Detection of ancient settlement mounds –

Archaeological survey based on the SRTM terrain model

B. H. Menze (1), J. A. Ur (2), A. G. Sherratt (2)

(1) Interdisciplinary Center for Scientific Computing (IWR), University of Heidelberg, GERMANY

(2) Department of Anthropology, Harvard University, Cambridge, MA USA

(3) Institute of Archaeology and Ashmolean Museum, University of Oxford, UK

Technical Report from

Multidimensional Image Processing, IWR, University of Heidelberg

ABSTRACT

In the present study we demonstrate the value of the SRTM 3 arc-second terrain model for a virtual survey of archaeological sites -- the detection and localization of ancient settlement mounds in the Near East. These so-called "tells" are the result of millennia of settlement activity within the period from 8000-1000 BC, and form visible landmarks of the world's first farming and urban communities. The SRTM model provides for the first time a systematic opportunity to identify and plot such sites. In order to map these cultural monuments for the purpose of settlement archaeology and conservation, we develop a machine learning classifier, which predicts probable tell sites from the terrain model. Point-like elevations of a characteristic tell shape, standing out for more than 5-6m in the DEM were successfully detected.

1. INTRODUCTION

The study of tell*s* is a fundamental category of archaeological research. Tells are settlement-mounds which are found in the Near and Middle East, in an arc from the Balkans to north-west India, and represent prehistoric and early historic villages and towns. These were typically occupied for long periods of time, often several millennia, during which the mud-based building technique caused building-debris to accumulate, and build up into a substantial mound, giving advantages of visibility and protection. These prominent landmarks of early human activity began to appear when the spread of farming in the Neolithic period (8000-6000BC) gave rise to permanent villages, and such sites continued through to the Bronze Age (3000-1000 BC), when some grew to the size of major urban centres. Although one or two are still inhabited, most of the mounds were abandoned two to three thousand years ago, and modern settlements exist on flat land nearby. These artificial mounds thus represent the remains of the earliest settlement systems, and a study of their spatial occurrence can reveal insights into the emergence, development and organization of the first complex human societies.

A comprehensive and accurate listing of these sites is thus a research priority, and has hitherto been achieved (for restricted areas) by survey on the ground. In consequence there is no overall picture of the distribution and relative density of these features, and the known pattern is largely a reflection of the differential intensity of investigation. Nor are the locations even of known sites well recorded, because of the unavailability of large-scale maps in these areas. Remarkably therefore, after more than 150 years of archaeological research, information about the locations even of major sites is notoriously imprecise, while most small mounds never made their way into the record. As a consequence, current compilations of published site positions are only available on coarse grids and contain no more than a subset of potential sites [Sherratt 04]. This is especially regrettable because the surviving population of ancient settlement-mounds is increasingly under threat: expansion of modern settlements, road-building, and the intensification of agricultural land-use has brought an unprecedented rate of destruction to these historical monuments. Not only are these sites of academic interest in the study of ancient settlement patterns, they are themselves valuable aspects of the cultural heritage. At the present time, they are probably disappearing faster than they are being recorded, and this is undoubtedly true of the smaller examples.

This unsatisfactory situation was transformed in 2004 by the publication of SRTM data. For the first time it became possible to observe topographic phenomena at the scale of tell settlements, which as artificial mounds stand out as well-defined anomalies from the kinds of flat lowland landscapes in which they are typically situated. The aim of the present work is to develop and assess a (semi-) automatic tell detection strategy, which is based on SRTM digital elevation models and will allow a virtual survey of ancient settlement mounds over wide geographic areas, under objective search criteria, and at low operational costs. This is already beginning to provide a representative list of sites for detailed assessment and ground-truthing.

In the following we will at first discuss possible approaches in the remote sensing of these archaeological sites (section 2). Then a machine learning-algorithm is proposed, which is trained and applied to the SRTM data of a test region (section 3.1). This algorithm is able to identify the typical conical tell pattern of a settlement-mound within the natural topographic variation (section 3.2), and will be used in a (semi-)automated survey for tell sites. The detection accuracy will be discussed (section 5) in relation to of the size and height of the tell mounds, as they derived from SRTM data and other sources of information (section 4).

2. REMOTE SENSING OF SETTLEMENT MOUNDS

From a simple physical point of view, tells are features of 5-50m height, 50-500m diameter and usually of conical shape. Also, they primarily consist of loam and mud-based materials (fig.4). Both features might be used in the identification of tell sites: Hyperspectral imagery, e.g. from LANDSAT, is a standard tool in the classification of soil types and ground cover. It is potentially helpful to identify the often uncovered and eroding tell sites. Digital elevation models reveal shape information. Their potential usefulness in the search for tells was identified by one of the authors shortly after the data was released [Sherratt 04]. Data for both approaches are available with high spatial resolution and wide coverage. Unfortunately, the spectral signature of known tell sites has so far proved to be too unspecific to serve as a diagnostic characteristic in an automated classification. Thus, the detection of tell sites falls back on an optimal processing of the DEM data, with the supplementary use of high-resolution satellite imagery and other georeferenced information.

The interesting regions of the Near East are covered by a number of digital terrain models, although most of them (e.g. GTOPO30, GLOBE) do not satisfy the required spatial resolution. In addition to the (3 arc-sec) SRTM data, only the ASTER-derived DEM has the potential to map elevations of tell size. However, a qualitative comparison of these two models has already demonstrates the major disadvantages of the latter (fig. 3,4). In areas with high gradients - mountain walls and ridges, rivers banks, but also some of the major tell sites - ASTER data do indeed reveal unique details at its maximum (30m) resolution. Unfortunately, artifacts in size and shape similar to a settlement mound characterize the ASTER DEM (fig.3). Comparison with LANDSAT images suggests a dependence of this effect on the type of agricultural land-use (fig.4). As the artifacts are primarily a problem in the flat plains - where the tells are expected - a use of this DEM in an automated or even visual search for tell sites is not too promising. In consequence, we concentrate in the following on the SRTM model and its use in digital survey for ancient settlement mounds.

3. SURVEYING THE DEM

Region Under Study

The data used in this study are from a test-area in the north of Mesopotamia. The upper Khabur catchment has a long settlement history, and saw the major expansion of nucleated settlements in the fourth millennium BC. It is one of the regions where archaeological settlement surveys were developed and which is still a focus of current research [Ur 04].

The basin lies mainly in Syria, and in adjacent parts of Turkey and Iraq (fig.1), measures roughly 320km from east to west, and 120km from north to south.

The Khabur basin is covered by six SRTM one-degree-tiles (36-38 N, 38-41E) at 3 arc-second resolution (90m) (fig.2). DEMs from several ASTER swaths were also available for the region. Landsat-7 ETM maps at 15m resolution were also used and, for the test region, declassified monochrome CORONA images at 2-6m resolution. Topographic maps derived from photogrammetric survey and produced at 1:5000 resolution were available for central parts of the test region to provide independent height information [Ashmolean], and 1:100000 topographic maps for the whole basin were used in the confirmation of tell sites [Berkeley].

Tell Detection

The algorithm for a machine-guided tell search was designed on data from a subset of five SRTM tiles and evaluated on ground truth for the sixth, southwestern tile (fig.2).

For training, 184 tell sites were identified visually in the SRTM data and were confirmed with help of LANDSAT images and topographic maps. These sites, together with 50 000 randomly chosen positions, served as training and validation data in the design of the classifier. As input for the training, the elevation data from circular regions of 1km diameter, centered around the training sites, was used. The height differences relative to the center in this image patch were concatenated to an 80 dimensional parameter vector.

The test data was acquired in an independent prior study [Ur 03]. CORONA images of the 1960s revealed the location of 133 sites with indications of settlement activity. The size of these sites ranges from one to 60ha (Tell Brak, fig. 5b) in area and less than 5m to more than 50m in height.

In the DEM, the tell pattern is superimposed on natural topographic variation. Though the geographic region under study is a relatively flat plain, natural variation of the land surface can be observed on various scales, ranging from slowly varying slopes to steep canyon walls (see fig. 11). In addition to this ‘background signal’, small conical hills of volcanic origin, more recent artificial accumulations and structured artifacts of the DEM are likely to be misclassified as human settlement mounds.

The standard approach for the detection of a known pattern in a varying (additive white gaussian) background signal is a matched filter. However, elevation data of a physical land surface are highly correlated and the ‘noise’ of non-tell surface-structures is only poorly approximated by the implied gaussian error model. So a direct application of a matched or template filter in the detection of characteristic tell like mounds results in an unacceptable number of false hits (fig. 11). Alternatively, ideas from high-level image analysis, namely from face recognition, can help to overcome the problem and to allow a reliable and automated screening of SRTM tiles for point like elevations (fig.6). A set of linear filters is learned from the training data to span a subspace for a multivariate decision rule. The purpose of these linear filters is a fast processing and low-dimensional mapping of the (80 dimensional) local image pattern in the first step, while the following multivariate classifier provides the desired adaptivity of the detection algorithm.

Partial least squares (PLS) filters were found to be superior to other linear methods (linear discriminant analysis, principal component analysis) [Menze 05a]. In combination with a subsequent classification by randomForest [Breiman] on a eight dimensional PLS subspace, a sensitivity of 95.4% could be reached at a specificity of 98.8% in a ten-fold-cross validation of the training error.

Applied to new data, the classification algorithm is able to provide ranked lists of positions with decreasing ‘settlement mound probability’ (fig.6). Although the specificity of these lists is high, conical elevations are not neccessarily due to settlement activity. Therefore other georeferenced modalities, such as digitized topographic maps or specific ground cover information, can be used either to separate elevations of different origin or to confirm the tell position and to localize the extensions of the sites in highly resolved satellite imagery in a final step.

4. RESULTS

Beside the detection of ancient settlements, the processing of a DEM also allows for the determination of a physical site parameter, which is new in the description of tell distributions: the height of the settlement mound above the surrounding area at the present time (though the actual mound may be deeper than the present land-surface, due to aggradation during or after occupation). On our test data, it was assessed as follows: A linear plane was fit repeatedly (20 times) by least-squares onto varying subsets (2/3) of the neighbouring SRTM pixel of each tell site, in order to serve as an approximation for the base of the tell. The tell height was estimated by the maximum difference between tell surface and ground plane. We find that the smallest sites are not higher than 2m in the DEM, while the biggest settlement mound rises as high as 30m above the surrounding area (fig. 9). While the variation of the height is often well below 1m - even for major tells - significant errors can be observed for small sites, which hardly stand out from the ground or on settlement mounds (fig. 8, 9).

For a subset of the tell sites the height from base to top could be determined from topographic maps (fig.8,12). A first order approximation between SRTM elevation and the base-to-top height of a settlement mound suggests that the latter is underestimated by a factor of two. Nevertheless, comparing the spatial average of the 3-arc-sec SRTM model with the extensions of a tell's peak, the base-to-top height should be considered a subpixel quantity and this results had to be expected.

When applying the classification algorithm to the SRTM test tile, it is possible to detect 85 out of the 133 test sites at a threshold, which results in 327 false positives for the 600*1200 pixels of the test region (northern half of the test tile). False positives are mostly due to natural elevations resembling tells in height and size, which occur frequently in the undulating slopes of the "jebel" in the southern half of the test tile, or are due to easily distinguishable artifacts in the presence of water surfaces. Obviously the first of these error sources sets natural limits to the presented application.

In order to understand the false negatives – the undetected tell sites - a more detailed view of the characteristic tell shape is helpful: In the earlier study of the training set, the spatial extensions of the tell sites had been outlined on CORONA imagery at a spatial resolution of 2-6m. In a comparison between this tell area and the height from SRTM data, a strong relation can be observed (fig. 9): diameter and height of the tells vary linearly (area = c* height^2), indicating a characteristic shape of a settlement mound over a wide range of size [Menze 05b]. Likely to be missed by the classifier are a low number of major sites, which do not show the typical height/area ratio. An extension of the training data set might alleviate this (current) drawback.

Overall, beyond features of the classifier and properties of the training set, a sharp increase in the detection probability can be noted for sites with a height of more than 5-6m in the SRTM data (fig. 9,10). This coincides very well with the limits of the SRTM data accuracy, though the base-to-top height of these sites might be somewhat higher both in real and at a subpixel resolution.

5. DISCUSSION

We envisage a programme of archaeological “virtual survey” for settlement mounds over a large part of the Near East, in the heartlands of the ancient civilisations and their neighbours. The methodological basis of this search is the automated processing of SRTM data by the point-detection algorithm. The spatial and vertical accuracy of the 3 arc-second SRTM data is sufficient to identify mounds standing out more than 5m from the surrounding at the resolution of the terrain model.

Currently the algorithm is being applied both to additional parts of the Near East where survey data offer a comparison with ground-derived records, and to areas not so far systematically subjected to ground-survey in Turkey, Syria, Iran, and Iraq, in order to compile a basic register of sites These investigations are yielding new settlement-sites, and providing precise positional information for ones already known. In addition, these listings will provide information of such physical parameters as height and area, as well as linking them to a growing body of increasingly high-resolution satellite imagery. These objectives would, of course, be greatly assisted by release of the 1 arc-second data, which would notably enhance the resolution of the model. Nevertheless the 3 arc-second SRTM data have already made a contribution in this field which is little short of revolutionary, and will be of major importance in studying the history of settlement in this important region, and in helping to conserve its remains.

REFERENCES

[Sherratt 04] A. Sherratt, Spotting Tells from Space, Antiquity, 2004

[Ur 03] J. Ur, CORONA Satellite Photography and Ancient Road Networks: A Northern Mesopotamian Case Study, Antiquity, 2003

[Menze 05a] B. Menze et al., Filtering linear subspaces for nonlinear point detection, technical report, University of Heidelberg, 2005

[Ashmolean] 1:5000 Photogrammetric survey of the Hasake region, (Copies in Ashmolean Museum, Oxford)

[Ur 04] Urbanism and society in the third millenium Upper Khabur basin, PhD thesis, University of Chicago, 2004

[Berkeley] Soviet topographic maps, UC Berkeley map collection

[Menze 05b] B. Menze et. al., Tell Spotting - surveying Near Eastern settlement mounds from space, Proceedings of the CIPA International Symposium, Torino, 2005

For additional material see *http://ArchAtlas.org* .

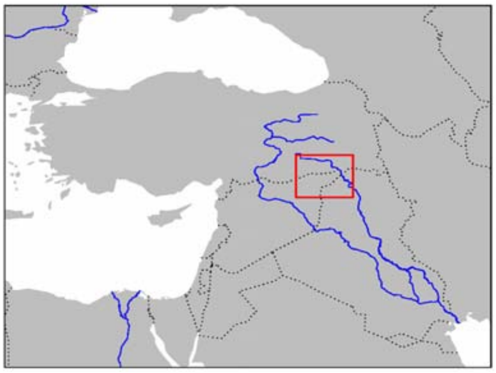

Fig. 1: Map of the region under study. The Khabur basin is in the north of Mesopotamia (Euphrates and Tigris indicated).

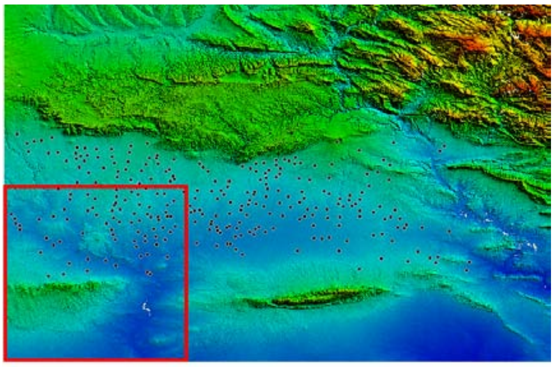

Fig. 2: Digital elevation model of the Khabur basin with settlement mount positions indicated (black dots). Tells of the southwestern SRTM tile (red bordered) serve as test sites.

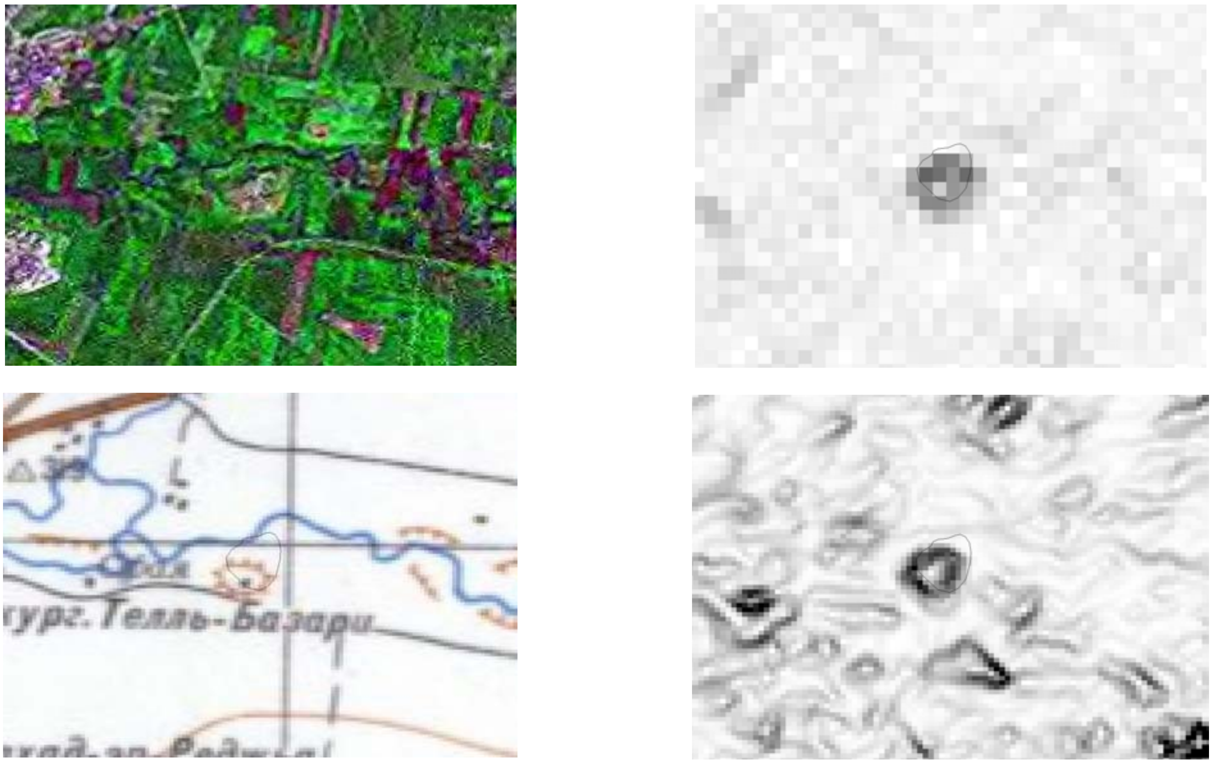

Fig. 2: Landsat ETM+ imagery (up left) was used for the confirmation of sites detected in the SRTM [up right, gradient image]. Topographic maps [below left] provide (modern) names of the sites and objective 'ground truth'. – Comparing ASTER DEM [below right] and SRTM data, the three times higher resolution of the first reveals detail of the site, which are in line with the Landsat observations. Artifacts in the ASTER model are often indistinguishable from elevations of tell size. -- Outlined are site extensions as assessed from CORONA imager, also compare fig. 11.

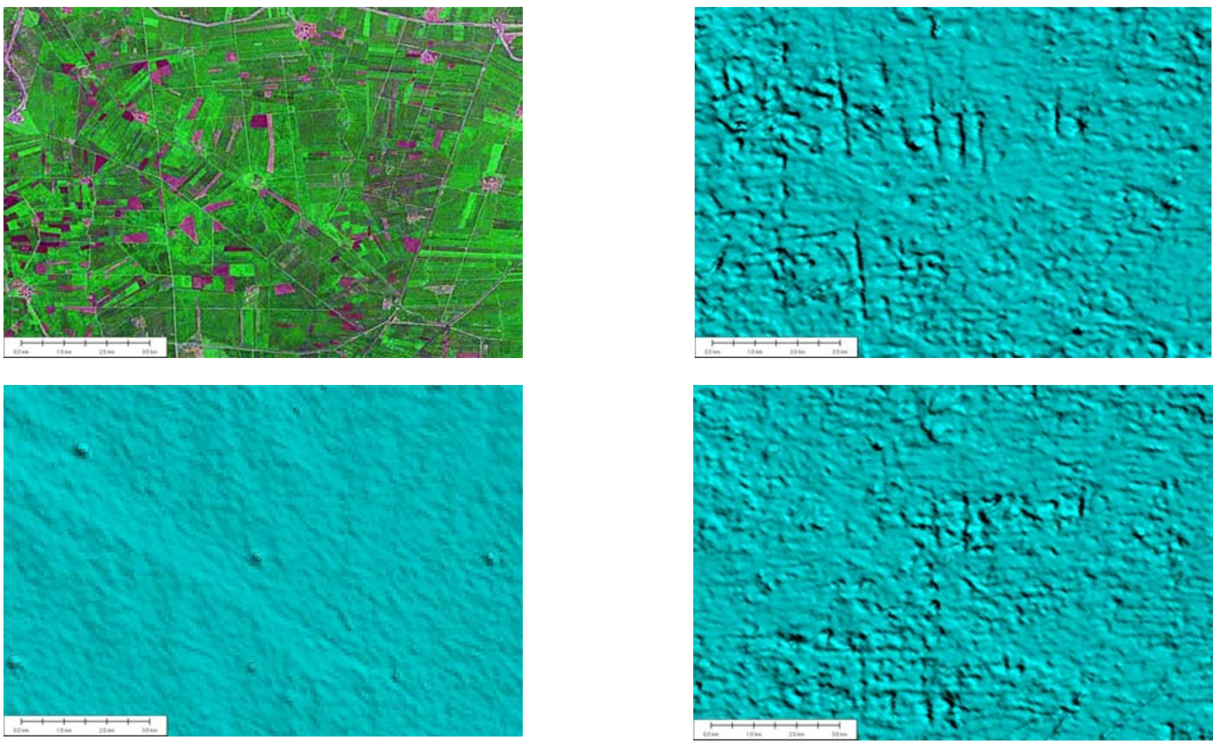

Fig. 3: View onto a central part of the test region. Tells can be easily spotted from the SRTM DEM [below left]. The Landsat view [up left] reveals the close vicinity between moden settlements and ancient settlement mounds. Although several ASTER Swaths [up& below right] are available for the generation of elevation models, they all suffer both from small scale (see Fig. 2) and extended artifacts, which often correlate with the type of ground cover, such as streets, channels and even crop fields. – Scale bar in the lower left image corner indicates 3km.

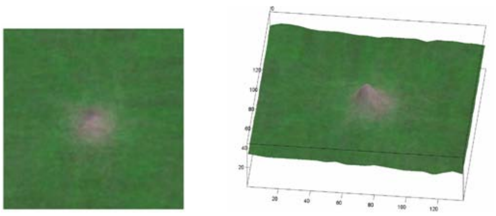

Fig. 3a: Landsat image and elevation model averaged over all sites of the test set. The majority of the tells is in direct vicinity to modern settlement, which is visible from the bright central spot.

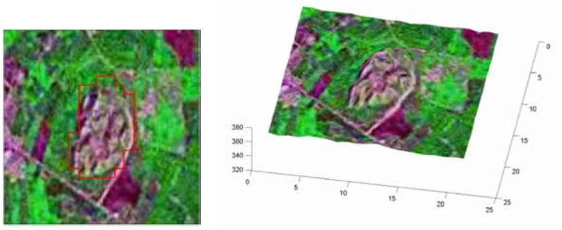

Fig. 3b: Tell Brak, major settlement site in the Khabur basin

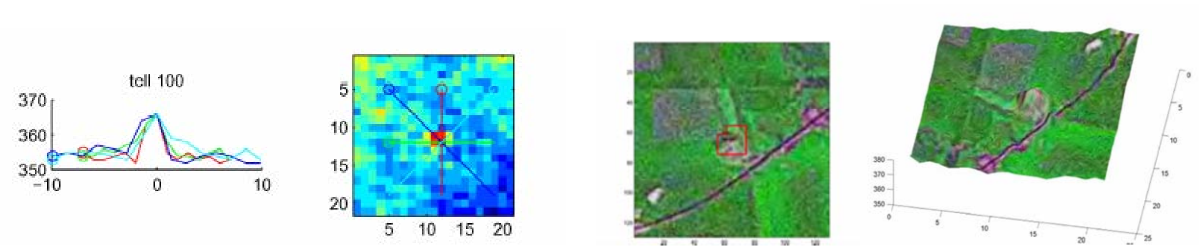

Fig. 4: Principle of point detection: The characteristic height profile of a tell site generates a typical point like pattern in the DEM. Under the resolution of the SRTM DEM, their elevation is often of nearly conical shape (height profiles left, as indicated in the DEM image patch).

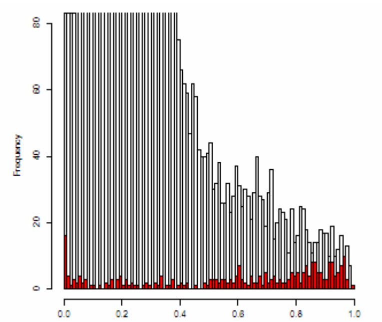

Fig. 6: Histogram of the cross-validated classification results on the training data. While the majority of the 50 000 'non-tell' pixels (white) are assigned to low probability values (frequency truncated at 80), most of the tell-pixels (red) are gathered at high values.

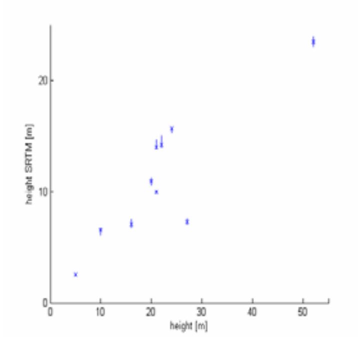

Fig. 7: Tell heights of the SRTM model tend to underestimate the real base-to-top height of the mound (from topographic maps) systematically.

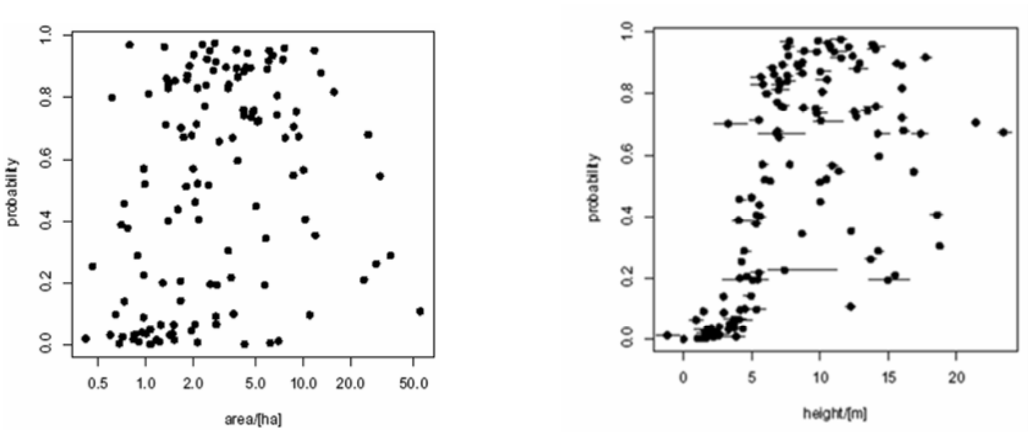

Fig. 8: The detection probabiliy of the tell sites in the test region increases for area sizes bigger than 1-2ha (approx 1-2 SRTM pixels) [left]. The elevation of a tell is the more important characteristic [right]: A 'transition' in the detection probability can be observed for mounds standing out higher than 5-6m from the ground (also see the discussion in the text).

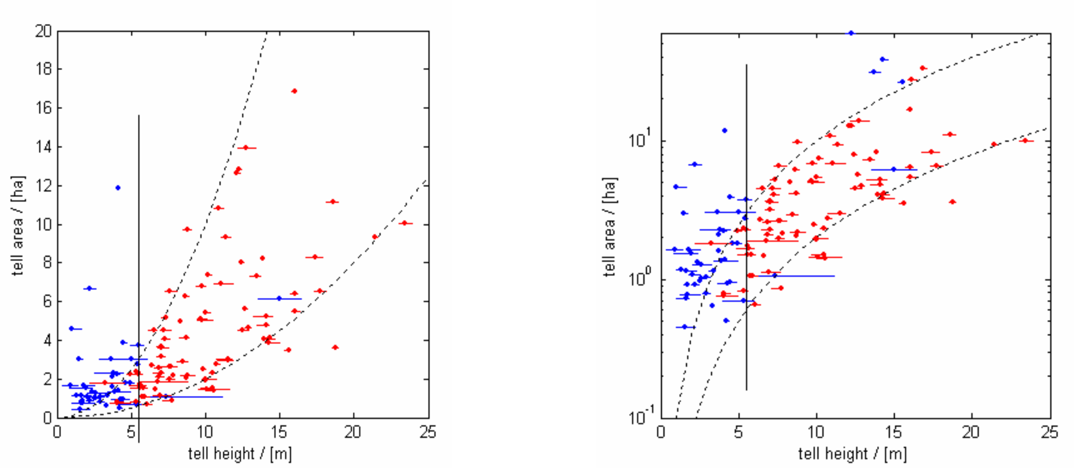

Fig. 9: Out of the 133 tell in the test set 85 could be detected (red points) at the chosen threshold of 0.3 (rf. fig. 8). The plot of tell area vs. tell height (left: linear area-scale, right: log area-scale) reveals that detected sites are typically higher than 5-6m (black line) and follow a quadratic area-height relation (dotted lines, see text).

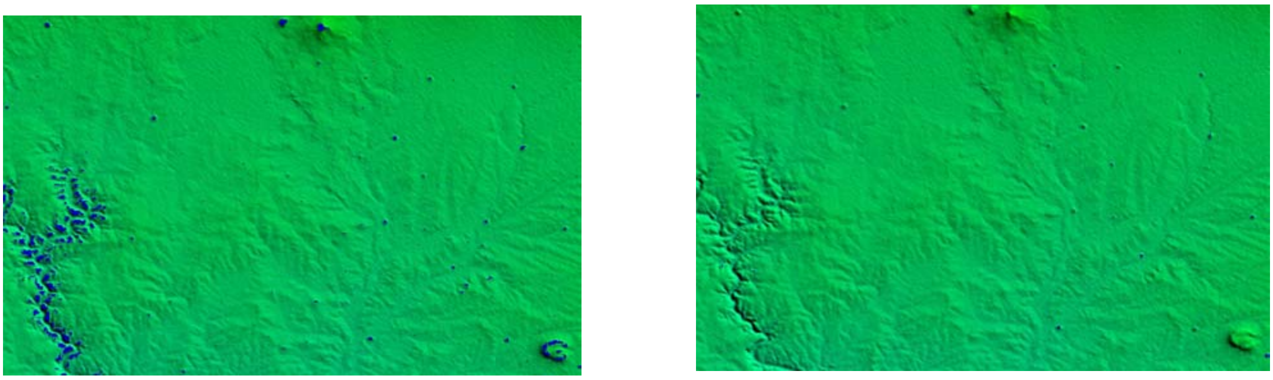

Fig. 10: While a standard matched filter (left image) marks a considerable number of false positive hits (see blue labeled pixels), preventing any computer-aided detection of settlement mounds, the designed classifier (right image) flags only a limited number of pixels with high specifity, allowing a time-efficient localization in high resolution LANDSAT images in a subsequent step.

Fig. 11: Topographic maps at 1:5000 resolution were used to determine ground truth in the height estimation. Shown is Tell Bazari, compare fig. 2.