

Roman Geisthövel

**Automatic
Swiss
style rock
depiction**

**Diss. *ETH* No. 24328
2017**

DISS. ETH NO. 24328

Automatic Swiss style rock depiction

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

ROMAN WOLFGANG GEISTHÖVEL

Dipl. Inf., University of Hamburg

born on 05. 07. 1975

citizen of
Germany

accepted on the recommendation of

Prof. Dr. Lorenz Hurni
Prof. Dr. Ralf Bill

2017

Abstract

The main objective of this thesis is the reproduction of the style of rock depiction featured in the Swiss National Map by digital means. Since the production of such rock hachures is still done manually by trained specialists, automating the process is going to benefit experts and laypersons alike, who want to include Swiss style rock hachures into their maps. Additionally, automation offers fast, reproducible and economically feasible results.

After describing the anatomy of rock hachures, and a brief historical survey, current approaches to automating rock depiction are discussed, resulting in the following principles for the development of the novel method.

Within a map, rock depiction is no isolated concept. It is intimately connected to the depiction of shaded relief. Both map elements represent a lighting model that is focused on preserving contrast across terrain edges like ridges and ravines. This model is quite different from analytical Lambertian shading employed in many current approaches. On a structural level, shaded relief is simpler than rock hachures: a variation of grey values versus multiple kinds of jittered hachures, i. e. broad form strokes and narrow contour strokes for indicating the rock skeleton, as well as fill hachures depicting light-modulated faces. From this structural disparity follows the primacy of the shaded relief in the workflow. The rock hachures are derived from the shaded relief. As with all map elements, the issue of generalization needs to be addressed for rock hachures and the shaded relief.

These principles are implemented as follows. The user supplied raster elevation model is generalized along slope lines using line integral convolution. For the purpose of determining terrain edges, two shaded reliefs are generated, using a main lighting direction and an orthogonal direction. The terrain edges are computed by merging the discontinuities of the two shaded reliefs, extracted using an edge detection algorithm. Grey values are assigned to either side of every terrain edge, while care is being taken that the assignments maintain contrast and generally observe the main lighting direction. The grey values of the remaining faces between terrain edges are computed by solving a Laplace equation, where the already assigned grey values serve as the boundary condition. The rock faces are covered with fill hachures, where the stroke width is derived from the grey values of the shaded relief. The form and contour strokes are derived by offsetting the terrain edges according to the gradient of the shaded relief.

The method is automatic in the sense that, apart from the raster elevation model and a mask indicating the raster cells covered by rock, the user is only supposed to enter declarative parameters, e. g. the main lighting direction. It follows that in using this method, expert knowledge is no decisive factor. Different users providing the same input will arrive at equal results.

Zusammenfassung

Das zentrale Thema dieser Arbeit ist die Nachempfindung der Felsdarstellung der Manier der Landeskarte der Schweiz mithilfe von digitalen Methoden, Imitation mittels Automatisierung. Liegt die Produktion solcher Schraffurdarstellungen nach wie vor in den Händen von geschulten Fachkräften, sollen durch die Automatisierung neben Experten auch Laien bei der Anwendung der Schweizer Manier in Karten unterstützt werden. Zudem ermöglicht die Automatisierung eine schnelle, reproduzierbare und kostengünstige Herstellung.

Nach der Beschreibung der Anatomie, sowie einem kurzen Überblick der historischen Entwicklung von Felsschraffen, werden bestehende Ansätze zur Automatisierung der Darstellung gesichtet, woraus sich folgende Grundsätze zur Entwicklung der neuartigen Methode ergeben.

Die Felsdarstellung existiert nicht isoliert in einer Karte. Sie ist engstens mit der schattierten Reliefdarstellung verknüpft. Beide Darstellungen sind Ausdruck eines Beleuchtungsmodells, welches die kontrastwahrende Abstimmung von Hell-Dunkel-Übergängen an Geländekanten wie Graten und Rinnen in den Vordergrund stellt. Dieses entspricht nicht dem zahlreich in bestehenden Ansätzen verwendeten Lambertschen Beleuchtungsmodell. Strukturell ist das schattierte Relief einfacher als Felsschraffen: eine Variation von Grauwerten im Gegensatz zu den verschiedenartigen verzitterten Schraffen, sogenannten breiten Formstrichen und schmalen Konturstrichen zur Verdeutlichung des Felsgerippes, sowie Füllschraffen zur lichtmodulierten Darstellung von Flächen. Aus diesem Gegensatz ergibt sich der Primat des schattierten Reliefs in der Methode. Die Felsschraffen können aus dem Relief abgeleitet werden. Wie bei allen Kartenelementen, bedarf es bei der Felsdarstellung, als auch beim schattierten Relief, einer Methode der Generalisierung im Hinblick auf den Kartenmassstab.

Diese Grundsätze werden folgendermassen umgesetzt. Ausgehend von einem Rasterhöhenmodell wird ein generalisiertes Höhenmodell errechnet. Die Generalisierung erfolgt längs von Falllinien mittels Linienintegralfaltung. Aus dem generalisierten Höhenmodell werden zum Zweck der Bestimmung der Geländekanten zwei schattierte Reliefs berechnet, in Lichthaupttrichtung und einer orthogonalen Richtung. Die Geländekanten entstehen aus der Zusammenführung der durch einen Kantendetektor ermittelten Unstetigkeiten der beiden Reliefs. Den Geländekanten werden beidseitig, kontrastwährend Grauwerte zugewiesen, unter Beachtung der generellen Einhaltung der Hauptlichtrichtung. Die restlichen Flächen zwischen den Geländekanten werden durch Lösung einer Laplace-Gleichung schattiert, wobei die bereits zugewiesenen Grauwerte an den Kanten die Randbedingung bilden. Die Felspartien werden mit Füllschraffen belegt, die Modulation der Schraffenbreite erfolgt dabei anhand der Grauwerte des Reliefs. Die Form- und Konturstriche entstehen durch den Versatz von Geländekanten anhand des Reliefgradienten.

Die Methode verläuft automatisch in dem Sinne, daß neben dem Höhenmodell und einer Maske, die angibt, welche Rasterzellen von Fels bedeckt sind, ausschliesslich deklarative Parameter, beispielsweise die Hauptlichtrichtung, angegeben werden müssen. Das bedeutet, daß Expertenwissen beim Gebrauch der Methode nicht ausschlaggebend ist. Verschiedene Benutzer/innen erhalten unter Eingabe der gleichen Parameter das gleiche Ergebnis.

Contents

1 Introduction and problem statement	1
2 The elements of Swiss style rock depiction	4
2.1 Strokes and hachures	4
2.2 Historical background	7
3 A review of prior work on the automation of rock depiction	10
3.1 Survey	10
3.2 Issues	14
3.3 Conclusions	16
4 Conceptual overview of the method	21
5 Implementation of the method	24
5.1 Generalization	24
5.2 Relief shading	29
5.3 Rock hachuring	45
6 Conclusion	51
6.1 Problems solved and unsolved	51
6.2 Adapting the method	54
References	57
A Appendix	62
A.1 Program code listings	62
A.2 Map settings	71
A.3 Images	71
A.4 Acknowledgments	90

1 Introduction and problem statement

Cartographic rock depiction lies in the hands of experts, literally. That is to say, in addition to knowledge encompassing such domains as cartographic design, map perception, and geomorphology, well-trained illustration skills are also indispensable for putting the monochrome strokes of the Swiss style rock hachure on a map. It comes as no surprise then, that while digitalization and automation thrive in cartography, rock drawing is still a handcraft, one stroke at a time ([Gilgen and Jenny, 2010]).

The presentation of relief in Swiss maps is held in high regard internationally, thereby also transcending the boundaries of the cartographic community ([Jenny et al., 2010], [Desimini and Waldheim, 2016]). Within the body of Swiss maps, the *Swiss National Map*, published by the *Swiss Federal Office of Topography (swisstopo)*, is one of its foremost exponents, cited for its clarity, expressiveness, and readability ([Tufté, 1990]).




Figure 1: A sample of Swiss style rock depiction (rock layer of Swiss National Map 1:25'000, © swisstopo).

The requirement for highly-skilled and trained cartographers as well as the very high expenditure of time to create such maps makes the business of map production a costly venture ([Gilgen, 1998]). Into the mid 1990s the time needed to create a rock engraving of a single high mountain sheet in the Swiss National Map was allocated at two thousand to six thousand hours, limiting dedicated resources to local revisions rather than a complete map overhaul ([Hurni, 1995]). Since the year 2000, the rock engraving technique has been superseded at swisstopo, giving way to software tools in the course of digitalization. However, the expenditures of time and cost remain, as rock depiction is still done manually ([Gilgen and Jenny, 2010]).

The objective of this thesis is to enable the creation of Swiss style rock depictions by way of automation. Two notions of automation are differentiated herein, designated by the words *automated* and *automatic*. The adjective *automated* refers to the fact that the process is carried out on a machine, a computational process solving a cartographic problem. It means that the

cartographic expertise, i.e. knowledge and skills, need to be modeled and formalized as algorithms and data structures, and implemented in a programming language to be interpreted by the computer. Furthermore, it is required for the computational process to be *automatic*, meaning that it operates on *descriptive* user input alone. This is to be contrasted with a semi-automatic approach, where the user is required to provide *prescriptive* solutions to subproblems that are not covered by the implementation, implying that different users arrive at different solutions. For the process to qualify as automatic, it must not be required by the user to instruct the machine *how* to do things. Instead, it has to be instructed *what* to do. An automatic process is always assumed to be automated.

The intent of automation, of course, is to relieve the user of the burden of having to fend for himself due to the lack of time, money, or domain expertise. With an automatic method at her disposal, the expert benefits from a supporting tool, while the non-expert can be enabled to perform tasks hitherto unfeasible. The dissemination of the workflow by means of a description, or an executable description, i.e. a computer program, helps in promoting the understanding and the routinely application of the method by a larger public. Automation can also be a building block for a unified mode of representation, a standardized depiction of rock in a map, which is easily recognized and interpreted.

In the automatic method devised in this thesis, given a *raster elevation model* and a *rock mask*¹, the user is merely expected to provide global map settings, for instance a *direction of illumination*, to be able to derive a Swiss style rock depiction from the elevation model. Since the location of rock cannot be inferred given just the elevation model, the user is required to provide the coverage by means of the rock mask.




Figure 2: Problem statement: develop a method to automatically derive a rock depiction from an elevation raster, a rock mask, and a set of parameters describing the desired output.

The choice of the input data format affects the range of applicable algorithms, and vice versa. Raster elevation models, or elevation *grids* for short, regular arrays of cells of altitude values, were chosen as a base in favor of other representations, such as triangulated surfaces, because the algorithms, existing as well as newly devised, were most easily implemented using image processing methods running on regular grids. Moreover, at the time of writing, high resolution raster elevation models are readily available via official agencies, covering the whole of

¹A binary raster indicating the cells covered by rock in the elevation model.

Switzerland, free of charge in an academic setting. See figure 2 for a conceptual overview of the problem statement. The ellipsis is going to be filled with content in the following chapters.

The thesis is structured as follows. In section 2, the purpose, the constituent elements, as well as the provenance of the Swiss style rock depiction are explained. This is followed by a review of prior work on the automation of rock depiction in section 3, resulting in a set of problems which need to be addressed in the further development of the method. In section 4, a high-level overview of the proposed solutions is given, outlining the workflow in conceptual terms, while the aspects concerning the implementation of the solutions are described in section 5. The discussion of results and implications concludes the presentation of the method in section 6. In the appendix, program code listings of key procedures, a collection the most important input parameters, as well as a series of enlarged figures are provided.

As a major part of this thesis, a software was developed implementing the method described herein. This text should be read along with a running copy of the software. The author hopes that the novel concepts presented in the following sections contribute to the development of the automation of rock depiction and that they will be applied in the creation of useful and beautiful maps, equating beauty with utility.

2 The elements of Swiss style rock depiction

Swiss style rock drawings are rendered in stroke manner, as patterns of monochrome hachures. The fabric of lines informs the map reader about the shape, quality, and inclination of the terrain [Federal Office of Topography swisstopo, 1990]. Hachures lend themselves to the depiction of minute detail. As it is shown in figure 1, the shapes are highly stylized, with the cartographer striving to cast the scenery into iconic triangular, columnar, or scaly shapes. This generalizing stylization helps to improve the recognition of typical shapes, thereby enhancing the readability of the map. Together with readability and clarity, topographical accuracy of the depiction is likewise important ([Gilgen, 1998]). However, readability and accuracy are not equally feasible in every setting, most notably when mapping precipitous terrain. In general, rock formations are to be portrayed in side view from the perspective of a terrestrial observer, rather than an aerial observer, to increase the utility of the map for map readers on the ground. This rule is difficult to follow in the depiction of steeply inclined rock faces due to their poor footprint in an orthogonal map view. In this case, the footprint is exaggerated, if this is compatible with adjacent map elements, in order to be able to create a more detailed image of the precipitous face ([Gilgen, 2008]).

2.1 Strokes and hachures

At least three kinds of stroke patterns are discernible in the Swiss National Map, corresponding to three kinds of rock quality, i.e. craggy rock, loose rock, and karst formations ([Jenny et al., 2014]). In this thesis, the focus will be on the depiction of craggy rock, since it is the most common one to be found in examples of the Swiss style, as well as the most comprehensive in terms of its constituent elements. Nevertheless, there is going to be an example of a depiction of loose rock in section 6.2.




Figure 3: **Left:** Enlarged rock hachures from Swiss National Map 1:25'000 rock layer (© swisstopo). **Center:** *Kanizsa triangle*, a subjective shape (after [Kanizsa, 1976]). **Right:** Relief inversion depending on the direction of illumination (from [Zwimpfer, 2001], © Verlag Niggli AG).

Apart from the quality of rock, exposition and slope are important properties to be depicted in a

map. *Slope hachures* are strokes in the direction of steepest descent that vary from long and thin where the slope is slight to short and broad where it is steep. An early example is the system of Lehmann in 1799 ([Hake et al., 2002]). *Shadow hachures* also follow the lines of slope, but the stroke width is varied according to light incidence under oblique lighting. In the Swiss style rock hachure, the effects of terrain slope and oblique light incidence are visualized in unison. In order to turn the hachures into graphic indicators of complex shapes of rock surfaces, they are modified subject to the illumination of an assumed light source, covering the surface with parallel rays, dividing the terrain into tones between light and shadow. The modifications include a change in stroke width as well as a change in stroke density, i.e. the number of parallel strokes per distance unit. Hachures on slopes directly facing the light source are drawn thinnest, whereas the stroke width is increased towards the areas located on the shadow side. The vertical inclination of the light source, usually fixed at an angle of forty-five degrees, results in the inclusion of the slope hachuring principle, although the strokes are not restricted to run in the slope direction. The light source is commonly positioned to illuminate the map from the top and left, i.e. from the north-west for northerly aligned maps. This preference of orientation is largely a habitual bias ([Imhof, 1965]), albeit failure to observe it may lead to relief inversion, the interpretation of hills as valleys and vice versa (see figure 3, right-hand side).

A second kind of modification is based on the concept of *aerial perspective*, i.e. the atmospheric effect that lets distant objects appear hazy, unsaturated, while near objects are sharply delineated and vividly colored. Applied to rock hachures, this principle is translated into a hypsometric effect. That is, contrasts are sharpened towards ridges, i.e. making strokes on the light facing side of a ridge even thinner while broadening the strokes on the adjacent side. Towards valleys, contrasts are diminished, resulting in the illusion of depth in analogy to the atmospheric effect.

Taken together, the differences in stroke width and density due to the modulation by illumination and aerial perspective translate into the perception of changes in reflectance, a powerful visual cue evoking the impression of three-dimensionality.



Figure 4: Stroke types used in rock hachuring. Form strokes and contour strokes make up the rock skeleton. Adding fill hachures completes the rock hachure (after [Gilgen, 1998], images from the Swiss National Map 1:10'000 ([Federal Office of Topography swisstopo, 2016]), © swisstopo).

There are three kinds of hachures present in a Swiss style rock depiction, which are functionally different. Refer to figure 4 for examples. *Contour strokes* are used to depict *skeletal lines*, i.e. sharp edges in the terrain, either positive, convex, like ridges and crests, or negative, concave, like ravines, gullies, or trenches. *Fill hachures* cover surfaces. *Form strokes*, broadly drawn hachures, are hybrids in that they both indicate edges and faces. A single broad form stroke is commonly used to depict the side of a trench facing away from the light, while the adjacent bright side is left blank. This juxtaposition of bright and dark creates a sharp contrast, incisively

suggesting a concave erosional structure. Referring to a well-known analogous visual effect, the practicability of using blank space as a modeling aid in creating the illusion of shape is illustrated in the central illustration of figure 3. By carefully assembling three indented circles on a white background, an observer interprets the scenery to be made up of four objects, including a triangle on top.

The form strokes and contour strokes are combined into the *rock skeleton*, outlining the overall shape of the rock formations and structuring its dominant features. The peripheral strokes of the rock skeleton delimit the rock to other parts of the terrain like scree, glaciers, et cetera. The rock skeleton is a style of rock depiction in its own right ([Imhof, 1965]), commonly used to depict rocky terrain on small scale maps, where space is too scarce to insert fill hachures.

Fill hachures are added to the rock skeleton in order to utilize the remaining blank areas between form and contour strokes to aid in the depiction of the local illumination conditions, enhancing the perception of the rock formations as three-dimensional objects. They are drawn using oriented, non-overlapping strokes with a uniform spacing for a given area. The manual for cartography apprentices by the Swiss Federal Office of Topography suggests placing seven strokes per two millimeters on average, leaning towards seven to eight strokes on the shadow side, and six to seven strokes on the light side ([Federal Office of Topography swisstopo, 1990]). Imhof ([Imhof, 1965]) and Spiess ([Spiess, 2014]) recommend to employ constant density, e.g. three strokes per millimeter for a 1:25'000 map, confining the modulation to a change in stroke width only. Fill hachures are usually terminated short of touching form strokes or contour strokes to enhance contrast, and thus the illusion of shape.

In addition to being a carrier of the illumination conditions, fill hachures are used to depict the degree of slope, being drawn vertically, i.e. along slope lines, to indicate steeply inclined faces, and horizontally, perpendicular to slope lines, for gently sloping terrain. On the light side, vertical fill hachures are tapered towards the ridges, fully connected via horizontal segments at the base, but only partially connected at the top. This scheme is reverted for vertical fill hachures on the shadow side, with the exception that hachures at the top need not be fully connected ([Federal Office of Topography swisstopo, 1990]).

There are cases when the depiction of sharp edges by solid contour strokes is omitted if the fill hachures of adjacent faces can be utilized to indicate the course of the edge. Refer to the leftmost image of figure 3 for examples like the rocky vertical pillar on the right-hand side or the formation directly above it. The effectiveness of the depiction is an impressive display of the gestalt laws of perception at work in this image ([Eysenck and Keane, 2000]). The *proximity* and directional *similarity* of the fill hachures indicate the different faces and their orientation, while the *continuity* of the ridge line is hinted at by the continuous sharp bend in the hachures. The artful omission of explicit lines incurs no loss in this image, since they are complemented by the viewer as a consequence of the perceptual process.

All kinds of hachures are to be drawn with a jittered pen to suggest the roughness and ruggedness of rock, as well as to set the strokes apart from isolines ([Federal Office of Topography swisstopo, 1990]). The compound hachure built from light-modulated form strokes, contour strokes, and fill hachures is the Swiss style *rock hachure* proper.

2.2 Historical background




Figure 5: **Left:** Enlarged, right-flipped extract of the Roverea and Grouner map of the Aigle district showing banded rock cliffs of the Diablerets mountain (from [Kraiszl, 1930]). **Right:** An enlarged sample of the Atlas Suisse by Meyer showing the Diablerets.

With the defining elements in mind, a brief retrospective of Swiss map products illustrates the provenance of the method. Map historiographer Willi Kraiszl considers the map of the Aigle district by *Roverea and Grouner*, published in 1788, to mark the transition from oblique to orthogonal map projection ([Kraiszl, 1930]). Rock formations are represented by short hachures in the direction of slope lines forming darkened horizontal bands. An extract of the map is shown in figure 5. In the *Atlas Suisse* by Johann Rudolf Meyer, published from 1796 to 1802, orthogonal projection is fully adapted ([Gilgen, 1998]). In this atlas, rock is depicted by darkly hachured slabs irregularly outlined with broad strokes or blank spaces, indicating structural lines or illuminated faces. Some of the slabs are covered by strokes with gradually changing widths hinting at curved surfaces by means of a lighting effect. The demarcation of rock against other parts of the terrain is not carried out consistently.




Figure 6: **Left:** Extract of the 1:100'000 scale Dufour map showing the Gross Ruchen mountain, in the edition of 1859. **Right:** The Gross Ruchen in the 1:50'000 scale Siegfried map, 1880 edition (both from [Federal Office of Topography swisstopo, 2016], © swisstopo).

Widespread adaption of the explicit depiction of rock guided by authoritative regulations begins with the publication of the 1:100'000 scale *Topographische Karte der Schweiz*, or *Dufour map*, starting in 1844, under the direction of the eponymous Guillaume-Henri Dufour. Enabled by

the first extensive topographic survey of high mountain terrain, the map features rock hachures modulated by oblique illumination ([Blumer, 1927], [Kraiszl, 1930]). Dufour advised his staff to aim at a vivid depiction free of mannerisms ([Spiess, 1970]). See figure 6 on the left-hand side for an example of the Dufour map.

In 1870 the first sheets of the *Topographischer Atlas der Schweiz*, or *Siegfried map*, named after Dufour's successor as head of the national mapping agency, Hermann Siegfried, were published. Contour lines replace the slope hachures in the display of topography. The specifications regarding the depiction of rock require to focus on recurrent typical forms and the complete depiction of skeletal lines. Assuming oblique illumination from the top left, rocks are prescribed to be hachured in a way as to portray their inclination and typical shapes like terraces or banded cliffs. The footprint of precipitous rock faces is to be exaggerated ([Lochmann, 1888]). Highly renowned cartographers like Held, Imfeld, Becker, Leuzinger, and Jacot-Guillarmod left their mark in the Siegfried map, adding elements such as aerial perspective or karst formation signatures to the depiction of rock ([Kraiszl, 1930]).

With the advent of photogrammetry around 1900, topographical surveying could be conducted with hitherto unavailable precision. The Swiss engineer Walter Blumer argued in favor of a mechanistic, geometrically correct, plane view rock depiction, rejecting the “artistic” ([Blumer, 1927]) style of the Dufour and Siegfried maps as imprecise, although conceding their “masterful” execution. In the works of Blumer, steep rocks are depicted by short vertical hachures that are drawn slightly broader on the shadow side under oblique illumination. Flat areas are covered with isolines interspersed with parallel, gap filling hachures. Leonhard Brandstätter adopted the method of Blumer in maps of the *Österreichischer Alpenverein* (Austrian Alpine Club), explicitly avoiding the footprint exaggeration of vertical faces, judging it detrimental to map understanding ([Brandstätter, 1968]).




Figure 7: **Left:** Extract of the Swiss National Map 1:25'000 showing the Gross Ruchen mountain, in the edition of 1963. **Right:** The same area in the 2016 1:10'000 revision (both from [Federal Office of Topography swisstopo, 2016], © swisstopo).

In the course of a legislation in 1935 for a renewal of the national map, comprising greater accuracy, increased detailing, as well as enhanced clarity and homogeneity in comparison to the Siegfried map ([Schneider, 1946]), rock depiction was further developed and standardized by the Swiss Federal Office of Topography ([Hurni and Neumann, 1999]). Publication of the new Swiss National Map commenced in 1938. Its rock depiction encompasses all the elements

discussed in this chapter. See figure 7 for examples. A notable example of the application of the Swiss style of rock depiction outside of institutionalized Swiss cartography is the 1:50'000 Mount Everest map issued by the National Geographic magazine in 1988, and described in [Washburn, 1988].

3 A review of prior work on the automation of rock depiction

In this section a survey is given of publications on the automation of rock depiction, the critique of which results in a set of open problems constituting the motivational basis for the development of the novel method. The prior work is scrutinized in particular regarding to its handling of the modulation of rock hachures in terms of the illumination model, to the degree of automation, as well as to the implementation of the generalization of rock formations.

3.1 Survey




Figure 8: Aspect based stroke width modulation in [Hurni, 1995]. **From left to right:** mean stroke width for light and shadow side, azimuth sectors, stroke width modulation for the light side, as well as for the shadow side.

Hurni developed a method ([Hurni, 1995]) for generating depictions of rocky cliffs in rock skeleton manner, where the cliffs are made up of top and bottom demarcating lines representing the head and foot of the cliff, as well as form strokes and contour strokes. Since it is a rock skeleton depiction, there are no fill hachures in this model. The top and bottom demarcations are given as polylines with equal numbers of vertices, vertical form strokes are drawn between each corresponding pair of vertices. Each line segment is turned into a polygonal shape whose wedged form is determined by an illumination model. In this model, stroke segments are classified according to the azimuthal exposition attribute of their vertices to fall into one of several discrete sectors of light incidence. The classification determines both the overall stroke width, as well as further stroke modulations depending on whether the segment resides on the light or shadow side (see figure 8). Moreover, a random process is used for varying the stroke width and to displace vertex positions within bounded limits to give the appearance of rough, jittered, hand-drawn hachures. The strokes may be bent to indicate terrain concavity or convexity. The generalization according to map scale is done manually by the map editor by selecting appropriate delineation lines and numbers of vertices. For an example of this kind of rock depiction refer to figure 9.

Originating in Hurni's illumination model, Dahinden developed a software for covering poly-




Figure 9: **Left:** Extract of the 1:25'000 Methana peninsula topographical map (© 1995 *Stiftung Vulkaninstitut Immanuel Friedlaender and Institute of Cartography, ETH Zürich*) featuring rock hachures after [Hurni, 1995]. **Right:** Mount Rainier map in [Hurni et al., 2001b].

gons with tapered and jittered fill hachures ([Dahinden, 2000], [Hurni et al., 2001a], [Dahinden, 2008]). Given a polygon, the segment between its first two vertices determines the orientation of the hachures, which is either parallel to the segment or perpendicular to it, depending on user preferences. Additionally, the segment, extended to an infinite line, serves as a ruler from which, given a unit distance, hachures are spawned with uniform spacing between them. Gently sloping terrain is hachured along contour lines, while steeply inclined rock faces are hachured along slope lines. Again, the stroke widths are assigned according to the illumination model in [Hurni, 1995]. Figure 10 shows a sample of this kind of rock depiction. The time needed to produce the rendering was stated to be approximately two hours ([Dahinden, 2008]).



Figure 10: Rock depiction using fill hachures in [Dahinden, 2008].

Yang et al. ([Yang et al., 2009]) extend the approach taken by Hurni to include equidistant slope break lines running perpendicular to the form lines for drawing cliffs made up of regular arrays of cells, i.e. rock faces bounded by vertical form lines and horizontal break lines, thereby assuming the appearance of rocky pillars. Within each pillar, the break lines run from one side

of the pillar, i.e. starting from a form line, to the other side of the pillar, just stopping short of touching the adjacent form line. This discontinuity effectively creates the impression of a trench between adjacent pillars. Automation is extended by only having to provide a single polyline across a cliff, while the upper and lower demarcation lines of the cliff are automatically created from the vertices of all triangles in the triangulated mesh elevation model that are intersecting this polyline. In figure 11 an example is shown.




Figure 11: **Left:** Hachured cliff in [Yang et al., 2009]. **Right:** Hachured doline in [Yang et al., 2015].

In another paper, Yang et al. ([Yang et al., 2015]) apply the method described above to the visualization of dolines, or sinkholes. The doline walls are segregated into quadrilateral cells by intersecting vertical ridge lines and horizontal concentric break lines. The illumination model is based on [Hurni, 1995]. See figure 11 for an example.




Figure 12: **Left:** Hachures for steeply sloping rock in [Gondol et al., 2008]. **Right:** Rock texture pattern and stylized hachures in [Christophe et al., 2016].

In a project commissioned by the French National Mapping Agency, IGN France, for the production of new basemaps at 1:25'000 and 1:50'000 scales, Gondol et al. ([Gondol et al., 2008]) extended automation towards the semi-supervised detection of rock surfaces from a combination of orthoimages, slope, aspect and elevation data, as well as (smaller scale) land cover data. For

visualization purposes, each rock area is put into one of thirty-six classes according to the aspect and slope properties of the area. The class number is then used to look up a prebuilt hachure texture to be pasted into the map. Special care is taken to differentiate textures indicating steeply sloping terrain from those indicating gently sloping terrain. Additionally, the boundaries between rock and other forms of land cover like scree are emphasized using dark strokes. On the left-hand side in figure 12, a patch of steeply sloping rock rendered in this manner is displayed.

The texture approach is also utilized in [Christophe et al., 2016], where rock faces are filled with single trapezoidal strokes or chains of trapezoids on the shadow side, signifying slabs of rock. The ridges are enhanced using dark patches. The textures are generated in a two-step process. At first, a *pattern* of vector paths is procedurally generated, which is *stylized* and rasterized in turn to create the texture. In figure 12 the pattern is displayed on the right-hand side at the top, while the final texture is shown at the bottom, with a stylization comprising colorization and random displacements of the vector vertices in order to convey the impression of an irregular, hand-drawn style. The pattern generation parameters include orientation and density, dependent on terrain slope and exposition for the area filling trapezoids, and dependent on slope, exposition and elevation for the ridge-enhancing patches.




Figure 13: Rock hachures in [Grünwald and Kriz, 2015].

Recently Grünwald and Kriz published a report ([Grünwald and Kriz, 2015]) on a workflow to generate rock hachures from raster elevation models. To create Swiss style rock hachures, fill hachures are automatically generated from the raster by using slope lines for hatching steeply inclined rock faces, while contour lines are used otherwise. Ridges and rock boundaries are added manually. Stroke width and density of the hachures are modulated according to terrain aspect. Ridge strokes are paired with white strokes on the light side to enhance contrast and thus the perception of three-dimensionality, while gullies are left blank. Furthermore, the hachures are given an irregular, jittered appearance using image processing methods. An example of hachures created in this way can be seen in figure 13 on the right-hand side. Likewise, two alternative styles of rock depiction are implemented in this workflow, an approach using vertical strokes modulated in width and density according to aspect, and a method using contour lines as horizontal hachures, exclusively, but switching to vertical hachures for steep faces.

To conclude the overview of the current state of research, a recent example of rock hachuring using plain lines which are not modulated by an illumination model is described. Lysák developed this automatic method ([Lysák, 2016]) to draw banded, columnar to tooth-like, stylized hachures inspired by the large scale cliff representation used in maps of the Land Survey Office of the Czech Republic (see figure 14). The hachures come in three variants as *upper symbols* representing the cliff head, *lower lines* along lines of slope demarcating the cliff columns, and *lower strokes* perpendicular to the lower lines serving as fill hachures and indicating cracks and crevices. Stroke vertices are displaced using a random process to suggest the irregularity of a hand-drawn sketch.




Figure 14: Large scale rock depiction after the method in [Lysák, 2016] (from [Czech Office for Surveying, Mapping and Cadastre, 2016], © Czech Office for Surveying, Mapping and Cadastre).

3.2 Issues

The publications mentioned above focus on the creation of hachure patterns to explore the use, utility and feasibility of automating the creation of rock depictions. Where hachures are modified in width or density in order to become shadow rock hachures², they are done so in a regular fashion as a function of local aspect, or aspect and slope, respectively. More specifically, the function is one of the light incidence of an assumed point light source sending parallel rays from north-west at an elevation of 45 degrees, which can be expressed in terms of aspect and slope as the ubiquitous *Lambertian hillshading* function ([Horn, 1981])

$H : (a, s) \rightarrow \max(\frac{\sqrt{2}}{2} (\cos(s) + \sin(s) \cos(\frac{3\pi}{4} - a)), 0)$, where a and s denote some values of aspect and slope³.

All the modes of hachure width and density modulation in these publications can be interpreted in terms of a function of H . In [Hurni, 1995], for example, s is set to a constant (i.e. it is not considered) and a is discretized into ten values, while in [Gondol et al., 2008] s and a are discretized to yield $|a||s| = 36$, et cetera. While this aspect-slope based lookup approach is easy to understand and to implement, it has some undesirable ramifications when it is applied in a general setting of mapping an arbitrary terrain at an arbitrary scale.

²I.e. in all of the publications mentioned above except for [Lysák, 2016].

³The general Lambertian hillshading function would be $H_{\alpha, \epsilon} : (a, s) \rightarrow \max(\cos(\frac{\pi}{2} - \epsilon) \cos(s) + \sin(\frac{\pi}{2} - \epsilon) \sin(s) \cos(\alpha - a), 0)$, for some constants α and ϵ denoting the azimuth and elevation of the light source, respectively.

Consider the interaction with *shaded relief*, for example, which is commonly combined with rock hachures, especially since it is much more eligible to depict the shape of terrain in place of contour lines ([Imhof, 1965]). Undoubtedly, shaded relief and hachures have to play well together, e.g. there must not be bold, dense hachures where the shaded relief conveys the impression of a brightly illuminated rock face. As Imhof remarks: “Die schattenplastische Felszeichnung aber ist bereits ein Element der allgemeinen kartographischen Schattenplastik. Die Lichtrichtungen für die Felsdarstellung und für die Schattentöne müssen miteinander übereinstimmen.”⁴ ([Imhof, 1965]). Unfortunately, however, aspect and slope alone are not sufficient indicators for a cartographically adequate shading. A few counter examples suffice to illustrate the pitfalls of naive hillshading.




Figure 15: **Left:** Hillshaded relief of an artificial wheel-like mountain structure. **Right:** Shading of a curved ridge, illustrating the inflection of the contrast direction (both after [Imhof, 1965]).

The left-hand side of figure 15 shows an artificial wheel-like mountain structure illuminated using the hillshading function H . It is clearly visible that the gray value contrast across the ridges parallel to the lighting direction is missing, mainly on the north-westerly and south-easterly spokes, but also at the tip of each spoke lying in a cardinal direction. This effect is to be avoided in cartographic relief shading, because it makes important features of the terrain indiscernible and leads to *expressionless* ([Imhof, 1965]) maps. Usually the flaw can be remedied by rendering the affected features under the illumination of a light source with a slightly deflected azimuth to make sure that contrast is maintained across all significant skeletal lines of the terrain at hand ([Imhof, 1965]). This procedure, however, poses a major deviation from the simple illumination model given by the hillshading function, particularly in terms of an algorithmic implementation.

Given a point of inflection of the contrast on a curved ridge, i.e. a change in sign of the gray value gradient, a *checkerboard* effect occurs, leading to the visual break-up of the incident faces of the ridge into a chiasmus of crossed bright and dark patches, as well as resulting in a zone of insufficient contrast around the inflection point (see figure 15, right-hand side). Imhof suggests ([Imhof, 1965]) to keep the checkerboard pattern, but to take advantage of minor terrain structures at either side of the inflection point to ameliorate the consequences of lacking contrast.

⁴“The three-dimensional shading of rocky regions is already, however, a part of three-dimensional shading in cartography as a whole. The direction of light used in the depiction of rocky areas must comply with the general shading tones.” ([Imhof, 2007])

Another option would be to make sure the contrast gradient is maintained consistently across the ridge by deflecting the lighting direction as described above.




Figure 16: Percentage of aspect blacked out for slopes above 45 degrees.

Another issue concerns the effect of slope in the hillshading functional term. Imagine a rotationally symmetric mountain containing slopes from 45 up to 90 degrees. In figure 16, the degree of slope is plotted against the percentage of aspect degrees for which the hillshading function yields zero light incidence. Above 45 degrees of slope, the percentage monotonously rises up to 50 percent for vertical precipices. That means, as slope increases, terrain facing away from the main lighting direction is more likely to yield zero reflectance when using the hillshading function, making it impossible to discern any detail.

3.3 Conclusions

These counter examples strongly suggest to opt against an illumination model based on naive hillshading as a means to derive shadow rock hachures, since it must fail at large to comply with the rules of cartographic relief shading.

Instead, the obvious solution would be to *equate* the hachure illumination model with the illumination model used in the creation of the shaded relief, provided that the shaded relief is cartographically sound. This is going to be the approach realized in this thesis. More specifically, the shaded relief *itself* will be the (reified) illumination model from which rock hachures are to be derived. It will be shown that, if there was a recipe to make a cartographically sound and pleasing shaded relief, a satisfying Swiss style rock depiction would not lie far behind.

The rationale for choosing shaded relief as the basis for rock hachures hinges on two arguments. As stated above, rock hachures don't exist by themselves on a map, but are commonly combined with shaded relief, especially in high mountain terrain. Thus, shaded relief is a requirement, anyway. More importantly, from a problem solving perspective, shaded relief has a simpler structure than hachures, a regular array of gray values compared to patches of jittered strokes.

Since in this thesis shaded relief is considered to be fundamental, even constitutive to rock hachures, its theory is going to be summarized, in consistence with the principles already outlined above in the discussion of rock hachuring theory in section 2. The summary follows Imhof's exposition of manual *oblique shading* ([Imhof, 1965]).

Starting off with fixing the rock skeleton lines, a direction ideally perpendicular to the major lines for an assumed diffuse light source is determined, usually lying towards north-west, west, or south-west. Where lines traverse the main lighting direction at an acute angle or where they are parallel to it, local changes in the direction of illumination should be introduced for the reasons discussed above. Every flat face is uniformly covered with a medium tone of the grayscale palette, as a reference tone to the remaining faces. Beginning with rendering the steeply inclined areas on the shadow side in the darkest tones, adjacent faces are subsequently shaded moving towards the light side. The slopes directly exposed to the light are left blank to appear the brightest. During the whole process, unfavorable illumination conditions, i.e. ones leading to missing contrast, are to be mended by local brightening or darkening. An area with constant slope and aspect should be filled with a uniform tone, except for a change in tone at its boundaries to enhance contrast towards adjacent areas. Light effects based on physics like specular highlights, reflection, or shadows are disregarded. However, an aerial perspective effect is introduced by sharpening the contrasts along major ridges and diminishing contrast towards valleys, thus enhancing visual cues for the perception of a three-dimensional form. In a key statement, Imhof writes “Besonders sorgfältig bearbeite man alle Schattierungsübergänge. *Die Form liegt an den Übergängen*”⁵. That is, the expression of form in shaded relief is essentially accomplished by the skillful design of the appearance of the discontinuities in the terrain. Sharp edges require sharp contrasts, while smooth gray value gradients are reserved for gentle transitions.

This principle – *form follows from discontinuities* – is regarded as fundamental in the process of creating a shaded relief in this thesis. A little thought experiment should help in the understanding of the rationale. Let A be the proposition *If the world was flat, there would be no relief* or, factually speaking *There are no discontinuities on a surface \rightarrow there is no relief on that surface*. The proposition is assumed to hold. By modus tollens one is able to assert that *There is a relief on a surface \rightarrow there are discontinuities on that surface*. The inverse proposition *If there was no relief, the world would be flat* is trivially taken for granted, which establishes, again via contraposition, together with A the necessity and sufficiency of surface discontinuities for the presence of relief. Taking literally the equivalence of discontinuities and relief, it is assumed that discontinuities are the loci of significance in a relief, and that in creating a shading of the relief the design of the contrasts across these edges is of primary importance. Moreover, it will be shown in section 5.2, that once the gray values at the edges have been fixed, the values of all remaining faces can be *functionally* derived from the tones at the incident edges.

This edge-based approach on relief shading is quite different from analytical, Lambertian hillshading regarding the notion of *locality*. In Lambertian hillshading each facet, be it a raster cell or a triangle in a three-dimensional mesh, is assigned a reflectance value based on the aspect and slope at that facet, i.e. the assignment is independent of the assignments of the neighbors of the facet. The shortcomings of analytical hillshading are rooted in this ignorance of one facet from another. In edge-based relief shading, however, each facet is by design a function of its environment.

In automatic relief shading, the concept of regarding terrain edges in the computation is not unknown. Kurt Brassel extended Pinhas Yoeli’s analytical hillshading method for raster elevation models ([Yoeli, 1965]) by introducing local adjustments of the angle of illumination ([Brassel,

⁵“Special care should be exercised with the transitional zones of shading tones. *The form depends on these transitions.*” ([Imhof, 2007])

1974]). Brassel’s approach requires a definition of the skeletal lines in the mapped terrain to be available, ridges and ravines, which are further divided into three categories with decreasing weights. For every grid cell a *direction coefficient* is calculated as a sum of orientation, weight, and distance of adjacent ridges and ravines in the four cardinal directions as well as the four diagonal directions, but only counting adjacent lines belonging to the highest category for the latter, in order to ensure that there are major lines included in the computation of the direction coefficient. This coefficient is used in the evaluation of a term including a user given maximal angle of deviation to yield a rotated direction of illumination which is substituted in the hillshading function for the original azimuthal lighting direction.




Figure 17: Relief shading with blurred diffusion curves from [Geisthövel, 2013].

Brassel’s approach incorporates the local environment of a facet in terms of neighboring structural lines, ameliorating the effects of naive hillshading, yet the computation of the reflectance value remains a function of local slope and aspect. However, these local properties do not provide enough structural information in order to fix the issues pertaining to the terrain edges, i.e. insufficient contrast as well as contrast inflection. By shifting the focus from facets to the edges themselves, Geisthövel presented a concept ([Geisthövel, 2013]) of relief shading based on *diffusion curves* ([Orzan et al., 2008]). In this paradigm, the modeling primitive is a two-dimensional spline curve representing a terrain edge. At either side of the curve, gray values are defined at arbitrary points along the spline. The gray values are linearly interpolated along the curve by arc length. Analogously, a set of linearly interpolated *blur values* is associated with either side. A smoothly shaded raster image I is calculated from the gray values associated with the spline curves by solving the Poisson equation $\Delta I = \nabla \cdot g$, where Δ is the Laplacian $\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$, $\nabla \cdot$ is the divergence operator $\frac{\partial}{\partial x} + \frac{\partial}{\partial y}$, and g is the gray value gradient, which is supposed to be nonzero across the curves. The equation is to be solved according to the boundary conditions $I(x, y) = v$, where (x, y) holds a gray value v associated with a curve. The resulting image I is blurred according to the blur values of each curve. With this focus on edges rather than faces it is easy to ensure that contrast is maintained throughout the image across important discontinuities in the terrain, thereby eliminating a serious deficiency of analytical hillshading. Furthermore, blur values can be utilized to implement the aerial perspective effect by diminishing contrasts across edges in valleys. While calculating the image given the curves is straightforward, the

challenge in using diffusion curves for shaded relief mainly lies in picking suitable gray values for each curve in the first place.

The idea of using diffusion curves for shaded relief was taken up by Marston and Jenny to render small scale reliefs ([Marston, 2014], [Marston and Jenny, 2015]). Utilizing hydrology algorithms, their method features automatic extraction of ridge lines and valley lines, as well as the identification of flat areas such as valley floors. The lines are vectorized to become diffusion curves. To determine the gray tones at the curves, aspect values are samples at each side of a curve and rotated using Brassel's method. The modulated aspect values are then mapped to gray values trailing the left and right of each curve. The shaded relief is produced by combining the smoothly shaded diffusion curve image with an analytically hillshaded relief using the overlay blend mode ([Porter and Duff, 1984]). Valley floors are masked out using a uniform medium gray tone.




Figure 18: **Left top:** Analytical hillshaded relief S . **Left bottom:** Diffusion curve image D generated from S . **Right:** Shaded relief as overlay of S and D with masked valley floors (all from [Marston and Jenny, 2015]).

In the relief shading algorithm presented in section 5.2, edges with associated gray values at either side are used as a graphical primitive. It features a novel method to select the gray tones of an edge where the maintenance of contrast across all edges is the primary objective, as well as a simple mechanism based on the Laplacian alone⁶ for deriving a smoothly shaded relief from these gray tones. The shaded relief is used in turn to modulate the width of form strokes, contour strokes, and fill hachures.

An important aspect to be included in this review is the issue of *generalization*. No map draft can do without generalization⁷. It is usually performed during multiple stages in the process of making a map. It is put into effect in relief shading both by lumping together uniform faces with a sweeping brush and etching out minute details at ridges or crests. It is most poignantly applied in Swiss style rock depiction in the stylization of forms. Surprisingly, given the importance of the concept, generalization is only mentioned explicitly in the context of rock depiction in [Hurni, 1995]. There it is performed by manually selecting suitable cliff demarcations and vertex counts. In the other publications, generalization is assumed implicit in certain actions like the selection of polygons to be filled with hachures ([Dahinden, 2000]), as well as by setting global parameters governing the scaling of width and density of hachures ([Yang et al., 2009], [Yang

⁶ $\Delta I = 0$, rather than the more complex Poisson equation.

⁷A notable exception would be the 1:1 scale map in Jorge Luis Borges' short story *Del rigor en la ciencia* (1946).

et al., 2015]), or introduced by means of a graphical revision step ([Grünwald and Kriz, 2015]). In [Gondol et al., 2008], generalization enters the workflow by means of the composition of the prebuilt hachure textures. Generalization is mainly achieved in this thesis by a modification of the elevation model.

To conclude this review, a set of pending problems that need to be addressed in the further development of the automation of rock depiction provides an outlook on the following sections. The foremost issue in automation is the removal of work which can not be effectively performed by the user due to limited resources or expertise. Thus, there must not be any interaction, pre- or postprocessing involved, beyond the capabilities of the user. The user is only required to specify the map parameters and check the results for compliance with the intent of the map. On the other hand, generalization, relief shading, as well as hachuring have to be done automatically. Having reduced the problem of generating modulated shadow rock hachures to what is essentially the problem of creating a shaded relief, serves as a guideline for the main task. Since the generalization step is regarded as fundamental, it needs to be realized in such a way that both the shaded relief and the rock depiction benefit from it.

4 Conceptual overview of the method

A useful dichotomy in science pertains to the distinction between *schema* and *instance* ([Wedekind et al., 2004], [Wedekind and Ortner, 2004]). An instance is a manifestation of a schema. In the Linnéan phylogenetic taxonomy for example, the species leopard (*Panthera pardus*) and lion (*Panthera leo*) are instances of the schema, or genus, *Panthera* ([Reece et al., 2010]). The genus *Panthera* is itself an instance of the schema *Felidae*, the family of cats. Thus, schemas and instances, or *concepts* and *implementations* as they are dubbed hereinafter, are organized in hierarchies, reversing roles at different levels of abstraction. Examples of high level concepts of interest in this thesis are processes like terrain generalization, relief shading, or rock hachuring. Implementations are concrete procedures executing these processes.

The dichotomy is useful because it helps in structuring problems, allowing to focus on concepts first and implementations later. The order of conduct is thereby predefined by the analogous distinction between asking *what is to be done* versus *how it is to be accomplished*. The primary concern in this thesis is to identify crucial concepts involved in the automation of rock depiction, to point out the problems associated with these concepts, as well as to provide solutions to these problems, both conceptually and with respect to implementations. The choice of concepts and implementations is motivated by the considerations in section 1 concerning the requirements of automation, by the central importance of shaded relief as a basis for rock hachuring as outlined in section 3, and by the adherence to software engineering principles of simplicity, interpretability, and robustness. Simplicity refers to the degree of complexity of the program structure. A simple program is likely to be easier to implement, to comprehend and to contain fewer bugs than a complex program. Interpretability relates to the capability of the user of predicting as well as understanding the result of running a program. Keeping the parameter count of a program small is one example of promoting interpretability, and simplicity, also. Heeding this rationale, a program without parameters would be deemed optimal, in particular with regards to automation. The property of robustness determines the runtime behavior of the program in the face of faulty input. A program being able to continue processing an elevation model containing sampling errors, like holes or spikes, for example, would be attributed as being robust. Making a program more robust expands the number of cases it can be applied, and enhances its usability by relieving the user of preprocessing the input data for validation purposes.

An implementation is likely to change but the concept it represents remains. While the implementations discussed in the following sections are mostly the simplest solution to a problem, considerations concerning the running time and memory consumption of a program often provide the impetus to change the implementation in order to be able to process data more efficiently. In this section a conceptual overview of the method is given, whereas implementations are discussed in the following sections.

The figure 19 illustrates the method by means of a data flow diagram. It is to be read from left to right and from top to bottom. The boxes represent static input or output data, while the circles denote dynamic processes. For every process there are some inputs and one output, as indicated by the arrows. The output of a processing stage is used as an input in later stages.

In the input column on the left-hand side there are three items, the gray boxes are rasters, the elevation model and the rock mask, while the white box designates the set of descriptive map




Figure 19: A conceptual overview of the method.

parameters. These map settings are used in every processing stage. Most of these settings pertain to the implementational aspects, and they are explained in the respective sections. There is also an overview of the parameters given in section A.2.

Generalization is an integral part of mapmaking, the adaption of the map content to the intent of the map and its scale, as well as to the level of detail of the input data. As Gilgen puts it in [Gilgen, 2008], the topographic abstraction used in rock depiction is in effect a caricature of shapes, thereby emphasizing important features, simplifying complicated forms, and discarding spurious detail. Its function is to render the landforms in such a way as to make their counterparts easily recognizable by a map reader on the ground. In order to incorporate this objective in the workflow presented herein, a *generalized elevation model* is derived from the input raster in the first processing stage. The accompanying map setting determines the degree of terrain abstraction, from low to high stylization. Having a separate, modified elevation model available means that the products derived from it, which are needed in later stages, are already generalized, all based on an equal footing. The generalization procedure, as implemented by the method described in section 5.1, yields an abstraction of the topography that closely mimics the manual generalization of surface formations as it is done by rock cartographers.

The creation of the shaded relief occupies the center of the figure, it is a pivotal element of the whole workflow, since the rock hachures are mainly derived from it. Therefore, the shading is performed on the generalized raster instead of the original one, in order to carry the generalization over to the rock hachures. As it has been stated in section 3, the shaded relief is modeled via the discontinuity edges in the terrain. The extraction and selection of these edges provides for an additional means of introducing generalization into the process. The map settings for relief shading include the main lighting azimuth and elevation, as well as the gray tone for valley bottoms and plains. This central stage is described in detail in section 5.2.

In the third processing stage the proper rock depiction is created. Its inputs are the rasters and

settings supplied by the user as well as the generalized elevation model and shaded relief generated in the previous stages. The rock fill hachures in gently sloping terrain are drawn at right angles to the gradient of the generalized raster, while the original elevation model is used to determine the actual slope values, since they are distorted in the smoothed grid. A threshold in the map settings indicates the slope value above which fill hachures are drawn vertically, i.e. in the direction of the gradient. The gray values of the shaded relief determine the stroke width of all fill hachures, from broad strokes for dark tones to thin strokes where the relief is bright. The rock skeleton strokes are derived from the edges in the shaded relief, the discontinuities with high contrast. The distinctive form strokes as well as the accompanying blank spaces are created by offsetting the edges corresponding to negative skeletal lines, i.e. ravines, against and along the direction of the edge gradient, using black and white as a pen color, respectively. The edge gradient magnitude determines the thickness of the form stroke, i.e. the higher the contrast the broader the strokes are drawn. Contour strokes for ridges are generated accordingly by offsetting the edges corresponding to positive skeletal lines. The rock mask is applied to confine the hachures to the areas covered by rock, resulting in the final output image. Since the rock depiction is essentially derived from the shaded relief, the two combine very well together, mutually intensifying the perception of shape in the map reader.

5 Implementation of the method

In this section, the conceptual triad of terrain generalization, relief shading, and rock hachuring is equipped with implementations, i.e. executable recipes, the combined effect of which results in the output of a Swiss style rock depiction. The order of presentation follows the three stages in the diagrammatic overview of figure 19.

5.1 Generalization

Generalization is not just employed to render map elements more general, but also to make certain elements more distinctive. In terrain generalization the emphasis of shapes is as important as the simplification and omission of features in the characterization of landforms ([Gilgen, 2008]). In rock depiction the exaggeration of the footprint of steep faces as described in section 2 is of special importance. In this section a generalization technique is described that is suitable for enhancing skeletal lines, being the indicators of the relief form, as well as for aggregating coherent slopes along slope lines. Moreover, it enables the footprint expansion of precipices.

5.1.1 Algorithm

The generalization of the elevation model is performed using *line integral convolution*, hereinafter abbreviated as *LIC*, a method originating in vector field visualization ([Cabral and Leedom, 1993]). The use of LIC for terrain abstraction is mentioned by Geisthövel and Hurni in [Geisthövel and Hurni, 2015]. In order to extend the exposition given there, the suitability as well as some ramifications of generalization based on LIC are discussed.




Figure 20: **Left:** Gradient vectors and contours on a hillshaded surface. **Right:** Cells affected in a single line integral convolution. Given an integration length of six and an origin (blue cell), the red cells indicate a walk along the positive direction of the vector field, while the green cells are covered by going in the negative direction.

In addition to the elevation grid Z , a vector field V of two-dimensional unit vectors defined on the grid, as well as an integration length l are given as inputs. The output is another grid S ,

where the cell at index $a = [i, j]$ is assigned to

$$\frac{Z_a + \sum_{b \in v_+} Z_b + \sum_{c \in v_-} Z_c}{2l + 1}.$$

This means that each cell in S is a weighted sum of values from Z . In the summation terms, v_+ and v_- denote two sets of indices into Z , v_+ is the set containing the indices of the cells starting at $[i, j]$ and going l steps along the directions indicated by the vector field V . Likewise, v_- contains all the cell indices starting at $[i, j]$, but proceeding in the reverse direction, with the signs of all vectors in V negated. Figure 20 illustrates the concept of going back and forth from a single cell by means of a vector field.




Figure 21: Hillshaded relief of the Mettelhorn in the canton of Valais (elevation model: swissALTI3D, 2 m cell size). **Left:** Original grid. **Right:** grid generalized using LIC. Integration length is set to 25.

If the normalized *gradient* of the elevation model is chosen as the guiding vector field, $V = \frac{(\frac{\partial Z}{\partial x}, \frac{\partial Z}{\partial y})}{|(\frac{\partial Z}{\partial x}, \frac{\partial Z}{\partial y})|}$, the integration is done along lines of slope. The gradient points in the direction of steepest ascent, if the slope is non-zero. Where the slope is nil, the magnitude of the gradient vanishes, and integration terminates. The normalized gradient is used in all examples of LIC generalization in this thesis. However, there is room for the use of alternative vector fields in terrain abstraction, for instance those generated by drainage algorithms used in hydrology, see [Geisthövel and Hurni, 2015] for examples. Figure 21 shows the hillshaded relief of an elevation raster and its generalization using LIC with the gradient as the guiding vector field. For a more vivid impression, take a look at the three-dimensional block diagram rendering of the grids in figure 60.

The integration length determines the amount of sampling of the elevation values. If it is set to zero, the output raster is equal to the input raster. For large values, heavy distortion of the topography incurs. See figure 57 for a progression of increasing integration lengths. The scale of the output map relates to the integration length in that smaller scales demand higher values,

i.e. greater amounts of generalization. Analogously, the cell width of the terrain raster affects the determination of the integration length, where smaller cell sizes require more integration steps for proper generalization. Thirdly, the value should be chosen depending on the quality of the input raster. Artifacts in the terrain model, e.g. due to the surveying method, may be removed by LIC, see figure 70 for an example. Along with the selection of the guiding vector field, the entire LIC generalization workflow is driven by this single integration length parameter, which keeps the procedure interpretable by the user.

Care needs to be taken when processing cells at the margin of the terrain model. The integration path may be truncated for cells within l steps of the model perimeter. When using LIC generalization on a set of adjacent tiles, for example, enough overlap should be allowed for, i.e. at least a strip of $2l$ cells in diameter, to make sure that the tiles can be merged without discontinuities. Additionally, the vector field tends to be more heterogeneous at the perimeter. For instance, when approximating the gradient with finite differences, the cells at the border need to be calculated using forward or backward differences instead of central differences, which yield more robust results ([Hengl and Reuter, 2009]). As with tiling, enlarging the perimeter provides a (local) solution in this case.

5.1.2 Discussion




Figure 22: Contour plots of the grids in figure 21 (**left**: original, **right**: smoothed).

In figure 22, the effects of LIC generalization are exemplified using a contour plot. At first glance, the lines in the generalized raster appear less curvy, more uniformly arranged, the spacing more homogeneous. In general, skeletal lines are accentuated in the generalized raster, turning arched contours into sharp angles, like the major ridge going from the central peak in a northerly direction, or the ravine at the center of the great catchment basin in the south-eastern quadrant, which is shaped like a leaf. Many details like minor peaks are lost, also. What is striking in particular is the lack of black patches in the generalized plot. It means that the inclination of

steep faces is reduced. While the maximum slope is lower in the generalized raster, the slope mean is increased in this example⁸. Most notably, due to the LIC smoothing, the precipitous faces at both sides of the south-westernly wedge have a larger footprint at the expense of the crest of the wedge and the lower slopes. However, this exaggeration of the footprint of steep faces is desirable, even encouraged, as a means to make room for a more detailed depiction of sheer cliffs in the side view perspective of rock depiction, as discussed in section 2.




Figure 23: **Left:** Difference of the original grid Z and generalized grid S mapped to a divergent color scale, superimposed with 20 meter equidistance contours of Z . The dark horizontal line is the trace of the profile on the right-hand side. **Right:** Profiles of Z (black) and S (red) at the top. At the bottom, each profile section is drawn separately to disentangle the lines. The vertical exaggeration is 1.5.

Taking the difference of the original and the generalized raster shows that there are gains as well as losses in height. In figure 23 on the left-hand side, red indicates a loss due to generalization, while blue indicates an increase in height. The footprint expansion of a steep cliff with the corresponding flattening of the surface is a consequence of both a lowering of the head of the cliff and a raising of its base. The combined effects of slope flattening as well as the emphasis of skeletal lines in the generalized elevation model are further illuminated by taking a look at the profile section in figure 23 on the right-hand side. Near the 200 meter mark, generalization has created a new, sharp ledge from which the cliff dips in an almost straight line, intersecting the original, steeper cliff about halfway. The peak in the profile at about 650 meters is a section of a ridge that is lowered in the generalized raster, but also more acute. Around the 725 meter mark, a ledge in the original terrain model has given way to a flat face in the generalized terrain. Apart from these regions, the profiles match closely.

In order to assess the suitability of the LIC method to recreate the generalization as applied by a cartographer, two manually created maps are superimposed with automatically derived 20 meter equidistance contours of both the original and the generalized elevation model. The first overlay in figure 24 is done using only the shaded relief layer of the Swiss National Map 1:25'000, while the second overlay in figure 25 is done using the regular Swiss National Map.

⁸The slope maximum (mean) is 84.4 (37.4) degrees for the original raster, and 80.5 (38.1) degrees for the generalized raster in the example of figure 21




Figure 24: Manually shaded relief of the Swiss National Map 1:25'000 (© swisstopo) superimposed with 20 m contours of the original (**left**) and generalized (**right**) elevation model.

For the shaded relief it can be said in general, that the contour lines show greater variation than the gray tones of the relief indicate, both for the original and generalized contours. However, the homogeneous spacing of the generalized contours corresponds much better with the smoothly shaded slopes, especially in precipitous areas. The sharp transitions of the relief at major ridges and ravines are commonly matched by sharply bent generalized contours, where the original contours appear more curvy and less stringent in depicting a discontinuity in the terrain.




Figure 25: Swiss National Map 1:10'000 (© swisstopo) superimposed with 20 m contours of the original (**left**) and generalized (**right**) elevation model (see enlarged version in figure 58).

The map in figure 25 contains mostly hachures, scree dots, and contour lines, conveying much more surface detail than the manually shaded relief. An enlarged version is shown in figure

58. The contour lines at 100 meter equidistance in rock areas and 20 meters everywhere else are barely visible, since they are matched quite well by the superimposed automatically derived contours of both the original and generalized elevation models. See figure 59 for an unimpaired view of the original map extract. In the rock areas, the derived contours of the original model are much more at odds with the fabric of the hachure strokes, while the generalized contours exhibit a high degree of consistency and alignment with the shapes of the rock formations, notwithstanding their unequal provenance of manually drawn strokes on the one hand in contrast to automatically generated lines on the other hand.

The large proportion of congruity among the manually drawn and the automatically derived contours, as well as the good alignment of the contours with the topography of both the shaded relief and the rock formations, indicate that LIC generalization is able to mimic the traits of the generalization performed by cartographers. Since the modified elevation model is used in the subsequent relief shading and hachuring steps, the topographical abstraction is passed on to the rock depiction.

As with most implementations discussed in this section, program code listings are available in the appendix A.1, where also an overview of the code syntax and recurrent programming constructs can be found. The exemplary implementation of line integral convolution is given in A.1.1. In this code, an additional step size ds is added as a further parameter to the procedure. It allows for the scaling of the step size to make sampling more robust in the face of errors introduced by the truncation of integration path coordinates to the integer coordinates of the grid. It is set to 0.5 for all examples in this thesis.

5.2 Relief shading

The purpose of creating a shaded relief at this point is to make the skeletal lines available, which are used in the final stage to derive the form strokes and contour strokes, as well as to be able to modulate the rock hachures in terms of the illumination model represented by the shaded relief. This approach begs the question why the skeletal lines should be provided by the shaded relief and not by an analysis of the elevation model, for instance. The equation of the skeletal lines with the discontinuities of the shaded relief is convenient, because it ensures the harmonic interplay of the rock depiction and the shaded relief when the two are combined. Moreover, the edges in the shaded relief can be easily extracted using standard image processing procedures. Thus, rather than applying one of the multitudinous methods based on differential calculus ([Eberly, 1996]) or topographic analysis ([Hengl and Reuter, 2009]) to extract skeletal lines, an edge detection method for gray value images is used instead.

As it is argued in section 3, the discontinuities of the terrain are the salient features, the formative elements of the relief. Therefore, when creating a shading of the relief, it is sufficient to focus on the design of gray value contrasts at the discontinuities, and let the remaining faces be interpolated by a suitable procedure. Making these edges the graphical primitives allows for the minute control of the gray values and the avoidance of the problems of low contrast or contrast inversion at the edges.

With these considerations in mind, the recipe for creating the shaded relief is set as follows:

- Determine the relief edges (skeletal lines)
- Choose the gray values at the edges
- Interpolate the shaded relief from the gray values of the edges

Given the intention to derive the skeletal lines from the shaded relief makes the very first line of the recipe seem like a dead end, since the shaded relief from which the edges could be extracted needs yet to be done. This seemingly contradictory puzzle is solved by bootstrapping the process, deriving the edges from a first approximation of the shaded relief, and then in turn using these edges in the creation of the proper relief.



Figure 26: Lambertian hillshaded relief, incident light indicated by arrow (**left, center**). **Right:** Blending of the two hillshaded images using Poisson editing. See also the enlarged rendering in figure 61.

As a prerequisite to understanding the steps involved in the making of the shaded relief, the image processing technique called *Poisson editing* needs to be introduced ([Pérez et al., 2003], [McCann and Pollard, 2008]). Poisson editing involves solving a Poisson equation $\Delta y = f$ for y given a function f , hence the name. It emphasizes the processing of an image in terms of its gradient, rather than its absolute values. The basic recipe is, given an image x , to modify it in the gradient-domain, and use the Poisson equation to calculate the image x' that corresponds to the modified gradient. Poisson editing enables the introduction of seamless local changes in contrast into an image while retaining the overall appearance of the original image. This property is taken advantage of in the remainder of this section by adjusting the gradients, i.e. gray value contrasts, of shaded reliefs in favor of a cartographically sound rendering. By varying the right-hand side of the equation, different effects can be implemented. But hereinafter, only a single variant is used, as explained below.

The walk-through of an example illustrates the procedure. Given the Lambertian hillshaded relief s of an elevation model with standard parameters, i.e. the light source being located in the north-west at an angle of 45 degrees, one way of enhancing the areas of low contrast is to locally illuminate the terrain using a different azimuth angle. The maximum gain in contrast is expected to occur at an azimuthal direction perpendicular to the standard direction, i.e. either oriented towards south-west or north-east. In this example, a Lambertian hillshaded image t , illuminated from the south-west, is chosen as an implementation of the rotated light source. To incorporate the local changes in illumination into the original image s , the right-hand side f of




Figure 27: Blending of two hillshaded images using Poisson editing. Enlarged extracts of the images in figure 26. **Left** column: s , **middle** column: t . Results are shown in the **right** column.

the Poisson equation is discretized as

$$f[i, j] := \sum_{[p, q] \in N_{i, j}} \begin{cases} s[p, q] - s[i, j] & \text{if } |s[p, q] - s[i, j]| > |t[p, q] - t[i, j]| \\ t[p, q] - t[i, j] & \text{otherwise} \end{cases}$$

where $N_{i, j}$ denotes the 4-neighborhood of the cell at integer coordinates $[i, j]$ on the two-dimensional grid, i.e. $N_{i, j} = \{[i + 1, j], [i - 1, j], [i, j + 1], [i, j - 1]\}$. That is, every cell in f contains the sum of the differences of the 4-neighborhood of s or t , whichever is greater in absolute value. The image f corresponds to the gradient modification step of the basic Poisson editing recipe outlined above. To solve the discretized $\Delta y = f$ for y , an iterative solution like the Gauss-Seidel method can be employed ([Press et al., 2007]). See the code listing A.1.2 for details.

The solution y contains an image where the highest contrasts from both s and t are seamlessly included. See figure 26 for the result, y is the image on the right-hand side. In figure 27 two enlarged extracts are shown, displaying more detail. As it can be seen in the second row of that figure, there is also the reversal of contrast across edges, i.e. from bright to dark in the original hillshaded relief on the left-hand side to a transition from dark to bright in the Poisson editing output y on the right-hand side. Note that in all examples, the values of y are clipped to $[0, 1]$, the gray value range of a shaded relief.

5.2.1 Determining skeletal lines

With this image editing machinery at hand, henceforth referred to as *Poisson-blending*, a first approximation of the shaded relief is created, from which the skeletal lines are derived. An




Figure 28: Determining skeletal lines from shadings. A Lambertian hillshade h with azimuth α is Poisson-blended with the sidelight $Sl(\alpha)$ to yield the enhanced hillshade s . Canny-edges are extracted and simple paths are determined from the edges. For the sidelight $Sl(\alpha + \frac{\pi}{2})$ edges are extracted and paths are determined, accordingly. Both sets of paths are merged using the straightness measure to yield the final set of skeletal lines.

overview of this process is shown in figure 28. To be more specific, the lines are going to be extracted from two images, a contrast enhanced Lambertian hillshaded image s and an image t with high contrasts in the direction perpendicular to the lighting direction of s to complete the set of lines missing in s . The two sets of lines are merged in a subsequent step to yield a unified set of skeletal lines.

The first image, s , is the blend of a Lambertian hillshade h with an illumination azimuth of α , and an image whose pixels are the dot product of the surface normal projected onto the base plane (i.e. $z = 0$) and the unit vector pointing in the direction of α , i.e. $N_x \cos \alpha + N_y \sin \alpha$, where $N = (N_x, N_y)$ is the unit projected surface normal. This image, mapped to the interval $[0, 1]$ where zero and one denote lowest and highest intensities, respectively, is also a hillshade, albeit not a Lambertian hillshade. It is brightest where N and the lighting direction coincide, i.e. where the dot product equals one, and darkest where N is facing in the opposite direction. Since this kind of shading is used several times in the following paragraphs, it is abbreviated as $Sl(\alpha)$, reminiscent of *sidelight*, i.e. $Sl(\alpha) := N_x \cos \alpha + N_y \sin \alpha$. Unlike Lambertian hillshading (with a fixed elevation), it satisfies the property $Sl(\alpha) = -Sl(\alpha \pm \pi)$, i.e. the absolute value is invariant under a rotation of the azimuth by 180 degrees. This property is essential in picking an unbiased perpendicular azimuthal lighting direction, as explained below. The sidelight features some contrasts lacking in h which are adopted in the Poisson-blend of h and $Sl(\alpha)$. The result s is almost a Lambertian hillshade, with some enhanced contrasts. Figure 29 shows h , $Sl(\alpha)$, and s .

The second image t needs to have high gradients perpendicular to the azimuthal direction of illumination α to indicate the relief edges missing in s . If a Lambertian hillshade with a per-




Figure 29: Blend using Poisson editing (**right**) of a Lambertian hillshade with illumination azimuth α (**left**) and the sidelight $Sl(\alpha)$ (**center**).

pendicular azimuth would be used, two images, with azimuths $\alpha + \frac{\pi}{2}$ and $\alpha - \frac{\pi}{2}$, respectively, would have to be generated in order to avoid a directional bias, since the sets of lines extracted from these would be distinct. Instead, $Sl(\alpha + \frac{\pi}{2})$ is used. There is no directional bias involved, since $Sl(\alpha + \frac{\pi}{2}) = -Sl(\alpha - \frac{\pi}{2})$, or, more verbosely, $N_x \cos(\alpha + \frac{\pi}{2}) + N_y \sin(\alpha + \frac{\pi}{2}) = -(N_x \cos(\alpha - \frac{\pi}{2}) + N_y \sin(\alpha - \frac{\pi}{2}))$. Thus, the sets of edges for $\alpha + \frac{\pi}{2}$ and $\alpha - \frac{\pi}{2}$ are equal when using the sidelight instead of a Lambertian hillshade. The image t is shown in figure 30 on the right-hand side.



Figure 30: **Left:** Lines extracted from hillshaded relief s (figure 29, right-hand side). **Center:** Lines extracted from t . **Right:** Sidelight $Sl(\alpha + \frac{\pi}{2})$, where α is the illumination azimuth of s . See an enlarged version in figure 63.

In order to extract the relief edges representing the skeletal lines from s and t , the Canny edge detector ([Sonka, 2008]), a standard image processing tool, is used. See figure 30 for the results of Canny edge detection in s and t . The sensitivity of the Canny detector is controlled in the map settings. The extracted edges are sets of connected pixels with multiple line ends and branches. See figure 31 for an example, where one edge is singled out in gray, its three ends being marked in red. It contains three *simple paths*, i.e. those with only two ends: one between the two red dots in the upper left corner, and one between each of the two and the red dot at the far right end.

To make sure that only one-pixel wide, simple paths are used in later steps, the Canny edges are preprocessed. At first, each edge is uniquely labeled using a standard algorithm to detect



Figure 31: Enlarged extract from Canny edge detection.

connected components in binary images ([Sonka, 2008]). The ends of each labeled edge are determined by extracting those pixels having at least six background pixels in a row in its 8-neighborhood. For every pair of line ends the shortest path lying on the edge is constructed. This is done using breadth first search ([Russell, 2010]), a method to trace paths between nodes in a graph. In this case, the nodes of the graph are the edge pixels and there is an arc between two nodes if pixels are adjacent in an 8-neighborhood. To construct the shortest path between a start and an end node, i.e. one containing the fewest nodes, the start node is put into a queue. While the queue is not empty, the node x at the head is dequeued. If it is the end node, the shortest path has been found and it is traced back to the start node. Otherwise, all its neighbors which have not been investigated are appended to the end of the queue and recorded as successors of x . See the code sample in A.1.3 for details.




Figure 32: Straightness measure for paths.

After the Canny edges from both images s and t have been split up into individual, simple paths, these are merged into a single set of paths representing the skeletal lines of the terrain. Since there are many pairs of paths lying on top or next to each other, they have to be filtered according to some predicate of path quality. The measure that is used here is the *straightness* of a path. The less bent a path is, the more likely it is to be selected in favor of nearby paths. This amounts to another generalization step next to the elevation model generalization.

Straightness is sampled at every pixel of a path by measuring the dihedral angle of two vectors extending from that pixel to pixels in either direction up to a given distance. See figure 32 for an example. Straightness is at its maximum if the angle is π . In the example figure, the dark gray path is straighter at the point of measure than the light gray path at its point of measure. In order to make the straightness measure more robust, a range of angle measures with a distance



Figure 33: **Left:** Canny edges from s (**top**) and t (**bottom**). **Right:** Final set of skeleton line paths from both s and t .

varying between a given minimal and maximal value is averaged for each pixel. For the merge operation, every path in s and t is examined and assigned a weight. The weight is calculated by iterating the pixels in a path and adding 1 if the straightness value of a pixel is higher than the straightness values of pixels of another path, both from s and t , in the 8-neighborhood. All paths are put into a list that is sorted by decreasing path weight. Iterating through the list, a path is chosen and put into the final set of skeleton lines if there is not already a path in its 8-neighborhood. It may be truncated or split up into several paths if it is only partially uninhibited by neighboring paths. The sorting order ensures that the straightest paths are selected first. Paths below a certain threshold length are discarded. In figure 33 the result of merging two edge sets into the final path set is illustrated.

5.2.2 Calculating the shaded relief

Having extracted the skeletal lines from the hillshades, one may turn towards the construction of the proper shaded relief. The crucial operation in this step is the selection of gray tones along and across the skeletal lines, from which the smoothly interpolated shading is finally computed. This selection process is guided by two rules. The first rule states that the gray value contrast across edges should be sufficiently distinctive. For every edge, one side of it should be discernibly lighter than the other. This requirement entails that there should be no points of contrast inflection on an edge, i.e. where the gray value gradient from one side to the other changes its sign. The second rule states that the gray values of an edge should be close to those of a given *enhanced* Lambertian hillshaded relief, as long as the first rule is not violated by this principle. This rule is used to incorporate the basic lighting model of both relief shading and rock depiction into the image, i.e. the tone variations by means of a global light source at a user supplied azimuth and elevation angle, while the intent of the first rule is to avoid the shortcomings of naive hillshading as explained in section 3. The adherence to the rules results in a set of constraints directing the rendering towards a cartographically adequate shaded relief. The relief shading workflow is depicted in figure 34.

By the *enhanced* shading, a Lambertian hillshaded relief $H_{\alpha,e}$ with a twofold modification is meant. The first one is the minor contrast boost gained by mixing in the sidelight $Sl(\alpha)$ by using Poisson editing as described above. The second modification is a gray value smoothing




Figure 34: Relief shading data flow. A Lambertian hillshade with azimuth α is Poisson-blended with the sidelight $Sl(\alpha)$ and smoothed using rotational LIC to yield an enhanced hillshade. In turn, this is blended with sidelights $Sl(\alpha \pm \frac{\pi}{2})$ and smoothed to yield images S^+ and S^- , respectively. The relief skeleton, made up of edges from S^+ and S^- , serves as the boundary condition for a Laplace equation, the solution of which is the shaded relief.

operation in order to approximate the gradually nuanced tone variations within homogeneous slopes in a manually shaded relief such as in figure 24. The smoothing also represents a further generalization step. It is achieved by applying a variant of line integral convolution, *rotational LIC* dubbed hereinafter, to the gray values of the hillshaded relief, rather than the altitude values of the elevation model. It differs from regular LIC in two ways. There is a user supplied number of iterations n , over which the integration results of each pass are averaged. After each iteration, the vector field is rotated by a constant angle. Secondly, there is no fixed integration length specified. Instead, integration terminates when a skeletal line is reached. The result of rotational LIC is a smooth shading of homogeneous faces bounded by skeletal lines. Figure 35 illustrates the basic concepts of rotational LIC.

More formally, the original gradient vector field is transformed by the rotation matrix

$$\begin{bmatrix} \cos i\theta & -\sin i\theta \\ \sin i\theta & \cos i\theta \end{bmatrix}$$

after the i th pass, where $\theta = \frac{\pi}{n}$ and i ranges from 1 to n . At the beginning of the last pass, the rotation angle is at a maximum value of $\pi \frac{n-1}{n}$. Within each pass, integration proceeds along the positive and negative directions indicated by the vector field just like in regular LIC, but only bounded by skeletal lines and the image perimeter. The greater the number n of iterations, the smoother the result is going to be, and the more generalized the slopes are going to appear. If n is too small, artifacts appear, depending on the distribution of the skeletal lines. See the progression of smoothed hillshaded reliefs with increasing iterations numbers in figure 65. An enlarged image is available in figure 64. This smoothing technique is akin to averaging images using anisotropic diffusion, where the diffusion terminates at high gradient magnitudes ([Perona and Malik, 1990]).




Figure 35: **Left:** Rotational LIC. The integration paths of four iterations for a single cell are shown. The white arrows indicate the positive vector field direction at the cell, while the light gray arrows indicate the negative direction. Iteration terminates at skeletal lines (black) or image boundaries. **Right:** Result of rotational LIC: a smoothed shading. Same extract as on the left-hand side.

In the selection of gray tones for the skeletal edges, the enhanced Lambertian shading S with an illumination azimuth and elevation angle of α and e , respectively, serves as the general reference image. However, the actual gray values are taken from another source, two Lambertian hillshades with the same illumination parameters as S , Poisson-blended with the sidelights $Sl(\alpha + \frac{\pi}{2})$ and $Sl(\alpha - \frac{\pi}{2})$, respectively. Both images, called S^+ and S^- hereinafter, are also smoothed using rotational LIC. They are displayed in figure 66. S^+ appears to be more in unison with S and the local contrast enhancements seem to be sound, while S^- looks awkward along some major ridges. This imbalance, of course, is due to the nature of this particular topography and the illumination parameters. For other terrains or illumination parameters, S^- would be preferred. It turns out that both S^+ and S^- are extensively congruent with S within large faces. The discrepancies are most severe at ridges and ravines, as it is to be expected. In figure 36 the differences of S to S^+ and S^- are mapped to a color range.

The two images S^+ and S^- represent the minimal unbiased selection from which to determine the skeleton edge gray values, since the original hillshaded relief $H_{\alpha,e}$ is included in both, and each contains a sidelight which is the negative mirror image of the other. This implies that Poisson-blending both S^+ and S^- in order to locally modify contrasts would be to no avail, since the equally strong gradients in $Sl(\alpha + \frac{\pi}{2})$ and $Sl(\alpha - \frac{\pi}{2})$ contained in S^+ and S^- , respectively, would annihilate each other. Instead, S^+ and S^- are going to be used as complementary palettes from which to paint the shaded relief, while S is going to be the reference model.

Prior to the assignment of gray tones to skeletal lines, the topology of these lines has to be determined. Up until now it is only known from which Canny edge a line is derived. In order to consistently assign gray values, the adjacency relations of all lines have to be identified. Line adjacency is established by sampling, at a regular interval, pixel coordinates on each line and submitting this set of coordinates to a Delaunay triangulation ([de Berg et al., 2008]). If coordinates of different lines are directly connected by an arc in the triangulation, the lines are defined as being adjacent. Coordinate sampling at pixel distances greater than one should be employed to avoid making the triangulation prohibitively large, i.e. not every pixel on a line

should be submitted to the triangulation.




Figure 36: Differences of shadings S , S^+ , and S^- . **Left:** $S - S^+$. **Right:** $S - S^-$.




Figure 37: A raster with a path (light gray) is superimposed with a sampling grid (broad strokes) with a cell diagonal of length 3, i.e. with a cell width and height of $\frac{3}{\sqrt{2}} \approx 2.11$. The dark gray pixels are admissible path samples given the sampling distance 3.

Thus, in order to sample the lines at a regular user supplied distance d , a sampling grid of square cells, each having a diagonal of length d , is superimposed onto the grid of the skeletal lines for fast checking of duplicate samples violating the minimal distance. See figure 37 for an example. To check if a pixel at coordinates $[i, j]$ is to be sampled, it is mapped to the respective sampling grid cell $[\lfloor \frac{i - (i \bmod c)}{c} \rfloor, \lfloor \frac{j - (j \bmod c)}{c} \rfloor]$ where c is the sampling grid cell width $\frac{d}{\sqrt{2}}$. If the cell is already occupied by another sample, the pixel can be discarded immediately, since it is known to be within distance d of the sample. Otherwise, if the cell is empty, its eight neighboring cells are examined. The pixel is accepted as a sample and recorded in the sampling grid, if $\sqrt{(i - p)^2 + (j - q)^2} > d$ for all neighboring samples with coordinates $[p, q]$. The sampling grid allows for constant time querying of the admissibility of a pixel as a sample, making the

entire sample selection process grow linearly in time with the number of path pixels an the image. For inspection of the accompanying program code listing, advance to section A.1.4. This sampling technique is adapted from [Jobard and Lefer, 1997], however, d is used for the length of the diagonal, instead of the width and height of the sample grid cells, as the authors do. The sampling grid method is utilized again in the drawing of evenly spaced fill hachures in section 5.3.




Figure 38: **Top:** Extract of Delaunay triangulation of sampled skeleton line paths. **Bottom:** Filtered Delaunay triangulation of sampled paths showing just the arcs between distinct paths and not crossing a third path. See an enlarged version with the actual paths in figure 68.




Figure 39: A subgraph of the skeleton line adjacency graph around a central vertex a with links to neighbors $b - f$. The skeleton line paths are underlaid in black, while the arcs of the filtered triangulation A are displayed in red.

Given the collection of skeleton line path samples, the Delaunay triangulation of the samples is computed. The arcs of the triangulation are filtered, such that only arcs between distinct paths remain, i.e. arcs within a single path are discarded. Additionally, arcs crossing a third path are excluded (see figure 38). The resulting triangulation subset is denoted by A , like *adjacency*, since the arcs in A define the adjacency relation on the skeleton lines. This relation is represented by an undirected graph $G := (V, L)$, where $V = \{v_1, v_2, \dots, v_n\}$ denotes the set of graph vertices and $L = \{l_1, l_2, \dots, l_m\}$ denotes the set of links between vertices. There is a vertex $v_i \in V$ for every skeleton line path, and a link $l_j = \{v_p, v_q\} \in L$ if there is an arc in A between

samples of paths denoted by v_p and v_q . See figure 39 for a close-up view.

This graph is used together with A to define an ordering relation $>$ on the skeletal lines interpreted as an indicator of the relative importance of these lines. A line x is deemed more important than a line y if $x > y$. Furthermore, the graph is used to query neighbors of a skeleton line during gray value assignment. The order $>$ is defined as follows. Given a link $l = \{v_i, v_j\}$, the *bound strength* of l , $\text{str}(l)$, is calculated as the sum of the reciprocal lengths of all arcs in A between the skeleton lines denoted by v_i and v_j . Adding over all arcs favors links with many arcs over links with few arcs, while taking the reciprocal of the length of an arc favors links with short arcs over those with long arcs, i.e. the close proximity of many paths increases the relative importance of a path. Based on the bound strength, the *weight* of a skeleton line designated by the vertex v is defined as $\text{wgt}(v) := \sum_{l \in L_v} \text{str}(l)$, where L_v denotes the set of links of v to its neighbors in G .

The weighting function establishes the ordering relation $>$ by $v_1 > v_2 \iff \text{wgt}(v_1) > \text{wgt}(v_2)$.

The order induced by $>$ on the skeleton lines is used to pick out the most important lines in the image, i.e. those having a major influence on the overall appearance of the final rendering, and to ensure that they are assigned gray values from either S^+ or S^- closest to the reference model shading S . Conversely, the less important lines according to $>$ may be assigned gray values at odds with S , but for the benefit of local contrast enhancement.




Figure 40: Expanded skeletal line for sampling gray values at either side.

In order to be able to compare the edges of different images for the selection of suitable gray values, the pixels making up the skeletal line paths are expanded at either side to form a buffer around the path. See figure 40. The buffer is made up of the coordinates of the original path P as well as the set of coordinates on the left side P_l and on the right side P_r . This *expanded path* is going to be filled with gray values of either S^+ or S^- for the final image interpolation.

The comparison of edge gray values in S^+ and S^- presupposes some predicates and functions. The *overall gradient* of an expanded path in an image is defined as the difference of the median of the gray values in P_l and the median of the gray values in P_r . A predicate is also defined to check for the equality of the sign of the overall gradient of an expanded path between two images. It is used to compare the general direction of contrast across an edge between two images. A third measure is the *similarity* of two images at an expanded path, which is defined as the sum of the absolute values of the differences of the images at every coordinate in the path. Lower values indicate higher similarity. The program code in section A.1.5 provides the details.

For every expanded path p a *preferred* assignment of gray values from either S^+ or S^- is deter-

mined. If the contrast of the reference model S at p is sufficient, the one of S^+ and S^- is chosen which is more similar to S , if it has enough contrast and the same gradient direction as S . The minimum contrast for this test is provided by the user. If there is not enough contrast in S at p , the one of S^+ and S^- with greater contrast is preferred.

Out of the multitudinous ways of assigning gray values to expanded paths, abbreviated just as *paths* hereinafter, the following scheme is suitable for producing a shaded relief that adheres to the reference model image S , except for local contrast improvements. This scheme is composed of a sequence of assignments with increasingly weaker preconditions at each step. In general, the assignment of gray values from the same source image is preferred for adjacent paths. At the end of the procedure, most of the paths of skeletal lines will be filled with gray values from either S^+ or S^- .

In the first pass all *major* paths are assigned their preferred gray values if the contrast of the path in S is sufficient. A major path is one which has no neighbor v in the graph G with a higher path weight $\text{wgt}(v)$ than itself. This pass ensures that the most influential lines in the image are close to the reference model.

In the second pass all unassigned paths are assigned the values from the same source S^+ or S^- as their neighboring paths, if all the neighbors have the same assignment. This pass locally spreads the assignments from the first pass, ensuring assignment of the same source image for adjacent paths. The iteration order is defined by the path importance relation $>$ defined above, i.e. more important paths are assigned first.

In the third round the remaining unassigned paths are processed, but this time the constraint that all neighbors need to have the same assignment is softened. The assignments of the neighbors v_i of a vertex v are weighted by $\frac{2\text{str}(\{v, v_i\})}{\text{wgt}(v) + \text{wgt}(v_i)}$, and the path represented by v is assigned the gray values of the source image with the highest weight. If a path has only unassigned neighbors, it is assigned the gray values from its preferred source. The iteration order is again determined by $>$, starting with the most important paths.

After three passes, *conflicting* assignments are common. These are assignments, where adjacent paths are assigned gray values from both S^+ and S^- . In the fourth pass, all conflicting assignments are nullified. This is done because conflicting assignments result in unusable shaded reliefs, where contrasts are either unnoticeable or too sudden. The conflicts are solved by removing an assignment from a path, if a neighboring path with a higher weight has a conflicting assignment. This method lets the more important paths keep their assignment, while the less important paths are reassigned later. Collecting the remaining assigned paths, an approximation of the shaded relief is generated which is used to derive the final shaded relief where the missing paths are integrated. The rendering of the assigned gray values on a white background looks like 41. It contains edges from both S^+ and S^- . This representation is dubbed *shaded relief skeleton*, inspired by the *Geripperelief* in [Hurni, 1989].

The mathematical machinery by which a shaded relief I is computed from the relief skeleton is the Laplace equation $\Delta I = 0$ over a bounded domain with Dirichlet boundary conditions ([Riley, 2006]). The domain is the collection of cells in I whose values are unknown. A Dirichlet condition designates a setup where the values at the boundary are fixed. In this case, the boundary is equated with the set of assigned gray values along the skeleton lines and a constant




Figure 41: **Left:** Shaded relief skeleton where paths with conflicting gray value assignments are discarded. **Right:** Location of the domain boundary (white) for the Laplace equation problem.

value at the image borders. These values represent that part of the solution which is already known. The solution of the Laplace equation for the unknown values in I , together with the values at the boundary, constitutes the desired output image. This formulation of interpolating unknown values in an image in terms of a boundary value problem is presented in [Fung and Parker, 1996]. The physical analogy to the smoothness of the solution of the Laplace equation is the smooth surface of a stretched elastic membrane held in place at its boundary (see figure 42). In the program code in section A.1.6 the location of the boundary is communicated in terms of a mask. An example of this mask is shown in figure 41 on the right-hand side. A progression of successive steps in the numerical solution of the equation is displayed in figure 43.




Figure 42: Two- and three-dimensional rendering of the smooth “membranes” between gray-value discontinuities (boundaries) as the solution of the Laplace equation.

Once the approximate shaded relief R is done, the missing paths are assigned. This is done by selecting the best fitting gray values, this time from one of all three images S , S^+ , or S^- , i.e. those values with enough contrast which are closest to R in terms of the sum of the absolute differences. The gray values from the chosen source image are linearly interpolated with the values from R to yield the final assignment for the missing paths. They are merged with the shaded relief skeleton of the first approximation and the Laplace equation is solved again, resulting in the final shaded relief (q.v. figure 67).



Figure 43: Progressive steps in the iterative solution of the Laplace equation with a given fixed boundary resulting in a smoothly shaded relief.

Flat regions in the terrain, like valley bottoms, bodies of water, or plains, are to be rendered in a medium gray plain tone ([Imhof, 1965]). In order to define these regions, the terrain slope is calculated, i.e. the magnitude of the gradient, and thresholded to yield a binary mask M of cells with a nonzero value for flat terrain, and zero everywhere else. Prior to thresholding, the slope values are optionally blurred using a Gaussian filter ([Sonka, 2008]) with a user specified standard deviation to smoothen the borders of the extracted regions. Additionally, holes in the regions may be filled using a morphological *closing* operation ([Sonka, 2008]). To incorporate the plain tone for flat regions in the shaded relief, the relief skeleton is erased where M is nonzero and filled with the user specified plain tone instead. Likewise, the flat regions are introduced into the mask defining the location of the boundary prior to solving the Laplace equation. An example can be found in figure 41 as the small roundish patch in the lower left quadrant.

5.2.3 Polishing the shaded relief

The shaded relief can be enhanced for display by various standard image processing tools like *histogram equalization* or *unsharp masking* for detail improvement ([Sonka, 2008]). A novel method for adding detail to a shaded relief is presented in [Geisthövel and Hurni, 2015], where Poisson editing is used to blend an image derived from the *mean curvature* of the terrain over the shaded relief, resulting in an enhanced display of structural features in the shaded relief. As an alternative to the mean curvature, the *convergence index* of the terrain ([Thommeret et al., 2010]), mapped to a gray scale image, can be utilized to similar ends. The convergence index is a morphometric tool to indicate areas of convergence and divergence of overland flow, i.e. channels and ridges. The gray value mapped convergence index image looks similar to the mean curvature, having low values at ridges and high values at channels, but it appears more

generalized and accentuated.




Figure 44: Terrain mean curvature (**left**) and convergence index (**right**), negated and mapped to gray value images.




Figure 45: **Left:** Shaded relief blended with the convergence index using Poisson editing. **Right:** Smoothed using rotational LIC, additionally.

In figure 44 the mean curvature and convergence index, both negated, are juxtaposed for comparison. The convergence index can be interpreted as an anisotropic, i.e. directionally independent record of structural lines, unlike the kind of structural lines extracted from shaded reliefs using an oriented light source. In this regard, enhancing a shaded relief with the aid of the convergence index is similar to the *texture shading* method ([Brown, 2010], [Patterson, 2014]), whereby the appearance of terrain features is accentuated without dependence on the exposition of the sur-

face. In figure 45 the blend of a shaded relief with the convergence index is shown, as well as a variant smoothed using rotational LIC.

5.3 Rock hachuring

With the shaded relief at hand, the rendering of the rock hachures can be undertaken. The fill hachure is the main carrier of the illuminational traits of the rock depiction. It is going to be drawn first, while the form strokes and contour strokes are added later, thereby painting over or erasing some of the hachure rendering.

5.3.1 Fill hachuring

Upon closer examination of manual rock depictions, it turns out that fill hachures are almost never drawn continuously over negative skeletal lines, i.e. trenches, ravines, or gullies, but are often seen crossing ridges, sometimes being a substitute for proper contour strokes (q.v. figure 3).




Figure 46: Negative skeleton lines (black), seed points (red), distance grid (background).

This principle is heeded by covering the entire terrain designated by the rock mask with equidistant horizontal hachures terminating at trenches. In order to extract these lines, the set of all skeletal lines, the paths computed from the Canny edges, is compared to the convergence index (q.v. section 5.2). Where this is above zero, i.e. displaying convergence rather than divergence, negative skeletal lines are recorded. The rendering of the horizontal, evenly spaced fill hachures, i.e. those indicating gently sloping terrain, drawn perpendicularly to slope lines, is performed using an algorithm originally developed for displaying streamlines in vector fields ([Jobard and Lefer, 1997]). In this algorithm, hachures are drawn commencing from seed points along the directions indicated by a vector field. The vector field V used to draw the horizontal fill hachures is the surface gradient rotated by $\frac{\pi}{2}$ to point in the direction of the level curves, and normalized to obtain unit vectors. In order to make sure that no space during hachure filling is squandered, an abundant set of points, like the coordinates of every second grid cell, is taken as the initial seed point set. Since hachure generation stops at trenches, the seed points most distant from these are preferable for the creation of new hachures. Thus, the initial set of seed points is put into a list which is sorted by decreasing distance to trench lines, i.e. the seed points most distant

to trenches appear first in the list. This distance is calculated on a binary grid, where the trench line cells are set to one and the background is set to zero. The *euclidean distance transform* of this grid ([Sonka, 2008]) yields the desired result. See figure 46 for a display of the negative skeletal lines, the seed points and the underlying distance transform image.




Figure 47: Seed points are generated from the central vertex of a hachure in the direction of the bisecting vector and its negative (dashed).

In order to grow hachures from seed points, a sampling grid like the one used in section 5.2 for sampling points on skeleton lines is created to enforce a minimal user supplied distance between fill hachures. This grid not only records seed points, but all vertices of hachures generated throughout the process. Iterating the list of seed points, each point is examined and looked up in the sampling grid if it is admissible for placement. If the point is rejected, because the space is already occupied by a vertex of another hachure, list iteration continues. Otherwise, the seed point is recorded in the sampling grid and a new hachure is created by putting the seed point into a new list and appending vertices at the end of the list by advancing the current point at the end of the list in the direction of the vector field V . This procedure continues until a trench line is reached or due to the proximity of other hachures as indicated by the sampling grid. Likewise, vertices are appended at the front of the list by advancing the points in the negated direction of the vector field similar to expanding the integration path in line integral convolution. The vertices in the list represent the hachure path. For efficiency reasons, a queue rather than a list may be used. Care must be taken to avoid overwriting the hachure by itself when hachuring around a peak without trenches to terminate the process. Once a hachure is finished, new seed points are not taken from the seed list, but from the hachure itself. This is achieved by translating the central vertex of the hachure in the direction of the vector bisecting the dihedral angle spanned by the vertex and its two neighbors. Since this kind of vector is referenced a few times in the following paragraphs, it is conveniently labeled as *bisecting vector*. See figure 47 for an example. The amount of translation is a user specified fill hachure distance. Likewise, the central vertex is translated in the negative direction of the vector, thus creating new seed points at both sides of the hachure at an even spacing. This process continues until no further hachures can be derived in this way, and picking seed points from the seed list is resumed.

Once every available space has been claimed by fill hachures, the inner vertices of the hachure paths are perturbed to lend the hachures the irregular, jittered appearance of fill hachures. This is done by iterating each path and translating the inner vertices, i.e. every vertex except the first and last, along a displacement bisecting vector. The displacement vector is a unit vector scaled by a sinusoidal with Gaussian random noise $s(v, t) = A \cos(\phi t) + G_\sigma$, where v denotes the vertex, t is the arc length at the vertex, A and ϕ are the user specified amplitude and frequency, respectively, and G_σ represents the noise source with a standard deviation of σ . It is important




Figure 48: Bisecting vectors at inner vertices used in the perturbation of the path. Note that the vectors need to point at the same side of the path.

for all displacement vectors to be oriented towards the same side of the hachure, like in figure 48, so that the alternating sign of the sinusoidal results in an even distribution of translations towards both sides of the hachure. The orientation of a vector $x = (x_1, x_2)$ in relation to a vector $y = (y_1, y_2)$ is calculated as the sign of $x_1y_2 - x_2y_1$. It is positive for $0 < \theta < \pi$ and negative for $\pi < \theta < 2\pi$, where θ is the angle between x and y . To ensure that all displacement vectors d_i point at the same side of the hachure, calculate the orientation of d_i and the vector from vertex v_{i-1} to v_i , and flip d_i if the sign does not match the specified side.




Figure 49: **Top:** A stroke polygon with an *intentional* constant stroke width (as indicated by the equal area circles at all vertices), but without a correctional scaling at inner vertices. **Bottom:** The same stroke path with a correctional factor applied at all inner vertices, resulting in an overall impression of constant stroke width.

Hereafter, the perturbed paths are converted into polygonal strokes of varying width to reflect the gray value modulations of the underlying shaded relief. The construction of the polygons follows [Northrup and Markosian, 2000]. The polygon boundary is defined by vectors originating from the path vertices. At the first and last vertex, these are the two unit vectors perpendicular to the incident path segment. For the inner vertices the unit bisecting vector and its negative are used. The polygon boundary is defined by connecting the end points of all the vectors in a clockwise fashion. To modulate the stroke width in terms of the shaded relief, these vectors are scaled for each vertex v by $\frac{\max((1-g(v))s_{\max}, s_{\min})}{2}$, where $g(v) \in [0, 1]$ is the gray value of the shaded relief at v , and s_{\min} and s_{\max} is the user specified minimum and maximum fill hachure stroke width, respectively. Thus, strokes are wide where the shaded relief is dark, and thin where it is bright. A correctional factor needs to be applied for the vectors at inner vertices to avoid a degradation of the desired stroke width at acute angles of consecutive path segments. See figure 49 for a stroke

path with an intentional constant stroke width, scaled both naively and using the correctional factor. This factor is $\max(|\vec{b}|/|\vec{b}\cdot\vec{n}|, 2)$, where \vec{b} is the scaled bisecting vector at the vertex v and \vec{n} is the unit vector perpendicular to the path segment between v and the next vertex. The factor amounts to one for straight consecutive path segments, and increases with the dihedral angle up to a maximum of a twofold scaling. This concludes the discussion of horizontal fill hachures.



Figure 50: Horizontal and vertical fill hachures depicting the rocky cliffs of the Weisshorn in the canton of Valais.

The vertical fill hachures, i.e. those indicating precipitous terrain, are drawn between selected horizontal fill hachure paths. The selection is done by sampling the slope in the *original* elevation model⁹ at every vertex of a fill hachure, and calculating the percentage of slope samples lying above a user specified threshold which indicates the transition between gentle and steep terrain. The original elevation grid is used here, since the slope values in the generalized model are diminished, as explained in section 5.1. If this percentage exceeds the minimum coverage value, also supplied by the user, the fill hachure is eligible for becoming the part of a vertical cliff hachure. If an adjacent hachure is also found to be eligible, the two are combined to form the upper and lower ledge of a vertical cliff, if they are roughly the same length, which is also verified by checking against the respective map setting. On each ledge an equal number of corresponding vertices is selected, depending on the fill hachure distance. A jittered stroke polygon is drawn between corresponding pairs of vertices of the upper and lower ledge, according to the rules of rock depiction. Reiterating the prescriptions given in section 2, the vertical strokes are tapered, i.e. downscaled towards the upper ledge on the light side, as well as towards the foot of the cliff on the shadow side. Gray values below 0.5 are considered to indicate the shadow

⁹Hence the arrow in the conceptual overview in figure 19 from *Elevation model* to *Rock hachuring*.

side. The upper and lower ledges are not drawn as continuous fill hachures, but as disconnected segments between the vertices of the vertical hachures. The rules dictate that on the shadow side the connectedness for upper and lower ledges is *strong* and *weak*, respectively, while it is *partial* and *full*, respectively, on the light side ([Federal Office of Topography swisstopo, 1990]). These adjectives are translated into probability values¹⁰ and ledge stroke polygon segments are inserted if a randomly picked value in the range [0, 100] from a uniform distribution falls below the respective probability. A rock depiction composed of both horizontal and vertical fill hachures, clipped against the rock mask, is displayed in figure 50.

5.3.2 Form and contour stroke rendering




Figure 51: Form stroke generation. **Left:** Negative skeletal lines in black. The arrows point in the opposite direction of the gray value gradient at the edge. **Center:** Negative skeletal lines with white intensity are translated in the direction of the gradient to become the characteristic *halos* of form strokes. **Right:** Form strokes and halos combined.

The form strokes and contour strokes are also derived from the shaded relief. This is achieved by translating and distributing the pixels of the positive and negative skeletal lines, respectively, in the *opposite* direction of the gradient of the gray values of the shaded relief. At the skeletal lines, the gray value gradient points in the direction of the steepest ascent, i.e. the brightest pixel. The pixels of the skeletal lines are interpreted in this process as units of black pigment. By translating these units for a number of times with a varying step width, strokes are constructed by accumulating fractions of blackness intensity on the darker side of the line. It is a purely raster based process, the form and contour strokes are not converted into polygons like the fill hachures. Where the gradient magnitude is high, i.e. where there is a strong gray value contrast, wide strokes are produced, while low contrast edges result in thin strokes. The number of iterations of this process is specified by the user. For the broad form strokes, that number needs to be greater than the number for the thin contour strokes. The form strokes are derived from the negative skeletal lines, since they are used to indicate trenches. The contour strokes, however, are derived from the positive skeletal lines alone, since there is no use in overwriting the already broader form strokes. For an individual pixel at integer coordinates $[i, j]$ on a skeletal line in the n th iteration, the displaced coordinates are floating point coordinates $(i + n \cdot ds \cdot gy, j + n \cdot ds \cdot gx)$, where ds is the step size and (gx, gy) is the gray value gradient at $[i, j]$. The blackness intensity is distributed in fractional parts to the destination pixel and its neighbors to avoid step-like stroke edges. This anti-aliasing technique is taken from [Wu, 1991]. See the program code listing A.1.8

¹⁰For instance *strong* \rightarrow 95 %, *weak* \rightarrow 10 %, *partial* \rightarrow 30 %, and *full* \rightarrow 100 %.

for details.




Figure 52: Rock depiction of the Weisshorn in the canton of Valais.

The characteristic blank spaces next to form strokes are created analogously, translating negative skeletal lines *in* the direction of the gradient. Prior to translation, the skeletal lines are clipped against the rock mask. The form strokes and contour strokes are incorporated into the fill hachure image by subtracting the blackness intensity of the strokes and clipping the result. In order to clear the blank spaces of hachures next to form strokes, the fill hachure image is divided and clipped by the inverted whiteness intensity b saturated shortly above zero to avoid a division error, i.e. $\max(0.01, 1 - b)$. This step finally completes the rock depiction. The result is shown in figure 52. In the appendix in figure 69, the rock depiction is combined with the shaded relief to demonstrate their close affinity.

6 Conclusion

Throughout the development of the proposed method, the following ideas have emerged as pivotal.

Rock depiction is a compound concept. The implementation of automatic rock depiction is a concerted effort in that it hinges on the availability and soundness of methods for terrain generalization and shaded relief. The relation of the latter two is such that there is no cartographic shaded relief without proper generalization. Thus, the linear dependency of generalization, shaded relief, and rock hachuring induces a sequential ordering of the processing steps.

Given a shaded relief, the derivation of rock hachures from the gray values of the relief is straightforward. This fact manifests the primacy of the shaded relief as the canonical illumination based terrain representation within the workflow.

The shaded relief is, first and foremost, a display of discontinuities, i.e. terrain edges. Thus, it is imperative to recognize the edges as the principal elements of consideration and to model the tones around these prior to handling the appearance of faces.

6.1 Problems solved and unsolved

A comparison of an automatically derived rock depiction R in figure 53 with an extract of the Swiss National Map 1:25'000 in figure 54, abbreviated as *SNM*, provides some insights into the workings of the method and highlights issues that need to be addressed in its further development.

The shaded relief and scree areas in *SNM* appear to be smoother and more generalized. The rock formations however, are more detailed and display a higher degree of stylization as compared to the rock hachures in R . Correspondingly, the batches of rock in the upper center of R appear somewhat flatter, less interrupted by vertical trench lines. Since these lines would be adhered to by the LIC generalization using the gradient as a guiding vector field, it is assumed that they are missed by the edge detection in the shaded relief. A remedy would be to incorporate other structural indicators in the edge detection process, like contour lines or the elevation model itself under the scrutiny of a topographical ridge and ravine analysis. Additionally, the border of the rock mask could be utilized as an indicator of structural surface features. In the current implementation, the rock mask is unduly used as a mere clipping device, yet it actually represents an important tool available to the user in determining the position, shape, and generalization of the rock areas. It could be put to more use in the depiction by emphasizing hachures approaching the border of the rock mask to indicate the foot or top of a cliff, for example. Another option lies in turning the rock mask border into an accentuated form or contour stroke itself, outlining the rock formation and making the area stand out from the surrounding map elements more clearly. Regarding the efficiency of hachuring, the rock mask should be introduced earlier into the workflow, while calculating the euclidean distance field for seed point placement and especially for pruning hachures outside of rock areas during hachure growing (q.v. section 5.3.1).




Figure 53: Shaded relief and rock depiction.

Obviously, there are many map elements in *SNM* which are missing in *R*. Scree is a carrier of the illumination intensity like rock hachures. It is placed along lines of slope due to the gravitational pull and the effects of erosion. Complementing the method for scree depiction in [Jenny et al., 2010], LIC generalization using the surface gradient which favors lines of slope would be an appropriate tool to generalize scree areas. As it has been indicated earlier, the shaded relief is a canonical common currency for all kinds of illumination based terrain representations. Thus, the modulation of scree according to lighting intensity could be derived straightforwardly from the shaded relief, analogously to the derivation of the rock hachures.

Another missing element, contour lines, could be supported easily by the sampling grid method. By placing contour lines into the grid prior to the insertion of the hachure vertices, they would not be touched or crossed by adjacent hachures.

The orientation of illumination matches well between *R* and *SNM*, except for the ridge of the




Figure 54: Swiss National Map 1:25'000 extract (© swisstopo, white arrow inserted by the author).

upper moraine rampart where the 2800 meter contour line crosses (indicated by white arrow in figure 54). In *SNM* a change of the orientation of illumination across the ridge takes place, where it is constant in *R*. The contrast flip in *SNM* leads to relief inversion and hinders the reading of the moraine formation as a whole, a strong advocacy in favor of the principle of making contrast maintenance across major skeletal lines a priority. Regarding the unburdened reading of recurring forms, it would be useful to have a higher order surface data structure available than just the path adjacency graph. For example, if a complete surface model was available, including edges and faces together with their relations, iconic shapes such as the triangular cones or rock pillars could be identified and rendered as units.

The vertical hachures coincide only partially in the images. In *R* the steepness of the terrain is sampled from the original elevation model. This means that hachures covering a precipitous face with an enlarged orthogonal footprint in the generalized elevation model tend to be wrongly classified as to be drawn horizontally (i.e. gently sloping), because they lie outside of the foot-

print of the original model. A better approach would be to project the original footprint onto the enlarged footprint in the generalized elevation model to benefit from a more detailed depiction using vertical fill hachures. The projection could be achieved in terms of mapping vertices of adjacent contour lines along with their slope values from the original elevation model to the smoothed model.

Generally speaking, the averaging capacity of simple LIC based generalization is very efficient at reducing representational artifacts. e.g. like visible triangles in an elevation grid derived from a triangulation. See figure 70 in the appendix for an example. However, the LIC generalization, as it is used here, is biased *against* slope break lines, since they lie perpendicular to the integration paths which are along slope lines. Thus, slope break lines may be smoothed out by the integration. As a remedy, an adaptive integration length could be introduced that is reduced as slope breaks are approached, i.e. at discontinuities in the profile curvature. Another approach lies in including integration barriers like they are used in rotational LIC at locations that must not be crossed during integration. Advancing the automation of the generalization using LIC requires finding a recipe to calculate the integration length depending on map output scale, raster cell size, and raster quality.

For an in-depth scrutiny of the usability of the maps generated in this novel fashion, a survey along the lines of the one conducted by Dahinden in [Dahinden, 2008] would be fruitful. The topics covered in this study include issues of orientability, like the ability to recognize the general shape of the illustrated surface, its slopes, as well as its constituent elements, like ridges, valleys, or peaks. Additionally, general issues like expressiveness and practicality are queried.

6.2 Adapting the method




Figure 55: **Left:** Swiss National Map extract showing an area with *loose rock* (© swisstopo). **Right:** Imitation of the *loose rock* style. The fill hachures are not accompanied by form strokes or contour strokes and are almost straightly drawn like contour lines. The yellowy-blueish coloration of the shaded relief is inspired by [Jenny and Hurni, 2006]. See figure 71 for an enlarged version.

Flexibility and extensibility are crucial in the handling of different use cases. Two more exam-

ples illustrate the versatility of the methods devised in this thesis.

The parameterization of the hachures is flexible. It allows for the imitation of the depiction of *loose rock* ([Jenny et al., 2014]) by discarding form strokes and contour strokes and drawing the fill hachures almost as straightly as contour lines. This is achieved by setting the perturbation parameters like amplitude or noise sigma to low values. See figure 55 for a comparison with an example taken from the Swiss National Map 1:25'000. Upon close inspection, the hachure strokes in the original are delicately bent as to indicate concavity of the surface, especially in the south-west face. This effect is lacking in the image of the imitation. Already a feature of Hurni's method ([Hurni, 1995]), the display of concavity or convexity of the terrain surface, illustrated via bent hachures, should be included in a future version by being able to translate hachure vertices depending on the curvature of the terrain.




Figure 56: **Top:** Shaded relief with diminished (0.75 scale, **left**) and increased (2.0 scale, **right**) gray value edge contrast, respectively. **Bottom:** Rock hachures derived from the shaded relief directly above.

Aerial perspective is an element of cartographic relief shading as explained in section 2, yet

it is not included in the implementation. However, a simple transformation of the gray values assigned to an edge allows for the scaling of contrast. This is how the aerial perspective effect is accomplished, contrasts at ridges are increased, while contrasts towards valley bottoms are diminished. The transformation is $\mu_g + s (g_i - \mu_g)$ for all gray values g_i of an edge, where μ_g is the mean gray value of the edge and s is a scale factor. The result is clipped to the range $[0, 1]$. If s is greater than one, the contrast is increased, while it is diminished for s less than one. In figure 56 the effects of contrast lowering or increasing are demonstrated.

This concludes the discussion of the method. The closing words, like carved in stone, belong to Imhof, *On the nature of cartographic representation* (in [Imhof, 2007]): “For all these reasons, the content and graphical structure of a complex, demanding map image can never be rendered in a completely automatic way.” The finality of this categorical statement is left to the reader to ponder on.

References

- [Blumer, 1927] Blumer, W. (1927). Die Felsdarstellung. Ein Beitrag zur neuen Landeskarte. *Die Alpen*, 3(12).
- [Brandstätter, 1968] Brandstätter, L. (1968). Baugerechte Felsdarstellung. *Bildmessung und Luftbildwesen*, 3.
- [Brassel, 1974] Brassel, K. (1974). A model for automatic hill-shading. *Cartography and Geographic Information Science*, 1(1).
- [Brown, 2010] Brown, L. (2010). A new technique for depicting terrain relief. Talk given at North American Cartographic Information Society Conference, St. Petersburg, FL., Oct. 2010.
- [Cabral and Leedom, 1993] Cabral, B. and Leedom, L. C. (1993). Imaging vector fields using line integral convolution. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. ACM.
- [Christophe et al., 2016] Christophe, S., Duménieu, B., Turbet, J., Hoarau, C., Mellado, N., Ory, J., Loi, H., Masse, A., Arbelot, B., and Vergne, R. (2016). Map style formalization: Rendering techniques extension for cartography. In *Expressive 2016 The Joint Symposium on Computational Aesthetics and Sketch-Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*. The Eurographics Association.
- [Dahinden, 2000] Dahinden, T. (2000). Weiterentwicklung des Programmsystems zur digitalen Felsdarstellung. Technical report, ETH Zürich, Zürich.
- [Dahinden, 2008] Dahinden, T. (2008). *Methoden und Beurteilungskriterien für die analytische Felsdarstellung in topografischen Karten*. PhD thesis, ETH Zürich.
- [de Berg et al., 2008] de Berg, M., Cheong, O., van Kreveld, M., and Overmars, M. (2008). *Computational geometry: algorithms and applications*. Springer, Berlin, 3rd edition.
- [Desimini and Waldheim, 2016] Desimini, J. and Waldheim, C. (2016). *Cartographic Grounds*. Princeton Architectural Press, New York.
- [Eberly, 1996] Eberly, D. (1996). *Ridges in image and data analysis*. Number 7 in Computational imaging and vision. Kluwer Academic Publishers, Boston, Mass.
- [Eysenck and Keane, 2000] Eysenck, M. W. and Keane, M. T. (2000). *Cognitive psychology*. Psychology Press ; Taylor & Francis, Hove, East Sussex, UK : Philadelphia, 4th edition.
- [Federal Office of Topography swisstopo, 1990] Federal Office of Topography swisstopo (1990). Topographische Kartographie Lehrlingsausbildung.
- [Federal Office of Topography swisstopo, 2016] Federal Office of Topography swisstopo (2016). Maps of Switzerland – Swiss Confederation – map.geo.admin.ch.

- [Fung and Parker, 1996] Fung, H. T. and Parker, K. J. (1996). Image interpolation as a boundary value problem. In *Visual Communications and Image Processing '96*. International Society for Optics and Photonics.
- [Geisthövel, 2013] Geisthövel, R. (2013). Towards automating swiss style rock depiction. Talk given at 26th International Cartographic Conference, Dresden, 2013.
- [Geisthövel and Hurni, 2015] Geisthövel, R. and Hurni, L. (2015). Automatic rock depiction via relief shading. In *Proceedings of the 27th International Cartographic Conference*, Rio de Janeiro.
- [Gilgen, 1998] Gilgen, J. (1998). Felsdarstellung in den Landeskarten der Schweiz. In *Hochgebirgskartographie – Silvretta '98, High Mountain Cartography*, number 11.
- [Gilgen, 2008] Gilgen, J. (2008). Aerial Photos + Photogrammetric Plot = swisstopo's Rock Representation: Caricatured Mountains? In *Proceedings of the 6th ICA Mountain Cartography Workshop*, Lenk.
- [Gilgen and Jenny, 2010] Gilgen, J. and Jenny, B. (2010). Digital rock and scree drawing in vector and raster mode. *Geographica Technica*, 12.
- [Gondol et al., 2008] Gondol, L., Le Bris, A., and Lecordix, F. (2008). Cartography of high mountain areas. In *6th ICA Mountain Cartography Workshop, Lenk, Switzerland*.
- [Grünwald and Kriz, 2015] Grünwald, M. and Kriz, K. (2015). Rock Depiction – A Semi-automated Approach. *Wiener Schriften zur Geographie und Kartographie*, 21.
- [Hake et al., 2002] Hake, G., Grünreich, D., and Meng, L. (2002). *Kartographie: Visualisierung raum-zeitlicher Informationen*. De Gruyter Lehrbuch. de Gruyter.
- [Hengl and Reuter, 2009] Hengl, T. and Reuter, H. I., editors (2009). *Geomorphometry: concepts, software, applications*. Elsevier.
- [Horn, 1981] Horn, B. (1981). Hill shading and the reflectance map. *Proceedings of the IEEE*, 69(1).
- [Hurni, 1989] Hurni, L. (1989). Verschiedene Felsdarstellungsarten für Gebirgskarten 1:25000. Technical report, Institut für Kartographie ETH Zürich.
- [Hurni, 1995] Hurni, L. (1995). *Modellhafte Arbeitsabläufe zur digitalen Erstellung von topographischen und geologischen Karten und dreidimensionalen Visualisierungen*. PhD thesis, ETH Zürich.
- [Hurni et al., 2001a] Hurni, L., Dahinden, T., and Hutzler, E. (2001a). Digital Cliff Drawing for Topographic Maps: Traditional Representations by Means of New Technologies. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 38(1-2).
- [Hurni et al., 2001b] Hurni, L., Jenny, B., Dahinden, T., and Hutzler, E. (2001b). Interactive Analytical Shading and Cliff Drawing: Advances in Digital Relief Presentation for Topographic Mountain Maps. In *Proceedings of the 20th International Cartographic Conference*, Beijing.

- [Hurni and Neumann, 1999] Hurni, L. and Neumann, A. (1999). Digitale Felsdarstellung für topographische Gebirgskarten. *Kartographische Nachrichten*, 49(1).
- [Imhof, 1965] Imhof, E. (1965). *Kartographische Geländedarstellung*. De Gruyter.
- [Imhof, 2007] Imhof, E. (2007). *Cartographic Relief Presentation*. ESRI Press.
- [Jenny et al., 2014] Jenny, B., Gilgen, J., Geisthövel, R., Marston, B. E., and Hurni, L. (2014). Design Principles for Swiss-style Rock Drawing. *The Cartographic Journal*, 51(4).
- [Jenny and Hurni, 2006] Jenny, B. and Hurni, L. (2006). Swiss-style colour relief shading modulated by elevation and by exposure to illumination. *Cartographic Journal*, 43(3).
- [Jenny et al., 2010] Jenny, B., Hutzler, E., and Hurni, L. (2010). Scree representation on topographic maps. *The Cartographic Journal*, 47(2).
- [Jobard and Lefer, 1997] Jobard, B. and Lefer, W. (1997). Creating evenly-spaced streamlines of arbitrary density. In *Visualization in Scientific Computing '97*. Springer Vienna.
- [Kanizsa, 1976] Kanizsa, G. (1976). Subjective contours. *Scientific American*, 234.
- [Kraiszl, 1930] Kraiszl, W. (1930). Historische Entwicklung der Felsdarstellung auf Plänen und topographischen Karten unter besonderer Berücksichtigung schweizerischer Verhältnisse. *Schweizerische Zeitschrift für Vermessungswesen und Kulturtechnik*, 1,2.
- [Lochmann, 1888] Lochmann, J. J. (1888). Instruktion für topographische Aufnahmen im Masstab 1:50,000. Eidgenössisches Topographisches Bureau, Bern, 1888.
- [Lysák, 2016] Lysák, J. (2016). An algorithm for automated digital rock drawing in the style used in Czech topographic maps. *AUC GEOGRAPHICA*, 51(1).
- [Marston, 2014] Marston, B. E. (2014). Improving the Representation of Large Landforms in Analytical Relief Shading. Master's thesis, Oregon State University, Corvallis.
- [Marston and Jenny, 2015] Marston, B. E. and Jenny, B. (2015). Improving the representation of major landforms in analytical relief shading. *International Journal of Geographical Information Science*, 29(7).
- [McCann and Pollard, 2008] McCann, J. and Pollard, N. S. (2008). Real-time gradient-domain painting. *ACM Transactions on Graphics*, 27(3).
- [Northrup and Markosian, 2000] Northrup, J. D. and Markosian, L. (2000). Artistic Silhouettes: A Hybrid Approach. In *Proceedings of the 1st International Symposium on Non-photorealistic Animation and Rendering*, New York, NY, USA. ACM.
- [Orzan et al., 2008] Orzan, A., Bousseau, A., Winnemöller, H., Barla, P., Thollot, J., and Salesin, D. (2008). Diffusion curves: A vector representation for smooth-shaded images. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2008)*, volume 27.
- [Patterson, 2014] Patterson, T. (2014). Enhanced shaded relief with terrain texture shader. www.shadedrelief.com/texture_shading, accessed March 2014.

- [Perona and Malik, 1990] Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7).
- [Porter and Duff, 1984] Porter, T. and Duff, T. (1984). Compositing digital images. In *ACM Siggraph Computer Graphics*, volume 18.
- [Press et al., 2007] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical recipes: the art of scientific computing*. Cambridge University Press, Cambridge, UK ; New York, 3rd edition.
- [Pérez et al., 2003] Pérez, P., Gangnet, M., and Blake, A. (2003). Poisson image editing. In *ACM Transactions on Graphics (TOG)*, volume 22. ACM.
- [Reece et al., 2010] Reece, J. B., Campbell, N. A., Urry, L., Cain, M., Wasserman, S., Minorsky, P., and Jackson, R., editors (2010). *Campbell biology*. Benjamin Cummings / Pearson, Boston, 9th edition.
- [Czech Office for Surveying, Mapping and Cadastre, 2016] Czech Office for Surveying, Mapping and Cadastre (2016). Map basis – Base map of the Czech Republic at 1 : 10 000 – www.cuzk.ch.
- [Riley, 2006] Riley, K. F. (2006). *Mathematical methods for physics and engineering*. Cambridge University Press, Cambridge ; New York, 3rd edition.
- [Russell, 2010] Russell, S. J. (2010). *Artificial intelligence: a modern approach*. Prentice Hall series in artificial intelligence. Prentice Hall, Upper Saddle River, 3rd edition.
- [Schneider, 1946] Schneider, K. (1946). Die Landeskarte der Schweiz 1:50000. *Geographica Helvetica*, 1:17–19.
- [Sonka, 2008] Sonka, M. (2008). *Image processing, analysis, and machine vision*. Thompson Learning, Toronto, 3rd edition.
- [Spiess, 1970] Spiess, E. (1970). Zur Felsdarstellung in unseren Gebirgsblättern. *Der Uto. Nachrichten der Sektion Uto des Schweizer Alpen-Club*, 48(10).
- [Spiess, 2014] Spiess, E. (2014). Topographisches Felszeichnen. Ein Rückblick auf die Exkursion "Topographisches Felszeichnen" nach Gimmelwald und Mürren, im Rahmen des Kartographiekongresses Interlaken im Juni 1996.
- [Thommeret et al., 2010] Thommeret, N., Bailly, J. S., and Puech, C. (2010). Extraction of thalweg networks from DTMs: application to badlands. *Hydrology and Earth System Sciences*, 14(8).
- [Tufté, 1990] Tufté, E. R. (1990). *Envisioning information*. Graphics Press, Cheshire, Conn.
- [Washburn, 1988] Washburn, B. (1988). Mount Everest: Surveying the Third Pole. *National Geographic*, 174(5).
- [Wedekind and Ortner, 2004] Wedekind, H. and Ortner, E. (2004). Toward universal literacy: from computer science upward. *Communications of the ACM*, 47(6).

- [Wedekind et al., 2004] Wedekind, H., Ortner, E., and Inhetveen, R. (2004). Informatik als Grundbildung: Teil 1. *Informatik-Spektrum*, 27(2).
- [Wu, 1991] Wu, X. (1991). An efficient antialiasing technique. In *ACM SIGGRAPH Computer Graphics*, volume 25. ACM.
- [Yang et al., 2009] Yang, N., Guo, Q., and Shen, D. (2009). Automatic modeling of cliff symbol in 3d topographic map. In *Proc. SPIE 7492, International Symposium on Spatial Analysis, Spatial-Temporal Data Modeling, and Data Mining*.
- [Yang et al., 2015] Yang, N., Wan, L., Zheng, G.-z., and Yang, J. (2015). Using Hachures to Construct a 3d Doline Model Automatically. *Cartographica*, 50(2).
- [Yoeli, 1965] Yoeli, P. (1965). Analytical hill shading. *Surveying and Mapping*, 25(4).
- [Zwimpfer, 2001] Zwimpfer, M. (2001). *2d Visual Perception*. Niggli AG.

A Appendix

A.1 Program code listings

This section contains listings of procedures mentioned throughout the text. In these listings, code clarity and conciseness are favored over efficiency issues. They are written in a procedural style using a syntax very similar to the Python¹¹ programming language. The syntax should be fairly self explanatory. However, frequently used constructs are summarized below.

```
def adder(a, b=3.1415):
    """
        This is a procedure definition named ADDER with two
        parameters A and B (with a default value of 3.1415)
        and a multiline documentation.
        The procedure returns the sum of the parameters.
    """
    # This is a comment; the sum of a and b
    # is assigned to a local variable
    res = a + b
    return res
```

Variables do not have to be declared, before they are used, just defined. The assignment syntax is flexible, thus multiple assignments per = sign are possible as well as the swapping of values in one line.

```
H, W = height, width # H = height; W = width
I, J = J, I # Swap i and j
```

Compound data types are tuples and dictionaries (key-value maps).

```
(2,4,8) # This is a tuple
d = { 1 : 2, 3 : 4 } # Dictionary. 1 is mapped to 2, 3 is mapped to 4
d[1] # Dictionary access. This yields 2
d[5] = 6 # Dictionary assignment. The key 5 is assigned the value 6
1 in d # Dictionary query. This yields TRUE
2 in d # This yields FALSE; only keys can be queried in this way, not values
```

The availability of a two-dimensional grid data type with zero based indexing and row-major order is assumed, along with some auxiliary operations summarized in the table below.

Name	Description
shape(X)	Dimension of X as a tuple (# rows, # columns)
copy(X)	A copy of X
fill(X, value)	Fill X with a value
X[i, j]	Access X at row index i and column index j
X[p]	Access X by a tuple p, p=(i,j)

¹¹Homepage at www.python.org

Iteration constructs include `for` loops and `while` loops. The assignment rules apply for the setting of the iteration variable in `for` loops.

```
# Looping over a tuple; s is 6 at the end of the loop
s = 0
for i in (1,2,3):
    s += i

# Same result as above; range(n) iterates from 0 to n-1
s = 0
for i in range(3):
    s += i+1

# Same result as above; range(start, end) iterates from start to end-1
s = 0
for i in range(1,4):
    s += i
```

A.1.1 Line integral convolution

```
def line_integral_convolution(z, Vx, Vy, L, ds):
    """
        Z          Elevation raster
        VX, VY     Vector field components
        L          Integration length
        DS         Integration step size
    """
    z2 = copy(z) # The output raster
    H, W = shape(z) # Height and width
    for i in range(H):
        for j in range(W):
            s = z[i,j]
            # Counter for the actual number of samples. Should be
            # 2*L+1 for cells more than ds*L/2 units away from boundary
            n = 1
            # Forward integration
            p, q = i, j
            y, x = i+0.5, j+0.5
            for k in range(L):
                if 0 <= p < H and 0 <= q < W:
                    y += ds*Vy[p,q]
                    x += ds*Vx[p,q]
                    p, q = int(y), int(x)
                    s += z[p,q]
                    n += 1
            # Backward integration: walking along negated vector field
            p, q = i, j
            y, x = i+0.5, j+0.5
            for k in range(L):
                if 0 <= p < H and 0 <= q < W:
                    y -= ds*Vy[p,q]
                    x -= ds*Vx[p,q]
                    p, q = int(y), int(x)
```

```

        s += z[p,q]
        n += 1
    z2[i,j] = s / n
return z2

```

A.1.2 Poisson editing

```

def poisson_edit(dst, src, eps=1e-7):
    """
        Poisson editing of DST using SRC

        DST          2d grid
        SRC          2d grid
        EPS          Residual error threshold for termination
    """

    H, W = shape(dst) # height, width

    # f is the right-hand side of Poisson equation.
    f = copy(dst)
    fill(f, 0)
    for i in range(1,H-1):
        for j in range(1,W-1):
            f[i,j] = blend_max(dst, src, i, j)

    # Initial solution and output
    R = copy(dst)

    # Iterative solution using Gauss-Seidel method
    while True:
        residual = 0

        for i in range(1,H-1):
            for j in range(1,W-1):

                Rold = R[i,j] # Save old state

                R[i,j] = (R[i+1,j] + R[i-1,j] +
                        R[i,j+1] + R[i,j-1] - f[i,j]) / 4

                residual += (Rold - R[i,j])**2 # ** denotes exponentiation

        if residual < eps:
            return R

def blend_max(dst, src, i, j):
    """
        Right-hand side of Poisson eq. Sum of absolute
        maximum of DST and SRC in 4-neighborhood.

        DST      2d grid
        SRC      2d grid
        I        Row index
        J        Column index
    """

```

```

"""
H, W = shape(dst) # height, width
r = 0
for p,q in ((i-1,j), (i+1,j), (i,j-1), (i,j+1)):
    s = src[p,q] - src[i,j]
    d = dst[p,q] - dst[i,j]
    if abs(s) > abs(d):
        r += s
    else:
        r += d
return r

```

A.1.3 Shortest path breadth first search

```

def path_search(grid, start, goal):
    """
    Search a shortest path on a GRID between two
    coordinates START and GOAL. The GRIDS cell values
    are zero for background pixels and nonzero otherwise.
    """

    todo = Queue() # define todo as an empty queue
    enqueue(todo, start) # put start into queue

    # Dictionary recording successor relation.
    # For a : b, a succeeds b.
    trace = { start : start } # start is its own successor

    H, W = shape(grid) # Height and width
    while not empty(todo):
        x = dequeue(todo) # get head element of queue
        if x == goal:
            path = queue() # reconstruct path from start to goal
            while x != start:
                enqueue(path, x)
                x = trace[x]
            enqueue(path, start)
            return path

        i,j = x
        for p,q in ((i-1,j-1), (i-1,j), (i-1,j+1),
                   (i,j-1), (i,j+1),
                   (i+1,j-1), (i+1,j), (i+1,j+1)):
            if 0 <= p < H and 0 <= q < W:
                if grid[p,q] != 0 and (p,q) not in trace:
                    trace[(p,q)] = x
                    enqueue(todo, (p,q))

```

A.1.4 Sampling grid

```

def make_sampling_grid(src_grid_shape, d):
    """

```

```

    Return a sampling grid with cell diagonal length D.
    Cells have a width and height of D / sqrt(2).
    The distance D does not need to be an integer.
    SRC_GRID_SHAPE is the dimension tuple (height, width) of
    the underlying grid for which sampling is to be done.
    The returned grid is to be used in conjunction
    with the procedure 'sample_grid_insert' below.
    """
H, W = src_grid_shape
cellsize = d / sqrt(2)

# Sampling grid height and width
# int(x) is a truncating integer conversion of x
sg_height = int(H / cellsize) + 1
sg_width = int(W / cellsize) + 1

# Make a new grid of given dimensions with cell type boolean
sgrid = Grid((sg_height, sg_width), bool)
fill(sgrid, False) # Initially, all cells are empty
return sgrid

def sample_grid_insert(sample_grid, hit_map, d, point):
    """
    SAMPLE_GRID    A sampling grid as returned by 'make_sampling_grid'.
                   It is modified inside this procedure.
    HIT_MAP        A dictionary mapping grid cell indices
                   to sample coordinates. Modified inside this procedure.
    D              The sampling distance. Same D as in 'make_sampling_grid'
    POINT         A tuple (x,y) with coordinates in the grid to be sampled.

    If the POINTS does not violate the minimal distance to existing samples,
    it is inserted into the SAMPLE_GRID and HIT_MAP and TRUE is returned.
    Otherwise, only FALSE is returned.
    """

    x,y = point
    cellsize = d / sqrt(2)

    # Map point coordinates to sample grid cell coordinates [row, col]
    # % is the floating point modulo operator
    row = int((y - (y % cellsize)) / cellsize + 0.5)
    col = int((x - (x % cellsize)) / cellsize + 0.5)

    H, W = shape(sample_grid) # height, width

    if 0 <= row < H and 0 <= col < W:
        if sample_grid[row, col] == True:
            # Sample grid cell is already occupied -> point is discarded
            return False
        else:
            # Check 8-neighborhood of sample grid cell for collisions
            for i,j in ((-1,-1), (-1,0), (-1,1), (0,-1), (0,1),
                       (1,-1), (1,0), (1,1)):
                ri = row + i
                cj = col + j
                if 0 <= ri < H and 0 <= cj < W:

```

```

        if sample_grid[ri, cj] == True:
            # Neighbor is occupied by a sample
            # Retrieve sample coordinates(v,w) and test distance
            v,w = hit_map[ri*W + cj]
            if sqrt((v-x)**2 + (w-y)**2) < d:
                # Coordinates too near -> discard
                return False

            # OK, no occupants in neighborhood or
            # all coordinates further away than d -> insert point as sample
            hit_map[r*W + c] = point
            sample_grid[r, c] = True
            return True
    else:
        # Outside sampling grid
        return False

```

A.1.5 Relief skeleton auxiliary functions

```

def overall_gradient(grid, Cl, Cr):
    """
        Return the overall gradient of an extended path, represented
        by the coordinates CL and CR on both sides of the path. The
        The gradient is calculated as the difference of the
        median of the gray values of CL and CR in GRID.
    """
    return median(grid[p] for p in Cl) - median(grid[p] for p in Cr)

def overall_grad_sign_cmp(grid1, grid2, Cl, Cr):
    """
        Return TRUE iff the signs of the overall gradients of the
        extended path represented by CL and CR (as in overall_gradient)
        is equal in grids GRID1 and GRID2.
    """
    og1 = overall_gradient(grid1, Cl, Cr)
    og2 = overall_gradient(grid2, Cl, Cr)
    return signum(og1) == signum(og2)

def edge_gray_compare(grid1, grid2, C, Cl, Cr):
    """
        Similarity measure for extended paths given the coordinates
        C, CL and CR between two grids GRID1 and GRID2. The smaller
        the result, the more similar are the gray values.
    """
    s = 0
    for p in (C, Cl, Cr):
        s += abs(grid1[p] - grid2[p])
    return s

```

A.1.6 Laplace shader

```
def smooth_shader(rskel, mask, eps=1e-7):
    """
        Derive a smoothly shaded relief image from a relief skeleton.

        RSKEL      Relief skeleton.
        MASK       Indicates the location of the relief skeleton
                   (i.e. Dirichlet boundary) by values other than 0.
                   A value of 0 indicates the background.
        EPS       Threshold value of residual error for
                   terminating the computation.
    """

    H,W = shape(rskel) # Height and width

    # Output image and initial solution
    R = copy(rskel)
    fill(R, 0.5)

    # Fix the boundary
    for i in range(H):
        for j in range(W):
            if mask[i,j] != 0:
                R[i,j] = rskel[i,j]

    # The output R is the solution of a Laplace equation with
    # the relief skeleton as a boundary condition.
    # The equation is solved using the Gauss-Seidel method.
    while True:
        residual = 0

        for i in range(1,H-1):
            for j in range(1,W-1):
                if mask[i,j] == 0:
                    # Save old state
                    Rold = R[i,j]
                    R[i,j] = (R[i-1,j] + R[i+1,j] + R[i,j-1] + R[i,j+1]) / 4
                    residual += (Rold - R[i,j])**2

        if residual < eps:
            return R
```

A.1.7 Convergence index

```
def convergence_index(expo):
    """
        Calculate the convergence index from the surface exposition.

        EXPO      A grid containing the exposition in radians
    """

    # Output grid
```

```

R = copy(expo)
fill(R, 0)

# Offsets to 8-neighborhood, contour-clockwise from east
N = [(0,1), (-1,1), (-1,0), (-1,-1), (0,-1), (1,-1), (1,0), (1,1)]

H, W = shape(expo)

for i in range(1,H-1):
    for j in range(1,W-1):
        for n in range(8):
            di,dj = N[n]
            T = expo[i+di,j+dj] - n*Pi/4 # Pi = 3.1415...

            if T > Pi:
                T -= 2*Pi
            elif T < -Pi:
                T += 2*Pi

            R[i,j] += abs(T)

        R[i,j] = R[i,j] / 8
        R[i,j] -= PI/2

return R

```

A.1.8 Form and contour stroke generation

```

def pixel_wu(x, y):
    """
        Return the brightness and distribution
        around a pixel at floating point coordinates (X,Y)
        as a tuple with elements (i, j, brightness),
        where i and j are integer coordinates.
    """
    # Truncate x and y to integer values
    ix = int(x)
    iy = int(y)

    # Fractional parts
    fx = x - ix
    fy = y - iy

    # Inverted fractional parts
    invfx = 1.0 - fx
    invfy = 1.0 - fy

    return ((ix, iy, invfx * invfy),
            (ix+1, iy, fx * invfy),
            (ix, iy+1, invfx * fy),
            (ix+1, iy+1, fx * fy))

```

```

def translate_lines(gx, gy, line_coords, num_iter):
    """
        Translate skeletal line coordinates a number of times. Values
        are accumulated in the output grid.

        GX          2D grid with the gray value gradient x-component
        GY          2D grid with the gray value gradient y-component
        LINE_COORDS Tuple of coordinates [i,j] of skeletal lines
        NUM_ITER    Number of iterations
    """

    # Output grid
    R = copy(gx)
    fill(R, 0)

    DS = 0.5 # Step size

    H, W = shape(gx) # Height and width

    for n in range(num_iter):
        for i,j in line_coords:
            for dj, di, b in pixel_wu(n * DS * gx[i,j],
                                     n * DS * gy[i,j]):
                if 0 <= di < H and 0 <= dj < W:
                    R[di, dj] += b

    return R

```

A.2 Map settings

An overview of input parameters, together with default value and unit, if applicable.

A.2.1 Generalization

Name	Description
LIC length	Number of steps in line integral convolution
Line threshold	Minimal length for skeletal lines, shorter are discarded (5 pixels)
Plain threshold	Slope value below which terrain is considered flat (2°)

A.2.2 Shading

Name	Description
Azimuth	Light source planar angle (315°)
Elevation	Light source elevation angle (45°)
Canny sigma	Gaussian standard deviation for Canny edge detector (1)
Canny low/hi	Canny edge detector hysteresis thresholds (0.05/0.15)
Plain tone	Gray value used to shade flat areas (0.8)
Poisson epsilon	Residual below which to terminate iterating Poisson solution ($1e-7$)
Laplace epsilon	Residual threshold for Laplace equation solution ($1e-7$)
Rotation step	Step width in rotational LIC (5°)

A.2.3 Hachuring

Name	Description
Distance	Minimal distance between hachures
Stroke size	Minimal and maximal stroke width
Steepness	Slope threshold above which to perform vertical hachuring (35°)
Amplitude	Hachure vertex perturbation amplitude
Frequency	Perturbation sine wave frequency
Sigma	Perturbation standard deviation for Gaussian noise
Stroke iterations	Number of iterations in form and contour stroke generation

A.3 Images

This section contains enlarged or alternative versions of images mentioned in the text, as well as three more examples illustrating the method.




Figure 57: Hillshaded reliefs of a progression of L1C generalized elevation rasters with increasing integration length, line by line, from top left: 5, 10, 15, 25, 50, 100. The raster cell size is 2 meters.




Figure 58: Swiss National Map 1:10'000 superimposed with dark 20 meter equidistance contours of the original raster (**left**) and generalized raster (**right**, enlarged version of 25).




Figure 59: Swiss National Map 1:10'000 extract (© swisstopo).




Figure 60: Block diagrams of original and generalized elevation model.



Figure 61: Lambertian hillshaded relief, incident light indicated by arrow (**top left, right**).
Bottom: Poisson-blending of the two hillshaded images.




Figure 62: Blend using Poisson editing (**bottom**) of a Lambertian hillshade h with illumination azimuth α (**top left**) and the sidelight $Sl(\alpha)$ (**top right**).




Figure 63: Hillshaded reliefs s and t (**top**) with their respective Canny edges (**bottom**).




Figure 64: **Left:** A hillshaded relief. **Right:** Smoothed version of the hillshaded relief using rotational LIC.



Figure 65: A progression of hillshades smoothed with rotational LIC with an increasing number of iterations (1, 2, 4, 8, 16, 32, from top left to bottom right).




Figure 66: **Left:** S^+ . **Right:** S^- .





Figure 67: **Left:** Shaded relief skeleton with the final distribution of gray values around skeleton lines. **Right:** Shaded relief interpolated from this skeleton.




Figure 68: Skeleton lines and filtered Delaunay triangulation of line samples.



Figure 69: Shaded relief and rock depiction of the Weisshorn in the canton of Valais.




Figure 70: **Top:** Hillshaded elevation grid of the Lhotse mountain. The grid with a cell size of 10 meters is derived from a triangulation. **Bottom:** Shaded relief of the generalized elevation model using LIC with an integration length of 25. The triangulation points were provided by the Institute of Geodesy and Photogrammetry, ETH Zürich; the grid itself and the rock mask were provided by the Swiss World Atlas, ETH Zürich.




Figure 71: **Right:** Swiss National Map extract showing an area with *loose rock* (© swisstopo). **Left:** Imitation of the *loose rock* style. The fill hachures are not accompanied by form strokes or contour strokes and are almost straightly drawn like contour lines.




Figure 72: Shaded relief of the high cliffs above the Walensee.



Figure 73: Shaded relief and rock depiction of Nuptse, Lhotse and Mount Everest.



Figure 74: Maybe the most famous peak of Switzerland. The image features a peculiar rock mask, ending just behind the national border. It offers a fine opportunity to state the ontological condition of rock as it is understood in this thesis: *rock* is that object which is designated by the rock mask.

A.4 Acknowledgments

I thank the following persons or groups of persons: Angeliki Tsorlini for providing the Methana map, Annelies Häcki Buhofer, Julia Häcki and Hans-Rudolf Bär for proofreading, the Swiss World Atlas team for collaboration on the Everest map, the Atlas of Switzerland team for sharing office space, as well as my supervisors Lorenz Hurni and Ralf Bill for their constructive support.