Some safe and sensible shortcuts for efficiently upscaled updates of existing elevation models.

Thomas Knudsen
Danish Geodata Agency
thokn@gst.dk

Allan Aasbjerg Nielsen
Technical University of Denmark
aa@space.dtu.dk

EGU General Assembly, Vienna, April 2013

Introduction

The Danish national elevation model, DK-DEM, was introduced in 2009 and is based on LiDAR data collected in the time frame 2005–2007. Hence, DK-DEM is aging, and it is time to consider how to integrate new data with the current model in a way that improves the representation of new landscape features, while still preserving the overall (very high) quality of the model.

To estimate the magnitude of potential improvements, and to devise efficient and effective ways of integrating the new and old data, we currently carry out a number of case studies based on comparisons between the current terrain model (with a ground sample distance, GSD, of 1.6 m), and a number of new high resolution point clouds (10-70 points/m^2).

Preliminary results reveal that for very high resolution data in smooth terrain (which is the common case in Denmark), using local mean (LM) as grid value estimator is only negligibly worse than using the theoretically "best" estimator, i.e. ordinary kriging (OK) with rigorous modelling of the semivariogram. The bias in a *leave one out* cross validation differs on the micrometer level, while the RMSE differs on the 0.1 mm level (see below).

This is fortunate, since a LM estimator can be implemented in plain stream mode, letting the points from the unstructured point cloud (i.e. no TIN generation) stream through the processor, individually contributing to the nearest grid posts in a memory mapped grid file.

Algorithmically this is very efficient, but it would be even more efficient if we did not have to handle so much data.

Another of our recent case studies focuses on this. The basic idea is to *ignore data that does not tell us anything new*. We do this by looking at anomalies between the current height model and the new point cloud, then computing a *correction grid* for the current model. Points with insignificant anomalies are simply removed from the point cloud, and the correction grid is computed using the remaining point anomalies only.

Hence, we only compute updates in areas of significant change, speeding up the process, and giving us new insight of the precision of the current model which in turn results in improved metadata for both the current and the new model.

Location and test sites

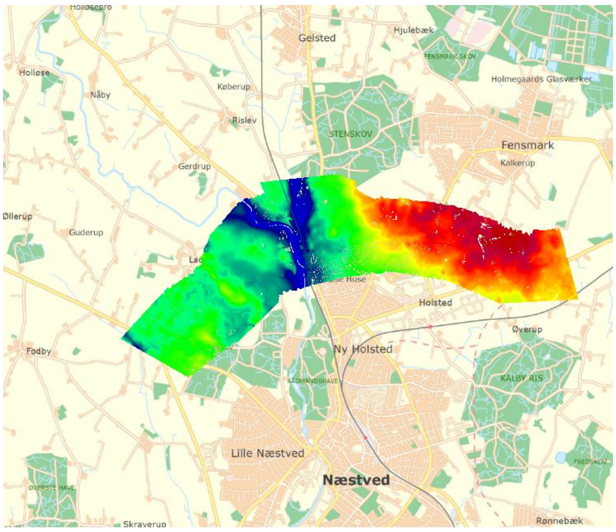

The full areal coverage of the test survey between Næstved and Fensmark in eastern Denmark. Heights range from 3 m (blue) to 35 m (red). Test sites are selected from a partial block starting roughly one quarter into survey, measured from its westmost point.

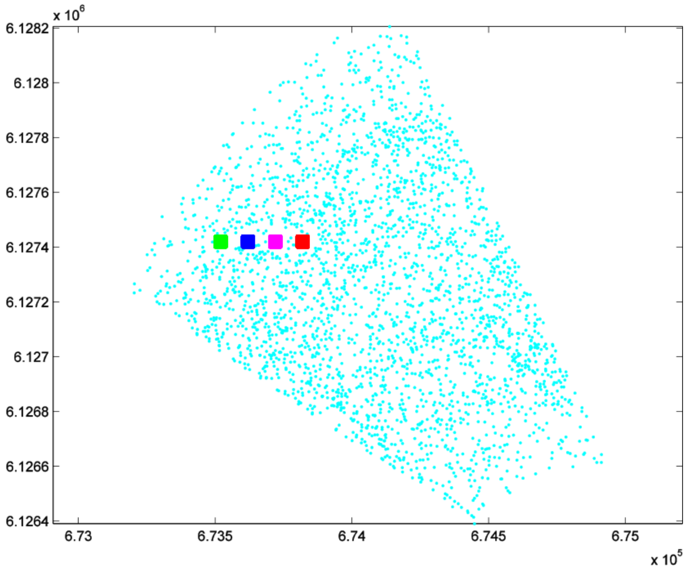

Test sites: site A is shown in red, site B in green, site C in blue, and site D in magenta. Axis values are UTM zone 32 coordinates.

Danish Ministry of the Environment
Danish Geodata Agency

Scan characteristics

The test data set is collected using a helicopter mounted instrument, which results in a characteristic sampling pattern.

This example shows the south-eastern most 10 m by 10 m of site A which contains 1369 observations.

Some striping due to flight lines is visible.

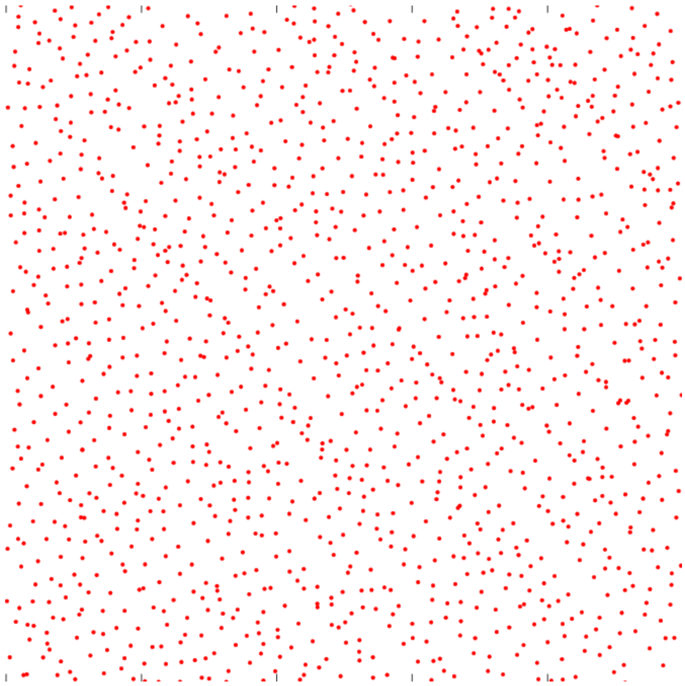

Data Characteristics: Variograms

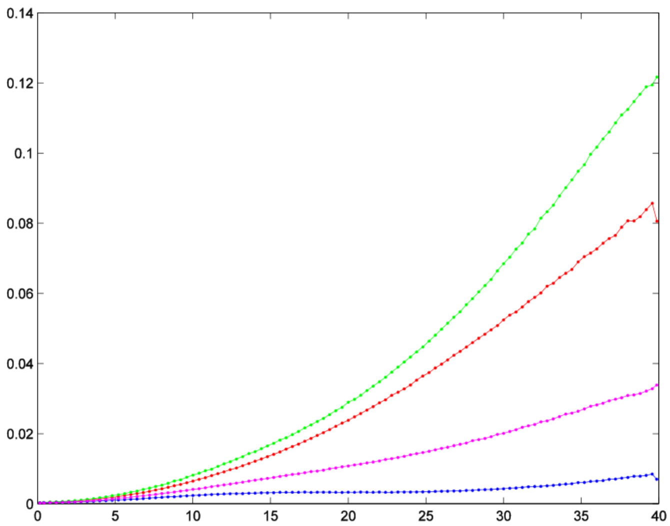

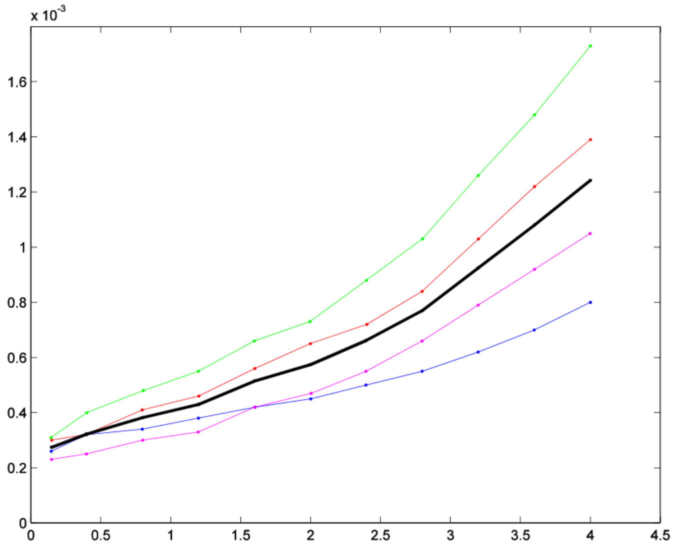

Left: Isotropic (omni-directional) semi-variograms for sites A (in red), B (in green), C (in blue) and D (in magenta), the lag separation value is 0.40 m.

For large lag values the four sites are very different, for shorter lag values they seem much more similar.

Right: Isotropic (omni-directional) semi-variograms for sites A (in red), B (in green), C (in blue) and D (in magenta), lag values up to 4.0 m. The average semi-variogram is shown in black.

The zoom on smaller lag values shows that the four sites mostly differ in the scale of the semi- variograms. *This scaling does not influence the resulting kriging weights in the interpolations below.*

For lag values up to 2.0 m we fit an isotropic (also known as omni-directional), linear model to the average semi-variogram. The resulting intercept is 0.0002529 m2 and the slope is 0.0001600 m.

With the sampling density used here, we will not use data points beyond 2.0 m in the interpolations below.

Leave one out (LOO) cross validation results for 5 different estimators

Site A: Bias and RMSE for measurements and LOO cross-validated estimates; 0.40 m lag separation distance. Number of observations is 24 687.

Method	Bias [mm]	RMSE [mm]
Ordinary Kriging	-0.187	18.47
Nearest Neighbour	-3.245	24.67
Local Means	-0.174	18.43
Inverse Distance Weight	-0.819	18.71
Inverse Square Distance Weight	-1.395	19.40

A closer look at kriging versus local means

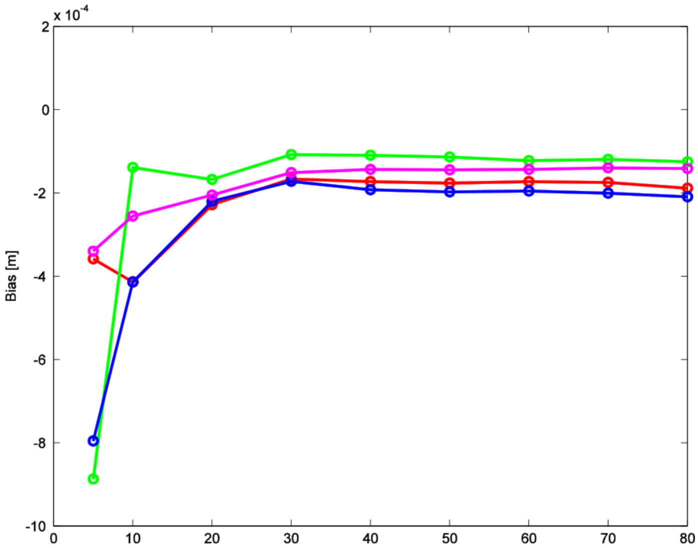

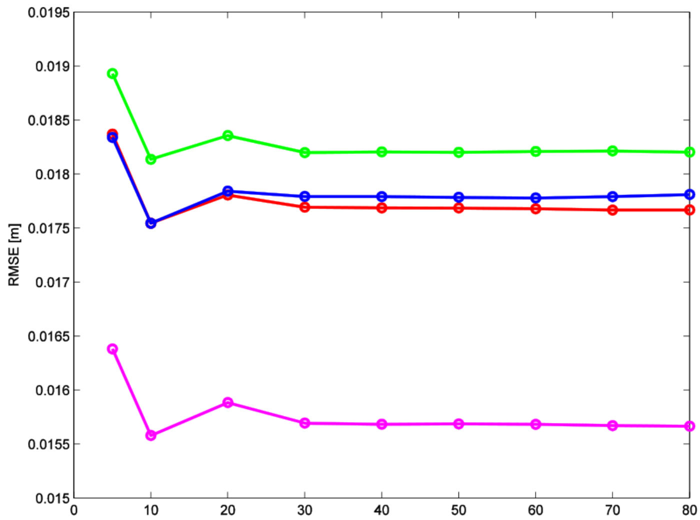

Left: Bias for OK, for sites A (in red), B (in green), C (in blue) and D (in magenta), as a function of number of nearest observations used by the estimator.

Right: RMSE for OK, for sites A (in red), B (in green), C (in blue) and D (in magenta).

We see that for OK the bias in all cases is negative, i.e., our interpolated values are systematically too low.

Also, for OK both bias and RMSE seem to stabilize at optimal values for around 30 nearest observations used in the interpolation.

This is due to the so-called screening effect in kriging for semi-variogram models with little nugget effect. (In all sites RMSE is smallest for 10 neighbours but here bias is mostly high.)

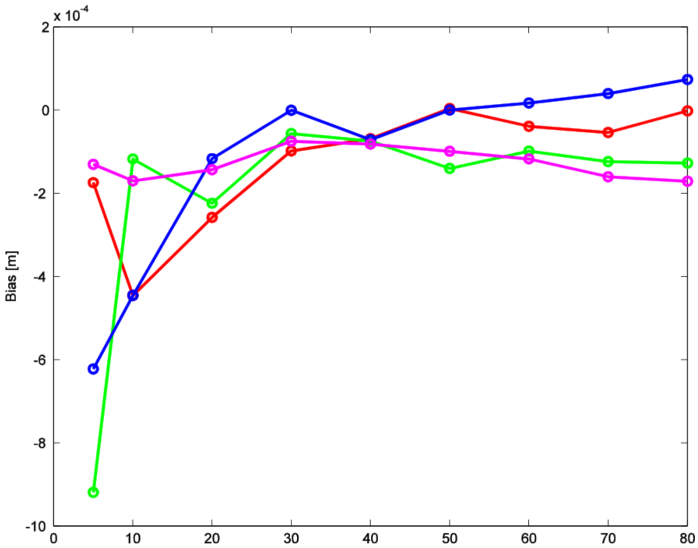

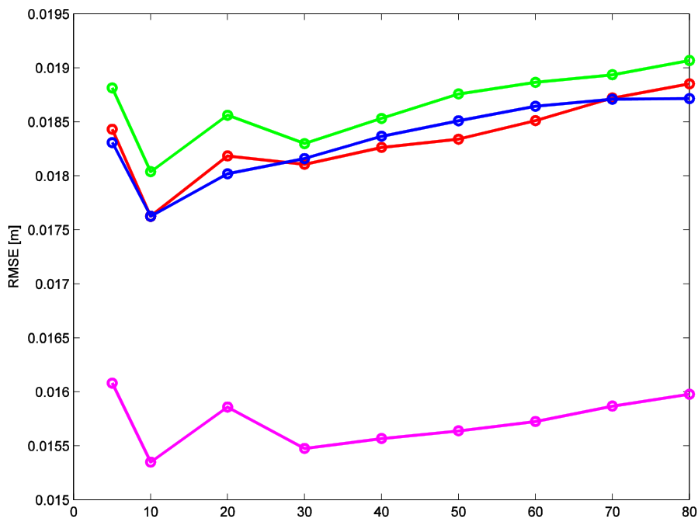

Left: Bias for LM, for sites A (in red), B (in green), C (in blue) and D (in magenta).

Right: RMSE for LM, for sites A (in red), B (in green), C (in blue) and D (in magenta).

As for OK, there is a tendency to negative bias, especially for a low number of neighbours (not pronounced for sites B and D). Also for LM 30 nearest observations seem optimal for these data.

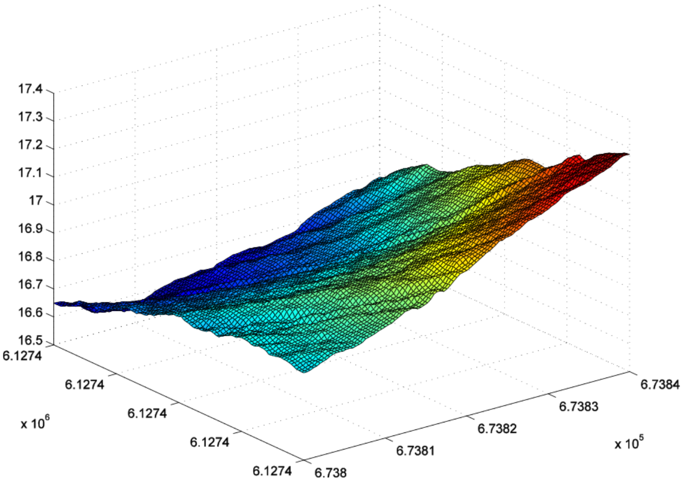

LM interpolated heights, site A. North is up and left.

Hence, runs with different numbers of neighbours used in the interpolation show that in this data material there is a tendency, that for a small number of neighbours used in the interpolation (here five to ten), *LM performs at least as good as OK* with a anisotropic double spherical semi-variogram model based on 0.10 m lag separation.

For a higher number of neighbours (here above 30), based on RMSE OK performs marginally better than LM, whereas based on bias, LM performs better. (This is an example of the well- known bias-variance trade-off.)

Outroduction

Currently we focus on simple approaches for creating a smooth update process for integration of heterogeneous data sets.

A basic example is to *ignore data that does not tell us anything new*.

In the figure below, we do this by looking at anomalies between the current height model and the new point cloud, then computing a *correction grid* for the current model. Points with insignificant anomalies are simply removed from the point cloud, and the correction grid is computed using the remaining point anomalies only.

Hence, we only compute updates in areas of significant change, speeding up the process, and giving us new insight of the precision of the current model which in turn results in improved metadata for both the current and the new model.

On the other hand, as years go by and multiple generations of data become available, more advanced approaches will probably become necessary (e.g. a multi campaign bundle adjustment, improving the oldest data using cross-over adjustment with newer campaigns).

But to prepare for such approaches, it is important already now to organize and evaluate the ancillary (GPS, INS) and engineering level data for the current data sets. This is essential if future generations of DEM users should be able to benefit from future conceptions of "some safe and sensible shortcuts for efficiently upscaled updates of existing elevation models".