# Benchmarking Knowledge-assisted Kriging for Automated Spatial Interpolation of Wind Measurements

**Zlatko Zlatev, Stuart E. Middleton, Galina Veres**
IT Innovation Centre
University of Southampton
Southampton, UK
{zdz,sem,gvv}@it-innovation.soton.ac.uk

*Abstract - We have benchmarked a novel knowledge-assisted kriging algorithm that allows regions of spatial cohesion to be specified and variograms calculated for each region. The variogram calculation itself is automated and spatial regions created via offline automated segmentation of either expert-drawn Google Earth polygons or NASA altitude data. Our use-case is to create interpolated wind maps for input into a bathing water quality model of microbial contamination. We benchmark our knowledge-assisted kriging algorithm against 7 other algorithms on UK met-office wind data (189 sensors). Our wind estimation results are comparable to standard ordinary kriging using variograms created by an expert. When using spatial segmentation we find our kriging error maps reflect better the known spatial features of the interpolated phenomenon. These results are very promising for an automated approach allowing on-demand datasets selection and real-time interpolation of previously unknown measurements. Automation is important in progressing towards a pan-European interpolation service capability.*

**Keywords:** Data Fusion, Kriging, Spatial Interpolation, Wind Speed, Wind Direction, OGC, WPS

## 1 Introduction

In-situ meteorological sensor measurements are generally recorded by sensor hardware at point locations, requiring some form of spatial interpolation should estimates at other locations be needed. Many spatial interpolation methods exist, both deterministic and geo-statistical, with accuracies dependent on the nature of the observed phenomena, spatial density of sensors, temporal frequency of sampling and the consistency and accuracy of measurement.

In the SANY project [6] we have developed an interpolation algorithm for handling wind measurements. Our use case is to generate spatial grids of estimated wind measurements for continual input into a bathing water quality model [13] and subsequent live prediction of microbial contamination levels of bathing water at beaches. Microbial contamination is important information for coast guards when making the decision as to if, and when, to close public beaches. Our algorithms are phenomena independent, and in SANY we have also successfully applied them to air pollution and ground displacement measurements.

We present a knowledge-assisted kriging algorithm applied to historical wind measurements from the UK meteorological office (UKMO) dataset archives [9]. We report cross-validated results for wind speed mean error (ME), root mean squared error (RMSE) and the range of estimated values. These results are compared directly to results for seven alternative interpolation methods reported in [7] on the same dataset.

Our knowledge-assisted kriging algorithm takes as input knowledge about known areas of spatial cohesion, either identified by an expert or heuristically calculated from the CGIAR-CSI GeoPortal SRTM (90m resolution) Digital Elevation Dataset [8]; multi-region ordinary kriging is then used to compute variograms for each region. The computation of the variograms is assisted by metadata representing wind phenomenon characteristics. This algorithm is hosted within an Open Geospatial Consortium (OGC) sensor service framework [5], showing how sensor processing services can be setup to 'plug and process' different sensor measurement datasets on-demand.

We outline in section 2 relevant related work, describe our algorithm in section 3 and experimental results in section 4. We discuss and conclude in sections 5 and 6.

## 2 Related work

In addition to ordinary kriging [1], which is outlined in the next section, there are a number of variants such as universal kriging [3], where a trend in data is assumed, and co-kriging [11], where covariates are provided in the same area as the primary sampled measurement to assist in prediction. Our novel automated variogram selection and multi-region ordinary kriging approach is another kriging variant, making use of expert knowledge about the phenomenon characteristics and the spatial cohesiveness in known regions to improve prediction accuracy.

Many general spatial interpolation techniques [12] exist, with most relevant ones including trend surface analysis, inverse distance weighting, local polynomial interpolation and thin plate spline. Until [7] few works had benchmarked these techniques thoroughly for wind

measurement interpolation, and our work builds on these results for direct comparison.

# 3 Knowledge-assisted ordinary kriging

Kriging [1] is a statistical technique for interpolation of random phenomenon that uses a linear combination of measured values at observed spatial locations to estimate the value at an unobserved location of interest. In contrast to other interpolation techniques, like inverse distance weighting and thin plane spline, it uses a model of the phenomenon's spatial correlation encoded as a corrrelogram or a semi-variogram. Kriging is an exact interpolator, so it respects the observed values at the observed locations, which differentiates it from interpolation techniques like local polynomial interpolation and trend surface analysis. For the purposes of our work we have used ordinary kriging with a semi-variogram correlation model. When using a semi-variogram, the spatial correlation of the phenomenon is quantified by means of a semi-variance function (1), where $N(h)$ is the number of pairs of observed locations $s_i$ a distance $h$ apart, and $z(s_i)$ is the observed value at $s_i$.

$$\hat{\gamma}(\mathbf{h}) = \frac{1}{2N(\mathbf{h})} \sum_{i=1}^{N(\mathbf{h})} \left[ z(s_i + \mathbf{h}) - z(s_i) \right]^2 \quad (1)$$

The semi-variogram can be a function of both distance and direction, so it can account for direction-dependent variability (anisotropic spatial pattern). For kriging, a smooth semi-variogram is required and we use smooth parametric functions as approximate models of the semi-variogram [2]. These smooth functions are fitted to the semi-variance function, or to the experimental semi-variogram obtained by averaging the semi-variance function over a set of distance lags, and after a goodness of fit analysis the best fitting function is selected [3] [4]. We refer to this function as a 'variogram model' and the experimental semi-variogram as an 'experimental variogram'. In our automated variogram model selection we use eight different families of variogram models: spherical, exponential, Gaussian, linear, power, generalised Bessel, sine hole-effect and cosine hole-effect. It is usual that a phenomenon expert will select a particular variogram model based on experience and fit the selected model to the available observations. Human expert intervention in the interpolation process is expensive and makes the interpolation process phenomena specific. Our solution, ordinary kriging with automated variogram model selection (AVMS), tackles this problem by utilising metadata representing high level phenomenon characteristics and automating the variogram model selection.

## 3.1 Automated interpolation workflow

The various offline and real-time stages of our knowledge-assisted interpolation process is shown in Figure 1.

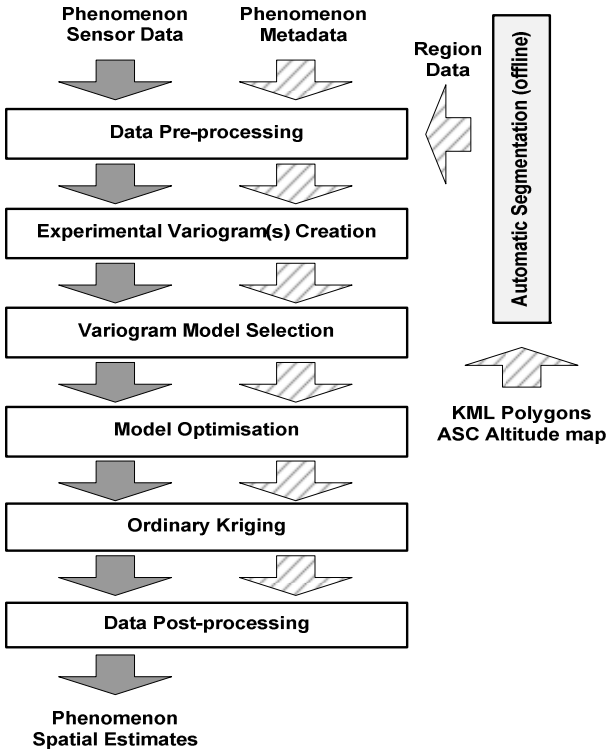


Figure 1. Ordinary kriging with automated variogram model selection procedure.

The first stage is the data pre-processing stage, where data cleansing, normalisation and all necessary data transformations are performed. The pre-processing stage includes input of knowledge-based descriptions for regions of spatial cohesion and offline preparation using segmentation techniques and heuristics prior to run-time. In the data post-processing stage, data de-normalisation and reverse transformations are performed. The core stages of ordinary kriging with AVMS are the experimental variogram creation, variogram model selection, model optimisation and ordinary kriging. A variogram is created for each region of spatial cohesion. Phenomenon metadata is setup via expert defined profile configurations.

## 3.2 Offline identification of knowledge about spatially coherent regions

Region calculation is performed automatically from either expert drawn Google Earth spatial polygons (KML format) or from NASA altitude data (ASC format) from the CGIAR-CSI GeoPortal SRTM (90m resolution) Digital Elevation Dataset [8]. The polygons / altitude maps are rendered as greyscale images and standard image processing techniques (colour reduction, binary mask per colour, pixel blur, labelling and edge identification) used to segment maps into unique regions suitable for input into the kriging process. Region segmentation is executed offline, via an automated web service, as part of the initial configuration stage prior to on-demand kriging.

We expect over mountainous ground the mean daily wind speed will have lower levels of spatial correlation than flat land. The land/sea boundary will also have an effect, especially since our sensors are land-based. Figure 2 shows the altitude segmentation, which we found to be somewhat over-segmented and too fine-grained for our wind phenomena. Figure 3 shows the region segmentation for the expert drawn polygons, which produced the best results and are used in the experiments later in this paper.

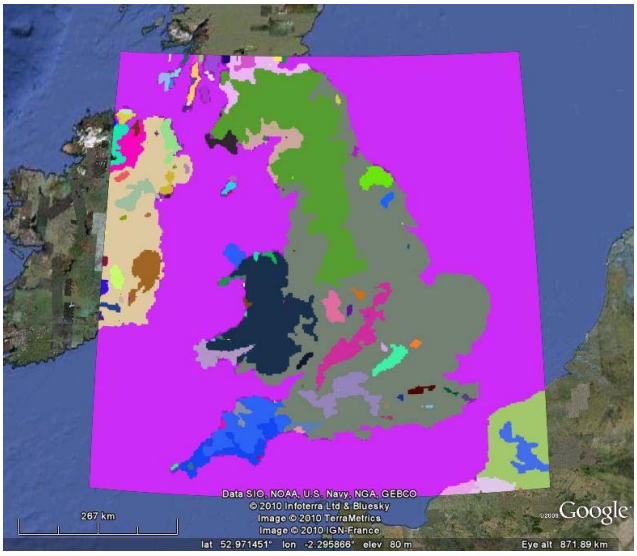


Figure 2. Region segmentation based on CGIAR-CSI GeoPortal SRTM altitude data

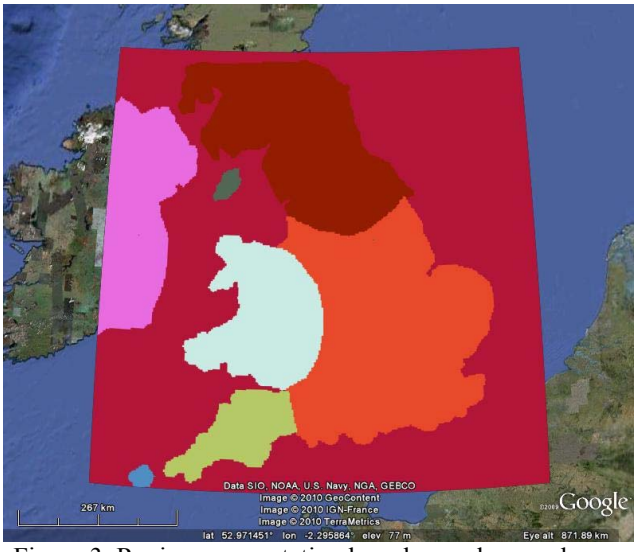


Figure 3. Region segmentation based on polygons drawn by an expert for land/water and flat/hilly areas

The output of the offline region segmentation process are a set of comma separated variable (CSV) maps containing region labels for every interpolation point, and inter-region neighbour linkage which is used later by the knowledge-assisted ordinary kriging algorithm.

### 3.3 Real-time automated variogram model creation

The most critical part of the experimental variogram creation stage is the selection of lags. Lags need to be selected so they contain an optimal number of points, do not smooth out physical phenomenon characteristics and avoid noise. Generally the initial slope of the variogram carries the most information, so our first few lags contain a smaller number of points. If no hole-effect is expected lags subsequent to the initial slope may contain a large number of points. If a hole-effect is expected these lags shall contain a lower number of points so the hole-effect is not smoothed out. The relative number of points in a lag is specified in the metadata supplied to the interpolation.

In the next stage, the variogram model selection, we fit our eight variogram models to the experimental variogram. The model shape is governed by a subset of the following parameters: nugget, correlation range, power, hole and sill. We use a weighted least squares fitting method to select a model that best fits the experimental variogram. We introduce expert knowledge about phenomenon by constraining the fitted model types and parameters' values and thus obtain a variogram model reflecting the characteristics of the phenomenon of interest.

For daily mean wind speed we expect low levels of noise to be present in the observations because of the averaging used to calculate the means (see section 4.1) and a low level of rapid phenomenon fluctuations. In the phenomenon metadata we thus set the nugget to be constrained to maximum of 20% of the sill. We expect the spatial correlation of the daily mean wind speed to decrease very slowly with increasing the distance. In the phenomenon metadata we thus set the lower and upper bounds of the correlation range parameter to be relatively high, respectively 25% and 75% of the maximum distance in the experimental variogram. We don't expect a hole-effect so do not select models with a hole-effect.

After selecting the variogram model, parameter optimisation is performed by minimising the mean error (ME, see equation (2)) and the root mean square error (RMSE, see equation (3)) calculated over 10-folds of cross-validation. We use simplex optimisation with a loss function of sum of ME and RMSE. The phenomenon constraints are reflected in the loss function by adding a high penalty value when the model parameters are outside the ranges specified in the phenomenon metadata.

Next, ordinary kriging with a moving neighbourhood is performed, using the optimised variogram model, and kriging mean and error computed. The kriging mean is the estimated value of the phenomenon at the location of interest. The kriging error is a measure of how good the observations' location configuration is for estimation at the location of interest.

Ordinary kriging with a moving neighbourhood is able to pick out variations of the phenomenon mean along the interpolated area. As the interpolated area is relatively

large in respect to the phenomenon scale we expect variations in the mean. Thus, in the phenomenon metadata we have specified a constrained moving neighbourhood consisting of the 11 closest observations to be used; the value of 11 neighbours was chosen after experimentation comparing results with different values. This choice can be automated by including an additional parameter in the model optimisation stage, which we intend to implement in future versions.

A point map showing region labels for the interpolated area is supplied to our ordinary kriging procedure so we know which region each interpolated point belongs to. The observations belonging to a particular region are used for working out the variogram of that region. If there are not enough observations available for a given region, observations from the neighbouring regions are used in order to build the raw variogram. The minimum number observations for the automated variogram estimation is *argmin(n*(n-1)/2>3* $N_{lags}$)*, where *n* is the number of observations and $N_{lags}$ is the number of lags in the variogram, i.e. we have minimum of 3 raw variogram points per lag. When computing estimates requiring observations from a neighbouring region we average the variograms of the home and the neighbouring regions. In the phenomenon metadata we have an inter-region correlation factor, *r*, ranging between 0 and 1, where 0 indicates minimal correlation between the regions and 1 indicates high correlation. We modulate averaged inter-region variograms by increasing the variogram values by a quantity equalling 1-*r* times the variogram sill, up to a value not larger than the sill. If the variogram doesn't have a sill, such as the power model, then variogram values are increased by a quantity equalling 1-*r* times the current variogram value. In this way we introduce additional knowledge of factors influencing the phenomenon spatial correlation.

# 4 Experiments and results

In the SANY project we have used our ordinary kriging with AVMS for interpolating micro-scale phenomena (e.g. ground displacement caused by underground tunnelling) and mini / meso-scale phenomena (e.g. wind speed and direction) yielding meaningful results.

In this paper we directly compare the performance of our algorithm against 7 well known algorithms and implementations on a benchmark data set. A definitive work in this context is [7], where the performance of seven interpolation methods are compared on a dataset of daily mean wind speeds obtained from sensor data supplied by the UK meteorological office (UKMO).

Similarly to [7] we have used leave-one-out cross validation and calculated the mean error (ME) and root mean square error (RMSE) as follows:

$$ME = \frac{1}{N}\sum_{i=1}^{N}\left[\hat{z}(s_i) - z(s_i)\right] \quad (2)$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left[\hat{z}(s_i) - z(s_i)\right]^2} \quad (3)$$

where $N$ is the number of validation folds, $\hat{z}(s_i)$ is the estimated value and $z(s_i)$ is the observed value.

## 4.1 Data

For our wind speed experiments we used the same dataset as [7] for 27th of March 2001, which was originally obtained from the UKMO. At this date 189 sensors in England, Wales, South of Scotland and Northern Ireland reported at least 12 times, the readings averaged to obtain daily mean wind speeds. The daily mean wind speeds vary from 2.2 to 13.6 m/s. Figure 4 depicts the locations and magnitudes of the daily mean wind speed observations, where the size of the dots is relative to the magnitude of the speed.

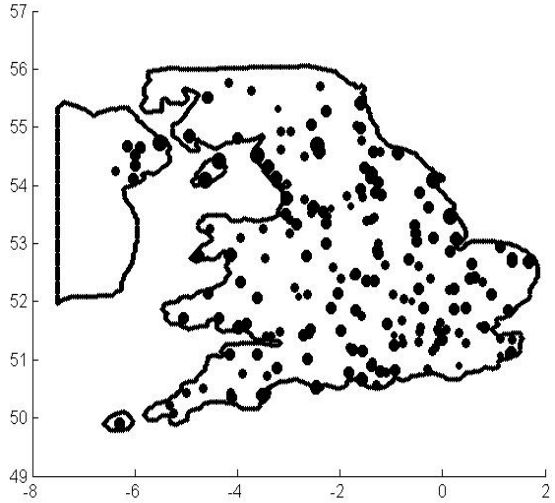


Figure 4. Locations and relative magnitude of daily mean wind speed observations for 27th of March 2001.

For our wind direction experiments we used 27th March 2001 data from the UKMO dataset archive [9]. It contains 157 observations with daily mean wind directions. Figure 5 depicts the locations of the daily mean wind directions observations, where the vectors show the direction of the wind.

We apply our ordinary kriging with AVMS to this data and interpolate on a grid with a resolution of 5km by 5km. For daily mean wind speed we compare our results against the results published in [7]. In [7] no results are presented for daily mean wind directions (just wind speed) so we just state our experimental results without comparison as a benchmark for other researchers.

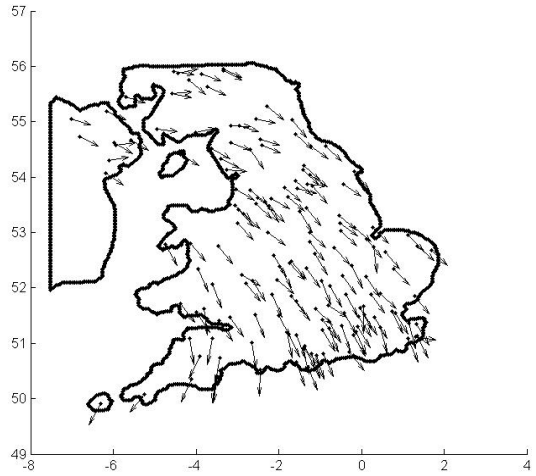

Figure 5. Locations of daily mean wind direction observations for 27th of March 2001.

## 4.2 Daily mean wind speed estimates benchmarking

In this experiment we compare our ordinary kriging with AVMS algorithm against the 7 algorithms evaluated in [7]. We set our algorithm with the following parameters:

1. Experimental variogram lags percentiles: 5, 10, 15, 20, 30, 40, 50, 60, 70, 80, 90, 100.
2. Nugget boundaries: 0 to 20% of the sill.
3. Correlation range boundaries: 25% to 75% of the maximum variogram distance.
4. Variogram models to use: spherical, exponential, Gaussian, linear, power or generalised Bessel.

***Ordinary kriging - no region knowledge used***

The first experiment is using kriging without any knowledge of regions and inter-region cohesiveness. The daily mean wind speed surface produced, shown in Figure 6, is consistent with the input values over land, and is very similar to the results found in [7]. High wind speeds are accurately estimated over the north, the north-east and north-west of England, and also over the Isle of Man. Low wind speed are accurately estimated over the Midlands, and the west and south-east of England. The kriging error, show in Figure 7, is small where we have high concentration of observations and larger where there are smaller number of observations.

The RMSE, ME and the min and max interpolated values suggest that the performance of ordinary kriging with AVMS is slightly worse than the ordinary kriging in [7] and slightly better than universal kriging and local polynomial (see table 1). It should be noted that our ordinary kriging with AVMS needs only a minimal input and configuration from users, unlike the classic manual approach used in [7].

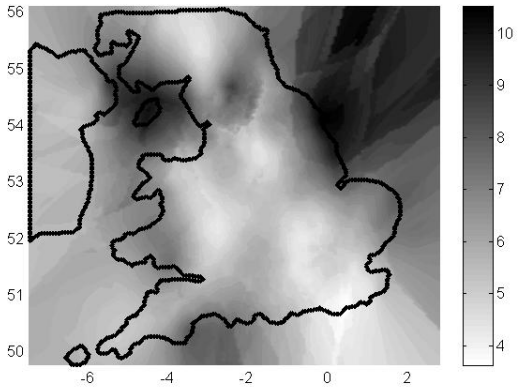

Figure 6. Daily mean wind speed estimates [189 sensors, no regions, 5km x 5km grid resolution, unit m/s]

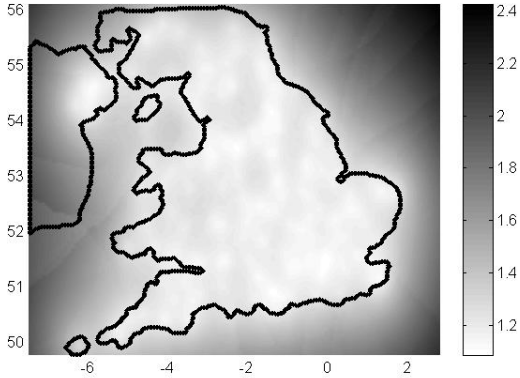

Figure 7. Kriging error of wind speed estimates [189 sensors, no regions, 5km x 5km grid resolution]. Notice how reported error is smoothed across coastline and land.

***Ordinary kriging - coastline region knowledge provided***

Since all sensors are land-based sea value estimates are in all probability not very reliable. Supplying expert-drawn coastline region boundary knowledge can help in this case.

For our second experiment we performed coastline segmentation on Google Earth drawn coastline polygons and supplied the resulting regions to our ordinary kriging with AVMS. For this experiment, the inter-region correlation factor *r* was set to 0.1 to indicate that there are minimal similarities in the way sea and land measurements influence each other. The estimated mean wind speed and the kriging error surfaces are shown in Figure 8 and Figure 9 respectively.

The in-land wind speed estimates change only slightly, with the only notable differences in estimated land values over the Isle of Man. Comparing figures 6 and 8 one can see that now there are more accurate high wind speed values predicted as the mainland observation's interference is diminished somewhat by the low inter-region correlation factor. This result is also reflected in the algorithm's overall performance metrics shown in table 1.

The kriging error for the Isle of Man is very different from the previous kriging error without using region

knowledge; this is clear comparing figures 7 and 9. There is a clear jump in the kriging error on the land-sea borders indicating correctly that the wind speed estimates over the sea areas are unreliable.

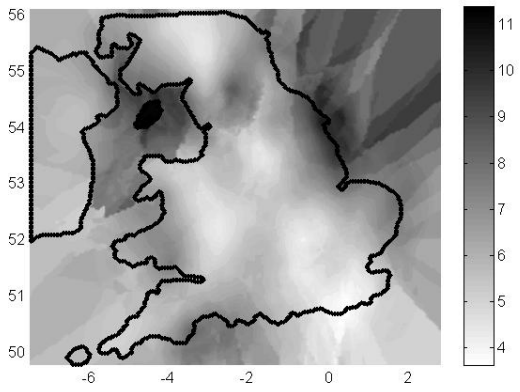


Figure 8. Daily mean wind speed estimates [189 sensors, coastline regions, 5km x 5km grid resolution, unit m/s]

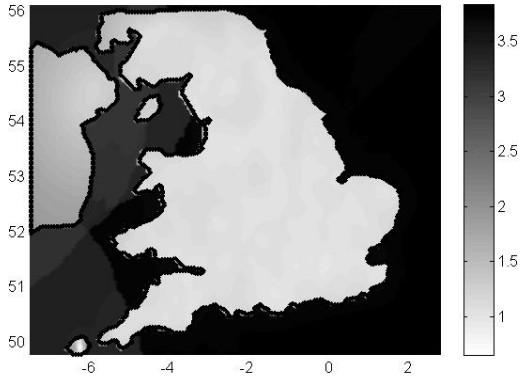


Figure 9. Kriging error of wind speed estimates [189 sensors, coastline regions, 5km x 5km grid resolution]. Notice the sharp error boundary at the coastline.

***Ordinary kriging - flat/hilly region knowledge provided***

For our last experiment regarding wind speed we provided domain knowledge about the different terrain topology (i.e. flat land, hilly land, mountainous land), again segmented into regions. We used a higher inter-region correlation factor of 0.9 since these new regions are in-land.

Results with these new regions are again improved and are comparable to ordinary kriging in [7] (see table 1). The bias, ME, is decreased as is the RMSE and the minimum and maximum of the estimated values are also closer to those present in the raw data. The estimated wind speed surface sown in Figure 10 suggests that with appropriate regionalisation the extremities are more likely to be identified but not smoothed out.

The kriging error map, Figure 11, now shows distinct patterns in each of the different regions. For example the kriging error over the Midlands, South and South-East region is very smooth, which we interpret as not needing very dense observations to achieve good estimates in this region. In contrast the spotty kriging error pattern in the northern region, and in Wales and south-west of England, we interpreted as needing much more dense observations to achieve good estimates in these regions.

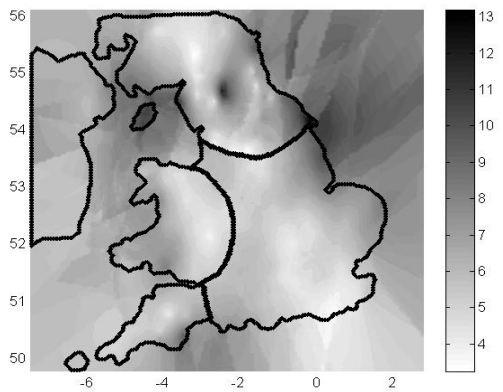


Figure 10. Daily mean wind speed estimates [189 sensors, flat/hilly regions, 5km x 5km grid resolution, unit m/s]

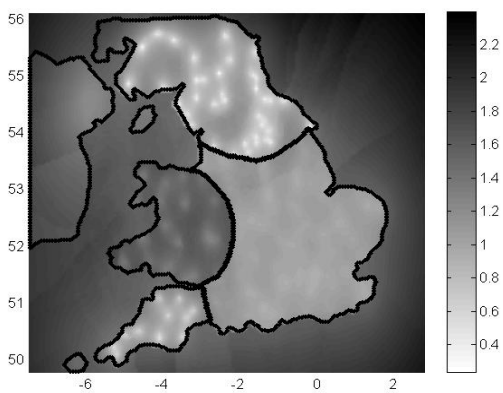


Figure 11. Kriging error of wind speed estimates [189 sensors, flat/hilly regions, 5km x 5km grid resolution]. Notice the sharp error boundaries between spatial features and different inter-region error profiles.

## 4.3 Daily mean wind direction estimation

Using the same ordinary kriging with AVMS setup as for the wind speed experiments we ran tests on the UKMO dataset for 27th March 2001 with no regions defined. In our implementation of ordinary kriging we have a special procedure [10] for handling periodic values like wind direction, which includes vector rotation and Cartesian transformation and simulation.

Estimated wind direction vectors are visualised in Figure 12. We use a 28km x 28km grid to make visualisation easier for arrow vectors. The ME and RMSE are given in table 2. We find the ME, of -0.60 degrees, and RMSE, of 13.66 degrees, in relative terms very low considering the worse case could be 180 degrees. The mean relative absolute error (MRAE) is 6%. For our wind speed experiments MRAE varies from 22% to 23%.

For wind speed we calculate MRAE as:

$$MRAE = \frac{1}{N}\sum_{i=1}^{N}\left[ |\hat{z}(s_i) - z(s_i)| / z(s_i) \right]$$

and for wind direction we calculate MRAE as:

$$MRAE^{(periodic)} = \frac{1}{N}\sum_{i=1}^{N}\left[ norm(|\hat{z}(s_i) - z(s_i)|)/180 \right]$$

where $N$ is the number of validation folds, $\hat{z}(s_i)$ is the estimated value, $z(s_i)$ is the observed value, $|.|$ is the absolute value and *norm*() yields a normalised wind direction in the range of -180 to +180.

The wind direction kriging error is shown in Figure 13, with normalised values in the range of 0 to 360. Kriging error quantifies the observations configuration suitability for estimating at a particular location and is not a measure of the estimation error in itself.

We are unaware of any other benchmark wind direction interpolation results so publish here in the hope of providing other researchers with a benchmark result set.

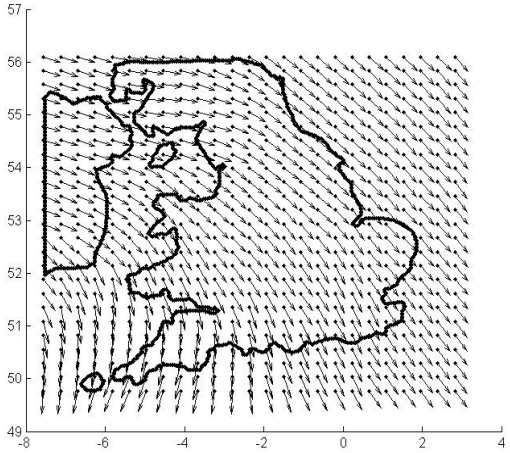


Figure 12. Daily mean wind directions [157 sensors, no region, 28km x 28km grid resolution, unit degrees]

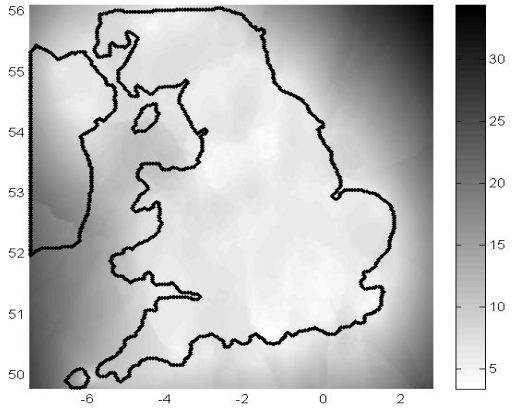


Figure 13. Kriging error for wind direction estimates [157 sensors, no region, 28km x 28km grid resolution]

| Method | Number of Regions | Inter-region correlation | Exec-time (mins) |
|---|---|---|---|
| Kriging_1 | 0 | n/a | 1 |
| Kriging_2 | 4 coastline | 0.1 | 8 |
| Kriging_3 | 7 flat/hilly | 0.9 | 10 |
| All methods use ordinary kriging with AVMS<br>Computer spec : 2 CPU's (2.4 GHz) with 4 Gbyte RAM | | | |

| Method | Min estimate m/s | Max estimate m/s | RMSE m/s | ME m/s |
|---|---|---|---|---|
| Kriging_1 | 3.6 | 10.7 | 1.67 | -0.03 |
| Kriging_2 | 3.6 | 11.6 | 1.67 | -0.03 |
| Kriging_3 | 3.24 | 13.40 | 1.60 | -0.01 |
| Cokriging [7] | 2.6 | 13.6 | 1.47 | -0.01 |
| Ordinary kriging [7] | 2.6 | 13.4 | 1.61 | -0.01 |
| Local polynomial [7] | 0.1 | 11.9 | 1.69 | 0.01 |
| Universal kriging [7] | 3.2 | 9.8 | 1.71 | 0.01 |
| IDW [7] | 3.1 | 12.1 | 1.74 | -0.09 |
| TPS [7] | 1.5 | 13.8 | 1.89 | -0.05 |
| TSA [7] | 4.6 | 9.2 | 1.93 | -0.02 |

Table 1. Comparison of knowledge-assisted kriging to benchmark interpolation results reported in [7]

| Method | ME degrees | RMSE degrees |
|---|---|---|
| Ordinary kriging with AVMS | -0.60 | 13.66 |

Table 2. Wind direction estimate ME and RMSE [157 sensors, no region, 28km x 28km grid resolution]

## 5 Discussion

In our experiments we have investigated methods to automatically select variogram models and ways to introduce expert knowledge into variogram calculation process. Our results show the impact using different levels of expert knowledge, about both the phenomenology and spatial areas of interest, can have on estimation error. This automated approach is in contrast to existing methods [7] that rely on manual variogram tuning by experts. We think our approach offers a realistic route to providing on-demand kriging services that can interpolate measurement datasets selected by users and that are unknown until run-time. This level of automation is becoming more important as we see European environment agencies sharing increasing amounts of sensor data under the INSPIRE directive [15]. Such automation, and a move away from manual tuning and configuration, increases scalability and could allow truly dynamic pan-European interpolation for wind measurements and other phenomenology.

The performance of our ordinary kriging with AVMS algorithm is slightly worse than that reported by ordinary

kriging in [7] with expert tuned variogram. Partly this is due to the expert's skill in manually tuning the variogram. However, our implementation is also only isotropic, which accounts for some of the performance compared against the anisotropic ordinary kriging algorithm in [7]. We intend to implement anisotropic kriging in future releases.

Where our knowledge-assisted approach offers the most improvement, compared to basic ordinary kriging, is in the enhanced spatial definition of the kriging error maps, and therefore improved confidence in them. This is most clear in areas where no sensors are located (e.g. offshore in the sea). This result can be applied more widely to other spatial features than hills and coastlines, at a variety of different spatial scales, such as building footprints and river outlets. Though not the focus of our wind interpolation paper, we have successfully interpolated measured phenomenon including ground displacement, water salinity and turbidity.

## 6 Conclusion

We have benchmarked a novel knowledge-assisted kriging algorithm that allows knowledge of regional spatial cohesion to be specified and variograms calculated per region. The variogram calculation itself is automated and phenomenon specific metadata allows us to configure kriging for more than just wind phenomenon. Spatial regions are created offline automatically by segmentation of either expert-drawn Google Earth polygons or NASA altitude data.

Our use-case is to create wind interpolation grids for input into a bathing water quality model of microbial contamination, and subsequent decision support for beach attendants to make bathing risk assessments.

We benchmark our knowledge-assisted kriging algorithm against 7 other algorithms using the same UK met-office wind measurement dataset reported in [7]. Wind speed estimation results are comparable, but not better than ordinary kriging, but the kriging error maps are much sharper and reflect the known spatial features better. We also provide results for wind direction interpolation.

These results are very promising considering it is an automated approach allowing on-demand datasets to be selected and real-time interpolation of measurements unknown a-priori. Automation is important in moves towards a pan-European interpolation service capability making use of European environment agency data shared in compliance with the European INSPIRE directive [15].

This work was funded by the European Commission's IST Programme under contract FP6-IST 0033564 SANY [6]

## References

[1] D.G. Krige, *Two-dimensional weighted average trend surfaces for ore-evaluation*, Journal of the South African Institute of Mining and Metallurgy 66: 13–38, 1966

[2] N.A.C. Cressie, *Fitting variogram models by weighted least squares*, Journal of the International Association for Mathematical Geology 17: 653–702, 1985

[3] N.A.C. Cressie, *Statistics for Spatial Data*, John Wiley and Sons: New York, 1993

[4] P.A. Burrough and R.A. McDonnell, *Principles of Geographical Information Systems*, Clarendon Press: Oxford, 1998

[5] S.E. Middleton (ed) et al, *SANY Fusion and Modelling Architecture*, OGC Discussion paper 10-001, 2010

[6] SANY project, http://www.sany-ip.eu

[7] W. Luo, M.C. Taylor and S.R. Parker, *A comparison of spatial interpolation methods to estimate continuous wind speed surfaces using irregularly distributed data from England and Wales*, Int. J. Climatol, 28: 947–959, 2008

[8] CGIAR-CSI, *GeoPortal SRTM (90m resolution) Digital Elevation Dataset*, http://srtm.csi.cgiar.org

[9] UK Met Office, *MIDAS Land Surface Observation Stations Data*, http://badc.nerc.ac.uk/data/ukmo-midas/

[10] Z. Zlatev, S.E. Middleton and G. Veres, *Ordinary kriging for on-demand average wind interpolation of in-situ wind sensor data*, EWEC 2009, France, March 2009

[11] A.G. Journel, C. Huijbregts, *Mining Geostatistics*, Academic Press: New York, 1978

[12] N.S. Lam, *Spatial interpolation methods: a review*, The American Cartographer 10: 129–135, 1983

[13] Z.A. Sabeur (ed) et al, *Development of environmental information tools for the prediction of water quality risks in bathing waters*, Final Report to the Environment Agency of England and Wales, Interreg ICREW project, Pilot Action 4: Forecasting bathing water quality. Final Technical Report/FINAL. 21st June 2006, pp128, 2006

[14] R.B. Stull, *Meteorology for Scientists and Engineers*, 2nd Edition, Earth Sciences series, Cengage Learning Publisher. ISBN-13: 9780534372149, 1999

[15] INSPIRE directive, *Directive 2007/2/EC of the European Parliament*, 14 March 2007