

TERRAIN GENERALIZATION WITH MULTI-SCALE PYRAMIDS CONSTRAINED BY CURVATURE

B. Jenny, L. Hurni

Institute of Cartography, ETH Zurich – (jenny, hurni)@karto.baug.ethz.ch

KEY WORDS: Terrain generalization, DEM frequency filtering, Laplacian pyramid, Terrain Equalizer

ABSTRACT:

This paper presents a method for terrain generalization using Laplacian pyramids. The method pre-processes digital terrain for generating cartographically generalized 3D maps. Map authors can attenuated or amplify selected frequency bands of the terrain with a graphical interface imitating an audio equalizer. Ridge lines and valley lines are localized by curvature indices, and their characteristic shapes are preserved or emphasized using additional sets of equalizer controls. Frequency bands are adjusted separately for the foreground and the background to remove visually disturbing terrain detail in the background. This is relevant for 3D maps in central perspective projection considerably compressing distant terrain features. The proposed generalization method was implemented in Terrain Equalizer, a free and open-source application providing a graphical interface with interactive 3D previews (available at <http://www.terraincartography.com/>). Using this application, disturbing high frequency details can be easily removed and major mountain forms can be accentuated. The level of generalization can be adjusted seamlessly from the foreground to the background of the 3D map. Topographic break lines, such as ridge or valley lines, are successfully preserved, which is important for conveying the characteristic shape of a generalized terrain.

1. INTRODUCTION

Expert cartographers know how to adequately generalize terrain for 3D maps. Unnecessary and visually distracting details are removed, important landforms are accentuated, and characteristic terrain features and typical shapes are preserved. However, generalization methods for 3D maps are most often technology driven, i.e. they are often not targeted at cartographic generalization, but intend to reduce the amount of terrain details in order to achieve responsive frame rates when rendering terrain in video games and other interactive environments. Such methods may introduce artefacts, such as artificial edges and flat triangle structures, or overly smoothed mountain ridges and valleys.

3D maps most often depict terrain in a central perspective projection, compressing terrain features in the background. The terrain compression increases with the distance to the viewer: the foreground is depicted at a large scale, while the background is rendered at a very small scale. This results in an excessively detailed background, especially when using high-resolution terrains. Such highly detailed representations are often visually inefficient because the main landforms are not clearly discernable due to the many distracting details. The dense details obscure macro topography—it is impossible to see the forest for the trees. Hence, the level of terrain generalization must seamlessly increase with the distance to the viewer.

Cartographic generalization is an inherently visual task. A graphical environment is required for adjusting the generalization parameters in a trial-and-error approach to the spatial resolution of the terrain, the landscape features and the purpose of the map. A problem plaguing authors of 3D maps, however, is the currently available software for generalizing terrain, which

sometimes is difficult to comprehend and control, or does not offer a WYSIWYG preview mode for interactive manipulation in real-time.

The research presented in this article aims at making a contribution to the solution of these problems. Visually disturbing small details are to be removed from digital terrain, while sharp edges and mountain ridges are to be retained. The goal is to preserve the characteristic appearance of a terrain, and to seamlessly adjust the amount of generalization from foreground to background. We are aiming at a generalization method that is easy to comprehend and control by authors of 3D maps, and provides feedback in real time.

For our prototype, an equalizer metaphor was chosen as basis for the user interface. In audio engineering, the equalizer filter is used to alter the frequency response of an audio source. Most users should be familiar with equalizers from audio playback software and physical audio equipment. The generalization method and its related user interface were implemented in Terrain Equalizer, a free and open-source software application.

2. RELATED WORK

In this overview of related research, we focus on the generalization of terrain models using related frequency-based methods. We neglect the generalization of representations (such as contour lines) derived from terrain models. Also, we concentrate on cartographic generalization aiming at visual effectiveness to generate clear and legible maps, and do not treat generalization to improve data storage or to accelerate rendering, analysis or transmission of terrain data.

2.1 Filtering

The term filtering is understood in the context of signal processing, *i.e.* the application of an operator (the filter) that removes frequency components from the signal (the terrain). Such filters are known from image processing, and are easily applied to gridded terrain models. Low-pass filters remove high-frequency details, whereas high-pass filters emphasise discontinuities. While high-pass filters are rarely applied to digital terrain (Weibel and Heller, 1991), low-pass filters are the simplest and thus most commonly used filters (Li, 2008). However, if prominent topographic breaks are involved, cartographic practice tells us to preserve their character and not simply smooth them. The applicability of global smoothing or low-pass filtering, therefore, is limited to smooth rolling terrain and minor scale reductions (Weibel, 1992).

2.2 Wavelet transforms

Relatively few attempts have so far applied the wavelet transform on digital terrain (Gallant and Hutchinson, 1996). Wavelet applications on DEMs are mainly used for feature extraction (Amgaa, 2003; Kalbermatten, 2010) and terrain simplification by filtering the wavelet coefficients (Mahler, 2001; BJORKE and Nilsen, 2003). The related Fourier transform has been found to be less adequate for terrain data (for a review see Kalbermatten, 2010, p. 47).

2.3 Skeleton and contour lines

Weibel (1992) proposed the generalization of terrain based on structure lines that are extracted from the terrain model to be generalized in the process, in some cases using techniques developed for line generalisation. Weibel (1992) proposes to use the drainage network and the ridge network, both extracted from the terrain model, depending on the characteristics of the terrain. Automatic tools exist to characterize terrain patches and steer this process. Such methods can also be applied to morphometric analysis, for example, to selectively remove valleys (Jordan, 2007).

2.4 Multi-resolution terrain

Research in computer graphics has resulted in a variety of methods for reducing the level of detail (LOD) of terrain models. Many of those methods remove terrain details from the background of the scene, aiming at achieving interactive frame rates in games and other applications with interactive terrain visualization. In such applications the main goal is the minimization of data storage, which leads to the application of TIN structures or nested regular grids at varying resolution (for an overview see Pajarola and Gobbetti, 2007). The goal, however, is a purely technical one: The amount of triangles rendered by the computer graphics hardware is reduced, such that visual artefacts are minimized and an interactive frame rate is achieved. This goal is fundamentally different from cartographic generalization for 3D maps, where irrelevant and disturbing micro topography has to be removed, and characteristic terrain structures need to be accentuated.

Multi-resolution TINs can be constructed using various methods, for example, the line simplification method by Douglas and Peucker applied to terrain (Fei and He, 2009), or techniques inspired by Morse functions (Danovaro *et al.*, 2006). Besides terrain visualization, multi-resolution terrain is also applied to filtering LIDAR data (Silván-Cárdenas and Wang, 2006).

3. FREQUENCY MANIPULATION WITH LAPLACIAN PYRAMIDS

The main idea of our method presented in this paper is to consider a terrain model as a two-dimensional input signal, where selected frequencies are enhanced or attenuated. A Laplacian pyramid is used to separate consecutive frequency bands. Individual frequency bands are amplified or attenuated to be combined into a new terrain synthesized from the weighted frequency bands.

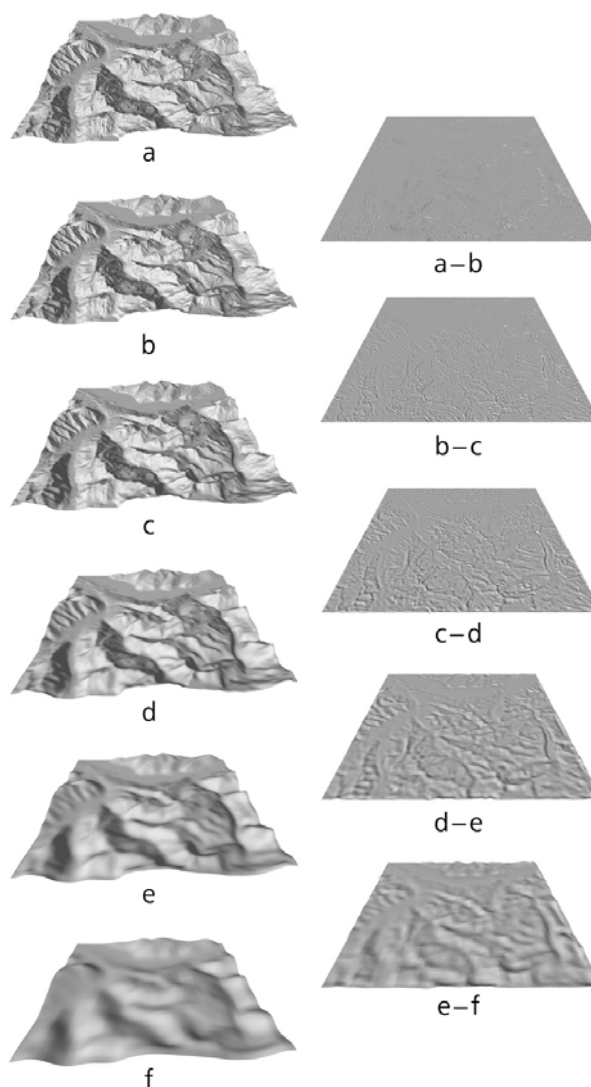


Figure 1. Gaussian pyramid (left column) and Laplacian pyramid (right column) of a terrain model. Levels b to f are enlarged to the size of level a.

Laplacian pyramids are commonly used in computer graphics for image blurring, image mosaicing, 3D texturing or for applications in computer vision. The construction and application of Laplacian pyramids for imaging was pioneered by Burt (1981) and his colleagues (Burt and Adelson, 1983; Burt and Edward, 1983). The construction of Laplacian pyramids is identical for both raster images and raster terrain models. Only the data type used by the algorithm varies, *i.e.* 8 bit or 16 bit integer numbers for each image channel *vs.* floating point values for gridded terrain models.

A Laplacian pyramid is derived from a Gaussian pyramid. A Gaussian pyramid consists of a series of raster grids, which are blurred using a Gaussian filter and scaled down (left column in Figure 1). Blurring and scaling are applied multiple times, creating a stack of successively smaller grids, with each pixel containing a local average that corresponds to a neighbourhood on a more detailed level of the pyramid. Hence, a series of Gaussian low-pass filters are applied, successively removing high-frequency details (a to f in Figure 1).

The Gaussian pyramid is then used to derive the Laplacian pyramid. Each level of this new Laplacian pyramid consists of the difference between two consecutive levels in the previous Gaussian pyramid (right column in Figure 1). To compute one level of the Laplacian pyramid, two levels of the previous Gaussian pyramid are first brought to the same grid size and then corresponding values of the two grids are subtracted. The resulting levels of the Laplacian pyramid encode successive frequency bands. For example, the level labelled a–b in Figure 1 contains the terrain’s highest frequency range.




Figure 2. Removing high-frequency details and amplifying a mid-frequency band. Left: original terrain; right: filtered terrain.

The original terrain model can be reconstructed from the Laplacian pyramid by summing all its levels, and adding the base level of the Gaussian pyramid. Using the labels of Figure 1, this can be expressed by:

$$a = f + (e - f) + (d - e) + \dots + (a - b) = f + \sum (L_{i+1} - L_i) \quad (1)$$

where a = reconstructed grid
 $L_i = a \dots f$ = levels of the Gaussian pyramid

The reconstructed grid a is identical to the original unfiltered grid. Equation 1 implicitly assigns a unary weight to each level of the Laplacian pyramid. To selectively amplify or attenuate frequency bands, a weighted reconstruction is needed:

$$a' = w_f f + \sum w_i (L_{i+1} - L_i) \quad (2)$$

where a' = reconstructed grid
 f = smallest level of the Gaussian pyramid
 w_f = weight coefficient for f
 L_i = levels of the Gaussian pyramid
 w_i = weight coefficients for L_i

The user interactively adjusts the weight coefficients w_{ij} . Figure 2 shows the visual effect when attenuating high frequency bands and exaggerating a mid-frequency band. The weight coefficients in this particular example are: $w_{ab} = 0\%$, $w_{bc} = 25\%$, $w_{cd} = 200\%$, and $w_f = w_{ef} = w_{de} = 100\%$. The amplification w_{cd} of the mid-frequency band accentuates valleys.

4. CURVATURE FOR EDGE PRESERVATION

Filtering high-frequency bands removes small details, but also overly smooths characteristic sharp-edged break lines, such as ridge and valley lines. Curvature indices can be used for dealing with this problem. The basic idea is to automatically identify topographic break lines and apply different weight coefficients to their neighbourhood. This is in line with Weibel (1992), who notes that the basic model of low-pass filtering can be improved further by including a second operation to preserve discontinuities and thus restrict simplification to the regions between them. The proposed method modifies the traditional reconstruction of Laplacian pyramids. The traditional reconstruction as described in the previous section applies isotropic weight coefficients to the frequency bands, *i.e.* coefficients do not change with location. This is expressed in Equation 3, where a frequency band is added to the sum of the lower frequency bands:

$$sum_{0 \dots i+1} = sum_{0 \dots i} + w_i B_i \quad (3)$$

where $sum_{0 \dots i+1}$ = weighted sum of bands 0 to $i + 1$
 $sum_{0 \dots i}$ = weighted sum of bands 0 to i
 $B_i = L_{i+1} - L_i$ = band of the Laplacian pyramid
 w_i = weight coefficient for B_i

In our approach for terrain generalization, the coefficient for each frequency band changes spatially with curvature. Curvature

coefficients are appropriate for modelling geomorphometric elements like ridges and valleys (Wood, 1996), and have been applied for the generalization of terrain for relief shading (Leonowicz et al., 2010). We use the plan curvature index, which measures the rate of change of aspect along a contour line in the horizontal plane; it differentiates between convex and concave forms, and defines sharp and clear lines of ridges and valleys (Wilson and Gallant, 2000). Equation 3 is extended as follows to take curvature into account when reconstructing the terrain from the Laplacian pyramid:

$$sum_{0...i+1} = sum_{0...i} + (w_i + w_{ci} curv_{0...i}) B_i \quad (4)$$

where w_{ci} = ridgeline weight coefficient for B_i
 $curv_{0...i}$ = local plan curvature of $sum_{0...i}$

w_i and w_{ci} are anisotropic weight coefficients, while $curv_{0...i}$ is locally variable. The plan curvature coefficient $curv_{0...i}$ is computed from the sum of frequency bands with lower frequencies than the current band B_i . This approach has the desired effect of accentuating ridgelines: At locations where lower frequencies contain concave or convex shapes, terrain curvature is further accentuated with details of the added band B_i (Figure 3).




Figure 3. A second set of weight coefficients for preserving break lines. Left: original terrain; middle: global filtering; right: local filtering steered by curvature.

Plan curvature is calculated with the Evans–Young method, which uses a quadratic trend surface for each cell, fitted to the 3×3 local elevation matrix. For formulas see Pennock et al., 1987; Wood, 1996; Shary et al., 2002. Two transformations are applied to the Evans-Young coefficient to find $curv_{0...i}$ in Equation 6. First, coefficients are computed for all terrain cells, then stored in a grid and blurred by convolving with a two-dimensional Gaussian bell curve. Without this blurring, spiky structured would appear along break lines in the reconstructed terrain. The second transformation is an exponentiation applied to the blurred

curvature coefficients. An exponent close to 0 results in mountain slopes with many details, while an exponent close to 1 adds high frequency details along break lines. The user can adjust the size of the blurring kernel and the exponent.

The procedure outlined so far concentrates on convex ridgelines and ignores the generalization of concave valley lines. However, a third set of user-defined weight coefficients can be applied along concave lines in a similar way. The value of w_{ci} in Equation 4 is either chosen from the user-defined coefficient set for ridges or from the one for valleys – depending on whether the frequency band B_i adds convex or concave details. To determine what type of curvature is added, the local profile curvature index of the frequency band B_i is computed (Wilson and Gallant, 2000). Figure 3 illustrates the effect of a second set of weight coefficients for mountain ridges (*i.e.* convex curvature), and a third set of coefficients for valleys (*i.e.* concave curvature). The figure in the middle shows attenuated high and mid-frequency bands. The figure on the right restores details along ridge and valley lines by amplifying the coefficients of the frequency bands attenuated before.

5. VIEW DEPENDENT FILTERING FOR 3D MAPS



Figure 4. Increasing filtering from foreground to background. Left: original terrain; right: filtered terrain, with ridges accentuated in the background.

As outlined in the introduction, the level of terrain generalization must increase with the distance to the viewer of the scene when central perspective projection is used. The proposed method can

be extended to variably adjust the level of generalization in the foreground and the background. Two sets of weight coefficients are defined, one for the foreground and one for the background. The weight coefficients w_i and w_{ci} in Equation 4 are then linearly interpolated between the foreground and the background. Equation 5 exemplifies the linear interpolation for the coefficient w_i . The interpolation of the coefficient w_{ci} is identical.

$$w_i = w_d w_{ifore} + (1 - w_d) w_{iback} \quad (5)$$

where w_i = the blended weight coefficient
 w_{ifore}, w_{iback} = foreground and background coefficients
 w_d = weigh factor: 1 in foreground, 0 in background

Figure 4 shows the type of generalization effect that can be achieved. On the right side of Figure 4, details in the highest frequency band are slightly attenuated in the foreground, and ridges in the highest frequency band are moderately accentuated. In the background, the three highest frequency bands are completely removed, except for details along ridges, which are accentuated to preserve the spiky character of this alpine terrain.

6. MASKED FREQUENCY FILTERING

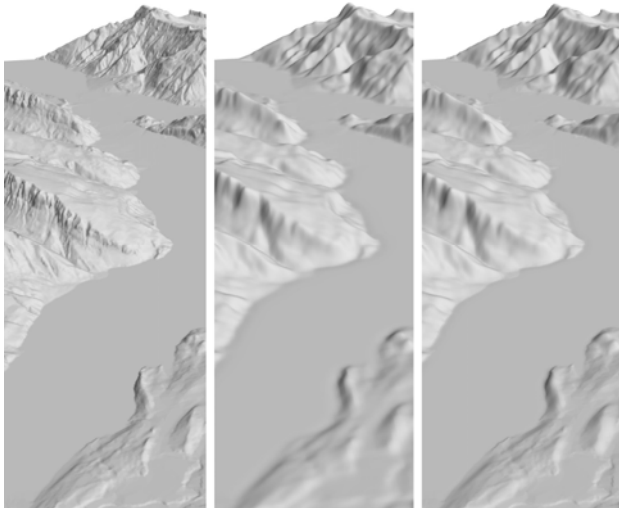


Figure 5. Masks for avoiding artefacts in flat areas. Left: original terrain; middle: globally filtered terrain; right: filtering limited to steep areas.

The amplification of mid-frequency bands accentuates ridges and valleys, but also creates disturbing concave depressions along valley bottoms, especially where steep slopes adjoin flat planes. This problem can be alleviated with a mask for flat areas. One approach is to extract slope indices, and limit the amplification of mid-frequency bands to flat terrain. Other types of masks can be used as well, for example, a cartographically designed mask with lakes and valley bottoms. Whatever type of mask used, a simplistic combination of (a) the terrain reconstructed from a weighted Laplacian pyramid with (b) the original terrain would

result in discontinuities in the combined terrain. Visually disturbing stepped structures would appear, because the reconstructed terrain might have been locally shifted in vertical direction due to the values of the coefficients w_i . Hence, masked areas must be taken into account when reconstructing the terrain from the Laplacian pyramid:

$$sum_{0 \dots i+1} = sum_{0 \dots i} + [(w_i + w_{ci} \text{curv}_{0 \dots i} - 1) w_{mi} + 1] B_i \quad (6)$$

where w_{mi} = local mask value [0 ... 1]

The mask is first transformed to a Gaussian pyramid, and w_{mi} is the corresponding level i of this pyramid. The mask weight w_{mi} locally varies between 0 and 1. If w_{mi} is 0, the frequency band B_i is added with no filtering applied, and locally the unfiltered original terrain model is reconstructed. If w_{mi} is 1, Equation 6 simplifies to Equation 4. Figure 5 illustrates the masking effect when slope values mask flat areas: artefacts are reduced where the steep slopes adjoin the flat plane.

7. IMPLEMENTATION AND EVALUATION

The described method for generalizing digital terrain was implemented in Terrain Equalizer, a free and open-source application available at <http://www.terraincartography.com/>. Terrain Equalizer was written in Java and runs on all major platforms. For rendering previews of the generalized terrain model, Terrain Equalizer builds on JOGL, a Java wrapper for OpenGL (JOGL, 2009; OpenGL.org, 2008). Terrain Equalizer's interface features sets of vertical sliders for adjusting the weight coefficients of individual frequency bands, using an equalizer metaphor.

As outlined in section 2, various alternatives exist to Laplacian pyramids for selectively manipulating frequency bands, such as low-pass filters or wavelet transforms. Compared to these alternatives, the method presented here is easy to comprehend and control by the user. Also, the algorithms are relatively simple to implement and fast enough for interactive manipulation of terrains with up to several million cells.

Weibel (1992) notes that frequency-based generalization is somewhat limited. Terrain can only be simplified or emphasized, which is also the case for our approach. Other generalization operators, such as displacement or combination operators, are not addressed by our research. Weibel (1992) also suggested generalization operators imbedded in an interactive graphical environment, since the objective of cartographic terrain generalization is not so much noise removal as it is graphical simplification. Terrain Equalizer offers such a user interface. We applied it for 3D mapping using terrain models of varying resolution and origin. We were able to successfully remove disturbing high-frequency details, while preserving important break lines; and we successfully enhanced important terrain structures by amplifying mid-frequency bands. Filtering parameters could easily be adjusted to the terrain character and to various levels of generalization. Furthermore, it proved to be easy to adjust the level of generalization to the distance from the viewer. However, we encountered problems where steep slopes

and flat valleys meet, when amplifying mid-frequency bands. These problems are at least partially solved by the inclusion of masks – nevertheless, their integration in Terrain Equalizer could be improved and easier to control.

Various experts have tested Terrain Equalizer and provided encouraging feedback. A more detailed quantitative analysis of the method is difficult. As with all research in cartographic generalization, only qualitative impressions can be gained from visual analysis to assess the suitability of a generalization operator. Only the cartographic results can tell whether the generalization operation was successful (Weibel, 1992).

Parameters are seamlessly interpolated from the background to the foreground throughout a 3D map. The portability of such a terrain with a generalization level varying from foreground to background is obviously limited, as it can only be viewed in a single predefined direction. This limitation could be overcome for interactive terrain visualizations – where the user can freely choose the position and direction of view – by generalizing the terrain in real time. Interactive frame rates could possibly be achieved by porting the Laplacian pyramid to modern GPUs, as proposed by Strengert (2006) for image processing.

Terrain Equalizer and the underlying algorithms were developed for generalizing digital terrain of mountainous areas at large and medium scales, specifically for the representation as 3D maps. Alternative applications of Laplacian pyramids constrained by curvature remain to be explored for terrain generalization and visualization. Particularly, this approach could be used to remove artefacts from SRTM or LIDAR terrains, and it could also be applied to terrain before generating 2D visualizations, such as contour lines or shaded relief.

8. REFERENCES

Amgaa, T., 2003. Wavelet-based analysis for object separation from laser altimetry data. Master's thesis, ITC, Enschede, NL.

Bjorke, J., Nilsen, S., 2003. Wavelets applied to simplification of digital terrain models. *Int. J. Geographical Information Science*, 17(7), pp. 601–621.

Burt, P. J., 1981. Fast filter transform for image processing. *Computer Graphics and Image Processing*, 16(1), pp. 20–51.

Burt, P. J., Adelson, E. H., 1983. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4), pp. 532–540.

Burt, P. J., Adelson, E. H., 1983. A multiresolution spline with application to image mosaics. *ACM Trans. Graph.*, 2(4), pp. 217–236.

Danovaro, E., De Floriani, L., Papaleo, L., Vitaliet, M., 2006. A multi-resolution representation for terrain morphology. In: GIScience 2006, LNCS 4197, pp. 33–46.

Fei, L., He, J., 2009. A three-dimensional Douglas-Peucker algorithm and its application to automated generalization of

DEMs, *International Journal of Geographical Information Science*, 23(6), pp. 703–718.

Gallant, J. C., Hutchinson, M. F., 1996. Towards an understanding of landscape scale and structure. In: Third International Conference/Workshop on Integrating GIS and Environmental Modeling, Santa Fe, NM, USA.

Hopf, M., Ertl, T. 2000. Hardware accelerated wavelet transformations. In: *Proceedings EG/IEEE TCVG Symposium on Visualization VisSym '00*, pp. 93–103.

JOGL, 2009. Java binding for the OpenGL API Wiki, <http://kenai.com/projects/jogl/pages/Home> (accessed 13 September 2010).

Jordan, G., 2007. Adaptive smoothing of valleys in DEMs using TIN interpolation from ridge elevations: an application to morphotectonic aspect analysis. *Computers & Geosciences*, 33(4), pp. 573–585.

Kalbermatten, M., 2010. Multiscale analysis of high resolution digital elevation models using the wavelet transform. PhD, EPFL Lausanne.

Leonowicz, A.M., Jenny, B., Hurni, L., 2010. Automated reduction of visual complexity in small-scale relief shading. *Cartographica*, 45(1), pp. 64–74.

Li, Z., 2008. Multi-Scale Digital Terrain Modelling and Analysis. In: *Advances in Digital Terrain Analysis*, Lecture Notes in Geoinformation and Cartography, Springer Berlin Heidelberg, pp. 59–83.

Mahler, E. 2001. Scale-dependent filtering of high resolution digital terrain models in the wavelet domain. Master's thesis, Departement of Geography, University of Zurich.

OpenGL.org, 2008. OpenGL – the industry's foundation for high performance graphics, <http://www.opengl.org> (accessed 13 September 2010).

Pajarola, R., Gobbetti, E., 2007. Survey on semi-regular multiresolution models for interactive terrain rendering. *The Visual Computer*, 23(8), pp. 583–605.

Pennock, D. J., Zebarth, B. J., de Jong, E., 1987. Landform classification and soil distribution in hummocky terrain, Saskatchewan, Canada. *Geoderma*, 40, pp. 297–315.

Shary, P. A., Sharaya, L. S., Mitusov, A. V., 2002. Fundamental quantitative methods of land surface analysis. *Geoderma*, 107, pp. 1–32.

Silván-Cárdenas, J. L., Wang, L., 2006. A multi-resolution approach for filtering LiDAR altimetry data. *ISPRS Journal of Photogrammetry & Remote Sensing*, 61, pp. 11–22.

Strengert M., Kraus, M., Ertl, T., 2006. Pyramid methods in GPU-based image processing. In: *Proceedings of Vision, Modeling, and Visualization 2006*, pp. 169–176.

Weibel, R., 1992. Models and experiments for adaptive computer-assisted terrain generalization. *Cartography and Geographic Information Science*, 19(3), pp. 133–153.

Weibel, R., Heller, M., 1991. Digital terrain modelling. In: *Geographical Information Systems: Principles and Applications*, pp. 269–297, Longman, London.

Leicester.

Wilson, J. P., Gallant, J. C., 2000. *Terrain Analysis: Principles and Applications*. New York: Wiley.

Wood, J. D., 1996. The geomorphological characterisation of digital elevation models. PhD, University of