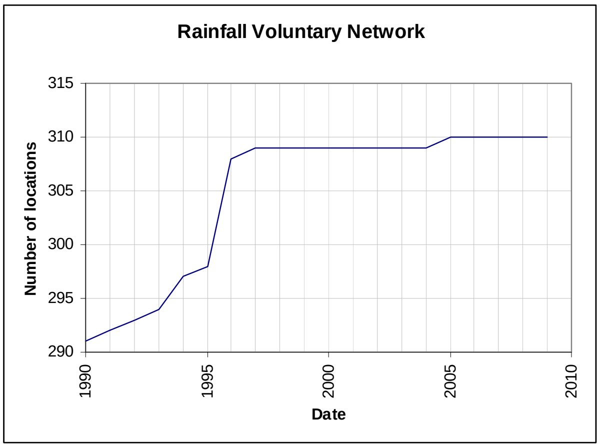

Number of observation locations through time.

Appendix 2 - Leave One Out Cross Validation (LOOCV)

Leave One Out Cross Validation (LOOCV) has been done according to the method as described by Bivand et al. (2008). It was first performed on the whole data set, including days with very low precipitation. See Figure A1-b. In this case, the average R^2 value is higher for ordinary Kriging but the range of R^2 values is much larger for ordinary Kriging than for standard IDW. Standard IDW is overestimating low precipitation, resulting in better LOOCV R^2 values for low rain situations.

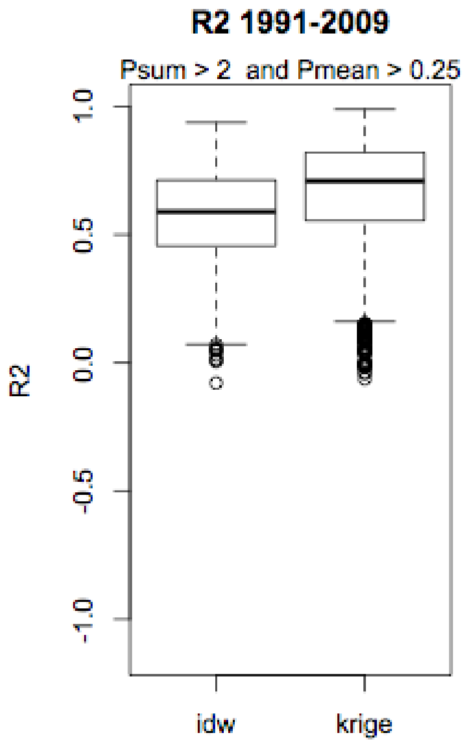

Figure A1-a

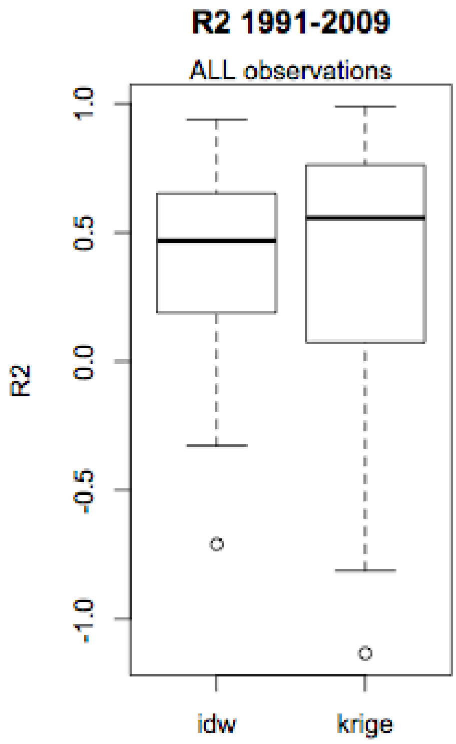

Figure A1-b

If days with low precipitation are left out (figure A1-a), ordinary Kriging performs better than IDW: the average R^2 value is higher and the range is smaller, although there are more outliers.
The limits for inclusion, Psum>2 and Pmean>0.25, are based on figure A2 and A3. This figure shows underlying data for April 2004. The grey bars in the background of the upper image display the percentage of stations with zero precipitation. The boxplot represents the distribution of precipitation values. Only three days have no precipitation at all. The limits Psum and Pmean were chosen rather arbitrarily, it would be better to have a more objective method. A trial to connect the R^2 with the *sserr*, which gives the sum of squares (standard output from `autofitVariogram`) of the variogram model, did not succeed: the correlation was very low.

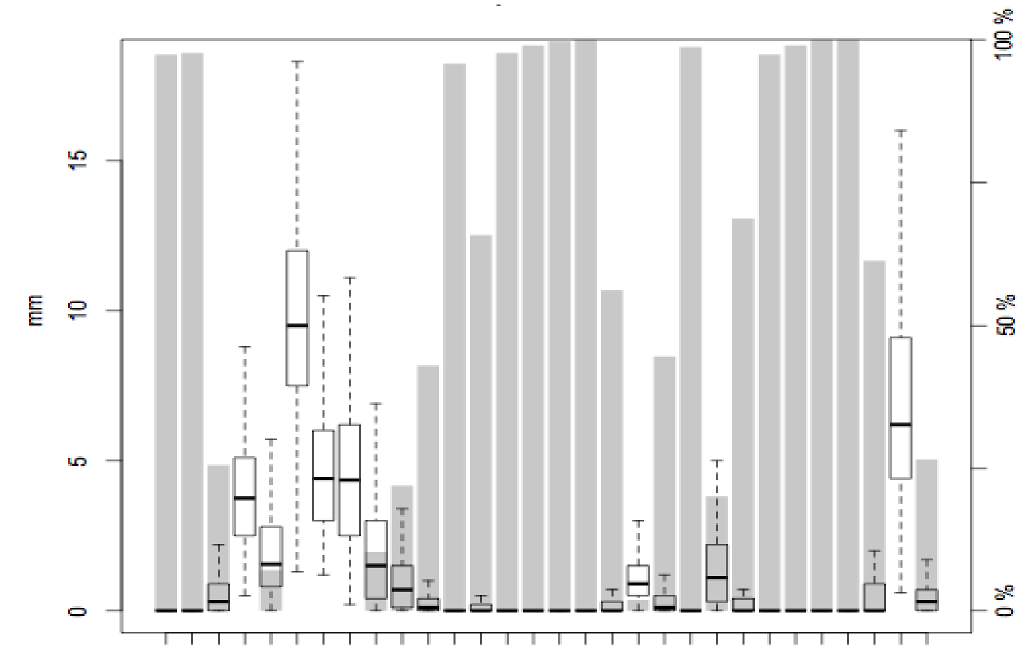

Figure A2 - Daily data for April 2004. Grey barplot: percentage of observations with zero precipitation. Boxplot: distribution of precipitation among observation stations.

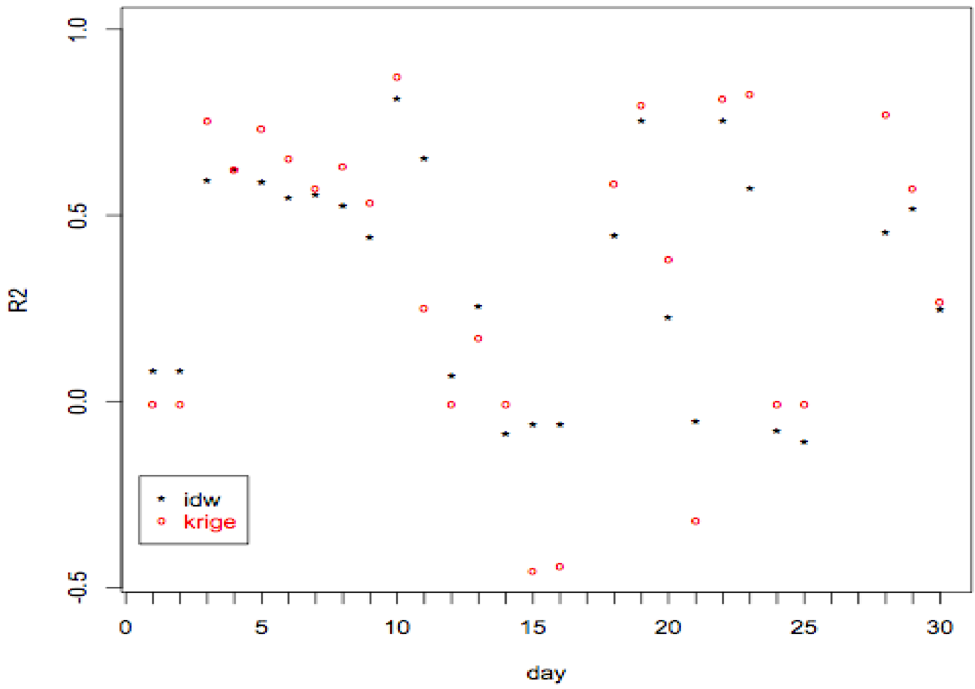

Figure A3 – LOOCV R^2 for standard IDW and ordinary Kriging.

Appendix 3 – Quantitative map comparison

Author: Maarten Plieger

The entire maps of the different interpolations results are compared using the mean and RMSE as indicators for where the predictions are deviating. All calculations are done for the areas where information in all grids is available: the Dutch coverage of the HYRAS-REGNIE interpolation (section 2.3.1 and figure A4). This coverage is also used for the sample points. At the locations where the grids do not contain information, the value of the stations is ignored.

Grid average
The mean represents the average of all pixels in the grid field. This calculation is applied for the pixels where all grids contain information. If a grid contains NA at a specific location, this location will not be used in all underlying grids.

RMSE – Root Mean Squared Error
The root mean squared error has been calculated by taking the root of the sum of all squares of the differences between each sample of grid A and grid B, divided by the total number of samples:

$$RMSE = \sqrt{\frac{\sum (A_i - B_i)^2}{n}}$$

Where *i* represents each individual pixel and n is the total number of overlaying pixels.

A low RMSE means that there is little variation between the datasets while a high RMSE means that there is more variation between the datasets. If the datasets have exactly the same spatial distribution and values the RMSE will be zero. The lower the RMSE, the lower the spatial difference between the datasets is. A low RMSE can therefore be interpreted as a better result for the interpolation.

Global mean and RMSE for precipitation
Besides information from the interpolation of the KNMI (ordinary Kriging) and DWD, information from the KNMI precipitation radar is also used for the comparison. Figure A4 shows the four datasets used.

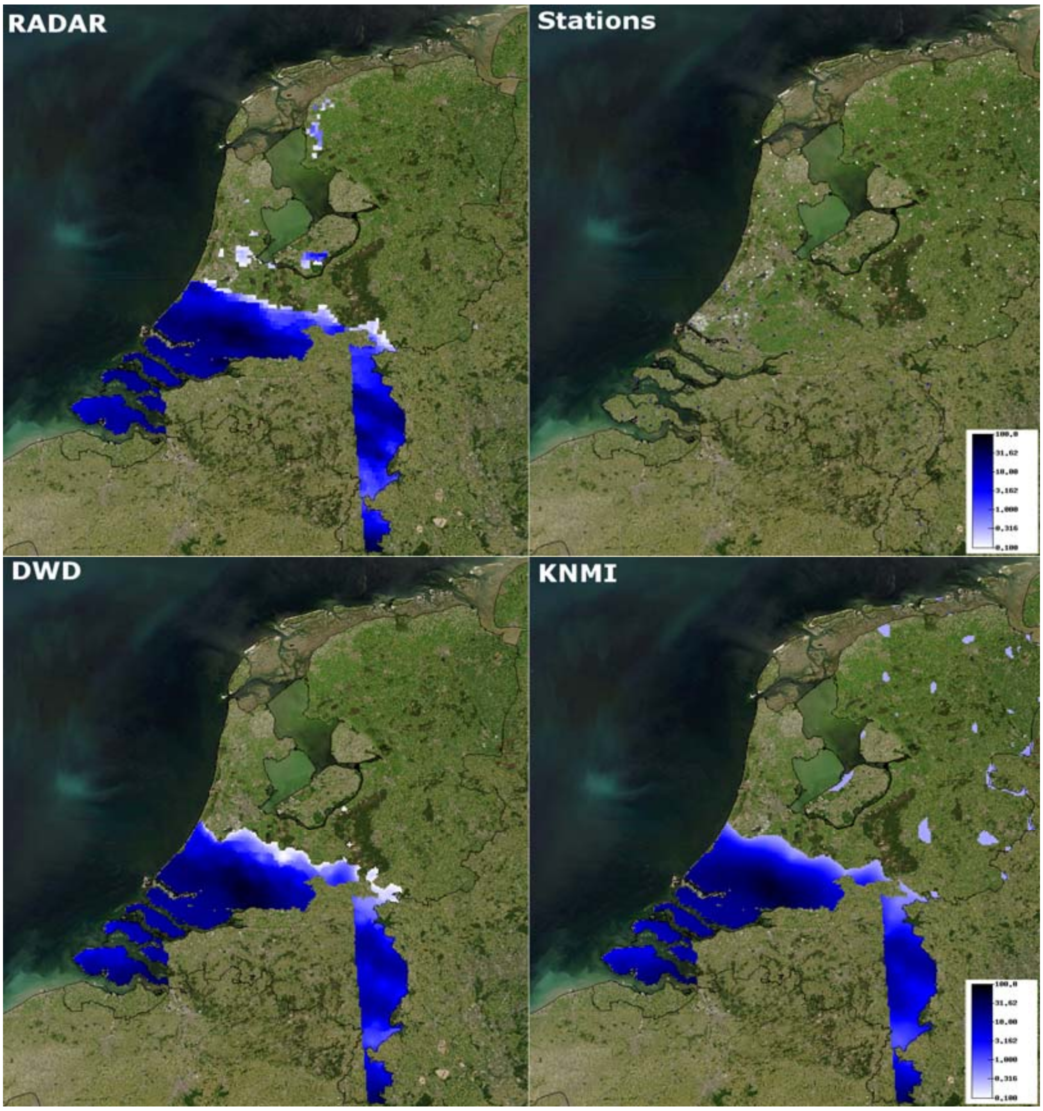

Figure A4 - Spatial coverage of the four precipitation datasets for May 1 2004: radar, station measurements, HYRAS-REGNIE (DWD) and ordinary Kriging (KNMI).

As can be seen from figure A5, the mean precipitation is very similar for all datasets. The trends in time are also similar for all datasets. The RMSE on the other hand is displaying strong differences between the datasets (figure A6). The RMSE between the two interpolations from DWD and KNMI is the lowest of all, while RMSE of the interpolations with the radar is higher. radar data provides a higher spatial variability at small spatial scales than the interpolations. In general the interpolations seem to have a smoother surface. The higher RMSE for the interpolations compared to radar can be explained by the lower spatial variability at close ranges for the interpolations.

Conclusion

The RMSE between the KNMI and DWD interpolations is lower than the RMSE between KNMI and DWD interpolations with the radar due to the generally smoother surfaces delivered by the interpolations. The RMSE between the KNMI and DWD interpolation results is very low, so we may conclude that both interpolation methods are robust for calculating the general patterns. However the detailed patterns differ as shown in section 3.3.3. Radar provides the best prediction for precipitation pattern and there is still (little) space for improvement in the interpolation algorithms.

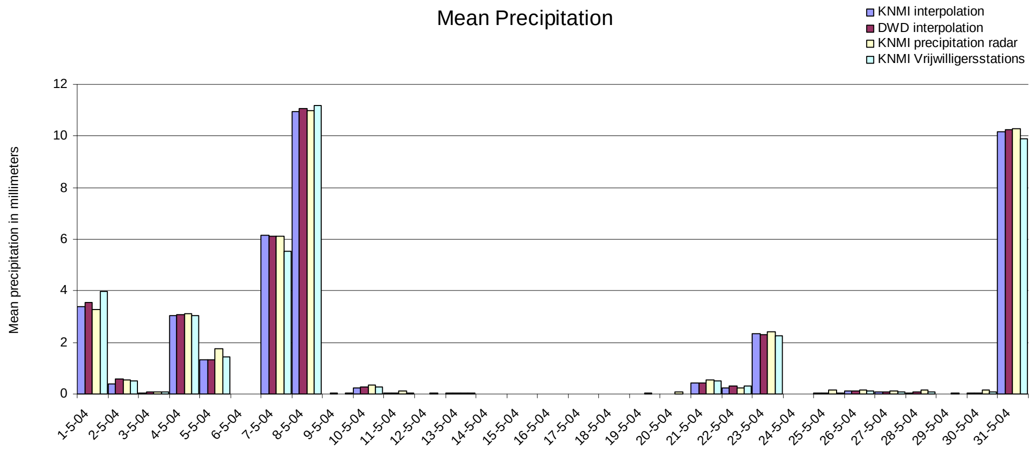

Figure A5 - Mean precipitation for ordinary Kriging (KNMI), HYRAS-REGNIE (DWD), radar and observations (Vrijwilligersstations).

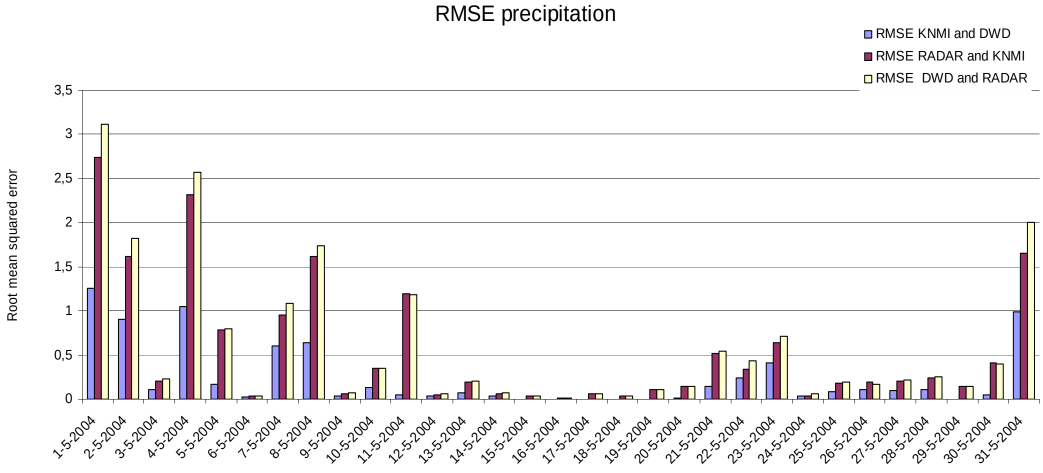

Figure A6 - Root mean squared error (RMSE) between the three grids: ordinary Kriging (KNMI), HYRAS-REGNIE (DWD) and radar.