UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Architettura, Urbanistica e Rilevamento

SCUOLA DI DOTTORATO DI RICERCA
IN
STUDIO E CONSERVAZIONE DEI BENI ARCHEOLOGICI E ARCHITETTONICI

INDIRIZZO
IN
SCIENZE E TECNOLOGIE PER I BENI ARCHEOLOGICI E ARCHITETTONICI

CICLO XXI

ADVANCED METHODOLOGIES FOR ACQUISITION, INTEGRATION, ANALYSIS, MANAGEMENT, VISUALISATION AND DISTRIBUTION OF DATA IN THE FRAMEWORK OF ARCHAEOLOGICAL AND ARCHITECTONICAL HERITAGE

Direttore della Scuola: Prof. Giovanni Leonardi

Supervisori: Prof. Vladimiro Achilli

Prof. Thomas H. Kolbe

Dottorando: Giorgio Agugiaro

31 luglio 2009

In order to apply for the Doctor Europaeus Ph.D. label, this thesis was refereed by:
Prof. Norbert Pfeifer, Technische Universität Wien, Austria
Prof. Javier Finat, Universidad de Valladolid, Spain

Revised version, October 2009

Problems increase in geometric ratio,
solutions in arithmetic ratio.
ISSAWI-WILCOX PRINCIPLE

Abstract

Use of high resolution sensors like laser scanners is becoming more and more common for surveying tasks. Nearly every single application field is gaining benefits from these booming methodologies. Production of high detailed models is becoming part of the standard workflow in many disciplines – architecture, cultural heritage, urban planning, etc.
At the same time, new related problems have emerged concerning the management of the resulting large datasets and the integration of data coming from different sources.
Reasons for integration inconsistencies are multiple: there could be geometric, topological or semantic incompatibilities, or “similar” datasets could originate from different sensors and have different resolutions and accuracies. Finally, data could have been acquired at different times. Of course, a combination of the above mentioned causes is also possible.
In this thesis a new deterministic approach is presented, which allows for integration between a laser-scanner-acquired, high resolution model and a lower resolution digital terrain model by means of a transition surface.
Using extra data around the laser-scanner-acquired model, which represents a sort of “collar”, the goal is to create a transition surface between the two models, which connects them smoothly without modifications to the actual high resolution object and permits a transition also in terms of point density.
Some experimental results, obtained from real datasets coming from archaeological and cultural heritage sources, are presented. They show that the developed approach is suitable to integrate two datasets although little information about their quality and accuracy is known *a priori*.
Open issues and possible improvements are finally discussed.

Riassunto

La possibilità di poter acquisire velocemente e a costi sempre inferiori grandi quantità di dati geometrico-geografici e, di conseguenza, di disporre di modelli ad alto livello di dettaglio sta caratterizzando sempre più il "normale" flusso di lavoro di ogni disciplina (architettura, archeologia, urbanistica, ecc.) grazie alla diffusione sempre più capillare, in anni recenti, di nuove metodologie di rilievo ad alta precisione come il GPS e la scansione laser aerea o terrestre.
Di solito, se un'applicazione è sviluppata specificatamente per un certo problema, può essere sufficiente servirsi di un solo modello ad una determinata risoluzione. Tuttavia, se l'analisi e la modellazione di un fenomeno richiedono l'uso di molteplici set di dati, è necessario definire dei criteri che ne permettano l'integrazione.
Integrando due o più modelli diversi si possono infatti verificare errori di vario genere, dovuti ad una vasta gamma di fattori quali, per esempio, incompatibilità di tipo geometrico, topologico o semantico. Altri errori si hanno talvolta con dati "simili", ma provenienti da sensori diversi e pertanto con precisione e risoluzione diverse. Può accadere di dover trattare dati acquisiti in periodi di tempo diversi, o – nel caso più generale – le ragioni che portano ad incompatibilità sono una combinazione dei suddetti ed altri problemi ancora.
In questa tesi viene descritta una metodologia che permette l'integrazione di un modello ad alta risoluzione in un modello digitale del terreno a risoluzione inferiore per mezzo di un'opportuna superficie di transizione.
La superficie ottenuta permette non solo di congiungere i due modelli, ma garantisce anche una transizione graduale tra alta e bassa risoluzione. Le caratteristiche di entrambi i modelli quali la geometria, la topologia, la densità di informazioni sono messe in relazioni evitando brusche discontinuità.
L'idea alla base della metodologia proposta consiste nell'utilizzare dati "aggiuntivi" attorno all'oggetto rappresentato nel modello ad alta risoluzione. Tali dati sono comunemente presenti nei modelli acquisiti con laser scanner e, anziché venire eliminati – come di solito avviene – nella fase di editing, sono utilizzati per modellare la superficie di transizione.
I risultati dei test eseguiti su dati sperimentali dimostrano che l'approccio deterministico proposto si presta all'integrazione di dati di cui non sono note informazioni *a priori* sulla qualità o sulla precisione.

Index

Chapter 1

INTRODUCTION

1.1 Motivation and objectives of the work

Use of high resolution sensors is becoming more and more common for surveying tasks, while nearly every single application field is gaining benefits from these booming methodologies.

As geodata are becoming more and more ubiquitous and applications based on it are literally flourishing in a wide array of fields, from the most specialised to everyday life ones, it can be assumed, in very general terms, that the whole terrestrial globe has been mapped so far – even if at different levels of accuracy.

Moreover, the continuously increasing availability of multiple models of the same "object" at different resolutions or acquired from different sensors leads to an array of still unsolved problems concerning their integration in terms of precision, data structure, or simply geographical extension.

Projects like Google Earth are starting to provide world-wide coverage with multi-scale and multi-format geodata, but are still facing problems regarding, for example, embedding of high resolution objects (landmarks, city models, etc.) in a generally low resolution Digital Terrain Model (DTM).

The framework of cultural heritage, in particular the archaeological and architectural one, has experienced a similar technological leap in the past decade. The possibilities offered by fast and high resolution devices like laser scanners or digital cameras enable to create complete and extremely detailed three-dimensional models which have proven suitable for multiple purposes: they range from specific ones like detailed documentation and record of changes for future restoration, to more general ones like education or *digital* tourism.

For simple and fast visualisation, a simplified and coarse model may suffice, but highly detailed representations are required for architectural analyses, conservation studies, and close-up visualisation.

The opportunity to have a consistent representation plays therefore a major role in this application field, which necessitates of large amounts of data to be acquired using different techniques and heterogeneous sensors. In addition to the aforementioned variety of scale (it may vary from a small archaeological find to the entire historical centre of a medieval village) the high resolution surveying of archaeological and architectural objects is characterised by the peculiar complexity of their geometries: data acquisition may be hindered due to unreachable locations, occlusions or particularly detailed decorations – just to name some.

If, on one hand, this has led to a profitable broad experimentation of different approaches tailored at specific aspects, on the other hand the drawback is that no established universal procedure is currently available for all surveying, integration,

management, visualisation, and distribution steps of this type of models. Nevertheless, in the course of years, a wealth of literature has already been published, where good practices and solutions to specific issues are presented.
An overview for 3D reconstruction from images can be found for example in REMONDINO *et al.* (2008), some issues and possible solutions related to laser scanning applied to cultural heritage are described in RUTHER (2007), while FUENTES *et al.* (2006) and FINAT *et al.* (2005) deal with data acquisition and fusion applied to urban spaces.
After data have been collected, a large spectrum of methods is necessary to seamlessly combine the different input models: overlapping zones must be removed, gaps where information may be missing are to be filled if one unique, integrated model is to be created.
Another important aspect is tied with spatial resolution. According to the desired level of detail, the type of employed devices and the absolute size of the real surveyed object, several strategies must be employed. In general, the final model will have different resolutions according to the complexity of its features.
Another parameter influencing the final model complexity is the level of realism in the visualisation: details should then be added accordingly. This is an open problem which can be actually found at any scale. For example, some particularly elaborated details of an architectural object or archaeological site may have been acquired singularly and need to be properly merged into a more general model.
It is indeed not so uncommon to create “holes” in a low resolution model and to replace the deleted data with selected details acquired at a higher resolution. A set of common rules can be summarised when applying this strategy: a) there should be a perfect connection between geometric primitives and the surfaces along the adjacent border of the hole, b) gaps, redundant surfaces or intersecting features should be avoided, in order that c) the overlying textures do not suffer from irregular geometries and d) the textures of adjacent models should integrate seamlessly [REMONDINO *et al.*, 2009].
A final step may consist in the visualisation of the generated three-dimensional model into the surrounding landscape. Digital terrain models and ortophotos at varying resolutions are used. Of course, if the heterogeneous input models vary significantly in accuracy and resolution, and the resulting discrepancies must not be underestimated.
Research work is therefore still being carried out in order to develop efficient methodologies aiming at a perfect combination of different three-dimensional models. However, a definitive solution has not been reached yet, a rather representative example is shown in Figure 1.1.
In the framework of data integration, Figure 1.1 represents also a good example of the main topic of this thesis: as soon as two or more heterogeneous models need to be combined, the task may not reveal itself as an easy one due to several possible issues concerning geometric, topologic or semantic inconsistencies. The reasons for such inconsistencies might be the underlying geometry structure (a mesh or a point cloud, a boundary-representation or Constructive Solid Geometry model), or the related topology (2D, 2.5D, 2.8D, 3D etc.), or simply the diachronic origin of datasets representing the same object.

Figure 1.1: Integration of heterogeneous models may not be flawless: the model of the Colosseum sinks into the surrounding terrain model.

More specifically, according to LAURINI (1998) errors resulting from the merge of different spatial datasets can be grouped into two main groups: *layer fragmentation* errors, which originate in case of datasets covering the same region but containing different feature classes (e.g. a DTM and a "flying" 3D building); *zonal fragmentation* errors, which refer instead to datasets containing the same feature class but covering spatially disjoint regions (e.g. overlaps or gaps at the borders of two adjacent DTMs).
This thesis focuses specifically on zonal fragmentation. Geometric and topological issues are considered with regards to data integration between a laser-scanner-acquired model and a lower resolution DTM. Attention is paid to the overlapping zones: a prerequisite is that the high resolution model has been acquired with some additional data around it – a common situation in laser scanner point clouds –, so that one can think of a sort of "collar" around the scanned object.
Using the collar data, the aim is to create a transition surface between the two models, which connects them without modifications at the actual high resolution object and which permits a transition also in terms of point density.
Experimental results coming from real datasets are presented, too. They cover distinct application cases, especially in the field of archaeological and architectural heritage.

1.2 Structure of the work

In chapter 2 a general overview about terrain modelling is given. Different approaches used for the representation of terrain surfaces are shortly described; their main characteristics, their strengths and weak points are presented. Today's state-of-the-art terrain models can be considered 2.5D, which means that only one height value is associated to a couple of planar coordinates. However, it is present subject of research how to overcome this intrinsic limit. Some recent approaches are described, which deal with the problem of how to extend the 2.5D limit.

Another important aspect related to terrain models is the definition of parameters which can be used to assess the accuracy of a model. Thus, the second part of the chapter concentrates on quality parameters and draws attention to the importance to use different approaches to quantify the quality of a terrain model, both globally and locally.
Chapter 3 deals with issues concerning the integration of spatial data. Some motivations are presented, followed by some definitions which will be used throughout the thesis. A general overview is given about integration problems in Information Technology, then the specific issues, originating when heterogeneous spatial data are combined, are brought into focus. Different solutions to specific problems are shown, with particular attention to the geometric homogenisation process between two or more overlapping datasets. Finally, some still unresolved issues and the research subject of this thesis, i.e. the definition of a transition surface, are introduced.
In chapter 4 a methodology is described to obtain a transition surface which connects a high resolution, laser-scanner-acquired model to a lower resolution digital terrain model. The transition surface must guarantee continuity in terms of geometry, topology and point density. The initial assumptions and working hypotheses are introduced, then the desired characteristics of the transition surface are discussed. Finally, the steps needed for its calculation are described.
In chapter 5 the developed methodology is applied to integrate data coming from two archaeological sites and to a small bas-relief. Results are presented, while in chapter 6 they are discussed together with the open issues and the possible further improvements of the current methodology.
Chapter 7 contains the concluding remarks.

Chapter 2

DIGITAL TERRAIN MODELS

In this chapter a general overview about terrain modelling is given. Different approaches used for the representation of the terrain surface are shortly described, their main characteristics, their strengths and weak points are presented. Some recent approaches which deal with the problem of extending the 2.5D limit to so-called 2.8D or 3D will be presented. Another important aspect related to terrain models is the definition of parameters which can be used to assess the accuracy of a model. Thus, the second part of the chapter concentrates on quality parameters and draws the attention to the importance of using different approaches to quantify the quality of a terrain model, both globally and locally.

2.1 Terrain surfaces

A Digital Elevation Model (DEM) is a digital representation of ground surface topography or terrain. In general, DEMs are obtained from observations of the Earth's surface (or portions of it) and represent the bare terrain at some level of detail, in contrast to Digital Surface Models (DSM), which may also include features like buildings, vegetation and roads, as well as the natural terrain.

DEMs are used in numerous disciplines, ranging from engineering to geodesy, from architecture to geoinformation sciences – just to name some. DEM-based applications cover a wide array of cases, like terrain/surface analyses, creation of relief maps, modelling of water flow or mass movements, visualisation of 3D scenes, along with any process where decision making is involved.

Digital elevation models are obtained in many ways, depending on the technology and the methodology employed and, of course, the extents and the resolution of the surface to be modelled. However, they are frequently obtained by remote sensing rather than direct survey. Stereophotogrammetry consists in analysing and correlating two images (e.g. of the same terrain portion), which have been acquired from different angles, like from an airplane or a satellite.

In recent years, improvements in laser scanning have found an increasing adoption in the surveying disciplines, since laser scanners are efficient and economical measuring devices which allow for data acquisition ranging from cultural heritage objects to large area DEMs, with a previously unknown level of detail. Airborne Laser Scanning (ALS) systems can quickly acquire millions of points over the scanned surface and produce a highly detailed terrain model.

Other methods of DEM creation may involve interpolation of datasets directly sampled on the terrain, e.g. by means of GPS surveying, or from contour line maps previously produced, although the latter is a somehow “older” method which is used today only in case other methods fail, e.g. in mountain regions or due to scarcity of acquired data.

In general, DEMs are also widely known as Digital Terrain Models (DTM), but since a DEM implies that elevation is available continuously at each location in the study area, contour line data or any other irregularly sampled elevation datasets are not, strictly speaking, DEMs, but may be considered digital terrain models. Anyway, for the sake of simplicity, only the term DTM will be used from now on.
A quick review of the most common terrain models is given. Different approaches exist, their characteristics and their differences will be mentioned, as well as their limitations.

2.2 DTM types and characteristics

As mentioned before, the goal of a digital terrain model is the proper representation of the bare terrain surface, where features like trees or buildings have been sorted out in some way. The most common types of models used nowadays describe the height as a function of the location. This means that a unique value z_i is given for every planimetric point of coordinates $P_i : (x_i, y_i)$. In mathematical terms, this corresponds to defining the surface as a graph of a bivariate function over a compact domain, and this domain is generally the xy-plane. This approach generates what is generally called a 2.5D model, and its name hints at the same time at its major drawback: 2.5D is obviously less than 3D! Since for every planimetric location only one height value is possible, shapes like overhangs or tunnels cannot be modelled, since this would require a function that is not bijective (e.g. in a tunnel more height values should be mapped from the same location $P_i : (x_i, y_i)$: the floor of the tunnel, its roof and the terrain on top of it).
From a topological point of view, a 2.5D model corresponds to a two-manifold of genus 0. While this actually still holds for surfaces with overhangs, the inclusion of objects like tunnels or bridges causes a change in the genus, i.e. in the number of “handles” that are present in the surface. The shape of a mug can be taken as an example for a topological surface of genus 1.
In § 2.4 some approaches, whose goal is to overcome the limitations of a 2.5D terrain model, will be presented. With the above mentioned properties and limits in mind, several types of digital terrain models exist.

2.2.1 Raster and grids models

Raster and grids belong to the family of models and data structures where height values are regularly distributed, at fixed intervals, over the xy-domain.
A raster model is perhaps the easiest model which can be encountered. It is basically a regular tessellation of the domain space, i.e. a partition of the plane with non-overlapping and covering faces, where every face is generally a rectangle or a square, even if other shapes are possible but seldom used.
To each of these elements, called cells, a height value is associated, which can represent the mean value of all heights falling into a single cell (or the maximum, or the minimum or any other statistically defined value). The neighbourhood relations between the cells are known implicitly and, due to their structure, step-like discontinuities along the borders of the cells can be possible: two points close to each other but lying on different cells, might have similar heights in reality, but different values in the DTM.

A raster model can be assumed to be similar to an image: a two-dimensional matrix, where a grey value (the height) is associated to each pixel (the xy-coordinates). Its data structure is therefore a matrix of elevations and is a very simple model. A raster DTM is usually defined over a rectangular domain, and the parameters are the number of cells identified by rows and columns, their dimensions, the position of one reference point (generally the upper or the lower left corner) and, if needed, a rotation parameter with regards to the coordinate system axes. DTM resolution is strictly dependent of the cell fixed size: DTMs with a variety of grid sizes have been generated in order to meet the practical needs of different disciplines operating at different scales, since the requirement on the cell size is one of the major issues, particularly for complex types of terrain.
DTMs based on grid models are conceptually very similar to raster models, although they belong to the vector model family. The main difference consists in height values being associated to points instead of cells. These points are regularly spaced on the xy-plane and are called grid points. Grid points are not directly measured data, but their heights are instead evaluated from the surrounding surveyed points, according to the selected grid spacing, using methods for surface interpolation or approximation.
Generally, neighbouring grid points are connected through line segments, called edges, which lead to a mesh; since the grid points are regularly spaced, the resulting mesh is called a rectangular mesh. Every inner grid point is connected directly with four neighbouring points and indirectly with four other points placed over the diagonals.
Since, unlike in a raster model, height values are available only point-wise, the value at an unknown location between the grid points must be calculated by some interpolation methods. Possible strategies consist in a triangulation of every mesh (four grid points originate two triangles), or an insertion of one central point in every mesh followed by a triangulation (leading to four triangles in every mesh). Otherwise, it is possible to define a local interpolating surface for each mesh which uses only the four surrounding grid points or even some neighbouring ones for a better prediction.
Grid models can sometimes be “enriched” by means of other vector elements such as points, lines or polygons. These elements are called breaklines and used to represent particular discontinuities in the mesh, so that the tangent plane is not constant along them anymore. The lines and their vertices are included in the model, as well as any intersection between the lines and the edges of the mesh. If the model is going to be triangulated, breaklines lead to a constrained triangulation, as explained in the next paragraph.

2.2.2 Triangulated Irregular Networks (TINs)

Rasters and grid models store information placed on a regular pattern. As mentioned before, the stored values are generally not direct measures but have been obtained by means of some mathematical method from the measured data. A different approach is represented by TINs, also called triangulations.
Height information is given either in form of an attribute joined to two-dimensional points, or as three-dimensional points. It is the case of 3D point clouds, a typical product of laser scanning or automatic digital photogrammetry, where point density may vary

across the dataset and is not constrained anymore to a fixed grid. Moreover, the shape of the terrain is described by original measurements themselves. Different point densities are allowed in the same model, thus higher resolutions allow for a more precise representation of details in the terrain, a lower resolution can be used in areas that are less variable or less interesting.

Since point clouds lack information about topology, the latter can be added by means of triangulation. Adjacent points, called nodes, are connected by means of linear segments, called edges, and the result is a TIN.

Different algorithms exist to generate a triangulation, i.e. to define which points are connected together as neighbours. The greedy triangulation algorithm, for example, inserts edges one by one by their increasing length from the shortest one, thus minimising the total length of edges. If the to-be-inserted edge breaks the triangulation constraints it is discarded and the next longest one is considered, otherwise it is accepted and added to the existing edges.

In case of a DELAUNAY triangulation, it corresponds to the dual graph of the VORONOI diagram[1] for a given point set. The algorithm maximises the minimum angle of all the angles of the triangles in the triangulation, avoiding "thin" triangles, in that no point must be inside the circumcircle of any triangle. Since a circumcircle of a triangle is formed by three points from the original point set, two special cases can happen: a) for a set of points on the same line there is, obviously, no DELAUNAY triangulation and b) for four points on the same circle the DELAUNAY triangulation is not unique. For example, in case of an inscribed rectangle, there are two possible triangulations that split it into two triangles.

Analogously to the aforesaid "enriched" hybrid grid models, also in triangulations special lines or polylines representing linear features such as ridgelines or stream courses can be inserted, so that the edges of those lines are also edges in the triangulation. In this case the result is called a *constrained* triangulation.

Once the TIN topology has been established by means of a triangulation algorithms, it is necessary to define a proper interpolation method for the height of any arbitrary point belonging to the triangle formed by the three nearest. Since only the height values at the vertices of the triangle are known, these values could be used. Generally, the simplest way consists in a linear interpolation, leading to a global surface made of many connected, non overlapping planar triangles. However, other techniques have been developed in order to add geometrical continuity in terms of smoothness between the different triangles. A curved surface can be defined, for example, on the domain of every triangle. Further constraints can be added, so that tangential continuity or even curvature continuity can be guaranteed over the neighbouring triangles. The global surface is therefore made of such curved surfaces.

Other interpolation strategies consider not only the triangle vertices, but more neighbouring points.

[1] A VORONOI diagram is a type of space tessellation. Given a set of input points in the plane, a convex (but sometimes unbounded) polygon is associated to every point P. This polygon is called VORONOI cell and represents all the points in the plane that are closer to point P. Edges of the VORONOI cell are all the points in the plane that are equidistant to two input points. Nodes of the VORONOI cell are the points equidistant to at least three input points.

TIN models are not so widely available as raster models and tend to be more expensive to build and process because of their more complex data structure. TINs are typically used in engineering applications for high precision modelling of smaller areas, while grid data area generally available for wider areas.
Although not strictly limited to the underlying structure model like rasters, nowadays state-of-the-art vector DTM models are still 2.5D. Several different strategies, whose goal is to move toward and fully integrate the third dimension, are present subject of research.

2.2.3 Other terrain models

The above described models represent the most common types that are used in Geographical Information Systems (GIS) for terrain modelling, but they are not the only ones. In the following, some others are briefly presented although they may be restricted to particular disciplines or applications.
The shape of the terrain can be described in cartography by contour lines, which are curved or straight lines describing the intersection of the terrain surface with one or more horizontal planes. Along a contour line all points representing the terrain surface have a constant value. However, if contour lines are the only available information, and if the underlying function defining the shape of the terrain is not known *a priori* (or was never obtained), some problems arise. Unlike a bivariate function, whose parameter domain corresponds to the "whole" xy-plane, in a contour line map it consists only of lines with an associated height value, but no other information is available for any arbitrary point between them. Since contour lines have been existing for long time, it is however not surprising that their digitalisation has been the next "natural" step toward the creation of digital terrain models and they are indeed used for some applications. Objects, whose geometries intersect contour lines, can be "draped" on top of them using height values at the intersections. On the other hand, in case a geometry item never intersects a contour line, no height information can be directly obtained and associated to it unless some interpolation strategy is adopted.
An extension of a raster cell or pixel in the third dimension is represented by its volumetric counterpart: a volumetric pixel, also called voxel. A voxel is a volume element representing the scalar (or vector) value on a regularly spaced three-dimensional grid. The boundaries of a voxel are exactly in the middle between neighbouring grid points and generally are set to be tetrahedrons or a cuboid.
Voxels are frequently used in the visualisation and analysis of medical and scientific data; in case of a terrain model, the terrain surface can be seen as the interface between air and land as solid earth. A voxel is therefore defined as an element which belongs to the solid part or not [PFEIFER, 2002]. Due to its inner structure, such a volumetric model can be used to represent fully three-dimensional objects like overhangs, caves, bridges or even "holes" underneath a surface.
In case of a boundary representation (B-Rep) model, a volume is described by its bounding surface, e.g. a solid is represented as a collection of connected surface elements, which define the boundary between solid and non-solid. B-Rep models are composed of topology and geometry, where the topological items consist of vertices,

edges and faces. A vertex corresponds to a point, an edge is a bounded piece of a curve and a face is a portion of surface bounded by edges. With such concepts in mind, it is therefore easy to recognise that a TIN is actually a two-dimensional case of B-Rep model.

2.3 Spatial interpolation

According to [BURROUGH & MCDONNELL, 1998], interpolation is the process to predict the values of a certain variable of interest at unsampled locations based on measured values at points within the area of interest; it is to be differentiated from extrapolation, which deals with the prediction of values of a certain variable *outside* the sampling area. Generally, a DTM is coupled with an interpolation method used to generate a finer model based on the given sample elevation points (gridded or not). Its goal is to convert point data to surface data.

The fundamental idea behind interpolation is that “near” points are more related than “distant” points and, therefore, near points generally receive higher weights than far away points. The obtained surface can pass through the measured points or not. In this case interpolation methods are classified into exact and inexact.

In case of an exact interpolator, a predicted value at a sample location coincides with the measurements values at the same location, otherwise it is the case of a inexact interpolator: predictions are different from the measured values at sampled locations and their differences are used to give a statement about the model quality.

The very large number of existing interpolation models allows to define different classification criteria, according to their characteristics.

A distinction can be made between deterministic and geostatistical interpolation methods. The first are based on mathematical functions that calculate the values at unknown locations according either to the degree of similarity or to the degree of smoothing in relation with neighbouring data. Typical examples of these interpolation family are Inverse Distance Weighting (IDW) or radial basis functions (e.g. thin-plate spline, spline with tension).

The latter use both mathematical and statistical methods, in order to predict values and their probabilistic estimates of the quality of the interpolation. These estimates are obtained using the spatial autocorrelation among data points.

In addition, interpolation methods can be classified into global or local with regards to whether they use all the available sample points to generate predictions for the whole area of interest, or only a subset of them, respectively.

In case of a global approach, then a proper weighting is generally necessary. Closer points should have a more important role than further points, moreover some interpolation models applied at global level might show unwanted side-effects (e.g. ROUNGE’s phenomenon by polynomials of high degree). Some algorithms with global behaviour include kriging, polynomial trend analyses, spline interpolation and Finite Element Method (FEM), for example, and these methods can be used to evaluate and separate trends in the data.

In case of a local approach, the predicted value is instead obtained only from known points within a certain distance, whereas the concept of distance does not refer strictly

to the Euclidean one only, but more generally to neighbourhood. Algorithms belonging to this class include, for example, nearest neighbour and natural neighbour interpolation. Global algorithms can also be applied at local level. The entire dataset is first split up into smaller areas, like “tiles”, then the global algorithms are applied locally to every tile, and the resulting terrain model is reassembled, as a sort of puzzle, from the models of every single tile. In order to guarantee the needed continuity (with regards to height and tangents) the extents of every tile are augmented, so that overlaps zone are created which exceed into the neighbouring tiles. The size of the overlaps is generally chosen in a way that points lying outside the overlaps have no influence on the surface inside the tile.
Regardless of the interpolation model adopted, it is hardly possible to define a mathematical expression that precisely describes the complex nature of the real-world terrain surface. Each of the existing methods of interpolation is, therefore, only an approximation method under certain assumptions. Error and quality parameters to describe how much a DTM adheres to the real world are therefore essential. Some will be presented in § 2.5.

2.4 DTM: going from 2.5D to 3D

Presently, a DTM is generally based on a 2.5D approach. However, such a model is unable to describe objects like overhangs, caves or complex structures. Since, for example, three-dimensional city models are quickly spreading [KOLBE *et al.*, 2008] and cannot be fully represented by a 2.5D approach, it is nowadays commonly accepted that a gradual step toward implementation of a “real-world” third dimension is inevitable. This means that it is necessary to gradually develop new applications able to handle full 3D information, and this represents one of the major subjects of current research.
However, simply adopting a global 3D approach would not be the best solution, due to the complexity of fully three-dimensional algorithms, since they tend to perform not as quickly as the established 2.5D algorithms. In the following, some recent approaches dealing with this topic will be shown.

2.4.1 Parametric patch surfaces

Starting from a triangulated surface, PFEIFER (2002) has presented an approach whose goal is to overcome some of the implicit limitations of a TIN: no vertical elements or three-dimensional objects like bridges or overhangs can be modelled and, moreover, triangulated surfaces are generally not smooth. In his work he presents two methods to obtain a smooth surface over a triangulation.
In the first one, a smooth surface, called patch, is defined for each triangle. Every patch is a small curved surface defined over the domain of one triangle; its shape is determined both by local properties of the surface, like curvature, and by continuity requirements along the patch edges. Along common edges, neighbouring patches share common borders and common tangent plane fields. In brief, the approach works as follows: for every vertex of the triangulation a normal vector is calculated; then curves are defined for every edge and along these curves a normal vector field is

estimated. Finally patches that interpolate the boundary curves and approximate the normal vector field are calculated, see Figure 2.1.

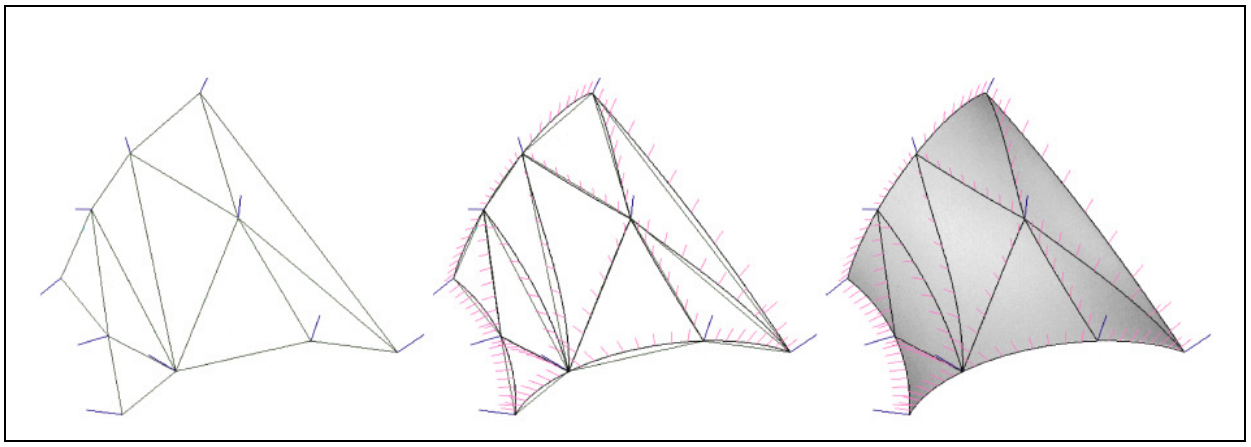

Figure 2.1: Construction steps in PFEIFER's approach for triangular patchworks. [Left] Triangular mesh and vertex normals. [Centre] Boundary curves and normal vector fields. [Right] Patches for each triangular face. Image source: [PFEIFER, 2002].

Since it is possible to define a patch for every triangle of the initial mesh, the resulting global surface is obtained joining all triangular patches, thus obtaining a so-called "patchwork".
The second approach uses subdivision patches to refine terrain surface. The idea is to insert new vertices, thus new edges, into the triangulation in order to gradually reduce the size of the triangles. If this operation is repeated several times, the more new points and edges are added, the smoother the surface becomes. Besides creating in both ways smooth surfaces, one of the advantages of PFEIFER's approaches is that his methods operate locally, since the shape of the surface at a specific location is influenced only by the neighbouring points.

2.4.2 2.8D maps

Another approach to improve a "classical" 2.5D terrain model without introducing extra complexity related to a real 3D model has been presented by GRÖGER & PLÜMER (2005). They analyse which kind of three-dimensional objects can be represented using models and topology from the 2D world. They extend a standard 2.5D DTM to contain vertical walls and overhangs (thus allowing for more *z* values for each planar point) so that geometry can be three-dimensional, but the underlying topology remains two-dimensional.
The model is based on the notion of a two-manifold embedded in 3D space: in this way, the model is "more" than 2.5D, but is "less" than 3D, hence the reason why the authors call it "2.8D". It is still not possible to model objects like tunnels and bridges, which would lead to a surface which is topologically equivalent to a two-manifold of genus 1. However, their approach permits to detect and to handle these cases in order to guarantee consistency inside the model.

The advantages of GRÖGER & PLÜMER's work consist in the conceptual similarity and simplicity to the well-known 2D approach, coupled with the relatively small effort required to extend 2.5D tools to 2.8D.

2.4.3 3D enclaves

OTEPKA, BRIESE & NOTHEGGER (2006) have presented a combined 2.5D and 3D approach. The idea consists of a 2.5D base model which can contain spatially delimited 3D objects, called enclaves, as shown in Figure 2.2. Existing algorithms can be employed for 2.5D areas, while each enclave object is provided with several extra functions, requiring that the object, for example, be able to represent itself as triangulation or to reduce itself to 2.5D.
An important aspect of this approach is the possibility to obtain a smooth transition between the enclave and its 2.5D surrounding area. In order to achieve the goal, first an interpolation across the complete base model domain is performed, whereas information from the enclave object is considered, too. The surface obtained is then accepted or rejected by the enclave object. In case it is accepted, the surface inside the enclave boundary, as well as height and tangential information along it, is saved and this can be used to model the Earth's surface below the enclave object (i.e. the surface under a bridge) and to guarantee a smooth transition between base map and enclave. Next, using also the previously stored geometrical continuity information, each enclave object is modelled by its own interpolation strategy.

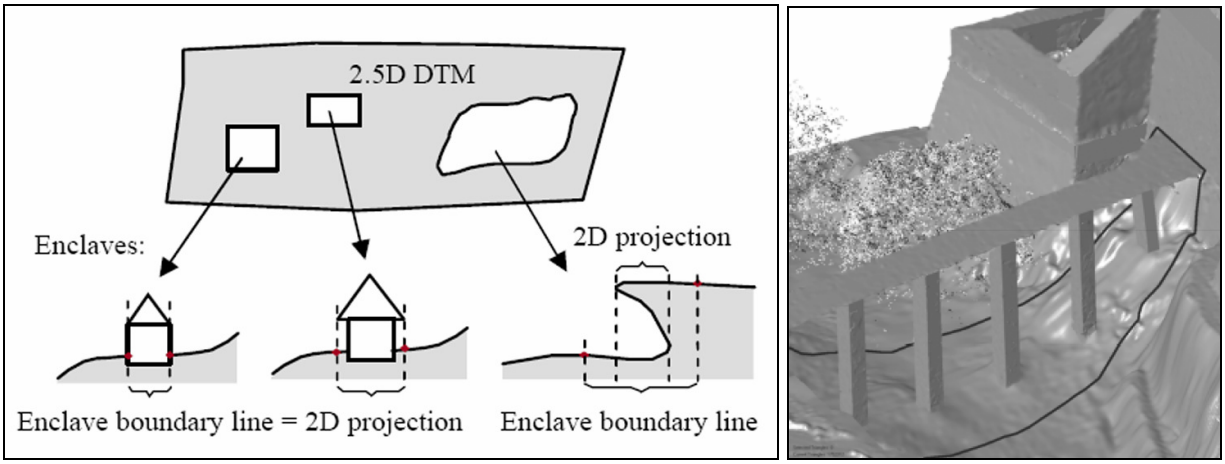

Figure 2.2: Examples of various 3D enclaves including their 2D projection and boundary line. Image source: [OTEPKA, BRIESE & NOTHEGGER, 2006].

2.5 DTM quality assessment

Given the continuously increasing importance that digital terrain models are gaining in recent years due to their wide field of applications, it is of great relevance to define and to assess the quality of the information they contain. Moreover, DTMs can be produced nowadays at shorter time intervals and for larger areas, and the production process requires quality controls so that a DTM can be checked and updated properly. For example, changes in the landscape caused by anthropic intervention or natural events

make update operations of a DTM necessary, since the geometric accuracy of the DTM influences very much the quality of the derived products.
Several interrelated factors contribute to define the quality of a DTM, like for example the original terrain geomorphology, the methods of data acquisition, the methods employed in generating the DTM itself, the vertical and the horizontal accuracy of the grid points.
Digital terrain models can be obtained by different methods and each method has its own characteristics. As an exemplification, a comparison between DTMs obtained from contour lines, laser scanning and automated photogrammetry can be made taking the standard deviation σ, one of the most used deviation descriptors, as a comparable parameter.
Scanning and interpolating contour lines from topographic maps at a given scale, for instance 1:25000, in order to obtain grid points, leads to height values representing the terrain with a vertical standard deviation of $\sigma_h \approx 1.5\text{-}3$ m and a horizontal resolution of 10 to 25 m.
More recent methodologies provide height values that are referred to the surface on top of buildings and vegetation: they need to be filtered and reduced to the underlying terrain. The final products can have a denser horizontal grid spacing, up to 1 m or even less, and a vertical accuracy $\sigma_h \approx 0.15\text{-}0.30$ m for laser scanning or $\sigma_h \approx 0.02\text{-}0.03\%$ of the flying height for photogrammetry (e.g. $\sigma_h \approx 0.8\text{-}1.2$ m at 4000 m). All methods of DTM generation may contain blunders, and proper algorithms must be applied to detect and eliminate them.

2.5.1 Global parameters

In general terms, the quality of a DTM should express: a) how accurate a height value is at a certain location, i.e. its absolute, external accuracy, and b) how accurately the morphology is presented, i.e. its relative, internal accuracy, also defined as precision. It is consequently clear that a single σ value does not suffice and has to be differentiated according to whether it refers to precision or accuracy [MIKHAIL & ACKERMANN, 1976].
Over the years, several approaches have been proposed for DTM quality assessment. An estimation of external accuracy should be obtained by comparing original elevations with the elevation values in a DTM surface.
Every DTM represents in fact only an approximation of the real terrain and if its surface is considered as a bivariate function $z_{DTM} = f_{DTM}(x, y)$, it can be compared to the real terrain $z = f(x, y)$. The error of the DTM can be expressed as the difference between the above mentioned surface at selected, tested points:

$$\Delta z = z - z_{DTM} \qquad [1]$$

or, in the continuous domain as

$$d(x, y) = f(x, y) - f_{DTM}(x, y) \qquad [2]$$

These error values between two sets of elevation data can be analysed with conventional statistical methods such as the standard deviation or the mean, as well as many other statistical measures that are effective to describe a frequency distribution (e.g. central tendency and dispersion measures).
For example, in the continuous domain, the variance of the DTM error is expressed as

$$\sigma^2 = \frac{1}{A}\int\int d(x,y)^2 dxdy \qquad [3]$$

where *A* is the area of the integration domain on the xy-plane.
However, the real terrain surface is unknown in practice. The variance between the estimated and the real terrain cannot be calculated and is generally approximated with the difference between the elevation value of *n* check points and the corresponding elevation value in the DTM. This leads to the most widely used error parameter for DTM: the Root Mean Square Error (RMSE), which measures the dispersion of the frequency distribution of differences between the original elevation data and the DTM data. Mathematically it is expressed as:

$$RMSE = \sqrt{\frac{\sum (z_i - z_{i,DTM})^2}{n}} = \sqrt{\frac{\sum \Delta z_i^{\,2}}{n}} \qquad [4]$$

and it means that the larger its value is, the greater the discrepancy between the two datasets is. One of the reasons for the wide adoption of the RSME as error parameter lies in its easy computation and simple concept.
However, some authors [WOOD, 1996; KRAUS, 2004] have expressed doubts about its validity. First of all, the RSME is essentially a single global measure, which means that the error is the assumed to be the same everywhere in the area covered by the DTM, but this is of course not necessarily true. Secondly, there is the assumption that there is no bias in the error, which is often invalid [LI, 1993].
Besides, the RSME does not give any information about the mean deviation between the two measures of elevation, neither the form of frequency distribution and its skewness. Moreover, the magnitude of the RMSE value is also influenced by the variance of a true elevation distribution and depends on the relative relief and scale of measurements.
Some other commonly used global quality parameters, which may consider blunders, are presented in Table 2.1. The threshold value for blunders is not fixed and actually changes according to the type of landscape. It is generally accepted to be set at 3 times the RMSE error.

Table 2.1: Some of the most common accuracy measures for DTMs. Source: [EuroSDR, 2006]

Vertical Accuracy	Parameter
Height difference from reference data	Δz
Number of points in dataset	n
Maximum height difference	$\lvert\Delta z_{\max}\rvert$
Definition of blunders	$\Delta z_{blunder} > 3 * RMSE$
Number of blunders	N
Number of point without blunders	$n' = n - N$
Mean	$\mu = \frac{\sum \Delta z}{n'}$
Standard deviation	$\sigma = \sqrt{\frac{\sum (\Delta z - \mu)^2}{n' - 1}}$
Horizontal accuracy	$\sigma_p = \sqrt{(\sigma_x^2 + \sigma_y^2)}$

However, measuring the error of a DTM is actually often nearly impossible, since the true values for all geographic features represented in a geographic dataset cannot be always determined completely [GOODCHILD *et al.*, 1994]. In some cases check points cannot be used for the determination of quality parameters, and another, more precise DTM is used as reference instead. An alternative consists in excluding a small part of the input data from the DTM creation process and in successively using it for accuracy tests. By means of this form of cross-validation it is possible to have a quality statement about both the DTM and the input data used for its generation, however the information generated should be interpreted using the methods for internal accuracy (precision) described in the following.

Finally, acquisition of external data of superior quality can require additional observations in the field and can lead to higher costs; depending on the specific methodology used for DTM processing, some empiric models have been presented in the past years that basically permit the estimation of DTM height accuracy *a posteriori*. KRAUS (2004), has presented the following formula for stereophotogrammetry, which can be used for open terrain:

$$\sigma_H = \pm\left(0.00015h + 0.15\frac{h}{c}\tan\alpha\right) \qquad [5]$$

with

σ_H : standard deviation of DTM,

h : height of flight,

c : focal length [mm],

α : terrain slope.

For Aerial Laser Scanning (ALS) data, KAREL & KRAUS (2006) propose instead the following formula:

$$\sigma_H[cm] = \pm\left(\frac{6}{\sqrt{d}} + 50\tan\alpha\right) \qquad [6]$$

with
σ_H : standard deviation of DTM,
α : terrain slope,
d : point density (points per square metre).

The last two formulae are supposed to be used *a posteriori*, but they can indeed be "reversed" to give information *a priori*. In [5] *h/c* represents the image scale; if the approximate maximum terrain slope α and the flying altitude h are known, it is possible to define *a priori* the proper image scale for the DTM requirements.
In a similar way, [6] can be used to obtain the optimal flying altitude, although only point density and terrain slope is included in the formula. However, data density does depend on the flying altitude, since it is related to the angle step width of the scanner.
For a DTM resulting from cartographic digitising, WENG (2002) proposes that its uncertainty be related to three aspects: the elevation measurements for the topographic map, the sampling and measurement error and the interpolation method. If every source of uncertainty is uncorrelated, the total DTM uncertainty $RMSE_{total}$ is computed as follows:

$$RMSE_{total} = \sqrt{(RMSE_m)^2 + (RMSE_s)^2 + (RMSE_i)^2} \qquad [7]$$

with
$RMSE_m$: uncertainty of the digitised source map,
$RMSE_s$: the sampling and measurement error,
$RMSE_i$: the interpolation method error.

WENG suggests to create list table in which the estimate of the uncertainty from every distinct source is reported, in order to provide some guidelines in the creation of DTMs with the cartographic digitising method. However, he specifies that several other factors must be taken care of in the process. Equation [7], for example, does not count for the correlation between uncertainty in the source map and the uncertainty caused by the interpolation (e.g. irregular terrain tends to concentrate high uncertainty in both error sources). The scale of a source map has a direct impact on the quality of DTM generated, since the accuracy of terrain parameters and features derived from DTMs exhibits scale dependencies [HUTCHINSON & GALLANT, 1999]. Finally, particular care must be taken in the process of digitalisation, because point density and point distribution influence the accuracy of the interpolation method. The morphological complexity of terrain and the different data acquisition patterns during digitalisation should be considered, e.g. random vs. systematic, significant points vs. contouring.

2.5.2 Local parameters

The previously presented quality parameters are often used to give a global quality value for the entire DTM, assuming a uniform error value for the DTM surface, but this assumption is often not satisfactory. The distribution of errors in a DTM can show some form of spatial pattern which can be best observed in a graphical representation – creating, for example, an accuracy surface.

Global quality parameters have the advantage to describe the whole area of interest with few parameters only. On the other hand, local ones describe the quality of a DTM at a higher level of detail. Therefore much effort has been put in the definition of adequate local quality parameters which can be used throughout the whole DTM creation process: from the original raw input data used for DTM generation, up to the definition of quality measures that are coupled with the obtained models.

KAREL, PFEIFER & BRIESE (2006) give a deeper insight of the topic and give an overview of several local quality assessment methods. Some of their considerations are presented.

Regarding the input data, DTMs generally result from datasets consisting of point and lines. Aerial laser scanning or digital photogrammetry can produce millions of surface points, randomly distributed and without topologic information. Accuracy information in height (σ_z) might be sometimes provided, as well as further planimetric accuracy values (σ_p) and metadata like sensor model, flight date, flight altitude, or other accuracy values of the matching process. However it is not uncommon that, if ever, only one global quality parameter is given for height accuracy. Nevertheless some initial investigations can be carried out: if σ_z and σ_p are given, maps showing their distribution are created.

Also data distribution can be considered: density maps, which show the number of points per area unit, depend only on the horizontal position of the data. They can be easily computed and visualised and can be helpful at identifying zones that are covered by too few or no data at all.

Distance maps can represent another helpful tool: the distance between the central coordinates of every pixel and the closest surrounding point is computed and represented. By means of a two-dimensional VORONOI tessellation over the input data points, an estimation of a DTM model reliability can be pursued, since reliability grows according the proximity of an interpolated point to its nearest data point. On the other hand, once a certain distance threshold has been set, interpolated points that are too far away from a measured point should be considered more properly as an extrapolation.

Finally, if DTMs are computed from different, overlapping datasets, a check over the consistency of the used data could help at highlighting possible systematic errors within the common areas, which may originate from imprecise sensor orientation or as a result of a sensor fault.

When it comes to the DTM itself, both precision and accuracy must be considered in order to assess its quality. In the process of DTM production, given the nowadays large quantities of input points, one initial problem consists in lack of topology, which is needed as soon as the simple point geometry does not suffice anymore and, for example, surface information is needed. A simple solution consists in performing a triangulation of the point data, but the quality of the resulting surface is generally very

poor due to the measurement errors that are included into the surface. The surface tends not to be smooth enough and does not provide precision estimations for the model.

Precision, defined above as internal accuracy, expresses how well the model adheres to the input data and is tightly dependent on data redundancy. The higher the density of the input data, the higher the level of detail which is likely to be obtained in the DTM, and, at the same time, the higher the possibility to eliminate potential random errors[2]. Thus, in order to generate better quality DTMs, algorithms which use redundancy, e.g. kriging or FEM, are generally preferred to identify terrain points from off-terrain points, to filter them out and to compensate random measurement errors.

The filtering process allows to obtain smooth, regular surfaces; at the same time the drawback consists in smoothing even areas where potentially relevant changes in gradient or curvature are present, typically along geomorphological structures. For this reason breaklines can be added to the model, leading to an enriched DTM [BRIESE, 2004].

Once the DTM surface has been interpolated from the filtered point cloud, a further simplification though data reduction can be carried out. A DTM could be stored as a regular grid, in this case a proper discretisation interval and the grid size are chosen and only the grid points hold precise height information from the original surface, while heights are computed from the simplified DTM.

Quality from the original and the simplified DTM will therefore differ, although remarkable height differences may imply a too large sub-sampling cell size. A gridded DTM can nevertheless be improved, in that additional vector data like peak points or breaklines are added to form a hybrid DTM.

Once the DTM surface has been obtained, measures of its precision can be generated with several approaches. Some use error propagation techniques, which allow to estimate the standard deviation in height, σ_h, as well as the precision of derivative values like slope and curvature. Kriging is for instance a widely used interpolation method, however its results are influenced by decisions that must be made by the user with regards to the input parameters.

Another approach to quantify the DTM precision consists in computing and analysing the height differences between the original points and the interpolated surface. Some of the global parameters seen before: the maximum, the mean, the median or the RMSE can still be used, but at local level adopting a cell approach.

KAREL, PFEIFER & BRIESE suggest furthermore that residuals, normalised by their standard deviation *a priori*, should be preferred in case of observations holding different weights or accuracies *a priori*.

Another approach to the determine local quality parameters is given by KAREL & KRAUS (2006). Provided that both the DTM and the input dataset, where blunders and points not belonging to the terrain surface have already been filtered out, are given, it is possible to estimate DTM accuracy, although no information is known on its

[2] According to the NYQUIST–SHANNON sampling theorem, a sampled analog signal can be reconstructed if the sampling rate is greater than 2*f* samples per second, with *f* the highest frequency in the original signal.

interpolation method. The presented empirical stochastic approach is based on characteristics like point density, the DTM curvature and point alignment. It provides accuracy values for every distinct grid point *G* of the DTM, facilitating visualisation by means of quality surfaces.
First, the point density $\overline{n}$ is calculated using the original input data and an overlaying regular grid. Then the distance s_i between each grid point *G* and the data point next to it is calculated using the Chamfer function. After curvature values at each grid point have been computed along the grid lines, the maximum, the minimum main curvatures, the curvature $1/r_{\alpha_i}$ in any α_i direction are obtained.
For the surroundings of every grid point G an accuracy value can be computed in the form of a weighed $RMSE_w$ with the expression

$$RMSE_w = \sqrt{\frac{\sum p_i d_i^2}{\sum p_i}} \qquad [8]$$

where d_i stands for residuals between original data points and DTM surface and p_i for the weights. If the $RMSE_w$ is smaller than the $\sigma_{a\ priori}$, then it is replaced by that value. Grid cells with few terrain points are susceptible of this phenomenon; another possibility is that random errors have not been filtered out from the DTM.
The weights have following expression:

$$p_i = \frac{1}{1 + \frac{s_i^2}{r_{\alpha_i}^2}} \qquad [9]$$

with
s_i : distance between grid point *G* and original point P_i,
r_{α_i} : curvature radius of the DTM at the original point P_i in direction of the grid point G,

and are meant to emphasise data points that are closer to the grid points, as well as those with smaller curvature, see Figure 2.3. Moreover, a threshold can be set in order to restrict the size of the neighbourhood area around a grid point G and therefore the number of points to use for the successive fitting.
The interpolated DTM surface is in fact not known, but it suffices to approximate it with a tilted plane centred at every grid point. The accuracy of DTMs suffers only marginally from the interpolation, which on the other hand can greatly influence the geomorphologic quality of the interpolated surface [KRAUS *et al.*, 2004]. The tilted plane is thus obtained by fitting the surrounding original data points with least squares and permits therefore to estimate the DTM accuracy for the grid point as

$$\sigma_{z,DTM} = RMSE_w \sqrt{q} \qquad [10]$$

where q is cofactor in height of the least squares adjustment of the tilted plane.

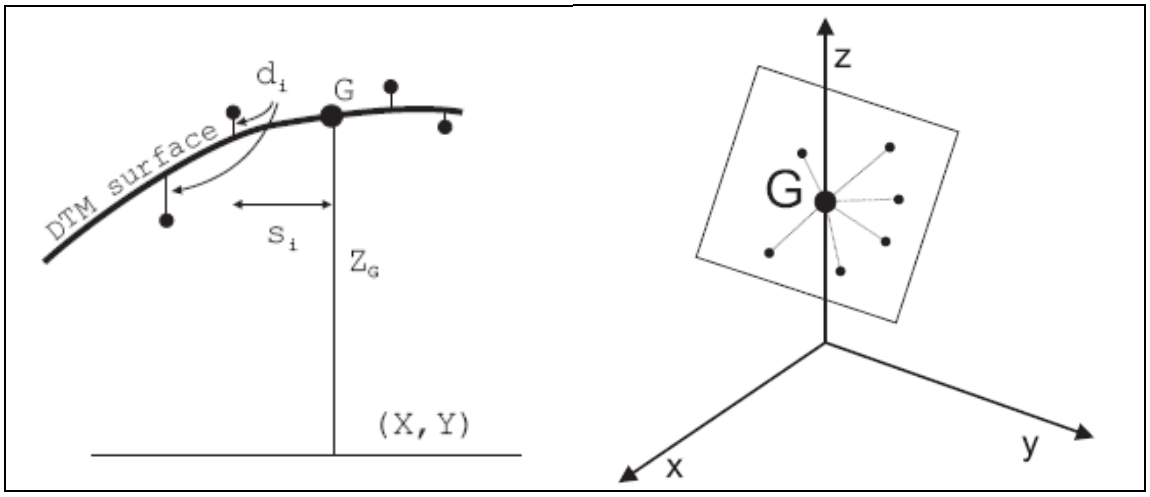

Figure 2.3: Geometric parameters to compute local quality parameters. Image source: [KAREL & KRAUS, 2006].

If $q < 1$ then the grid point G has a higher accuracy than the surrounding data points. If $q = 1$ then the G has the same accuracy as the surrounding data points. If $q > 1$ then G has a lower accuracy than the neighbouring data points, possibly due to large distances or high curvatures to the next data points. In this case, areas of the DTM which are located too distant from original data are identified as unusable.

2.5.3 Open issues

In the previous paragraphs several quality parameters for DTM have been presented, due to continuously growing importance of terrain models and their related applications. It is therefore desirable that modern and future DTMs will be provided not only in terms of height models, but also with adequate quality information: global parameters for a synthetical description *and* local parameters for a more detailed representation. It will pertain to the user to employ the quality information, depending on the application.
It must be noted that it still remains an open question which quality measures must be eventually chosen and how they will influence the decisions based on them.
In addition, the rapid technological improvements and the associated decreasing costs in data acquisition will surely contribute to quicken the production of modern, more precise DTMs, but it must not be forgotten that large amounts of previous data still exist. For these “older” DTMs, often only one single height accuracy value is given, but sometimes quality information is completely missing.
Local update or enhancement operations, restricted to some portions of a dataset only, or integration of heterogeneous DTMs represent therefore another field of research where no definitive solutions have been found yet, and new investigations are required.

Chapter 3

SPATIAL DATA INTEGRATION

This chapter deals with issues related to integration of spatial data. Some motivations are initially presented, followed by some definitions which are used throughout this and the following chapters. A general overview is given about data integration problems in Information Technology; then specific issues, which originate when heterogeneous spatial data are combined, are highlighted. Different solutions to specific problems are shown, with particular attention to the geometric homogenisation process between two or more overlapping datasets. Finally, some still unresolved issues and the research subject of this thesis are introduced, acting as a bridge to the next chapter.

3.1 Motivation

Generally speaking, it can be assumed that nowadays the whole Earth's surface has been mapped – up to different levels of detail, of course. Several methodologies using different sensors and techniques have been used, thus leading to representations of the Earth's surface which have different accuracies depending, for example, on the origin and the goal for which the data were acquired.
It can moreover be assumed that for specific portions of surface multiple models exist, which may represent the same object in different ways, or simply refer to the same geographic area over the course of time.
These multiple models may be the result of different surveying campaigns and, as stated before, may differ from each other for a wide range of reasons: resolution, accuracy, and data structure are only few of the possible factors.
The relatively recent quick development of newer sensors and faster and more accurate acquisition methodologies, along with the booming possibilities offered by *geo-enabled* applications and the growing amount of distributed geodata which can be accessed, are nevertheless pushing forward in the direction of multi-resolution, multi-format data integration. Datasets created for a specific task are therefore being merged into larger, integrated, multi-purpose geodatabases.
The reasons for this convergence are multiple and often related to the specific application framework, but some general principles can be deduced. First, data acquisition and especially the subsequent elaboration and analysis steps still holds a relevant role among the cost factors: integration of existing data into a common environment may reduce costs avoiding the need for new data acquisition. Sometimes older datasets might even be needed on purpose. Secondly, integration has the welcome result to help reducing inconsistencies that are caused by different object modelling, surveying and production methods. Again, the possibility to verify outdated datasets that can be checked against newer ones should not be underestimated.
A very good example of data integration trend is offered by projects like Google Earth or Microsoft Virtual Earth, see Figure 3.1, which have experienced a steady growth in

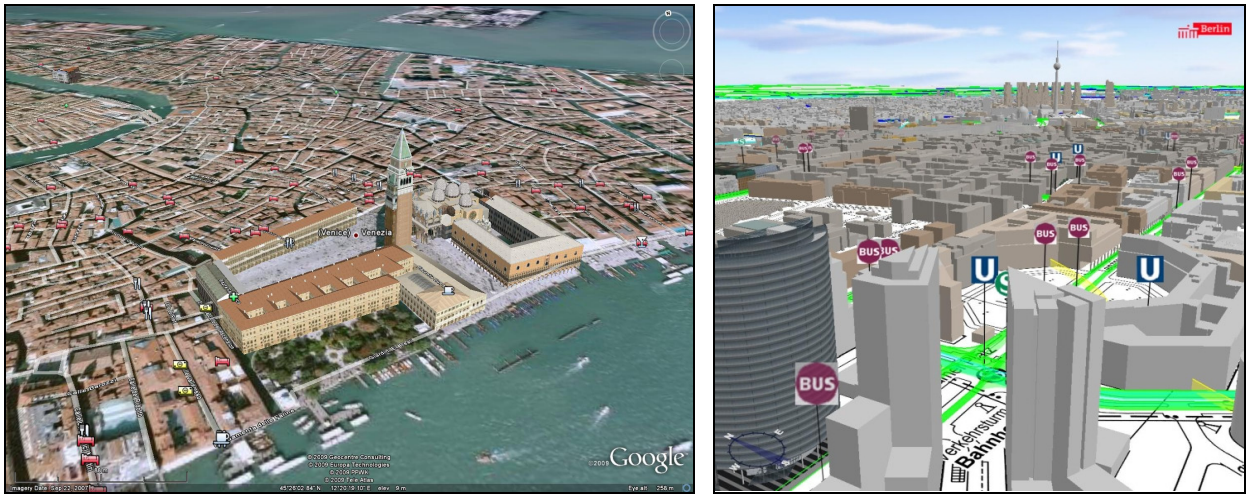

Figure 3.1: Examples of spatial multi-resolution and multi-format data integration: [left] Google Earth project and [right] the official Berlin 3D city model. Right image source: [BERLIN3D].

popularity and amount of offered services since their launch few years ago. What is more, such applications can deliver effective visualisations of large scale models to a broad audiences. Another relatively recent example are 3D city models, which have been literally booming in the last few years: integration of different models like DTMs, buildings, transportation facilities etc. is presently a hot topic which is gathering worldwide large research efforts [KOLBE, 2008]. Virtual 3D city models represent in fact an important and cost effective tool which is useful in those fields where multidisciplinary approaches are of great importance, like urban planning and disaster management, but architectural design and cadastre can profit from them, too.

However, this does not hold true only for "modern" datasets! The last decade has experienced an equal boom in the surveying methodologies applied to archaeological cultural heritage. Similarly to the technological improvements experienced by the DTMs, cultural heritage is one of the areas which has seen the most advances with regards to data acquisition, integration, management and visualisation. Recently, the Rome Reborn project, see Figure 3.2, has been implemented in Google as Ancient Rome 3D layer, see Figure 3.3 [GUIDI, FRISCHER & LUCENTI, 2007], while projects like Virtual Rome represent an example how to jointly offer 3D cultural information and landscape reconstruction about the territory of Rome during Roman imperial times, as shown in Figure 3.4 [PESCARIN, 2009].

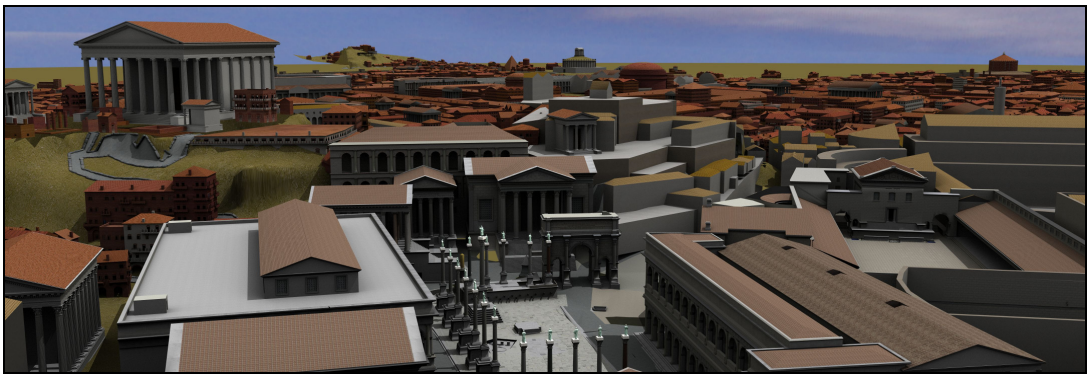

Figure 3.2: Example of the Rome Reborn project…

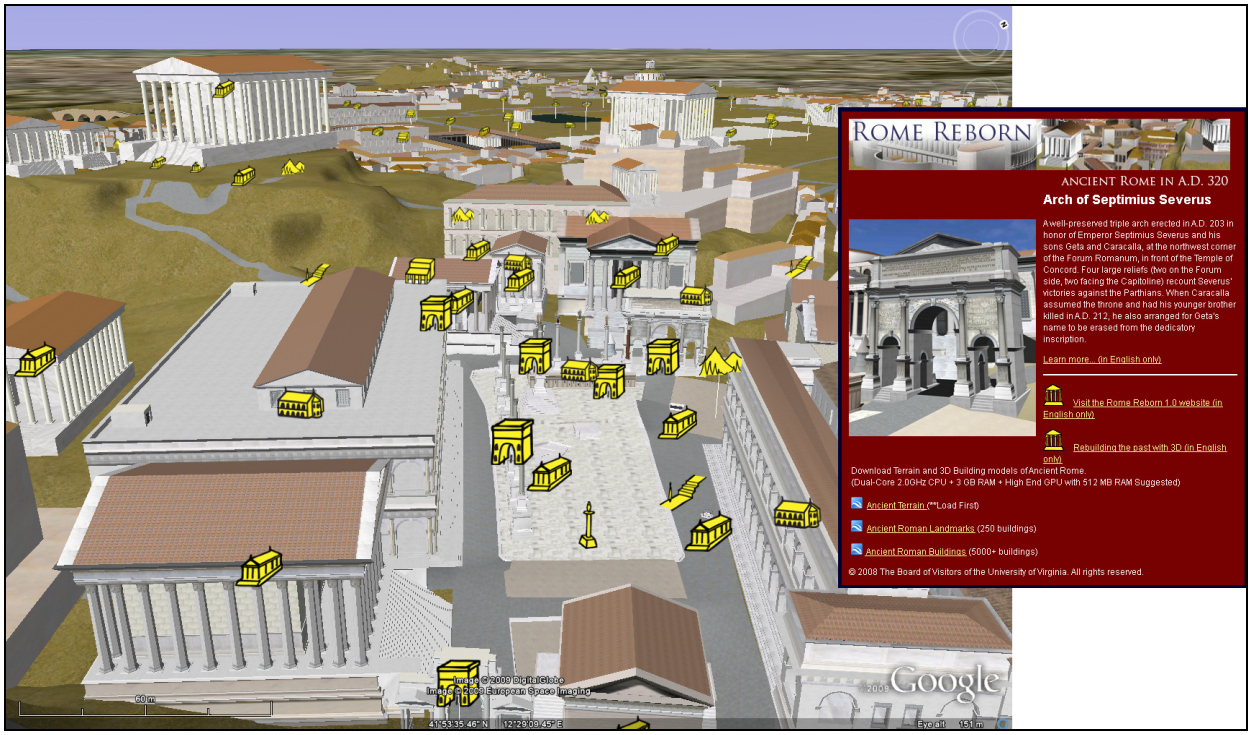

Figure 3.3: …and its implementation in Google Earth.

Figure 3.4: Example of the Virtual Rome project.

Google Earth and 3D city models are not the only players in the field of geoapplications: portable devices which are able to deliver real-time spatial information (mobile telephones, GPS navigation systems, etc.) have become a daily commodity, so that the evident benefits brought by the aforementioned trend in data convergence is quite indisputable.

At the same time, new related problems have emerged: as long a specific task can be carried out using one single model, or a set of models with the same characteristics,

there may be no apparent reason to consider how a model relates to the other *different* ones. However, as soon as two or more heterogeneous datasets need to be simultaneously used, problems related to their integration might arise.

3.2 Some definitions

In this and the following paragraphs, after some general definitions have been given, problems related to (spatial) data integration and data interoperability will be discussed.
According to the Institute of Electrical and Electronics Engineers (IEEE), *interoperability* is defined as the ability of two or more systems or components to exchange information and to use the information that has been exchanged [IEEE, 1990]. It is possible to distinguish several kinds of interoperability: for a computer program, data interoperability means that it can run independently from the data formats; for a dataset, program interoperability means that the dataset can be used by different types of computer programs. In general, interoperability regards both data level and program level [LAURINI, 1998].
In Information Technology, with *information integration* is meant the operation of merging information from disparate sources with differing representations [LENZERINI, 2002]. Data integration should provide the user with a unified view of the data.
The adoption of databases in the last decades has led to the need of sharing or merging existing repositories, and therefore issues resulting from the combination of heterogeneous data sources under a single query interface must be analysed and resolved.
Data fusion and *multi-sensor data fusion* are a subset of information integration and represent the combination of sensory data (or data derived from sensory data) from heterogenous sources. The resulting information is in some way better than it would be possible when these sources were used individually. The term better can mean more accurate, more complete, more dependable, or refer to the result of an emerging view, such as stereoscopic vision [MITCHELL, 2007].
In the geospatial domain data fusion is often synonymous with data integration. Diverse datasets are combined into a unified dataset, whereas the fused dataset is different from a simply combined superset; information in the fused dataset contains attributes and metadata which may not have been included in the original dataset.
According to ZIEGLER & DITTRICH (2004), integration operations can take place at several levels in a database architecture. The goal of this work it not to give an omni-comprehensive overview about data integration: some problems regarding data interoperability will be mentioned here, however the discussion will be focused on integration of spatial data and on the most frequent problems related to the subject.
Assuming a single relational database on a single computer, the structure of a database is represented by its schema, which basically defines how and where data are stored: the tables, the fields in each table, the relationships between fields and tables.
Horizontal fragmentation occurs when structurally analogous records, so-called tuples, are split over different tables. For example, a list of LEGO models could be split for some reasons into separate tables, according to the year of release. In Structured Query Language (SQL), the full list can be obtained by means of a `UNION` operation.

In case of *vertical fragmentation*, data is instead split vertically, i.e. some attributes are located on a table and the other attributes on another table. Data can be retrieved from both tables by a SQL `JOIN` operation using a necessary reference key, see Figure 3.5.
With both horizontal and vertical fragmentation, the case can be referred as *mixed fragmentation*.

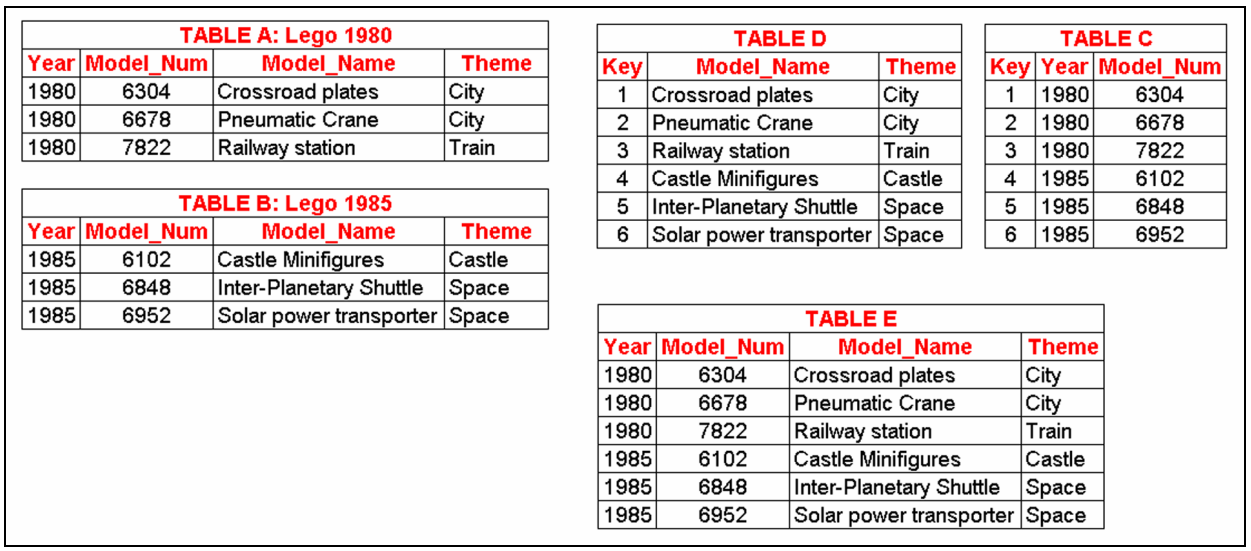

TABLE A: Lego 1980

Year	Model_Num	Model_Name	Theme
1980	6304	Crossroad plates	City
1980	6678	Pneumatic Crane	City
1980	7822	Railway station	Train

TABLE B: Lego 1985

Year	Model_Num	Model_Name	Theme
1985	6102	Castle Minifigures	Castle
1985	6848	Inter-Planetary Shuttle	Space
1985	6952	Solar power transporter	Space

TABLE D

Key	Model_Name	Theme
1	Crossroad plates	City
2	Pneumatic Crane	City
3	Railway station	Train
4	Castle Minifigures	Castle
5	Inter-Planetary Shuttle	Space
6	Solar power transporter	Space

TABLE C

Key	Year	Model_Num
1	1980	6304
2	1980	6678
3	1980	7822
4	1985	6102
5	1985	6848
6	1985	6952

TABLE E

Year	Model_Num	Model_Name	Theme
1980	6304	Crossroad plates	City
1980	6678	Pneumatic Crane	City
1980	7822	Railway station	Train
1985	6102	Castle Minifigures	Castle
1985	6848	Inter-Planetary Shuttle	Space
1985	6952	Solar power transporter	Space

Figure 3.5: Data fragmentation: horizontal fragmentation [left] and vertical fragmentation [right]. Data in tables A and B can be merged by a `UNION` SQL statement. Data from tables C and D can be merged using a `JOIN` SQL statement on the reference key. Both operation result in table E.

A distinction can be made between *structural* and *semantic incompatibilities* among datasets, and also in this case a combination is possible. A simple example can be given with two distinct tables containing a list of telephone numbers.
In case of structural incompatibilities, the structure of the data differs: in a table one field is used for the whole number, while in the other table two fields are used for the area code and the local number, respectively.
In case of a semantic incompatibility, both tables may use the same data structure for data fields, but with different field names. In either case, SQL operations like `UNION` or `INTERSECT` might lead to errors in the resulting output table, see Figure 3.6.

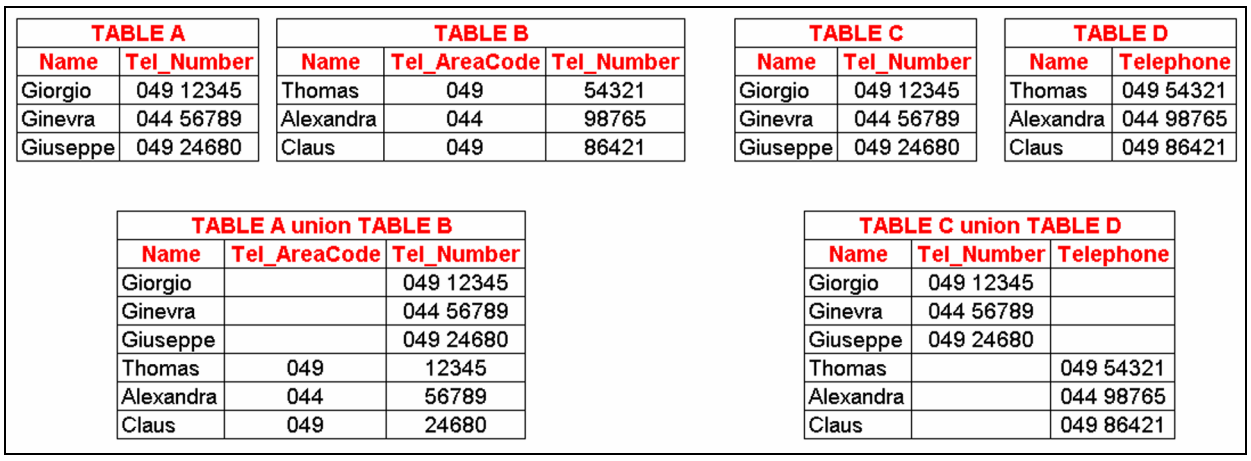

TABLE A

Name	Tel_Number
Giorgio	049 12345
Ginevra	044 56789
Giuseppe	049 24680

TABLE B

Name	Tel_AreaCode	Tel_Number
Thomas	049	54321
Alexandra	044	98765
Claus	049	86421

TABLE C

Name	Tel_Number
Giorgio	049 12345
Ginevra	044 56789
Giuseppe	049 24680

TABLE D

Name	Telephone
Thomas	049 54321
Alexandra	044 98765
Claus	049 86421

TABLE A union TABLE B

Name	Tel_AreaCode	Tel_Number
Giorgio		049 12345
Ginevra		044 56789
Giuseppe		049 24680
Thomas	049	12345
Alexandra	044	56789
Claus	049	24680

TABLE C union TABLE D

Name	Tel_Number	Telephone
Giorgio	049 12345	
Ginevra	044 56789	
Giuseppe	049 24680	
Thomas		049 54321
Alexandra		044 98765
Claus		049 86421

Figure 3.6: Examples of structural [left] and semantic [right] incompatibilities between tables, and the result of a simple SQL `UNION` statement.

A common strategy to solve semantic incompatibilities involves the use of ontologies. These provide a shared vocabulary, such that information can be exchanged unambiguously. Ontologies are used as a form of knowledge representation in order to model the type of objects or concepts, their properties, and relations. This approach represents ontology-based data integration.

The basic definitions given above refer to data contained in a single database on a single computer. It is however clear that the same concepts still hold if data are distributed over several heterogeneous databases running on one or multiple, distributed machines.

The term *site* can thus refer to a single computer with a single database, to several databases located on the same computer or several computers connected such that they appear as a single one from outside.

Similarly to the already defined structural and semantic incompatibilities, the concepts of *syntactic* and *semantic heterogeneity* (or its vice versa: interoperability) can be applied to databases. In the former, aspects related to the possibility of two or more systems to communicate and exchange data are considered (e.g. use of different file formats, access protocols, query languages etc.); the latter stands for the capability to correctly interpret the exchanged information with the goal of delivering useful results as defined by the users.

When it comes to data integration spread over multiple sites, one solution consists in the so-called *data warehousing*. Data from several sources are extracted, transformed and loaded by the warehouse system into a single, unified schema. The advantage of this approach is that data resides together in a unique repository; problems can yet arise when an original data source is updated while the warehouse still contains the older data an therefore the extraction, transformation and load operation must be executed again.

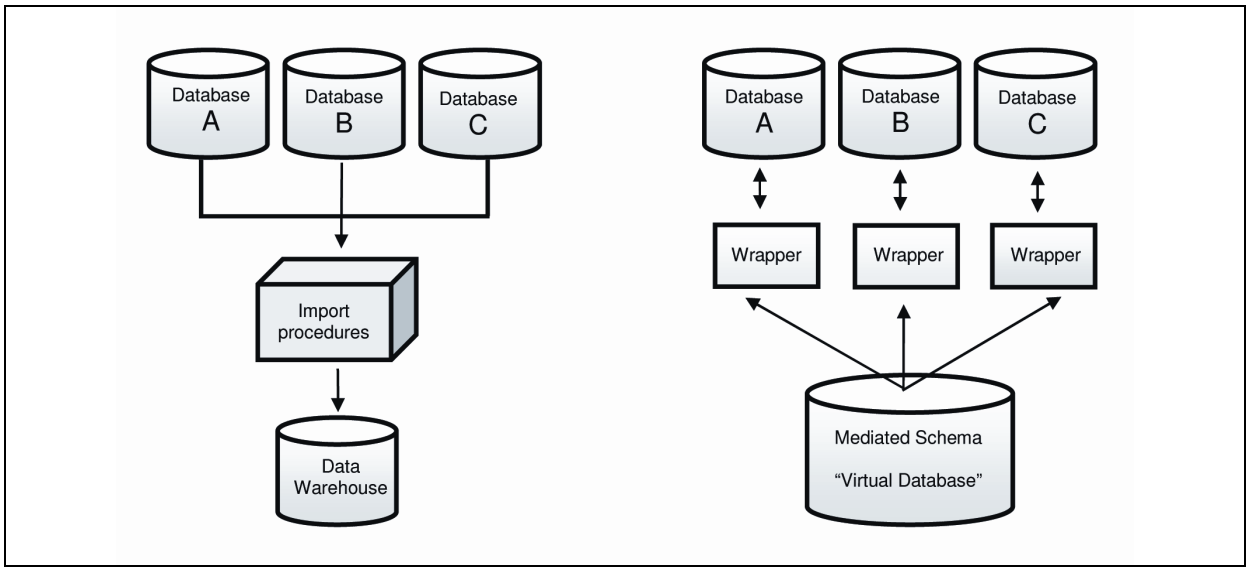

Figure 3.7: Data integration strategies. [Left] Data are extracted, transformed and loaded into a unique schema. [Right] A "virtual" database is created, which is laid over the existing ones.

More recently, the trend in data integration has leaned towards loosening the coupling between data. A uniform query interface is instead provided over a *mediated schema*, thus transforming a query into specialised queries over the original databases.
The advantage of this approach lies in the simplicity involved in answering queries that are issued over the mediated schema, which functions as a view at an intermediate level between the sources and the user. Unfortunately, it is necessary to rewrite the view for the mediated schema whenever a new source is integrated, or an existing source changes its original schema. A schematic example is presented in Figure 3.7.

3.3 Spatial data integration: issues

All definitions already seen about data interoperability can also be applied to spatial data, since what makes spatial data different is the presence of a geographic reference and the annexed topology, which can nevertheless be stored together with the other standard attribute values.
Organising data into information layers in a geographical information system – or, more generally, in a *site* – is one of the most common ways a user can merge and combine information. Data are structured into maps where each layer describes a particular aspect of the real world. This kind of data organisation attempts to achieve efficiencies in data storage and manipulation, but spatial queries involving data on different layers can be computationally expensive, especially when data are distributed among remote databases.
Syntactic and semantic heterogeneity are in fact still valid concepts. Digital terrain models can be considered as an example of a variety of data models applied to *one* of all possible spatial products, but many considerations still apply if they are generalised to other spatial datasets.
Horizontal and vertical fragmentation are still present, although for geographic databases another distinction is preferred. According to LAURINI (1998) inconsistencies resulting from the merge of different geographic datasets can be grouped into two groups: with the term *layer fragmentation*, discrepancies are meant that originate in case of datasets covering the same region but containing different feature classes, for example a DTM and a topographic map; *zonal fragmentation* refers instead to errors which originate from datasets containing the same information but covering spatially disjoint regions, for example the DTMs belonging to different neighbouring countries. Because of the inexact matching at the boundaries some difficulties occur in order to ensure geometric and topological continuity between the different datasets. Layer fragmentation and zonal fragmentation are indeed very frequent, considering that different institutions have different kinds of information on different zones and they might need to integrate them (Figure 3.8).

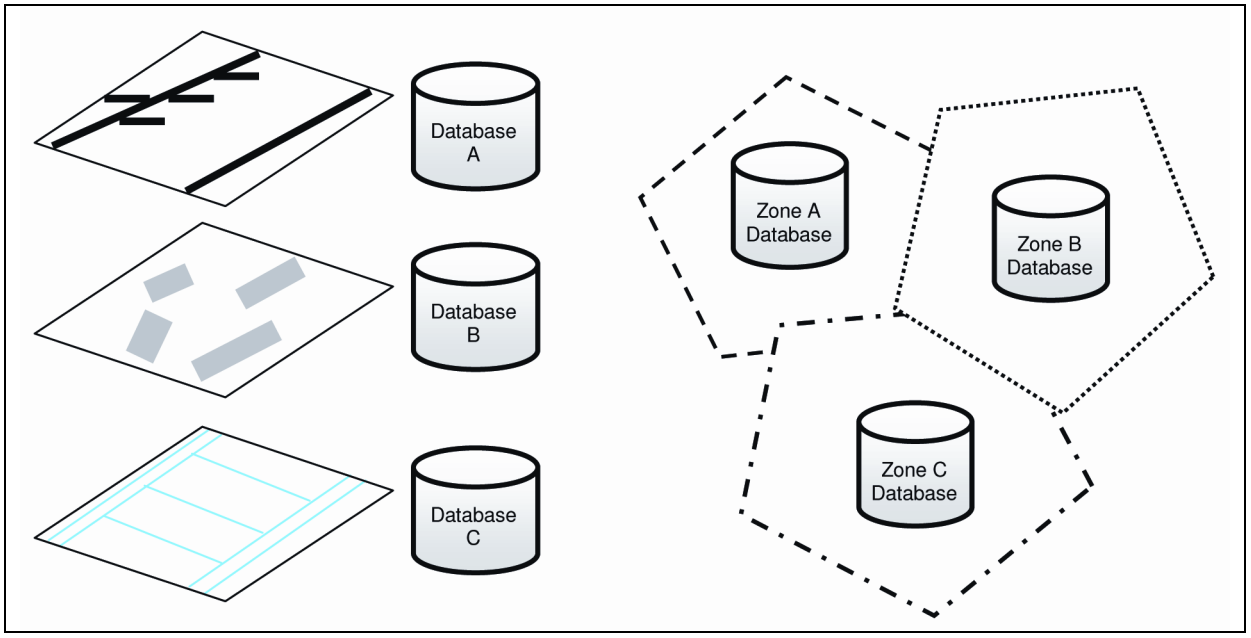

Figure 3.8: Examples of layer and zonal fragmentation among spatial datasets.

Structural and semantic incompatibilities can be exemplified by some of the ways a road can be represented: in a raster model its surface is represented in a discrete way by regular cells; in a vector map the same object can be represented by:

- a single, two-dimensional polyline with one or more attribute fields for width and height values,
- two polylines within a certain offset distance,
- an areal object (a polygon).

Moving to the third dimension, more possibilities are offered, e.g:

- a 3D polyline,
- a patch of 3D polygons, with or without attribute fields.

It is clear that interoperability cannot be achieved just by rough conversion of the spatial data structures, since even more discrepancies exist for geographic information in addition to the conventional ones.

3.3.1 Problems from the "surveyed" world

This relatively general category embodies problems related to data that results from any kind of surveying techniques. These data are generally managed by means of geoinformation systems. Objects may be acquired, for example, from images, laser scanners, terrestrial measurements or remote sensing techniques. With this kind of information, the exact spatial reference of every geoinformation layer is a common characteristic and it is of great importance.

Every map layer can have different geodetic Coordinate Reference Systems (CRS) in terms of coordinate systems or different projections; tools exist for coordinate transformations and permit to combine heterogeneously georeferenced objects. However, a conversion from a coordinate system to another can be a source of errors. The arithmetic precision of computers, which is limited, should not be underestimated when storing and processing geometric features. Different values can be associated to

the same item which is located in different sites, or different spatio-temporal sampling can lead to discrepancies.
Map layers can contain homologous objects which are then represented on multiple datasets. The same object can result from measurements carried out with similar sensors having different properties as resolution and accuracy, or with different sensors. It is not uncommon, for example, that the geometries of city models are obtained from both terrestrial (walls and façades) and aerial (roofs) surveying methodologies. One more possibility is that the same object originates from different modelling paradigms: as a B-Rep model from a surveying approach or a as CAAD (Computer Aided Architectural Design) object – this specific case will be discussed later in the next paragraph.
Accuracy and scale of a map layer are aspects to deal with. Although each layer may be reasonably accurate within its scale limitations, differences in input errors between heterogeneous maps might lead to mismatches, as well as different resolutions or level of generalisation.
Any coordinate is in fact coupled with measurement errors: two maps seldom match along common boundaries, and the same holds for zonal fragmentation, where boundaries between two databases do not match planimetrically. Similar thoughts apply to the third dimension, since neighbouring DTMs may have different height values associated to the same planimetric position.

3.3.2 Problems from the "constructed" world

With the recent growing importance of multi-format, multi-purpose projects like Google Earth or the progressive diffusion of three-dimensional city models, another set of integration problems has gained attention. Two different domains have been brought in tighter connection: the GIS domain – belonging to the "surveyed" world – and the AEC/FM (Architecture, Engineering, Construction and Facility Management) domain, whose major difference consists in the way data are produced.
Nevertheless integration of 3D CAAD and GIS models is increasingly being carried out, but some major differences still need to be overcome. The following points can be better understood bearing in mind, as an example, the ways a building can be represented.
Regarding geometry, the CAAD building is based upon Constructive Solid Geometry (CSG), which means that its geometric shape consists of volumetric and parametric primitives that are combined using the operations like union, intersection and difference. This kind of representation reproduces the construction process in disciplines like architectural design, but differs completely from a B-Rep model, typically obtained from a surveying campaign, which is defined as the collection of all surfaces enclosing its volume. Currently, in 3D GIS programs only B-Rep models are implemented, thus incorporating CSG models can be a difficult task due to the lack of proper editing tools which follow the constructive approach of CSG.
Another point of divergence consists in georeferencing: architectural objects are generally first created using local Cartesian coordinate systems and then, optionally, georeferenced by the explicit application of an affine transformation which translates,

scales and rotates all objects in a project to their world coordinates [KOLBE & PLÜMER, 2004].
The advantages rely on the possibility to use prototype-objects which can be created once and repeated at different locations by a simple transformation. This feature is inhibited in GIS, where it is necessary to explicitly represent the absolute coordinates of every object. In the case of a city model, for example, every streetlight, traffic light, tree and other vegetation object must be modelled singularly.

3.3.3 Problems from the "bureaucratic" world

Briefly, one last category of possible problems for geodata integration will be mentioned here. Geographic databases can contain heterogeneous data which belongs to different owners. Therefore, any user must take copyright related aspects into account since the access right to data might be limited or denied *a priori* by the owner.
Moreover, when integrating data, problems might arise in case of update operations (who is the owner of the update, the user who found an error or the data owner?), or in case data are mixed (who is responsible for a derived map which mixes information from different existing models?).
Finally, one limit to integration could be caused by accounting issues among different institutions. If datasets belonging to one institution and an application belonging to another one are being used, who, how, and how much is going to be paid?
It is therefore clear that these and other similar problems can be solved only if some consensus exists among all institutions sharing data and application.
In 1998 LAURINI stated that the only solution to these problems would be "to set up or negotiate an inter-organisational protocol which has to be signed by all multi-database partners. A nice possibility is to create a sort of agency in charge of enforcing this protocol".

3.4 Integration strategies

In the previous three paragraphs problems impeding spatial data interoperability have been presented. Some solution strategies and approaches will be introduced here. Following a reverse order when it comes to the solutions, interoperability for both programs and data can profit greatly from the definition of common specifications. Great improvements have been achieved in recent years thanks to the diffusion and adoption of comprehensive set of standards up to ISO level.

3.4.1 Standardisation

Regarding the GIS world, the Open Geospatial Consortium (OGC) is an international industry consortium of 385 companies, government agencies and universities participating in a consensus process to develop publicly available interface standards [OGC].
Through its OpenGIS standards it is possible to develop interoperable geo-enabled solutions such that the spatial information and services are accessible from all kinds of applications. In particular, an XML-based standard for the representation of geo-objects

has been released with the Geography Markup Language 3 (GML3), which itself is based on different 191xx ISO standards.
In the AEC/FM domain, Industry Foundation Classes (IFC) have been developed, and their technical specifications are maintained by buildingSMART International, formerly International Alliance for Interoperability (IAI).
IFC are an open and standardised specifications, registered as an ISO Publicly Available Specification, ISO/PAS 16739 and currently in the process of becoming an official International Standard ISO/IS 16739.
They provide support for data exchange among AEC/FM software applications, description of spatial elements, building elements and other components that constitutes a building or a facility, as well as the relationships of such components between each other.
A building and its components are modelled both in terms of geometric and semantic properties: definitions exist for buildings, storeys, walls, roofs, rooms, stairs, etc.
It is evident that international standards have been spreading and are being adopted more and more, but it also emerges that they result from the work of distinct organisations which have cooperated with different ISO Technical Committees. A conversion between models can for example lead to problems, since GML3 is restricted to B-Rep, while IFC supports both CSG and B-Rep geometry but has only little CRS support.
Although cooperation projects between the Open Geospatial Consortium and buildingSMART International exist, it can be expected that all interoperability problems between the GIS and the AEC/FM world will not be solved soon, due to the substantial differences between the two approaches.

3.4.2 Semantic modelling

One of the keys to the integration of geo-objects from different domains is represented by the semantic modelling. If geodata reside on fragmented datasets, the integration process must deal first with the identification of the various representations of corresponding objects, then their geometric inconsistencies can be corrected and/or homogenised.
As soon as an ontology has been developed for every domain, all terms, the meaning of different objects and their interrelationships are thus defined. The mapping of two different ontologies allows then to interpret objects from one domain in the context of the other[1]. Therefore, the more information is provided by the semantic layer, the less ambiguities remain for the geometric integration.

[1] Such ontology already exists for the AEC/FM domain, and the future IFC release IFC2x4 (due by spring 2010) is expected to include, among other improvements, the possibilities of links to GIS models [IFC].
In the GIS domain, an equivalent ontology has been defined in the context of city modelling by CityGML. CityGML represents an information model for the representation of 3D urban objects. It has been adopted as an official OGC Standard in August 2008. It is implemented as an application schema for GML3 and it defines the classes and relations for the most relevant topographic objects in cities and regional models (buildings, streets, water bodies, vegetation, terrain, etc.) with respect to their geometrical, topological, semantic and appearance properties. It includes generalisation hierarchies between thematic classes, aggregations, relations between objects, and spatial properties [CITYGML]. BENNER *et al.* (2005) have

STADLER & KOLBE (2007) discuss the importance of coherently structured semantics and geometry in a model and analyse their correspondence, referred to as spatio-semantic coherence. Essentially, the more relations can be mapped from the geometry hierarchy to the semantics hierarchy (and vice versa), the higher the level of coherence is.
They distinguish up to six categories of spatio-semantic coherence, ranging from models that are completely defined geometrically *or* semantically, to models structured both semantically *and* geometrically. In the latter, the structure is considered fully coherent since all semantic components correlate to geometric components on the same level of the hierarchy.
Furthermore, they examine the benefits of spatio-semantic coherence and the role it plays in the integration of distributed datasets. Three objects are considered: a terrain model, a path and a building model.
In case of fully coherent structural modelling, all semantic information can be used and it helps to identify objects by their properties and to define common points for geometric adjustment. Since the location and the function of the building's door are known, the terrain can be aligned to the doorstep. The path can be embedded into the terrain and correctly connected to the doorstep, as shown in Figure 3.9, left.
In case the spatio-semantic coherence is only partial, simple integration strategies can be implemented, but they might rely only on approaches that depend on data quality (i.e. data acquisition methods and their accuracy). A possible integration strategy could consist in adjusting a less accurate terrain model around a more detailed building's ground surface, however the limited spatio-semantic coherence could hinder a correct result if the building has a cellar: the terrain would not know how much the building must actually sink and at which height to connect. Similarly, the path could be embedded into the terrain model, but its connection to the doorstep would not be possible due to the lack of information about the exact location of the door.

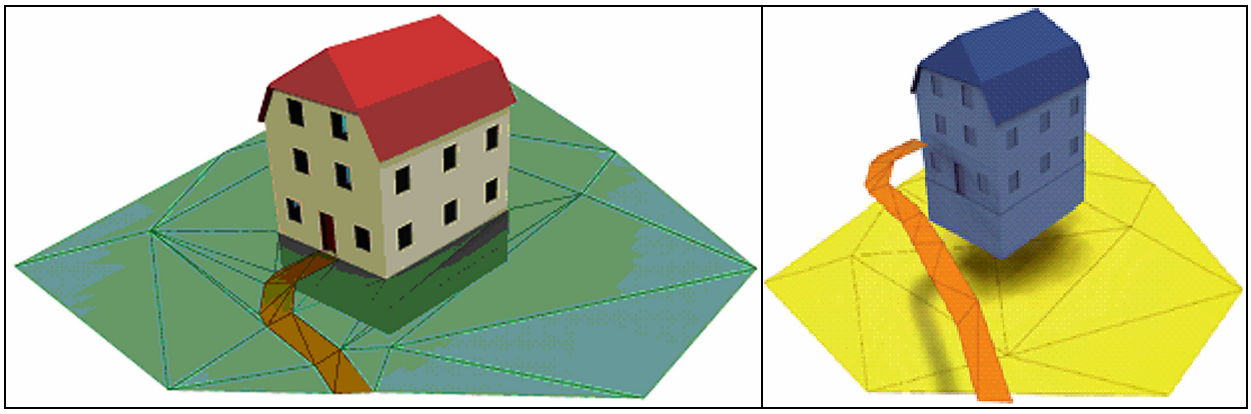

Figure 3.9: Examples of semantic data integration between three datasets representing a terrain model, a path and a building. [Left] Correct data integration by means of full spatio-semantic coherence among the datasets. [Right] Erroneous/missing integration due to the lack of semantics. Images courtesy of A. STADLER, TU Berlin.

presented how to map the semantics and CSG geometry of an IFC model to a spatio-semantic coherent B-Rep representation in CityGML.

If there is no semantics at all or semantics and geometry are uncorrelated, integration can be carried out purely on the basis of geometric shape: the harmonisation cannot rely on connecting features to correlate geometries. The information about the terrain, the path and the building is not modelled and, generally, it cannot be presumed automatically from the data themselves. Therefore, an automatic integration process would fail to connect the path to the door, and the building would probably float over the terrain, because it is not known that the shapes represent a path, a building and a terrain (Figure 3.9, right).

An example of semantic integration between a DTM and an existing topographic vector datasets has been presented by KOCH (2005). His approach does not use only geometric, but also semantic information in order to enhance DTMs using data coming from a two-dimensional topographic vector dataset.

The basic idea is that, although the topographic objects of the vector data are often only two-dimensional, some objects contain indeed implicit height information. This extra information can be contained in form of attributes, which might define the road width, type, etc. Another possible example could be a lake, which represents a horizontal plane, and all height values inside its area extents must be the same. Thus, the intrinsic height information of these geographical features can be employed through a proper set of constraints: the slope of a road must not exceed a certain maximum value, the banks around a lake must be higher than the water level, etc.

KOCH'S approach consists of several steps. In the first one, two-dimensional features are classified according to their semantic characteristics. One class contains for example objects which represent horizontal planes (a soccer field, a lake, etc). Another class contains objects representing tilted planes: roads for instance present slopes both in driving direction and perpendicular to it. Moreover, linear object like streets and railways are mostly modelled by lines in topographic datasets, thus they must be buffered using additional information before integrating them with a DTM.

For each object class a proper constraint can be formulated: in case of a horizontal object like a lake, all height values z_i of the i points inside the polygon delimiting the feature must ideally obey following equality:

$$z_i = z_{lake} \qquad [11]$$

In an analogous way, all j points belonging to the polygon itself must be interpolated from the neighbouring DTM triangle vertices (l,m,n), so a similar constraint can be given:

$$z_j(x_j, y_j, z_l, z_m, z_n) = z_{lake} \qquad [12]$$

Also the neighbouring terrain around the lake is influenced: the banks of the lake shore are not allowed to be lower then the lake itself (otherwise water would flow out) and, secondarily, this influence decreases with the distance from the object.

In case of roads, thus objects belonging to the class of tilted planes, constraints can be expressed by means of inequalities. For a road segment with endpoints *a* and *b*, the maximum slope *s* in driving direction can be defined as

$$\left| \frac{z_a - z_b}{d_{ab}} \right| \leq s_{\max} \qquad [13]$$

where d_{ab} is the horizontal length of the segment. In a similar way, height values for bridges are related to other objects as a river or another road, so that their values can be constrained as

$$z_{brigde} > z_{river} \qquad [14]$$

Once the constraints have been formulated, a constrained DELAUNAY triangulation is carried out using the input DTM points and the geometries of the classified objects. The height values of the inserted 2D objects are initially interpolated from those of the DTM. Finally, the previously defined equality and inequality constraints are considered in an optimisation process. For this purpose, the approach is based on a inequality constrained least square adjustment algorithm.
Some simplifications are made: the planimetric coordinates of the topographic vector dataset are introduced as error-free, and those of structure elements are not considered in the adjustment process, potentially leading to deletion of objects inside area based objects. Ultimately, a requirement is that the neighbouring terrain morphology of the objects be indeed considered, but improvements of the heights must be small. The implemented algorithm leads to improved height values and to integration of the input datasets with regard to the initially defined requirements, see Figure 3.10.

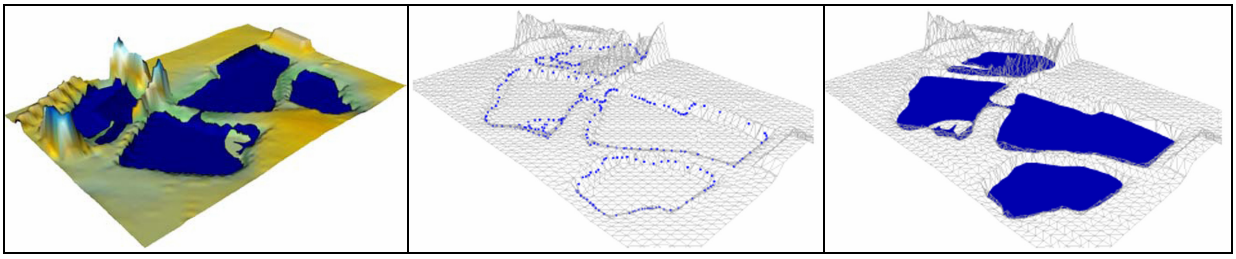

Figure 3.10: Example of semantic integration between a DTM and a topographic, two-dimensional, vector dataset. [Left] Without considering semantics, the lake reaches over its banks. [Centre] The TIN of the DTM, and the topographic vector data draped on top of it, with interpolated heights. [Right] TIN of the integrated datasets using semantics: the lake is contained correctly inside the banks. Image source: [KOCH, 2005].

3.4.3 Geometric homogenisation

From the discussion in the previous paragraph it has become clear that semantic information can facilitate and improve the quality of data integration, on condition that its relation to the geometry is known and the spatio-semantic coherence is full. Only then known relations between specific feature classes can be used, in order to identify

corresponding objects, to find connecting elements and to provide threshold values for the geometric homogenisation.
Some questions still remain partially unanswered: what happens when semantic integration is not possible? Which strategies can be considered for the “pure” geometric harmonisation of fragmented datasets?
As mentioned before, errors may be present in fragmented datasets. In case of layer fragmentation, multiple representations of the same object might be found across different layers. Due, for example, to surveying errors, homologous buildings belonging to database A and Bmay not perfectly coincide when overlapped, showing geometric discrepancies in form of sliver polygons, see Figure 3.11.

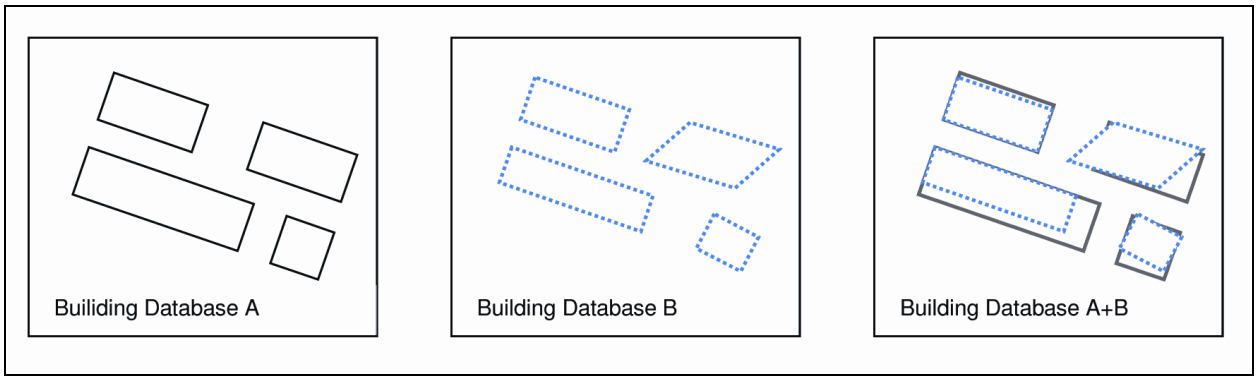

Figure 3.11: Geometric inconsistencies with layer-fragmented datasets.

As yet, the identification of a generic and automated overlaying process is an unresolved issue and a present (and future) subject of research. However, provided that the matching between two geometric objects has been carried out in some way (e.g. by the user, by defining common tie point, or semi-automatically by means of fuzzy logic), one possible solution could be to correct *all* databases. This operation could be, however, inappropriate and, actually, “virtually” impossible, since homogenising (i.e changing) the original coordinates in database A and B could lead to further problems whenever a third database C, a fourth database D etc. are added.
A more common solution consists in leaving the original databases untouched and performing the geometric homogenisation only at query level. Generally, a real-time *rubber-sheeting* technique is used, where the term rubber-sheeting stands for a transformation that allows some geographic elements such as points or lines to be force-fitted by means of a sort of elastic transformation. This transformation is not necessarily linear, and objects could be distorted. The image of rectangular building could be no more rectangular, for example, unless specific constraints are taken into account.
Similar strategies can be applied in case of zonal fragmentation, too. Errors at the boundaries of two datasets are frequent, since there are often mismatched elements. As a solution, a rubber-sheeting transformation can be performed. If constraints are added, buildings can be kept rectangular, roads aligned, and so on.

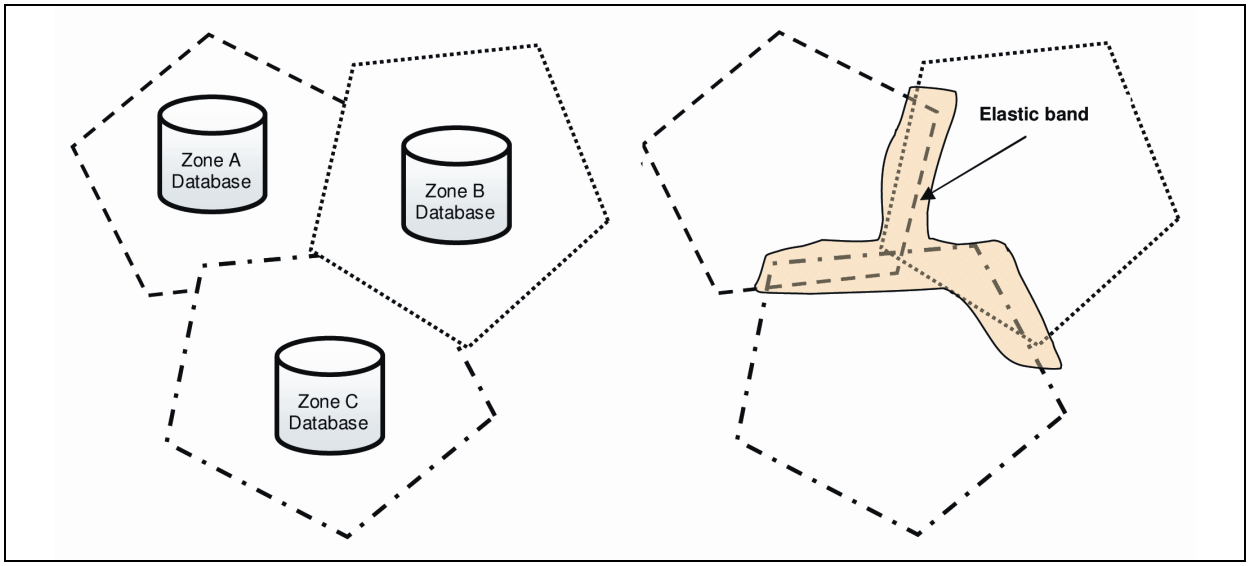

Figure 3.12: Elastic band over the boundaries of adjacent datasets.

The idea is that a proper “elastic band” is defined along the boundary (Figure 3.12) upon which the transformation will be applied. This band represents therefore a *transition zone* between the two datasets and the “elastic band” transformation should take place progressively from the external limits of the band toward its middle. In other words, the maximum geometric distortions should be located in the middle of the transition zone, with a seamless transition to the original dataset along the borders of the transition zone, see Figure 3.13.

It shall be said that so far no specific mention has been made of the type of mismatched elements between the fragmented datasets. The examples in Figure 3.11 and in Figure 3.12 refer to two-dimensional objects (lines, polygons), but the same holds also for the third dimension: as mentioned before, an example could be the height differences at the boundaries of two neighbouring digital terrain models.

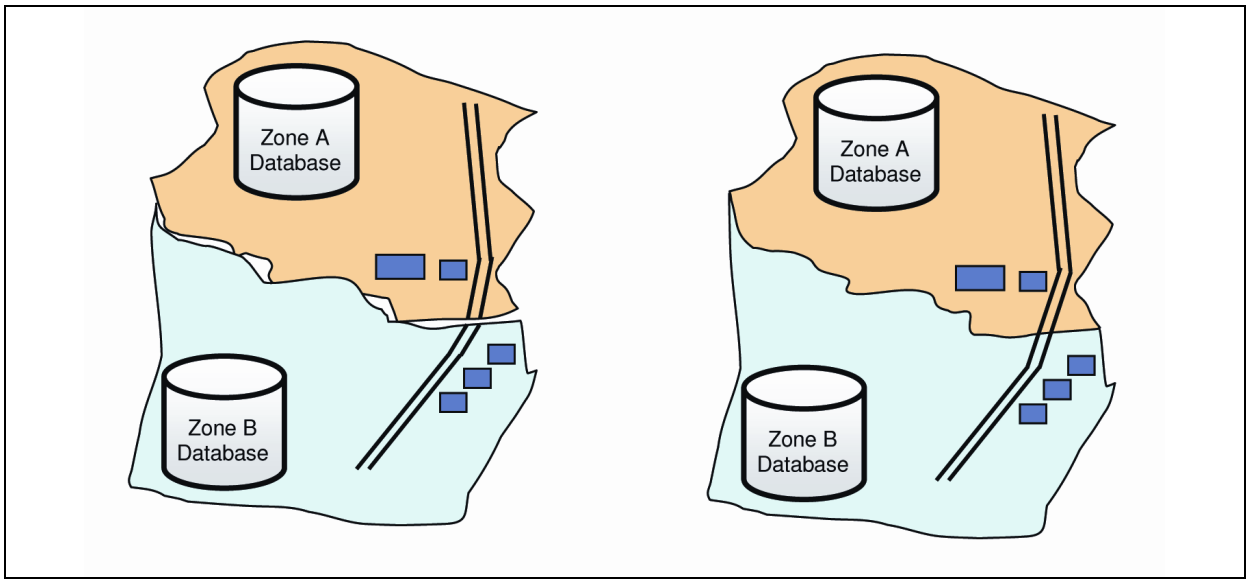

Figure 3.13: In zonal fragmentation some geometric corrections are necessary at the boundary.

When it comes to the size of the transition zone, several application-specific rules of thumb exist. However, an omni-comprehensive, general approach has not yet been presented and is therefore subject of research. Sometimes the extents of the elastic band are set by the user by means of a visual interface. Alternatively, a common, fixed value for the size of transition zone can be defined without no other rules other than a certain offset distance from the boundaries. The risk connected to such an arbitrary delimitation is that an important object can be cut in an elastically transformed part and an unchanged part. If this effect cannot be considered negligible, then the user has to define another limit.

3.5 DTM integration

In this paragraph data integration regarding fragmented digital terrain models will be brought into focus. The idea behind DTM integration is simple: a combination of multiple independent datasets should result in a new dataset containing more information and details than any input dataset, allowing for a better description of the geomorphology of the modelled surface.

Every dataset, even of lower resolution, can be useful for the production of higher quality DTMs. Potential data sources are not only all available DTMs in the study area at local, regional, and global scale, regardless of their spatial precision and accuracy, but also other information layers such as boundary points, databases of buildings, cadastre databases, geodetic network points, hydrological and road network measurements, etc. In addition, integration of these data layers may simplify the description of the temporal modifications on the terrain surface due to landslides, highway construction, etc.

In a similar way to other spatial datasets, integration of digital terrain models must face both layer fragmentation and zonal fragmentation. Sometimes the terms vertical and horizontal mosaicking can also be found, respectively.

3.5.1 DTMs and layer fragmentation

When it comes to the layer fragmentation, there exist a wide spectrum of approaches. Here only a few will be mentioned that are related to the work described in this thesis. Some deal with integration of a DTM and a two-dimensional topographic vector dataset: the height values are taken from attributes, or, for example, grids and triangles are jointly used to enrich the DTM with information from topographic objects. LENK (2001) gives an overview about existing methods.

FELUS & CSATHÓ (2000) describe how to obtain a DTM from patches of different local DTMs, starting from five different height models with different formats (contour lines, grid points, GPS points), different density and different height accuracy, in the range of 0.01 to 100 m. They propose a two step process for data integration.

First, they create a primary DTM from two low resolution datasets, USGS 1:50000 and 1:250000 contour maps. They perform an overlay and clip the common area from the 1:250000 dataset, which is substituted with the 1:50000 data. After setting a grid size of 100 m, they interpolate the whole area using a thin plate spline technique, which allows

to take abrupt changes in terrain such as streams and ridges into account. In order to improve data continuity along the clipping boundaries and to have a seamless elevation grid, they define a transition zone of 1 km across the boundaries.
In the second step, they refine the base gridded DTM with more accurate elevation data in specific areas, by means of geostatistical interpolation.
With the diffusion of newer technologies many authors have concentrated on the aspects related to integration (fusion) of data coming from aerial laser scanning. REISS (2002) proposes a workflow in order to combine laser scanner data with additional photogrammetric and terrestrial measurements, in order to check the data obtained from laser scanning.
SCHENK & CSATHÓ (2002) deal with aspects related to merging aerial imagery and ALS data, too. They focus on the data alignment problem by using sensor-invariant features such as breaklines and surface patches, in order to achieve a more complete surface reconstruction.
PODOBNIKAR (2005) proposes a method, defined "weighted sum of data with geomorphologic enhancement", to sequentially combine different datasets. He first combines individual datasets according to their weights and then applies a geomorphologic enhancement. The weights depend on the quality of the spatial datasets.
In particular, a unique cell size for all data sources is first set, and this also corresponds to the resolution of the final product. Vector datasets containing linear and area features, or other data formats, are consequently transformed or converted.
Each grid point is denoted with a quality parameter (a random error σ), so that the weights w_1 and w_2 of two individual datasets are respectively

$$w_1 = \frac{1}{\sigma_1^2} \qquad w_2 = \frac{1}{\sigma_2^2} \qquad [15]$$

and the height values of the resulting DTM, derived from both data sources, are

$$z_{DTM} = \frac{w_1 z_1 + w_2 z_2}{w_1 + w_2} = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} z_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} z_2 \qquad [16]$$

In case of more than two datasets, height values are added to the previous ones iteratively. The associated random error is obtained according to HEUVELINK (1998):

$$\sigma_{DTM} = \sqrt{\left(\frac{\delta z_{DTM}}{\delta z_1}\right)^2 \sigma_1^2 + \left(\frac{\delta z_{DTM}}{\delta z}\right)^2 \sigma_2^2} = \frac{\sigma_1 \sigma_2}{\sqrt{\sigma_1^2 + \sigma_2^2}} = \frac{1}{\sqrt{w_1 + w_2}} \qquad [17]$$

Moreover, the error decreases on every iteration, as shown in

$$\sigma_{DTM} = \frac{1}{\sqrt{w_1 + w_2}} \leq \min\left(\frac{1}{\sqrt{w_1}}, \frac{1}{\sqrt{w_2}}\right) = \min(\sigma_1, \sigma_2) \qquad [18]$$

When combining the datasets, PODOBNIKAR proposes to start from data of the lowest quality and finish with the best ones. Since the obtained DTM tends to be smoother than the highest quality data source, at the end he performs an enhancement by applying the geomorphologic details of the most appropriate data source to each grid cell.
As shown in Figure 3.14, assumed that a proper filter has been found, a low frequency trend surface and the respective relative high frequency height differences ($\Delta z_1, \Delta z_2$) can be calculated for every dataset and for the resulting weighed DTM. If the dataset with subscript 2 is supposed to be the one with the best geomorphologic details, then its relative height differences Δz_2 are added to the trend surface of the weighed DTM.
In this way the final, enhanced DTM is produced, which is to some extent worse than the purely weighted one, but describes the high resolution variations in the shape of the terrain in a better way. Therefore, the optimal solution lies in between statistical quality and geomorphologic accuracy.

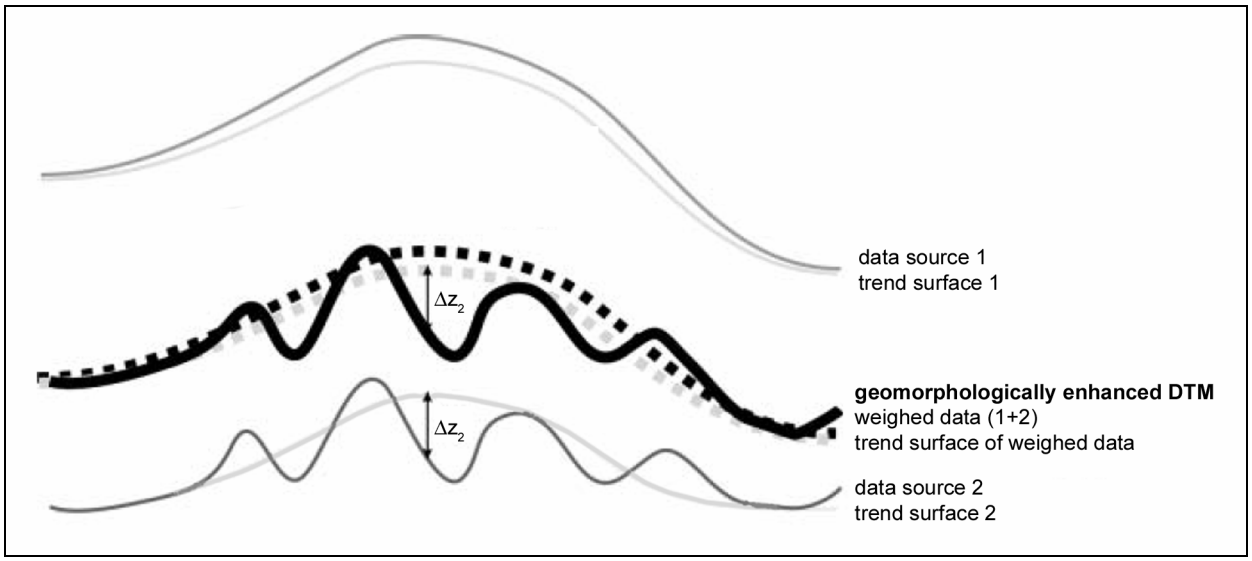

Figure 3.14: Geomorphologic enhancement according to PODOBNIKAR. Image source: [PODOBNIKAR, 2005].

3.5.2 DTMs and zonal fragmentation

When it comes to heterogeneous DTM datasets that are characterised by zonal fragmentation, they can cover spatially disjoint regions with the typical result of possible overlaps or gaps at the borders of the adjacent models.
A first fundamental problem resides in the so-called matching process. In the most general case, the goal is to establish the best transformation which aligns two datasets, so that one single common reference system is used and height differences along the boundaries are minimised. For the operation, also called registration, different techniques have been developed according to the DTM data structure, aiming at a seamless homogenisation of the entire mosaicked DTM.
If the DTM has a regularly spaced structure as raster cells or grid points, neighbouring datasets can be joined using image mosaicking procedures. This is the traditional method adopted in stereophotogrammetry where a feature on the ground is reconstructed from two overlapping aerial images. This requires the identification of

ground features in both images and the exterior orientation of the images. In case the input images have different resolutions, which nevertheless should be generally avoided, decisions must be taken regarding a over- or sub-sampling operation, according to the required output unique resolution. Height values are treated as colour information and a smoothing function can be used along the overlapping zones.

In case of three-dimensional point clouds originating from terrestrial and aerial laser scanning, the analogous task consists usually in treating a surface as a fixed reference model and matching the remaining datasets with the reference one. Since this operation can be complicated in most applications, a wide range of solutions has been proposed to address this problem.

To this extent many algorithms already exist. Working with laser scanner data, point clouds or meshes, the Iterative Closest Point (ICP) algorithm is one of the most widely used for geometric alignment. It first establishes a correspondence between pairs of features in the two datasets that are going to be aligned, based on proximity; then it estimates the best rigid transformation that permits to map the first dataset onto the second using a cost function, and finally applies it. The steps are iteratively repeated until convergence is reached, however the algorithm needs a good estimate of the initial relative positions. The basic algorithm, presented by BELS & MCKAY (1992), has been successively extended, and today several variants exist that affect all steps, from the selection and matching of points to the minimisation strategy. A comprehensive overview of surface matching strategies can be found in GRÜN & AKÇA (2005) and in MILLER (2008).

Registration models and the related issues are not in the scope of this work, however it suffices to underline a few aspects that play a role in the following chapters:

- the matching algorithms represent the optimal solution to the rotation, translation and scale operations between two overlapping datasets. However, they are not necessary exact: locally, height differences in the overlapping zones can be still found;
- with laser scanner datasets, point density in the overlapping zone tends to be higher than in the input models. Even if this is not necessarily a disadvantage, it represents indeed a type of discontinuity and it may lead to problems in the following texturing phase, so that a local data decimation is necessary [AGUGIARO *et al.*, 2008];
- decisions regarding the dimension of the overlapping zone are generally tied to rules of thumb (“the larger, the better”). In the case of the above mentioned merging operation, FELUS & CSATHÓ opt for a 1 km overlap, their explanation being that this value is about ten times bigger than the biggest height difference in the two datasets.

Although not strictly a problem of registration models, little literature seems to exist which deals specifically with “elastic functions” applied to already georeferenced datasets (DTMs), which indeed are affected by fragmentation. In such cases, the height values of a digital terrain model need to be properly “warped” in order to guarantee a sufficiently smooth transition. In the following, some ideas which have influenced this work will be briefly introduced.

WARRINER & MANDLBURGER (2005) discuss how to locally update an older DTM, obtained from contour lines and spot heights, by inserting newer, more precise data only at selected places, namely urban and floodplain areas. Their intention is to target only specific areas where high resolution data is most useful, and to avoid the creation of a new global terrain model.
Since newer data come from several sources, the authors describe an intermediate merge operation between two ALS and photogrammetry datasets, before the older DTM is updated. Let's call for simplicity DTM1 and DTM2 these two datasets coming from ALS and photogrammetry.
They have similar quality levels and they must be blended in such a way that there is a continuous transition from one model to the other. In order to obtain a smooth transition, WARRINER & MANDLBURGER define a tolerance band of constant width (i.e. LAURINI's "elastic zone") around either DTM1 or DTM2, see Figure 3.15, left. A weighted average of both height values is performed inside it.

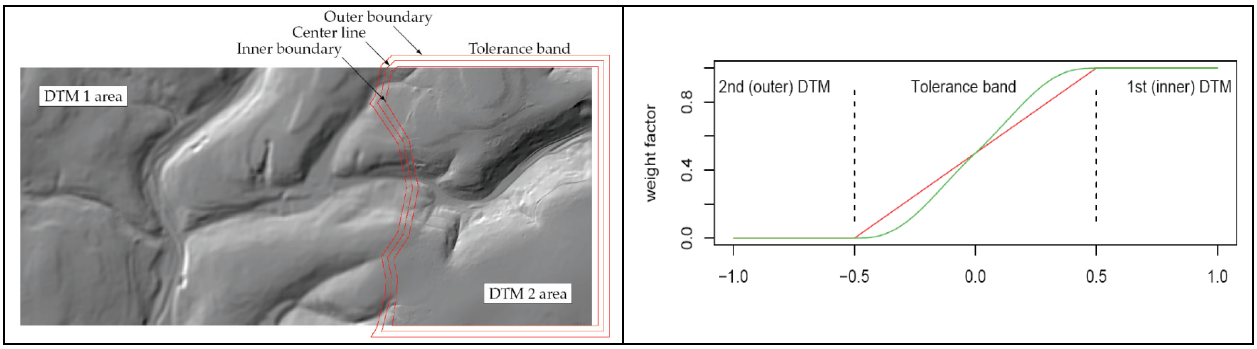

Figure 3.15: [Left] Integration of two DTMs by means of a tolerance band in which a weighed sum of the height values is carried out. [Right] Weights vary according to the distance from the centreline of the tolerance band. Image source: [WARRINER & MANDLBURGER, 2005].

The weights depend on the distance from the centreline of the tolerance band. The weight w_1 for the DTM1 gradually increases from 0 to 1, moving from the outer border (DTM2) to the inner border of the tolerance band (DTM1), as shown in Figure 3.15, right. In a similar way, the complementary weight w_2 for the DTM2 is defined as

$$w_2 = 1 - w_1 \qquad [19]$$

The height values can thus be calculated as

$$z = w_1 z_1 + w_2 z_2 \qquad [20]$$

According to the weight function used, different results can be achieved. WARRINER & MANDLBURGER discuss advantages and disadvantages of a linear, curved and jump weight function and come to the conclusion that the linear function represents the best solution for most applications. A curved function allows for horizontal tangential continuity at the edges of the tolerance band, at the cost of more abrupt transition in the middle of the tolerance band; the jump function, being actually an extreme case, does not allow in fact any blending at all. Some results can be seen in Figure 3.16.

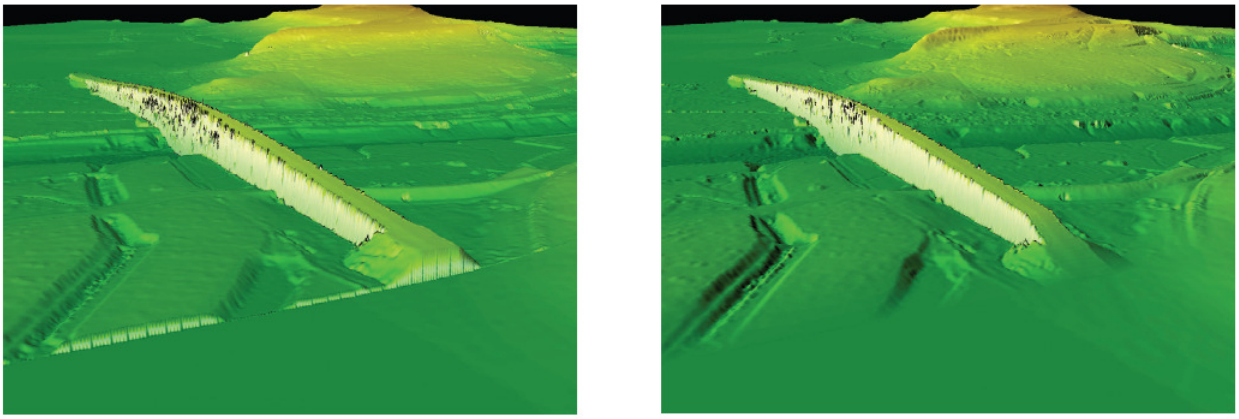

Figure 3.16: Experimental results of DTM integration using different weight functions. Jump function [left] and linear function [right]. In the left image, the two different datasets are clearly recognisable, due to the abrupt discontinuity caused by the jump function. Image source: [WARRINER & MANDLBURGER, 2005].

Another approach to overcome zonal fragmentation is presented by LATHAM & BURNS (2006). They deal with the problem of adjusting a terrain database for visual simulation. Given a set of precise objects (called targets), their goal is to smoothly fit the terrain in order to avoid targets floating above or hidden below the terrain surface. The approach consists in adopting a correction function over the whole terrain database that permits to correct height values of the terrain according to the gap existing between a target and the terrain surface. Examples for targets could be floating houses or other urban furniture objects like street lamps.
Since, potentially, even a small number of targets can affect large areas of the terrain database, there is the need to modify the surrounding terrain locally and in a smooth way, so that the integration of a target can influence only the nearby zones and leave the far away zones of the DTM unchanged.
A requirement for their model is that data need to be organised in layers, with a layer S containing the terrain surface, a layer T the targets and a layer C the correction values. The targets consist of vector features which may be represented by points, lines or areas, defined by three-dimensional coordinates. The correction layer C is a TIN obtained from the targets' coordinates defining the points of contact with the terrain, and it is independent of the underlying terrain database.
Provided a set of targets t_i having coordinates x_i, y_i, z_i in layer T, the height differences Δz_i between t_i and the terrain in layer S at (x_i, y_i) are computed as

$$\Delta z_i = z_i - S(x_i, y_i) \qquad [21]$$

These values are considered always correct and unchangeable; they correspond to the correction values to be applied to the terrain at (x_i, y_i) in order to guarantee integration with the targets, and they represent the height values of the TIN vertices in correction layer C.

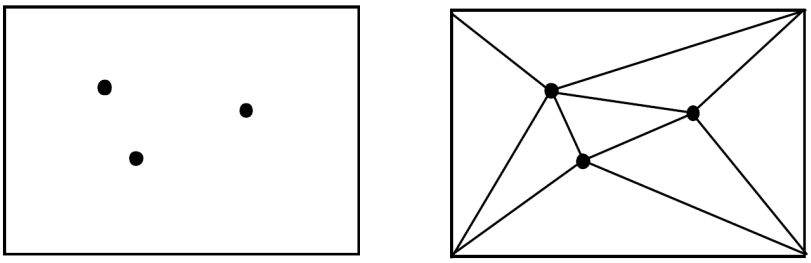

Figure 3.17: Example of three point targets and their triangulation, which defines the correction layer C. Image source: [LATHAM & BURNS, 2006].

Once the correction layer C has been created, the impact of the height changes needs to be softened in the surrounding terrain areas. Taking as an example Figure 3.17 (right), the four TIN corners represent the extents of the terrain database area whose height values are being modified. The target vertices and thus their height values are all specified; the vertices at the corners should be assumed to be zero. In case of small corrections (small values of Δz_i), a linear interpolation could work well; in case of large corrections, new long and unnatural ridge lines or other abrupt changes could be produced. Therefore, for every target t_i an influence function is introduced

$$f_i(x, y) = \Delta z_i \exp(-d^2 / k^2 \Delta z_i^2) \qquad [22]$$

with

$d^2 = (x - x_i)^2 + (y - y_i)^2$: Euclidean distance from a target,

k : slope constant (k is nominally 4. Increasing values of k correspond to smoother height profiles).

By this two-dimensional Gaussian function, some geometrical continuity aspects can be guaranteed: the apex is horizontal, which allows the flattened terrain to fit buildings. Moreover, for increasing values of d, height and slope values tend to zero, allowing the correction values to the terrain model to disappear gradually.

In case of multiple targets, the value of the global correction function $f(x, y)$ within each triangle of correction layer C is determined as the weighed sum of the influence functions $f_i(x, y)$ at the three vertices. Function $f(x, y)$ can thus be applied to each of the existing terrain vertices; there might not be enough points to accurately represent the smooth hills and valleys created by the influence functions. Therefore, each of the original terrain polygons and the new ones created by adding the targets are subdivided if necessary: a too large triangle is split into four triangles by connecting the midpoints of each edge. The process is then repeated on each of the four smaller triangles until all triangles are small enough to adequately approximate the surface, as shown in Figure 3.18.

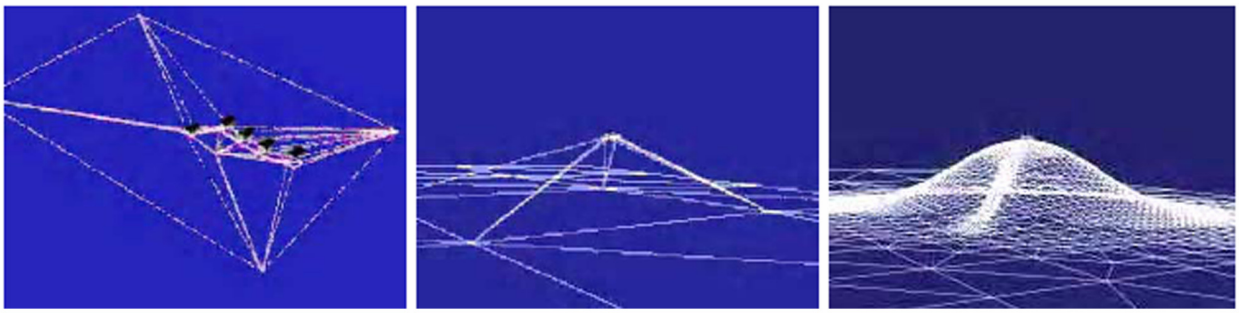

Figure 3.18: Steps in creating the correction function. [Left] Triangulation of correction layer C. [Centre] The correction function is applied to the target positions and [right] the smoothing of the correction function is performed by the influence functions. Height values in the figures are intentionally exaggerated. Image source: [LATHAM & BURNS, 2006].

3.5.3 Open issues

It has been shown in the previous paragraphs that the growing availability of spatial data is leading to new problems and to new solution approaches for data integration. In particular, although the modern technologies may favour creation from scratch of new, wide area and high resolution models, for some specific tasks this could not be possible or desired. WARRINER & MANDLBURGER and LATHAM & BURNS propose application-tailored solutions to solve zonal layer fragmentation by means of different warping functions, but some questions still remain open and will be schematically listed here:

- WARRINER & MANDLBURGER's method is thought for two DTMs of analogous quality level. What happens, for example, if qualities vary greatly?
- the transition zone has a fixed, regular shape: what happens in case of varying width? Can a rule be defined to set a proper width of the transition zone?
- what happens if a second dataset contains objects other than a DTM?
- LATHAM and BURNS provide a solution which leaves the vector targets unchanged and warps the surrounding terrain by means of a Gaussian function. The influence of the terrain distortion is, on the other hand, not distinctly delimited, although, in practice, the correction terms become negligible with the distance from the targets. Can this and WARRINER & MANDLBURGER's idea be somehow joined?
- what happens to data density and data precision inside the transition zones, after the geometric continuity has been guaranteed? Is the question trivial or should it be further investigated?

A further application example is given by CityGML, where multiple DTMs can already be stored and represented according to each Level of Detail [GRÖGER, KOLBE *et al.*, 2008]. DTMs can in fact be composed of heterogeneous models – TINs, grids, 3D breaklines and 3D mass points – and each component may be restricted to be valid in a specific region only by providing a so-called validity extent polygon. Validity extent polygons can have holes so that nested DTMs are allowed, see Figure 3.19. Nevertheless, a method for their integration along their boundaries has not been presented yet.

In addition, the field of cultural heritage could directly benefit from this kind of integration, since multiple heterogeneous and diachronically acquired models could be merged

together. Data coming from an archaeological excavation site is a well representative example.

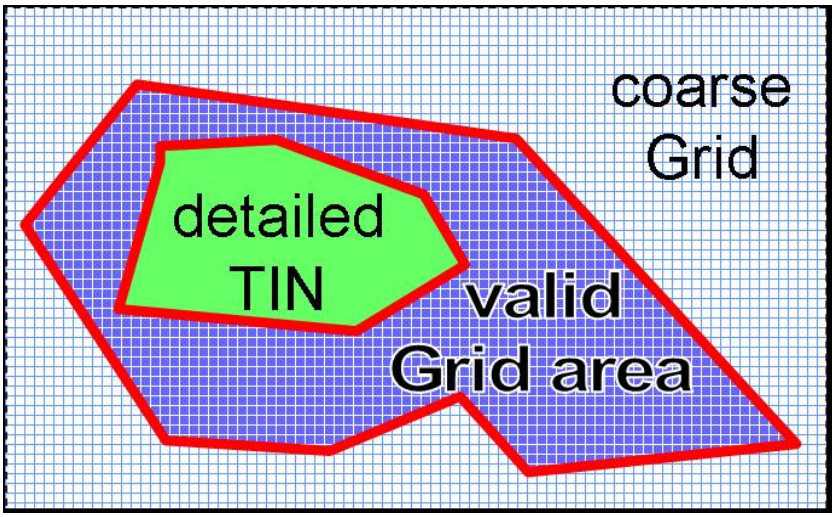

Figure 3.19: Example of nested DTMs delimited by a validity extent polygon (in red). Image courtesy of T. KOLBE, TU Berlin.

3.6 Problem definition and subject of the thesis

This thesis focuses on zonal fragmentation among different spatial datasets and deals with the open issues listed in § 3.5.3. Special attention is paid to the world of archaeological and architectural cultural heritage, but some issues can be easily extended to other disciplines and application fields.

Picking up again the idea of an elastic zone between neighbouring datasets, this work aims at the integration between a low resolution DTM and a laser-scanner-acquired, high resolution model. A proper transition surface and its underlying warping function are searched, which allow for a gradual transition between the two datasets.

Topological and geometric continuity must be guaranteed, and also a gradual transition in terms of data density is sought. Finally, the size of LAURINI's "elastic band" is allowed to be of any size, although its shape is still tied to the horizontal plane.

The characteristics of the transition surface, the working hypotheses and the developed method to calculate it will be described in the next chapter.

Chapter 4

TRANSITION SURFACE

This chapter describes how to obtain a transition surface which connects a high resolution model to a lower resolution digital terrain model. This surface must guarantee continuity in terms of geometry, topology and point density.
The initial assumptions and working hypotheses are first introduced; then the required characteristics of the surface are presented. In the end, the steps needed for its computation are described and discussed.

4.1 Initial assumptions

Zonal fragmentation has been defined in the previous chapter as one of the possible sources of inconsistencies when heterogeneous spatial datasets are integrated. Typical errors consist in cracks and overlaps at the borders of datasets containing the same spatial objects, but covering spatially disjointed regions.
This work focuses specifically on these errors and aims at creating a transition surface between the two models, which connects them without modifications at the actual high resolution object.
It is supposed that a high resolution object must be embedded into the surrounding DTM. It may have been obtained by different surveying techniques. The DTM data may represent a portion of the terrain surface that does not exist anymore and is therefore outdated. An example might be a DTM created for an archaeological site *before* excavation begins, whereas the high resolution model represents a pit that has been dug successively and is more recent in time, as shown in Figure 4.1, left.

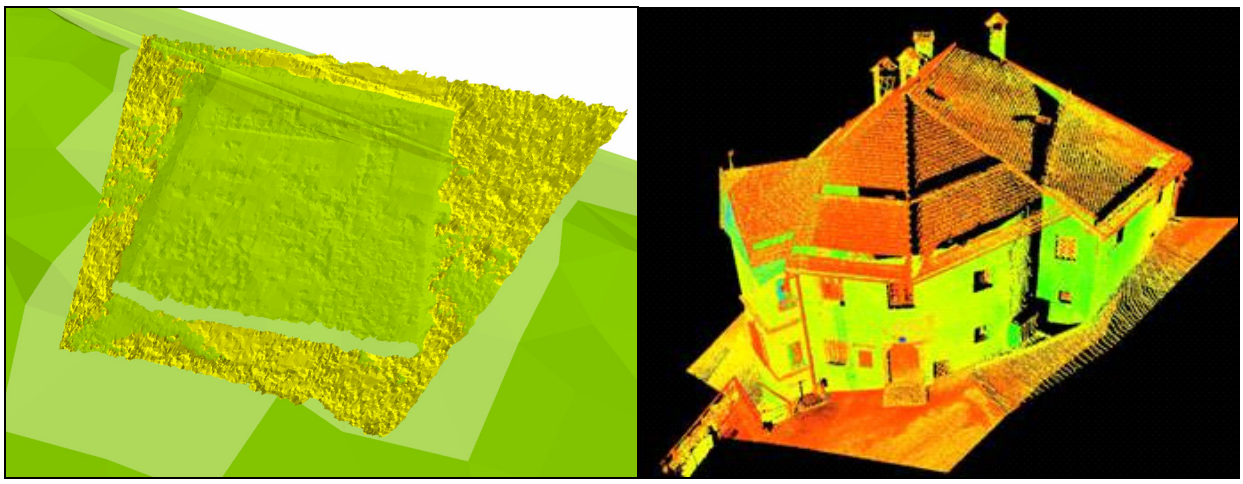

Figure 4.1: [Left] The high resolution model (yellow) needs to be embedded into a lower resolution DTM. The archaeological excavation pit is more recent than the DTM. Therefore some portions of the DTM refer to a surface which does not exist anymore. [Right] A laser scanner model of a house, which has been acquired with some extra data of the surrounding terrain (the "collar"). Right image courtesy of D. BRAGAGNOLO, Università di Padova.

The primary assumption of this work is that some extra information has been acquired around the high resolution object, so that its extents are actually larger than strictly needed. The extra information, which is actually common in laser scanner point clouds, can be thought as a sort of "collar"[1] and it is generally pruned during the point cloud editing which follows the acquisition phase (see, for example, Figure 4.1, right). Instead of deleting the data contained in the collar, the collar will be used for the modelling purposes while the actual high resolution object will be left unchanged.
Both datasets are assumed to be triangulated meshes, without blunders; they have already been previously aligned and georeferenced, so they share the same reference system – this implies that some previous data processing has taken place.
A simple example is shown in Figure 4.2: two synthetic models are represented. They are a low resolution DTM of a field (in green colour) and a high resolution object (in yellow and black colours), respectively. The two triangulated meshes share the same coordinate system, the same feature classes (triangular faces), but the transition from one dataset to the other is not correct due to geometric errors (triangle intersections) and topology discontinuities (gaps: the resulting model is not *one* orientable two-manifold surface).
As soon as the two models are overlapped, three zones can be distinguished and described as follows:

- Zone a), the high resolution object *per se*. In this zone, the low resolution DTM data refer to a portion of the DTM which will not be included in the final, merged model. Therefore they will be eventually discarded;
- Zone b), the outer, low resolution DTM features, which are not overlapped by any high resolution data and which will remain the same in the final model;
- Zone c), the overlapping zones, which contain data from both high and low resolution models. Points in this zone refer to the same object (the collar); some slight differences can be seen with regards to the height profile.

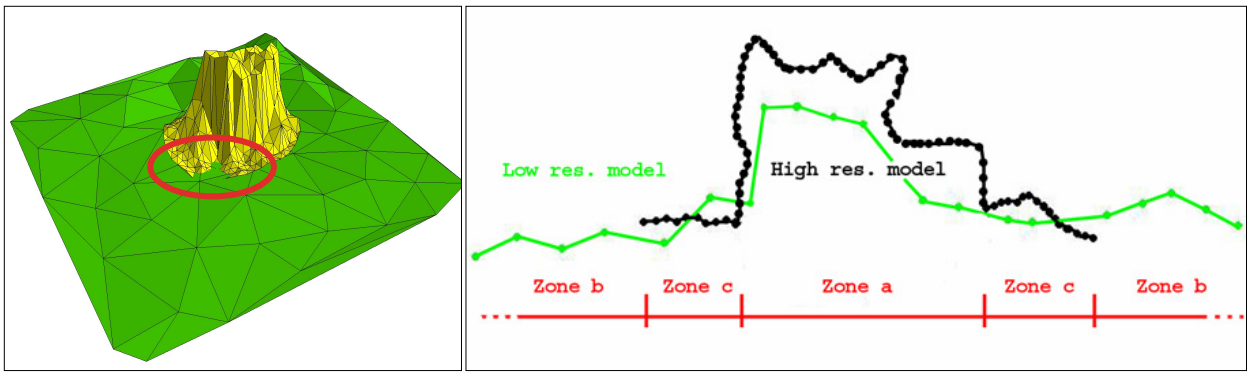

Figure 4.2: Geometric and topological inconsistencies due to layer fragmentation between two datasets: overlaps and gaps exist at the borders of two overlapped synthetic models: a low resolution DTM (in green) and a high resolution model (in yellow and black).

[1] The idea of a collar is already being explored by other authors: SCHMITTWILKEN *et al.* (2007) attempts to capture semi-automatically detailed building models aiming particularly at the building collars, i.e. the transition from facades to the digital elevation model, using hybrid data sources.

Zone C represents LAURINI'S "elastic band" zone – see § 3.5.2 – where the new connecting surface will be calculated. Two borders delimit it: the *inner border* joins the high resolution model to the collar and to the transition surface. It is a closed 3D polyline obtained from adjacent triangle edges of the high resolution mesh. This polyline separates the high resolution object from the collar and is intended to remain unchanged.
Similarly, the *outer border* connects the surrounding low resolution DTM to the transition surface. It is a closed 3D polyline which originates from the high resolution mesh (it is the outer limit of the collar) and is then projected on the z-axis onto the low resolution mesh. It is assumed that both 3D polylines delimiting the collar are given as input. Otherwise they can be interactively selected by the user, see Figure 4.3.

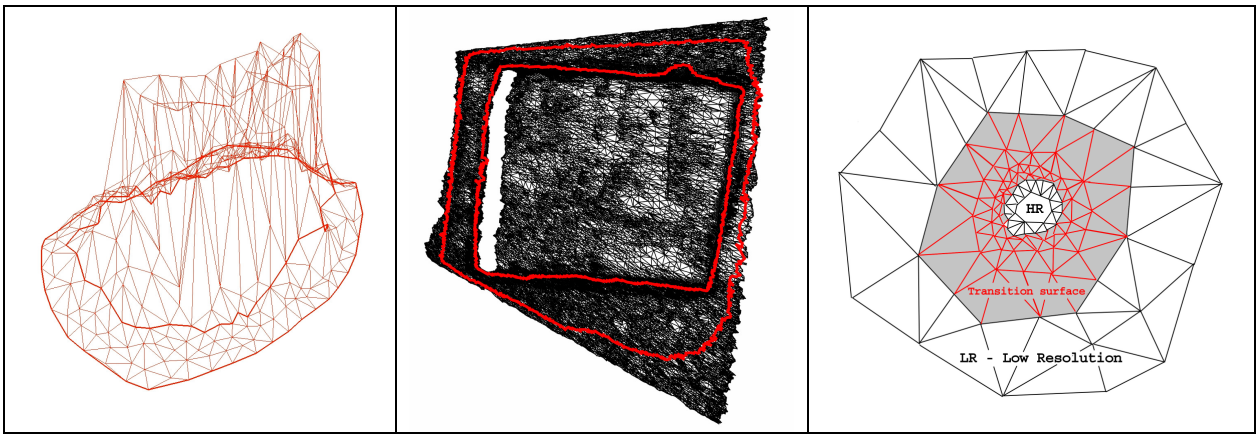

Figure 4.3: Every high resolution model is provided with an inner and outer border delimiting the collar. [Left, centre] Examples of a synthetic dataset and a laser-scanner-acquired model. [Right] Inside the transition surface (highlighted in grey) triangles of growing size from the inner to the outer border permit a transition also in terms of point density.

However, some restrictions apply to the connectivity of the meshes. They are required to be equivalent to a two-manifold surface. This means that every vertex is adjacent to a set of triangles which form a single, complete cycle around the vertex. In addition, every edge is adjacent to exactly two triangles. An exception is represented by the borders, where a chain of connected edges may be adjacent to one triangle only.
Non-manifold meshes are indeed quite frequent, as a result of the meshing process from point clouds or from a conversion from a CSG to B-Rep representation. A mesh can also be *mostly manifold* with some isolated non-manifold portions, or at the other extreme, connectivity information could be completely absent. In that case, the set of triangles is called a "triangle soup".
When it comes to the high and low resolution meshes needed for the transition surface, they are supposed to be two-manifold meshes *at least* in the overlapping zone. Strictly speaking, there is no real restriction to the type of mesh for the high resolution object, since that data will not be used: the point cloud does not even need to be triangulated. A simple point cloud or any other derivate product could be used instead. The only condition is that the "inside" is correctly connected to the collar in terms of topology and geometry.

In the overlapping zone, both meshes are required to be at least 2.5D, i.e. for each planimetric coordinate, only one height value is given.
No other requirements are done at this stage: there is no information about mesh segmentation or classification, no accuracy values of the surveyed points, etc.
Basically, only the three-dimensional coordinates of the mesh nodes, their topological relations to the neighbours and the extents of the collar are known.

4.2 Properties of the transition surface

As stated before, this work focuses specifically on the overlapping zone C. Inside it, a surface is sought which connects both models, preserves the characteristics of both high and low resolution datasets and, at the same time, allows for a progressive transition between them.
As a result, the sought surface will be itself a triangulation and, in order to ensure a geometrically and topologically correct connection, the following properties must be ensured:

- The inner and the outer borders mark the extents of the transition surface. Nothing happens outside the borders, i.e. no changes are made to the height values of the surrounding DTM and of the high resolution object, which therefore remains unchanged. This is meant to hinder that local height modifications can influence areas distant from the transition surface. In the case of a large DTM triangle with one vertex in zone C and the other two outside, a large triangle of this type must be therefore split.
- Both high and low resolution data contained in the overlap can be used to obtain the transition surface.
- At the borders there must be height continuity (in the following referred as condition C0). Triangles belonging to the transition surface must close gaps and allow for topology correctness.
- At the borders, tangential continuity (in the following referred as condition C1) must be taken into account. Triangles of the transition surface near the inner border must be “closer” to the high resolution dataset, while triangles near the outer border must be “closer” to the low resolution dataset. With regards to tangential continuity, the term “closer” means that a smooth, seamless transition is sought. Often higher degrees of continuity (C2 for curvature continuity, C3, etc.) can be added for better, visually more appealing results. However, C2 smoothness for the height profile is beyond the scope of this work and will not be considered.
- A transition in terms of point density is nonetheless required: triangles of growing size should characterise the surface while moving from the inner border to the outer border, and vice versa. A visual, qualitative example is given in Figure 4.3, right.

4.3 Procedure overview

The procedure to obtain the transition surface has been divided into successive steps. First, a selection of the data needed for the modelling purposes is carried out; subsequently, global and local parameters which permit an initial description of the data are calculated.
In the second step, the outer border is first projected onto the low resolution mesh, then a new constrained triangulation is performed inside zone C using the borders as breaklines.
In the next step, a height interpolation model is defined and applied, allowing geometric and topological continuity to be added.
Finally, a progressive mesh simplification is carried out, which allows to achieve in the so far calculated surface the desired transition also in terms of desired point density. Each single step will be described separately in the following paragraphs.

4.3.1 Setting up the environment

For the working environment, it has been decided to use as much already available open-source software and free development tools as possible. Implementation has been carried out mainly in the PostgreSQL object-relational database (version 8.2), in tight conjunction with its “extension” PostGIS (version 1.3), which adds support for geographic objects, allowing the PostgreSQL server to be used as a back-end spatial database for geographic information systems. PostGIS follows the OpenGIS “Simple Features Specification for SQL”.
Server functionality has been extended with one of the built-in procedural languages (PL/pgSQL), which allows for easy query scripting and grouping inside the database server, thus joining the power of a procedural language and the ease of use of SQL, but with considerable savings because there is less client/server communication overhead. Otherwise, every SQL statement should be executed individually by the database server.
Quantum GIS (from version 0.8 upwards) has been chosen as the visualisation and inspection tool, given its support for the PostgreSQL/PostGIS datasets. Furthermore, it offers an easy integration with GRASS GIS, which has been used for 3D views (on GNU/Linux), while proprietary ESRI ArcScene has been the 3D visualisation choice on Windows.
For mesh inspection and light editing tasks, MeshLab has been used (version 1.x), a free and open-source extensible platform to process and edit of 3D triangular meshes; advanced mesh editing has been done with RapidForm 2004 by Inus Technology. For few other specific tasks, external software has been used. Reference will be given in the following when needed.
For data import and export from and to PostgreSQL, the PLY file format, also known as Polygon File Format, has been adopted. It is a well-know, documented [PLY] and open format designed to store three-dimensional data from 3D scanners. It allows to store the polygons of a mesh together with optional attributes for the vertices or for the faces (e.g. surface normals, texture coordinates and data confidence values). Finally, it is one of the supported input/output formats of MeshLab.

Given the two input PLY files, one for each model and containing the vertex coordinates and the face topology, they are imported separately into two distinct PostgreSQL schemas (namely HR_MODEL and LR_MODEL). For each schema, data are saved into a POINTS and a FACES table.
A third schema called MODEL is also created, where most of the data processing takes place and all the needed tables, views and developed SQL functions are contained. A description of its objects will be done in the following, when necessary.
This three-schema-structure is intended to keep the original datasets separated from each other and from the to-be-calculated transition surface. Moreover, it is meant to ideally simulate two distinct data sources that may be accessed remotely.
It must be noted that PostGIS does have a topology support which consists in a schema model and the accessory functions to handle the topological element faces, edges and nodes; it is still in pre-alpha stage (as of July 2009) and it was not adopted.

4.4 Step 1: Data import and characterisation

Upon data import in PostgreSQL, the first step consists in selecting the input data in the overlapping zone from both datasets HR_MODEL and LR_MODEL and to copy it into schema MODEL of the database. If the three-dimensional inner and outer borders are projected on the xy-plane, they define an horizontal “annular”-shaped domain which defines the planar spatial extents of the query.
A distinction is still made according to the origin (data are saved in distinct tables); all high and low resolution triangles that do not lie within zone C or are not intersected by the inner or outer borders are ignored. Thus only a smaller number of triangles and relative nodes are used. The quantity of geometric features to be processed is greatly reduced through this initial simple spatial query. From now on, unless differently stated, it will be referred to the imported data in MODEL as the high and low resolution datasets.

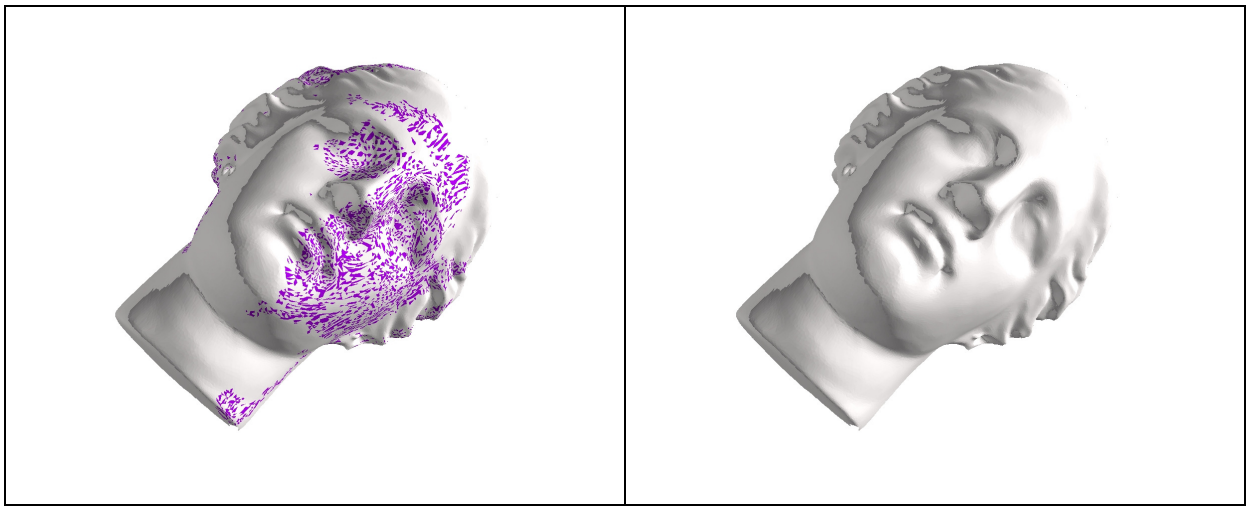

Figure 4.4: Topological check for mesh orientation. [Left] Violet-highlighted triangles are not correctly oriented: they are identified and flipped, in order to obtain an orientable surface where all face normals point outwards [right]. Tests for the implemented algorithm have been carried out using RapidForm’s three-dimensional model “Venus” and comparing results for reference.

Upon conclusion of data import, topology is completed with the creation and the population of the **ARCS** table. A topology check is performed on all faces of both meshes, in order to verify their orientation. For a 2.5D triangulated mesh like a TIN, a triangle is considered to have a positive orientation if its normal vector points ”outwards”; in other words, its vertices must be ordered in an anticlockwise order. In case triangles are wrongly ordered, they are flipped. The implemented algorithm – which works also with 3D surfaces – picks an initial "seed" triangle, checks its normal vectors and propagates this information across the TIN surface, so that all neighbouring triangles have always an opposite orientation along their shared edges, see Figure 4.4.

Afterwards all nodes and faces are characterised by means of local parameters whose goal is to give an initial description of the quantities listed below.
The parameters for every face/triangle are:

- indices of edges from table **ARCS**,
- indices of neighbouring triangles,
- area,
- normal vector $n:(n_v,n_v,n_v)$[2].

The parameters for every node/vertex are:

- number of neighbouring vertices, also called vertex valence,
- average, max, min three-dimensional distance to the neighbouring vertices,
- height difference between the meshes, obtained as

$$\Delta z = z_{highresolution} - z_{lowresolution} \qquad [23]$$

The height difference is calculated on a per point basis. High resolution points are vertically (z-axis) projected onto the low resolution faces and vice versa. A vertex $V:(x_v,y_v,z_v)$ and its projection $V^p:(x_v,y_v,z_v^p)$ on triangle ΔABC, with $A:(x_1,y_1,z_1)$, $B:(x_2,y_2,z_2)$ and $C:(x_3,y_3,z_3)$, share the same planimetric coordinates. The projected height value can be obtained from

$$n \bullet (A - V^p) = 0 \qquad [24]$$

which can be expanded and rewritten as

$$z_v^p = \frac{n_x(x_v - x_1) + n_y(y_v - y_1)}{-n_z} + z_1 \qquad [25]$$

If $n_z = 0$, then the triangle is parallel to the z-axis and there is no unique z value. However, the condition for both meshes to be 2.5D inside the collar prevents from this case.

[2] For a generic triangle ΔABC, the normal vector is obtained as $n = (B - A) \times (C - A)$

Finally, some global parameters are calculated from the Δz values, such as $\max(\Delta z)$, $\min(\Delta z)$, $\text{avg}(\Delta z)$ and

$$\text{RMSE} = \sqrt{\frac{\sum \Delta z^2}{n}} \qquad [26]$$

4.5 Step 2: New triangulation

In order to close the gaps and to permit C0 continuity between adjacent triangles at both sides of the borders, a new constrained DELAUNAY triangulation is performed.

Inside the xy-plane annular domain all high and low resolution points are selected; a check is performed to pick only once high or low resolution points which may share common planimetric coordinates. The borders are treated as breaklines. The external points of overlapping low resolution triangles are also included, although the latter choice is optional. A topologically continuous surface is created because the points lying on the outer border are now connected to the low resolution ones. Moreover, this enables to split large triangles which may be partially inside and partially outside the overlapping zone.

No other constraints are considered at this stage. This can be seen in Figure 4.5, right: the heights of the points in the transition zone are not smoothed, although the two meshes are now correctly merged.

The free and open-source program Triangle [SHEWCHUK, 1996] has been used for the triangulation operation. It has been developed specifically in order to create two-dimensional finite element meshes, but it can also perform simpler tasks such as DELAUNAY triangulations, simple or constrained.

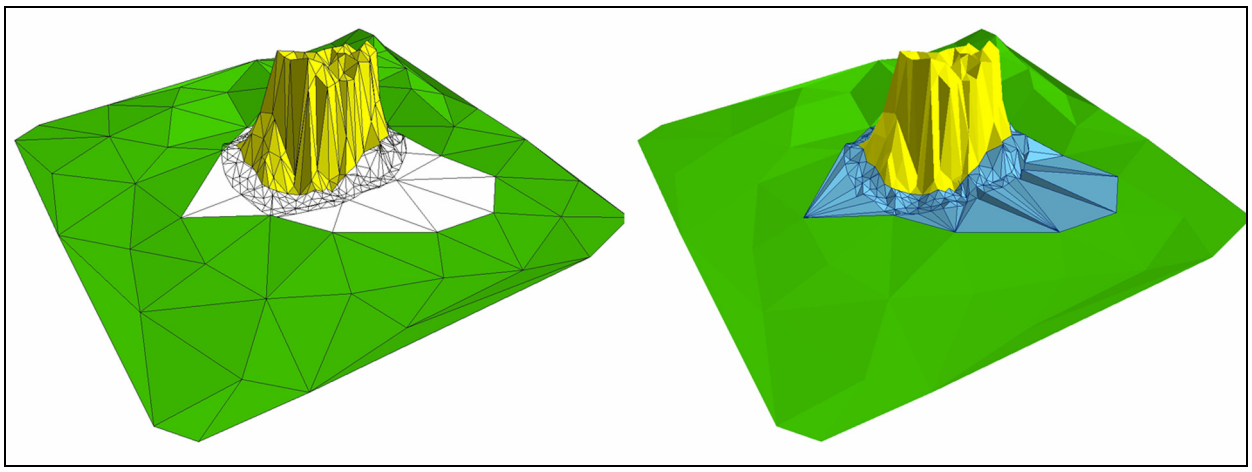

Figure 4.5: [Left] Selection of the triangles in the overlapping zone, represented in white, from both the high and low resolution models, in green and yellow, respectively. [Right] The new constrained DELAUNAY triangulation is carried out. The resulting mesh in represented in azure.

4.6 Step 3: Height interpolation

Once the new constrained triangulation has been carried out, topological continuity has been achieved. The height profile inside the transition surface has not been modelled yet: ideally, height values should be more "similar" to the DTM the closer they get to the outer border and – vice versa – they should be more "similar" to the high resolution object, the closer they get to the inner border, see Figure 4.6, left.
For this purpose, a weighed height interpolation function is formulated. Inside the overlapping zone, every point will have an intermediate z_{new} height value given by

$$z_{new} = z_{lowresolution} + w\Delta z \qquad [27]$$

with

$z_{lowresolution}$: height value in the low resolution mesh (for high resolution points, it corresponds to the projected z_v^p, as in [25]),
Δz : point-wise distance between the two meshes, as in [23],
w : distance-dependent weight, as discussed in § 4.6.1.

4.6.1 Distance-dependent weight function

The accompanying distance-dependent weight w must be formulated. On the x-axis a [0,1] domain is defined, which represents the range of a normalised distance parameter. Similarly, on the y-axis a [0,1] codomain is defined, representing the range of the weight values. The height value of the low resolution model at the outer border corresponds to the origin of the axis (0,0), while the height value of the high resolution model at the inner border corresponds to point (1,1). Different weight functions can be modelled.
The simplest, linear function $y = f(x) = x$ allows for C0 continuity: points lying on the outer border ($x = 0$) have the same height as the DTM ($y = 0$); points on the inner border have the same height as the high resolution mesh. However this linear model does not guarantee C1 continuity at the border points (0,0) and (1,1), where the transition surface is joined with the models.
An exponential function $y = f(x) = x^a$ (with $a > 0$, $a \neq 1$) allows for C0 and some C1 continuity, however the tangent is acceptable only at either of the border points, depending on the value of parameter a.
A piecewise polynomial curve like a uniform spline of second degree can provide the required level of continuity. Its formulation is as follows:

$$y = f(x) = \begin{cases} x^2 & 0 \le x \le 0.5 \\ -2x^2 + 4x - 1 & 0.5 < x \le 1 \end{cases} \qquad [28]$$

Some functions are plotted in Figure 4.6, right.

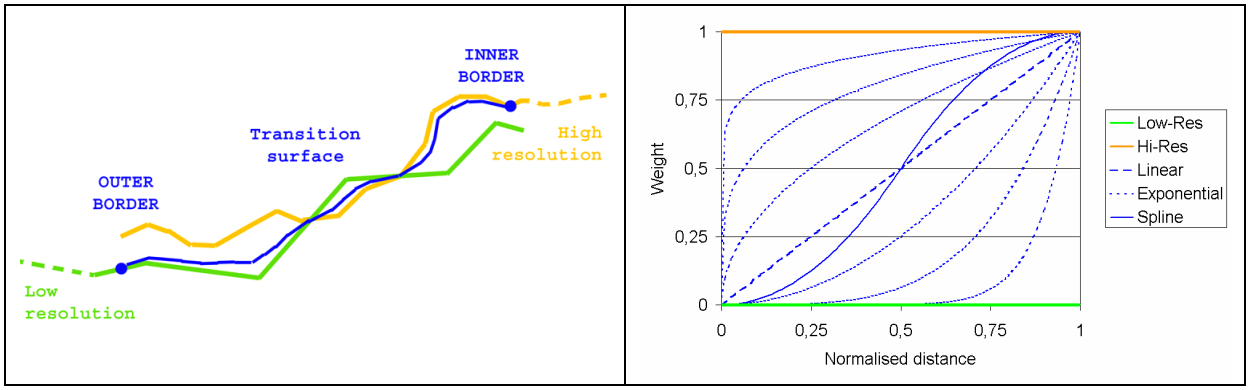

Figure 4.6: [Left] Desired qualitative height profile of the transition surface,in blue colour. Height values near the outer border should be more "similar" to the DTM; vice versa, they should be more "similar" to the high resolution object, the closer they are to the inner border. [Right] Different weight functions: linear, exponential (with varying exponential values) and spline.

4.6.2 Normalised distance parameter

A proper formulation must still be given to the x parameter, which represents the normalised distance between inner and outer border for any point belonging to the transition surface. Hence, it can be so defined:

$$x = \frac{d_{outerBorder}}{d_{outerBorder} + d_{innnerBorder}} \qquad [29]$$

with

$d_{innerBorder}$: shortest path between a point and the inner border with respect to the annular topology of the domain shape in the xy-plane,

$d_{outerBorder}$: shortest path between a point and the outer border with respect to the annular topology of the domain shape in the xy-plane.

With regular convex shapes, computation of distance can be generally performed with simple Euclidean distance functions. PostGIS itself offers some standard out-of-the-box GIS functions like **ST_DISTANCE**, which returns the smaller distance between two geometry entities. However, in case of extremely irregular shapes (which may show concavities of the borders, presence of multiple holes, etc.) computation of distance in such way can lead to errors. Some situations are shown in Figure 4.7: the left picture represents, in orange, a regular annular shape of the overlapping zone on the xy-plane. The central hole stands for the high resolution object, the outer zone is the surrounding DTM. The right image generalises the shape to an extremely irregular one, where multiple holes are present and several concavities can be found, both on the inner and the outer border. Although it is an artificially obtained shape, it is meant to simulate a case where two high resolution objects share the same collar. In addition, its irregular shape permits to recognise unsuitable results of the Euclidean distance calculation.
In both pictures, the shortest path from a point to the inner border inside the orange area is depicted in red. Green is used for the shortest path to the outer border. It is easy to see that simple Euclidean distance functions like **ST_DISTANCE** can provide a correct value for regular shapes. However, in the right example, results given by the

ST_DISTANCE function and represented by the red and green solid lines are not always acceptable, since they may cross several time the borders and do not respect the annular topology of the orange shape. Dashed lines are also provided to show a qualitative result of the correct computation for the shortest path from the points to the borders.

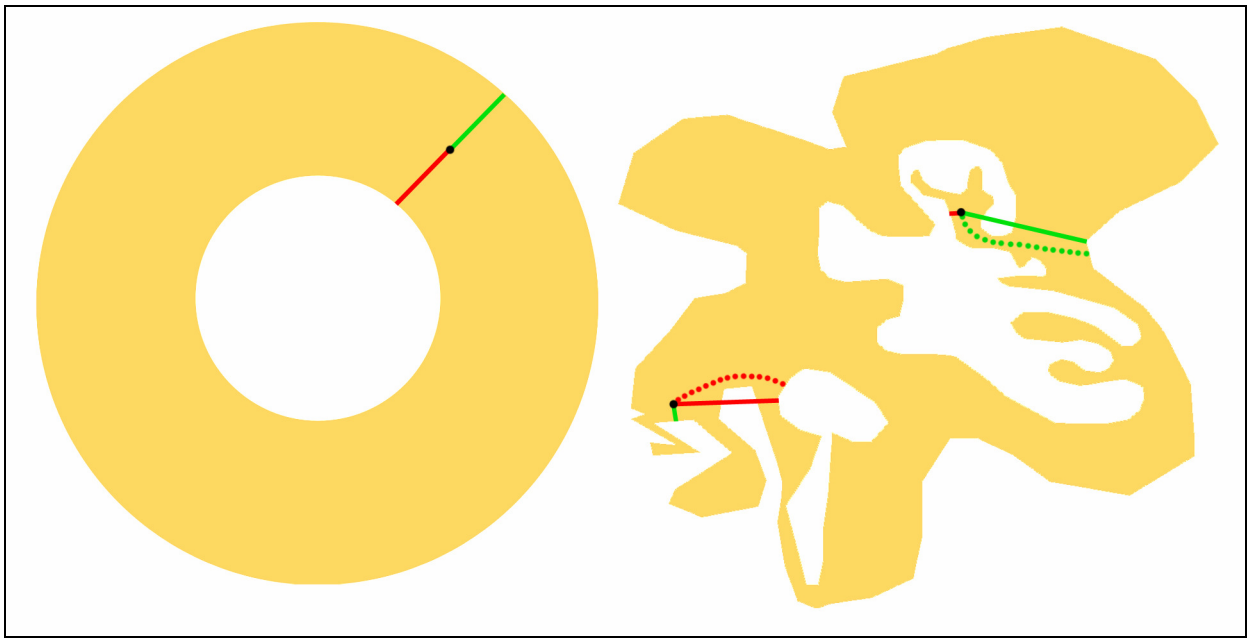

Figure 4.7: [Left] Inside the overlapping zone (in orange), standard Euclidean distance functions provide correct results for the distance from a point to the inner and outer border (in red and green solid line, respectively) if the shape is regular. [Right] In case of irregular shapes they may deliver unsuitable results, since distance is computed across the borders. Qualitative correct results are shown by the dashed lines.

A general solution to this problem consists in calculating distance values from the inner and outer borders separately, and then combining the two partial results. The distance values $d_{innerBorder}$ and $d_{outerBorder}$ are calculated through progressive iterative buffering. Once a buffer width is chosen, buffers are gradually calculated from the inner border outwards and, vice versa, from the outer border inwards. A combination of the two resulting maps, according to [29], yields the normalised distance from the outer border for any point inside the transition surface domain, see Figure 4.8.

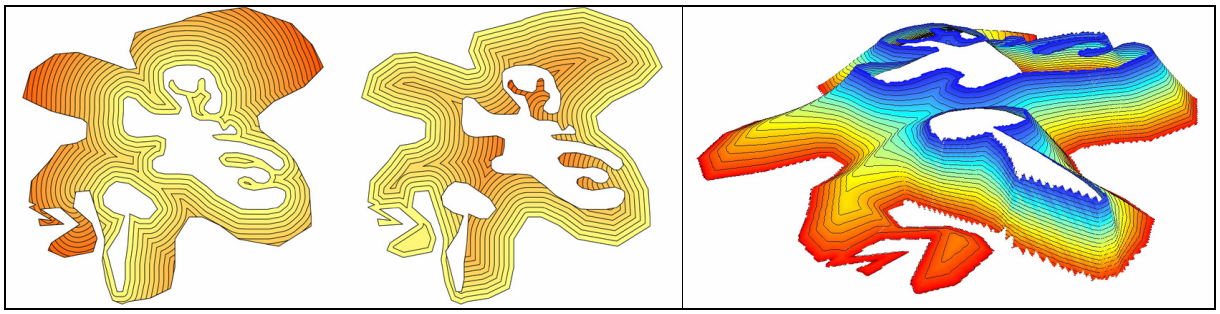

Figure 4.8: [Left, centre] Example of progressive buffering from the inner borders outwards and vice versa, inside an irregularly shaped domain. [Right] Three-dimensional representation of the combination of the two buffer maps, which yields the normalised distance from the outer border for any point inside the irregular shape.

Once the x parameter and the weight function $w = y = f(x)$ have been determined, C0 and C1 continuity are now added at the borders. The height profile inside the transition surface can be obtained from [27] as

$$z_{new} = z_{lowresolution} + f(x)\Delta z \qquad [30]$$

Depending on the used weight function, linear or spline, different results are obtained for the transition surface. Some examples are presented in Figure 4.9.

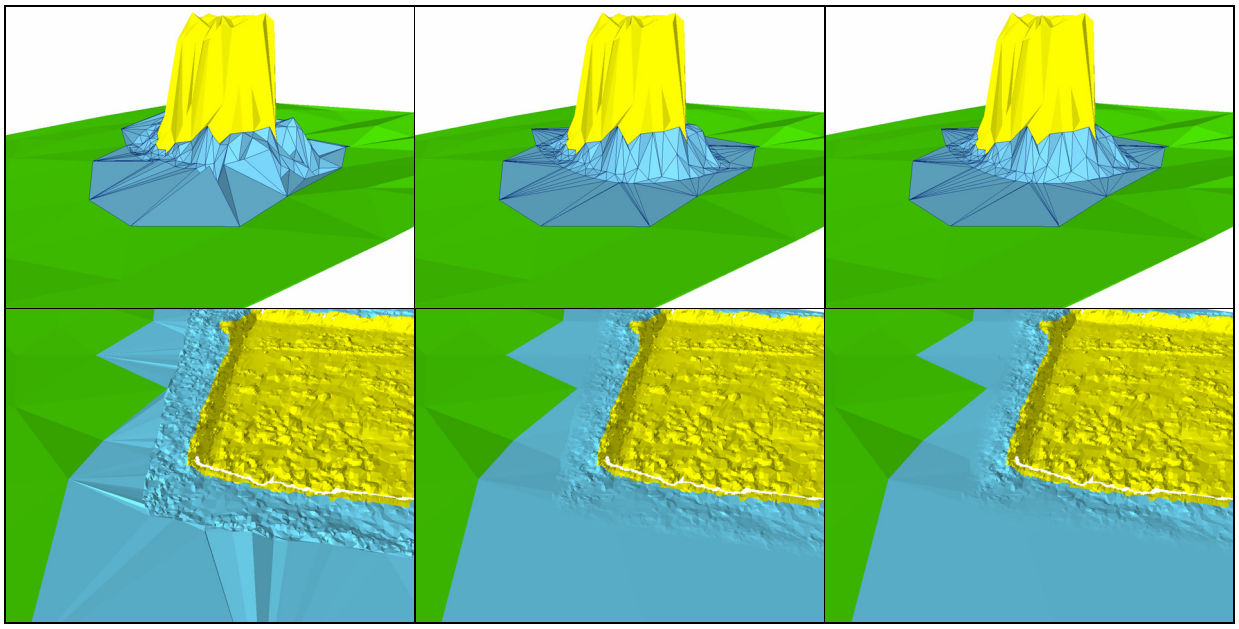

Figure 4.9: Transition surface creation steps: before applying the height interpolation model [left], with a linear weight function [centre] and with a spline function [right]. Examples refer to synthetic data [top] and to real data coming from an archaeological excavation site [bottom].

4.7 Mesh simplification

Although topological and geometric continuity have been obtained so far, a transition in terms of point density has not been modelled yet: density should decrease gradually, moving from the inner to the outer border of the transition surface. Hence the mesh must be simplified accordingly.

Simplification is the process to automatically reduce the complexity of a given model [LUEBKE *et al.*, 2002]. There is not a single best simplification algorithm, since the characteristics of the input model and the application for which the simplified output model is obtained play an important role in the selection of the proper simplification technique.

As a rule of thumb, a compromise between detail richness and a reduced amount of computational resources, coupled with a detail decrease, seems to be the only way when looking for the proper representation. Ideally, smaller or less important objects should be modelled using fewer, thus larger, triangles than the more important and detailed objects.

Several mesh simplification algorithms have been developed in the last twenty years in the field of Computer Graphics, which permit to eliminate selected elements of the mesh within a certain error value. Giving an extensive and detailed overview of all published algorithms is well beyond the scope of this work due to the quantity of existing literature, a general overview of the main types and their common classifications will be given, yet emphasis will be set on those aspects which have been crucial to select the simplification algorithm to be implemented.

4.7.1 A quick overview

Several classifications for the existing mesh simplification algorithms have been proposed on the basis of different criteria. VARSHNEY (1994) distinguishes with regard to the way the optimal solution is sought. If an input model is given with a complexity value n (measured in terms of number of vertices or faces) and an ε error value, a simplification algorithm may try to reach the maximum simplification (i.e. the smallest number of vertices or faces) with respect to the ε error. As an alternative, given a target level of complexity, the algorithm may try to achieve it minimising the error ε in the process. Another distinction can be made according to the type of input: some algorithms require a two-manifold triangulated surface, while others can work with more general triangle sets.
The ability to preserve, or not, the input topology also in the simplified models is another important characteristic to take into consideration. The connectivity structure of the output surface is intended to remain unchanged, such that a triangle shares edges and vertices with the neighbouring triangles only: any other self-intersections are forbidden. In this case the simplification algorithm is said to preserve *local topology*. If connectivity and no self-intersections are guaranteed for the whole model, then it preserves *global topology*.
Finally, local or global mesh simplification operators can be distinguished. The former simplify the geometry and connectivity in a local region of the mesh, the latter operate over larger regions and can help to simplify the mesh topology.
Since a topology-preserving algorithm is sought for the transition surface, only local simplification operators and their effects on the resulting mesh will be further discussed.

4.7.2 Local simplification operators

Every local operation reduces the number of vertices or faces by a small amount and is repeated until the final simplification level is obtained. The most common local operations are presented here.
By *vertex removal*, a single vertex is removed from the mesh together with all *n* adjacent triangles. Therefore, this operation creates a hole with *n* sides that must be closed by means of a new triangulation with *n*-2 triangles. Its closure implies that a new triangulation has to be chosen from a finite number of possibilities, see Figure 4.10. More precisely, the number of unique ways to triangulate a convex planar polygon with *i*+2 sides is provided by the Catalan sequence $C(i)$, which gives the upper limit for a number of non-self-intersecting triangulations of a hole in 3D [PLOUFFE & SLOAN 1995]:

$$C(i) = \frac{1}{i+1} \bullet \binom{2i}{i} = \frac{1}{i+1} \bullet \frac{(2i)!}{i!(2i-i)!} = \frac{1}{i+1} \bullet \frac{(2i)!}{i!(i)!} = \frac{(2i)!}{(i+1)!i!} \quad [31]$$

By *edge collapse*, two vertices of an edge are merged into a single vertex. In the process, all the neighbouring triangles are stretched while the two triangles sharing the edge are removed. Unlike the vertex remove operation, where there is a finite number of triangulation to chose from, the edge collapse must select the coordinates of the new vertex from a continuous domain. Generally, the start point, the end point or the midpoint of the collapsed edge are used, but any other point on the collapsed edge or in its surroundings can be chosen. Care must be taken of this operator, because it could lead to mesh foldovers or to topology inconsistencies. A mesh foldover consists in a great change in the angle of the triangle normal, usually greater than 90°. Besides, a non-manifold edge could be also created.

Face collapse is similar to the edge collapse operation: an entire triangle collapses into a single vertex, thus reducing the number of faces by four (including the three surrounding triangles). The new vertex can be one of the triangle vertices or can be calculated *ex novo*. Compared to edge collapse, this operation may contribute to a faster simplification algorithm, however it delivers coarser results since the error tends to accumulate more quickly in comparison to edge collapse. It is also possible to further generalise the process, as larger connected portions of the input model can be collapsed at the same time.

By a *vertex cluster* operation, several nearby vertices are merged into a single vertex, whose coordinates may coincide with one of the removed vertices or newly calculated in order to minimise the error. The search radius to find nearby triangles can be set in different ways. Regular space partitioning – e.g. a grid or an octree – can be used. In these cases, vertices inside one spatial unit (a cell or a voxel) are merged together. One of the strengths (and at the same time drawbacks) of vertex clustering is that it relies only on the geometry of the input model, whereas topology is ignored (and not required). This means that it can be applied to arbitrary sets of triangles, but there is no guarantee about the connectivity between the faces in the simplified model. For this reason, this operation and those who derive from it (i.e. generalised edge collapse: a combination of edge collapse and vertex cluster, etc.) will not be further investigated.

Polygon merging consists in merging nearly coplanar and adjacent polygons into larger polygons which are then triangulated. This operation is a more general version of vertex removal, since it can also use polygons other than triangles. Besides, several vertices can be removed at once with the possible result of merged polygons with holes.

A general geometric replacement operator has been proposed by DEFLORIANI *et al.* (1997) and is able to use both edge collapse and vertex removal operations. Edge flip is also possible, which replaces the common edge of two triangles with the edge connecting the two other opposite vertices. A group of adjacent triangles is substituted by another, numerically smaller set of triangles which share the same boundary.

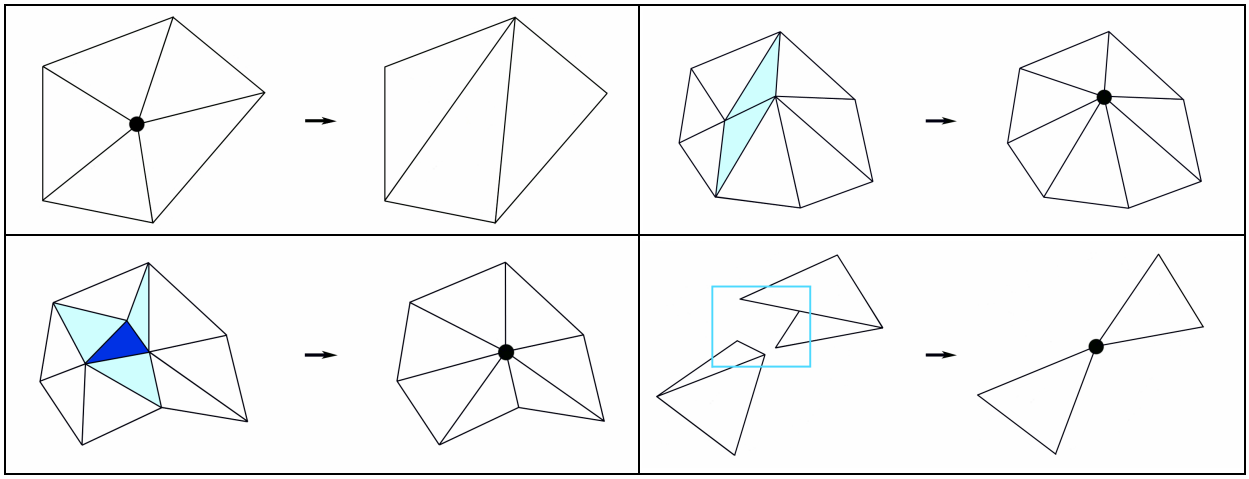

Figure 4.10: Examples of mesh simplification operators: vertex removal [top left], edge collapse [top right], face collapse [bottom left] and vertex clustering [bottom right].

4.7.3 Error metrics

Every simplification operation reduces the complexity of an input model by a certain amount. However, a measurement of the output quality is fundamental: the necessity to quantify the errors made in every simplification step, and how much the final model differs from the input one, has lead to several approaches to measure the errors. Some are tailored to geometry, others concentrate on how attributes of a mesh (colours, normals, textures) can be preserved in the final model. In this paragraph geometry error metrics and some relative simplification approaches will be discussed. Two formulations of the distance between two surfaces will be given: the HAUSDORFF distance and the mapping distance.

Given two meshes built upon point sets A and B, the HAUSDORFF distance d_H measures the distance between them. It is defined as the maximum of the minimum distance between points in the two sets. For every point x in A the closest point y in B is found, and vice versa, which leads to

$$d_H(A,B) = \max\left(\max_{x\in A}\min_{y\in B}(d(x,y)), \max_{y\in B}\min_{x\in A}(d(y,x))\right) \qquad [32]$$

with

d : a distance function, i.e. $d(x,y) = \|x-y\|$.

From Figure 4.11 it can be seen that $d(x,y)$ and $d(y,x)$ are not symmetric. Although the HAUSDORFF distance is constructed to consider both the one-sided distances and to keep the maximum, one of its drawbacks is that no neighbourhood information is used to establish a correspondence between pairs of points, thus leading to possible discontinuities: points from A might have multiple corresponding points on B, or none – and vice versa.

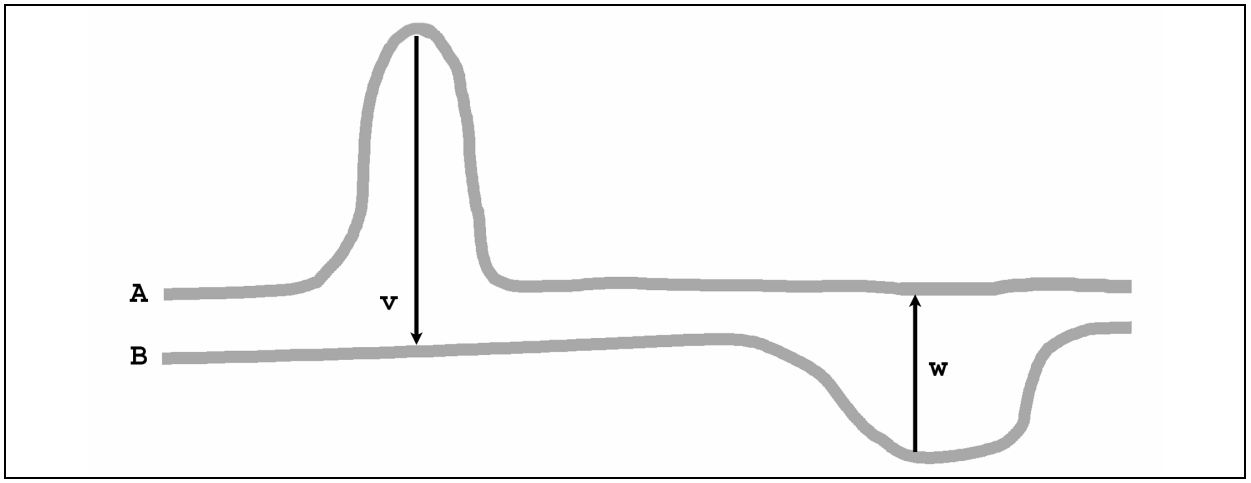

Figure 4.11: One-sided HAUSDORFF distance between two surfaces A and B is not symmetric: $\|v\| = \max_{x \in A} \min_{y \in B} (d(x,y)) \neq \max_{y \in B} \min_{x \in A} (d(y,x)) = \|w\|$ **. Two-sided HAUSDORFF distance is instead** $d_H(A,B) = \max(\|v\|, \|w\|) = \|v\|$**.**

If a point-to-point continuous mapping function[3] is instead defined, distance between the two surfaces can be calculated with respect to this function. For a function $f : A \to B$ the distance between corresponding points in A and B can be defined as

$$d(f) = \max_{x \in A} \|x - f(x)\| \qquad [33]$$

but since there are many possible mapping functions f, the general minimum distance is simply

$$d_{\min} = \min_{f \in F} d(f) \qquad [34]$$

where F stands for all continuous mapping functions. Although $d_{\min}$ and its associated mapping function may be difficult or impossible to explicitly compute, any of the functions provide an upper limit to $d_{\min}$ as well as on the HAUSDORFF distance [LUEBKE, REDDY *et al.*, 2002].
The distance between two surfaces can be therefore measured with either formulations as the maximum of the point-wise distances. Instead of the maximum, the average or other aggregate functions could be used. However, the maximum guarantees that the error will never be greater than that value, and this is often a desired property.

4.7.4 Error measurement approaches

Once an error metric has been established, one last algorithm classification can be done with regards to the strategies adopted to measure the distances between the original and the simplified model.

[3] In topology, a function $f : X \to Y$, where X and Y are topological spaces, is continuous at x ($x \in X$) if for any neighbourhood V of $f(x)$ there is a neighbourhood U of x such that $f(U) \subseteq V$.

The vertex-vertex approach is the simplest one. The distance between the original vertices and the simplified vertices is measured, however the result depends on which vertex correspondences have been established and which simplification operators have been used. In case of an edge flip between to adjacent triangles, without any other modifications, the surface may change, although the vertex distance could remain zero.
A second approach consists in measuring the distance between a vertex and a plane. Given for example a point $P:(x,y,z)$ and a plane $\pi:(n,q)$ defined by its normal vector $n:(n_x,n_y,n_z)$ and the signed distance from the origin q, the distance between them can be written as

$$d(P,\pi) = n_x x + n_y y + n_y y + q \qquad [35]$$

Every vertex in a mesh has a certain number of supporting planes: one for every adjacent face [RONFARD & ROSSIGNAC, 1996]. If the vertex is collapsed, the supporting planes change and increase in number. Indicating the vertex with $v:(v_x,v_y,v_z,1)$, it is therefore possible to formulate the error as

$$E_v = \max_{\pi \in planes(v)} (v \cdot \pi)^2 \qquad [36]$$

This vertex-plane parameter is helpful in order to give a priority to the edge collapse operations, but it may underestimate the maximum deviation of the simplified surface from the original.
GARLAND & HECKBERT (1997) have improved the method introducing error quadrics. Instead of the maximum, the error parameter yields the sum of the squared vertex-plane distances.

$$E_v = \sum_{\pi \in planes(v)} (v \cdot \pi)^2 = \sum_{\pi} (v^T \cdot \pi)(\pi^T \cdot v) = v^T \left(\sum_{\pi} \pi\pi^T \right) v = v^T \left(\sum_{\pi} Q_\pi \right) v = v^T Q_v v \qquad [37]$$

Every plane contributes with a 4x4 matrix $Q_\pi = \pi\pi^T$ which can be added to propagate the error whenever vertices are merged. Moreover, the optimal position of the vertex can be found by minimising the error E_v and solving the system of linear equations.
Further improvements have been proposed in the course of time [e.g. GARLAND & HECKBERT, 1998; HOPPE, 1999], but the metric which can indeed provide the most reliable error bound is the surface-surface distance approach. Every point in the input surface and in the simplified model is considered to determine an error. Such simplification strategies concentrate on minimising the maximum error, such that anywhere in the simplified model the distance to the original model can be guaranteed to be always shorter or equal to it.

4.8 Step 4: Gradual point density reduction

Bearing in mind the previously presented local simplification operators and the error metrics strategies, it is shown that the topology of the original mesh is not changed by

operations such as vertex remove or edge collapse. For the purpose of the transition surface, the former approach has the advantage of eliminating existing points (i.e. observations) from the mesh, without having to calculate or insert new ones like, in general, it happens during vertex collapse operations.
Finally, even if a surface-surface error measurement approach is desirable, the simplification algorithm should also allow to set a variable error throughout the transition surface. The idea is: the further one moves from the inner border to the outer border, the bigger the error is allowed to be, i.e. smaller near the high resolution zone, greater near the low resolution zone.
Among the existing different approaches, "Simplification envelopes" [COHEN, VARSHNEY *et al.*, 1996) seems to fulfil the requirements in terms of topology preservation and varying error value.
Simplification envelopes consists of two offset surfaces which are no more than a user-defined ε value distant from the original surface. The "outer" surface is created by a displacement along the normal vector of every vertex by ε and the "inner" surface is created by displacing it by $-\varepsilon$. Since the envelopes are not allowed to self-intersect, the simplified model surface will lie between the offset surfaces.
The ε value can be set to be either constant or variable on a per point basis, so a variable approximation is allowed, i.e. details which are not to be simplified beyond a certain level can be preserved. Furthermore, point reduction occurs only through simple deletion of selected points.
Mesh simplification is obtained iteratively: a mesh vertex is removed, a hole is therefore created which is then triangulated again with fewer triangles. If all new triangles do not intersect the offset surfaces or the other surface faces, the point deletion and the new triangles are accepted, and the algorithm continues with the next vertex.
The most important steps needed for this simplification algorithm are explained in the following paragraphs.

4.8.1 Offset computation

The offset surfaces lie no more than a given ε from the input surface, both inwards and outwards. They build a domain volume which contains the input model and the simplified model. In order to compute the offset surfaces, for every vertex *v* a normal vector n_v is obtained as the area-weighed sum of all neighbouring face normals

$$n_v = \frac{\sum_{i \in faces(v)} a_i n_i}{\sum_{i \in faces(v)} a_i} \qquad [38]$$

A copy of every triangle is gradually moved along its vertex normals, in one direction for one offset surface, in the opposite direction for the other one. During each step triangle collision tests are carried out to avoid intersections with other triangles. For this purpose, the fast triangle-triangle algorithm by MÖLLER (1997) has been implemented, which permits to check whether two triangles intersect in 3D. A special case is represented by

the neighbouring triangles, for which the collision test has to be adapted. A graphical example for a synthetic model is given in Figure 4.12.
In addition, in order to reduce the global number of tests over the whole surface, a simple octree data structure is implemented and used.

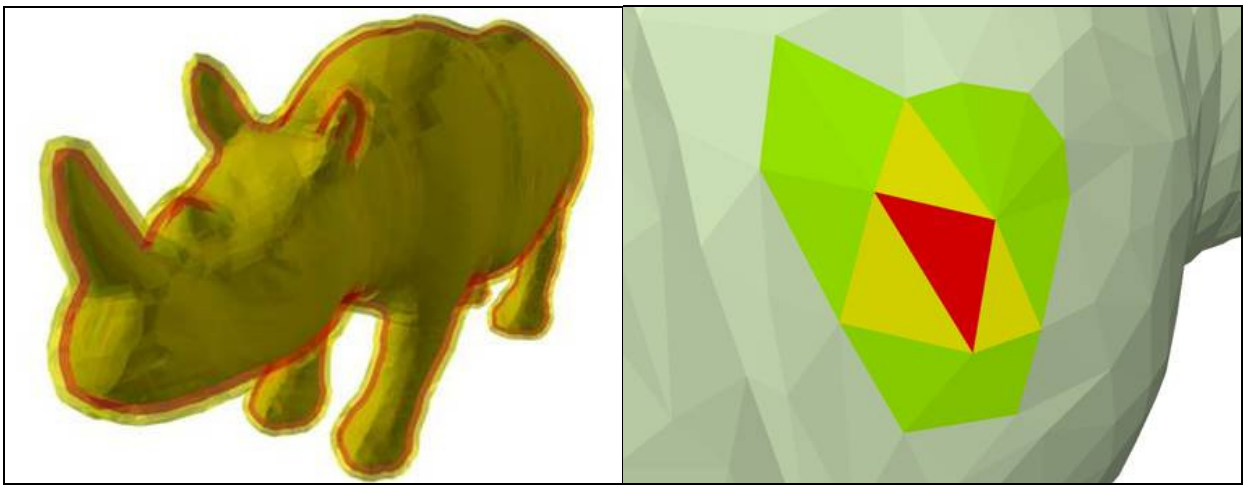

Figure 4.12: [Left] The resulting offset surfaces and an input model can be seen overlaid, with transparent colours. [Right] Neighbouring triangles for which a modified collision test is required: given the red triangle to be gradually offset, neighbouring triangles are allowed to touch a side (yellow triangles) or a vertex (green triangles). For these triangles the collision test has to be loosened, since adjacency "collisions" at those specific points and edges are actually correct. Tests for the implemented algorithm have been carried out using RapidForm's three-dimensional model "Rhino".

4.8.2 Vertex removal

Once the two offset surfaces have been obtained, the actual simplification process can start. Iteratively, a vertex is removed and a hole is therefore created with *n* sides. A new triangulation is carried out inside the hole, resulting in *n-2* new triangles (see vertex removal, § 4.7.2). All the vertices of the new triangles lie within the offset surfaces, however the triangles must themselves lie within the offset surfaces and must not intersect existing triangles of the simplified surface. Every candidate new triangle is therefore tested for intersections with the envelopes and with the surrounding triangles. Again, the process can be sped up if an octree is used.
If all collision tests succeed (i.e. no intersections), then the new triangulation of the hole is accepted and the algorithm proceeds to the next vertex.

4.8.3 Vertex error modelling

The ε parameter defines the maximum distance of the offset surfaces from the input model, in other words, the "thickness" of the envelopes: the bigger the ε value is, the thicker the envelopes will be and, as a consequence, the further the simplification of the mesh will proceed. However, due to the envelopes geometry, all resulting triangles will be no further then ε from the original surface. The ε value is therefore a quality parameter for the output mesh.
From a geometric point of view, the ε value has many common traits with the maximum distance between the original curve and the simplified curve in the DOUGLAS-PEUCKER

algorithm: for a given curve composed of approximating segments, a similar curve with less points is sought, consisting of a subset of the original points. All points in the simplified curve are no further than an input distance $\varepsilon_{DP} > 0$ from the original line. On the other hand, all discarded points are closer than ε_{DP} to the nearest approximating line segment. A graphical example is given in 2D in Figure 4.13, with the original unsmoothed curve on the top, and the final output curve on the bottom right.

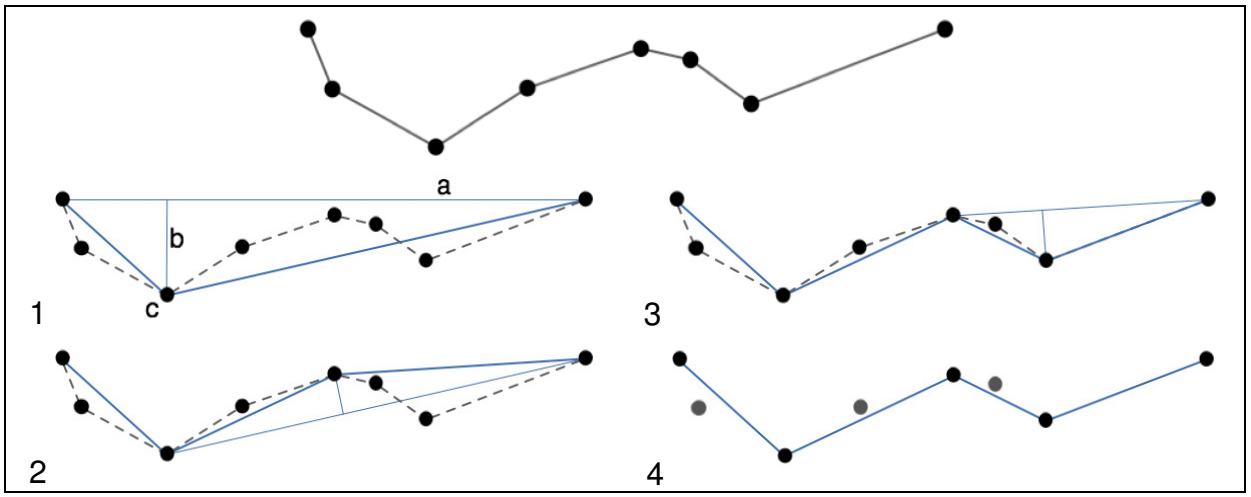

Figure 4.13: The DOUGLAS-PEUCKER algorithm allows to smooth a piecewise linear curve. Steps are shown from the input curve [top] to the final smoothed curve [bottom right]. Image source: [DOUGLAS-PEUCKER].

When it comes to the complexity reduction (i.e. the number of triangles), the simplification envelopes algorithm can be successfully used with datasets coming from laser scanners: redundant data of over-sampled models are reduced, while the global shape is preserved within a certain error bound. It is useful to distinguish two mutually dependent aspects which characterise the data reduction process.

As an example, let's think of a mesh obtained from a point cloud of a smooth vertical wall: even with a small ε value, the mesh can be greatly simplified, since only few triangles are required to preserve the shape of the object which is originally made of nearly coplanar triangles. Ideally, in case of a perfect wall without any noise in the point cloud, simplification could progress to as few as two triangles with a positive ε≈0, without detail loss.

From the signal processing point of view, a smoothing (filtering) of the high frequency details on the mesh surface is carried out. The complexity of a surface can be described by means of concepts like irregularity or roughness. Its representation can be transformed from the space domain to the frequency domain by means of a FOURIER transformation. Every surface is then characterised in its frequency domain by its own frequency spectrum. In general, a spatial dataset consists of three components: regional variations, local variations and random noise [LI, ZHU & GOLD, 2004]. The first component defines the basic shape of an object and represents the low frequency part of the spectrum, the second one varies according to the scale: important for large scales since it contains information about details, but redundant at small scales because it adds unneeded information to the basic shape.

If a filter is able to separate the low frequency component from the remaining part, it is called a low-pass filter. It is however difficult to clearly define and separate high from low frequency components. Noise is by contrast always problematic, since it causes distortions in the appearance of both regional and local variations. However, the noise part is considered absent by hypothesis in this work, since both meshes used are assumed to be “clean” and outlier-free.
From this point of view, Simplification envelopes represents an analogy to a variable low-pass filter: for large values of ε, the frequency threshold of the digital filter is lowered (only very low frequencies are kept), while for small values of ε the threshold is raised and only very high frequencies are cut.
When it comes to the transition surface, the simplification is desired to grow *gradually* from the inner border to the outer border. Bearing in mind the previous remarks about the ε value, a correspondence can be set between ε and the normalised distance parameter *x* as computed in [29], see § 4.6.2. By setting *ε=0* at *x=0* for the inner border points and letting it grow toward the outer border up to an $\varepsilon_{\max}$ value at *x=1*, a gradual transition can be modelled in terms of detail reduction.
Similarly to the weight function used for the height interpolation, a distance-dependent ε function can be shaped, from the inner to the outer border, see Figure 4.14, and is defined as

$$\varepsilon = g(x, \varepsilon_{\max}) = f(1-x) \cdot \varepsilon_{\max} \qquad [39]$$

with
$f(1-x)$: a spline function, as in [28],
x: distance parameter, as in [29],
$\varepsilon_{\max}$: maximum displacement value of the offset surfaces.

The resulting simplification level of the transition surface will vary according to the growing values of $\varepsilon_{\max}$. An example is given in Figure 4.15: for $\varepsilon_{\max} = 0$ there is no simplification in the transition surface. As long as $\varepsilon_{\max}$ increases from 1 cm to 10 cm, triangles grow gradually in size (i.e. points are reduced), from the inner border to the outer border, according to the irregularity of the dataset. In addition, $\varepsilon_{\max}$ gives the upper bound for the quality of the mesh: triangles near the outer border will be no further than 10 cm from the mesh before simplification.

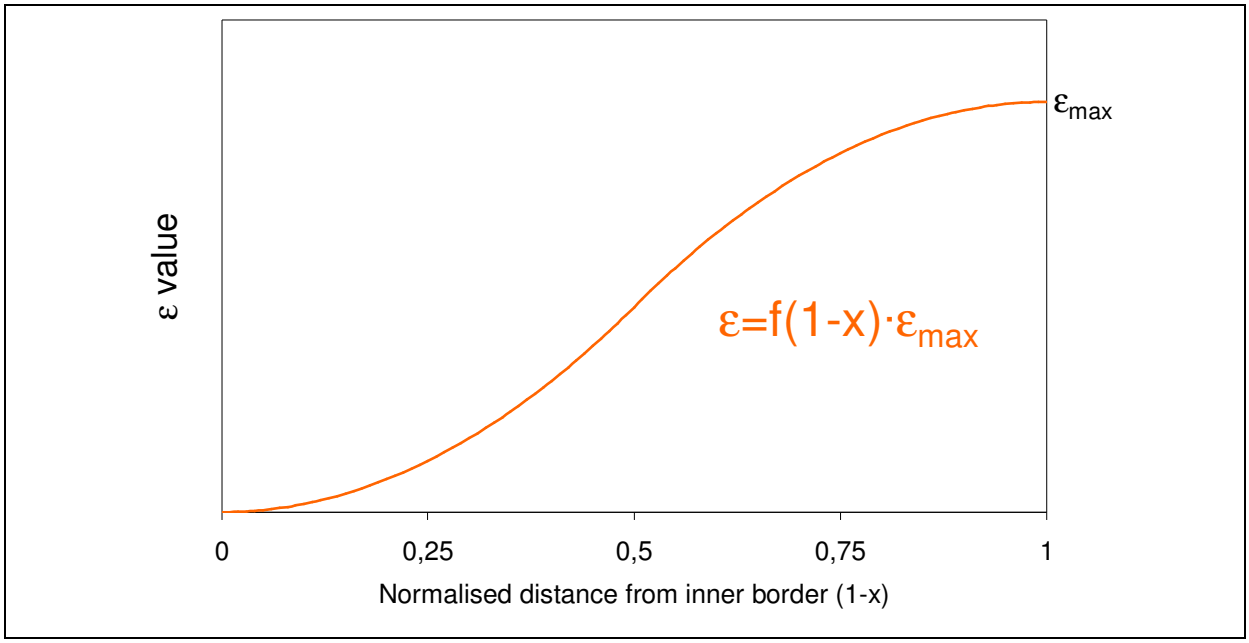

Figure 4.14: Similarly to the weights for the Δz values, a spline function is used to define the ε values for the points inside the transition surface.

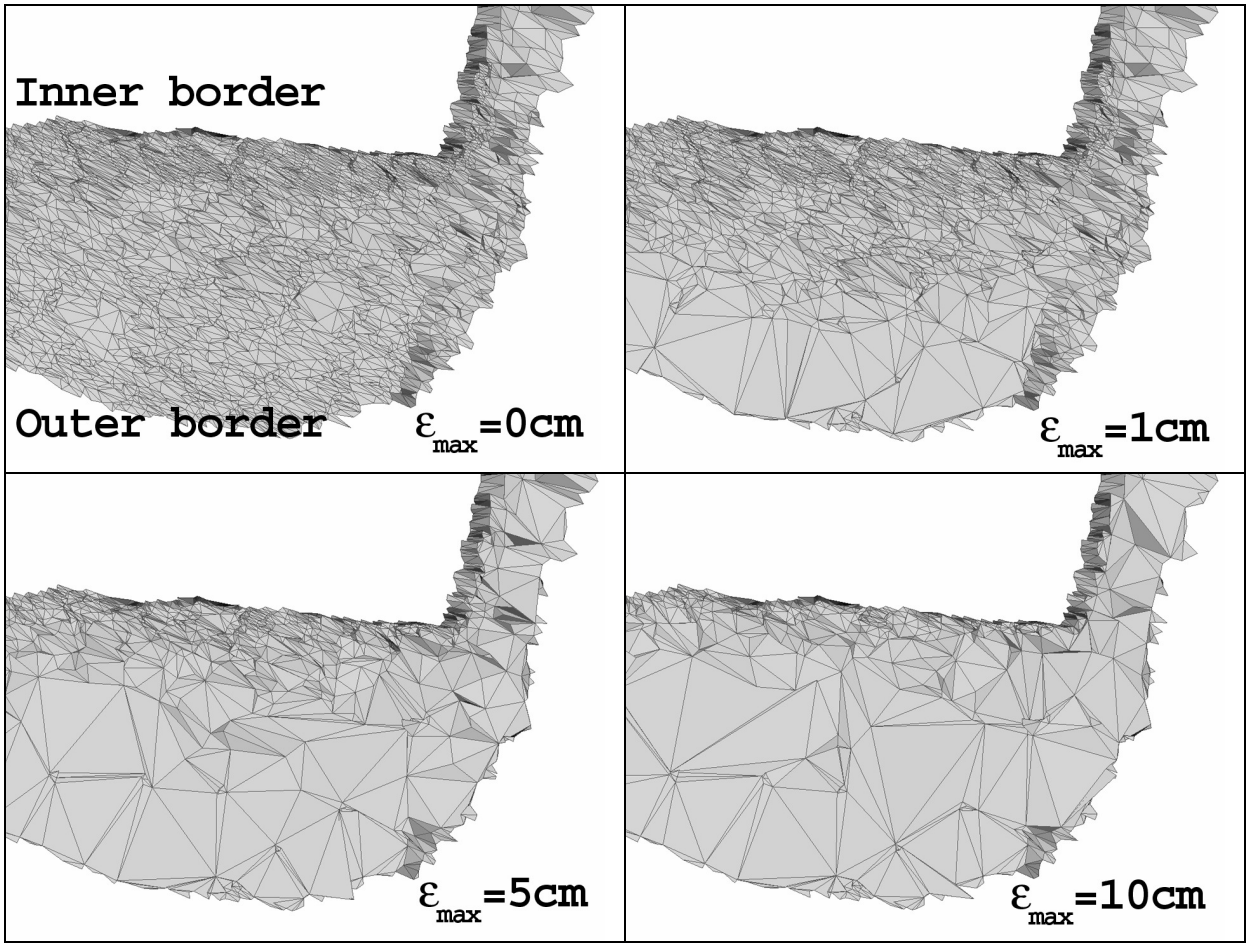

Figure 4.15: Progressive simplification in the transition zone, between the inner border and the outer border, with growing values of $\varepsilon_{\max}$.

4.9 Choosing the suitable error value

It has been shown in the previous paragraph that the level of simplification across the transition surface grows according to the $\varepsilon_{\max}$ value. It still remains to be explained how this parameter can be chosen and which the conditions are in order to set the most suitable value.

Since $\varepsilon_{\max}$ affects the junction between the outer border of the transition surface and the low resolution dataset, it should express the characteristics of the DTM at – or in the surroundings of – the outer border. The parameter $\varepsilon_{\max}$ should have a global or a local connotation. This depends on a variety of factors like the accuracy and the geometry of the DTM, and, most of all, on the presence or absence of such information. If nothing is known about accuracy, at least a global value should be assumed on the basis of the surveying technique used for the low resolution DTM. If this information or similar metadata are also missing – a rather common situation in the archaeological framework –, a characterisation could be performed using geometry only. If no suitable $\varepsilon_{\max}$ can be obtained in this way, some heuristics could be employed.

One of the advantages of the ε is its multiplicity: although it is a single value, it can be interpreted in different ways, as seen in the previous paragraph. In the following, some possibilities will be discussed to define it, together with some critical cases the integration approach developed in this work is not currently able to solve.

4.9.1 $\varepsilon_{\max}$ value as maximum displacement error

Bearing in mind the properties of $\varepsilon_{\max}$ in term of maximum error bound for the simplification envelopes, the simplified triangles of the transition surface are no further than $\varepsilon_{\max}$ from the surface before simplification.

If a height accuracy value σ_H for the DTM points is known, a first solution consists in applying an epsilon-band around every face of the low resolution DTM. In a conceptually analogous way to simplification envelopes, every DTM triangle can be thought as contained between an upper and a lower “enveloping” triangle, whose distances are given by σ_H. Therefore, setting $\varepsilon_{\max} = \sigma_H$ could be an acceptable solution to bridge the gap at the junction between transition surface and surrounding DTM.

4.9.2 $\varepsilon_{\max}$ value as terrain irregularity parameter

With no information about DTM height accuracy, some strategies can be applied to define $\varepsilon_{\max}$ on the basis of the low resolution terrain irregularity.

Bearing in mind the characteristics of the ε value in terms of high frequency filter, and its resemblance to the DOUGLAS-PEUCKER algorithm to smooth a polygonal curve, this analogy can be used if the curve is imagined as the height profile of a DTM, which consists of triangles.

The DOUGLAS-PEUCKER algorithm produces a simplified, smoothed polyline curve from a given ε_{DP} which is no further than ε_{DP} from the original polyline. If the principle on which the algorithm relies is somehow “inverted”, starting from a smoothed polyline a minimum $\varepsilon_{DP}^{r} \geq 0$ can be calculated, from which the polyline starts to be simplified. For any value in the range between 0 and ε_{DP}^{r}, no further smoothing can be carried out. The

concept, exemplified by a two-dimensional polyline shown in Figure 4.16, can be extended to the third dimension.

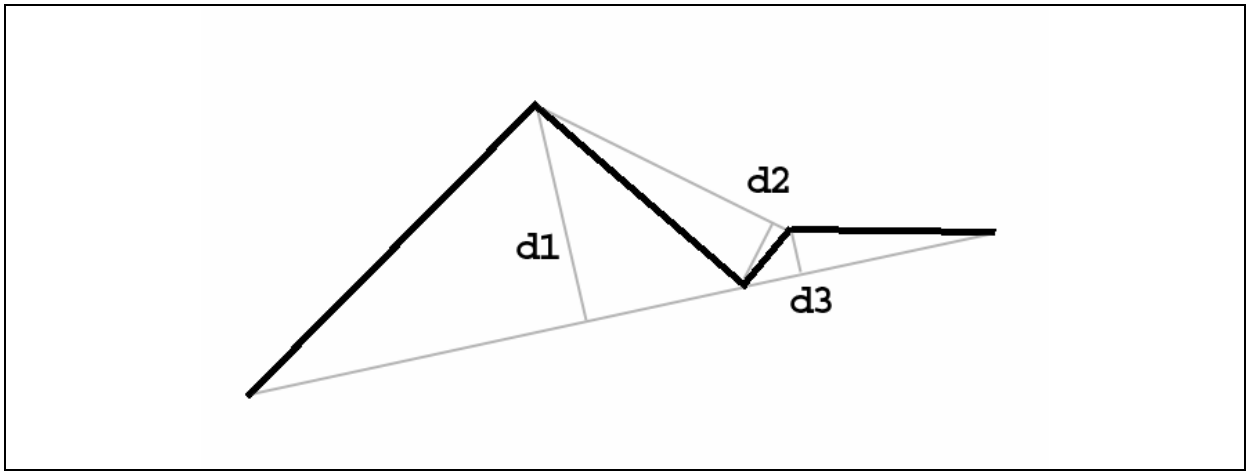

Figure 4.16: Given a polyline, its simplification starts for any ε_{DP}^{r} ≥d3.

Provided that an ε_{DP}^{r} exists and is greater than 0 (the worst case scenario $\varepsilon_{DP}^{r} = 0$ will be discussed in the next paragraph), then this ε_{DP}^{r} can be used as the geometric connecting value between the low resolution DTM and the outer border, thus setting $\varepsilon_{\max} = \varepsilon_{DP}^{r}$.
It remains to be explained how to obtain the ε_{DP}^{r} value. One feasible way is to iteratively apply the simplification envelopes algorithm to the DTM triangles surrounding the transition surface. A variable, growing $\varepsilon_{\max}$ is used, starting from a positive $\varepsilon_{\max} \approx 0$, until the first simplification takes place, i.e. the number of triangles is reduced.

4.9.3 Critical cases

The possibility to "interpret" the ε value in terms of maximum displacement or as a parameter of the terrain irregularity is an advantage, but this simplification has of course its costs in terms of flexibility. The drawback is that it is not universally employable, and critical cases exist for which this approach is unsuitable. For some an alternative is proposed, others represent critical cases for which there is not yet a solution.
Let's start with ε_{DP}^{r}. For simplicity, the discussion will refer to a two-dimensional polyline representing a terrain height profile, but analogous concepts hold for the real, three-dimensional case. The ε_{DP}^{r} parameter needs at least two segments to be determined. The two segments represent two triangles in the low resolution DTM. During the initial data import process in Step 1, only those DTM triangles are selected that are completely or partially covered by the overlapping zone (see § 4.4 and Figure 4.5, left).
If only one large triangle has been imported, ε_{DP}^{r} cannot be computed. In the two-dimensional case, this corresponds to a single segment of the polyline. In order to overcome this drawback, a solution might consist in importing at least the neighbouring triangles and then in calculating the ε_{DP}^{r} value. However, it may be argued that this workaround breaks the initial hypothesis to use only information in the surroundings of the transition zone. In addition, it may be questionable to obtain a TIN irregularity

parameter from adjacent, but distant triangles, which are not directly related to the overlapping zone.

A second problematic case is if $\varepsilon_{DP}^{r} = 0$, i.e. when two segments of the polyline are collinear. This can be imagined as a case of over-sampling: the geometry of the polyline does not change even if the point is removed. Theoretically, one of the strengths to use a TIN as DTM is that triangles can be used to describe a terrain with varying level of detail: plane areas with no height changes can be represented with very few large triangles (ideally, only two triangles for a square or rectangle area); for vertically irregular areas, more and smaller triangles are used.

This is, however, not always true in real world TINs: a certain amount of redundant data may be still present. A TIN obtained by triangulation of a regularly gridded DTM – indeed a quite common practice in many applications – could have a variable amount of redundant points, especially in very smooth areas.

If information is known about the accuracy of the gridded DTM, a solution could consist in performing a data reduction of the DTM within the known accuracy. An example is shown in Figure 4.17 to Figure 4.20: a tool in Esri ArcView's 3D Analyst has been used to transform a raster DTM into different TINs within a selectable height tolerance. The cell centres from the input raster are added as TIN points until the desired vertical accuracy is achieved. The closer the tolerance is to 0, the more points will be included in the TIN.

It must be noted that the solution proposed for the over-sampling problem has broken the initial assumptions again: "nothing happens to the low resolution DTM, except for the triangles lying in the overlapping zone". Moreover, the proposed workaround requires prior knowledge about the initial raster accuracy – but this is not always the case –, and secondarily it requires a certain degree of control in the creation of the low resolution DTM.

If the only source of information is the over-sampled DTM itself, and no other external data can be used without breaking the initial hypothesis, then a possible solution consists in choosing the $\varepsilon_{\max}$ heuristically, since setting $\varepsilon_{\max} = \varepsilon_{DP}^{r} = 0$ would impede the mesh simplification inside the transition surface.

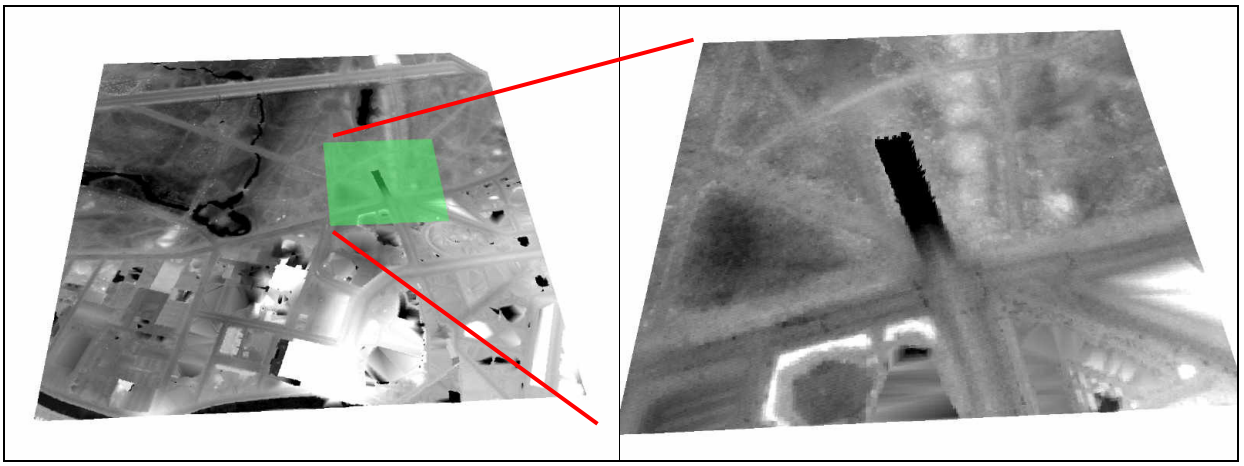

Figure 4.17: Excerpt from the Berlin DTM with raster resolution of 1 m. The southern entrance of the Tiergarten tunnel is highlighted, a zoom can be seen in the right picture.

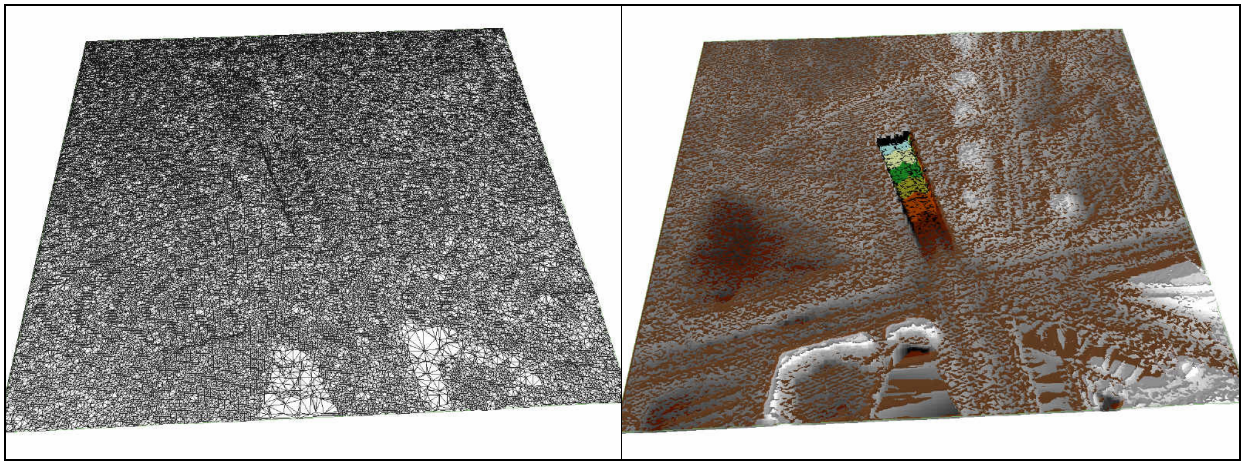

Figure 4.18: [Left] The TIN obtained in Esri Arcview from the DTM with a maximum height error of 5 cm. [Right] Overlay between the DTM (in grey colours) and the TIN.

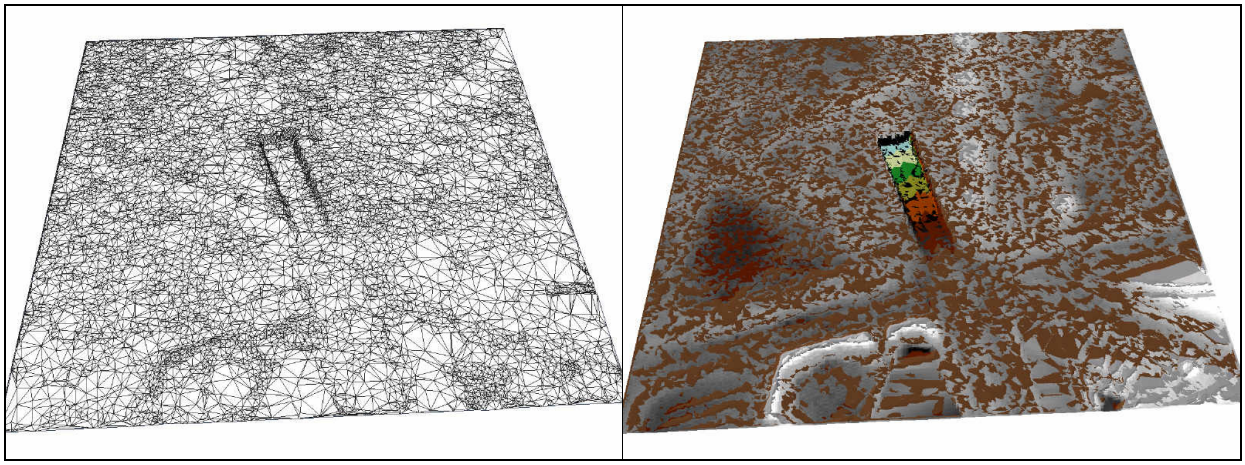

Figure 4.19: [Left] The TIN obtained in Esri Arcview from the DTM with a maximum height error of 15 cm. [Right] Overlay between the DTM (in grey colours) and the TIN.

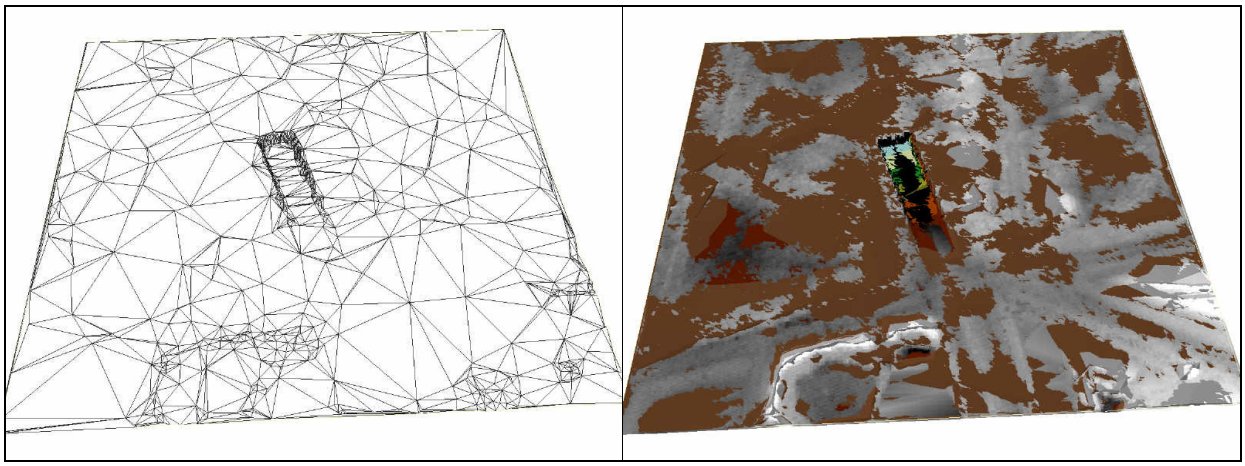

Figure 4.20: [Left] The TIN obtained in Esri Arcview from the DTM with a maximum height error of 50 cm. [Right] Overlay between the DTM (in grey colours) and the TIN.

Before discussing the heuristic workaround, some considerations about the point reduction in the simplification process are necessary. The simplification envelopes algorithm removes all vertices until the remaining triangles cannot be further simplified without intersecting the offset surfaces. In the simple case of a unique constant offset value defining the thickness of the envelope, the simplification algorithm preserves irregularities and simplifies (i.e. smoothes) regular planar shapes as much as possible, within the offset surfaces. Point reduction is, therefore, dependent on the irregularity of the transition surface.

The currently implemented simplification algorithm has a limitation which can be explained with the following example. A perfectly planar low resolution DTM with redundant points is considered, i.e. with more and smaller triangles than strictly needed. A high resolution object is placed in the middle of the DTM. The collar of the high resolution model is perfectly planar, too. The high resolution object remains unchanged: it is inside the inner border. Inside the transition surface, before mesh simplification starts, there is a even higher point density due to the new constrained triangulation (§ 4.5) performed using points from both high and low resolution data. Since $\varepsilon_{\max} = \varepsilon_{DP}^{r} = 0$, due to all DTM triangles being coplanar, no simplification can theoretically take place. However, for any positive $\varepsilon_{\max} \approx 0$, simplification could start resulting in the elimination of *all* redundant points inside the transition surface. Geometrically and topologically the resulting surface would still be correct, the point density would be lower than the high resolution object, *but also lower than the low resolution DTM*. Such a transition surface would not fulfil the initial requirement in terms of gradual point density transition. A possible solution to this problem will be discussed in chapter 6.

Although this critical case remains so far without solution, it represents a quite extreme situation: with data coming from surveying techniques, some operational margin is indeed left that leads to set a certain positive $\varepsilon_{\max} \neq 0$.

The iterative approach is similar to the one explained in § 4.9.2: multiple simplified transition surfaces are obtained starting from the same unsimplified mesh, but with growing values of $\varepsilon_{\max}$. The number of remaining triangles after each simplification is plotted versus the growing $\varepsilon_{\max}$ values. The resulting curve is called “simplification curve” of a transition surface (or, more generally, of a model). An example is given in Figure 4.21.

According to the experimental tests done on real datasets, a value of $\varepsilon_{\max}$ has been empirically found to deliver acceptable results if it is chosen in the interval $0 \leq \varepsilon_{\max} \leq \bar{\varepsilon}_{\max}$, where $\bar{\varepsilon}_{\max}$ represents an upper bound and the point where the simplification curve assumes a fairly linear trend and the simplification process stabilises. In the example in Figure 4.21, this point is at circa $\bar{\varepsilon}_{\max} = 8$ cm.

If the iterative process is matched with a continuous visual inspection of the resulting simplified meshes, this operation can be, of course, of great help.

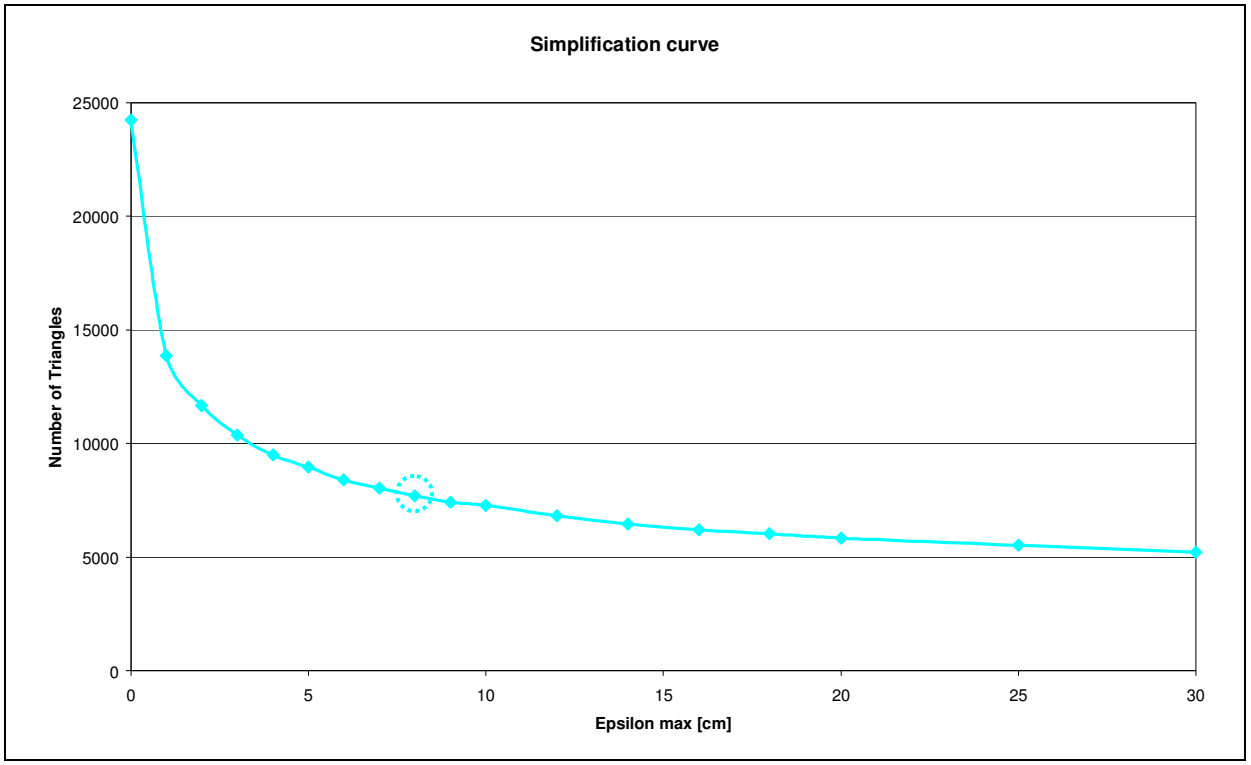

Figure 4.21: Number of remaining triangles in the simplified meshes for growing values of $\varepsilon_{\max}$. The highlighted point $\overline{\varepsilon}_{\max} = 8$ cm represents the upper bound for the interval [0, $\overline{\varepsilon}_{\max}$] where $\varepsilon_{\max}$ is to be chosen.

In the next chapter the results from the application of the presented approach to real datasets will be presented. All three methods described in § 4.9 will be adopted, according to the quality of the data and the quantity of accompanying information.

Chapter 5

EXPERIMENTAL RESULTS

In this chapter experimental results, carried out on models coming from real data, are presented. High resolution objects of varying size and characteristics have been embedded into lower resolution models. Three examples are described: two use data acquired at archaeological sites in Italy, the last one slightly transcends the concept of low resolution DTM and applies the developed methodology in the context of small, highly detailed cultural heritage objects.
In the first example, integration is performed among models acquired with different sensors at different resolutions; in the last two examples, datasets originating from the same sensor but at different resolutions are used.
The examples adopt different strategies to compute the transition surface; depending on accuracy information for each low resolution input dataset, a different approach is chosen and tested. The examples and the results are presented in the following paragraphs in decreasing order of available accompanying information, in a similar way as the strategies to set the $\varepsilon_{\max}$ value are described in § 4.9.
A standard colour coding, as already seen in the previous chapter, will be used to simplify the identification of the models throughout the chapter: the low resolution dataset is coloured in green, the transition surface in azure and the high resolution models in yellow. For consistency and a better understanding, the simplification curves will be coloured in green when obtained from low resolution data, in azure when obtained from the iterative simplification of the transition surface.
All the models are based on triangulated meshes; they have already been previously aligned and are topologically well-behaving: no mesh self-intersections, no foldovers, etc.

5.1 Example 1: Archaeological site in Montegrotto Terme

The datasets refer to an archaeological site located in Montegrotto Terme (near Padova, Italy) and have been acquired by the members of the Laboratorio di Rilevamento e Geomatica in cooperation with the Department of Archaeology, University of Padova. Site excavations began 2006 and are still work in progress; it was necessary to set up a reference network for present and future surveying campaigns. A geodatabase was created to store all created data; even older, previous datasets were added, although very little or no information about data quality was available.
During the surveying campaign in 2006, an integrated approach with GPS, laser scanner, and “classical” techniques was adopted. The result was the creation of a three-dimensional, georeferenced, textured model of the archaeological site [ACHILLI, AGUGIARO, SALEMI *et al.*, 2008].
Before the excavation campaign started, a general DTM of the archaeological site was created by stop-and-go GPS surveying. In Figure 5.1, left, an aerial image of the site is

presented, where the approximate positions of the reference points are marked. In particular, point 5000 was used for the base receiver. About 4200 points were measured with the rover over a surface of about 100x120 m and with acquisition time 30 s. During data post-processing, only points with a maximum global accuracy of 2 cm were kept. From the resulting points, a TIN of the area was obtained, see Figure 5.1, right. The irregular sampling pattern in the northern part is due to inaccessibility of the area during the surveying campaign. The denser sampling in the central part was originally intended to represent with more detail a still unexplored sector, where irregularities in the height profile might indicate underlying structures.

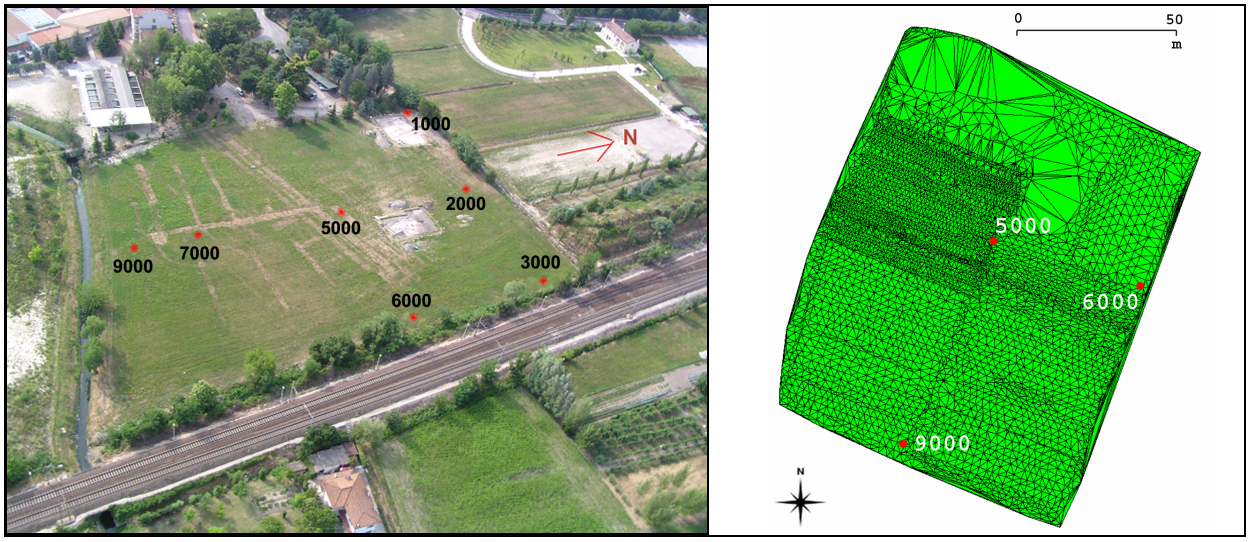

Figure 5.1: Archaeological site in Montegrotto Terme, near Padova, Italy. [Left] Archive view of the archaeological area (around 2004). The reference points and their approximate positions are marked. [Right] Digital terrain model of the archaeological site before the 2006 excavation campaign. The GPS reference point 5000 was used as fixed point for the base receiver. Points used for the TIN were acquired by stop-and-go GPS surveying.

Upon completion of the archaeological campaign 2006 and before the excavated trenches and pits were partially covered again, a laser scanner acquisition was carried out with a Leica HDS 2500 scanner placed on an aerial platform. The global, registered and georeferenced point cloud of the main trench and of the two side pits is shown in Figure 5.2, left.

The proprietary software Cyclone from Leica Geosystems was used to obtain the triangulated meshes from the global point cloud. Some editing was necessary to remove self-intersection and other topological inconsistencies in the high resolution models. The two pits and the southern end of the main excavation trench were chosen as test models to be embedded into the low resolution DTM. The three models are called “esedra pit”, “eastern pit” and “western pit” and are presented in a simple overlay with the DTM to show their relative position, see Figure 5.2, right.

A more detailed, three-dimensional views of each pit are given in Figure 5.3 to Figure 5.5. The overlapping triangles of the DTM are depicted in light transparent green, partially showing the underlying models in yellow.

Regarding the integration procedure for their embedding into the DTM, a description will be given separately for each model.

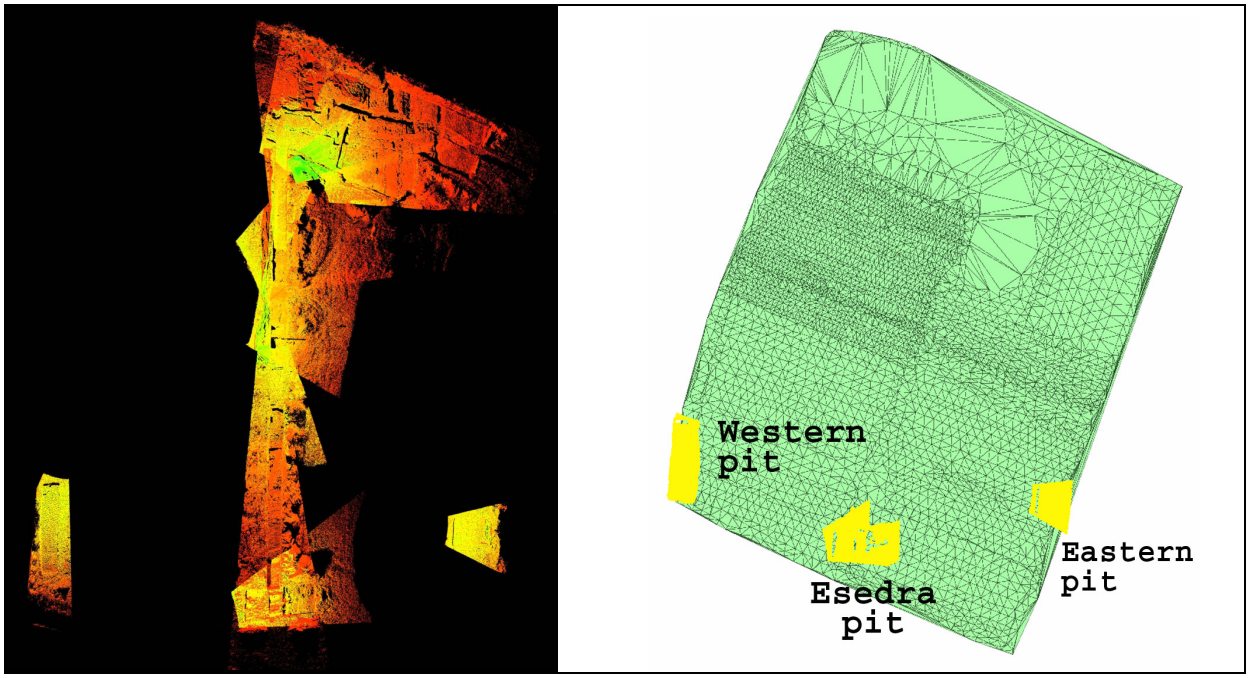

Figure 5.2: [Left] The global, merged and georeferenced point cloud of the entire excavation area after the 2006 archaeological campaign. [Right] Low resolution DTM with overlay of the three models chosen as test datasets to be embedded into the DTM. Left image courtesy of D. BRAGAGNOLO, Università di Padova.

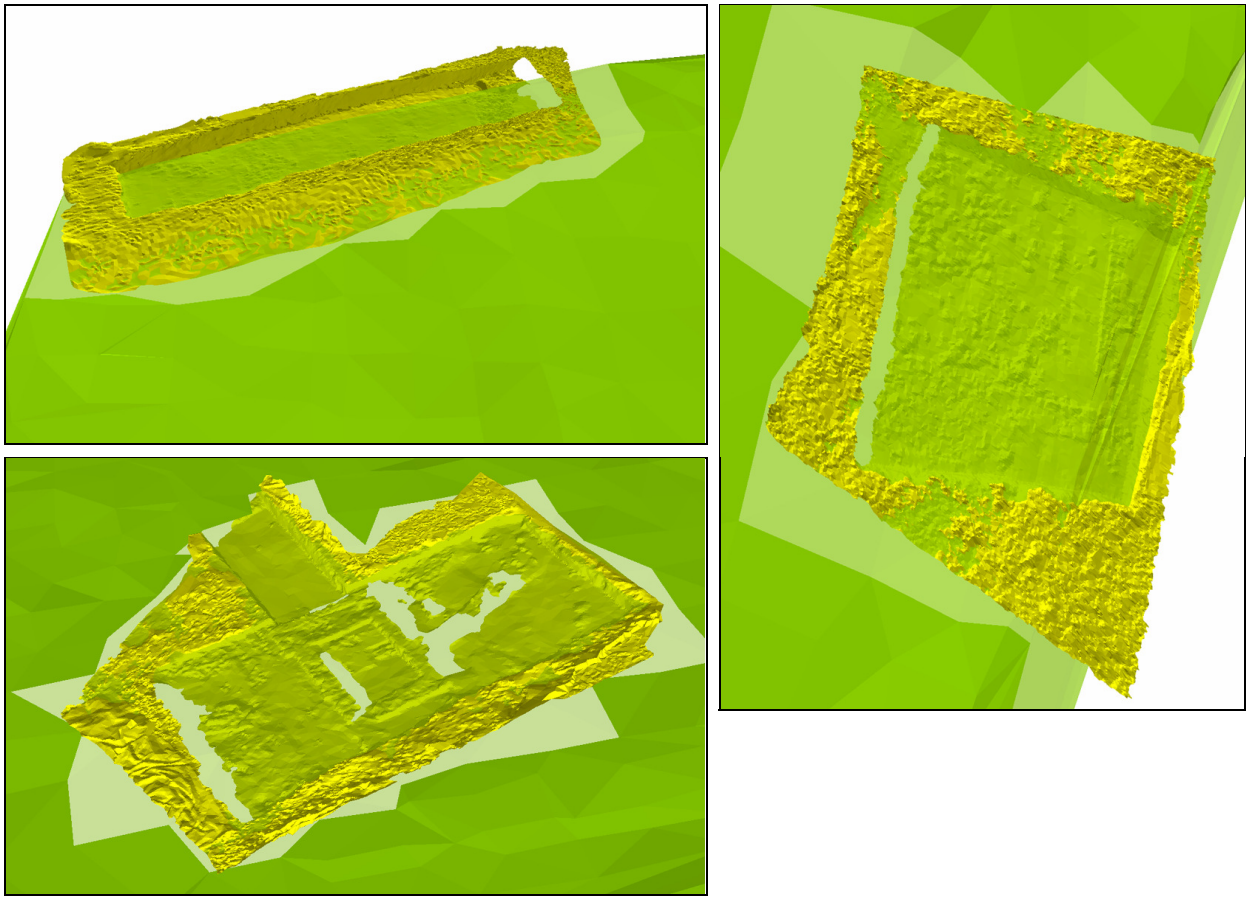

Figure 5.3: Three-dimensional view of the low resolution DTM, in green, with the three high resolution models of the excavation pits to be embedded. For easier visualisation, the overlapping triangles of the DTM are coloured in light transparent green.

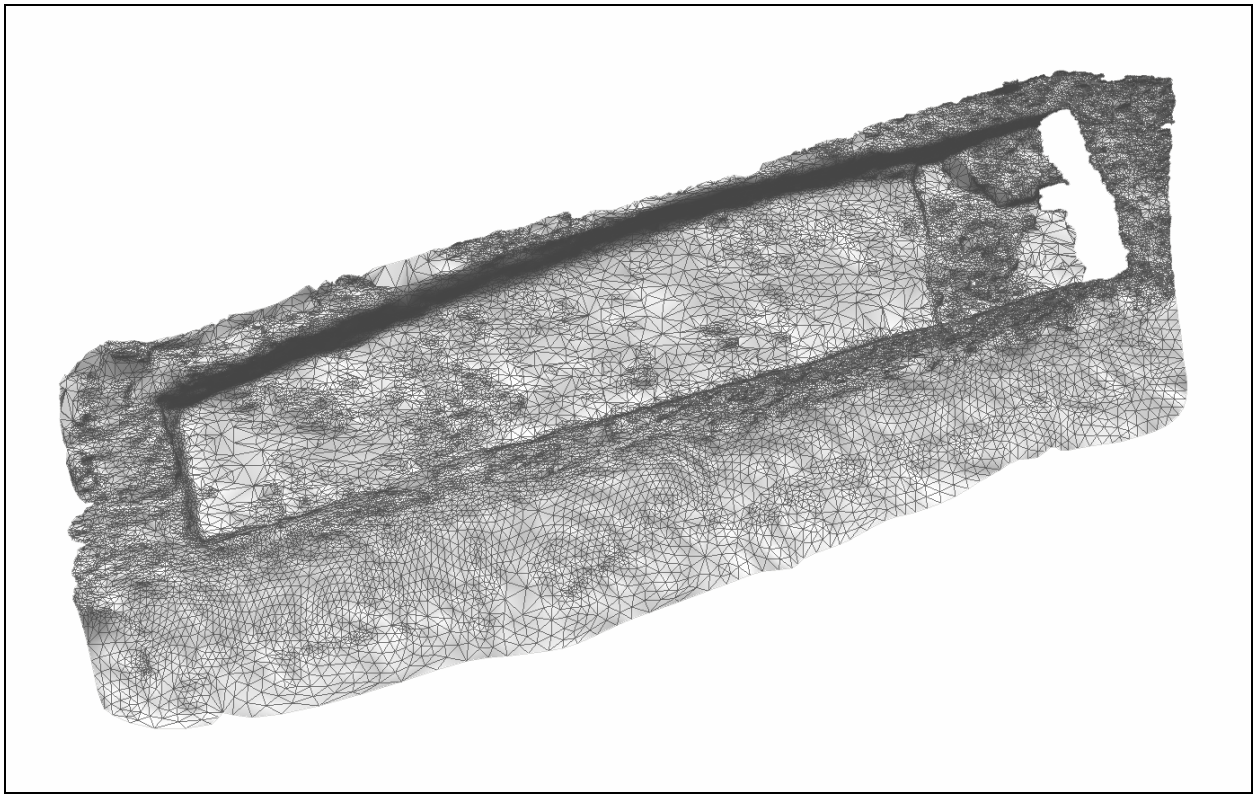

Figure 5.4: The triangulated mesh of the so-called "western pit".

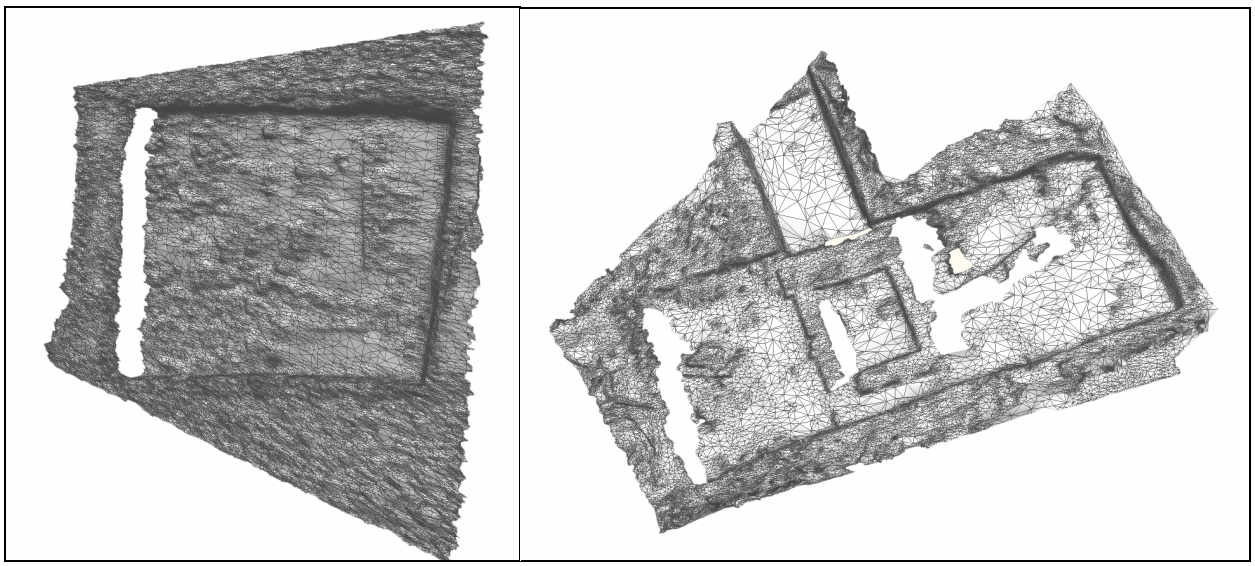

Figure 5.5: The triangulated meshes of the so-called "eastern pit" [left] and "esedra pit" [right].

5.1.1 The esedra pit

The esedra pit represents the ideal case for the integration procedure developed in this work. It is completely surrounded by the low resolution DTM, except for the short side where it was cut from the whole excavation trench. It has a user-defined collar consisting of extra laser scanner data, it has been georeferenced (all related systematic errors are supposed to be negligible) and the mesh is a well-behaving two-manifold surface.

In addition, a global value of accuracy for the low resolution DTM is known from the GPS post-processing ($\sigma_{\max} = 2$ cm); the adopted strategy to set the $\varepsilon_{\max}$ value, which is needed as a maximum global displacement error for the simplified mesh across the transition surface, is $\varepsilon_{\max} = \sigma_{\max} = 2$ cm. The method is explained in § 4.9.1.

The inner and outer border are given as input 3D polylines consisting of high resolution triangles edges, see Figure 5.6, left. Their projection on the xy-plane identifies an annular shape, which is the overlapping zone C as explained in § 4.1, see Figure 5.6, right. The overlapping zone has a varying size, it is limited inwards by the edges of the excavation pit and outwards by nearly all available laser scanner data of the unexcavated terrain.

Inside the overlapping zone C, both high and low resolution data will be used to compute the transition surface, but no height changes are allowed outside the outer border (in the surrounding DTM) and inside the inner border (the excavation pit).

The procedure steps presented in chapter 4 are:

- Step 1: Data import and characterisation,
- Step 2: Projection of the outer border and new triangulation in the zone C,
- Step 3: Distance-weighed height interpolation,
- Step 4: Progressive mesh simplification.

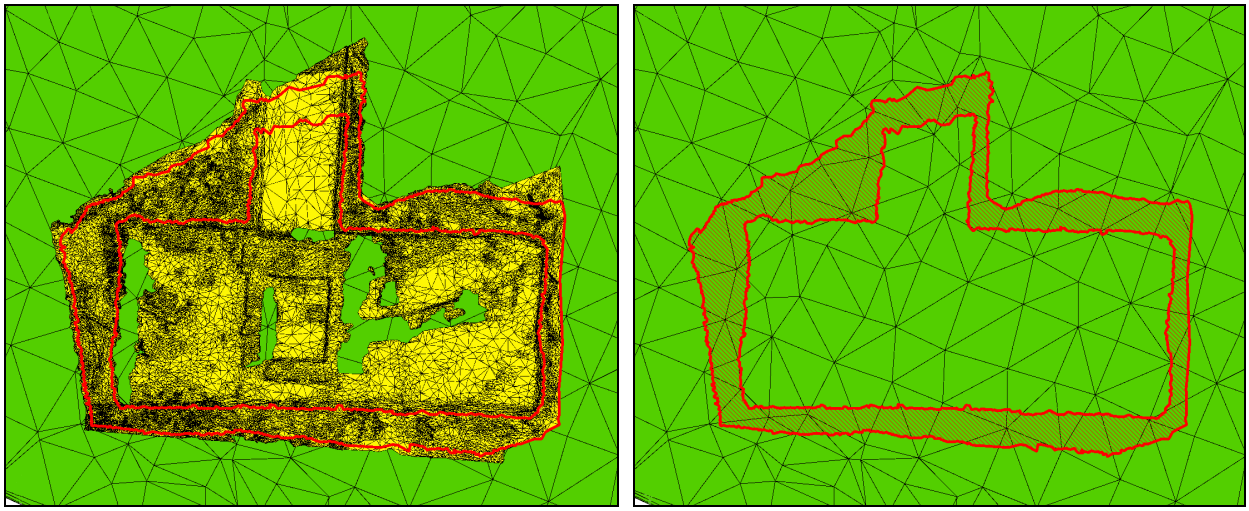

Figure 5.6: [Left] The high resolution model with the inner and the outer border, highlighted in red, are given as input information and identified on the high resolution mesh. [Right] The inner and the outer border, projected vertically on the xy-plane, represent the extents of the overlapping zone C inside which both high and low resolution data will be used to compute the transition surface, see § 4.1.

They will be only quickly referenced here, whereas the visual results accompanying each step will be emphasised.
The DTM and the esedra PLY input file are imported into the PostgreSQL **LR_MODEL** and **HR_MODEL** schemas, respectively, see § 4.4.
Once the overlapping zone C has been obtained, a selection of the input data can be performed in order to import only the data strictly needed for the transition surface. As shown in Figure 5.7, the high resolution triangles in the collar and the low resolution triangles which overlap the annular shape are selected and imported into PostgreSQL schema **MODEL**, thus reducing the quantity of used data.
Upon completion of the data import process, the orientation of the triangles is tested. If any of them have normals pointing "downward", they are flipped. Global and local quality parameter are computed for the low and high resolution data, respectively; the global parameters are presented in Figure 5.7. The height differences are shown in Figure 5.8: the black points have not been used for the Δz statistics since they do not represent the same object in the two models: low resolution as DTM surface, high resolution as trench floor.

Table 5.1: Global quality parameters for the high and low resolution datasets. The values for l_{3D} and Δz are calculated only on imported data. Trench points, see Figure 5.7, are not considered for Δz statistics.

Transition surface DTM - Esedra pit					
Triangle count		*3D edge length l_{3D}*	*[m]*	*Height diff. Δz*	*[cm]*
LR triangles (total)	8163	LR l_{3D} range	0.718 ÷ 3.667	Δz range	-18.9 ÷ 16.7
LR triangles (import)	90	LR avg(l_{3D})	2.272	avg(Δz)	4.8
HR triangles (total)	52370	HR l_{3D} range	0.004 ÷ 1.254	avg($\lvert\Delta z\rvert$)	5.2
HR triangles (import)	23414	HR avg(l_{3D})	0.089	RMSE	6.8

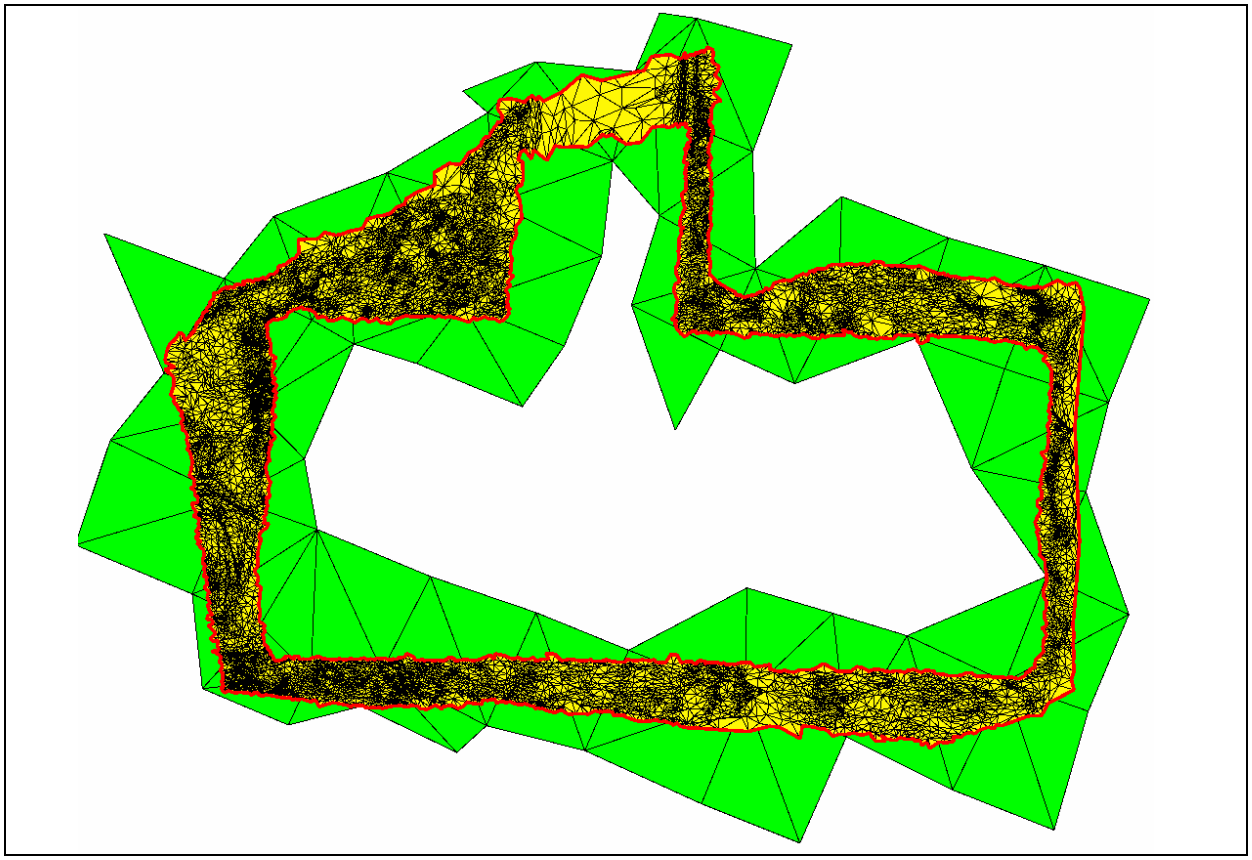

Figure 5.7: Step 1. Data import: by means of spatial queries, only triangles which lie within or intersect the borders (highlighted in red) are selected and imported, thus reducing the global amount of data to process.

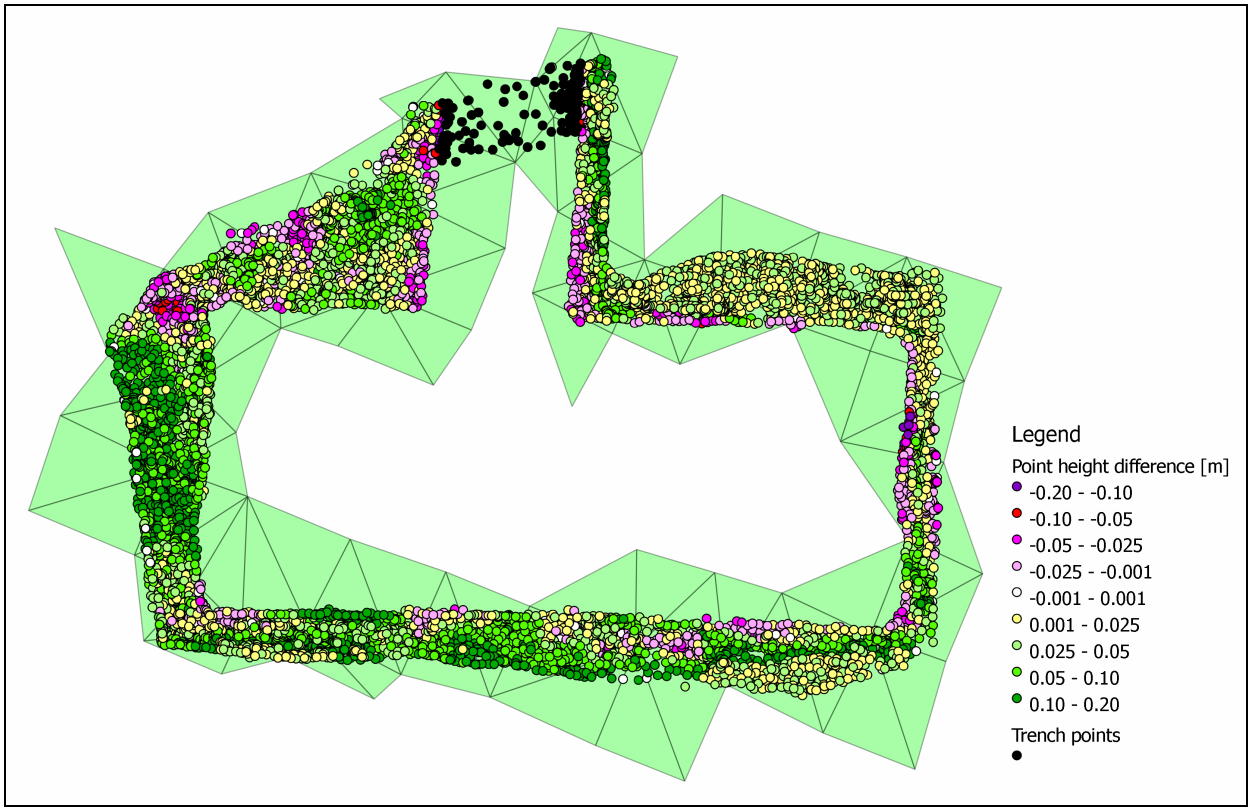

Figure 5.8: Step 1. Height differences between the high resolution model and the DTM. Trench points, coloured in black, are not considered for Δz statistics.

The 3D polyline of the outer border is then projected onto the low resolution DTM. The height values of endpoints of every segment of the polyline are interpolated on surface of the underlying[1] low resolution triangle. If a polyline segment crosses one or more DTM triangle edges, intersection points are calculated and added to the projected 3D polyline, see Figure 5.9, left. The resulting projected polyline is the outer border of the transition surface which connects it to the surrounding DTM. A new constrained DELAUNAY triangulation is performed using all points inside the borders, the borders as breaklines, and the external vertices of the low resolution triangles, see Figure 5.9, right. As explained in § 4.5, the new triangulation adds topological continuity to the two models, however no C0 and C1 conditions are met yet.

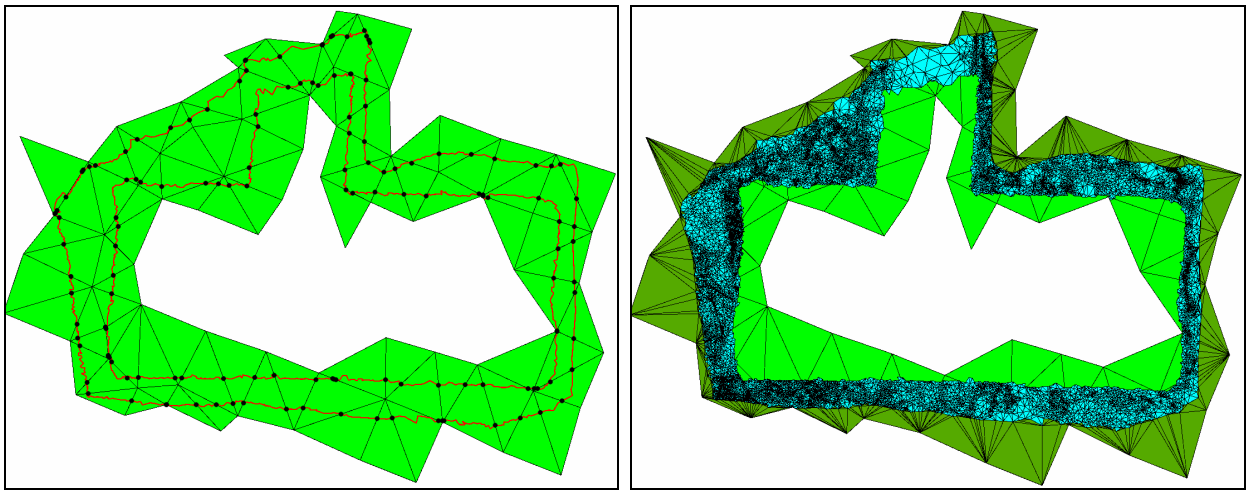

Figure 5.9: Step 2. [Left] Projection of the outer border onto the low resolution DTM. Intersections with the DTM triangle edges are computed and inserted into the resulting 3D polyline. [Right] New constrained DELAUNAY interpolation using all imported points inside the borders, the borders as breaklines and the outer vertices of the low resolution triangles. Low resolution triangles are cut: the light green portions inside the inner border are not used.

The height interpolation model requires the normalised distance parameter from the outer border to the inner border to be computed, since the spline weight function depends on it, see § 4.6.
The planar annular shape of the overlapping zone, on the xy-plane, is used. Progressive buffering is performed inside its shape from the inner border to the outer border, and vice versa, as shown in Figure 5.10. The normalised distance parameter is obtained by combining the two buffer maps, see Figure 5.11.
Once the distance parameter is determined, the height interpolation model adds C0 and C1 conditions to the transition surface, see § 4.6, while a gradual point density reduction is carried out in the successive progressive mesh simplification, see § 4.8. As stated before, the simplified mesh is obtained by setting $\varepsilon_{\max} = \sigma_{\max} = 2$ cm.
In the following Figure 5.12 to Figure 5.17, the procedure steps are shown three-dimensionally in order to achieve a better visualisation of the transition surface until the final integrated model is obtained.

[1] The word “underlying” is used for explanatory simplicity only. Of course, the low resolution triangle can also lie *over* the high resolution point whose height is to be interpolated.

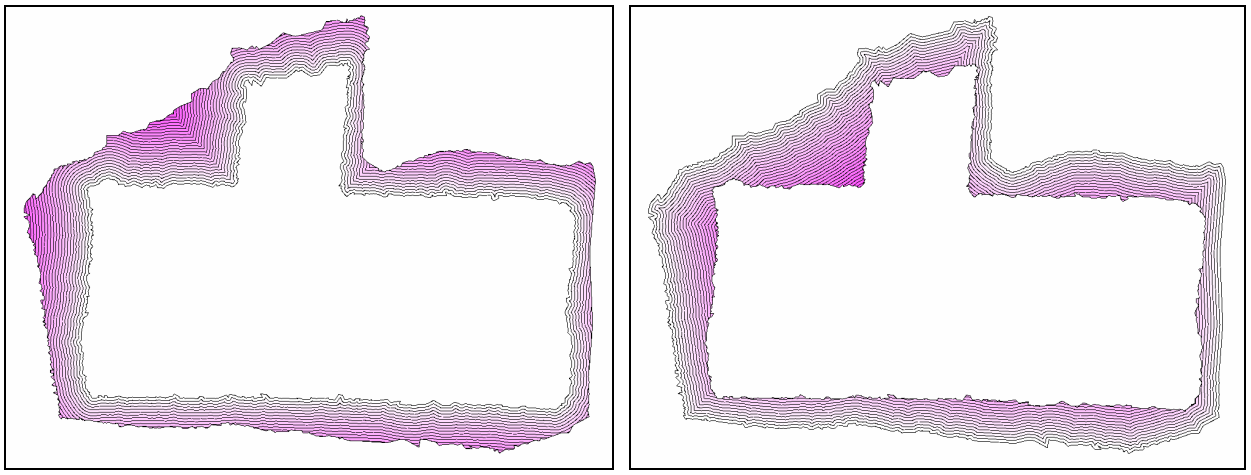

Figure 5.10: Progressive buffering from the inner border to the outer border [left] and from the outer border to the inner border [right].

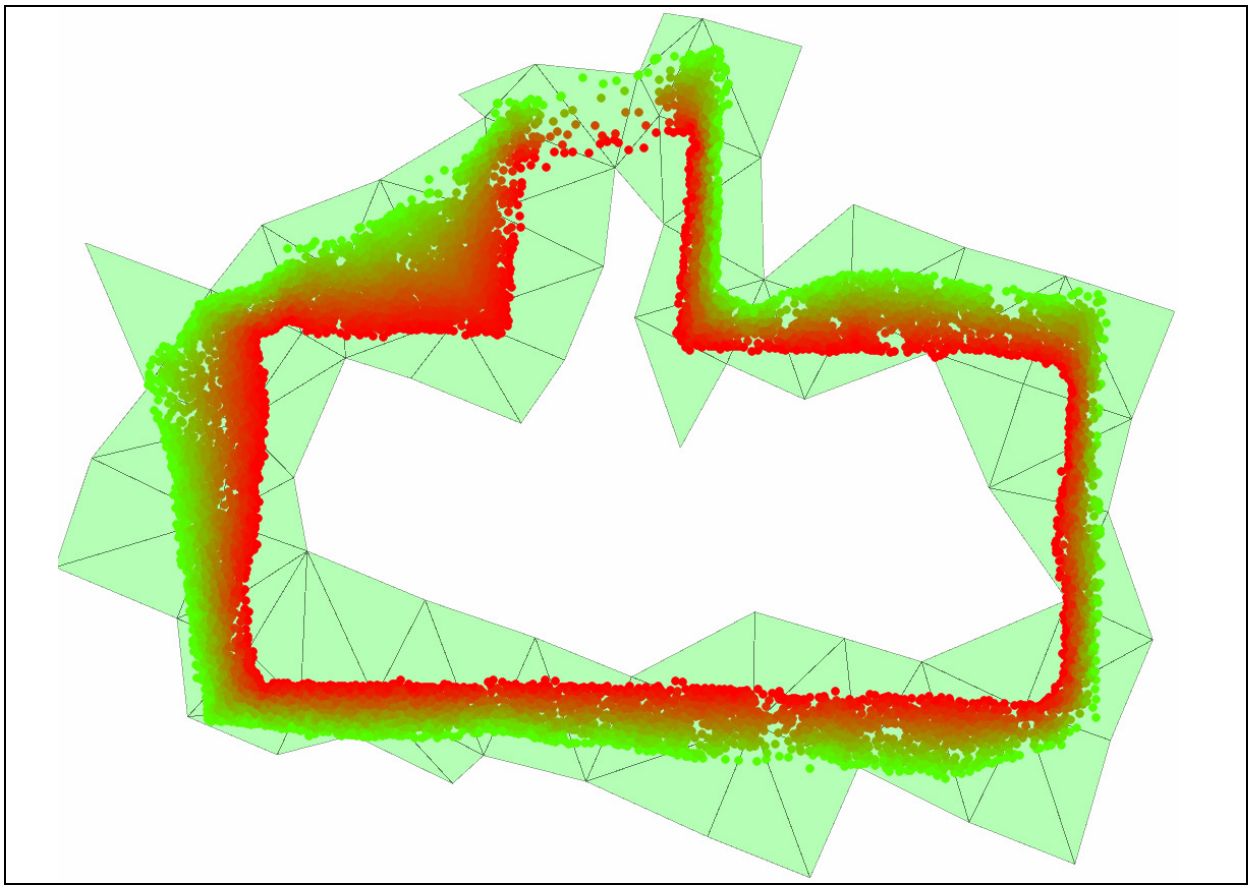

Figure 5.11: By combining the maps in Figure 5.10, it is possible to determine the parameter for the normalised distance from the outer border. It is here represented in a colour range from green (0) to red (1).

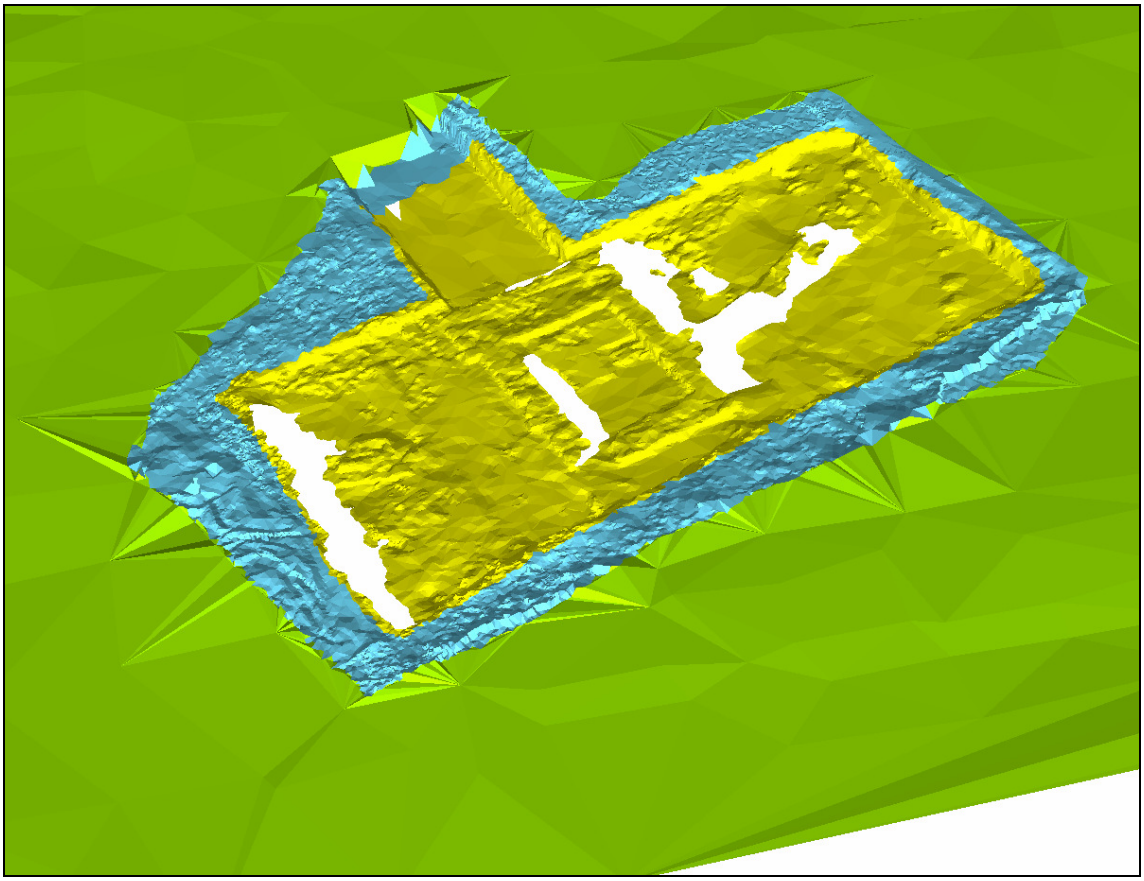

Figure 5.12: Step 2. Transition surface after applying the new constrained DELAUNAY triangulation.

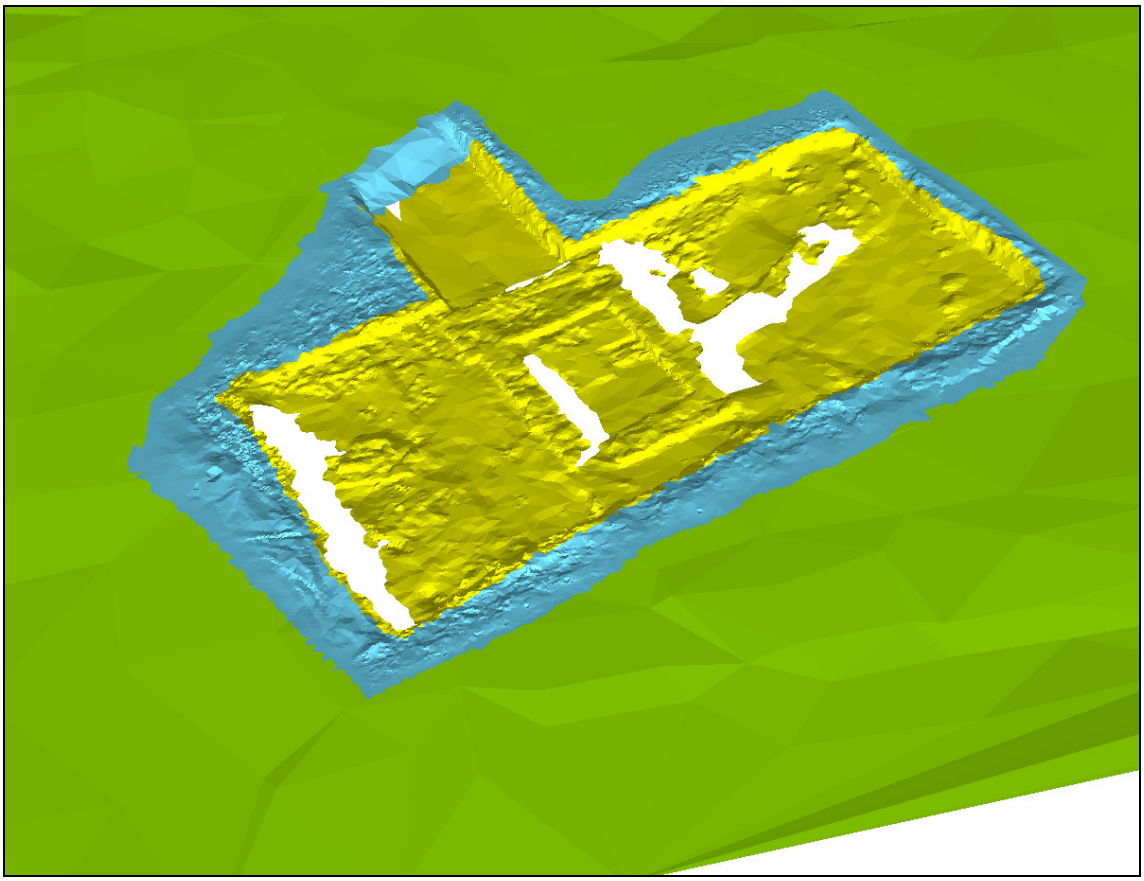

Figure 5.13: Step 3. Transition surface after applying the spline height interpolation model.

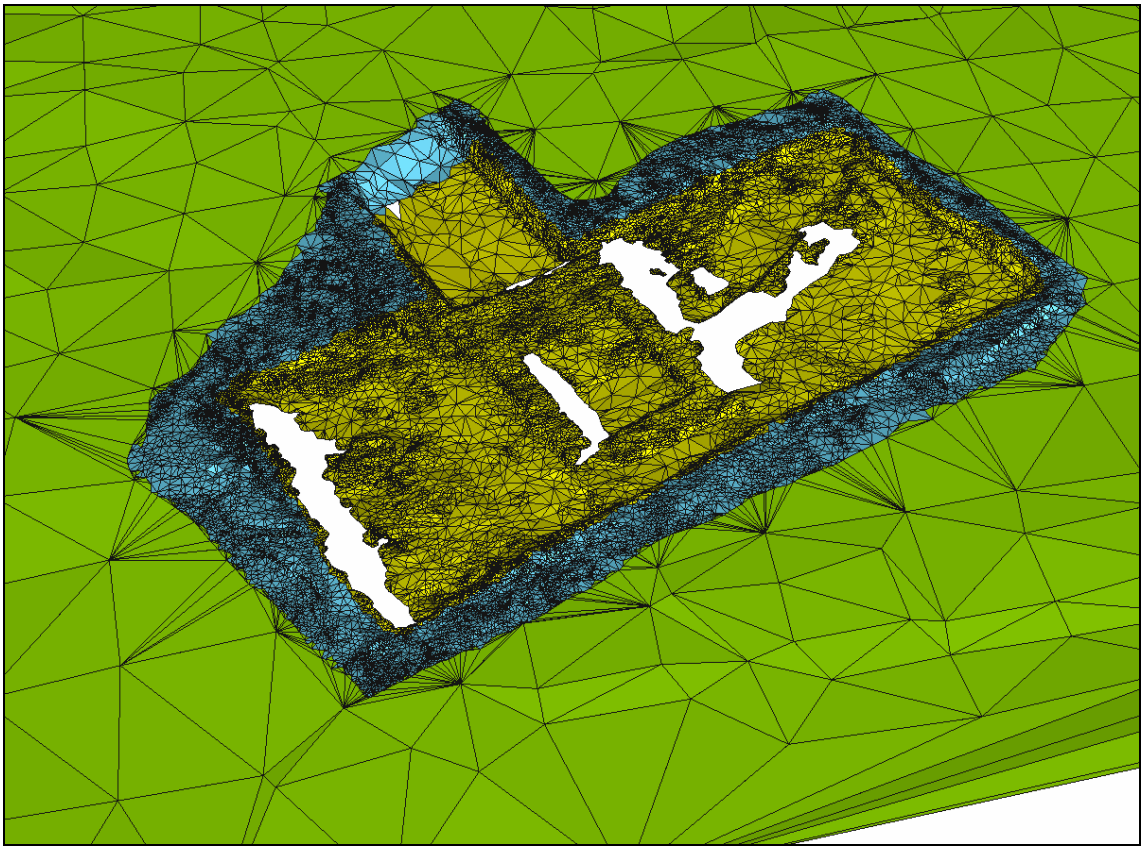

Figure 5.14: Step 3. Transition surface as Figure 5.13, with highlighted triangle edges.

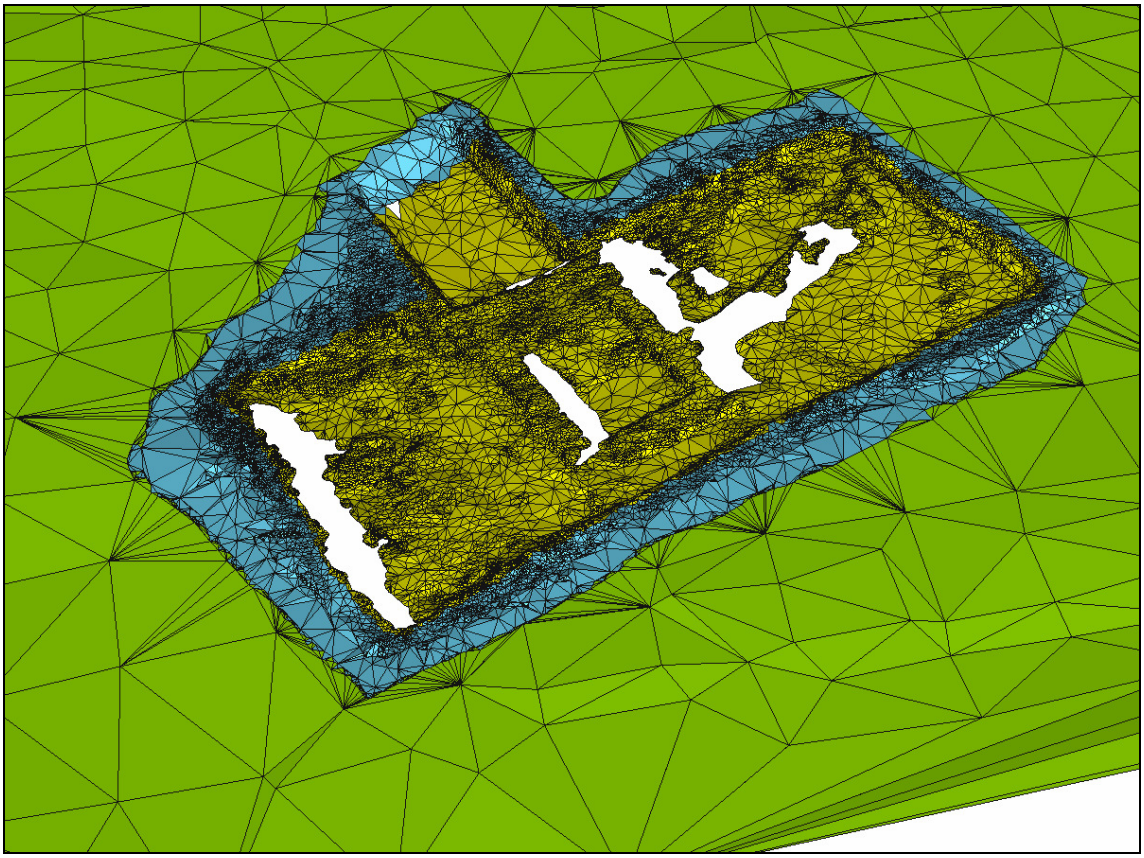

Figure 5.15: Step 4. Simplified mesh in the transition surface, using $\varepsilon_{\max} = 2$ cm. Note the growing size of the triangles in the azure zone.

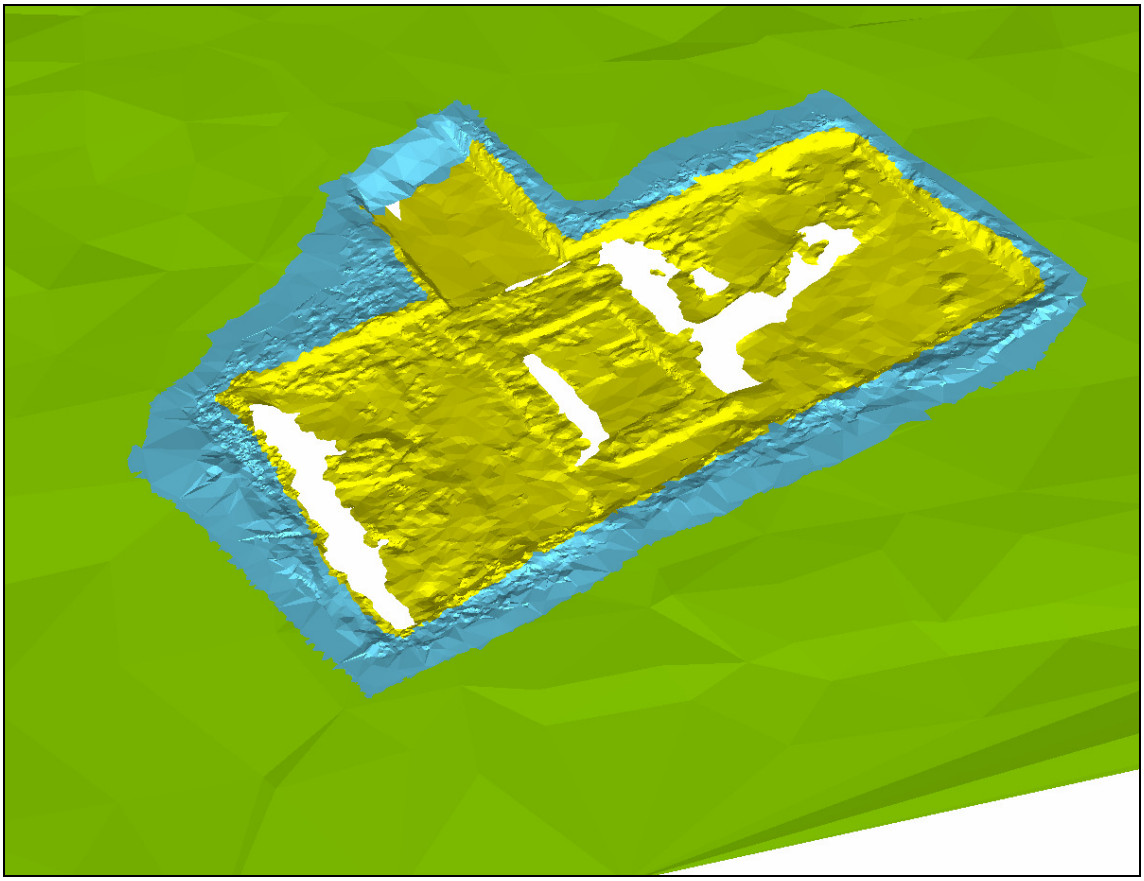

Figure 5.16: Step 4. Transition surface as in Figure 5.15, without highlighted triangle edges.

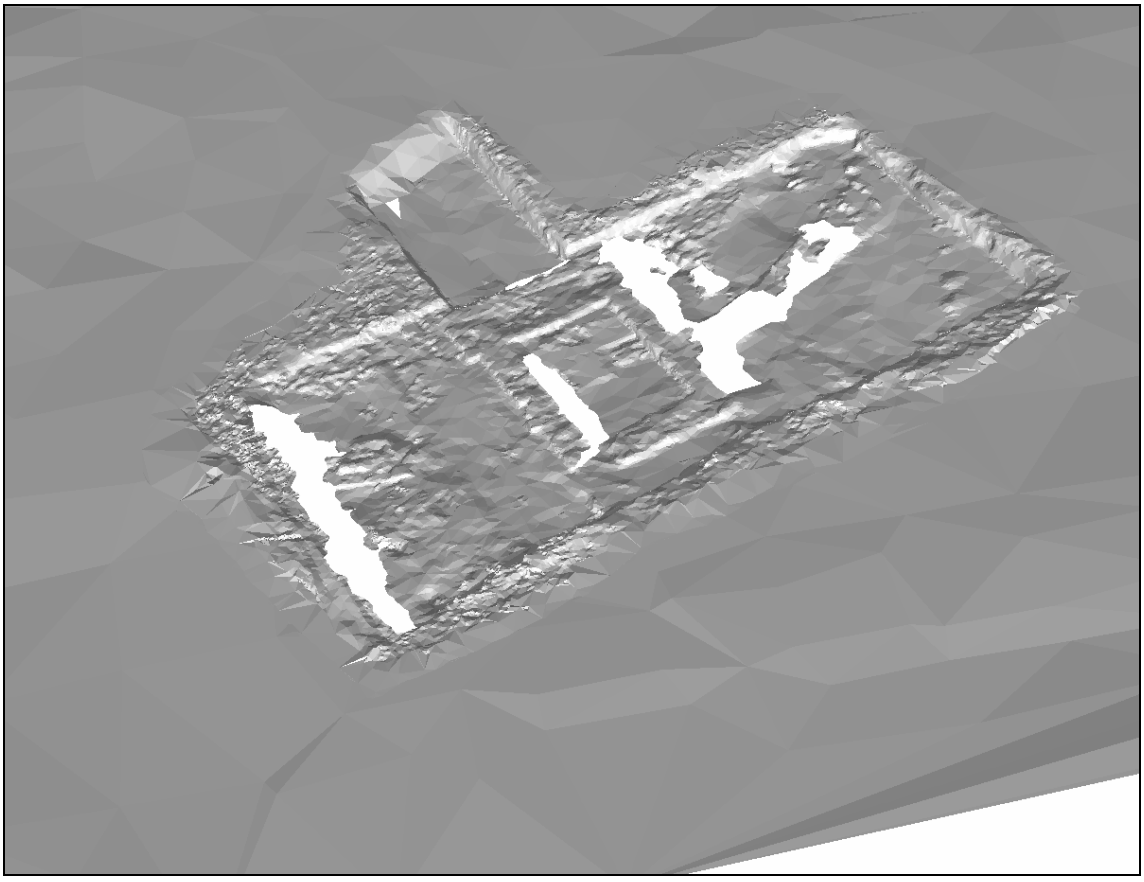

Figure 5.17: The final, integrated model.

5.1.2 The eastern pit

Computation of the transition surfaces for the eastern and the western pit has been carried out in an similar way to the esedra pit. Therefore, only few introductory, distinctive remarks will be done here while the results of the procedure steps will be presented in similar figures and by the accompanying captures.
In general, the eastern and the western pits share common properties with the esedra pit. Their inner and outer borders have been identified similarly to the esedra, though no trench was cut and the inner borders correspond to the pit margins.
The only prominent difference resides in the shape of the overlapping zones. For the eastern pit, part of the high resolution collar lies outside the DTM, see Figure 5.18.
The situation for the western pit is even worse: nearly half of the whole model does not overlap the DTM, see Figure 5.24. These two non-ideal cases have been indeed chosen to test the developed method with unconventional situations.
One question may arise, since it must be decided what happens to the high resolution triangles outside the DTM: they are part of the collar, however no low resolution data is available. In this case their height difference is set $\Delta z = 0$, such that all the following steps can proceed as stated, i.e. their height will simply remain unchanged throughout all remaining steps.
Finally, it remains to be decided whether a mesh simplification must take place within the transition surface outside the DTM or not. Since the ε value can be set on a per point basis, it is possible to set $\varepsilon = 0$ locally and impede mesh simplification in those areas. However, modelling the ε value analogously to the other points over the DTM permits to achieve a more regular – and visually more appealing – transition surface. For the eastern and the western pit, the ε value has been modelled as usual, i.e. mesh simplification in the zones outside the DTM has not been impeded.

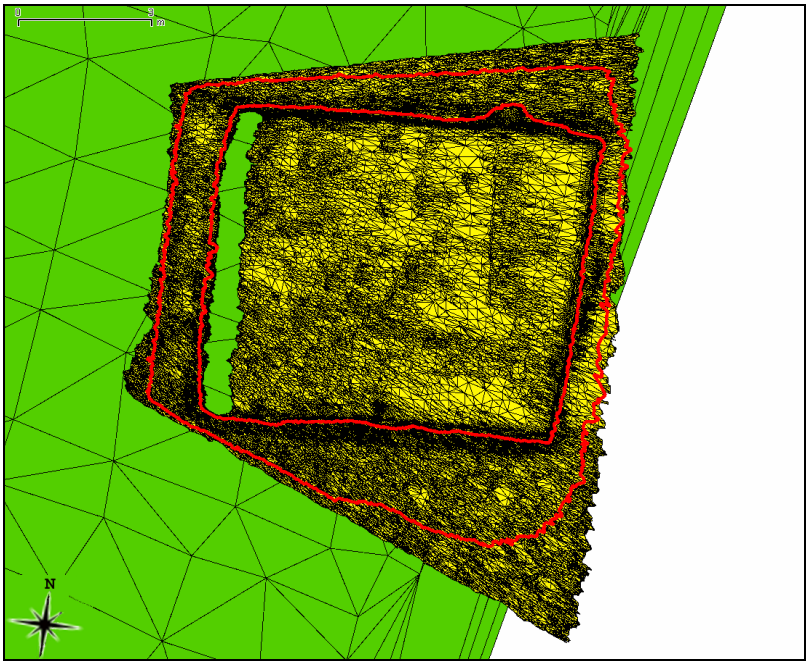

Figure 5.18: The high resolution model with the inner and the outer border (highlighted in red) are given as input information and identified on the high resolution mesh. Unlike the esedra pit, part of the high resolution collar lies outside the DTM.

Table 5.2: Global parameters for the high and low resolution datasets. The values for l_{3D} and Δz are calculated only on imported data. Points outside the DTM are not considered for Δz statistics.

Transition surface DTM - Eastern pit					
Triangle count		*3D edge length l_{3D}*	*[m]*	*Height diff. Δz*	*[cm]*
LR triangles (total)	8163	LR l_{3D} range	0.364 ÷ 110.474	Δz range	-19.1 ÷ 19.7
LR triangles (import)	51	LR avg(l_{3D})	6.327	avg(Δz)	3.1
HR triangles (total)	62344	HR l_{3D} range	0.001 ÷ 0.569	avg($\lvert\Delta z\rvert$)	4.3
HR triangles (import)	24338	HR avg(l_{3D})	0.07	RMSE	5.5

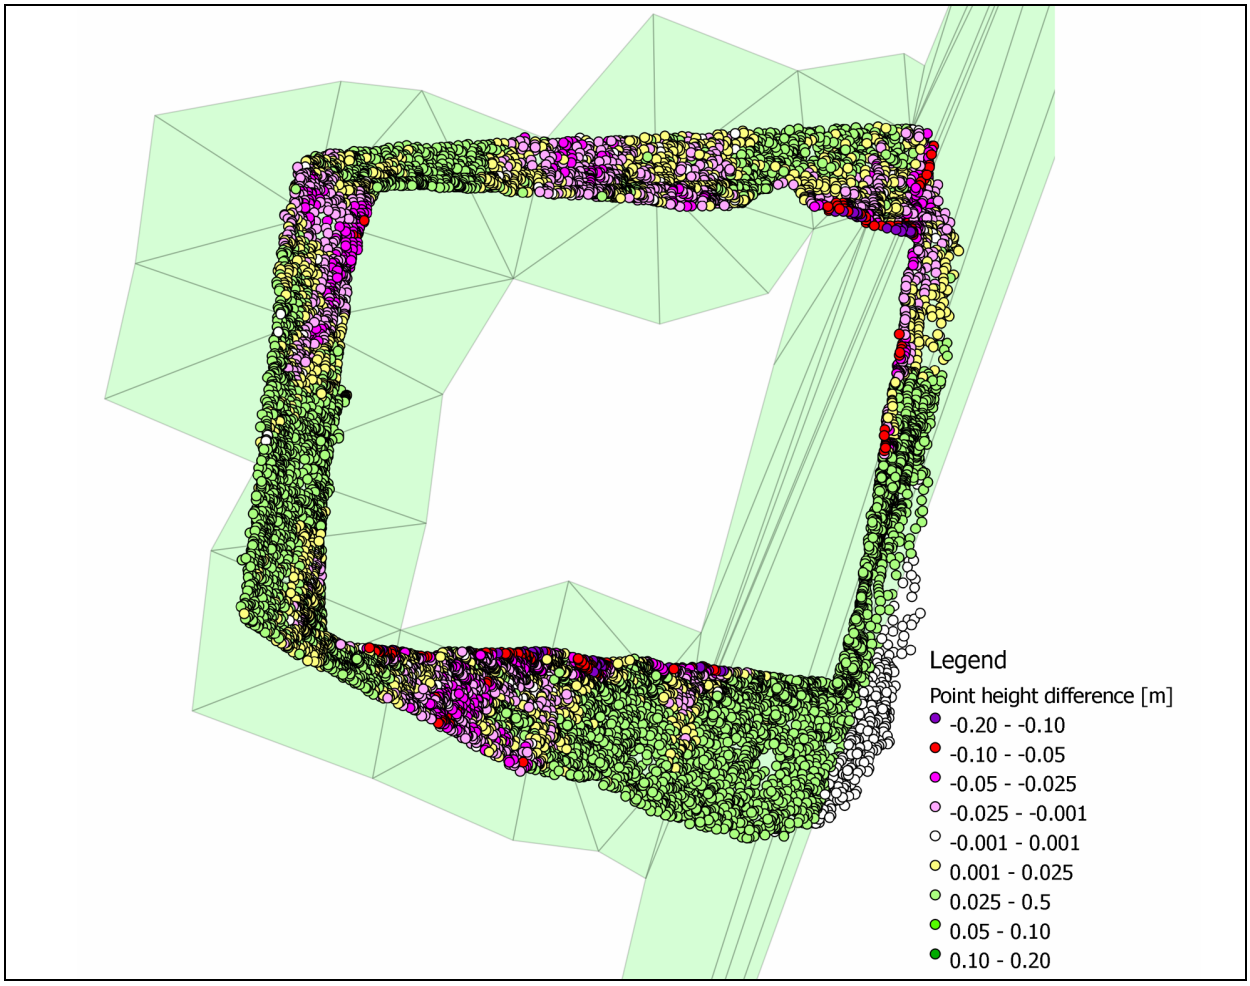

Figure 5.19: Step 1. Height differences between the high resolution model and the DTM. For points lying outside the DTM the height differences are set to zero, such that no height change will be caused by the following height interpolation model.

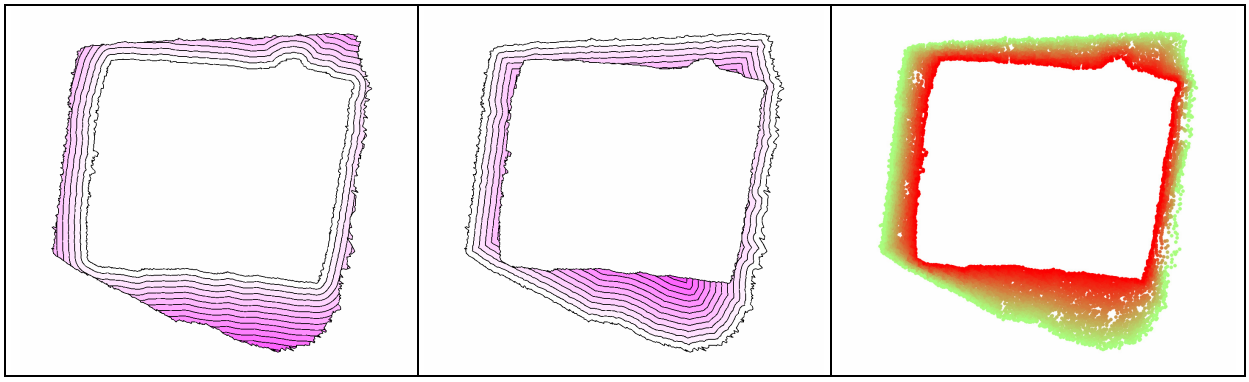

Figure 5.20: Step 2. Progressive buffering from the inner border to the outer border [left] and from the outer border to the inner border [centre] allows to determine the parameter for the normalised distance from the outer border. It is represented in a colour range from green (0) to red (1).

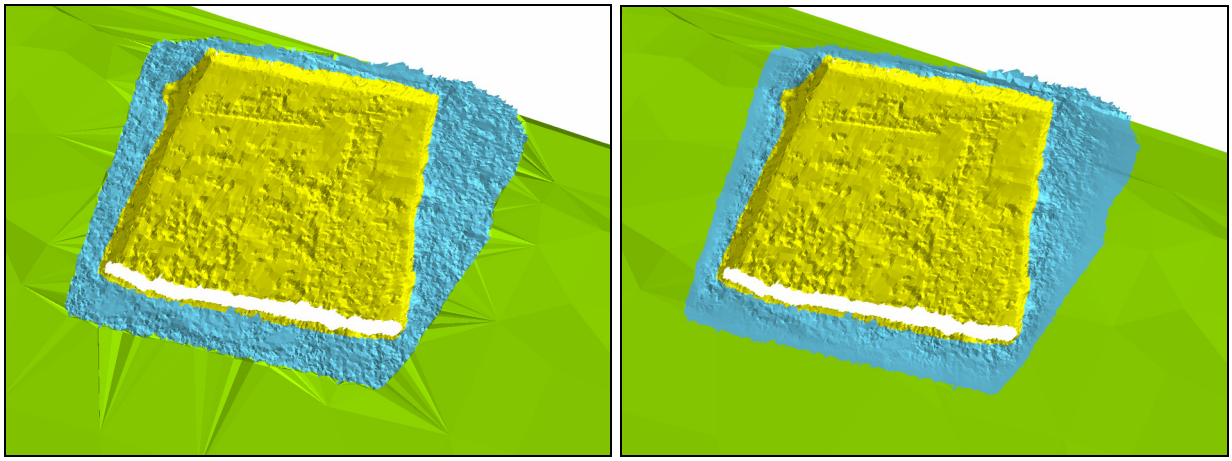

Figure 5.21: [Left] Step 2. New constrained DELAUNAY triangulation in the transition zone. [Right] Step 3. Height interpolation with spline function.

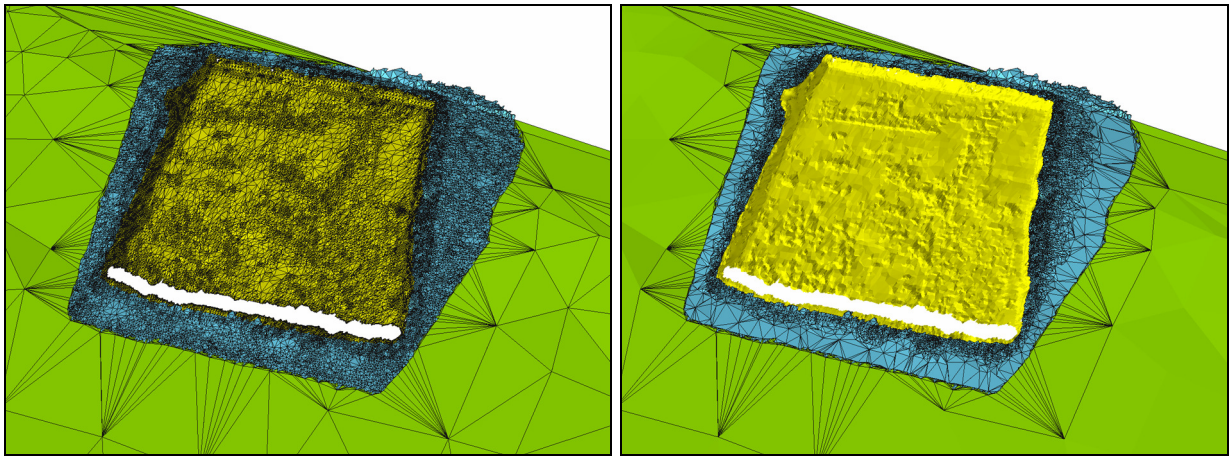

Figure 5.22: Step 4. Transition surface before [left] and after [right] the mesh simplification has been applied with $\varepsilon_{\max} = 2$ cm.

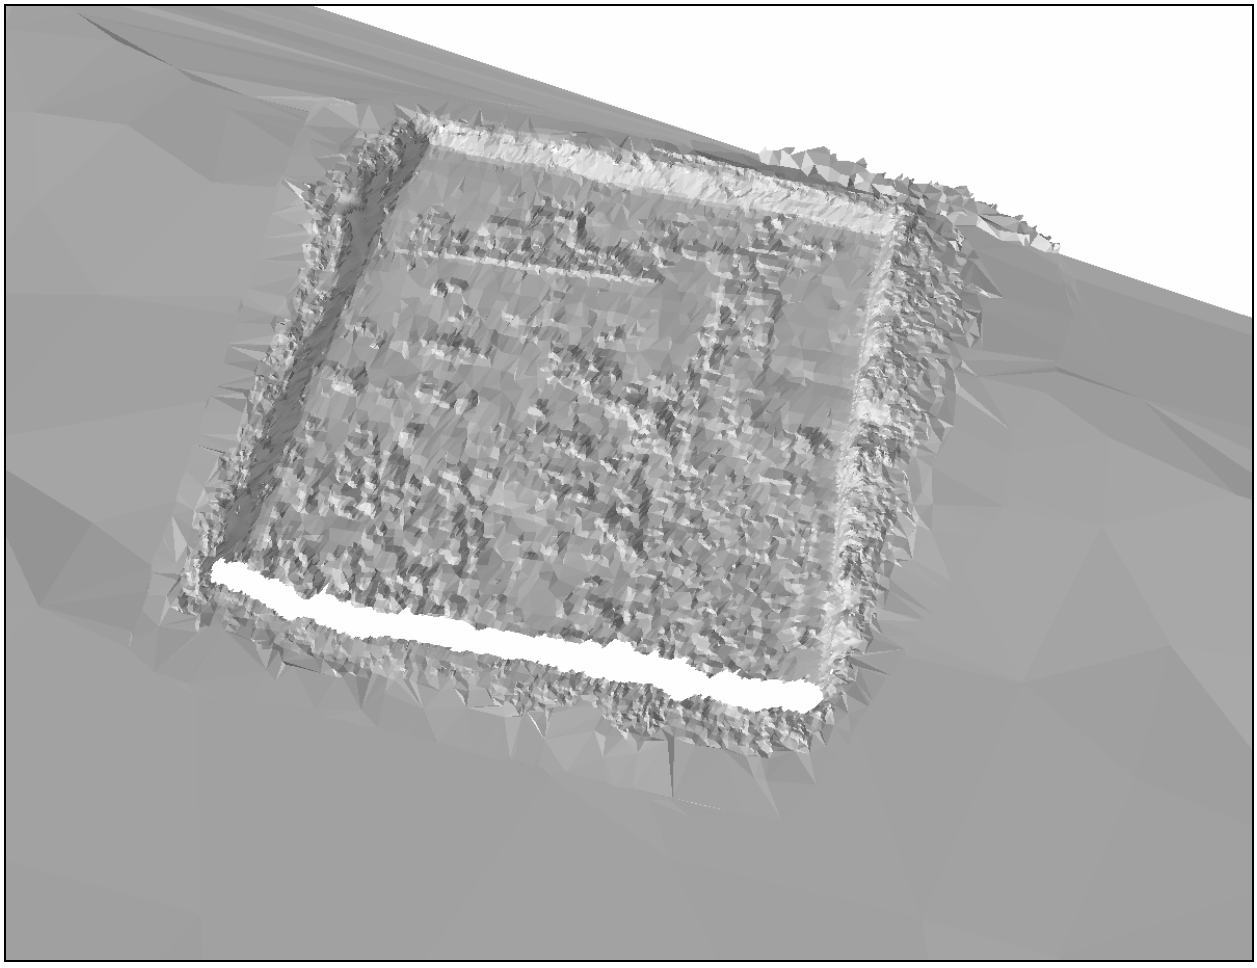

Figure 5.23: The final, integrated model.

5.1.3 The western pit

As stated before, the western pit shares many common properties with the prior two excavation pits. The biggest difference resides in its position: nearly half of the whole model does not overlap the DTM, see Figure 5.24. For all the points outside the DTM, the Δz value has been set to zero, analogously to the eastern pit. This said, the procedure steps remain the same; results are presented in the following figures and the accompanying captions.

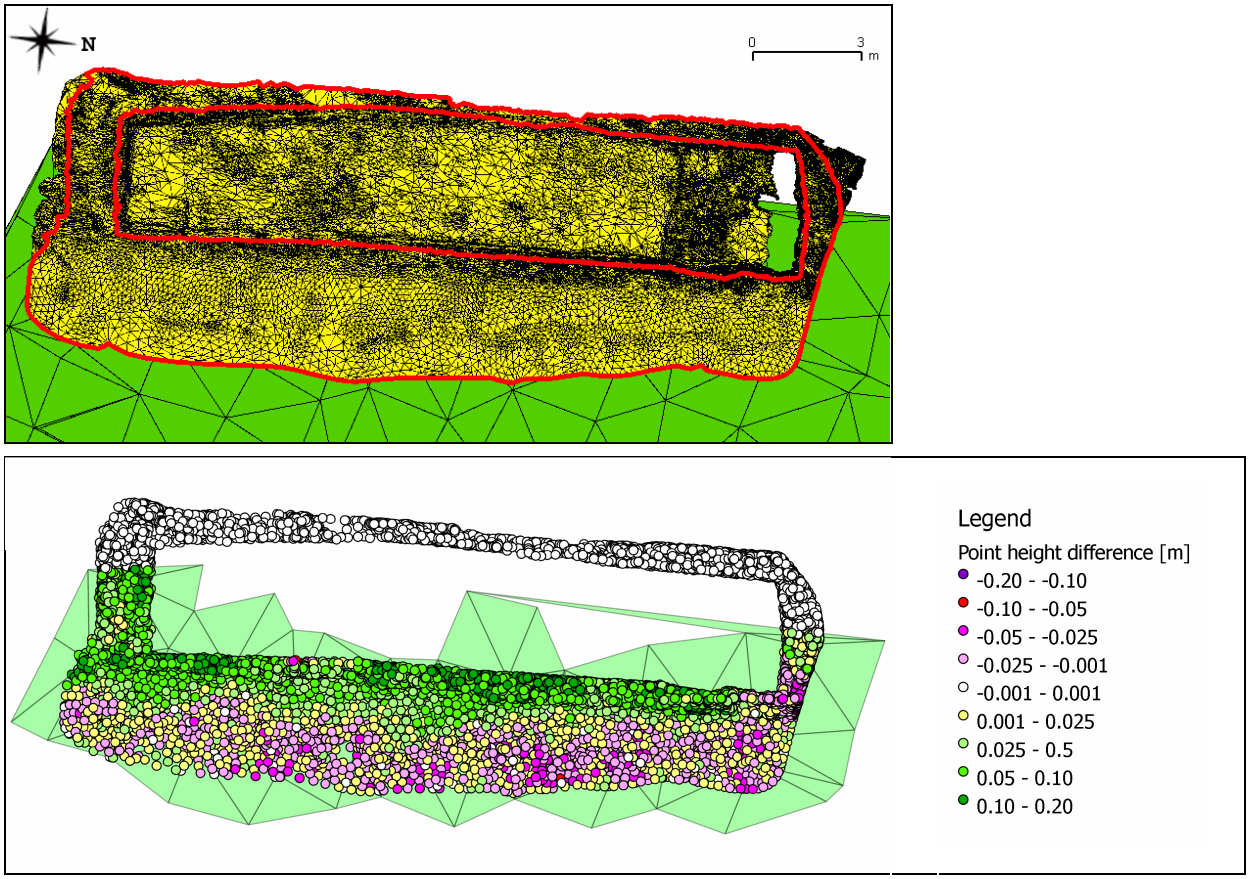

Figure 5.24: [Top] The high resolution model with, highlighted in red, the inner and the outer border, are given as input information and identified on the high resolution mesh. [Bottom] Step 1. Height differences between the high resolution model and the DTM. Since nearly half of the high resolution model lies outside the DTM, height differences are set to zero, such that no height changes will be caused by the following height interpolation model.

Table 5.3: Global parameters for the high and low resolution datasets. The values for l_{3D} and Δz are calculated only on imported data. Points outside the DTM are not considered for Δz statistics.

Transition surface DTM - Western pit					
Triangle count		*3D edge length* l_{3D}	*[m]*	*Height diff.* Δz	*[cm]*
LR triangles (total)	8163	LR l_{3D} range	0.488 ÷ 12.469	Δz range	-8.9 ÷ 24.6
LR triangles (import)	77	LR avg(l_{3D})	2.275	avg(Δz)	4.6
HR triangles (total)	48158	HR l_{3D} range	0.002 ÷ 0.818	avg($\lvert\Delta z\rvert$)	5.1
HR triangles (import)	22212	HR avg(l_{3D})	0.091	RMSE	6.3

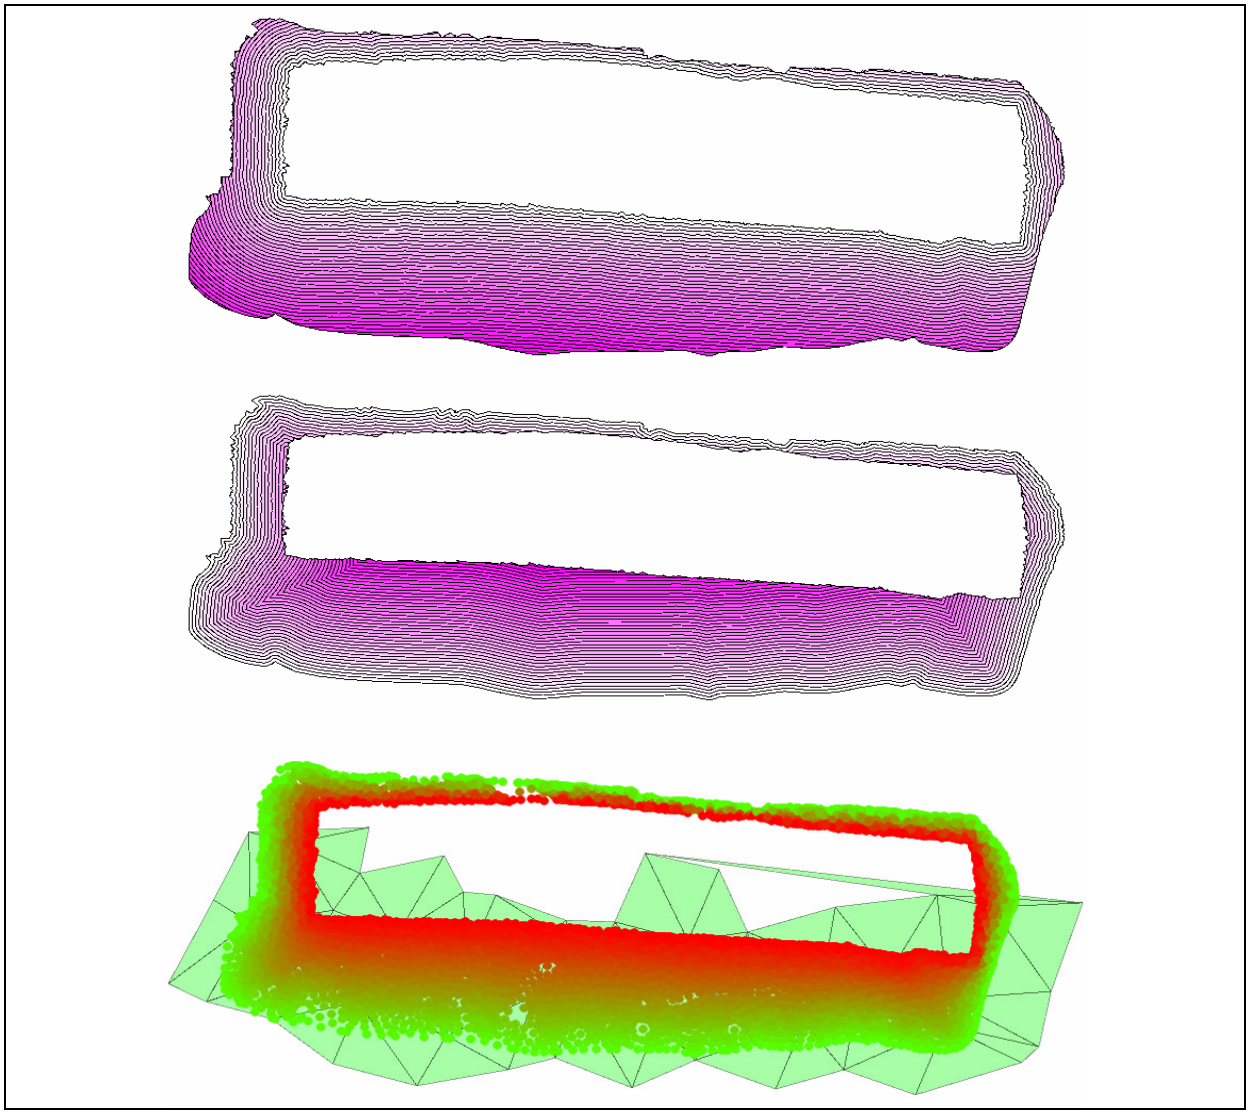

Figure 5.25: Step 3. Progressive buffering from the inner to the outer border [top] and from the outer border to the inner border [centre] allows to determine the parameter for the normalised distance from the outer border. It is represented in a colour range from green (0) to red (1) [bottom].

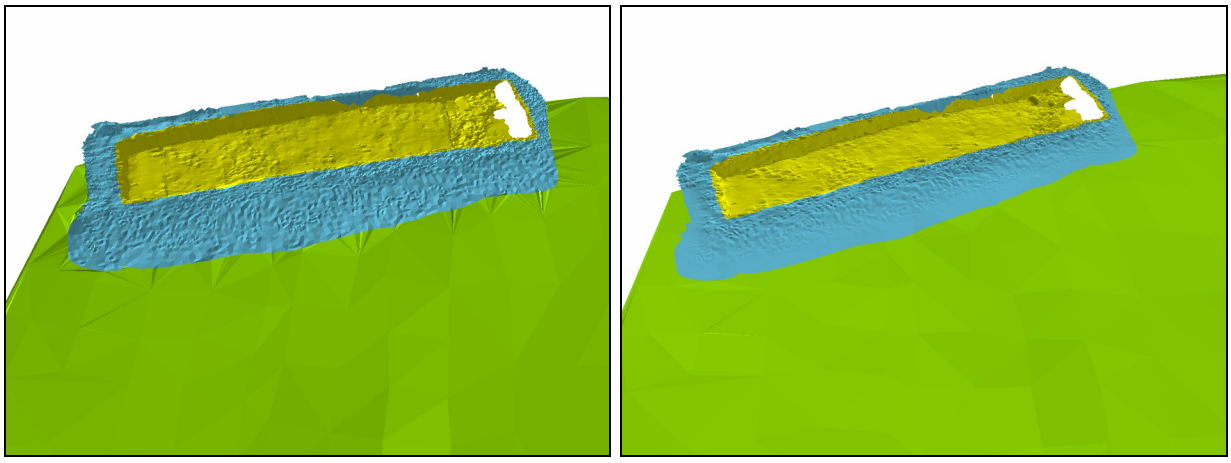

Figure 5.26: [Left] Step 2. New constrained DELAUNAY triangulation in the transition zone. [Right] Step 3. Height interpolation with spline function.

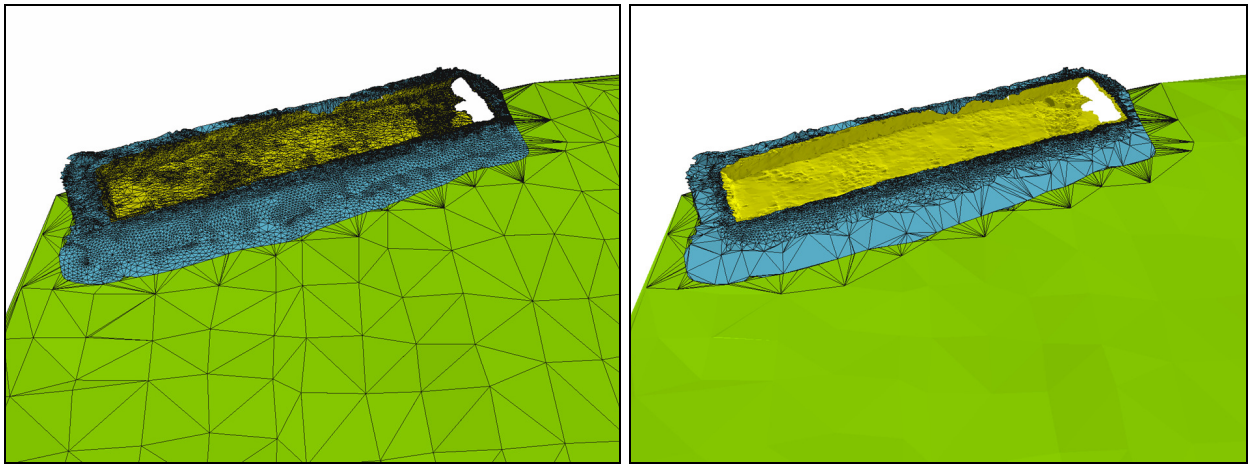

Figure 5.27: Step 4.: Transition surface before [left] and after [right] the mesh simplification has been applied with $\varepsilon_{\max} = 2$ cm.

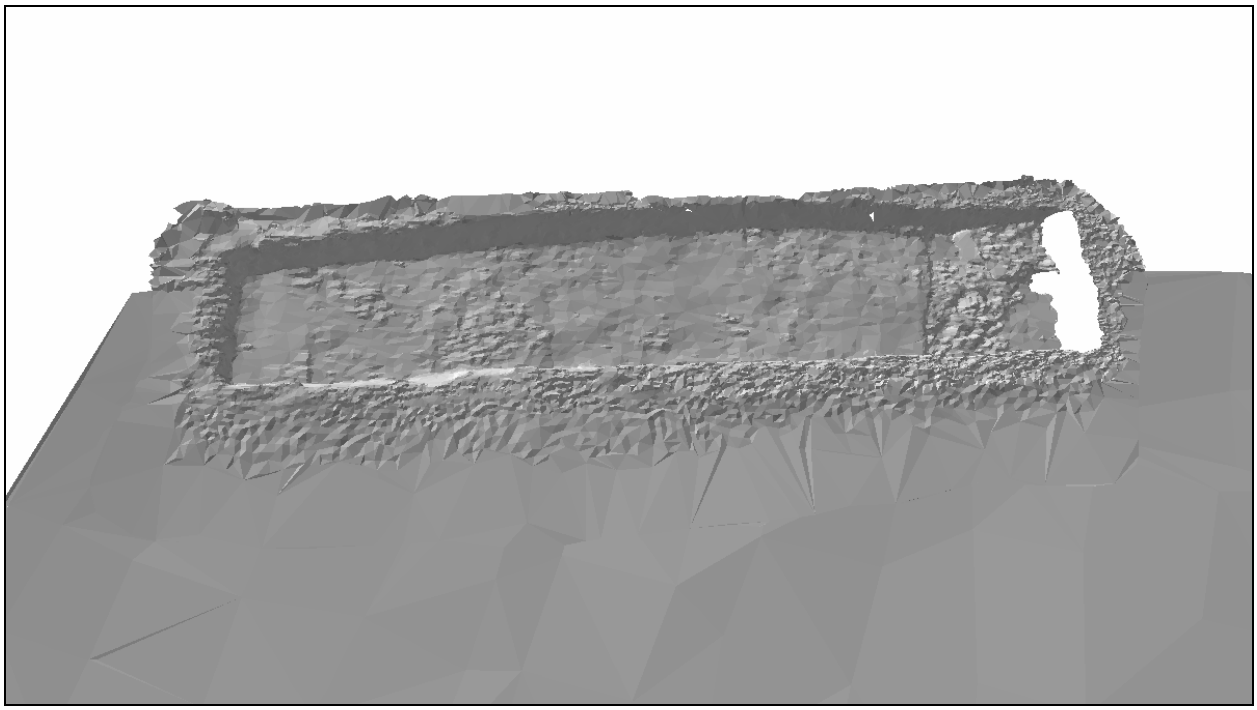

Figure 5.28: The final, integrated model.

5.2 Example 2: Archaeological site in Nora

The datasets used in this example refer to the archaeological site of Nora, an ancient Roman and pre-Roman town located on the southern coast of Sardinia, near Cagliari, Italy.

In summer 2008 an archaeological campaign was started by the Department of Archaeology in cooperation with the Laboratorio di Rilevamento e Geomatica, University of Padova. A model of the whole area was obtained using a Leica HDS 2500 laser scanner. From the global scan, obtained after registration of all point clouds, several models of the temple of Eshmum/Aesculapius' area were created with decreasing levels of resolution. Finally, after manually removing all surveyed architectural elements, a coarse model was created serving as low resolution DTM of the area (about 30x50 m).

As test datasets, three models, and the resulting triangulated meshes, have been chosen: the DTM (mesh-1), a model of the whole excavated area of the temple (mesh-2) and a detail model of the apse in the southern part of the building (mesh-3), as shown in Figure 5.29 and Figure 5.30.

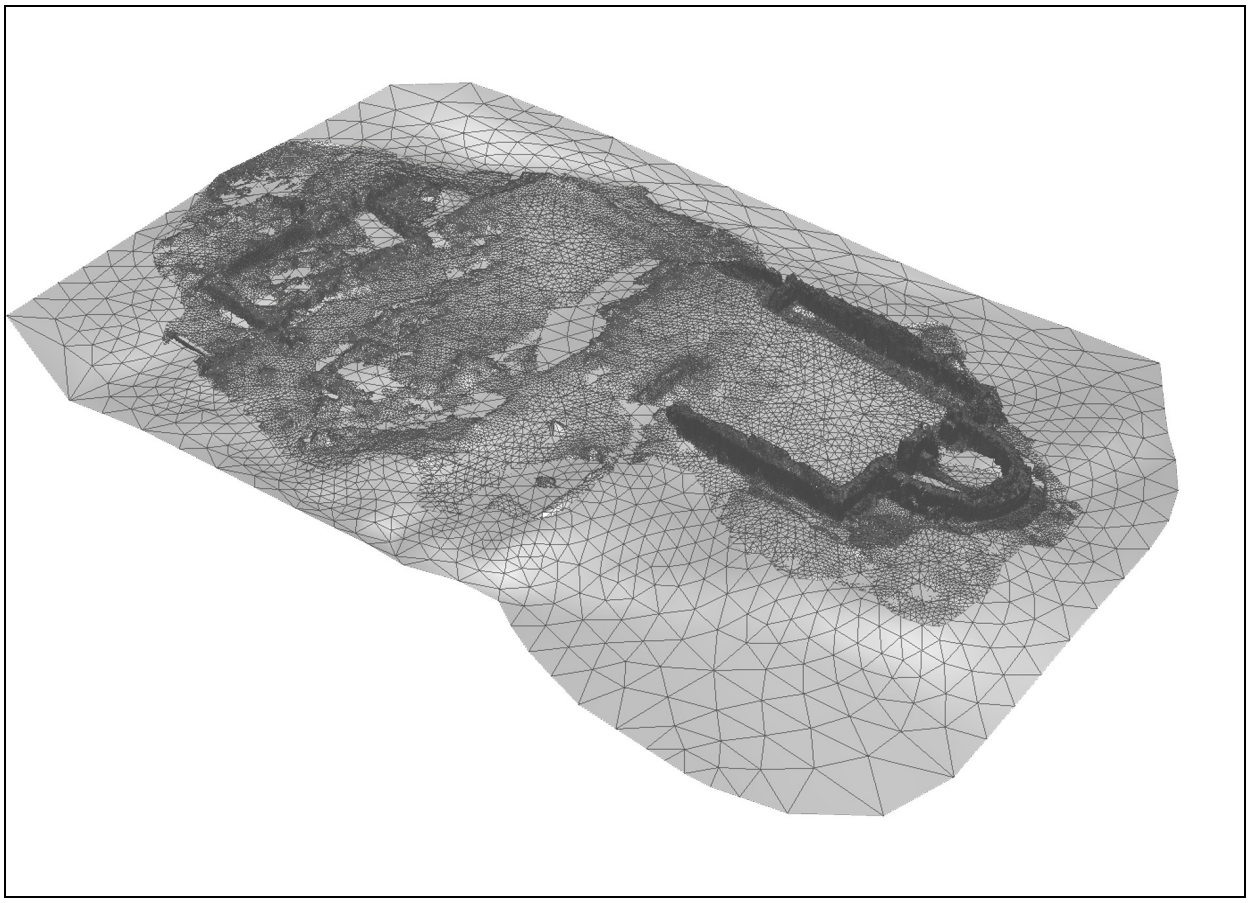

Figure 5.29: Global view of the not yet integrated input models. Datasets refer to the archaeological site of Nora, near Cagliari, Italy.

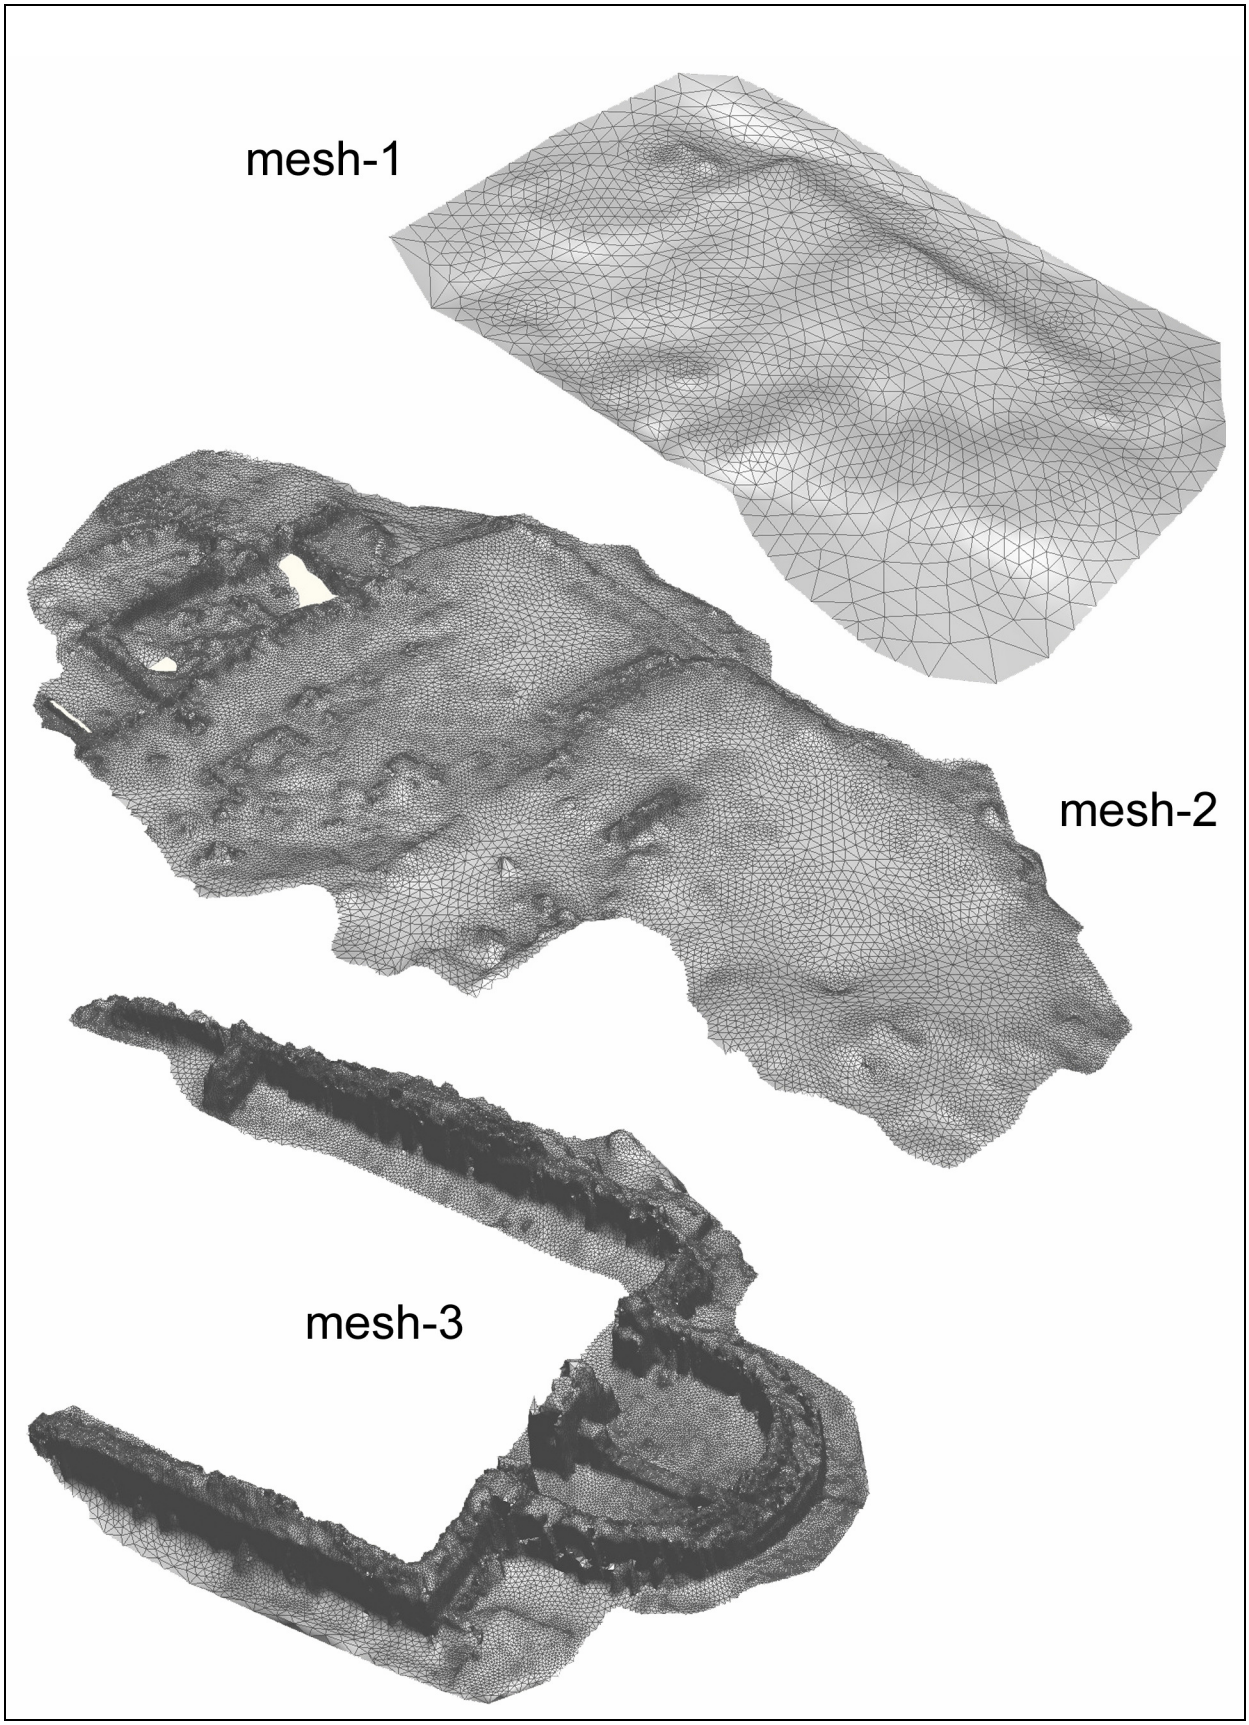

Figure 5.30: Detail view of the input models mesh-1, mesh-2 and mesh-3 (scale varies).

Preliminary editing work has been necessary on the point clouds, so that the obtained meshes comply with the required working hypotheses (i.e. “clean” meshes). The high resolution object depicted in mesh-3 should be visible in a lower resolution version in mesh-2, too, but it has been pruned from mesh-2, although the necessary collar has been preserved. This operation has been chosen for practical reasons (faster operations) and it actually corresponds to having a low resolution model with a hole instead of zone A as explained in § 4.1, thus it does not affect the final result.
Integration among the three models has been carried out in two stages: first, mesh-1 is integrated with mesh-2, obtaining mesh-12, then mesh-3 is embedded on top of the resulting model.

5.2.1 Integration of mesh-1 with mesh-2

The first three steps regarding data import, data characterisation, new triangulation and the height interpolation coincide with the previous example, therefore the graphical results will be presented only briefly in the following pictures with the accompanying captures, see Figure 5.31 to Figure 5.33.
Once the height interpolation model has been obtained, the $\varepsilon_{\max}$ value must be chosen in order to allow for the mesh simplification. However, unlike the previous example, global accuracy of the low resolution model is not known beforehand, so it is not possible to set $\varepsilon_{\max} = \sigma_{\max}$.
The second method, as explained in § 4.9.2, is chosen to set $\varepsilon_{\max} = \varepsilon_{DP}^{r}$. The ε_{DP}^{r} value corresponds to the minimum value which causes the low resolution dataset to be further simplified. A “reverse” simplification is carried out on the imported low resolution triangles: the simplification envelopes algorithm is iteratively applied for a gradually growing $\varepsilon_{\max}$, starting from $\varepsilon_{\max} = 0$.
The number of remaining triangles after each simplification is plotted versus the $\varepsilon_{\max}$ values. If a step in the so-called simplification curve is recognisable, then the associated $\varepsilon_{\max}$ is chosen as connecting value between the outer border of the transition surface and the surrounding DTM. Regarding the imported low resolution triangles from mesh-1, a remarkable drop in the number of triangles takes place at $\varepsilon_{\max} > 1$ cm, this is visible in the graph of Figure 5.34. Therefore, the value $\varepsilon_{\max} = 1$ cm has been adopted for the successive mesh simplification of the transition surface. In the following Figure 5.35 to Figure 5.37 results are presented. The integrated model is called mesh-12.

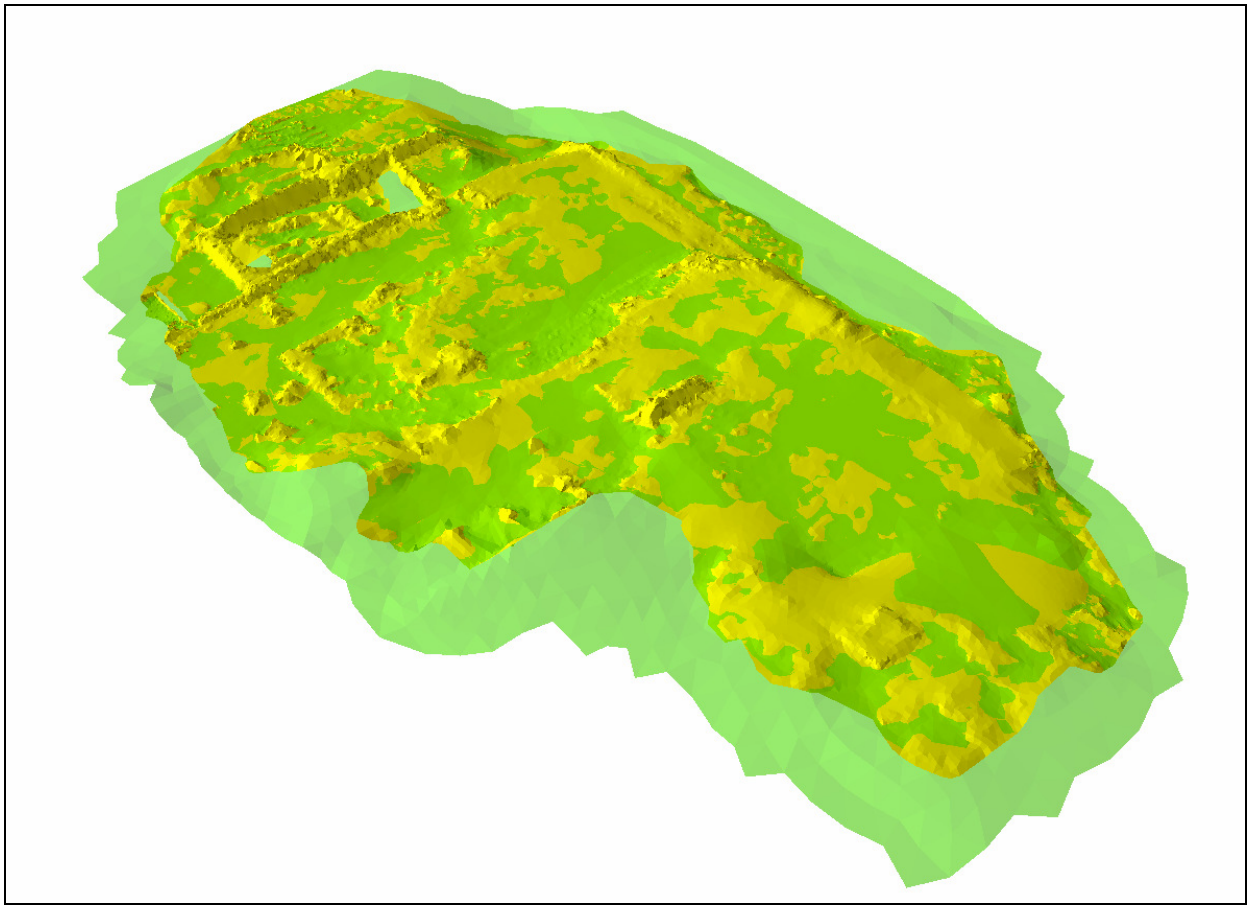

Figure 5.31: Three-dimensional view of mesh-1, in light transparent green, with the overlapping mesh-2, in yellow.

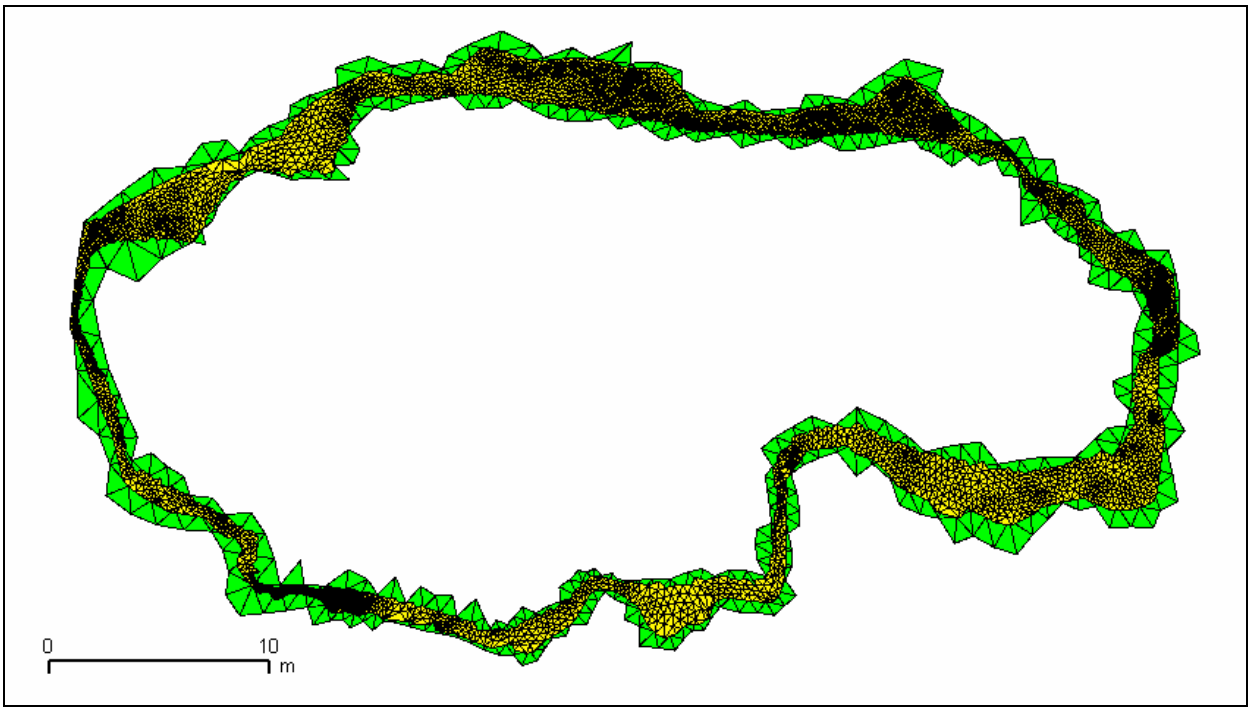

Figure 5.32: Step 1. Data import.

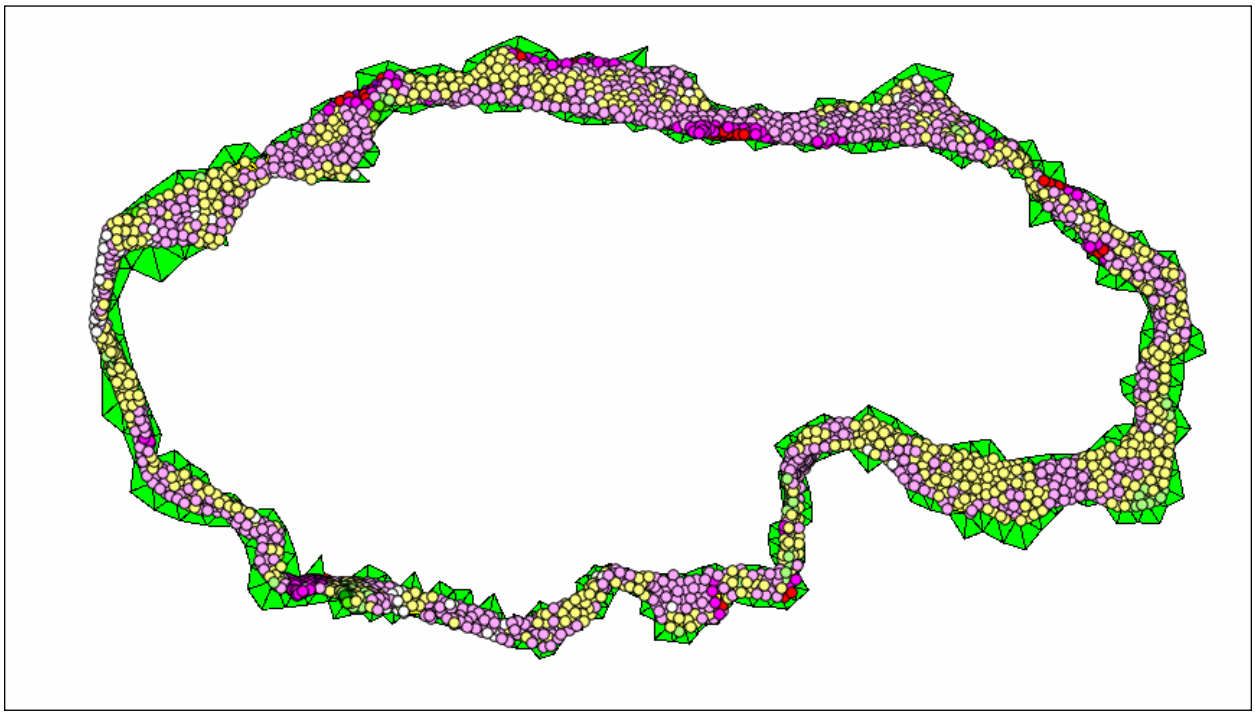

Figure 5.33: Step 1: Height differences between mesh-1 and mesh-2.

Table 5.4: Global parameters for mesh-1 and mesh-2 models. The values for l_{3D} and Δz are calculated only on imported data.

Transition surface mesh-1 - mesh-2					
Triangle count		*3D edge length* l_{3D}	*[m]*	*Height diff.* Δz	*[cm]*
LR triangles (total)	3748	LR l_{3D} range	0.392 ÷ 2.336	Δz range	-43.9 ÷ 42.8
LR triangles (import)	848	LR avg(l_{3D})	0.905	avg(Δz)	-0.5
HR triangles (total)	83053	HR l_{3D} range	0.012 ÷ 1.058	avg($\lvert\Delta z\rvert$)	4.5
HR triangles (import)	9653	HR avg(l_{3D})	0.195	RMSE	7.2

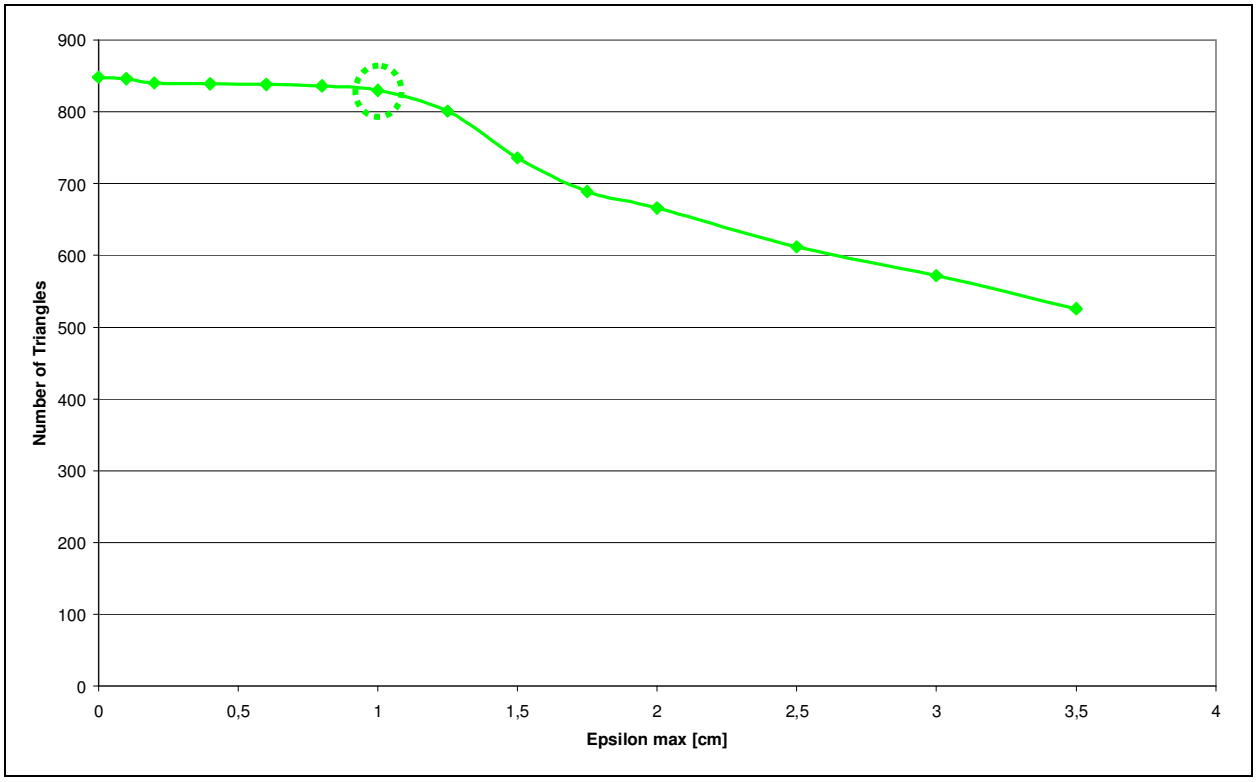

Figure 5.34: Simplification curve of the imported low resolution dataset. The threshold is chosen at $\varepsilon_{\max} = 1$ cm.

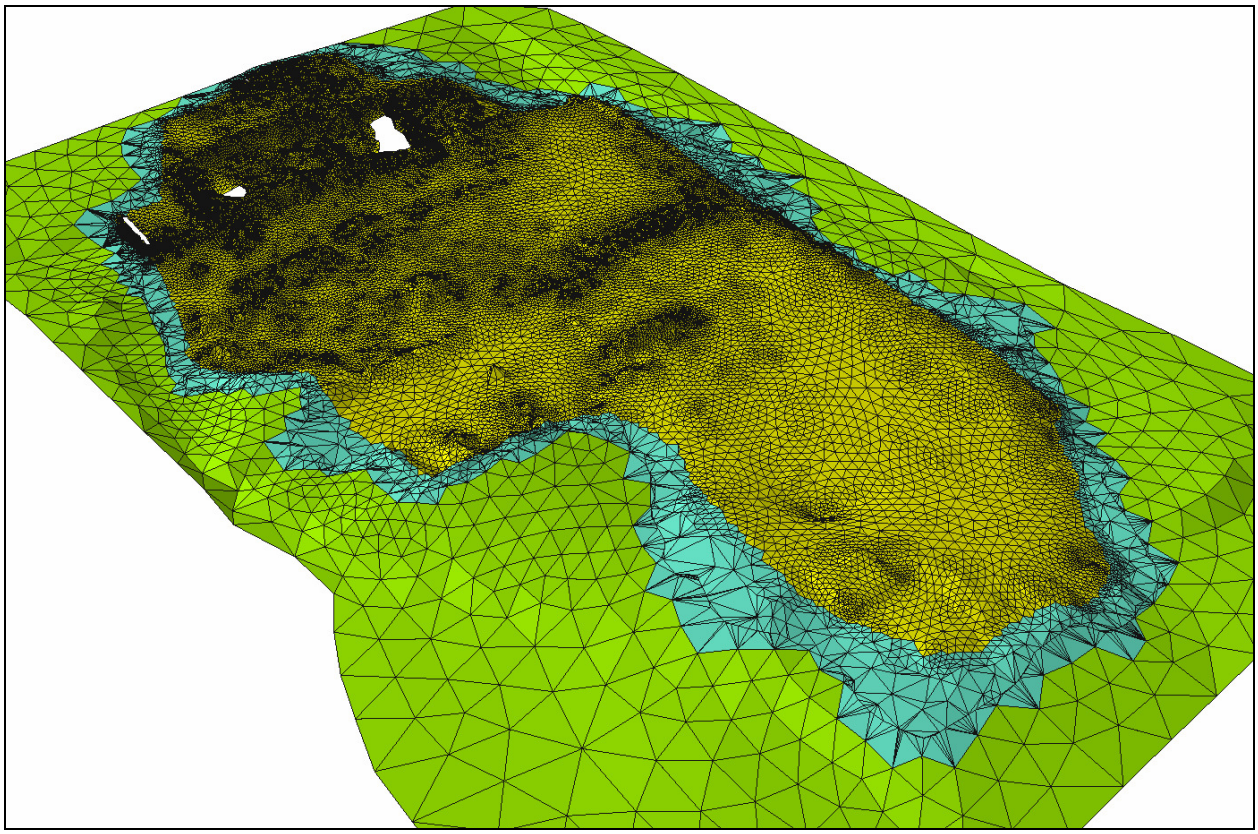

Figure 5.35: Transition surface between mesh-1 and mesh-2 obtained by setting $\varepsilon_{\max} = 1$ cm.

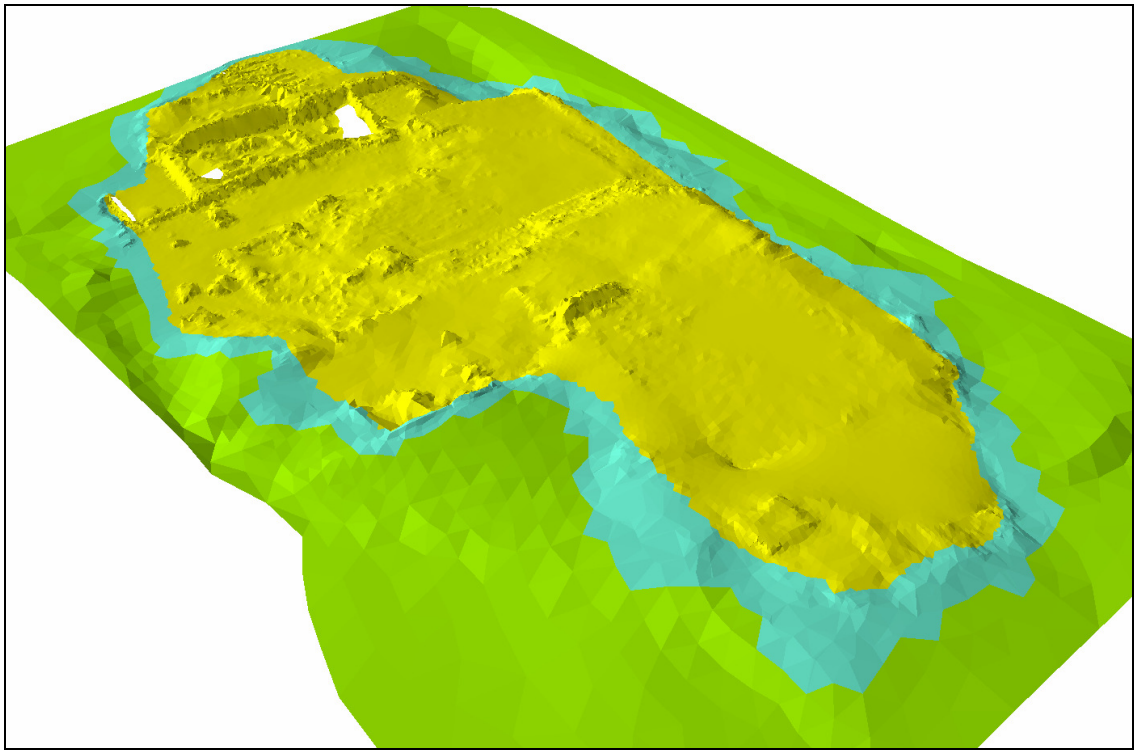

Figure 5.36: Transition surface as in Figure 5.35, without highlighted triangle edges.

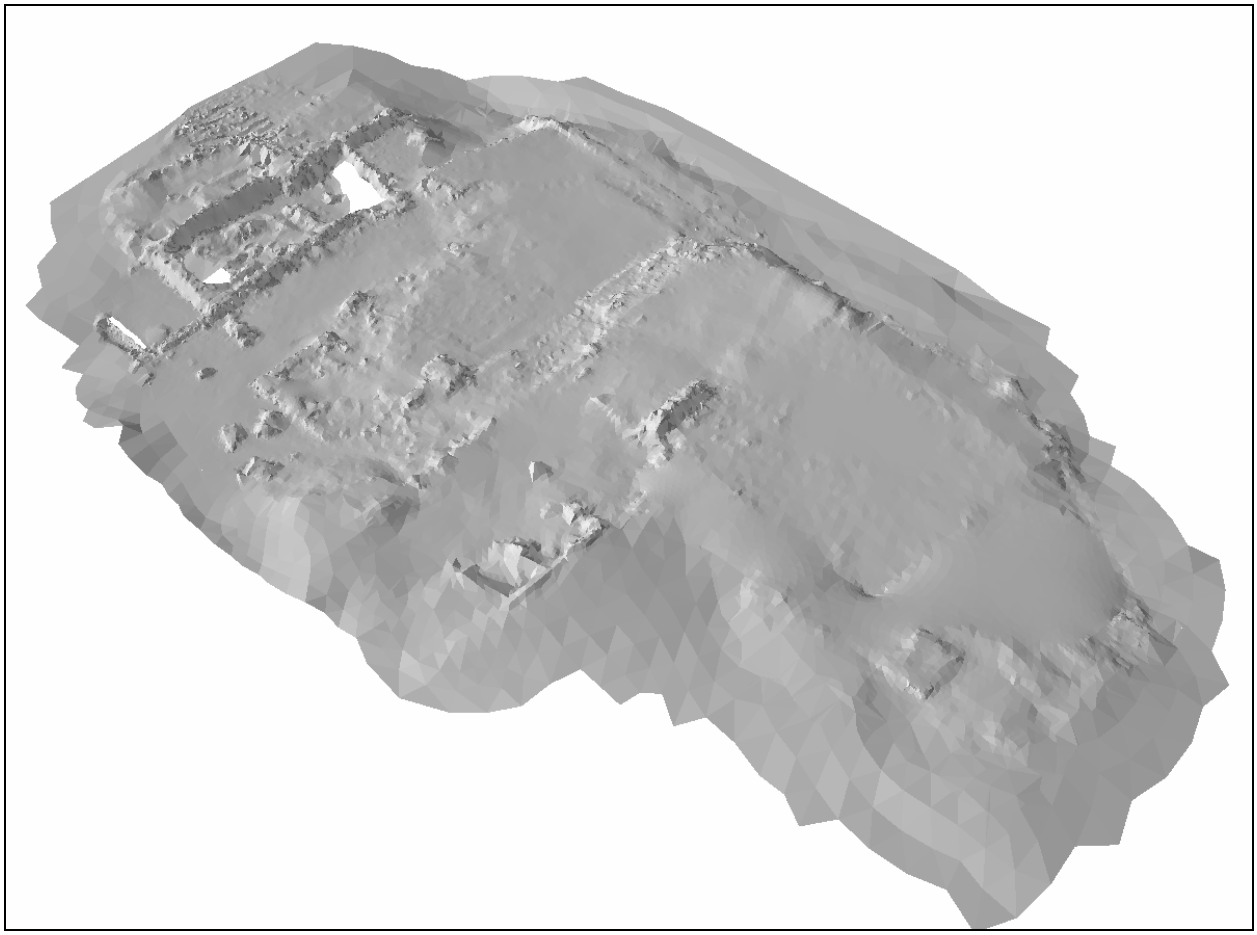

Figure 5.37: The final integrated model, called mesh-12.

5.2.2 Further integration with mesh-3

Upon completion of the first integration between the DTM and the excavated area of the temple, the third model, i.e. the high resolution apse, has been embedded. An overlay of mesh-12 and mesh-3 is shown in Figure 5.38.

An analogous geometric approach by iterative simplification has been performed on the imported low resolution triangles of mesh-12. Unlike the previous case, no clear indication has emerged from the iterative approach, as shown in Figure 5.39: no threshold is recognisable in the graph of the simplification curve; moreover simplification starts even for small values of positive $\varepsilon_{\max} \approx 0$. As explained in § 4.9.3, this is due to nearly coplanar set of triangles: if a vertex is removed, the new triangles reduce the complexity of the mesh but its geometry does not change since a previously planar area is simply filled with less triangles. This is indeed the case of the dense planar triangulation surrounding the inner walls of the apse, a detail view is given in Figure 5.40.

A possible solution to define the $\varepsilon_{\max}$ resides in a heuristic approach: the mesh simplification algorithm is applied iteratively to the transition surface itself (and not to the imported low resolution triangles). From the graph of the simplification curve, an approximate upper limit for $\varepsilon_{\max}$ can be identified as $\bar{\varepsilon}_{\max} = 2$ cm, i.e. as the point where the simplification curve tends to stabilise. It is left to the user to select a proper $0 \le \varepsilon_{\max} \le \bar{\varepsilon}_{\max}$, depending on the point density of the surrounding low resolution data. In this example, the value chosen is $\varepsilon_{\max} = 0.5$ cm. Results are presented in the following Figure 5.41 to Figure 5.45.

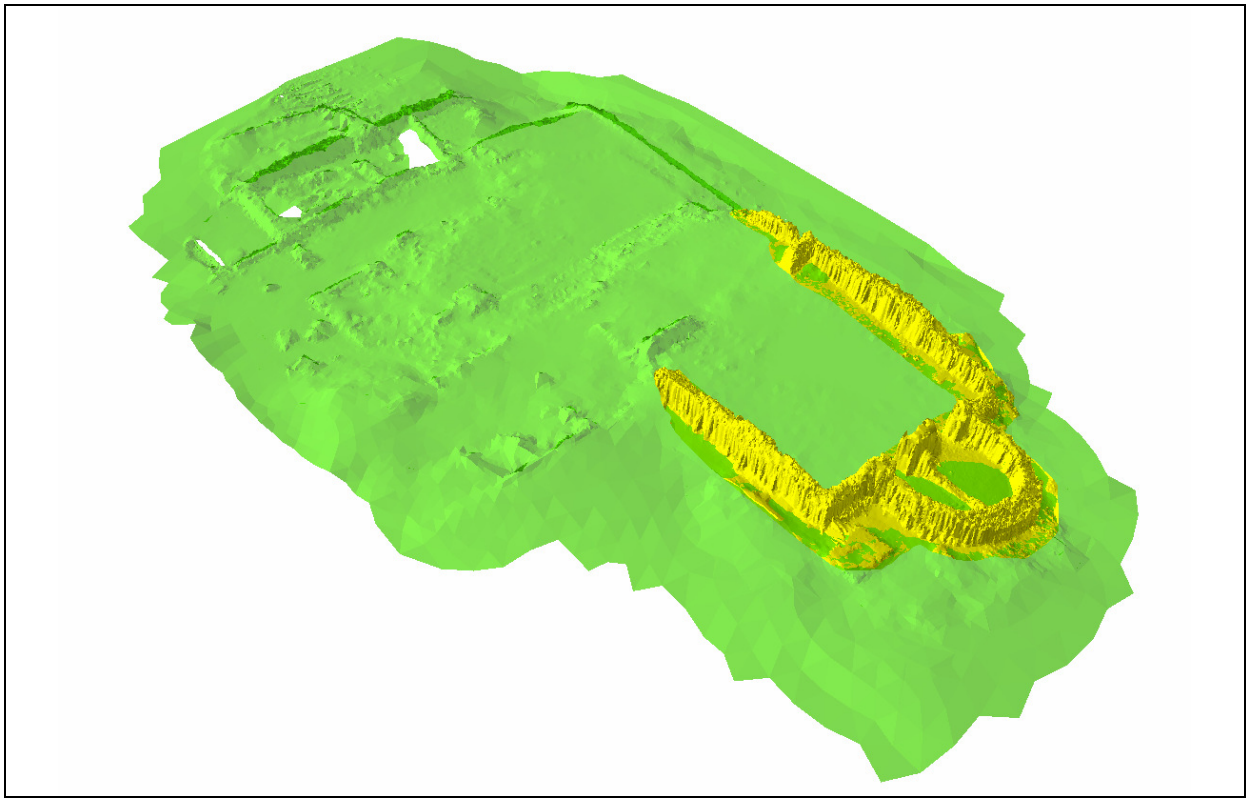

Figure 5.38: Three-dimensional view of mesh-12, in light transparent green, with the overlapping mesh-3, in yellow.

Table 5.5: Global parameters for mesh-1 and mesh-2 models. The values for l_{3D} and Δz are calculated only on imported data.

Transition surface mesh-12 - mesh-3					
Triangle count		*3D edge length l_{3D}*	*[m]*	*Height diff. Δz*	*[cm]*
LR triangles (total)	54404	LR l_{3D} range	0.018 ÷ 0.989	Δz range	-31.6 ÷ 24.0
LR triangles (import)	3156	LR avg(l_{3D})	0.211	avg(Δz)	0.2
HR triangles (total)	173279	HR l_{3D} range	0.004 ÷ 0.893	avg($\lvert\Delta z\rvert$)	1.9
HR triangles (import)	17477	HR avg(l_{3D})	0.073	RMSE	5.3

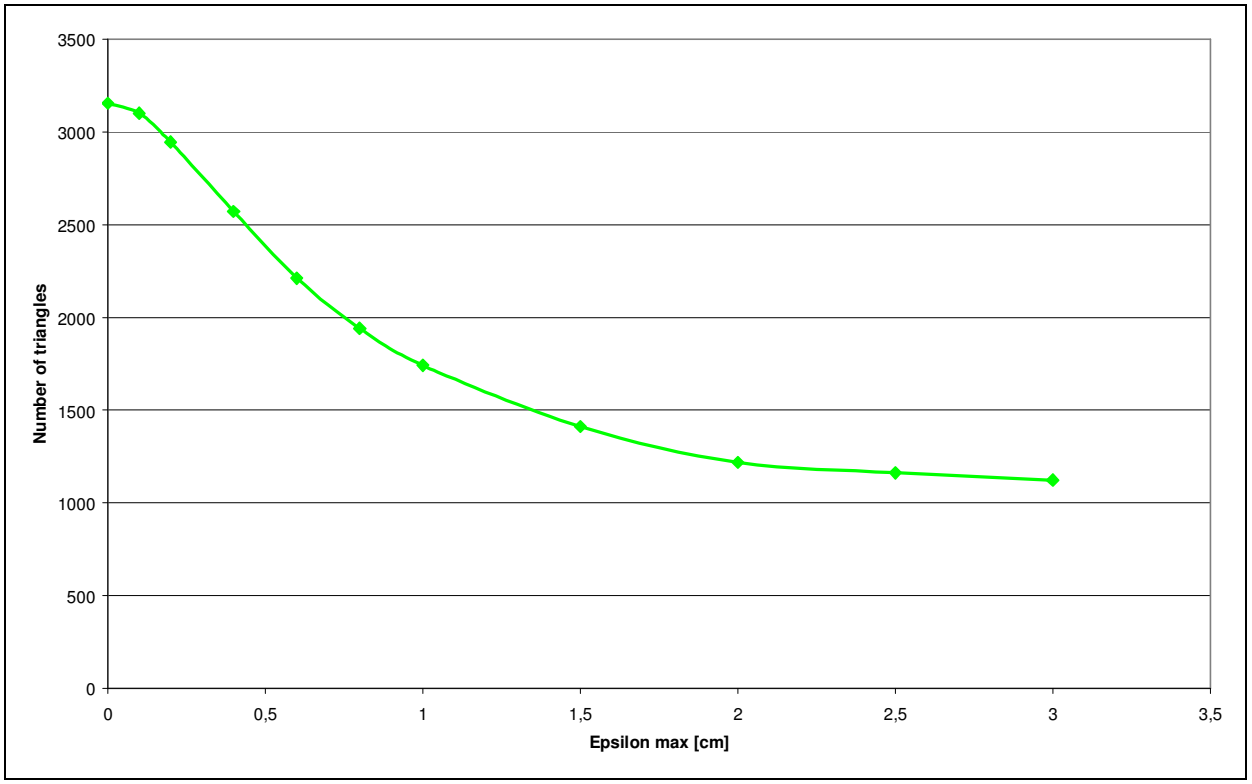

Figure 5.39: Simplification curve of the imported low resolution triangles. No threshold value is recognisable.

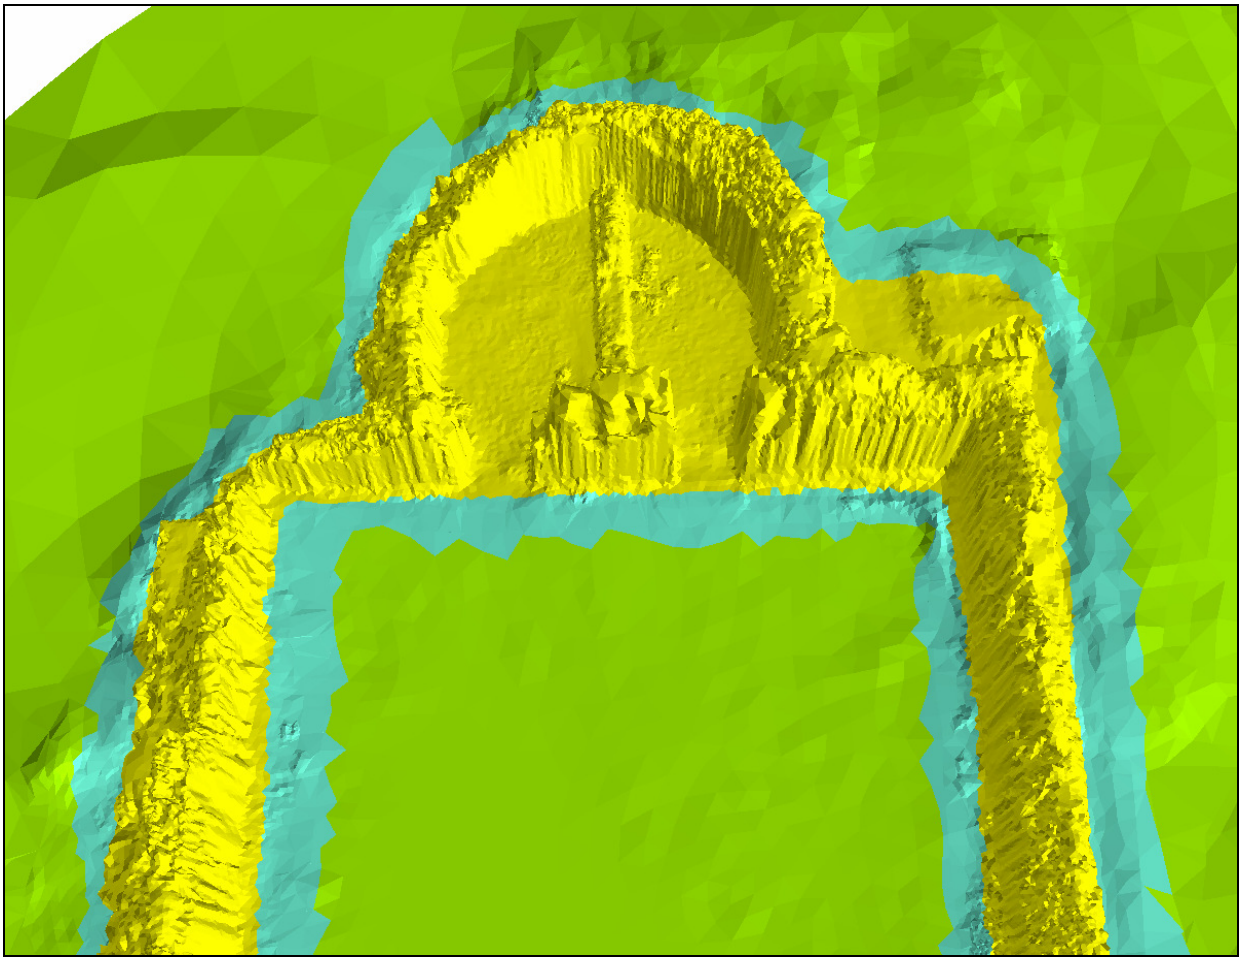

Figure 5.40: Detail view of the nearly planar area near the inner walls of the apse.

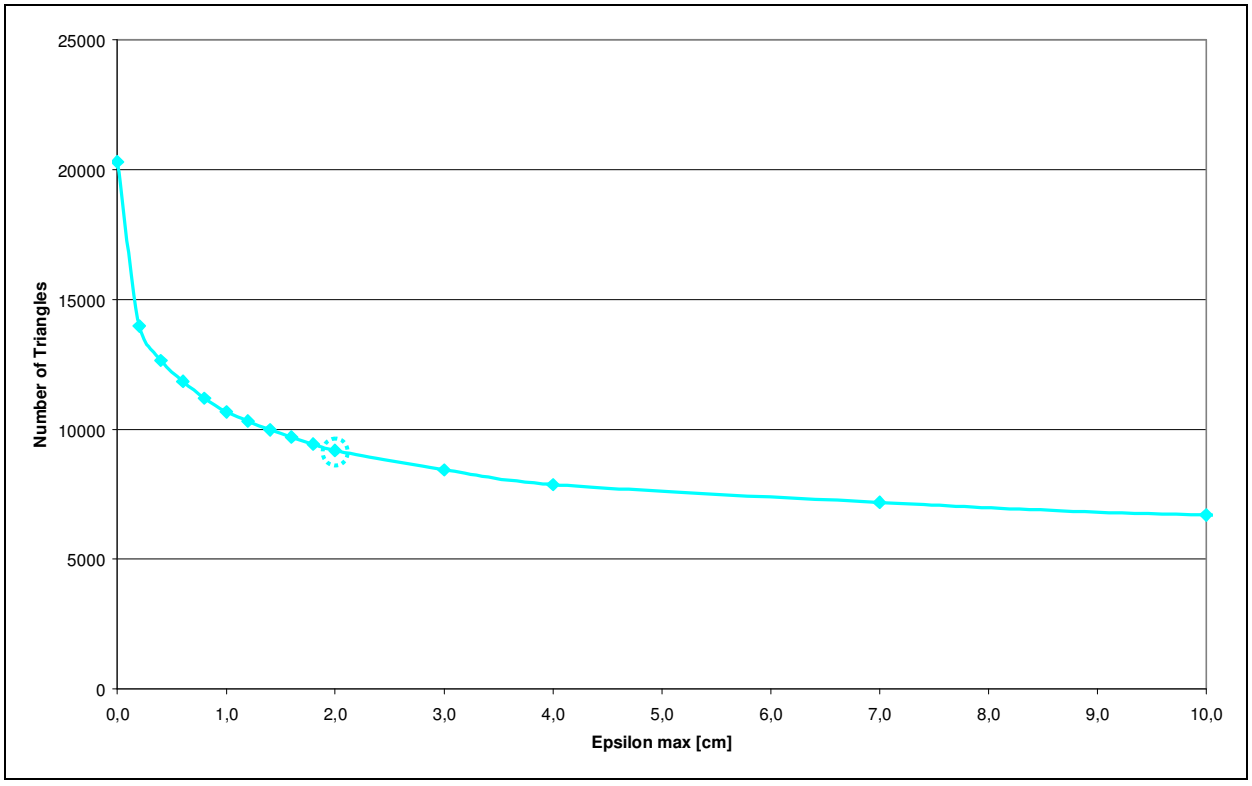

Figure 5.41: Simplification curve for the transition surface between mesh-12 and mesh-3. The interval for the $\varepsilon_{\max}$ is $0 \le \varepsilon_{\max} \le \bar{\varepsilon}_{\max}$. Value $\bar{\varepsilon}_{\max} = 2$ cm is highlighted on the curve.

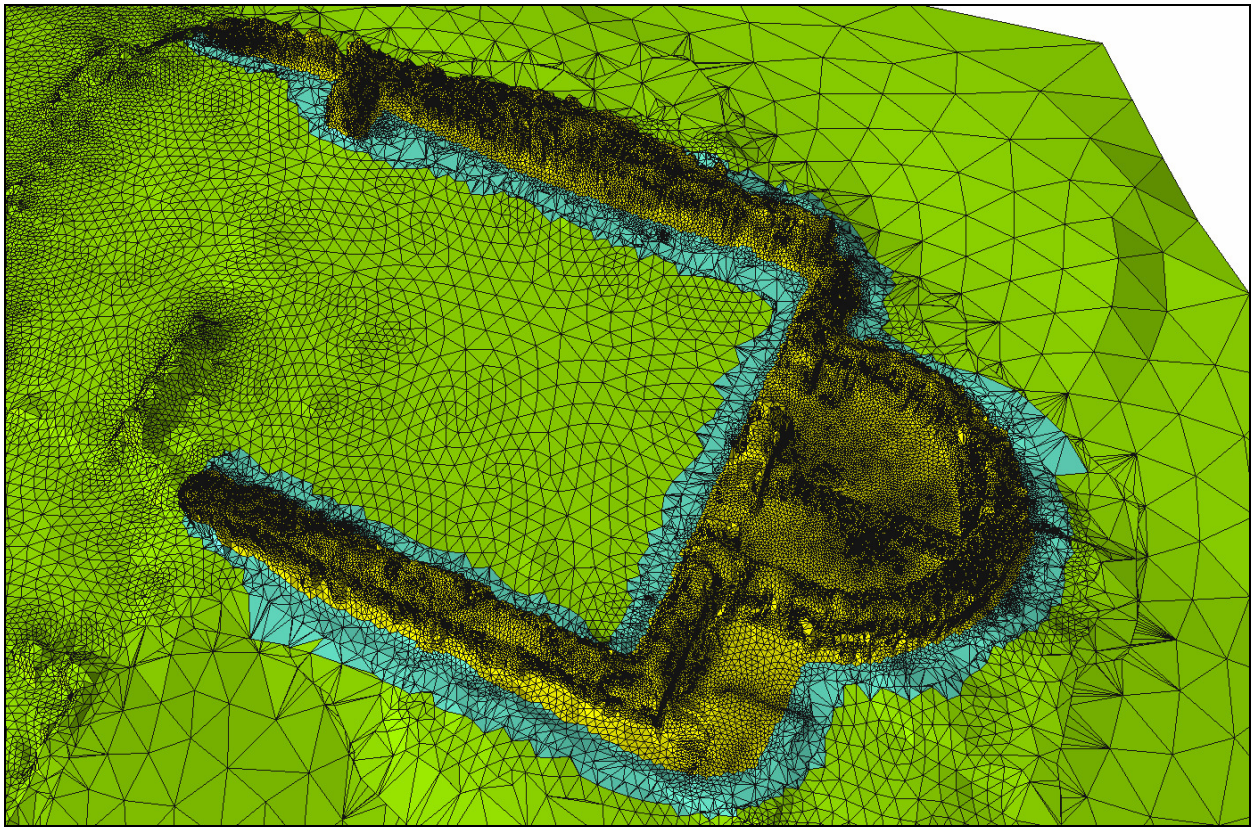

Figure 5.42: Transition surface between mesh-12 and mesh-3 obtained by setting $\varepsilon_{\max} = 0.5$ cm.

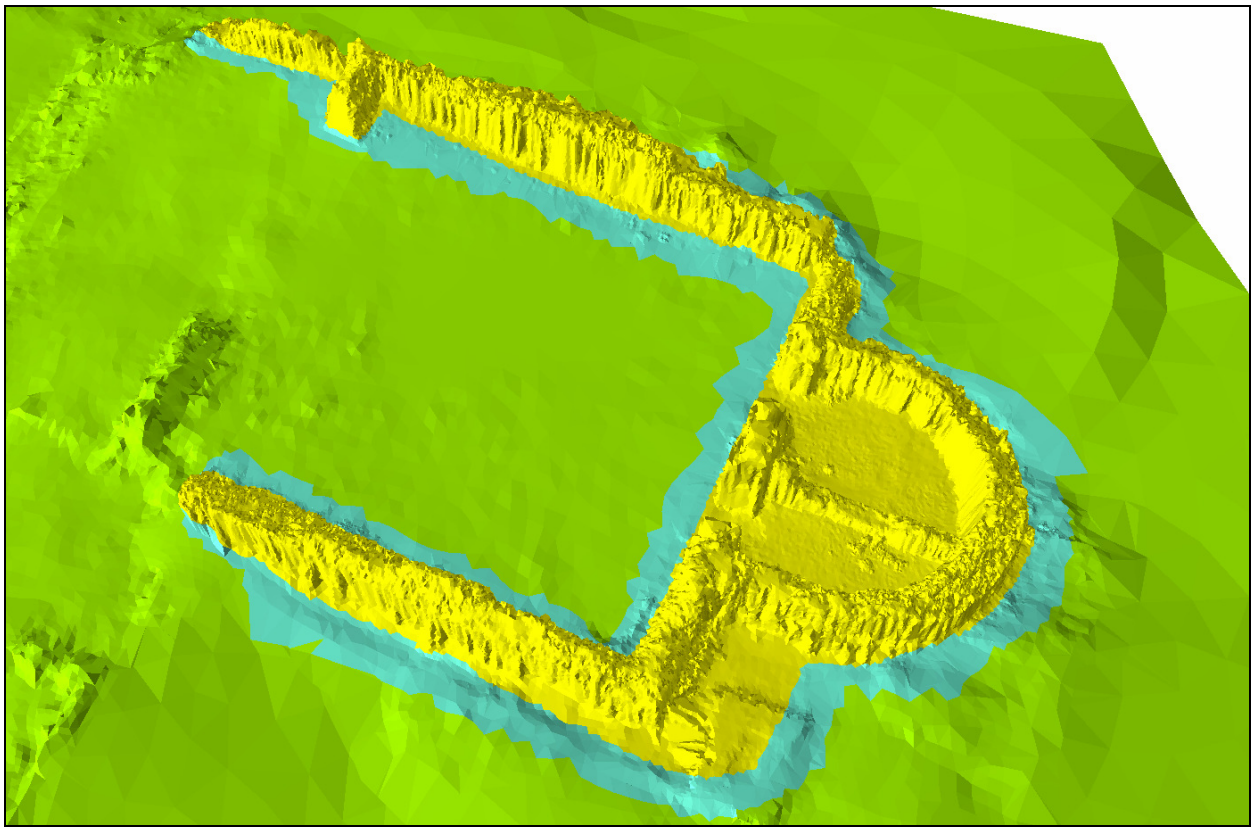

Figure 5.43: Transition surface as in Figure 5.42, without highlighted triangle edges.

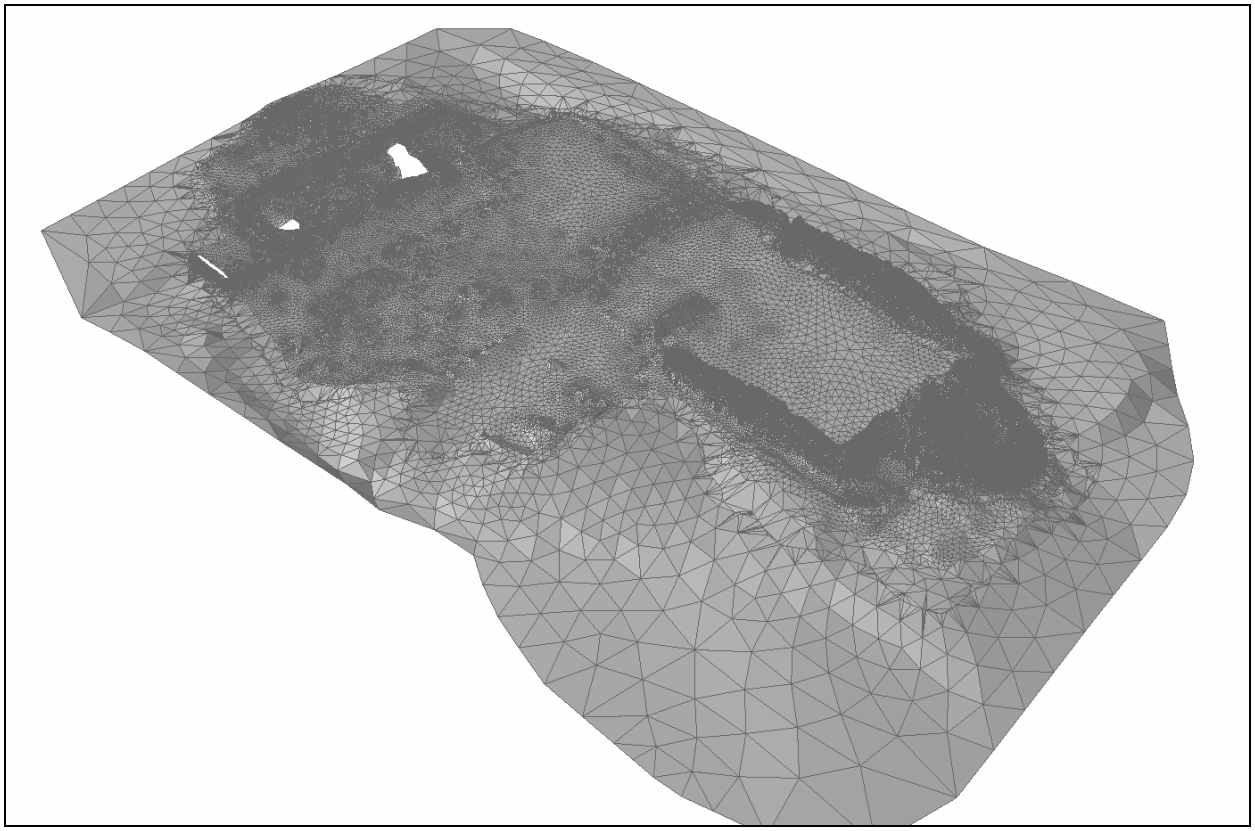

Figure 5.44: The final integrated model resulting from mesh-1, mesh-2 and mesh-3 with highlighted edges.

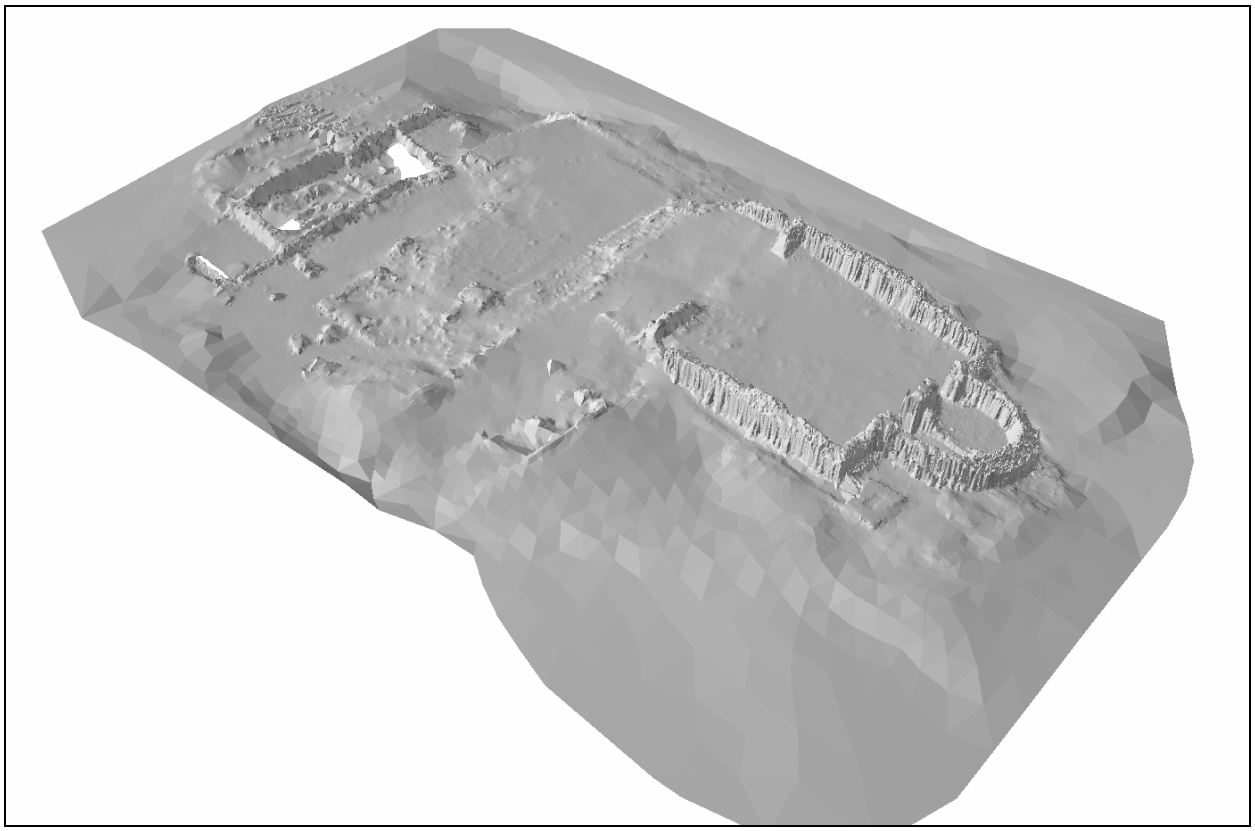

Figure 5.45: The final integrated model resulting from mesh-1, mesh-2 and mesh-3.

5.3 Example 3: Application to a small, detailed bas-relief

This last example has been chosen to test how details are preserved in the transition surface. The datasets used do not actually belong to the category of models previously seen: there is no DTM, although a dominant plane can be identified and used as xy-plane after a coordinate transformation. On the other hand, this example simulates a common problem in the framework of cultural heritage surveying: different models of a highly detailed object may exist at different resolutions; a particularly elaborated decoration may have to be embedded into the surrounding lower resolution model.
The datasets used in this example refer to the Scaliger tombs, see Figure 5.46, left, a group of five Gothic funerary monuments celebrating the Scaliger family, who ruled the city of Verona, Italy, from the XIII to the late XIV century. The model in Figure 5.46, right, shows one of the family emblems. It was acquired using a Konica Minolta Vivid 910 laser scanner.

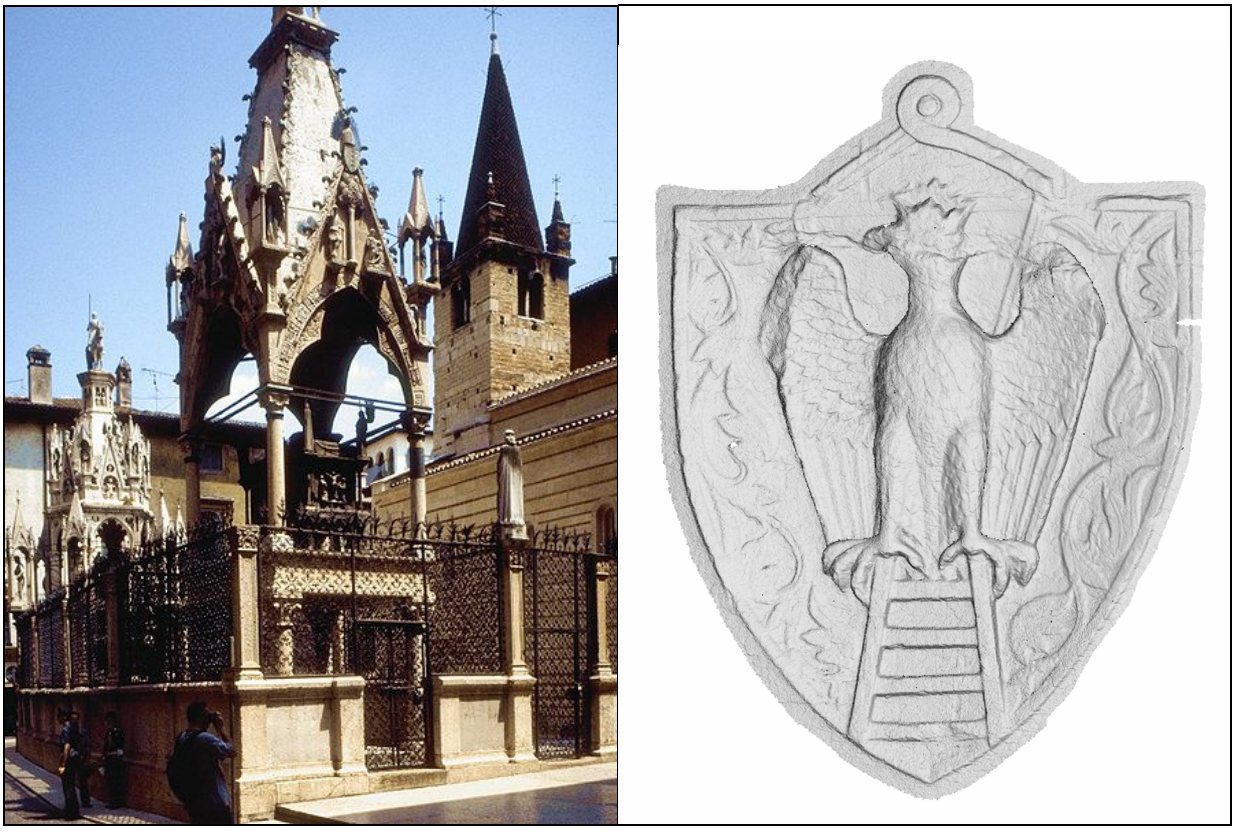

Figure 5.46: [Left] View of the group of the Scaliger Tombs in Verona, Italy. [Right] The laser scanner model of one of the Scaliger family emblems. It is located on the side of one of the sarcophagi. The stair motif is recognisable, in reference to the Italian meaning of the family name: "della Scala". Left image source: [Scaliger Tombs].

From the original model, two models were taken (lr_mesh and hr_mesh): the low resolution one was obtained using the point decimation built-in function of the Konica Minolta Polygon Editing Tool software. For the high resolution model, the eagle was cut out, the intention being to embed it into the low resolution model. Both datasets are shown in Figure 5.47 and Figure 5.48. Thr preliminary steps, from data import to the

computation of the unsimplified transition surface, are shown in the following pictures and briefly described in the accompanying captures.
The initial method chosen to set the $\varepsilon_{\max}$ by means of the geometric approach on the low resolution triangles delivered no clear indication, since no “step” can be seen in the green simplification curve of Figure 5.49.
A simplification curve has been obtained for the transition surface itself and is shown in Figure 5.50 (azure curve). The value $\bar{\varepsilon}_{\max} = 10$ mm was chosen as the approximate upper bound of the interval in which to set the $\varepsilon_{\max}$ value for the transition surface. For this example, it was set $\varepsilon_{\max} = 6$ mm.
Results are presented in the following Figure 5.51 to Figure 5.53.

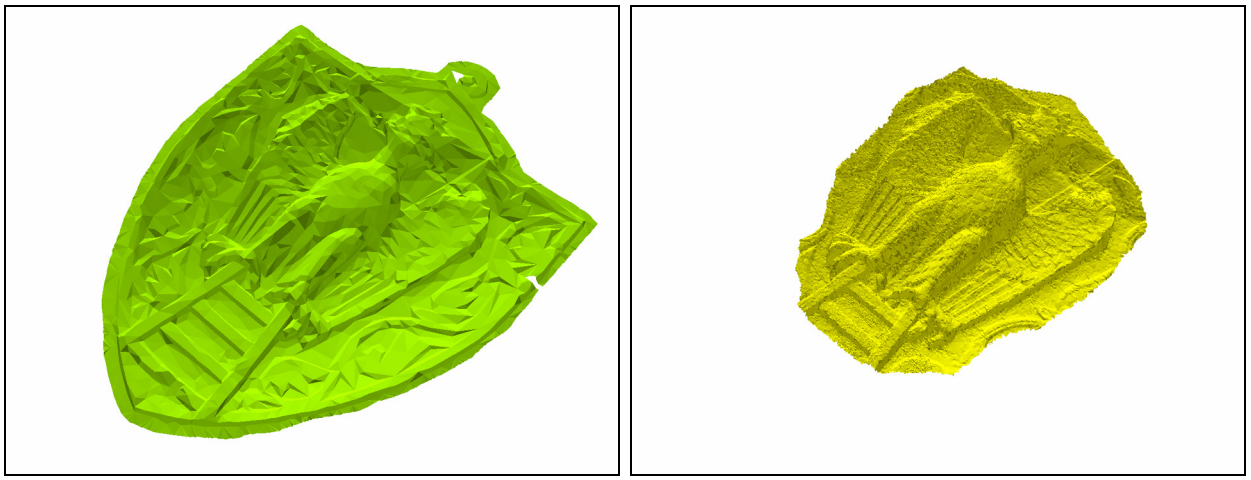

Figure 5.47: [Left] The low resolution dataset of the whole emblem. [Right] The high resolution detail of the eagle.

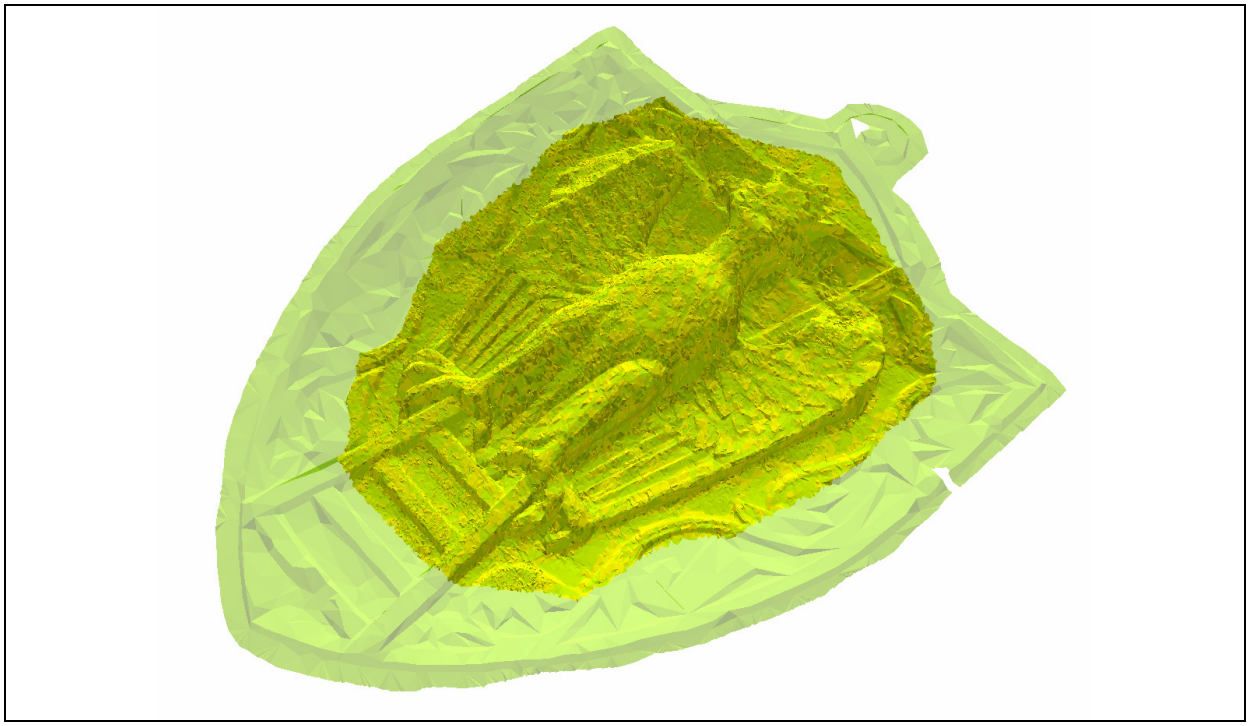

Figure 5.48: Three-dimensional view of the low resolution model, in light transparent green, with the overlapping high resolution model, in yellow.

Table 5.6: Global parameters for lr_mesh and hr_mesh models. The values for l_{3D} and Δz are calculated only on imported data.

Transition surface lr_mesh – hr_mesh					
Triangle count		*3D edge length* l_{3D}	*[mm]*	*Height diff.* Δz	*[mm]*
LR triangles (total)	4676	LR l_{3D} range	4.9 ÷ 69.1	Δz range	-14.5 ÷ 14.3
LR triangles (import)	692	LR avg(l_{3D})	19.9	avg(Δz)	-0.1
HR triangles (total)	71734	HR l_{3D} range	0.2 ÷ 15.2	avg($\lvert\Delta z\rvert$)	0.4
HR triangles (import)	16834	HR avg(l_{3D})	3.4	RMSE	0.7

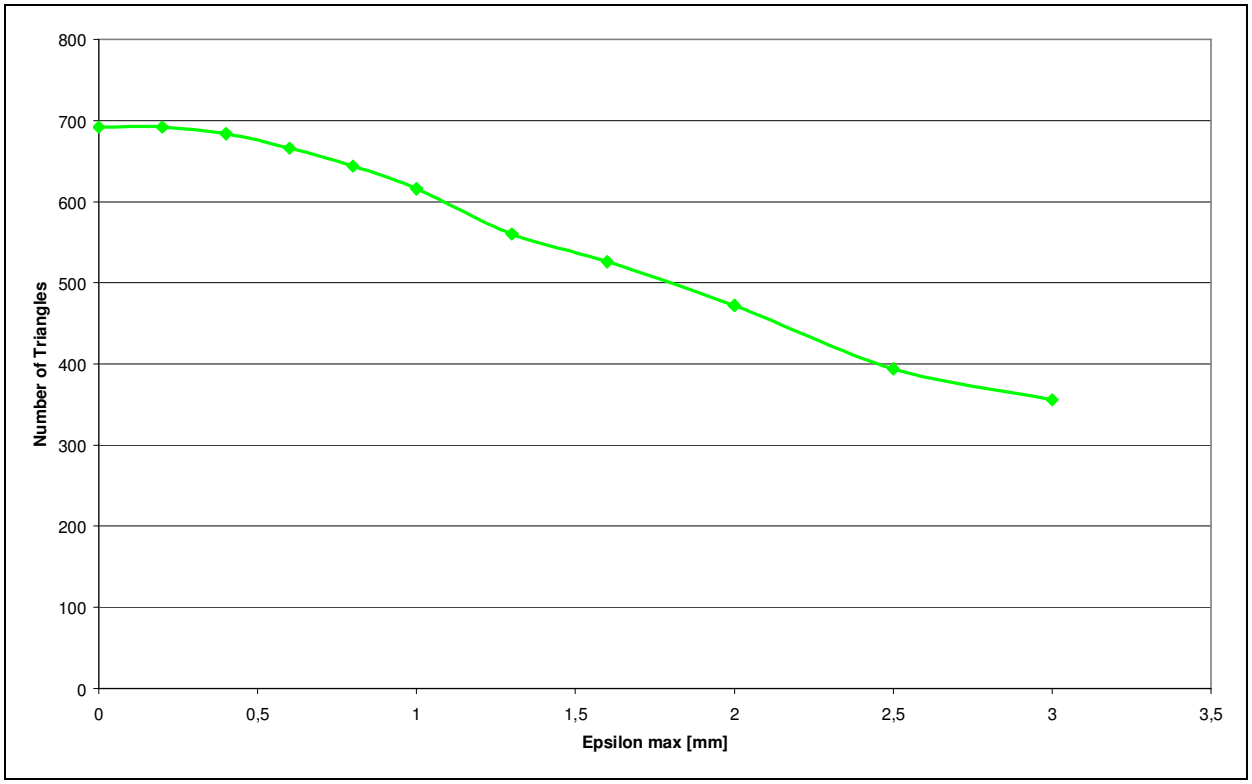

Figure 5.49: Simplification curve of the imported low resolution triangles. No threshold value is recognisable.

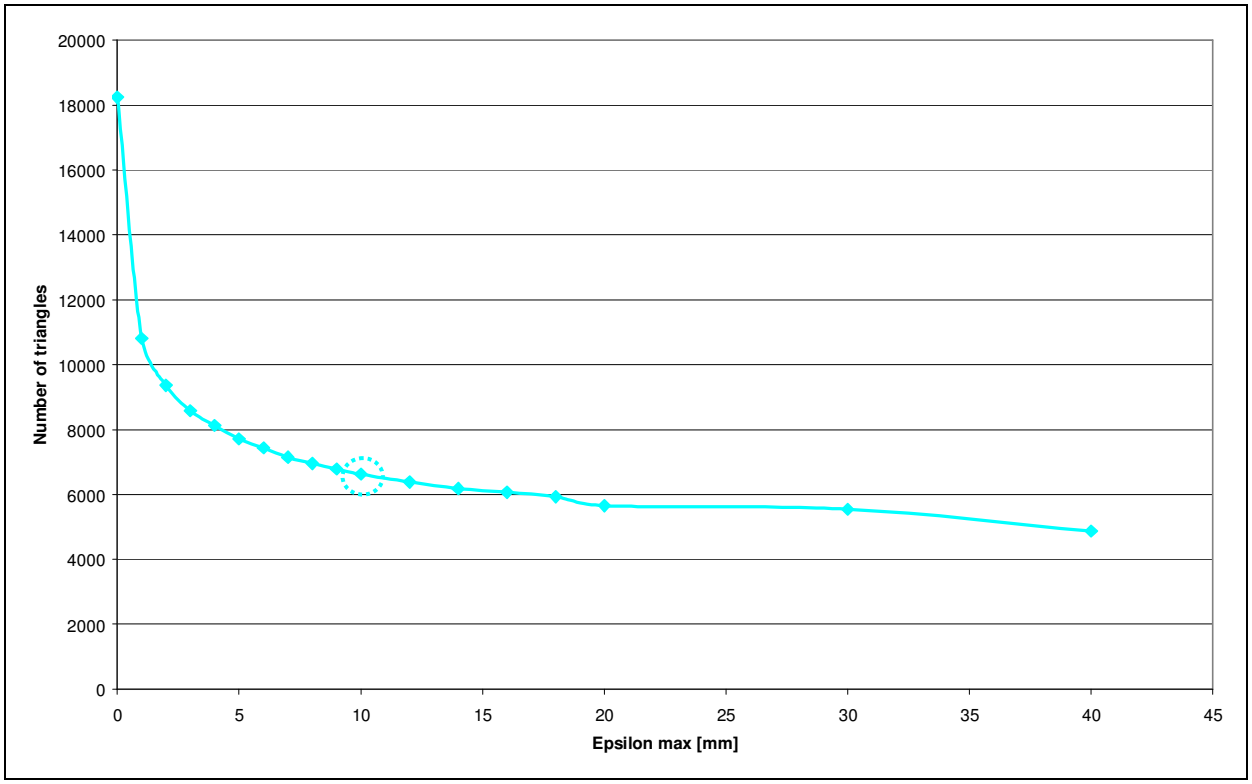

Figure 5.50: Simplification curve for the transition surface. The interval for the $\varepsilon_{\max}$ is $0 \le \varepsilon_{\max} \le \bar{\varepsilon}_{\max}$. Value $\bar{\varepsilon}_{\max} = 10$ mm is highlighted on the curve.

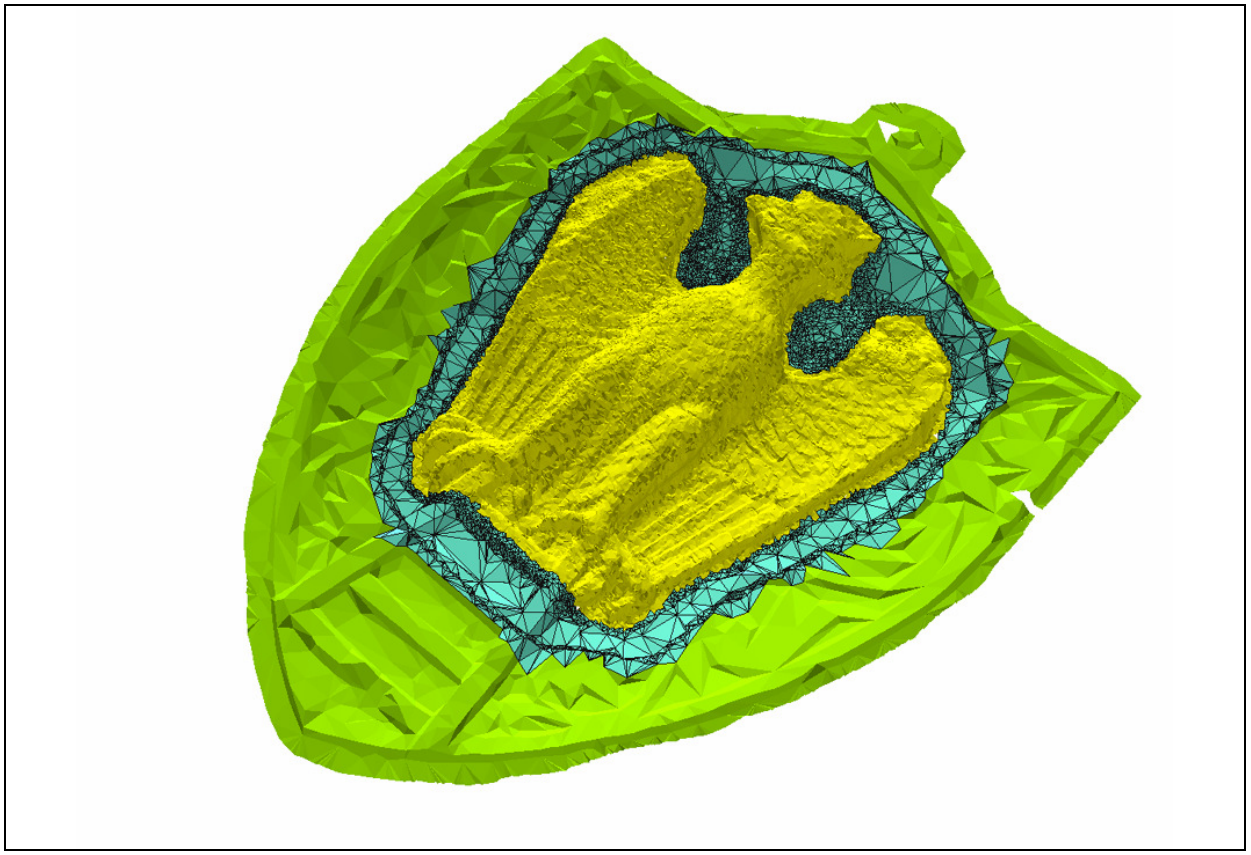

Figure 5.51: Transition surface obtained by setting $\varepsilon_{\max} = 6$ mm.

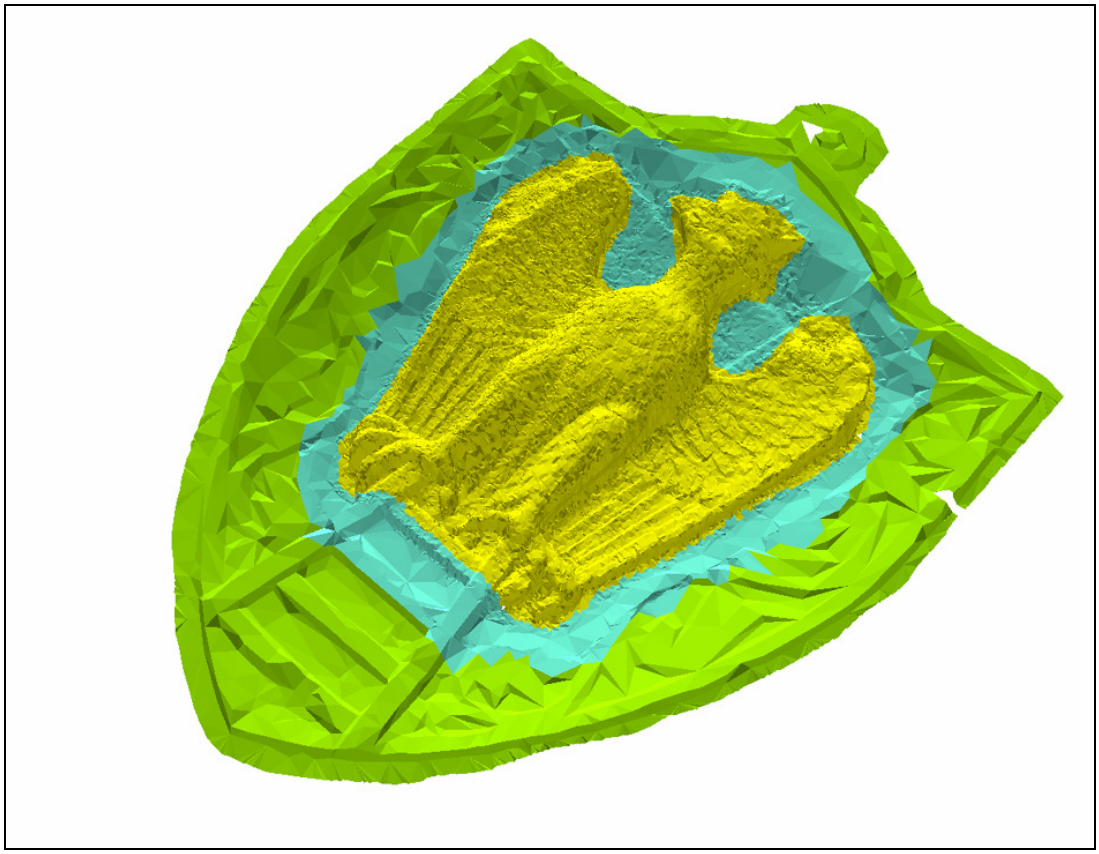

Figure 5.52: The same model as in Figure 5.51, without highlighted triangle edges.

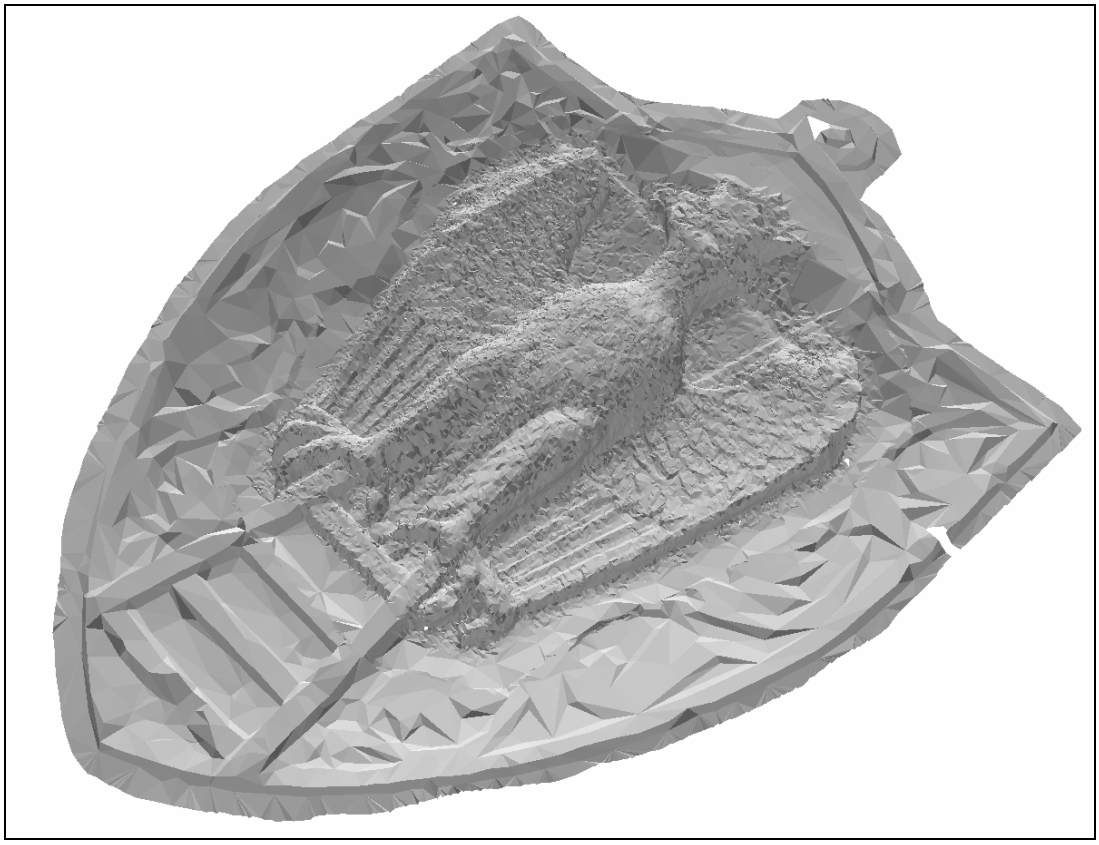

Figure 5.53: The final, integrated model.

Chapter 6

FURTHER REASONING

In this chapter experimental results obtained from the tests are discussed, as well as open issues and further improvements of the developed methodology. Possible future research topics are briefly presented.

6.1 Comments on the experimental results

In the previous chapter the developed methodology has been applied to several datasets with different characteristics. Bearing in mind the initial assumptions (§ 4.1) and the desired properties of the transition surface (§ 4.2), results show that integration can be successfully performed between two heterogeneous datasets acquired at different resolutions and with different surveying techniques, although laser-scanner-derived datasets play a major role in the presented examples.

It has been shown that a transition surface connecting a low resolution DTM to a high resolution model can be created, and the initial continuity requirements in terms of topology, geometry and point density can be fulfilled under the working hypotheses.

In the overlapping zone (the so-called zone C), the transition surface uses information from both datasets and allows for a gradual transition between them. Height and tangential continuity is guaranteed at both delimiting borders: points near the outer border lie closer to the DTM, points near the inner border lie closer to the high resolution object. Inside zone C, the height profile lies in-between. Moreover, a progressive mesh simplification can be achieved: point density decreases gradually from the inner border to the outer border, i.e. triangles grow accordingly in size. Planimetrically, no new points are created, only existing points are removed.

The methodology has also been positively tested with “unfavourably” lying datasets, i.e. in case the copresence of high and low resolution data inside the whole zone C cannot be guaranteed.

The obtained transition surface permits therefore a local update of a DTM through a geometrically and topologically correct insertion of a high resolution model, provided that some extra information, the so-called “collar”, is present in the high resolution object and that it can be used.

The possibility to leave the high resolution unchanged in zone A and to insert it into the DTM yields some analogies with the idea of the 3D enclaves presented by OTEPKA, BRIESE & NOTHEGGER (§ 2.4.3). Provided that information is given for the junction (the inner border) and for the transition surface (the collar), a 2.5D DTM can be locally enhanced. Since no strict requirements are given for data *inside* the inner border, it may well be a topologically different model. Whether another 2.5D DTM – e.g. a nested DTM in CityGML –, or a 2.8 model – à la GRÖGER & PLÜMER (§ 2.4.2) –, or a completely 3D object, this is not explicitly against the initial assumptions; it may, however, depend on

other potential limitations like, for example, the capability of the database to handle heterogeneous spatial features.

Another advantage of the methodology is that, unlike LATHAM & BURNS' approach (§ 3.5.2), the distant low resolution DTM remains unchanged: there are no height changes for any data outside the outer border of the transition zone. Limiting the extents of the modifications to a fixed domain, i.e. inside the overlapping zone C, recalls LAURINI's "elastic band" idea and prevents that local changes influence distant areas of the DTM. If necessary, low resolution triangles are cut by the outer border, and the vertices outside zone C are connected to the outer border in order to guarantee topological continuity, but their height does not vary. Cutting the triangles is actually an optional step, even if this operation has been carried out thoroughly in this work. If topological continuity is not required and the transition surface is not going to be integrated into the final model, C0 and C1 continuity at the outer border is achieved anyway by the height interpolation model.

Also in WARRINER & MANDLBURGER (§ 3.5.2) the integration approach permits a smooth transition of the height profile between two datasets by a chosen weight function. However, in this thesis the overlapping zone is not defined by a buffer of constant width; it can be instead of any shape and can potentially have multiple inner borders in it, as described in § 4.6.2. The experimental results confirm that the limitations of the simple Euclidean distance functions on irregular planar shapes have been overcome.

Regarding the simplification of the transition surface, different strategies have been adopted in the examples, due to the multiple interpretation which is possible for the simplification parameter $\varepsilon_{\max}$ (§ 4.9) and according to the availability of information about the surrounding DTM.

The first example (§ 5.1) uses the global point accuracy of the DTM to set the $\varepsilon_{\max}$ value. The simplified transition surface is guaranteed to have triangles that lie no further than $\varepsilon_{\max}$ from the unsimplified surface. It must be noted that the maximum surface displacement can be reached only on the outer border since $\varepsilon_{\max}$ decreases gradually toward the high resolution dataset.

As explained in § 4.7, the $\varepsilon_{\max}$ value defines the thickness of the simplification envelopes (i.e. the offset surfaces), which are computed using the three-dimensional normal vectors of the surface vertices; in short $\varepsilon_{\max} = \|n\|$, with $n:(n_x, n_y, n_z)$ normal vector of the transition surface on the outer border.

In case that only the global height accuracy σ_H of a DTM is known – a rather common case indeed – it may be debatable whether the first method, for which $\varepsilon_{\max} = \sigma_H$, is correct: a vertical accuracy value is set to be the norm of a three-dimensional vector. However, this is an approximate but conservative approach. A vector $\sigma:(0,0,\sigma_H)$ can be assumed such that $\varepsilon_{\max} = \sigma_H = \|\sigma\|$. Since $\varepsilon_{\max} = \|n\|$, with regards to the z-axis it holds $n_z = \varepsilon_{\max} \cos\alpha \le \varepsilon_{\max}$, with α angle between the normal vector and the z-axis. This leads to $n_z = \varepsilon_{\max} = \sigma_H$ only for $\alpha = 0$, i.e. only in the case the normal vector itself is vertical.

Although a more detailed discussion about accuracy will be done in the next paragraph, this is the method which should be preferred for data integration whenever accuracy information is given.

This is however not always the case, due to several reasons which have been mentioned in the previous chapters: the low resolution model may be a derivative product where accuracy information is not available or has been lost for some reason in the data processing pipeline [GOODCHILD, 2008]. For example, working with terrestrial laser scanner data, it is not uncommon that existing point cloud processing software offers tools for mesh creation and editing, as well as various other procedures for point decimation, mesh simplification and smoothing, etc. However, these tools operate sometimes as black boxes where little control is left to the user, once a global initial optimisation parameter has been set.

For such worst case scenarios – actually not so infrequent, and the second and third examples in § 5.2 and § 5.3 belong to this category – the alternative method of § 4.9.2 has been tested, which seeks a characterising parameter of the surrounding low resolution model using only geometric information. The $\varepsilon_{\max}$ value is calculated iteratively as the minimum value beyond which the low resolution dataset starts to be further simplified, in the assumption that the low resolution model is itself a product of a previous unknown data reduction process. If this value can be calculated, then it is used to simplify the transition surface (§ 5.2.1). For the case that this method cannot be successfully applied and an $\varepsilon_{\max}$ value cannot be determined, two critical cases are also presented (§ 5.2.2 and § 5.3), where a heuristic approach is instead adopted.

As mentioned in § 4.9.3, a limitation of the current implementation of the simplification algorithm resides in the lack of control over the vertex removal. Basically, all vertices are removed till the transition surface cannot be further simplified. In order to stop the simplification process at a desired intermediate level, a further check for the vertex removal operation should be implemented: besides the current geometrical constraints, another condition should be added which stops the removal operation if this causes a drop in point density beyond a certain threshold value. Moreover, the threshold value should vary itself according to the distance from the inner and outer border where the respective point density values must be respected.

It must be noted that so far the term “point density” has been used in this work quite freely, without giving a specific definition. The concept of point density is generally straightforward when referred to the xy-plane, but this is not the case when considering the point density on the surface itself. Which surface should be used? The mesh surface is itself an approximation of the (unknown) real surface. One alternative could consist in using the mutual distances among connected points, but by now this is only a idea for a future improvement.

Regardless of the method to set the $\varepsilon_{\max}$ value, results show that the simplification algorithm acts indeed as a variable low pass filter: details are progressively smoothed moving from the inner border to the outer border, although relevant features are preserved also near the low resolution dataset, as shown especially in the last example (§ 5.3).

The methodology has been tested also outside the specific framework of DTM integration, since the used datasets bear a clear resemblance to the “standard” ones seen before, and their integration problems can be assimilated to those of a DTM.

Although the obtained results show that the methodology can deliver acceptable results, several initial assumptions have been done. Moreover some decisions must still be met by the user, such that a complete automatic integration process is not possible, yet. In the following paragraph possible enhancements of the methodology and open issues will be discussed.

6.2 Open issues and further improvements

In order to reduce the initial working hypothesis and extend the field of applicability of this work, some existing limitations need to be overcome.
Concerning topology, some operational steps currently hinder the adoption of a full three-dimensional approach: the height difference values between the two input meshes are, for example, still calculated on the z-axis, however this restricts the datasets in the overlapping zone C to be 2.5D models only. This approach is in fact standard and well-accepted in the framework of DTM, but yields some limitations when it has to be applied to topologically more complex surfaces, like an overhang in a DTM, or to complex surfaces which generally originate in the context of cultural heritage and where a dominant plane is not present – for example, how to measure the distance between the low resolution model of a column and the to-be-embedded high resolution capital?
In § 4.8 some metrics have been presented from the framework of mesh simplification algorithms which could represent the initial step for an improvement of this aspect.
A second working hypothesis is that the computation of the border distance parameter *x* occurs on the basis of buffers created on the xy-plane. This is a reasonable approximation as long as the annular domain of the overlapping zone on the xy-plane does not differ too much from the overlying polyhedral surface(s). Dealing with irregular surfaces means however that the absolute distance obtained on the xy-plane underestimates the actual distance between the borders on the 2.5D (or 3D) surface. An improvement should consider calculating geodesic shortest paths on the actual polyhedral surface. Its computation is a common operation in many computer graphics applications, since geodesic paths establish a surface distance metric, see Figure 6.1.
In SURAZHSKY *et al.* (2005), for example, several practical methods to compute both exact and approximate distances on a triangle mesh are presented. The geodesic paths can cut across faces in the mesh and are, therefore, not found by the traditional graph-based DIJKSTRA algorithm for shortest paths.
However, in the case of the transition surface, applying one of these improved distance functions may lead to a circular problem, because geodesic distances should be calculated on a surface… that has still to be calculated, since it first requires the distance parameter for the height interpolation model!
A possible solution to this tricky egg-and-chicken problem could reside in a iterative procedure which should, hopefully, converge to the desired transition surface.

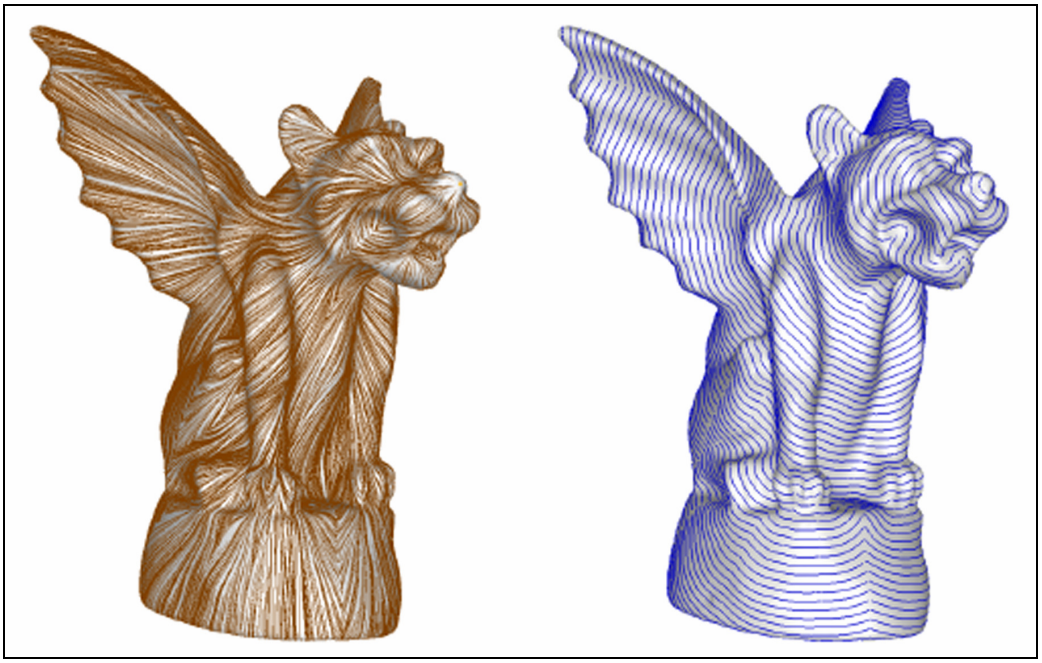

Figure 6.1: [Left] Geodesic paths from a source vertex. [Right] Isolines of the geodesic distance function. Image source: [SURAZHSKY *et al.*, 2005].

One major planned improvement of the methodology resides in a general overhaul of the current deterministic approach with a stochastic one that accounts for spatial data uncertainty. According to one initial assumption, the input datasets are required to be aligned and all resulting errors from this operation are considered negligible. In reality, this simplification does not always hold true and systematic errors should be indeed considered and included into the model.

More in general, uncertainty is part of a wider group of quality parameters which includes for example positional and attribute accuracy, logical consistency, and completeness. They are commonly referred as Spatial Data Quality (SDQ).

At least two further problems need to be addressed: a) how to model (i.e. to store) variation of spatial data quality in a database and b) how to model (i.e. to calculate) variation of spatial data quality for different GIS primitives such as points, lines and polygons.

For the former problem, several approaches exist: SDQ information can be stored separately from the spatial database through links. This leads to a so-called *external representation*, where however spatial data quality is harder to update, and often contains only aggregated quality values for the entire spatial database, ignoring *de facto* spatial variation.

In case of the so-called *integrated representation*, SDQ information is integrated into the spatial database, thus letting the user to represent and query more easily spatially varying quality on a per feature basis. It can be easily understood how a set of quality parameters can be associated to every point feature.

Nevertheless, both approaches suffer from an important drawback: variation in quality *within* an object, sometimes defined as sub-feature variation, cannot be represented. In other words, how to represent spatial data quality along a line or inside a polygon?

For completeness it must be added that a third, *hybrid approach*, has been recently proposed by SADIQ & DUCKHAM (2009) in which a possible strategy is presented to overcome the aforementioned drawback. Nonetheless, the overall problem is still subject of research.

The latter open issue regards how to calculate variation of spatial data quality for different GIS primitives. The discussion will be done on the basis of some examples and, in the context of this work, only accuracy related matters will be mentioned. As stated in § 2.5.1, absolute accuracy should be distinguished from internal accuracy (i.e. precision).

A statement about the quality of the final, simplified transition surface should be therefore added to the developed methodology. In other words, how can its internal and external accuracy be addressed and quantified? And to which original dataset should the latter be compared?

Several factors must be considered; it may be convenient to recall some of the operations performed to obtain the transition surface. First, it is assumed that the input models consist of triangulated meshes: the surface is described by triangular planar features and, as stated before, this is itself an approximation. When high resolution points are projected onto the low resolution DTM (and vice versa), height values are interpolated on the low resolution triangle surface.

Provided that standard deviation values are given for the vertices of the mesh, what happens to the points inside the triangle face? Or, in the two-dimensional case of a line defined by two known endpoints, how varies uncertainty for a point inside the line?

It can be easily shown by error propagation that if the coordinates of the endpoints are spatially uncorrelated and hold the same standard deviation value σ, the standard deviation decreases along the segment and reaches its minimum at the midpoint, in which it is $\sigma/\sqrt{2} \approx 0.707\sigma$.

Some interesting characteristics can be observed: one would expect the value to increase as long as it gets further from the known measured endpoint, and eventually to reach its maximum at the midpoint.

In addition, the σ value is not influenced by the length of the segment. The minimum value at the midpoint is always circa 0.707σ, although one would expect unknown points between two near measured positions to be more precise than between two more distant measured positions.

This apparent paradox is well known is the GIS community and has been subject of research by several authors. In practice, the segment line is itself a fiction, and deviations of the truth from the straight line will tend to grow away from the endpoints, more than compensating for this effect [GOODCHILD, 2008]. Analogous concepts hold in case of a point interpolation inside a triangle, of course. A common strategy to describe uncertainty in linear objects, which has also been adopted in this work, consists then in the so-called epsilon-error bands.

The issue remains currently not completely resolved and is subject of further research, a solution may come from the adoption of interpolation algorithms like kriging, which permits to estimate prediction errors and take spatial correlation into account. However, this would not necessarily reduce the intervention of the user: obtaining a correct

empirical variogram requires, for example, that parameters like the lag size, the variogram fitting model or the definition of the searching neighbourhood are set. Moreover, a basic assumption of kriging is that the random field must be stationary, but this assumption does not correspond to the truth for most landscapes. So, it is necessary either to separate a trend model (but for which there is no correct definition) or to use a non-stationary variogram [KAREL, PFEIFER & BRIESE, 2006]. Furthermore, the variogram is calculated on the basis of distances obtained on the xy-plane. Its parameterisation on top of a triangulated surface is theoretically possible, but not immediately straightforward…
Finally, one further enhancement to this work should allow to overcome the initial assumption of the collar borders given as input parameters. This remains however another subject of research which the stochastic approach should eventually help to improve.

6.3 Outlook

In the previous paragraphs the proposed methodology, its strengths, its weak points and the possible improvements have been discussed together with some specific and other more general open issues. Bearing in mind the initial assumptions and the goals which were set at the beginning of this work, integration obtained so far with this deterministic approach already provides satisfactory results.
Looking forward and assuming that most of the previous open issues may be solved sooner or later, three topics will be mentioned here which could represent interesting further developments.
In the framework of digital terrain modelling, it could be worth to investigate how to integrate multiple zonally fragmented datasets. A border point may be shared by three countries having each a DTM with different characteristics. It could be therefore worth to extend the dual approach of the current methodology (high resolution vs. low resolution) to a more complex, simultaneous combination of datasets.
In the previous chapters it has been hinted at the analogies between signal processing and the characterisation of a DTM. A complex terrain can be decomposed into low- and high-frequency components in the frequency domain. A terrain that changes smoothly in the spatial domain contains information of lower frequency, while a terrain with accentuate spatial variations contains information of higher frequency. Even if spectral analysis may not have found its way to practice in the framework of digital terrain modelling, it could be indeed worth to create a variable filter for the transition surface using these techniques.
Finally, data integration problems in the framework of cultural heritage bear many common traits with those of the DTMs. On the other hand it could be extremely useful to integrate heterogeneous three-dimensional models with a fully three-dimensional approach which can preserve topology, geometric continuity, and guarantee a gradual transition between the surfaces of the different models.

Chapter 7

CONCLUSIONS

This thesis has dealt with integration problems of spatial data. Whenever two heterogeneous datasets share common geometric objects but have different extents, errors may generate at the borders. Typical errors are overlaps or gaps, which hinder continuity as soon as the datasets are merged.
In this work particular attention has been paid to how to embed a high resolution model into a low resolution digital terrain model. A deterministic methodology has been developed which allows to integrate two zonally fragmented datasets. It is possible to overcome the aforementioned inconsistencies by means of a transition surface which restores topological and geometrical continuity in a user-defined overlapping zone. In addition, a smooth transition in terms of height profile and point density is guaranteed.
Implementation has been carried out in a relational database management system using mainly open-source software. Tests have been performed on datasets at different resolutions in the framework of cultural heritage.

7.1 Scientific contribution

Although topological and geometric continuity is already achieved by other authors, in this work it is possible to gradually vary the point density inside the transition surface. The developed methodology can also handle overlapping zones with annular irregular shapes and containing multiple inner borders.
Since only data in the overlapping zone are used, creation of nested DTMs is enabled, and nothing happens elsewhere: the high resolution model and the low resolution model are left unchanged inside the inner border(s) and outside the outer border, respectively.
As a result, it is possible to obtain local updates of a low resolution DTM at specific areas, paying attention only to the borders of the overlapping zone.
This characteristic meets, for example, the needs of city modelling, where it could be necessary to add particular features to a DTM, like the entrance to an underground station.
In the framework of architectural and archaeological cultural heritage, the possibility of local updates is moreover appreciable, where, in consideration of the large amounts of available heterogeneous data, it is crucial to have a consistent representation and coherent visualisation.
Drawing on the presented examples, why not integrate, at different levels of detail, a small object of artistic interest with an architectural scale object, and then embed the resulting model into an urban scale model and continue forth with a wider area DTM?

References

ACHILLI V., AGUGIARO G., BRAGAGNOLO D., SALEMI G., STEVANATO G., 2008: *Modello digitale 3D dell'area archeologica di Montegrotto mediante integrazione di dati GPS e laser scanner in ambiente GRASS GIS*. Bollettino della Società Italiana di Fotogrammetria e Topografia n. 1, Anno 2008, pp. 83-99.

AGUGIARO G., SALEMI G., ACHILLI V., CERVATO C., CUPPARI F., 2008: *Il rilievo laser scanner con texturizzazione HDR dei portali della chiesa di Mattia a Budapest*. Recupero e conservazione n. 83 – Anno XIV, deLettera Editore, pp. 40-45.

BELS P.J., MCKAY N.D., 1992: A *method for registration of 3-D shapes*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 14(2), pp. 239-256.

BENNER J., GEIGER A., LEINEMANN K., 2005: *Flexible generation of semantic 3D building models*. In: (Gröger, Kolbe eds.) 1st International ISPRS/EuroSDR/DGPF-Workshop on Next Generation 3D City Models. Bonn, Germany, EuroSDR Publication N. 49.

BERLIN3D, The Official Berlin 3D City Model, http://www.3d-stadtmodell-berlin.de.

BRIESE C., 2004: *Three-dimensional modelling of breaklines from airborne laser scanner data*. International Archives of the Photogrammetry, Remote Sensing, and Geoinformation Sciences XXXV(3), pp. 1097-1102.

BURROUGH P. A., MCDONNELL R. A., 1998: *Principles of geographical information systems*. Oxford University Press, Oxford.

CITYGML, CityGML, www.citygml.org.

COHEN J., VARSHNEY A., MANOCHA D., TURK G., WEBER H., AGARWAL P. K., BROOKS JR F. P., WRIGHT W. V., 1996: *Simplification Envelopes*. Proceedings of Siggraph, New Orleans, LA.

DE FLORIANI L., MAGILLO P., PUPPO E., 1997: *Building and traversing a surface at variable resolution*. Proceedings of IEEE Visualization '97. pp. 103-110.

DOUGLAS-PEUCKER, Image source: http://en.wikipedia.org/wiki/File:Douglas_Peucker.png (http://creativecommons.org/licenses/by-sa/2.0/de/deed.en).

EUROSDR, 2006: *The EuroSDR Test "Checking and Improving of Digital Terrain Models"*. (Höhle and Potuckova eds.), Gopher, Utrecht, The Netherlands.

FELUS Y.A., CSATHÓ B., 2000: *Multi-source DEM evaluation and integration at the Antarctica transantarctic mountains project*. International Archives of Photogrammetry and Remote Sensing, XXXIII(B1), pp. 117-123.

FINAT J., FERNÁNDEZ MARTÍN J.J., FUENTES L., GONZALO M., MARTÍNEZRUBIO J., SAN JOSÉ J. I., 2005: *Ordering criteria for information management in 3D laser surveying of small urban spaces*. In: (Remondino, El-Hakim, Gonzo eds.) 3D-Arch'05, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVI-5/W17, ISSN 1682-1777.

FUENTES L., FINAT J., FERNÁNDEZ J. J., SANJOSÉ J. I., 2006: *Using Laser Scanning for 3D Urban Modeling*. Proceedings of 25th Urban Data Management Symposium, Aalborg, Denmark.

GARLAND M., HECKBERT P., 1997: *Surface Simplification using Quadric Error Bounds*. Proceedings of SIGGRAPH 97, pp. 209-216.

GARLAND M., HECKBERT P., 1998: *Simplifying Surfaces with color and texture using quadric error metrics*. Proceedings of IEEE Visualization 1998, pp. 263-269.

GOODCHILD M.F., 2008: *Imprecision and spatial uncertainty*. In (Shekhar & Xiong eds.) Encyclopedia of GIS, Springer, ISBN: 978-0-387-30858-6.

GOODCHILD M.F., BUTTENFIELD B.P., WOOD J., 1994: *Introduction to visualizing data quality*. In: (Hearshaw, Unwin eds.) Visualization in Geographic Information Systems. John Wiley and Sons, New York, USA.

GRÖGER G., KOLBE T. H., CZERWINSKI A., NAGEL C., 2008: *OpenGIS City Geographic Markup Language (CityGML) Encoding Standard, Version 1.0.0*. OGC Doc. n. 08-007r1, Open Geospatial Consortium.

GRÖGER G., PLÜMER L., 2005: *How to get 3-D for the price of 2-D. Topology and consistency of 3-D Urban GIS*. Geoinformatica, Vol. 9(2), , Springer, Netherlands, ISSN 1384-6175, pp. 139-158

GRÜN A., AKÇA D., 2005: *Least squares 3D surface and curve matching*. ISPRS Journal of Photogrammetry and Remote Sensing, Vol. 59, n. 3, pp. 151-174.

GUIDI G., FRISCHER B., LUCENTI I., 2007: *Rome Reborn - Virtualizing the ancient imperial Rome*. In: (Remondino, El-Hakim, Gonzo eds.) 3D-Arch'07, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVI-5/W47.

HEUVELINK G. B. M., 1998: *Error Propagation in Environmental Modelling with GIS*. CRC Press, ISBN: 978-0748407439.

HOPPE H., 1999: *New quadric metric for simplifying meshes with appearance attributes*. Proceedings of IEEE Vizualization '99, pp. 59-66.

HUTCHINSON M.F., GALLANT J.C., 1999: *Representation of terrain*. In: (Longley, Goodchild, Maguire, Rhind eds.) Geographical Information Systems, Vol. 1, Wiley and Sons, New York, USA, pp 105-124.

IEEE (INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS), 1990: *Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*. New York, NY.

IFC, buildingSMART International, http://www.iai-tech.org.

KAREL W., KRAUS K., 2006: *Quality parameters of digital terrain models*. In (Höhle eds.), Seminar on Automated Quality Control of Digital Terrain Models, Aalborg, Denmark. EuroSDR, ISBN: 978-9051794915.

KAREL W., PFEIFER N., BRIESE C., 2006: *DTM quality assessment*. ISPRS Technical Commission II Symposium 2006, International Archives of the ISPRS, XXXVI/2, ISSN: 1682-1750 pp 7-12.

KOCH A., 2005: *An Integrated Semantically Correct 2.5D Object Oriented TIN*. In: (Gröger & Kolbe eds.) 1st International ISPRS/EuroSDR/DGPF-Workshop on Next Generation 3D City Models. Bonn, Germany, EuroSDR Publication n. 49.

KOLBE T. H., 2008: *Representing and Exchanging 3D City Models with CityGML*. In: (Lee, Zlatanova eds.) 3rd International Workshop on 3D Geoinformation, Seoul, Korea. Lecture Notes in Geoinformation & Cartography, Springer Verlag.

KOLBE T. H., PLÜMER L., 2004: *Bridging the Gap between GIS and CAAD*. GIM International, 2004, n. 7.

KRAUS K., 2004: *Photogrammetrie. Band 1: Geometrische Informationen aus Photographien und Laserscanneraufnahmen*. Walter de Gruyter, Berlin, Germany.

KRAUS K., BRIESE C., ATTWENGER M., PFEIFER N., 2004: *Quality Measures for Digital Terrain Models*. In: (Altan eds.) ISPRS XX Congress, Istanbul, Vol. XXXV(B/2), pp. 113-118, ISSN: 1682-1750.

LATHAM R., BURNS D., 2006: *Dynamic Terrain Modification Using A Correction Algorithm*. Image 2006 Conference Proceedings, Scottsdale, Arizona USA.

LAURINI R., 1998: *Spatial multi-database topological continuity and indexing: a step towards seamless GIS data interoperability*. International Journal on Geographical Information Science, Vol. 12, n. 4.

LENK U., 2001: *2.5D-GIS und Geobasisdaten – Integration von Höheninformation und Digitalen Situationsmodellen*. Wiss. Arb. Fachrichtung Vermessungswesen, Universität Hannover Nr. 244; DGK bei der Bayerischen Akademie der Wissenschaften, Reihe C, n. 546. Diss., Universität Hannover.

LENZERINI M., 2002: *Data Integration: A Theoretical Perspective*. PODS 2002, pp. 233-246.

LI Z., 1993: *Theoretical models of the accuracy of digital terrain models: an evaluation and some observations*. The Photogrammetric Record 14(82), pp 651-660.

LI Z., ZHU Q., GOLD C., 2004: *Digital Terrain Modeling: Principles and Methodology*. CRC Press, ISBN 978-0415324625.

LUEBKE D., REDDY M., COHEN J., VARSHNEY A., WATSON B., HUEBNER R., 2002: *Level of Detail for 3D Graphics*. Morgan Kaufmann Publishing, ISBN 1-55860-838-9.

MIKHAIL, E. M., ACKERMANN F., 1976: *Observations and least squares*. IEP Dun-Donnelley, New York, USA.

MILLER P. E., 2008: *A robust surface matching technique for DEM integration in the context of coastal geohazard monitoring*. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 37(A3), pp. 147-154.

MITCHELL H. B., 2007: *Multi-sensor Data Fusion*, *an introduction*. Springer Verlag, Berlin, ISBN: 9783540714637.

MÖLLER T., 1997: *A Fast Triangle-Triangle Intersection Test*. Journal of Graphics Tools, Vol. 2, n. 2, pp. 25-30.

OGC, Open Geospatial Consortium Homepage: www.ogc.org.

OTEPKA J., BRIESE C., NOTHEGGER C., 2006: *First steps to a Topographic Information System of the next generation*. ISPRS Technical Commission Symposium, Goa, India; Symposium of ISPRS Commission IV - Geo Spatial Databases for Sustainable Development.

PESCARIN S., 2009: *Spatial data integration in real-time cooperative systems*. In: (Remondino, El-Hakim, Gonzo eds.) 3D-Arch'09, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII-5/W1, ISSN 1682-1777.

PFEIFER, N., 2002: *3D Terrain Models on the Basis of a Triangulation*. Ph.D. Thesis, Geoscience Report of the University Study Surveying and Geoinformation, Vienna University of Technology, ISBN: 3-9500791-7-3.

PLOUFFE S., SLOAN N. J. A., 1995: *The Encyclopaedia of Integer Sequences*. Academic Press 1995.

PLY, Polygon File Format, http://local.wasp.uwa.edu.au/~pbourke/dataformats/ply.

PODOBNIKAR T., 2005: *Production of integrated digital terrain model from multiple datasets of different quality*. International Journal of Geographical Information Science, Vol. 19, n. 1, pp. 69-89.

REISS P., 2002: High-quality DTMs: *Combining Laser Scanning, Digital Photogrammetry and Ground Survey*. GIM International, 16/11, pp. 40-43.

REMONDINO F., EL-HAKIM S., GRÜN A., ZHANG L., 2008: *Development and performance analysis of image matching for detailed surface reconstruction of heritage objects*. IEEE Signal Processing Magazine, Vol. 25(4), pp.55-65.

REMONDINO F., EL-HAKIM S., GIRARDI S., RIZZI A., BENEDETTI S., GONZO L. , 2009: *3D virtual reconstruction and visualization of complex architectures*. In: (Remondino, El-Hakim, Gonzo eds.) 3D-Arch'09, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII-5/W1, ISSN 1682-1777.

RONFARD R., ROSSIGNAC J., 1996: *Full-Range Approximation of Triangulated Polyhedra*. Computer Graphics Forum, Vol. 15(3), pp. 67-76.

RUTHER H., 2007: *Laser-Scanning and Heritage*. GIM International, Vol. 21, N. 5, pp. 14-17.

SADIQ M. G. S., GHOUSE M., DUCKHAM M., 2009: *Integrated storage and querying of spatially varying data quality information in a relational spatial database*. Transactions in GIS v.13, N.1, pp. 25-42.

SCHENK T., CSATHÓ B., 2002: *Fusion of LiDAR data and aerial imagery for a more complete surface description*. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXIV, Part 3A, pp. 310-317.

SHEWCHUK J.R., 1996: *Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator*. In: (Ming, Lin, Dinesh eds.) Applied Computational Geometry: Towards Geometric Engineering, Vol. 1148 of Lecture Notes in Computer Science, Springer Verlag, Berlin, pp. 203-222.

SCHMITTWILKEN J., SAATKAMP J., FÖRSTNER W., KOLBE T.H., PLÜMER L., 2007: *A Semantic Model of Stairs in Building Collars*. Photogrammetrie, Fernerkundung, Geoinformation PFG., pp. 415-428.

STADLER A., KOLBE T.H., 2007: *Spatio-Semantic Coherence in the Integration of 3D City Models*. In: Proceedings of the 5th International Symposium on Spatial Data Quality ISSDQ 2007, Enschede, Netherlands, ISPRS Archives.

SURAZHSKY V., SURAZHSKY T., KIRSANOV D., GORTLER S. J., HOPPE H., 2005: *Fast exact and approximate geodesics on meshes*. ACM Transactions on Graphics, Vol. 24(3), pp. 553-560, ISSN: 0730-0301.

VARSHNEY A., 1994: *Hierarchical geometric approximations*. Ph.D. Thesis, TR-050-1994, Department of Computer Science, University of North Carolina, Chapel Hill, NC 27599-3175.

VERONA, Arche Scaligere. Image source: http://en.wikipedia.org/wiki/Scaliger_Tombs.

WARRINER T., MANDLBURGER G., 2005: *Generating a new high resolution DTM product from various data sources*. (Dieter Fritsch Ed) Photogrammetrische Woche 2005, Wichmann Verlag, Heidelberg, pp. 197-206.

WENG Q., 2002: *Quantifying Uncertainty of Digital Elevation Models Derived from Topographic Maps*. In: (D. Richardson, P. van Oosterom eds.) Advances in spatial data handling, Springer, New York, pp. 403-418.

WOOD J., 1996: *The geomorphological charaterisation of Digital Elevation Models*. Ph.D. thesis, Department of Geography, University of Leicester, Leicester, UK.

ZIEGLER P., DITTRICH K.R., 2004: *Three Decades of Data Integration - All Problems Solved?*. WCC 2004, pp. 3-12.

Free and open-source software resources:

POSTGRESQL: http://www.postgresql.org

POSTGIS: http://postgis.refractions.net

QUANTUM GIS: http://www.qgis.org

GRASS GIS: http://grass.itc.it

SIMPLIFICATION ENVELOPES: http://www.cs.unc.edu/~geom/envelope.html

TRIANGLE: http://www.cs.cmu.edu/~quake/triangle.html

MESHLAB: http://meshlab.sourceforge.net

Acknowledgements

I owe gratitude to Fondazione Cassa di Risparmio di Padova e Rovigo and to DAAD (Deutscher Akademischer Austausch Dienst) which granted scholarships to carry out great part of this research work in Italy and Germany.

I would like to thank Prof. Norbert Pfeifer at TU Wien, Austria, and Prof. Javier Finat at Universidad de Valladolid, Spain, who kindly accepted to read and referee this thesis.

The datasets presented in this work have been acquired and partly pre-processed by the members of the Laboratorio di Rilevamento e Geomatica, University of Padova. My thanks to all the Italian colleagues for their valuable help, in particular to Prof. Giuseppe Salemi for his support during my absence from Padova.

My gratitude goes to all the staff at the Institut für Geodäsie und Geoinformationstechnik, TU Berlin, who has welcomed me from the very first day. In particular, I am thankful to Prof. Thomas Kolbe, who directly invited me to Berlin after we met for the first time in a rainy cold afternoon of November 2006. Thanks to him I could understand the meaning of the German word *Doktorvater* and enjoy with him fruitful discussions ranging from RANSAC algorithms to *warping donuts*. It is no mystery that the term *Donutology* became, for this thesis, a cheerful running gag along the banks of the river Spree.