Local Shading Analysis via Isophotes Properties

by

Radim Šára

A thesis submitted to the faculty of
Johannes Kepler University in Linz
in partial fulfillment of the requirements for the degree of

"Doktor der technischen Wissenschaften"

Department of Systems Theory and Information Engineering
Institute of Systems Science
Johannes Kepler University
A-4040 Linz / Austria

March 1994

The author was on the leave from Computer Vision Laboratory
Department of Control Engineering, Faculty of Electrical Engineering
Czech Technical University
Karlovo náměstí 13, CZ-121 35 Praha 2 / Czech Republic

To **Trurl** and **Klapacius**,
the two famous *constructors*.

Acknowledgements

First of all, I would like to thank my supervisor, Professor Franz Pichler, for his support over my post-graduate study in Linz. He gave me the opportunity to work with his department to complete the study and finish this thesis.

Many warm thanks also go to Dr. Wilhelm Burger—a brilliant scientist and a very good man—whose kind support let me overcome some initial difficulties and who let me do my research within his project. His comments and suggestions helped to clarify the text.

Many thanks to my co-adviser, Dr. Václav Hlaváč, the head of Computer Vision Laboratory in Prague whom I owe a lot for his constant support of my work and for his encouragement when I was in doubts. He also provided valuable comments on this work.

I would like to commemorate my colleague, Jaroslav Woska, who opened the world of computer vision for me during my undergraduate studies. He tragically died in Japan in the course of his own Ph.D. studies. Many thanks to him!

I also want to thank all my colleagues at the Institute of Systems Theory and Information Engineering, namely Thomas Müller-Wipperfürth, Herbert Prähofer, Rudolf Mittelmann, Christoph Schaffer, Josef Scharinger, and others. They created a stimulating and friendly working atmosphere in their laboratory.

My special thanks to my parents for giving me and my brother a warm home in which we learned, read, and grown. They did not give in to the dogmatic, mendacious, and intolerant world we all had to live in. Their supportive encouragement of good work at school and their honourable life made us what we are.

Finally, I thank all those who made this research possible. The research has been carried out at the Department of System Science at the Johannes Kepler University Linz within the framework of a cooperation with the Technical University in Prague and the Technical University in Vienna. I thank Professor Vladimír Mařík from Prague and Professor Walter Kropatch from Vienna for their support of this cooperation. The financial support for the research project I worked within was provided by the *Fonds zur Förderung der wissenschaftlichen Forschung* in Austria under the grant P8785 "Theory and Applications of Digital Image Processing and Pattern Recognition." The research has also been partially supported by the *Grant Agency of the Czech Republic* under the grant 102/93/0954 "3-D Reconstruction and Modelling Using Active Vision," by the Ministry of Economic Policy and the Czech Republic Development under the grant Z038 "Digitized Teeth Reconstruction," and by matching internal grants of the Czech Technical University in Prague 8093 and 8066.

Eidesstattliche Erklärung

Ich erkläre an Eides statt, daß ich vorliegende Dissertation selbstständig verfaßt, andere als die angegebenen Quellen und Hilfsmittel nicht benutzt und mich auch sonst keiner unerlaubten Hilfe bedient habe.

Linz, am 21. Jänner 1994 Radim Šára

1st *referee (1. Begutachter):*

O. Univ.-Prof. Dr. Franz Pichler
Institute of Systems Science
Johannes Kepler University
A-4040 Linz / Austria

2nd *referee (2. Begutachter):*

O. Univ.-Prof. Dr. Jens Volkert
Institute of Informatics
Johannes Kepler University
A-4040 Linz / Austria

Kurzfassung

Diese Arbeit behandelt die lokale Interpretation von Bildern beleuchteter Flächen mit dem Ziel, lokale Oberflächencharakteristika ohne Oberflächenrekonstruktion zu erhalten. Einschränkungen in Bezug auf die lokale 3-D Flächenform, welche durch veränderliche Schattierung gegeben sind, werden untersucht. Bekannten Methoden fehlt vielfach ein physikalisch motivierter Ansatz. Wir schlagen Linien gleicher Lichtintensität (Isophoten), die durch die Reflexionsphysik induziert werden, als geeignete und physikalisch relevante Strukturen vor. Sie schränken in hohem Maße die differentialgeometrische Flächenform ein, falls die Anordnung zum Zeitpunkt der Bildaufnahme bekannt ist.

Eine Analyse der Schattierung in der Nähe von Objektgrenzen, von polyhedralen Knoten und singulären Punkten in Bildern wird angegeben. Wir zeigen, daß, im Gegensatz zu früheren Annahmen, die lokale Oberflächenorientierung am Rande eines Eigenschattens eindeutig bestimmt werden kann. Wir erhärten ein Ergebnis von Saxberg, demzufolge an singulären Bildpunkten, welche durch Extremwerte der Oberfläche bedingt sind, mindestens zwei und höchstens vier Lösungen für die Hauptkrümmung einer nicht-entwickelbaren Fläche existieren. Weiters zeigen wir, daß der Flächenindex nach Koenderink an diesen Stellen eindeutig meßbar ist.

Oberflächendarstellungen basierend auf (1) einer Parameterisierung einer nicht-entwickelbaren Fläche "durch ihren Normalvektor," und (2) den Weingarten-Koeffizienten werden vorgestellt. Diese Darstellung erweist sich insbesondere für die Untersuchung der lokalen Schattierung als nützlich, da sie die Problematik beherrschbar macht ohne *a priori* Einschränkungen bezüglich der Art oder Ordnung der Fläche zu machen, wie es bei anderen Autoren üblich ist. Wir vergleichen diesen neuen Ansatz mit existierenden Methoden, sowohl durch theoretische Untersuchungen als auch anhand von Experimenten mit synthetischen Daten.

Schlüsselworte

Computer Vision, 3-D Geometrie, Oberflächenrekonstruktion, Flächenrekonstruktion aus Schattierung, Lokale Analyse der Schattierung, Flächenzüge, Isophoten.

Abstract

This work deals with local image interpretation of an illuminated surface in order to get local shape features without surface reconstruction. Constraints on local 3-D shape imposed by varying shading are studied. Known solution methods suffer from the lack of a physically motivated approach. We propose *isophotes* as suitable and physically relevant structures induced by reflectance physics. They strongly constrain the differential surface shape if the imaging geometry is known.

An analysis of shading near boundaries, polyhedral vertices, and image singularities is given. We prove that, contrary to previous beliefs, local surface orientation near a self-shadow boundary can be determined uniquely. We confirm Saxberg's result that at an image singularity due to an extremal surface point there are at least two and at most four solutions for the principal curvatures of a non-developable surface. We also show that Koenderink's shape index can be measured uniquely at these locations.

A surface representation based on (1) a parameterization of a non-developable surface "by its normal vector," and (2) the Weingarten coefficients is introduced. This representation is particularly suitable for local shading analysis, since it makes the problem tractable without *a priori* constraining the surface to any particular type or order, as usually done by other authors. We compare this new approach to existing methods by theoretical analysis and experimental results on synthetic data.

Key words

Computer Vision, 3-D Vision, Shape from Shading, Local Shading Analysis, Shape Features, Isophotes.

Contents

List of Figures . . . iii

List of Tables . . . v

List of Algorithms . . . vi

Mathematical Notation . . . viii

1 Introduction **1**

1.1 Methodology . . . 3

1.2 Previous Work Related to *LSA* . . . 5

1.3 Overview of the Thesis . . . 11

2 Basic Concepts **13**

2.1 Differential Geometry of Surfaces . . . 13

2.2 Model of Image Formation . . . 18

2.2.1 Surface Reflectance . . . 19

2.2.2 Image Projection and Natural Coordinate Systems . . . 22

2.3 Isophotes . . . 25

3 Shading Analysis near Boundaries and Singular Points **28**

3.1 Occlusions, Self-Shadow Boundaries and Surface Discontinuities . . . 28

3.1.1 Isophotes and Surface Orientation Near an Occlusion Boundary . 29

3.1.2 The Sign of the Gaussian Curvature on an Occlusion Boundary . 33

3.1.3 Surface Orientation on a Self-Shadow Boundary . . . 35

3.1.4 Shading and Surface Discontinuities . . . 39

3.2 Image Singularities . . . 44

3.2.1 Classification . . . 45

3.2.2 Properties Related to Surface Geometry . . . 46

3.2.3 Shading and Shape Near Isolated Singular Points . . . 50

4 Local Analysis of the Isophote Field **55**

4.1 Properties of Image Isophotes 55

4.1.1 T-invariance and R-invariance 58

4.2 Gaussian Parameterization of a Non-Developable Surface 60

4.2.1 Existence Theorem 63

4.3 The Difference Angle between Normal Projection and Image Gradient in $\mathcal{L}$-Plane 66

4.4 The Sign of the Gaussian Curvature on Occlusion Boundaries Revisited . 71

4.5 Shading on Developable Surfaces 72

5 Local Orientation and Local Shape Quality from Shading **75**

5.1 The Existence of a Second-Order LSA Algorithm 75

5.2 Recovering Local Orientation and Surface Curvature from Images of Spheres 77

5.3 A Reasonable Estimator of Local Surface Slant 79

5.4 Estimator of Local Tilt for a General Surface 80

5.5 A Relaxation Algorithm for Normal Tilt Recovery 81

5.6 When Does the Tilt Estimator Based on (4.24) Become Ill-Posed? 84

6 Computational Experiments **87**

6.1 Standard Test 87

6.2 Evaluation of LSA Algorithms 90

6.2.1 Tested Algorithms 94

6.2.2 Slant Estimators 95

6.2.3 Tilt Estimators 98

6.2.4 Gaussian Curvature Recovery 105

6.2.5 The Influence of Surface Curvature on the Tilt Estimate 109

6.2.6 Sensitivity of Tilt Estimators to Explicit Slant Knowledge 109

6.2.7 Sensitivity to Illuminant Position 112

6.2.8 The Influence of ξ and R on the Iterative Algorithm Convergence 115

6.2.9 3-D Surface Reconstruction from Needle Maps 115

6.2.10 Recovering Needle Maps of Developable Surfaces 122

6.2.11 Summary of the Proposed Algorithms Properties 124

6.2.12 Notes on Gaussian Curvature Recovery 126

7 Conclusions and Pointers **129**

7.1 Open Questions . . . 129

7.2 Summary . . . 131

Bibliography **135**

List of Figures

1.1 The crater illusion . . . 6

2.1 Local imaging geometry . . . 20
2.2 The projection plane of a pinhole camera . . . 23
2.3 Global imaging geometry . . . 24

3.1 Surface tangent plane at a point on an occlusion boundary . . . 30
3.2 Image gradient operator's behaviour on an occluding boundary . . . 32
3.3 The surface normal vector at a point on an occlusion boundary . . . 33
3.4 Occlusion boundary perception . . . 33
3.5 Basic vectors related to a special parameterization near the occlusion boundary . . . 34
3.6 A surface discontinuity as an arc of a spatial curve . . . 39
3.7 A junction of image edges as an image of a surface vertex . . . 41
3.8 Two types of image singularities due to a τ-parameter curve touching the parabolic line on the Gaussian sphere . . . 47
3.9 The moment of creation of a couple of apices on a parabolic line . . . 48
3.10 The moment of creation of a saddle and an extremum on a parabolic line 49

4.1 Slant σ and tilt τ of a vector $\mathbf{x}$ with respect to a vector $\mathbf{L}$ . . . 61
4.2 Gaussian parameterization of an ellipsoid . . . 62
4.3 The relation between the surface normal vector projected on $\mathcal{L}$-plane and the vectors $\mathbf{u}$ and $\mathbf{v}$ . . . 68

6.1 Standard views of the standard test bodies . . . 88
6.2 Synthetic image of the test ellipsoid . . . 91
6.3 Synthetic image of the test hyperboloid . . . 92
6.4 Maps of the true local surface tilt and the true Gaussian curvature of the test surfaces . . . 93

6.5 The heuristic slant estimator error as a function of increasing Pentland's heuristic radius . . . 96
6.6 Slant estimates for the test ellipsoid . . . 97
6.7 Slant estimates for the test hyperboloid . . . 98
6.8 Recovered tilt maps of the test ellipsoid . . . 100
6.9 Recovered tilt maps of the test hyperboloid . . . 101
6.10 Recovered needle maps of the test ellipsoid . . . 103
6.11 Recovered needle maps of the test hyperboloid . . . 104
6.12 Gaussian-curvature maps recovered from the needle maps of the standard ellipsoid . . . 106
6.13 Gaussian-curvature maps recovered from the needle maps of the standard hyperboloid . . . 107
6.14 Tilt estimation error as a function of the test ellipsoid eccentricity . . . 110
6.15 The influence of non-Lambertianity on the Pentland tilt estimates . . . 113
6.16 The influence of non-Lambertianity on the proposed tilt estimator results 114
6.17 The influence of the illuminant slant on the local tilt estimates . . . 115
6.18 The influence of the illuminant slant on the local orientation estimates . 116
6.19 Convergence of the relaxation algorithm run on the ellipsoid image . . . 117
6.20 Convergence of the relaxation algorithm run on the hyperboloid image . 118
6.21 The test ellipsoid reconstructed from the needle maps . . . 119
6.22 The test hyperboloid reconstructed from the needle maps . . . 120
6.23 The two test surfaces representing developables . . . 123
6.24 Absolute error of the recovered tilt map of a slanted cylinder . . . 124
6.25 Absolute error of the recovered tilt map of a slanted cone . . . 125
6.26 Gaussian curvature map recovered directly and from the needle map . . . 127

List of Tables

6.1 Local surface slant recovery by different algorithms 99
6.2 Local surface tilt recovery by different algorithms 99
6.3 Gaussian curvature recovery from the needle map 105
6.4 Gaussian curvature recovered directly and from the needle map 126

List of Algorithms

1 Surface Normal and Sign of K at an Occlusion Boundary 36

2 Surface Normal at a Self-Shadow Boundary 38

3 Surface Orientation near Vertices 44

4 Local Orientation and Surface Curvature at an Apex 53

5 Local Shading Analysis for Spheres 79

6 Single-Step Local Shading Analysis for General Surfaces 81

7 Relaxation Local Shading Analysis for General Surfaces 85

Mathematical Notation

$\mathbf{u} \times \mathbf{v}$	vector product of the vectors $\mathbf{u}$ and $\mathbf{v}$
$\mathbf{u} \cdot \mathbf{v}$	dot product
$\\|\mathbf{x}\\| = \sqrt{\mathbf{x} \cdot \mathbf{x}}$	Euclidean norm of a vector
$\mathbf{x}_u = \frac{\partial \mathbf{x}}{\partial u}$	partial derivative (of a vector function $\mathbf{x}(u, v)$)
$\mathbf{x}_{\mathbf{u}} = D_{\mathbf{u}}\, \mathbf{x}$	directional derivative in a direction $\mathbf{u}$
	Imaging geometry
$\mathcal{L}$, $\mathcal{V}$	light plane and image plane
$\sigma_{\mathcal{L}}$, $\tau_{\mathcal{L}}$	slant and tilt of the illuminant
$\mathbf{N}$	unit normal vector
$\mathbf{L}$	unit light direction vector
$\mathbf{V}$	unit viewer vector
$\mathbf{H}$	unit specular direction vector
	Surface photometry
R	reflectance function
α	surface albedo
	Image
I	image function
∇I	image gradient vector
$\mathbf{H}_I$	image Hessian matrix
${\mathcal{F}_I}^n$	local image function properties of n-th order
ρ, ρ^*	isophote density and relative isophote density
γ	image gradient direction
κ_{F}, κ_{I}	flow line curvature and isophote curvature
ν	relative variation of the isophote density
$\mathbf{u}$, $\mathbf{v}$	vectors parallel to negative image gradient and to image isophote

Differential surface properties

E, F, G	first fundamental form coefficients
$\mathbf{G}$	metric tensor
L, M, N	second fundamental form coefficients
$\mathbf{B}$	second fundamental form matrix
β_i^j	Weingarten coefficients
$\mathbf{S}$	shape operator
κ_n	normal curvature
κ_1, κ_2	principal surface curvatures
K, H	Gaussian and mean curvatures
Γ_{ij}^k	Christoffel symbols of the second kind
σ, τ	surface normal slant and tilt

Projection

$\widetilde{\mathbf{x}}$, $\widehat{\mathbf{x}}$	projection of a vector $\mathbf{x}$ on the light plane and on an arbitrary image plane
$\widehat{\kappa}_t$	curvature of the occluding contour

Chapter 1

Introduction

The visual system of all higher vertebrates is a rich, remarkably reliable, and robust source of information about the surrounding world. It enables fast orientation in the world to control and navigate egomotion, fast recognition of the objects of interest (whose variability is enormous), and the evaluation of the own actions' impacts on the outer world. It can work under a surprisingly wide range of lighting conditions—among other capabilities—and it is able to infer *global* 2-D or 3-D *shape* from a variety of *local* shape cues: occlusion contour and surface discontinuities, specular reflection, smooth shading, surface texture, stereopsis, motion parallax, and maybe some others. Shape as an intermediate-level abstraction and co-identification of the viewed object is necessary for *geometric reasoning* and *recognition.* The second aspect of vision is *location and relative position estimation.* A biological visual system infers depth from a variety of depth cues, some of which are the same as the shape cues: stereopsis, focus, eyeballs convergence, relative motion, to name a few.

Although a lot of research on the psychology and physiology of vision has been carried out using the classical methodology of biological sciences, mathematical and technical research—not necessarily motivated by the results from the physiology of natural visual systems—can also put some light on the problem so that the progress in the technical understanding of the problem can stimulate the progress in biological research.

One such 'technical research goal' is to study *shading as a cue to 3-D surface shape.* The goal of this research, which is also the topic of this work, is to gain a better understanding of the problem how to make *local qualitative surface interpretations from single monocular grey-level images.* The technical applications are also of high interest, since the method enables 3-D surface reconstruction from such images, which is important when there is no suitable direct measuring method or existing devices cannot be modified (such as the scanning electron microscope), or when the object of interest is not within mechanical reach (such as the surface of the Moon).

Here we have to outline the border between the topic studied here and what is commonly called *shape from shading.* Shape from shading, as it is understood in the literature, is formulated as a broader problem of *recovering 3-D surface* in explicit *depth* *form* from grey-level image data. The ill-posedness of such a problem is usually overcome by more or less explicit *global* regularization and smoothness constraints. The problem was first formulated by Horn [Hor75]. On the other hand, it is generally accepted that *qualitative shape* information is of particular importance for the purpose of recognition [LF84, MN87, LF86, JH88, BPYA85, BJ86] and reasoning about the viewed object's functionality [Baj90]. The quality of 3-D shape is related to the concept of *surface curvature.* This intuitive notion is formalized by the differential geometry of surfaces. And, as will be seen later, a grey-level image contains information that is related to the *local surface orientation and curvature* rather than to its relative depth. We therefore see the question of how the pattern of apparent local surface shading constrains the differential shape to be more important for the recognition aspect of vision and we see the dense relative-depth surface representation recovery from intensity images as a rather technical problem with nice potential applications, but not as a fundamental problem of vision. This is the reason why we emphasize the *locality* of the method. The approach conforms the Marr's theory of vision [Mar82] that stresses "bottom-up" or "data-driven" processing over the "top-down" or "conceptually-driven" approach. Marr demonstrated that a great deal of processing of images can be achieved *without* recourse to specific world knowledge. On the other hand, the latter approach can be used at a higher level when the results of local processing are grouped into structures whose elements represent primitives that have well-defined geometric properties (occlusion contours, for instance).

It seems that—at least at an intermediate interpretation level—the structure should model the topological structure of the object rather than the logical structure, as it has been proposed in early work on structural modelling and recognition motivated by advances in Artificial Intelligence research. The logical structure cannot be found without an interpretation and the interpretation cannot be found without the knowledge about the structure, which causes combinatorial explosion of such an approach. The topological modelling approach, on the other hand, does not require any semantic information related to the visible scene. A necessary prerequisite for implementing this approach is a good understanding of the image formation, which is what our subsequent work aims to deal with [Šár94].

Local methods possess several advantages over global ones. The interpretation they produce is independent of the interpretation of the whole scene.[1] Thus, they do not depend on non-local events as occlusion, separability, or the relative position of objects within the scene. A great advantage of a local method is that it requires smaller computational effort and, in addition, it can be well parallelized. A local analysis of shading can also better cope with non-constant surface features, such as the surface albedo (the reflectivity constant). On the other hand, the local approach is more sensitive to noise because it cannot simply include any global smoothness constraint (some ways around are suggested in [FC88, FL92], for instance). Natural visual systems do not seem to be

[1]This is not fully true: local properties of an intensity image depend on such non-local events like the interreflection among surfaces.

in contradiction with this principle. A large percentage of neurons in the retina and the visual cortex is devoted to an initial *local* analysis of the image [BG85].

It will be seen in Chapter 4 that the local shading pattern can provide enough information for recovering the *qualitative type* of the viewed surface *without the need for reconstructing the whole surface*. The ambiguity of the recovered information can be modelled explicitly and reduced by a non-local method, making use of general additional information. The qualitative labelling can be used for grouping into qualitatively homogeneous primitives that may form a basis for structural description of the viewed surface and for recognition, for instance.

1.1 Methodology

An intensity image[2] is an image of surface radiance that is strongly related to the local orientation of the viewed surface. The imaging process maps points in 3-D space to the 2-D image plane while the image intensity values are related to the local surface orientation at each surface point. The change of image value along an elementary arc in the image will depend on the surface orientation change along a corresponding pre-image of the arc on surface. This orientation change is related to the surface curvature. The surface curvature map *and* the local orientation, together with a fixed (image) coordinate system, constitute a unique description of the surface up to a rigid motion in depth, as will be seen later. Thus, the local properties (derivatives) of image function are somehow related to the geometric properties of the surface with a degree of ambiguity that depends on the imaging process.

A suitable methodology for the purpose of recovering local geometric surface characteristics from single monocular intensity images is *Local Shading Analysis (LSA)*, introduced in the classical work of Pentland [Pen84]. *LSA* studies how the intensity image of a visible surface constrains its actual *differential* 3-D shape. It is based on the assumption that there exists a relation between the surface differential properties and the local properties of the image function. *LSA* can provide surface-related data for shape from shading or for higher level vision algorithms from single monocular 2-D intensity images.

A set of independent local image function properties of n-th order ${\mathcal{F}_I}^n$ is supposed to be related to the set of surface differential properties $\mathcal{G}_S(\mathcal{T},\, \beta)$ and to imaging geometry parameters $\mathcal{G}_I(\mathbf{V}, \mathbf{L})$ by a mapping

$$\mathcal{P}\colon\ \mathcal{G}_S \cup \mathcal{G}_I \mapsto \mathcal{F}_I^n\,.$$

The differential properties include the local orientation of the tangent plane $\mathcal{T}$ and the Weingarten coefficients β (and possibly their derivatives, [Sto69]). The imaging

[2]This term will be described formally in Chapter 2.

geometry is given by the light direction vector $\mathbf{L}$ and the viewing vector $\mathbf{V}$. The (ill-posed) inverse mapping $\mathcal{P}^{-1}$ is the subject of local shading analysis aiming to get $\mathcal{T}$ and β from ${\mathcal{F}_I}^n$ with minimum ambiguity assuming known imaging geometry parameters $\mathcal{G}_{\mathcal{I}}$. One of the key problems is a suitable choice of the set ${\mathcal{F}_I}^n$. The choice is crucial to the tractability of the problem. The statement that the image function contains the required information and the minimum degree of the solution ambiguity will be discussed in Chapter 4. A more specific definition of the *LSA* subject will be given in Chapter 5 after the necessary concepts have been introduced.

The goal is reachable only under certain constraints. Without any assumptions about the lighting conditions, surface properties, and geometry of the scene, there is no hope for recovering any relevant information about the viewed surfaces: we can take nearly any arbitrary smooth surface and paint it to give the same image as another surface without paint. If we do not want to involve other cues into the analyzed vision process, the shading analysis problem must be significantly simplified.

Also, shading information alone is not sufficient. Without the knowledge about the imaging geometry, only little can be said about the surface. The problem is well illustrated by the 'crater illusion:' when the direction of illumination of a surface like the one in Fig. 1.1 changes, the apparent convexity/concavity of the surface also changes, thereby demonstrating that the two are inseparably linked [Pen84].

Proposition 1.1 (Standard *LSA* Constraints[3])

- *There is only a single distant isotropic illuminant of known position.*
- *In between the illuminant and the illuminated surface there is a homogeneous transparent medium.*
- *The surface is uniform, no texture markings or varying albedo (including colour) are allowed.*
- *No other surface reflects light towards the analyzed surface.[4]*
- *The surface is viewed from a distant position.*

Violation of either of these constraints leads to a certain degree of misinterpretation of the image data: the less these constraints are fulfiled the more the recovered parameters of the viewed surface deviate from their actual values. This does not mean that the *LSA* method fails, however.

There are non-local methods that, under certain additional assumptions, can esti-

[3]We do not include such trivial constraints like one that only the illuminated part of the visible surface can be analyzed.

[4]Strictly speaking, this means that concave objects are excluded, since the surface may reflect light on itself in the concavities.

mate the likely illumination direction [LR85, Pen82]. The problem of estimating shape from local texture information has already received a lot of attention [Ken79, Wit81, Pen86, BS90, Len91]. The problem of mutual illuminations was explored e.g. by Forsyth & Zisserman [FZ91]. The problem of image structure near edges and shadow borders is discussed in [LZ87]. A separate research field is *shape from contours.* Occluding contours are an important cue to shape [BT81, Ste81, BY84, Koe84, PK90]. Reconstructing shape from intensity images in the presence of interreflections is studied in [NIK91a].

Definition 1.1 (Proper Illumination Condition)

We define a properly-illuminated surface patch as one that satisfies the constraints in Proposition 1.1.

1.2 Previous Work Related to *LSA*

We will follow important concepts developed in the course of investigating the problem of shading analysis. The references given here concern only the topics that are related to the subject of this thesis. We do not aim to present a complete overview of the whole shape from shading literature. The concepts that are introduced but not fully explained in this section will be explained later in the text. This is in order to keep the text short. We help the reader so that the terms that are going to be properly explained later appear in slanted typeface *like this*$^{(1.2)}$, while the superscript in braces refers to the section where they are explained.

The early work of Mackworth [Mac73] dealt with polyhedral scenes. The author observed that the image of a polyhedral edge is perpendicular to the dual line in gradient space under orthogonal projection. The dual line connects the points in gradient space that correspond to surface gradients (normals) of the adjoining planar faces forming the edge. Moreover, if the gradients appear on the dual line in the same order as the corresponding surfaces appear in the image then that edge is convex, but if they are ordered in the reverse direction then it is concave. The gradients representing the faces are constrained to form a polygon in gradient space dual to the images of edges meeting at the polyhedral vertex. Three degrees of freedom remain: the position and the scale of the polygon. Mackworth used this observation for edge classification and global image interpretation.

Mackworth's observation can also be used for local surface reconstruction near polyhedral vertices if the shading near the junctions is taken into account. This was developed in the work of Horn [Hor77]. If the *imaging geometry*$^{(2.2)}$ and the *reflectance function*$^{(2.2)}$ of the surface are known, the shading of faces forming the junction provides an additional and sufficient constraint for recovering orientation of these faces. The result can be generalized to vertices of non-polyhedral surfaces. We will discuss this in more detail in Chapter 3.

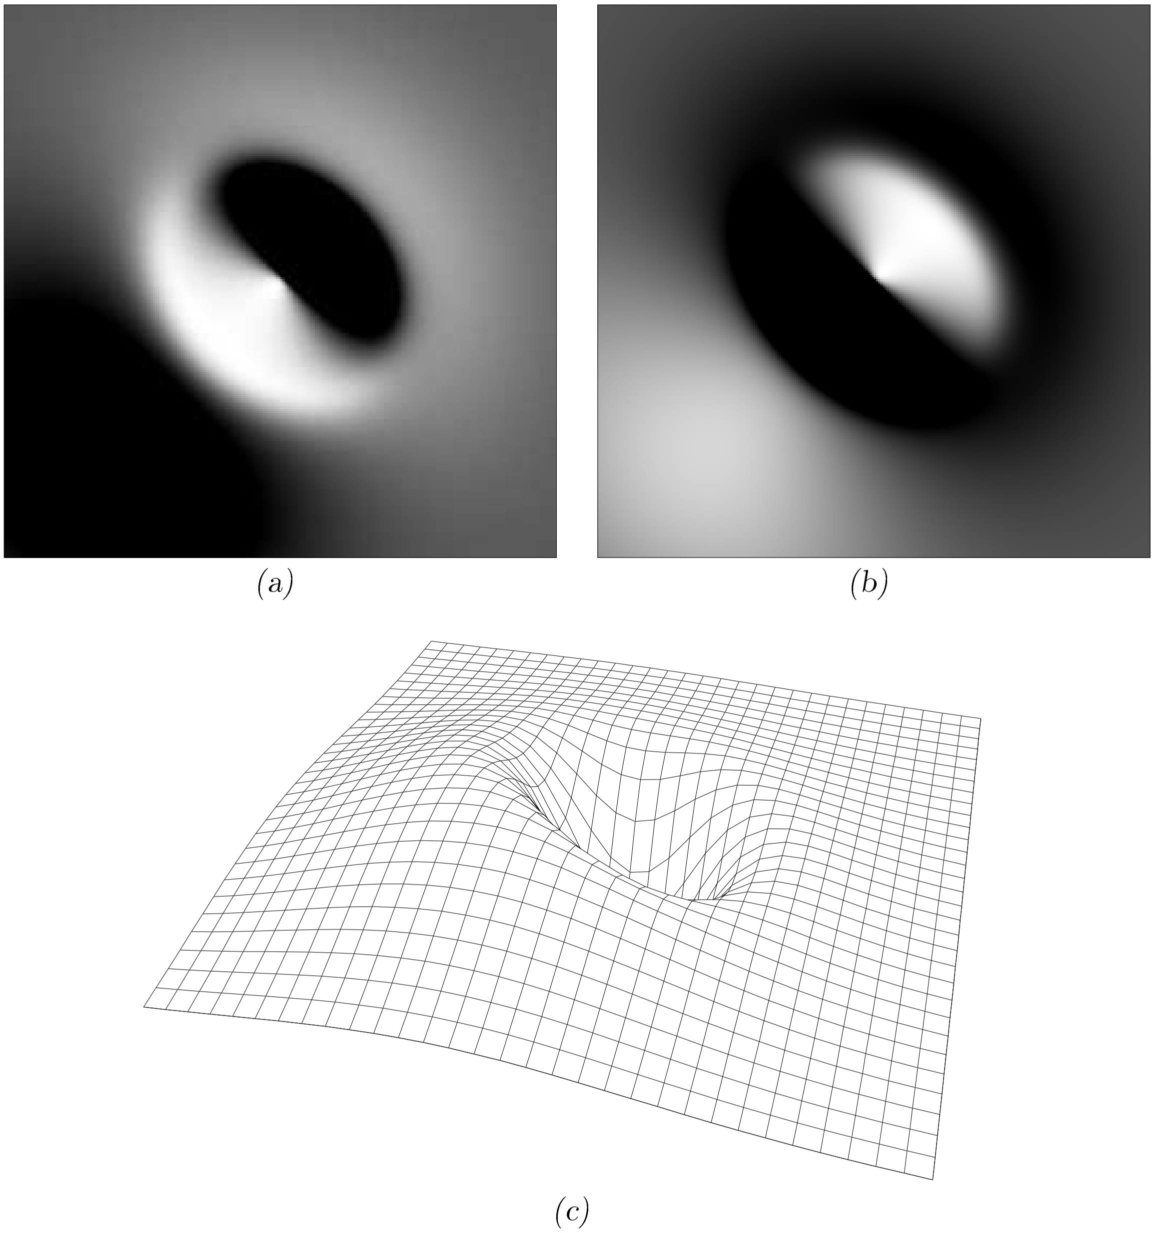

Figure 1.1: *The crater illusion: (a) is a surface that looks like a recession in a large hump illuminated by a sharp light; (b) is a surface looking like a hump inside a large hollow. Try turning the figure upside down. Now, the visual interpretation of (a) and (b) reverses (and the surfaces seem to be different). In fact, both (a) and (b) are images of the same asymmetric surface (c) illuminated from upper right and lower left, respectively. The illusion is due to our expectation that the light comes from the upper part of the visual field.*

The *photometric approach* to determining local orientation of a curvilinear surface from image brightness was first formulated by Horn [Hor75]. The shape from shading problem formulation has the form of a non-linear first-order partial differential equation in two unknowns.

Horn also introduced the concept of a reflectance map that characterizes the reflectance of a surface of a particular photometric type when the imaging geometry is known. The map represents the surface reflectance if the surface normal obtains a certain orientation (p, q) in gradient space, and it can be constructed from the known imaging geometry and the known photometric properties of the surface (from the so-called Bidirectional Reflectance Distribution Function [NRH+77, HS79]), or experimentally.

Horn and his colleagues [IH81] have analyzed the process of image formation and have developed several numerical integration schemes for using image intensity to recover the object's shape. The key structure appearing in the solution is the *characteristic strip.* It is an entity passing through any surface point of *known* surface orientation, allowing to find the orientation in its neighbourhood *in one direction.* Since a single characteristic strip passes through all points on a surface or there is a singularity— a source or a sink, the surface can be reconstructed by moving along these strips from points of known orientation. Suitable candidates for such starting points are the *singular image points*(3.1) with *maximum brightness,* since the surface normal of a (presupposed) *Lambertian surface*(2.2) is parallel to the light direction there, or the *specular points*(3.2) on specular surfaces, since the normal is parallel to the *specular direction*(2.2) there. Horn also stressed the importance of an *occlusion boundary*(3.1) which is a curve on the surface due to the surface rolling away from the viewer. The surface normal can be determined uniquely along these boundaries, since it is parallel to the image plane and oriented perpendicularly to the corresponding contour. This information can be propagated into a recovered surface patch from the boundary, provided that the surface is at least of class $\mathcal{C}^2$. Although the occlusion boundary constrains the surface orientation (and the Gaussian curvature) uniquely, it does not sufficiently constrain the solution in the explicit depth form [Oli91].

Horn et al. also discuss the properties of a *self-shadow boundary*(3.1), which is a locus of points on the surface where the tangent plane contains the light direction vector. Horn claims that the surface normal is undetermined with only a single degree of freedom there. We will see in Chapter 3 that this is not fully true, at least not for surfaces with space-invariant reflectance—in that case the orientation *can be determined uniquely* if the imaging geometry is known.

Woodham [Woo81] tried to relate the image gradient to the curvature of the underlying surface. If the surface is assumed in the form of a *Monge patch*(2.1) with $f(x, y)$ being the surface-generating function, the Hessian of f, $\mathbf{H} = [f_{ij}]$ is related to the image gradient $[I_x, I_y]$ by

$$\mathbf{H} = \begin{bmatrix} I_x & \mathrm{d}p \\ I_y & \mathrm{d}q \end{bmatrix} \begin{bmatrix} R_p & \mathrm{d}x \\ R_q & \mathrm{d}y \end{bmatrix}^{-1},$$

where [dp, dq] represents the change of surface orientation in gradient space corresponding to a small movement [dx, dy] in image, and R_p, R_q are the change of reflectance corresponding to a change in orientation. This equation has an infinite number of solutions. The Hessian is related to the surface curvature. When the properties of the surface are known, they can provide an additional constraint on $\mathbf{H}$. If a surface is known to be convex, the positive definiteness of $\mathbf{H}$ leads to the constraints $R_p I_x + R_q I_y \geq 0$ and $\mathrm{d}p\,\mathrm{d}x + \mathrm{d}q\,\mathrm{d}y \geq 0$, that constrain the solution. But there still remains an ambiguity.

Woodham noted that the singularity of $\mathbf{H}$ at a point on a developable surface implies that the brightness isocontour passing through that point becomes locally parallel to the direction of the surface ruling, i.e., the direction in which the surface normal does not change.

Woodham also proposed *photometric stereo* as a method that can recover surface orientation unambiguously, assuming a known reflectance function [Woo80]. The key idea of photometric stereo is to vary the direction of incident illumination between successive views, while holding the viewing direction.

Koenderink & van Doorn [Kv80] studied the field of constant image brightness contours (*isophotes*$^{(2.3)}$) and discussed some of their geometric features. The authors demonstrate that there are structural properties of the field of isophotes that bear an invariant relation to geometric properties of the object.

They show that the *Gaussian mapping*$^{(2.1)}$ maps the surface isophotes onto latitude circles on the unit sphere. The authors restrict themselves to Gaussian maps that are "stable" in the sense that their structure is not changed by small perturbations. In such a case there are two types of the Gaussian map singularities: *folds* (curves) and *cusps* (isolated points). The authors classify the extrema of the image function based on their behaviour under the displacement of either illuminant or the vantage point (or both). Not all singular points in the image come from extremal points of the surface. Koenderink & van Doorn show that there is a class of local image extrema that occur only on the pre-images of Gaussian map folds on surface (so called *parabolic lines*$^{(2.1)}$). A change of illuminant position causes these singularities to move over parabolic lines that are closed and nested curves on a closed $\mathcal{C}^2$ surface. The classification will be discussed in more detail in Chapter 3.

Haralick et al. [HWL83] define a set of features—*peaks, pits, ridges, ravines, saddles, flats,* and *hillsides*—in terms of the first and second image derivatives that can describe each point on the image as if it is a surface. The authors pointed out that the labelling is invariant to invertible image value transformations: peaks stay peaks, pits stay pits, and so forth. Pong et al. [PSH85] applied this representation to synthetic images of simple solids. The labelling produces patterns that are typical for certain surfaces. They propose a scheme for partial classification of 3-D object surfaces. They limit their consideration to five types of surfaces: *planar, developable, spherical, elliptical, and hyperbolic*$^{(2.1)}$. Since the structures formed by the labels (curves arcs and isolated points) are *sparse* there must be a knowledge about the patches in the image that likely come from the same surface patch. This limits the method.

The problem of invariance of the image function properties to invertible grey-level transformations is discussed in [FRKV91, RFK91]. The authors show that the isophote curvature κ and the flow-line (gradient curve) curvature μ are second-order invariants under the transformation. They also define a complete, irreducible set of 2-jet[5] invariants for the basic transformation group represented by spatial translation and rotation,

$$H = \{\gamma,\, \kappa,\, \lambda,\, \mu,\, \nu\}\,,$$

where λ is the image function value itself (invariant to the basic transformation group), γ is the image gradient magnitude (which is equivalent to the isophote density and invariant to the image value offset transform), and ν is the relative variation of the isophote density along the image flow-line (invariant to image value scaling). The invariance of certain image features to invertible grey-level transformations as well as the invariance to the reflectance function will be discussed in Section 4.1.

The fundamental contribution to *LSA* is the work of Pentland [Pen84]. He makes two heuristic assumptions: (1) the surface is restricted to be *spherical*, and (2) the reflectance function is restricted to be *Lambertian*. Pentland shows that, under such conditions, the local surface orientation, the surface albedo, and the light direction can be estimated up to the convex/concave case *without the knowledge of the illuminant position.* If the qualitative illuminant position[6] is known (like in the 'crater illusion' problem solution), the ambiguity can be resolved (as we have seen in Fig. 1.1). When these assumptions are violated, the resulting estimates degrade. In an extreme case, when $I_{xx}/I_{yy} < 0$, no solution exists. The author shows that this can happen on a parabolic surface.

Pentland proposes a classification of surface types based on second partial derivatives of the image function I. If there exists a locally rotated coordinate system x, y in the image, $\mathbf{H}_I = [I_{ij}]$ is the image Hessian, and

$$\begin{array}{rl}
\mathbf{H}_I = \mathbf{0} & \text{then the surface is } \textit{planar,} \\[1ex]
\mathbf{H}_I \neq \mathbf{0} \;\text{ and }\; \det \mathbf{H}_I = 0 & \text{then the surface is } \textit{cylindrical,} \\[1ex]
I_{xx}/I_{yy} < 0 & \text{then the surface is a } \textit{saddle,} \\[1ex]
\text{otherwise:} & \text{the surface type cannot be specified.}
\end{array}$$

This is also related to the above mentioned Haralick's topographic labelling.

[5] A local 2-jet is the equivalence class of images that locally coincide up to their second order derivatives, see [FRKV91].

[6] We mean namely the light direction in the visual field.

Later, Pentland proposed a unified approach to 3-D shape estimation that treats shaded and textured surfaces in a unified manner using the fractal model of the surface [Pen86].

Pentland observed that there are two natural components of the local surface orientation: the *slant* of the normal vector from the light direction, and the *tilt*$^{(4.2)}$ of the normal vector when projected onto a plane perpendicular to the light direction vector. He gives the solution for the unknown orientation in terms of those two parameters. The behaviour of Pentland's slant and tilt estimators will be discussed in Chapter 6.

The drawback of this method is that, when the two conditions mentioned above are violated, the recovered normals need not form a smooth (integrable) surface. Our experiments revealed that the recovered surface may significantly differ from the real surface. To compensate the lack of generality, Pentland proposed heuristic methods for surface orientation that are based on statistical methods [Pen86].

Frankot & Chellappa [FC88] suggest taking a set of non-integrable surface normals and constructing an integrable set by projecting the surface normals onto the nearest set of integrable normals, where the "nearest" is defined by a *global* integral distance measure defined in the gradient space coordinate system. The integrability is achieved when the derivatives satisfy $p_y = q_x$.

A different, more local approach is presented by Ferrie & Lagarde in [FL92]. They improve local orientation estimates so that the recovered surface curvature varies according to the expected local model of the surface. The model they choose is the *augmented Darboux frame*$^{(2.1)}$ $\mathcal{D}(P)$ that includes surface normal, principal directions, and principal curvatures at the surface point P. Such a model is invariant to the actual image coordinate system. The model is iteratively refined, based on functional minimizing the *local curvature consistency* defined as the distance of the current frame $\mathcal{D}(P)$ from all neighbouring frames $\mathcal{D}(P_\alpha)$ transported from P_α to P (by a *parallel transport* [Sto69]).

Lee & Rosenfeld [LR85] showed that the first-order image derivatives are sufficient to recover local orientation of a spherical surface *provided that the light direction is known.* They propose the image function gradient vector as a good tilt estimator. The authors argue that, *in the case of a sphere,* the image gradient vector is parallel to the projection of the surface normal $\mathbf{N}$ onto the plane perpendicular to the light vector $\mathbf{L}$. This vector is then related to the normal tilt in the viewer-centered coordinate system by a simple 3-D space transformation, whose parameters are known if the components of $\mathbf{L}$ are known in this system.

Ferrie & Levine [FL89] suggest that, under Lambertian conditions, the slant σ can be estimated as

$$\sigma = \arccos \frac{I}{\alpha},$$

where the proportionality constant α (surface albedo) can be either measured or determined. If we assume $\alpha = const$ throughout the image of a properly-illuminated surface patch, and if there is a point (u, v) on the surface such that $\mathbf{N}(u, v) = \mathbf{L}$, then there is a local maximum of the image function at (u, v) and $I(u, v) = I_{max}$. It follows that α can be estimated as $\alpha = I_{max}$ at a local image function maximum on a properly illuminated Lambertian surface patch.[7] This estimator will be discussed in more detail in Section 5.3.

Oliensis [Oli90] assumes that a neighbourhood of a surface point can be approximated by the *osculating paraboloid* at that point, $d^2\mathbf{x} \cdot \mathbf{N} = \mathrm{I\!I}$, where $\mathrm{I\!I}$ is the *second fundamental form*$^{(2.1)}$ of the surface. Then, the set of image derivatives up to second order and *known* light direction $\mathbf{L}$ suffice to estimate the surface normal parameters, the principal curvatures, and the surface albedo. If the vector $\mathbf{L}$ is unknown, it can be estimated by a suitable global technique, as proposed in [Pen84, LR85], for instance.

Arnspang [Arn91] relates the image function derivatives to local surface orientation and curvatures supposing the surface is in the form of second-order polynomial. His results do not differ qualitatively from other approaches.

Saxberg [Sax89] formulated the shape from shading problem as trajectories of a dynamic system in state-space. He showed that solution surfaces for the problem are invariant manifolds of the flow generated by the image dynamic system, and that the stable and unstable manifolds associated with certain critical points in the image play an important role in determining the solution surfaces. The role of critical points is of utmost importance. Saxberg showed that critical points of the image intensities due to critical points in the reflectance function give rise to critical points in the image dynamic system. The author showed that at such points there exist at least two and at most four local solutions for the surface consistent with the image.

Nayar et al. [NIK91a] show that, under Lambertian assumption, the shape of a concave object recovered from intensity data (*pseudo shape*) maps uniquely to the actual surface shape. This mapping is *illuminant invariant.* The authors propose a *global* relaxation algorithm recovering the shape and albedo from its pseudo shape. A drawback of this method is that the relaxation requires reconstruction of the surface in explicit depth form to evaluate the *interreflection kernel.*

1.3 Overview of the Thesis

The topological structure of surface events, like discontinuities and vertices, parabolic lines, planar points, etc., is important for surface modelling and recognition. The study of image formation, which we would like to contribute to, can reveal what geometric information *can be* recovered; although it does not answer the question of *how to model* the surface, it can answer the question *what kind of qualitative and structural information*

[7] Local maxima of the image function can also be located on images of parabolic lines on surface, where, in general, $\mathbf{N} \neq \mathbf{L}$ (see [Kv80]). In this case, $I \leq I_{max}$.

is available for modelling. Our research presented in this work concentrates on the topic of local image analysis by using the tool of differential geometry.

Our work should be a step forward in *LSA*. We do not use any kind of explicit local surface model. The tractability of the problem is re-gained by choosing physically relevant structures induced by the reflectance physics. We concentrate mainly on accurate estimation of surface normal tilt and on qualitatively correct Gaussian curvature recovery. We emphasize the role of image gradient as a direction related to the normal tilt. The true tilt differs from image gradient by an angle whose value depends on local image properties.

The goal of the thesis will be precised later as soon as the necessary concepts have been introduced in Chapter 2. This chapter aims to help the reader by giving a brief overview of the differential geometry. Section 2.2 presents the model of image formation that will be used throughout the thesis. The last section of this chapter introduces the concept of isophotes that will become important in Chapter 4.

Chapter 3 summarizes the problem of shading near singular image elements: namely, the occlusion and self-shadow boundaries and the singular image points. It will be shown that the local surface orientation and the sign of Gaussian curvature can be determined uniquely on the occlusion boundary. This is a known result. The new result is that the orientation on a self-shadow boundary can be determined uniquely as well, which will be discussed in Subsection 3.1.3. In Subsection 3.1.4 we present an algorithm that allows reconstruction of surface normals near vertices of a surface. Section 3.2 deals with image singularities: after a classification is given, shading near apices is studied and it is proven that the surface orientation can be determined uniquely at those points, while the signs of principal curvatures always remain undetermined.

The main theoretic result of this work is presented in Chapter 4. It is argued that isophotes are image structures whose properties are well-related to the underlying surface. Some basic properties of those curves will be discussed and Gaussian surface parameterization will be introduced. This parameterization offers a surface representation that is unique up to a proper Euclidean (rigid) motion and it is well-related to the image properties. In Section 4.3, the fundamental equation relating the local image properties to the differential properties of the viewed surface is derived and some of its features are discussed.

Based on the analysis given in Chapter 4, a series of local shading analysis algorithms are presented in Chapter 5. In Chapter 6, we then discuss results of detailed computational experiments with the aim of comparing the proposed *LSA* algorithms to some other *LSA* algorithms presented in the literature so far.

The final chapter gives conclusions and raises questions that remain unanswered in this thesis. It also proposes some possible but uninvestigated solutions to these problems.

Chapter 2

Basic Concepts

2.1 Differential Geometry of Surfaces

Differential geometry is a suitable framework allowing efficient study of an image function's local behaviour that is related to local geometric properties of the viewed continuous surface. We will briefly overview some of the basic concepts here. For a similar review, the reader can see [BJ86] or any textbook on standard differential geometry (e.g., [Sto69, Lip69]).

Concept of a Surface

When we speak about a surface, we mean *surface in the small,* i.e., only in some neighbourhood of one of its points. Surfaces *in the large* can be built up as a set of overlapping *patches,* each of which is a surface in the small.

A *regular surface patch of class* $\mathcal{C}^m$ $(m \geq 1)$ is a locus of points S in Euclidean three-space E^3 defined by the end points of a two-parametric vector $\mathbf{x} = \mathbf{x}(u, v)$ mapping an open set U in the uv plane onto S, so that

1. $\mathbf{x}$ is of class $\mathcal{C}^m$ in U.

2. The *coordinate vectors*[1]

$$\mathbf{x}_u = \frac{\partial \mathbf{x}(u, v)}{\partial u}, \quad \mathbf{x}_v = \frac{\partial \mathbf{x}(u, v)}{\partial v}$$

   are linearly independent, i.e.[2],

$$\mathbf{x}_u \times \mathbf{x}_v \neq \mathbf{0}\,.$$

[1] Subscripts are used to denote partial derivatives.

[2] The operator '$\times$' represents a cross-product of two vectors.

The last condition states only that there must be *unit normal* defined over the patch S:

$$\mathbf{N} = \frac{\mathbf{x}_u \times \mathbf{x}_v}{\|\mathbf{x}_u \times \mathbf{x}_v\|}.$$

We usually speak of the parameters u, v as determining a curvilinear coordinate system on the surface, and the curves $u = \text{const.}$, $v = \text{const.}$ corresponding to a net of regular space curves covering S. The vectors $\mathbf{x}_u$ and $\mathbf{x}_v$ are then tangent vectors to the *parametric curves* and they locally span the surface *tangent plane.* The normal $\mathbf{N}$ is then perpendicular to the tangent plane and determines its orientation uniquely.

Any surface patch can be locally represented by a Monge patch, such that

$$\mathbf{x}(u, v) = [u, \;\; v, \;\; f(u, v)]^{\mathrm{T}},$$

where f is the surface generating function. This may be viewed so that the surface is simply a 3-D "plot" of $f(u, v)$. This representation satisfies the conditions given above and the normal is

$$\mathbf{N} = \frac{[-f_u, \;\; -f_v, \;\; 1]^{\mathrm{T}}}{\sqrt{1 + f_u^2 + f_v^2}}.$$

First and Second Fundamental Forms

The length of a curve $\mathbf{x}(t) = \mathbf{x}(u(t), v(t))$ on the surface is[3]

$$s(t) = \int_{t_0}^{t} \|\mathbf{x}'(\tau)\| \, d\tau = \int_{t_0}^{t} \sqrt{\mathbf{x}'(\tau) \cdot \mathbf{x}'(\tau)} \, d\tau = \int_{t_0}^{t} \sqrt{\mathrm{I}(\tau)} \, d\tau \, ,$$

where

$$\mathrm{I} = \left(\frac{ds}{dt}\right)^2 = \mathbf{x}' \cdot \mathbf{x}' = E\left(\frac{du}{dt}\right)^2 + 2\,F\,\frac{du}{dt}\frac{dv}{dt} + G\left(\frac{dv}{dt}\right)^2$$

is a positive definite quadratic form called *first fundamental form.* We may arrange the coefficients to form a *metric tensor*

$$\mathbf{G} = \begin{bmatrix} E & F \\ F & G \end{bmatrix} = \begin{bmatrix} \mathbf{x}_u \cdot \mathbf{x}_u & \mathbf{x}_u \cdot \mathbf{x}_v \\ \mathbf{x}_v \cdot \mathbf{x}_u & \mathbf{x}_v \cdot \mathbf{x}_v \end{bmatrix}.$$

Since I is positive definite, the coefficients must satisfy

$$E > 0, \quad G > 0, \quad E\,G - F^2 > 0\,. \tag{2.1}$$

The first fundamental form (not the coefficients themselves) is invariant to a change in the surface parameterization and to translations and rotations of the surface. It is thus called an *intrinsic* property of the surface.

[3]The operator '·' represents the dot-product of two vectors.

If R is a region on a patch, its area A is

$$A = \iint_W \|\mathbf{x}_u \times \mathbf{x}_v\| \, du \, dv = \iint_W \sqrt{E\,G - F^2} \, du \, dv \,,$$

where W is the set of points in the parameter space mapping onto R. The area is also an intrinsic property.

Since the angle ϕ between the parametric curves is

$$\cos\phi = \frac{\mathbf{x}_u \cdot \mathbf{x}_v}{\|\mathbf{x}_u\| \, \|\mathbf{x}_v\|} = \cdots = \frac{F}{\sqrt{E\,G}} \,,$$

the curvilinear coordinates are orthogonal if and only if $F = 0$.

In contrast, the *second fundamental form* of a surface is dependent on the embedding of the surface in 3-D space and is therefore considered an *extrinsic* property of the surface. This form models the deviation of the surface from its tangent plane in the neighbourhood of the point of tangency P. We can imagine that there is a local coordinate system in P, defined so that its three independent base vectors are parallel to the vectors $\mathbf{x}_u$, $\mathbf{x}_v$, and $\mathbf{N}$ (we may imagine $\mathbf{N}$ as parallel to the usual "depth" or "z-coordinate"). If

$$Q\colon\ \mathbf{x}(u + du,\, v + dv) = \mathbf{x}(u,\, v) + d\mathbf{x} + \frac{1}{2}\, d^2\mathbf{x} + \mathbf{o}(du^2 + dv^2)$$

is another point in the neighbourhood of $P\colon \mathbf{x}(u,\, v)$ on the surface, its "z-coordinate" relative to Q will be given by the "projection" on the normal $\mathbf{N}$:

$$\Big(\mathbf{x}(u + du,\, v + dv) - \mathbf{x}(u,\, v)\Big) \cdot \mathbf{N} = \frac{1}{2} d^2\mathbf{x} \cdot \mathbf{N} + o(du^2 + dv^2) \,,$$

since $d\mathbf{x} \cdot \mathbf{N} = 0$. The quadratic form

$$\mathrm{II} = d^2\mathbf{x} \cdot \mathbf{N} = L\, du^2 + 2\, M\, du\, dv + N\, dv^2$$

is the second fundamental form. The coefficients may be arranged to form a matrix in a similar way as for the first fundamental form

$$\mathbf{B} = \begin{bmatrix} L & M \\ M & N \end{bmatrix} = \begin{bmatrix} \mathbf{x}_{uu} \cdot \mathbf{N} & \mathbf{x}_{uv} \cdot \mathbf{N} \\ \mathbf{x}_{vu} \cdot \mathbf{N} & \mathbf{x}_{vv} \cdot \mathbf{N} \end{bmatrix} = \begin{bmatrix} -\mathbf{x}_u \cdot \mathbf{N}_u & -\mathbf{x}_u \cdot \mathbf{N}_v \\ -\mathbf{x}_v \cdot \mathbf{N}_u & -\mathbf{x}_v \cdot \mathbf{N}_v \end{bmatrix}.$$

The rightmost matrix coefficients are obtained from the middle matrix and from the conditions $\mathbf{x}_i \cdot \mathbf{N} \equiv 0$ that lead to $(\mathbf{x}_i \cdot \mathbf{N})_j = 0, \quad i, j = u, v$.

The second fundamental form generates so called *osculating paraboloid.* It is a second-order surface that "has the best touch with the surface $\mathbf{x}(u,\, v)$" at P. The nature of this paraboloid determines the qualitative nature of the surface. It is appropriate to distinguish four cases, depending upon the sign of the determinant $L\,N - M^2$ of the second fundamental form II:

$$
\begin{aligned}
\text{elliptic case:} \quad & L\,N - M^2 > 0, \\
\text{hyperbolic case:} \quad & L\,N - M^2 < 0, \\
\text{parabolic case:} \quad & L\,N - M^2 = 0, \\
\text{planar case:} \quad & L = M = N = 0.
\end{aligned}
$$

The points on a general surface, where $L\,N - M^2 = 0$, are called *parabolic points.* They form a system of closed and nested curves on a $\mathcal{C}^2$ closed surface. Parabolic lines divide the surface into hyperbolic and elliptic regions.

Normal Curvature

The second fundamental form measures somehow the turn $d\mathbf{N}$ of the normal vector $\mathbf{N}$ after a change $d\mathbf{x}$ of the position on the surface $\mathbf{x}$, while the first fundamental form measures the length of the arc corresponding to $d\mathbf{x}$. The ratio of

$$
\kappa_n = \frac{\mathrm{II}(du,\, dv)}{\mathrm{I}(du,\, dv)}
$$

defined for a particular direction $du : dv$ on surface is called *normal curvature.* All curves through a point P on the surface, which are tangent to the same line through P, have the same normal curvature at P, since κ_n depends only upon the ratio du/dv, i.e., the direction of the tangent line at P.

There are two perpendicular directions in which the normal curvature values take on maximum and minimum κ_1 and κ_2, or the normal curvature is equal in all directions. These directions are called *principal* and have the property that the change $d\mathbf{N}$ of the normal in such a direction $d\mathbf{x}$ is

$$
d\mathbf{N} = -\kappa_i\, d\mathbf{x}\,,
$$

where $i = 1$ or $i = 2$. This relation is called *Rodrigues's formula* and the curvatures κ_1, κ_2 are called *principal curvatures.* There exists a surface parameterization, where the parametric curves are parallel to principal directions at each point on the surface. This parameterization always exists and the parameter curves are called *lines of curvature.* The u- and v-parameter curves are lines of curvature if and only if $F = M = 0$ at every point on the patch.

The principal directions and the normal vector define a unique moving orthogonal coordinate frame on the surface. This frame, together with the principal curvatures, form a so-called *augmented Darboux frame* [FL92].

Conjugated Directions

If u- and v-parameter curves are lines of curvature, it holds

$$
\mathbf{x}_u \cdot \mathbf{N}_v = 0 \quad \text{and} \quad \mathbf{x}_v \cdot \mathbf{N}_u = 0\,.
$$

A general direction $\delta u : \delta v$ at a point on the surface is said to be *conjugate* to a direction $du : dv$ if

$$d\mathbf{x} \cdot \delta\mathbf{N} = 0\,.$$

At an elliptic or hyperbolic point on a surface, every direction has a unique conjugate direction. At a parabolic point on a surface, every direction is conjugate to the zero-curvature principal direction. At a planar point, each two directions are conjugate. The direction that is self-conjugate is called *asymptotic direction.*

Gaussian Mapping

The normal vectors $\mathbf{N}$ of a surface patch S of sufficiently high class can be displaced to the origin of the global coordinate system. The locus of their endpoints is then a point set Ω on the unit sphere G and it is called the *Gaussian image* of the surface S. Parabolic lines on the surface map onto *folds* in Gaussian image (the Jacobian of the mapping vanishes there). The mapping is therefore not globally $1:1$. A linearized version of the Gaussian mapping expresses the relation between vectors in the tangent plane:

$$\begin{aligned} \mathbf{N}_u &= -\beta_u^u\,\mathbf{x}_u - \beta_u^v\,\mathbf{x}_v\,, \\ \mathbf{N}_v &= -\beta_v^u\,\mathbf{x}_u - \beta_v^v\,\mathbf{x}_v\,. \end{aligned}$$

These are called *Weingarten equations.* The *Weingarten coefficients* may be arranged into a matrix called *shape operator* $\mathbf{S}$ that is related to the matrices of the first and the second fundamental forms

$$\mathbf{S} = \begin{bmatrix} \beta_u^u & \beta_u^v \\ \beta_v^u & \beta_v^v \end{bmatrix} = \left(\mathbf{G}^{-1}\,\mathbf{B}\right)^{\mathrm{T}}.$$

Natural algebraic invariants of the shape operator—the determinant and the trace—are called Gaussian and mean curvatures:

$$\begin{aligned} K &= \det\mathbf{S}\,, \\ H &= \frac{1}{2}\,\mathrm{tr}\,\mathbf{S}\,. \end{aligned}$$

If the parametric curves are lines of curvature, then $\beta_u^v = \beta_v^u = 0$ and the curvatures become

$$\begin{aligned} K &= \kappa_1\,\kappa_2\,, \\ H &= \frac{\kappa_1 + \kappa_2}{2}\,. \end{aligned}$$

Existence Theorem

There is another system of similar equations, called the *Gauss equations*:

$$\begin{aligned}
\mathbf{x}_{uu} &= \Gamma_{11}^1\,\mathbf{x}_u + \Gamma_{11}^2\,\mathbf{x}_v + L\,\mathbf{N}\,,\\
\mathbf{x}_{uv} &= \Gamma_{12}^1\,\mathbf{x}_u + \Gamma_{12}^2\,\mathbf{x}_v + M\,\mathbf{N}\,,\\
\mathbf{x}_{vv} &= \Gamma_{22}^1\,\mathbf{x}_u + \Gamma_{22}^2\,\mathbf{x}_v + N\,\mathbf{N}\,,
\end{aligned}$$

where the quantities Γ_{ij}^k are called *Christoffel symbols of the second kind.* If we evaluate Γ_{ij}^k, we see that they depend only on the first fundamental coefficients and their derivatives.

Theorem 2.1 (Fundamental Theorem of Surfaces)

> *Let E, F, G be functions of at least class $\mathcal{C}^2$ and L, M, N be functions of at least class $\mathcal{C}^1$. Assume all six functions are defined on an open neighbourhood U of the point $(u_0,\, v_0)$. If all six functions satisfy the sign restrictions* (2.1) *and the following* compatibility equations,
>
> $$\begin{aligned}
L_v - M_u &= L\,\Gamma_{12}^1 + M\,(\Gamma_{12}^2 - \Gamma_{11}^1) - N\,\Gamma_{11}^2\,,\\
M_v - N_u &= L\,\Gamma_{22}^1 + M\,(\Gamma_{22}^2 - \Gamma_{12}^1) - N\,\Gamma_{12}^2\,,\\
L\,N - M^2 &= F\,\Big[(\Gamma_{22}^2)_u - (\Gamma_{12}^2)_v + \Gamma_{22}^1\,\Gamma_{11}^2 - \Gamma_{12}^1\,\Gamma_{12}^2\Big]\\
&+\ E\,\Big[(\Gamma_{22}^1)_u - (\Gamma_{12}^1)_v\ + \Gamma_{22}^1\,(\Gamma_{11}^1 - \Gamma_{12}^2) + \Gamma_{12}^1\,(\Gamma_{22}^2 - \Gamma_{12}^1)\Big]\,,
\end{aligned}$$
>
> *then there exists a unique surface patch $\mathbf{x}(u,\, v)$ of class $\mathcal{C}^3$, defined in U such that E, F, G and L, M, N are the first and second fundamental form coefficients, respectively. Uniqueness is determined up to a proper Euclidean motion.*

The first two compatibility equations are referred to as the *Mainardi-Codazzi equations.* The third equation states that the determinant of the second fundamental form depends only on the coefficients of the first fundamental form and its derivatives. It follows that the Gaussian curvature is a function of only E, F, G and their derivatives and is thus an intrinsic property of the surface. This is known as the *Theorema Egregium of Gauss.*

The existence theorem has an important corollary: *an arbitrary smooth surface shape is captured by six scalar two-parametric functions E, F, G, L, M, N.*

2.2 Model of Image Formation

Before the local shading on the viewed surface can be studied, a mathematical model of the image formation process must be developed. This is the aim of this section. We start with the concepts of surface reflectance and introduce a commonly accepted model for it. We then introduce some simplifications of the model for the purpose of tractability of the shading analysis that will be carried out in Chapter 4. Then, the image projection will be discussed and it will be argued that the orthogonal projection together with certain assumptions about the global imaging geometry model the image projection well.

2.2.1 Surface Reflectance

From the psychology of vision (which is confirmed by the technology of light-sensitive elements, of course) is known that the elementary perception of a particular illuminated surface point "brightness" is proportional to the physical quantity known as light flux per unit area, which is nothing but the amount of the corresponding light energy absorbed by the sensor. This energy is reflected by the surface in the direction of the elementary sensor and it is a portion of the total energy reflected by the illuminated surface element to the whole half-space. If L denotes the surface element *radiance* towards the sensor, the *response* from the sensor is

$$I = \int_{\lambda} s(\lambda)\, L(\lambda)\, d\lambda\,,$$

where λ is the light wavelength and $s(\lambda)$ is the spectral sensitivity of the sensor. The wavelength is related to colour perception, the quantity I to intensity of the perception. Let $\Phi_i(\lambda)$ denote the flux per unit area of the light beam striking the surface element and $\Phi_r(\lambda)$ the flux per unit area scattered by the surface element back to the *whole half-space.* The ratio

$$\rho(\lambda) = \frac{\Phi_r(\lambda)}{\Phi_i(\lambda)}$$

is called *reflectance coefficient* or, briefly, albedo. It defines the portion of incident energy that is reflected back to the half-space. It is an integral of the surface radiance over the solid angle ω representing the half-space:

$$\Phi_r = \int_{\omega} L\, d\omega\,.$$

We introduce a *reflectance function* R such that

$$\int_{\omega} R\, d\omega = 1\,,$$

where $d\omega$ is an infinitesimal solid angle around the viewing direction. This function models the spatial spread of the reflected light. The sensor response can be rewritten to

$$I = \int_{\lambda} s(\lambda)\, \frac{d\Phi_r(\lambda)}{d\omega}\, d\lambda = \int_{\lambda} s(\lambda)\, \rho(\lambda)\, R\, \Phi_i(\lambda)\, d\lambda = R \int_{\lambda} s(\lambda)\, \rho(\lambda)\, \Phi_i(\lambda)\, d\lambda\,,$$

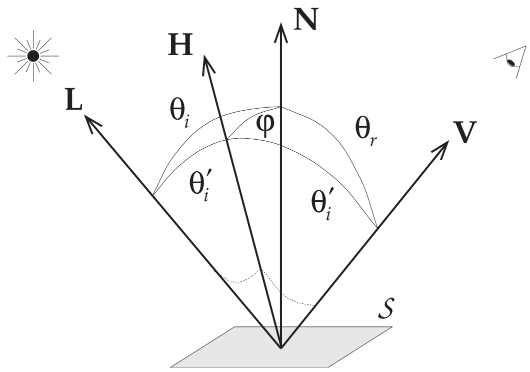

Figure 2.1: *Local imaging geometry: principal unit vectors whose mutual orientation at a surface facet S determine the radiance in viewing direction* $\mathbf{V}$. θ_i, θ_r, *and* φ *represent the angle of light incidence, the scattering angle, and the slope of micro-facets that reflect the light to viewing direction, respectively.* $\mathbf{L}$, $\mathbf{N}$, *and* $\mathbf{H}$ *are the light direction vector, the surface normal vector, and the specular direction vector, respectively.* $\mathbf{L}$, $\mathbf{V}$, *and* $\mathbf{H}$ *lie in the principal plane,* $\mathbf{H} = (\mathbf{L}+\mathbf{V})/\|\mathbf{L}+\mathbf{V}\|$.

It is supposed that the reflectance function models the influence of the local imaging geometry on the sensor response and it is thus invariant to the light wavelength λ. Color properties of the surface will be modeled by the reflectance coefficient $\rho(\lambda)$.[4]

The proper illumination condition, given on page 5 in Section 1.1, states that the surface reflectance coefficient is to be constant, as well as the spatial colour and intensity distributions of the light source. Therefore, the integral at the very right side of the preceding equation can be considered constant:

$$\int_\lambda s(\lambda)\, \rho(\lambda)\, \Phi_i(\lambda)\, d\lambda = s\, \rho\, \Phi_i = \alpha = \text{const}\,. \tag{2.2}$$

We assume, without loss of generality, that the product $s\,\Phi_i$ is unit and we call α *albedo* in what follows. We hope there will be no confusion with the term introduced above for the quantity ρ.[5]

From the principles of electro-magnetic wave reflection on a solid opaque surfaces follows that the surface radiance in a given spatial direction depends on the local surface tangent

[4]This need not be true for real surfaces: the colour properties of the scattered light need not be isotropic. Nevertheless, this is a reasonable simplification of the properties of most of the non-specular surfaces that appear in a natural environment.

[5]The reason for doing this is motivated by compatibility with other authors who implicitly make the same assumption about the reflectance function and albedo. They often use the symbol ρ for what has been introduced as α here. We use ρ for a different purpose.

plane orientation, on the micro-structure of the surface, and on the spatial distribution of the incident light [NIK91b]. The *multi-lobed reflectance model* has been widely accepted as a reasonable approximation of the reflection law on surfaces that appear in our world. This model can be derived from Maxwell's equations under certain restricting conditions[6] and it is well-confirmed by empiric observations. According to this model, the surface reflectance has three basic components: *a diffuse lobe, a specular lobe,* and *a specular spike* [NIK91b] (see also Fig. 2.1):

$$R = C_{dl}\,\cos\theta_i + C_{sl}\,\frac{G(\theta_i,\theta_r,\varphi)}{\cos\theta_r}\,e^{-\frac{1}{2}\left(\frac{\varphi}{\sigma_\varphi}\right)^2} + C_{ss}\,\delta(\theta_i-\theta_r)\,\delta(\varphi)\,,$$

where R is the reflectance function, C_{dl}, C_{sl}, and C_{ss} are the strengths of the three basic components, and $G(\theta_i,\theta_r,\varphi)$ is a *geometrical attenuation factor* that models the light masking and interception by neighbouring surface micro-facets:

$$G(\theta_i,\theta_r,\varphi) = \min\left(1,\ 2\,\frac{\cos\varphi\,\cos\theta_r}{\cos\theta_i'},\ 2\,\frac{\cos\varphi\,\cos\theta_i}{\cos\theta_i'}\right).$$

A detailed derivation of $G(\theta_i,\theta_r,\varphi)$ is given in [Td88].

In addition, some authors assume also *the backscatter lobe* which is a component

$$C_{bs}\,G(\theta_i,\theta_r,\varphi)\,\frac{p(\theta_i')}{\cos\theta_i\,\cos\theta_r}$$

due to $p(\theta_i')$ which is the phase distribution function of micro-facets located in a plane parallel to the mean surface [TP81]. The distribution of micro-facets (that are supposed to reflect specularly) is a key factor influencing the formation of all the lobes in the model. A good and brief introduction to the micro-facet surface model can be found in [NIK91b].

In general, the surface reflectance function is a scalar function of the following three dot-products:

$$R = R(\mathbf{N}\cdot\mathbf{L},\ \mathbf{N}\cdot\mathbf{V},\ \mathbf{V}\cdot\mathbf{L})\,. \tag{2.3}$$

If this function becomes independent of the viewing direction $\mathbf{V}$, it is called *space-invariant.* This means that the brightness of a particular visible surface point will not change under the change of the viewing direction. An example is the *Lambertian reflectance function:*

$$R = \frac{1}{\pi}\,\mathbf{N}\cdot\mathbf{L}\,.$$

Because of its simplicity, the Lambertian reflectance function has been widely accepted as a reasonable non-specular reflectance model for shape from shading.

[6] These restrictions are necessary to make the problem tractable.

The irradiance of an infinitely small light sensor located at position (x_1, x_2) in the image plane is equivalent to the surface radiance at (u, v)[7], under the assumption that there is no light attenuation in the optical medium between the surface and the sensor. The equation

$$I(x_1, x_2) = \alpha(u, v)\, R(\mathbf{N}(u, v) \cdot \mathbf{L},\, \mathbf{N}(u, v) \cdot \mathbf{V},\, \mathbf{V} \cdot \mathbf{L}) \tag{2.4}$$

is called *image irradiance equation* and $I(x_1, x_2)$ is called the *intensity image.* The scalar proportionality function $\alpha(u, v)$ represents the surface albedo.

There are two natural components of the image irradiance equation that form the image: (1) the reflectance function R and (2) the albedo α. Varying albedo gives rise to a visible texture on the surface (we can speak of texture only when the spatial frequencies of $\alpha(u, v)$ are sufficiently high), while the reflectance function gives rise to a (slowly varying) surface shading. The image formed by the reflectance component can be called *radiance map.*[8]

LSA studies local properties of radiance maps aiming to recover local geometric properties of the underlying surface. When we speak of an image in the following text, we will usually mean a radiance map.

The geometric relation between the location (u, v) on the surface where the particular shading pattern is observed and the corresponding location (x_1, x_2) in the image is given by a *projection* onto the image plane. The following subsection focuses on this issue.

2.2.2 Image Projection and Natural Coordinate Systems

An image is a composition of reflectance (mapping local surface orientation to a scalar brightness value) and image projection (mapping points on the viewed surface onto points on imaging surface S_I that need not necessarily be planar). Saxberg [Sax89] suggests decomposing the process of image projection $\mathbb{R}^3 \to S_I$ into two components: (1) the projection from $\mathbb{R}^3$ space to a set of (non-intersecting) rays running from points on the viewed surface to points on S_I, and (2) the intersection of the bundle of rays with S_I. This could allow analysis of the projection independent of the actual imaging surface, since it is not clear how this surface should look like: should it be a (probably deformed) sphere like in the mammalian retina or a plane like in a camera? Since the brightness of an image point would depend on the local orientation of the general imaging surface and we are going to study just the image brightness pattern in what follows, we restrict S_I to plane and speak about an *image plane.*

It has widely been accepted that a reasonable projection map is one corresponding to *perspective projection onto a plane.* A device doing this is called a pinhole camera

[7] u and v are the surface parameters.

[8] This term should not be confused with the concept of a *reflectance map,* which is a representation of the reflectance function in gradient space.

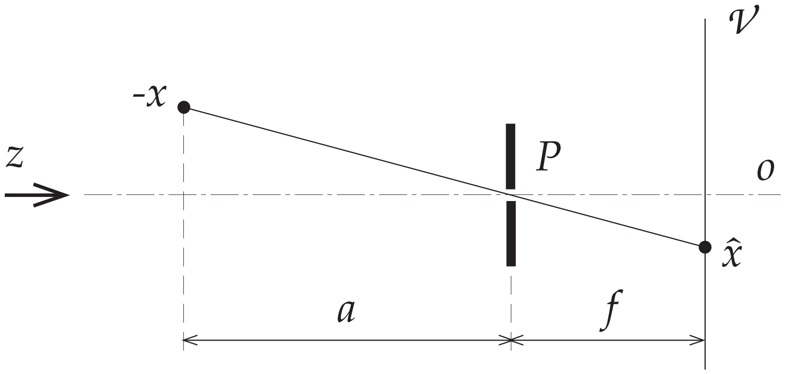

Figure 2.2: *The projection plane of a pinhole camera. A projection ray passing through the pinhole P connects the 3-D point with its image projected onto the image plane $\mathcal{V}$, which is a focal length f apart from P and perpendicular to the optical axis o.*

(see Fig. 2.2). A point $\mathbf{x}$ in space is mapped to $\hat{\mathbf{x}}$ in the image plane $\mathcal{V}$, such that the straight ray connecting $\mathbf{x}$ and $\hat{\mathbf{x}}$ passes through the pinhole P. The spatial points $\mathbf{x}$, $\hat{\mathbf{x}}$, and P form a plane containing the optical axis of the camera. We will call it a *projection plane*. If a denotes the perpendicular distance of $\mathbf{x}$ from the pinhole P in the projection plane, f the distance of the image plane from P, and x and $\hat{x}$ the distances of $\mathbf{x}$ and $\hat{\mathbf{x}}$ from the camera axis, it is simply found that

$$\hat{x} = -\frac{f}{a}\,x\,. \tag{2.5}$$

Since the viewing direction $\mathbf{V}$ (introduced in the previous subsection) is parallel to the camera's optical axis o, the projection mapping (2.5) can be rewritten into the full 3-D form as follows:

$$\hat{\mathbf{x}} = \left[-\frac{f}{a}\,x,\ \ -\frac{f}{a}\,y,\ \ f+a\right]^{\mathrm{T}} = -f\,\frac{\mathbf{x}-(\mathbf{x}\cdot\mathbf{V})\,\mathbf{V}}{\mathbf{x}\cdot\mathbf{V}} + (f+\mathbf{x}\cdot\mathbf{V})\,\mathbf{V}\,. \tag{2.6}$$

Real cameras approximate this model with a reasonable accuracy if the viewing distance a is much larger than the focal length f, $|f| \ll |a|$, which is almost always satisfied in reality [Hav56]. But the model can be simplified under this assumption. Suppose $\mathbf{x} = \mathbf{x}_0 + \Delta\mathbf{x}$, where $|\mathbf{x}_0\cdot\mathbf{V}| \gg |\Delta\mathbf{x}\cdot\mathbf{V}|$ and $\mathbf{x}_0$ lie on the optical axis, henceforth $\mathbf{x}_0 - (\mathbf{x}_0\cdot\mathbf{V})\,\mathbf{V} = 0$. Then

$$\begin{aligned}
\hat{\mathbf{x}} &= -f\,\frac{\mathbf{x}_0+\Delta\mathbf{x}-[(\mathbf{x}_0+\Delta\mathbf{x})\cdot\mathbf{V}]\,\mathbf{V}}{(\mathbf{x}_0+\Delta\mathbf{x})\,\mathbf{V}} + [f+(\mathbf{x}_0+\Delta\mathbf{x})\cdot\mathbf{V}]\,\mathbf{V}\\
&= -f\,\frac{\Delta\mathbf{x}-(\Delta\mathbf{x}\cdot\mathbf{V})\,\mathbf{V}}{\mathbf{x}_0\cdot\mathbf{V}+\Delta\mathbf{x}\cdot\mathbf{V}} + \mathbf{x}_0 + [f+\Delta\mathbf{x}\cdot\mathbf{V}]\,\mathbf{V}\,,
\end{aligned}$$

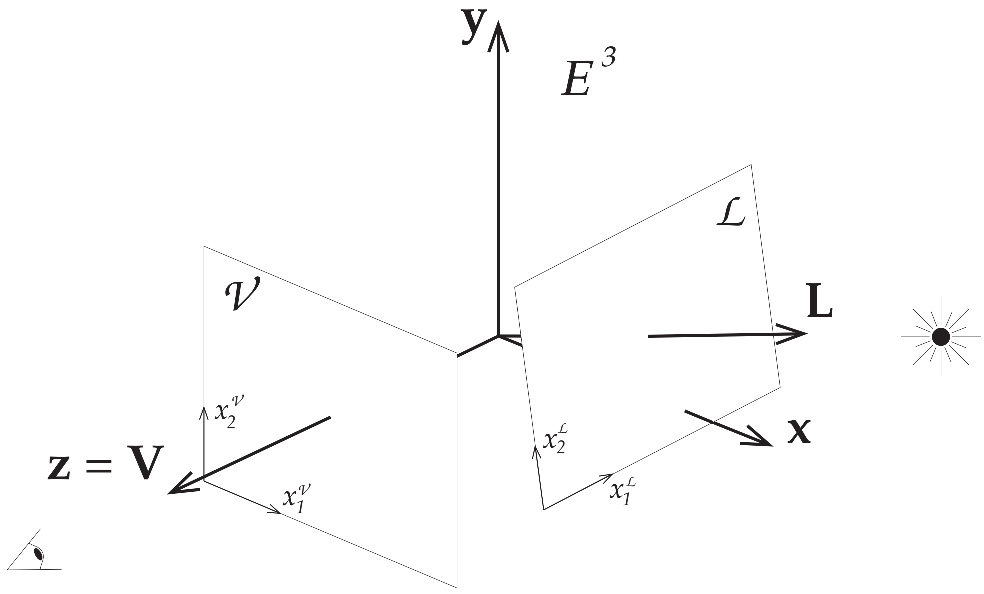

Figure 2.3: *Global imaging geometry. $\mathcal{V}$ is the image plane, $\mathcal{L}$ is the light plane. The principal plane is spanned by the viewing and light direction vectors* $\mathbf{V}$ *and* $\mathbf{L}$*. The principal vectors of local imaging geometry correspond to the vectors of the global imaging geometry, assuming both the illuminant and the viewer are infinitely distant from the illuminated surface.*

which becomes

$$\hat{\mathbf{x}} \approx \frac{-f}{\mathbf{x}_0 \cdot \mathbf{V}} \left(\Delta\mathbf{x} - \left(\Delta\mathbf{x} \cdot \mathbf{V} \right) \mathbf{V} \right) + \mathbf{x}_0 = m \left(\Delta\mathbf{x} - \left(\Delta\mathbf{x} \cdot \mathbf{V} \right) \mathbf{V} \right) + \mathbf{x}_0$$

for $|\mathbf{x}_0 \cdot \mathbf{V}| \gg |\Delta\mathbf{x} \cdot \mathbf{V}|$, where m can be thought as a magnification factor and the term $\Delta\mathbf{x} - (\Delta\mathbf{x} \cdot \mathbf{V})\,\mathbf{V}$ represents the *orthogonal projection* onto the image plane. The constant term $\mathbf{x}_0$ can be omitted or, equivalently, the image plane can be moved along the optical axis to the coordinates origin. Therefore, the projection of an object assumed to be far enough from the camera is considered orthogonal in what follows (if not explicitly stated otherwise). The x-axis of the image coordinate system is chosen so that the magnification factor m can be assumed *positive and unit.*

The viewing direction vector $\mathbf{V}$ can be considered constant since it is parallel to the optical axis o. If the illuminant is sufficiently far from the illuminated object, also the light direction $\mathbf{L}$ can be considered constant. We therefore define a *global imaging geometry* in terms of $\mathbf{L}$ and $\mathbf{V}$ given in a 3-D Euclidean space (see Fig. 2.3).

There exists a special position of the illuminant relative to the viewer—when they are placed near together. In this case $\mathbf{V} = \mathbf{L}$ and the $\mathcal{V}$-plane becomes perpendicular to the light direction vector. If we rotate $\mathbf{L}$ from this position, we can also rotate the associated perpendicular plane and assume there is still an 'imaginary image' projected onto it. This plane is called the *light plane* or, briefly, $\mathcal{L}$-plane. It will be seen later that the image projected onto the $\mathcal{L}$-plane has certain interesting properties related to the differential properties of the viewed surface.

Let us define a camera coordinate system in 3-D space so that the z-coordinate (depth) runs along the camera axis (with positive direction towards the camera), such that the x, y coordinates map to image coordinates $x_1^{\mathcal{V}}$, $x_2^{\mathcal{V}}$. This system is called a *viewer-centered coordinate system,* or, briefly, $\mathcal{V}$-system.

The illumination vector can be expressed in the $\mathcal{V}$-system as

$$\mathbf{L} = [\sin\sigma_{\mathcal{L}}\cos\tau_{\mathcal{L}},\ \ \sin\sigma_{\mathcal{L}}\sin\tau_{\mathcal{L}},\ \ \cos\sigma_{\mathcal{L}}]^{\mathrm{T}}\ ,$$

where $\sigma_{\mathcal{L}}$ is the vector slant from the viewing direction $\mathbf{V}$, and $\tau_{\mathcal{L}}$ is the tilt in the $\mathcal{V}$-plane, respectively (see also Fig. 4.1 on page 61).

The $\mathcal{L}$-plane setup generates another natural coordinate system: its depth axis contains the vector $\mathbf{L}$ and its (perpendicular) $x_1^{\mathcal{L}}$ and $x_2^{\mathcal{L}}$ axes are oriented so that the $x_1^{\mathcal{L}}$-axis has the same tilt in the $\mathcal{V}$-system as the tilt of the vector $\mathbf{L}$ (see Fig. 2.3). The $\mathcal{V}$-system and the $\mathcal{L}$-system are related by a constant linear mapping $T\colon \mathbf{x}^{\mathcal{L}} = \mathbf{T}\,\mathbf{x}^{\mathcal{V}}$ with a transformation matrix $\mathbf{T}$ [LR85]:

$$\mathbf{T} = \begin{bmatrix} \cos\sigma_{\mathcal{L}}\cos\tau_{\mathcal{L}} & \cos\sigma_{\mathcal{L}}\sin\tau_{\mathcal{L}} & -\sin\sigma_{\mathcal{L}} \\ -\sin\tau_{\mathcal{L}} & \cos\tau_{\mathcal{L}} & 0 \\ \sin\sigma_{\mathcal{L}}\cos\tau_{\mathcal{L}} & \sin\sigma_{\mathcal{L}}\sin\tau_{\mathcal{L}} & \cos\sigma_{\mathcal{L}} \end{bmatrix}. \tag{2.7}$$

The transformation is equivalent to a counterclockwise rotation of the image coordinate system around the vector $\mathbf{V}$ by the tilt angle $\tau_{\mathcal{L}}$, followed by a counterclockwise rotation of the new system by the slant angle $\sigma_{\mathcal{L}}$ around the transformed y-axis. This setup and the corresponding transformation have been proposed by Lee & Rosenfeld, who studied the mapping of the gradient of a *spherical surface* image projected on the $\mathcal{V}$-plane onto the gradient of the corresponding 'imaginary image' projected on the $\mathcal{L}$-plane [LR85].

Because of the properties of the image projected on the $\mathcal{L}$-plane, we will study this case in particular. We will therefore assume that $\mathbf{L} = \mathbf{V}$ throughout large parts of the following text (namely in Chapters 4 and 5). This restriction does not reduce the generality of the achieved results because, by the constant imaging geometry assumption, the transformation between the two systems is constant and known, and the image properties in the $\mathcal{V}$-system can be transformed into the $\mathcal{L}$-system.

2.3 Isophotes

In this section, general properties of certain structures induced on the surface by the reflectance physics will be overviewed. Later, in Chapter 4, the relation of local properties of the corresponding image structures to the differential properties of viewed surface will be discussed.

Let us suppose there is an intensity image of a $\mathcal{C}^2$ surface. *Image isophotes* constitute a family of non-intersecting curves in the image plane. These curves are equivalent to the *brightness isocontours* characterized (1) *globally* as curves along which

$$I(x_1,\, x_2) = I\,,$$

where I *identifies* the curve, and (2) *locally* by its tangent direction $\mathbf{v}$, such that the first directional derivative of the image function vanishes in this direction:

$$D_{\mathbf{v}}\, I = 0\,.$$

Surface isophotes are brightness isocontours located on the viewed surface. *The image projection maps the surface isophotes on the image isophotes,* because the projection transformation does not influence the intensity *value* (it only maps the locations on surface onto locations in the image plane). One of the reasons why we distinguish the surface and the image isophotes is that these two types of curves have different spatial geometry (the surface isophotes are spatial curves, while the image isophotes are planar).

The surface isophotes change when the illuminant changes its position or character, they may change with the movement of the viewer (depending on whether the surface reflectance function is space invariant or not). They are therefore induced (mainly) by the factor we have called global imaging geometry. Since inter-reflections can occur among the surface patches of non-convex bodies or among different bodies in the scene, this is also a co-factor that induces the isophotes. To avoid the problems with mutual illumination, we have to accept the simplified model introduced in Section 1.1, requiring that the surface is *properly illuminated.* This is—after all—a standard assumption for all methods that recover shape from a single image of light pattern induced on a surface by a (controlled) illumination.[9]

The image isophotes constitute a system of non-intersecting curves that are closed within the image function domain if they are sufficiently close to a non-saddle image singularity. The global topological structure of the system is governed by several singular elements (assuming a single body with $\mathcal{C}^2$ surface in the visible scene):

- the singular image points (i.e., the locations where the image gradient vanishes),
- the geometric contours (occlusion contours or surface discontinuities);
- the shading contours (shadow-edges or texture marking boundaries);

[9] A way around is proposed in [NIK91a], for instance. But this method is not local any more.

- the points on an occlusion contour corresponding to far and near points in the Gaussian map of the occlusion boundary;[10]
- the points of an ending contour.[11]

Some of these structures will be studied in Chapter 3. The global topological structure will not be studied here. For a good insight see [Koe90].

Another, closely related system of curves is the system of image gradient curves, which we call *flow-lines.* The flow-lines are perpendicular to isophotes at each point—their properties are therefore related to the properties of isophotes. They are locally characterized by a unit tangent vector $\mathbf{u}$ that is parallel to the image gradient

$$\mathbf{u} = -\frac{\nabla I}{\|\nabla I\|}\,.$$

The image flow-lines *are generally not* a projection of surface flow-lines. Nevertheless, the image flow-lines can be back-projected onto the surface forming a system of curves defined up to a certain ambiguity. It will be shown later that this ambiguity is related to the unrecoverable convexity versus concavity of the viewed surface.

Both systems of surface curves can be chosen as parametric curves of the surface. It is known from standard differential geometry that two systems of intersecting curves define any regular surface uniquely. The surface curves are related to the corresponding curves in the intensity image by a (perspective) projection. *A suitable, sufficiently rich set of their features will therefore constrain the differential properties of the surface.*

This concept will be elaborated upon in the following chapters. The novelty of this approach over the standard approaches is that it uses physically relevant structures for expressing the relation $surface \mapsto image$. It will be seen that the main advantage of the approach is that it makes the problem of theoretical *LSA* tractable *without imposing any a priori constraints on the differential properties or the order of the viewed surface.*

[10]The Gaussian map of the contour is non-planar, thus showing points "far" and "near" with respect to the north pole of the sphere: the Gaussian image just touches a sphere parallel there. The terminology is taken from [Koe90].

[11]An ending contour can appear in the image of a hyperbolic surface at the point located on an occlusion boundary viewed just from the asymptotic direction, see [Koe90] for details.

Chapter 3

Shading Analysis near Boundaries and Singular Points

In this chapter, we investigate local shading near singular image elements which are not covered by the model introduced in Chapter 4. The aims of this chapter are the following: (1) it should overview and explain some of the basic concepts appearing in previous work on analysis of local brightness distribution, (2) it should provide a background on concepts developed in Chapters 4 and 5 in more detail, and (3) it supports the overall aim to deal with the topic of local shading systematically.

The main theoretic results of this chapter appear in Section 3.1.3, where the uniqueness of recovered local surface orientation along a self-shadow boundary is proved, and in Section 3.2.3, where the existence of at most four solutions for the principal surface curvatures near a singular image point of the first type is proved.

The text is structured into two relatively independent parallel streams: (1) the main body dealing with the analysis of specific phenomena, and (2) a series of 'algorithms' summarizing the theoretical results and presenting them in a structured, brief, and well-arranged form that can provide the basis of an implementation. The two streams contain different kinds of information, but either of the streams can be read without reading the other.

3.1 Occlusions, Self-Shadow Boundaries and Surface Discontinuities

A *boundary* is basically a curve on a surface that divides it locally into two adjoining parts. In the image, the boundary projects to a curve which we call a *contour.* Boundaries are loci of certain local surface events. We distinguish between a *geometric boundary* and a *shading boundary,* according to its nature.

There are two types of a geometric boundary:

an occlusion boundary, which is a curve on the surface due to its smooth rolling away from the viewer, thus splitting the surface into visible and invisible parts (it is sometimes called 'self-occlusion boundary' or 'rim'),

a surface discontinuity, which is a locus of points on the surface where it is locally not smooth, i.e., of the order $\mathcal{C}^{\infty}$. We say there is a surface discontinuity of the order n if it is locally of $\mathcal{C}^{n}$: zero-order discontinuities are those in depth, first-order discontinuities are those in local surface orientation. Higher-order discontinuities are not assumed in this text.

The presence of a self-occlusion contour in the image also implies that a different body in the background may be occluded. We will not distinguish between an occlusion boundary and a self-occlusion boundary, since the difference cannot be recognized *without a global image interpretation.*

There are three types of a shading boundary:

a cast-shadow boundary, which is a locus of points forming the boundary between illuminated and non-illuminated areas due to a (partial) light source occlusion by a body in space,

a self-shadow boundary, which is an edge on the surface where one of the light sources illuminating the surface just sets below the local horizon,

a boundary of a texture marking, which is formed by a rapid change in surface albedo.

The following subsections focus on occlusion boundaries, self-shadow boundaries, and surface discontinuity boundaries. We will not study such cases as boundaries due to texture or colour change. Certain properties of local brightness pattern in the vicinity of cast-shadow boundaries, as well as the boundaries of texture markings and some other discontinuities have been studied in [LZ87].

3.1.1 Isophotes and Surface Orientation Near an Occlusion Boundary

An occlusion boundary is a general curve $\mathbf{o}(t)$ formed by points on the surface where the rays leaving a given vantage point just graze the surface. Its tangent vector $\mathbf{o}'(t)$ lies on the surface tangent plane and it is not necessarily perpendicular to the viewing direction $\mathbf{V}$, because an occlusion boundary is generally *not* a planar curve. Consider an arbitrary surface curve $\mathbf{y}(s)$ crossing the occlusion boundary $\mathbf{o}(t)$ at an angle α (see Fig. 3.1a). Its tangent vector $\mathbf{y}'(s)$ must lie on the surface tangent plane as well. By

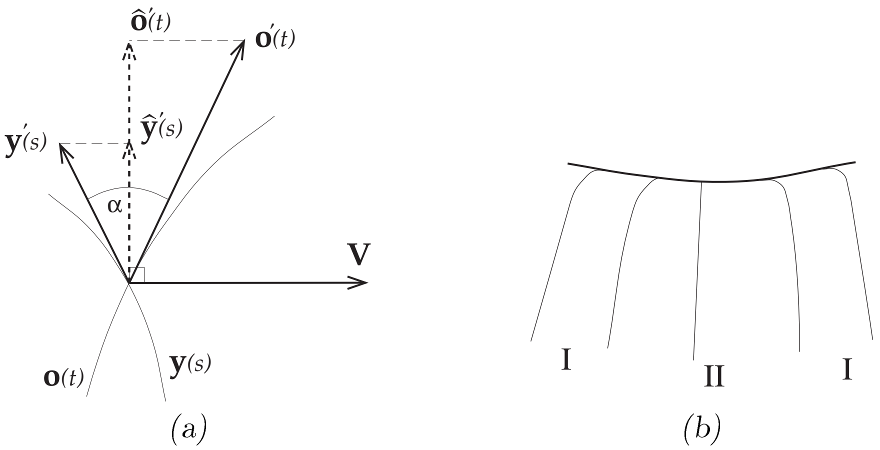

Figure 3.1: *Surface tangent plane at a point on an occlusion boundary (a). The tangent vector of a surface curve* $\mathbf{y}(s)$ *crossing the occlusion boundary* $\mathbf{o}(t)$ *at an angle* α *projects to a vector* $\hat{\mathbf{y}}'(s)$ *that is parallel to the occlusion contour tangent vector* $\hat{\mathbf{o}}'(t)$*. Projection of surface curves approaching the occlusion boundary (b). Images of arbitrary surface curves are almost always tangent to the contour* (I)*. The singularity* (II) *occurs when they are exactly perpendicular. (b) is adapted from [Sax89].*

the definition of the occlusion boundary, the tangent plane also contains the viewing direction vector $\mathbf{V}$. The projection $\hat{\mathbf{y}}'(s)$ onto the $\mathcal{V}$-plane is either perpendicular to $\mathbf{V}$ or a zero-length vector. Let us exclude the second case for a moment. Then, since both $\mathbf{V}$ and $\hat{\mathbf{y}}'(s)$ are perpendicular and lie on the tangent plane, as does $\hat{\mathbf{o}}'(t)$, $\hat{\mathbf{y}}'(s)$ and $\hat{\mathbf{o}}'(t)$ must appear parallel in the image. This confirms the analysis given in [Sax89].

Thus any surface curve crossing the occlusion boundary at an angle α *projects to a curve locally parallel to the corresponding occlusion contour. An exception is the case projecting to a curve crossing the contour perpendicularly (see Fig. 3.1b). The occlusion contour is thus an envelope of all images of surface curves approaching it non-perpendicularly. The exception projects to a singular point on the contour.*

Also surface isophotes can be considered as surface curves. This has two important consequences:

- isophotes are always parallel to an occlusion contour in its vicinity,
- isophotes are almost never parallel to a surface discontinuity contour.

This provides a clue about whether an image contour is due to an occlusion or due to a surface discontinuity. Although we did not perform systematic experiments, we believe that the decision about the type of the boundary based on the above property can be implemented with a sufficient robustness. The image gradient vector at the occlusion boundary—although tending to be of infinite magnitude—has nearly always a correct direction even when the image is blurred. Consider a locally straight occlusion contour, for example. It may be oriented vertically without loss of generality (see Fig. 3.2). Suppose the visible part of the surface approaches the contour from the right side. Then, if the background (left side of the contour) is uniform, the gradient operator will yield a correct orientation because the surface shading must be symmetric in the vicinity of the contour (isophotes are parallel to the contour). This is due to the response of the operator's y-component, which vanishes. When the support field becomes large, the surface shading need not be symmetric within the field. Also, the uniformity or the symmetry of the shading on the left side of the contour is not guaranteed. It seems that the response will not be correct in those cases. But the gradient magnitude is equivalent to the isophote density, which approaches infinity due to projective foreshortening near the occlusion contour. This causes the contribution of the left part of the support field to the total operator's response becomes relatively small. Also, the influence of the shading non-symmetry near the contour decreases because the effective part of the operator field becomes very narrow, oriented along the edge (see Fig. 3.2).

Blurring the input image is equivalent to the extension of the support field of the gradient operator: the gradient of the blurred image function can be written as

$$\nabla(G * I) = (\nabla G) * I$$

because of the properties of the convolution '$*$'. Extension of the Gaussian blurring kernel G thus gives rise to a more blurred image, which is equivalent to extending the support field of the gradient operator applied to the unblurred image.

The angle between the blurred image gradient and the local tangent to a geometric contour is therefore an important cue to the type of the underlying boundary. As the angle value gets closer to $\pi/2$, the more likely the contour is caused by an occlusion.

An occlusion contour is a locus of image points where the true 3-D underlying surface orientation *can be determined uniquely.* The image gradient direction plays a key role here: Let $\mathbf{o}(t)$ be the occlusion boundary curve. The negative image gradient vector $-\nabla I$ is perpendicular to the corresponding occlusion boundary tangent vector $\widehat{\mathbf{o}}'(t)$, since it is always perpendicular to the isophote direction at a point, and the isophote is parallel to the contour (see Fig. 3.3). This means that $-\nabla I$ *is parallel to the surface normal* $\mathbf{N}$, because both $-\nabla I$ and $\mathbf{N}$ are perpendicular to the tangent plane spanned by the vectors $\mathbf{V}$ and $\mathbf{o}'(t)$.

The analysis given so far assumed no surface texture, but it is valid even if the surface becomes textured: the texture projected onto the image plane is highly anisotropic due to projective foreshortening. The direction of a maximum texture variation becomes oriented perpendicularly to the occlusion contour *independent of the isotropy of the surface texture.* Thus, whatever has been said about the smooth shading near the

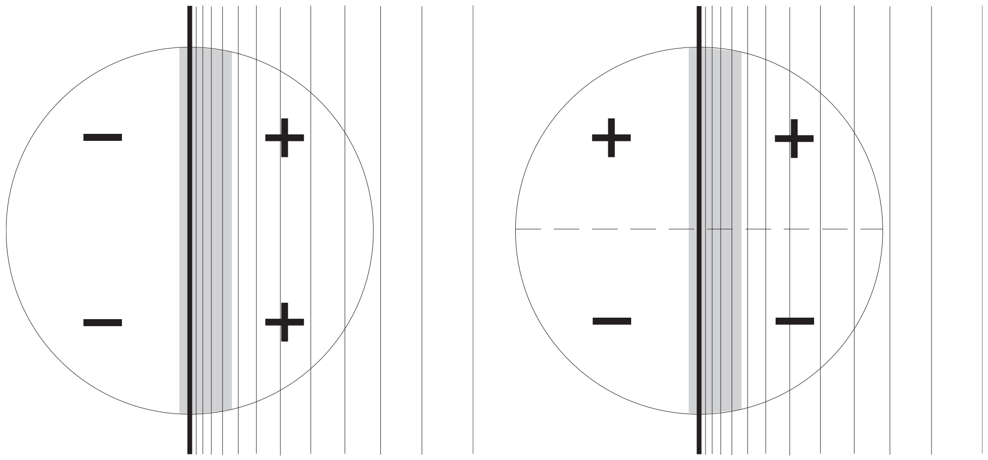

Figure 3.2: *Two independent components x and y of an image gradient operator determine the operator's behaviour on an occlusion contour. The (locally straight) contour divides the symmetric circular support field into two regions with equal area. The structure of the field is similar to the structure of receptive fields of simple cortical cells found in the mammalian visual cortex [BG85]. The part of the support field that becomes effective near the occlusion contour is shadowed. If the edge slants from the vertical orientation, the operator's response will not change.*

occlusion contour also holds for a textured surface: the shading may be viewed as a limiting case of a texture. In real, discrete images the variation in texture cannot approach infinity, however. The reason is that the spatial period of the changes becomes shorter than the sampling period near the occlusion. This limits the applicability of the image gradient as a local surface orientation detector at the occlusion contour.

It is known from the psychology of vision that there is a kind of simple receptive fields in the mammalian visual cortex having excitatory and inhibitory areas, whose structure is similar to the structure of the gradient operator [BG85]. Psychology also suggests that the occlusion boundary is crucial for 3-D structure acquisition because its properties are strongly related to the qualitative properties of the surface. This will be elaborated upon in the next subsection. Thus, the existence of a robust occlusion contour characteristics detector plays a substantial role in any robust visual system (see Fig. 3.4).

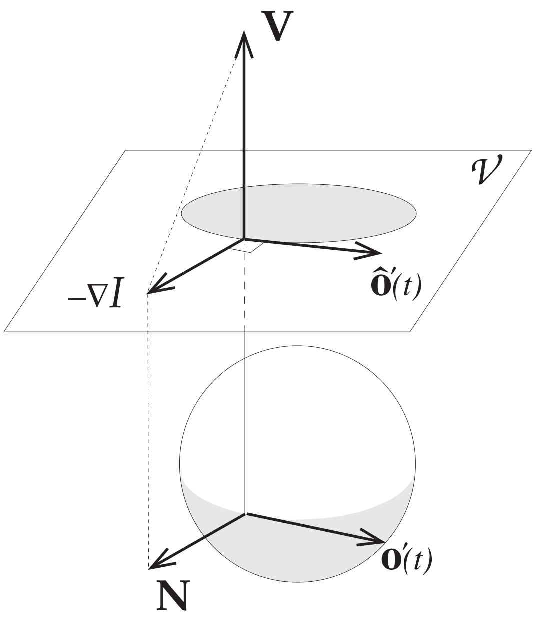

Figure 3.3: *The surface normal vector* **N** *at a point on an occlusion boundary* $\mathbf{o}(t)$*, the viewing direction* **V***, and the image gradient* $-\nabla I$ *at the corresponding image contour point all lie in a common plane.*

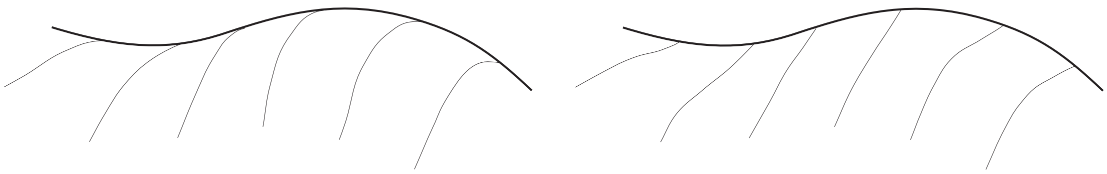

"a well drawn occlusion" *"a badly drawn occlusion"*

Figure 3.4: *Occlusion boundary perception. The left drawing is interpreted as an image of an occlusion boundary by your visual system. The right one is probably not because the contour is not an envelope of the surface curves images.*

3.1.2 The Sign of the Gaussian Curvature on an Occlusion Boundary

The sign of the Gaussian curvature K is an important surface feature because it is equivalent to the quality of the surface: hyperbolic surfaces have negative K, parabolic surfaces have vanishing K, and elliptic surfaces have positive K. The surface quality is important for geometric reasoning about the objects in the viewed scene. We will see that the occlusion contour is a rich source of this qualitative information because the quality of the contour (taken as a curve) is equivalent to the quality of the surface. No other element of an image of a smooth surface can probably yield this qualitative

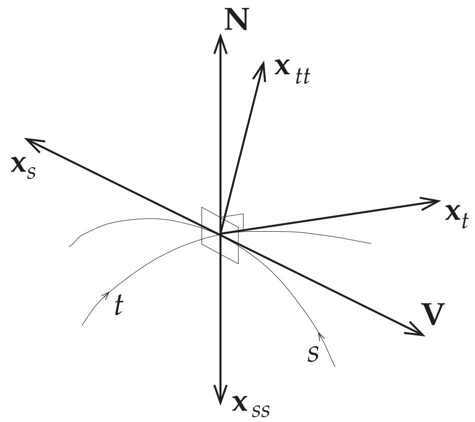

Figure 3.5: *The basic vectors related to a special parameterization* $\mathbf{x} = \mathbf{x}(s,t)$ *near the occlusion boundary.* $\mathbf{N}$, $\mathbf{x}_s$, $\mathbf{x}_{ss}$, *and* $\mathbf{V}$ *lie in the same plane and they are pairwise perpendicular. The angle between* $\mathbf{x}_{tt}$ *and* $\mathbf{V}$ *is arbitrary.*

information unambiguously.

As has been noted by several authors, the quality of a surface near its occlusion boundary can be determined uniquely [Koe84, BPYA85]. In the following text, we will follow the idea of the proof given by Brady et al. in [BPYA85], who prove that the sign of the Gaussian curvature on an occlusion boundary is equal to the curvature sign of the corresponding occlusion contour (image curve).

Let the occlusion boundary curve $\mathbf{o}(t)$ be a surface parameter curve. t is then a surface parameter. Let the other parametric curve on the surface have its osculating plane oriented so that it is spanned by the surface normal vector $\mathbf{N}$ and the viewing vector $\mathbf{V}$. This curve will be parameterized by a parameter s. The surface near the occlusion boundary is then represented as a parametric strip $\mathbf{x} = \mathbf{x}(s,t)$, and the tangent vectors to the parametric curves are $\mathbf{x}_s$ and $\mathbf{x}_t$, respectively (see also Fig. 3.5). The parameterization is chosen so that the tangent vectors have unit length *at the point in question.* This does not affect the generality of the result but it simplifies the subsequent formulas. We will also assume that the orientation of the s-parameter curve is chosen so that the curvature of this curve is non-negative. The properties of the parameterization are as follows:

property 1: $\mathbf{x}_{st} = 0$, since $\mathbf{x}_s$ does not change along the boundary, being always parallel to the viewing vector $\mathbf{V}$,

property 2: $\mathbf{x}_s \cdot \mathbf{x}_s = \mathbf{x}_t \cdot \mathbf{x}_t = 1$.

The Gaussian curvature is equivalent to the determinant of the shape operator

$$K = \det \mathbf{S} = \det(\mathbf{G}^{-1}\,\mathbf{B})^{\mathrm{T}},$$

where, due to property 2,

$$\mathbf{G} = \begin{bmatrix} \mathbf{x}_s \cdot \mathbf{x}_s & \mathbf{x}_s \cdot \mathbf{x}_t \\ \mathbf{x}_t \cdot \mathbf{x}_s & \mathbf{x}_t \cdot \mathbf{x}_t \end{bmatrix} = \begin{bmatrix} 1 & \mathbf{x}_s \cdot \mathbf{x}_t \\ \mathbf{x}_s \cdot \mathbf{x}_t & 1 \end{bmatrix}$$

is the matrix of the first fundamental form of $\mathbf{x}(s,t)$, and where, due to property 1,

$$\mathbf{B} = \begin{bmatrix} \mathbf{x}_{ss}\cdot\mathbf{N} & \mathbf{x}_{st}\cdot\mathbf{N} \\ \mathbf{x}_{ts}\cdot\mathbf{N} & \mathbf{x}_{tt}\cdot\mathbf{N} \end{bmatrix} = \begin{bmatrix} \kappa_s\,\mathbf{n}_s\cdot\mathbf{N} & 0 \\ 0 & \mathbf{x}_{tt}\cdot\mathbf{N} \end{bmatrix}$$

is the second fundamental form matrix, in which κ_s is the curvature of the s-parameter curve and $\mathbf{n}_s$ is its principal *unit* normal vector. But $\mathbf{n}_s\cdot\mathbf{N} = -1$ because the osculating plane of the s-curve contains $\mathbf{N}$ (see Fig. 3.5). We suppose that the occlusion is *smooth* and that we can assign a sign to κ_s so that it is negative when the principal normal vector $\mathbf{n}_s$ coincides with $\mathbf{N}$. It follows that the Gaussian curvature of the surface at the occlusion boundary is

$$K = \frac{\det\mathbf{B}}{\det\mathbf{G}} = -\kappa_s\,\frac{\mathbf{x}_{tt}\cdot\mathbf{N}}{1-(\mathbf{x}_s\cdot\mathbf{x}_t)^2}. \tag{3.1}$$

Assuming orthogonal projection, the occlusion contour tangent vector in the image is

$$\widehat{\mathbf{x}}_t = \mathbf{x}_t - (\mathbf{V}\cdot\mathbf{x}_t)\,\mathbf{V}.$$

We assign a sign to the occlusion contour curvature in a similar way as we did for the s-parameter curve (note that $\widehat{\mathbf{N}} = \mathbf{N}$). Then, after some algebra, and noting that $\mathbf{N}\cdot\mathbf{V} = 0$, and $\mathbf{x}_t\cdot\mathbf{N} = 0$, we get

$$\widehat{\kappa}_t = \frac{1}{\|\widehat{\mathbf{x}}_t\|}\,(-\widehat{\mathbf{N}})\cdot\frac{\partial}{\partial t}\left(\frac{\widehat{\mathbf{x}}_t}{\|\widehat{\mathbf{x}}_t\|}\right) = \ldots = -\frac{\mathbf{x}_{tt}\cdot\mathbf{N}}{1-(\mathbf{V}\cdot\mathbf{x}_t)^2}. \tag{3.2}$$

But $\mathbf{V} = -\mathbf{x}_s$. Therefore, from (3.2) and (3.1) follows

$$K = \kappa_s\,\widehat{\kappa}_t. \tag{3.3}$$

The curvature κ_s is unknown except for its sign, which is known because it must always be positive on a smooth occlusion. *Therefore, the sign of the Gaussian curvature near a smooth occlusion boundary is determined by the sign of the occlusion contour curvature.*

The procedure of recovering local surface orientation and quality along occlusion boundaries of a visible surface is summarized in Algorithm 1 on the next page.

3.1.3 Surface Orientation on a Self-Shadow Boundary

A self-shadow contour in an arbitrary image plane becomes a self-occlusion contour in the image projected on the light plane. Since we already know how to determine the local surface orientation near an occlusion boundary, there is a hope that we will be able to do the same on a self-shadow boundary *if there is a relation between gradients* of the corresponding images in the $\mathcal{V}$- and $\mathcal{L}$-planes.

Algorithm 1 (Surface Normal and Sign of K at an Occlusion Boundary)

For each pixel located at an occlusion contour do:

1. *Determine the image gradient $-\nabla I$ and complete the 2-D vector by adding a vanishing third component.*
2. *The result of step 1 normalized by its magnitude is the recovered unit surface normal $\mathbf{N}$ up to the $k\,\pi$ ambiguity, $k = 0, 1$.*
3. *If the occlusion contour becomes curved so that the image gradient is directed towards the center of its curvature, the surface has a negative Gaussian curvature K. The opposite direction of the image gradient implies a positive K, and a vanishing occlusion contour curvature implies a vanishing K.*

Theorem 3.1 (Mapping of ∇I from the $\mathcal{V}$-plane to the $\mathcal{L}$-plane)

Let T be a linear mapping representing the spatial rotation from the $\mathcal{V}$-coordinate system to the $\mathcal{L}$-coordinate system, P be an orthogonal projection from 3-D space onto the image plane, and let the surface reflectance function be space-invariant.

Then the gradient of an image $I^{\mathcal{V}}$ in the $\mathcal{V}$-plane (taken as a 3-D vector) maps on the gradient of the corresponding image $I^{\mathcal{L}}$ in the $\mathcal{L}$-plane by the composite mapping

$$\nabla I^{\mathcal{L}} = (P \circ T)\, \nabla I^{\mathcal{V}} .$$

Proof. Let

$$I^{\mathcal{L}} = I^{\mathcal{L}}(x_1^{\mathcal{L}}, x_2^{\mathcal{L}})$$

be an intensity image projected on $\mathcal{L}$-plane, and let

$$I^{\mathcal{V}} = I^{\mathcal{V}}(x_1^{\mathcal{V}}, x_2^{\mathcal{V}})$$

be the corresponding image projected on the $\mathcal{V}$-plane. The image coordinates $x_1^{\mathcal{L}}, x_2^{\mathcal{L}}$ and $x_1^{\mathcal{V}}, x_2^{\mathcal{V}}$ may represent two different parameterizations of the same surface. The brightness values observable at image points on $\mathcal{V}$ and $\mathcal{L}$ that are projections of the same particular surface point are identical because the surface reflectance is assumed to be invariant to the viewing direction.

The gradient of $I^{\mathcal{L}}(x_1^{\mathcal{L}}, x_2^{\mathcal{L}})$ can be expressed as

$$\nabla I^{\mathcal{L}} = \left[\frac{\partial I^{\mathcal{L}}}{\partial x_1^{\mathcal{L}}},\ \frac{\partial I^{\mathcal{L}}}{\partial x_2^{\mathcal{L}}},\ 0 \right]^{\mathrm{T}} = \ldots = \mathbf{P}\,\mathbf{J}\,\nabla I^{\mathcal{V}},$$

where $\mathbf{J}$ is the Jacobian matrix of the inverse mapping T^{-1}:

$$\mathbf{J}[i,j] = \left[\frac{\partial x_j^{\mathcal{V}}}{\partial x_i^{\mathcal{L}}}\right], \quad i,j = 1,2,3 \tag{3.4}$$

and $\mathbf{P}$ represents the projection onto the image plane. The spatial transformation T maps a vector $\mathbf{x}^{\mathcal{V}}$ in the $\mathcal{V}$-coordinate system onto a vector $\mathbf{x}^{\mathcal{L}}$ in the $\mathcal{L}$-system so that

$$\mathbf{x}^{\mathcal{L}} = \mathbf{T}\,\mathbf{x}^{\mathcal{V}}.$$

Then, the transformation matrix $\mathbf{T}$ may be viewed as

$$\mathbf{T}[i,j] = \left[\frac{\partial x_i^{\mathcal{L}}}{\partial x_j^{\mathcal{V}}}\right] = \left[\frac{\partial x_j^{\mathcal{V}}}{\partial x_i^{\mathcal{L}}}\right]^{-\mathrm{T}}, \tag{3.5}$$

since the mapping $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is linear and scale-preserving.

The matrix $\mathbf{T}$ is orthogonal, which means that $\mathbf{T}^{-1} = \mathbf{T}^{\mathrm{T}}$. Then, comparing (3.5) and (3.4), we can see that

$$\mathbf{P}\,\mathbf{J} = \mathbf{P}\,\mathbf{T}^{-\mathrm{T}} = \mathbf{P}\,\mathbf{T}. \tag{3.6}$$

Thus,

$$\nabla I^{\mathcal{L}} = \mathbf{P}\,\mathbf{T}\,\nabla I^{\mathcal{V}},$$

which is the desired result. □

This analysis confirms the observation of Lee & Rosenfeld published in [LR85]. The authors show in their article that the gradient of an image projected onto the $\mathcal{V}$-plane and mapped to the $\mathcal{L}$-system by a mapping $T : \mathcal{V} \rightarrow \mathcal{L}$ followed by a projection P is equivalent to the gradient of the image projected onto the $\mathcal{L}$-plane. Their result is derived under the assumption that the surface is spherical. Our result is more general and has important corollaries:

Corollary 3.1.1

Knowing the image gradient direction in an arbitrary $\mathcal{V}$-plane and the parameters of the imaging geometry, the gradient of the 'imaginary image' projected onto the $\mathcal{L}$-plane can be determined uniquely.

Which gives rise to

Corollary 3.1.2

The surface normal at a self-shadow boundary can be determined uniquely, assuming a space-invariant reflectance function and a known imaging geometry.

Algorithm 2 (Surface Normal at a Self-Shadow Boundary)

requirements

The illuminant position $\mathbf{L}$ is known. The reflectance function is spatially invariant.

1. *Compute the transformation matrix* $\mathbf{T}$ *transforming vectors from the viewer coordinate system to the light source coordinate system. Use the known slant* $\sigma_{\mathcal{L}}$ *and tilt* $\tau_{\mathcal{L}}$ *of the light direction expressed in the viewer coordinate system. See* (2.7) *on page 25.*
2. *For each pixel located at a self-shadow contour perform steps 3–6.*
3. *Determine the image gradient* $-\nabla I$ *and complete the 2-D vector by adding a vanishing third component. Assign the result to a vector-variable* $\mathbf{g}^{\mathcal{V}}$.
4. *Transform* $\mathbf{g}^{\mathcal{V}}$ *to the vector* $\mathbf{g}^{\mathcal{L}} := \mathbf{P}\,\mathbf{T}\,\mathbf{g}^{\mathcal{V}}$, *where* $\mathbf{P}$ *is an orthogonal projection matrix*
$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$
5. *The tilt component of the vector* $\mathbf{g}^{\mathcal{L}}$ *normalized by its magnitude gives the tilt* $\tau^{\mathcal{L}}$ *of surface normal in the* $\mathcal{L}$*-system. The slant component in the same system is* $\sigma^{\mathcal{L}} = \pi/2$. *Constitute a vector* $\mathbf{N}^{\mathcal{L}}$ *in the* $\mathcal{L}$*-system in the form of*
$$\mathbf{N}^{\mathcal{L}} := \left[\cos\tau^{\mathcal{L}},\ \ \sin\tau^{\mathcal{L}},\ \ 0\right]^{\mathrm{T}}$$
and transform it back to the $\mathcal{V}$*-system:* $\mathbf{N}^{\mathcal{V}} := \mathbf{T}^{-1}\,\mathbf{N}^{\mathcal{L}}$.
6. *The result* $\mathbf{N}^{\mathcal{V}}$ *of step 5 is the recovered normal expressed in the viewer coordinate system.*

This can be explained as follows. Consider only a single (and distant) light source illuminating the viewed surface. A self-shadow boundary is then a locus of points where the surface tangent plane contains the light direction vector (see Fig. 3.3 on page 33). It grows into an occlusion boundary as the vantage point moves to the vicinity of the light source. As we already know, the local surface orientation at the occlusion boundary can be determined uniquely. Because of the knowledge of the gradient of the 'imaginary image' projected onto the $\mathcal{L}$-plane, which is a simple and known transformation of the image gradient in a general $\mathcal{V}$-plane, *the surface normal at a point on the self-shadow boundary can be determined uniquely.* This result is in contrast to previous beliefs [IH81].

The procedure of recovering local surface orientation at self-shadow boundaries of a visible surface is summarized in Algorithm 2.

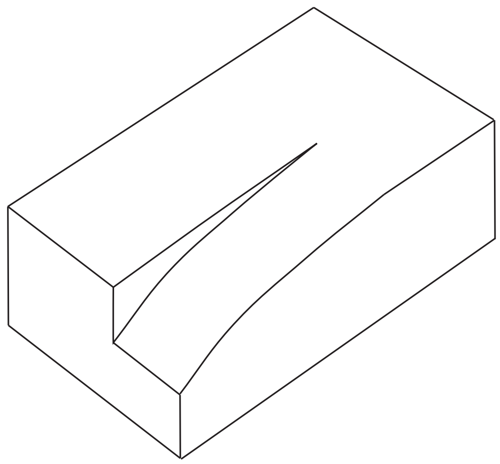

Figure 3.6: *A surface discontinuity is a regular arc of a spatial curve* $\mathbf{x}(s)$.

The conclusions of the discussion given so far lead us to a yet vague hypothesis:

Hypothesis 1

The image gradient plays a central role in local shading understanding.

This is also confirmed by an observation that, for a restricted class of surfaces, the gradient vector of an image projected onto the $\mathcal{L}$-plane is parallel to the projection of the surface normal vector onto this plane [Pen84, LR85]. More details on this issue will be given in Chapters 4 and 5.

3.1.4 Shading and Surface Discontinuities

We have already seen in Section 3.1.1 that the properties of shading near an occlusion boundary and near a surface discontinuity differ. We now ask the question what local information is available about the surface when a $\mathcal{C}^1$ surface discontinuity projects onto the image plane so that both its *illuminated* "sides" are visible.

We suppose that the first-order *surface discontinuities* form regular arcs which project onto *edges* in the image (see Fig. 3.6). Singularities separating the regular arcs will be called *vertices.* A vertex projects onto *a junction* of image edges. Thus the terms 'junction' and 'edge' are reserved for image structures, while the terms 'discontinuity' and 'vertex' are reserved for surface structures. Further, it is supposed that, in the vicinity of a vertex, the surface can be *locally modelled* by planar faces intersecting along the surface discontinuities and at the vertex. We will study the image of a vertex in particular.

We will develop a local $\mathcal{C}^1$ discontinuity model. Let $\mathbf{x}$ be a point on a surface discontinuity. $\mathbf{x}$ lies at the intersection of m surface tangent planes, each of them is tangent to one of the m adjacent surface patches (faces). $m > 2$ implies that the point must be a vertex. Such a vertex is incident with m arcs of regular spatial curves, each of them formed by the intersection of two neighbouring (curvilinear) faces. Along these arcs holds $m = 2$ and they are regular if the intersecting faces are regular along their intersection.

Suppose that $\mathbf{x}(s)$ is a locally regular discontinuity arc along which $m = 2$ and that $s \in I_s \subset \mathbb{R}$ is the parameter of the arc. Let $\mathbf{N}_i$ denote the normal vector of the i-th incident surface face taken *at the discontinuity.* A first-order local model of the discontinuity is a couple of intersecting planes, each of them being the tangent plane of either surface. Vector $\mathbf{t} = \dot{\mathbf{x}}$, tangent to the arc $\mathbf{x}(s)$, lies in both tangent planes. Consequently,

$$\mathbf{N}_i \cdot \mathbf{t} = 0, \qquad i = 1, 2. \tag{3.7}$$

In other words, $\mathbf{t}$ is perpendicular to the plane spanned by $\mathbf{N}_1$ and $\mathbf{N}_2$, $\mathbf{t} \sim \mathbf{N}_1 \times \mathbf{N}_2$.

Since both the surface faces intersect, the arc lengths *along the discontinuity* expressed in terms of the first fundamental forms of both surfaces must be equal:

$$\mathrm{I}_1 = \mathrm{I}_2\,. \tag{3.8}$$

Both (3.7) and (3.8) form the local surface discontinuity model. In fact, it is a local model for any surface of class $\mathcal{C}^n$, $n \geq 1$.

An analogous discontinuity model using the concept of gradient space is given in Mackworth's work on interpreting polyhedral scenes [Mac73]. The author argues that the image edge is always perpendicular to the line passing through points representing the intersecting planes in gradient space that is superimposed onto the image plane. It is also readily seen from (3.7): The gradient space representation of a plane given by its normal $\mathbf{N}_i = [-p_i, \ -q_i, \ 1]^{\mathrm{T}}$ is a point (p_i, q_i).[1] A line connecting gradient space representations of neighbouring planes i and $i+1$ is then parallel to the vector $\mathbf{d}_i = [p_i - p_{i+1}, \ q_i - q_{i+1}, \ 0]^{\mathrm{T}}$ in the gradient space. From (3.7) follows that this line is perpendicular to the (orthogonal) projection of the discontinuity edge tangent vector $\widehat{\mathbf{t}}$:

$$\mathbf{d}_i \cdot \widehat{\mathbf{t}} = (\mathbf{N}_i - \mathbf{N}_{i+1}) \cdot [\mathbf{t} - (\mathbf{t} \cdot \mathbf{V})\,\mathbf{V}] = -(\mathbf{t} \cdot \mathbf{V})\,(\mathbf{d}_i \cdot \mathbf{V}) = 0,$$

since the 2-D gradient-space can be embedded in the 3-D space so that it is parallel to image plane and therefore $\mathbf{d}_i \cdot \mathbf{V} = 0$.

Later, Mackworth's observation was re-applied by Horn [Hor77], who proposed a method for determining orientation of polyhedral faces from images of polyhedral edges and surface shading. Horn gives the solution in the form of a set of non-linear equations derived provided that the illuminant is located at the vantage point. We will shortly describe a generalized case and discuss it briefly.

[1] Note that $\mathbf{N}_i$ need not necessarily be a unit vector here.

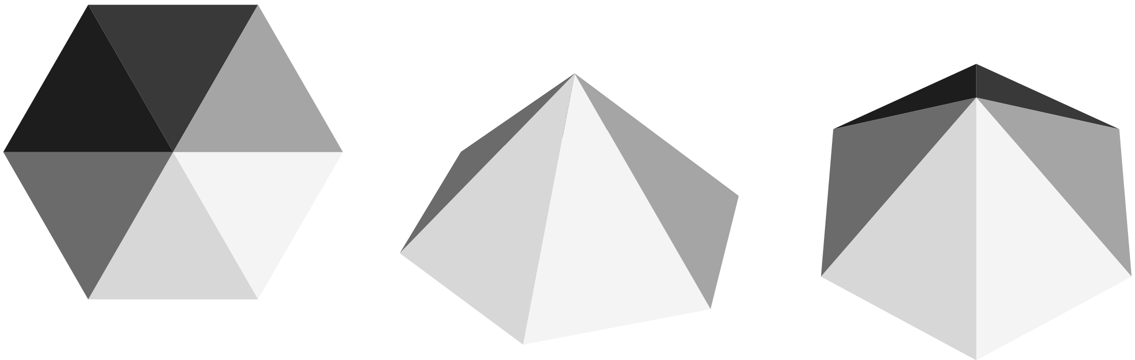

Figure 3.7: *A junction of image edges is an image of a surface vertex. In the leftmost view of the illuminated pyramid, there are three couples of parallel edges forming the central junction. In the center view, two of the six faces forming the top vertex are occluded, and in the rightmost view, there are two parallel edges among the six forming the central junction. No one of these views could be used to recover the spatial orientation of any of the pyramid faces (without the additional knowledge about the symmetry of this body).*

Let us suppose there are at least three discontinuity arcs joining at a common vertex on the surface. We also suppose that (1) all the incident surface faces are visible and that (2) there is no couple of parallel edges among those forming the junction in the image. This is necessary to ensure that an unambiguous solution will exist (see Fig. 3.7).

A face i incident with the vertex in question is represented by its normal $\mathbf{N}_i$ at the vertex, $\mathbf{N}_{i+1} \neq \mathbf{N}_i$. The discontinuity separating neighbouring faces i and $i+1$ has a tangent vector $\mathbf{t}_i$ at the vertex projecting onto $\widehat{\mathbf{t}}_i$ in the image plane. The index i is modulo m.[2] From (3.7) follows that

$$\mathbf{N}_i \cdot \mathbf{t}_i = 0, \qquad \mathbf{N}_{i+1} \cdot \mathbf{t}_i = 0, \qquad i \in N_m. \tag{3.9}$$

Furthermore, since $\mathbf{d}_i \cdot \mathbf{V} = 0$, there are $2\,m$ equations for $2\,m$ unknowns representing the surface normals $\mathbf{N}_i(p_i, q_i)$:

$$I_i = \alpha\, R(\mathbf{N}_i \cdot \mathbf{L},\, \mathbf{N}_i \cdot \mathbf{V},\, \mathbf{V} \cdot \mathbf{L}), \tag{3.10}$$

$$\mathbf{N}_i \cdot \widehat{t}_i = \mathbf{N}_{i+1} \cdot \widehat{t}_i, \qquad i \in N_m\,, \tag{3.11}$$

where I_i is the brightness of the face i observed in the vicinity of the vertex, assuming a known surface reflectance function R, albedo α, and imaging geometry ($\mathbf{L}$ and $\mathbf{V}$). Equations (3.10) are generally non-linear, which causes the algebraic solution to the

[2] We will denote this as $i \in N_m$, where N_m is the set of all integers modulo m.

system (3.10, 3.11) to be rather complicated. (3.10) can be expressed in the form of

$$f_i = R(p_i, q_i) - \frac{I_i}{\alpha} = 0, \qquad i = 1, 2, \ldots, m.$$

The equations (3.11) are linear in p, q and they can be written in the form of

$$f_{i+m} = (p_i - p_{i+1}) \cos\phi_i - (q_i - q_{i+1}) \sin\phi_i = 0, \qquad i \in N_m\,,$$

where ϕ_i is a known angle corresponding to the direction of the edge i in the image plane.

Multiple solutions must be expected since (3.10) is a system of non-linear equations. Some of the solutions will not be compatible with *the visibility and the illuminability constraints:*

$$\mathbf{V} \cdot \mathbf{N}_i > 0, \qquad \mathbf{L} \cdot \mathbf{N}_i > 0, \qquad i = 1, \ldots, m. \tag{3.12}$$

Nevertheless, the visibility constraint is always satisfied for a surface normal expressed in the form $\mathbf{N}_i = [-p_i,\ \ -q_i,\ \ 1]^{\mathrm{T}}$ in the viewer coordinate system where $\mathbf{V} = [0,\ \ 0,\ \ 1]^{\mathrm{T}}$.

In general, more than a single acceptable solution exists. The ambiguity of the solution is inherent in *any* local shading analysis.

Let us suppose now that the surface is Lambertian. Then,

$$R(p_i, q_i) = \frac{\mathbf{N}_i}{\|\mathbf{N}_i\|} \cdot \mathbf{L} = \frac{-l_1\, p_i - l_2\, q_i + l_3}{\sqrt{1 + {p_i}^2 + {q_i}^2}}\,. \tag{3.13}$$

If $\mathbf{L} = \mathbf{V}$, $R(p_i, q_i)$ becomes positive definite and symmetric. Then, if the vectors $\mathbf{P} = [p_i]$ and $\mathbf{Q} = [q_i]$ are a solution to (3.10, 3.11), also $\mathbf{P}' = [-p_i]$ and $\mathbf{Q}' = [-q_i]$ will be a solution. One of them will correspond to a convex and the other to a concave surface vertex. Generally, the problem (3.10, 3.11) has up to $2\,m$ possible solutions. The visibility and illuminability constraints (3.12) are always satisfied, so they cannot force any solution to be excluded. From the experience with local shading analysis follows that there are probably always up to four possible real solutions.

The problem (3.10, 3.11) can be solved using a numerical method. If we let $\mathbf{x}$ denote the entire vector of values p and q, $x_j = p_j$ for $j = 1, \ldots, m$ and $x_j = q_{j-m}$ for $j = m+1, \ldots, 2\,m$, then, in the neighbourhood of $\mathbf{x}$, each of the functions f_i can be expanded into a Taylor series

$$f_i(\mathbf{x} + \delta\mathbf{x}) = f_i(\mathbf{x}) + \sum_{j=1}^{m} \frac{\partial f_i}{\partial x_j}\, \delta x_j + O(\delta\mathbf{x}^2).$$

By neglecting $O(\delta\mathbf{x}^2)$ and higher terms, we obtain a set of equations for the $2\,m$-vector of corrections $\delta\mathbf{x}$, namely

$$\mathbf{D}\,\delta\mathbf{x} = -\mathbf{f}, \tag{3.14}$$

where

$$\mathbf{D}[i,j] = \left[\frac{\partial f_i}{\partial x_j}\right]$$

is a $2\,m \times 2\,m$ rectangular matrix and $\mathbf{f}$ is a $2\,m$-vector $\mathbf{f} = [f_i]$. A new estimate of $\mathbf{x}$ is obtained after a single Newton-Raphson step

$$\mathbf{x}^{k+1} = \mathbf{x}^{k} + \delta\mathbf{x}. \tag{3.15}$$

The matrix $\mathbf{D}$ is of the form of

$$\mathbf{D} = \left[\begin{array}{c|c} \mathbf{D}_{pN} & \mathbf{D}_{qN} \\ \hline \mathbf{D}_{pL} & \mathbf{D}_{qL} \end{array}\right],$$

where $\mathbf{D}_{pN}$ and $\mathbf{D}_{qN}$ are diagonal $m \times m$ matrices

$$\mathbf{D}_{pN} = \operatorname{diag}\left[\frac{\partial R(p_i,q_i)}{\partial p_i}\right], \qquad \mathbf{D}_{qN} = \operatorname{diag}\left[\frac{\partial R(p_i,q_i)}{\partial q_i}\right], \tag{3.16}$$

and $\mathbf{D}_{pL}$, $\mathbf{D}_{qL}$ are also $m \times m$, constant, circulant, always singular, and having the form of

$$\begin{bmatrix} d_1 & -d_1 & 0 & \dots & 0 \\ 0 & d_2 & -d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & d_{m-1} & -d_{m-1} \\ -d_m & 0 & \dots & 0 & d_m \end{bmatrix}, \tag{3.17}$$

where the known elements are $d_i = \cos\phi_i$ in the case of $\mathbf{D}_{pL}$, and $d_i = \sin\phi_i$ in the case of $\mathbf{D}_{qL}$, respectively.

The vector $\mathbf{f}$ has the form of

$$\mathbf{f} = \Big[\, R(p_1,q_1) - \frac{I_1}{\alpha},\ \dots,\ R(p_m,q_m) - \frac{I_m}{\alpha};\ 0,\ \dots,\ 0 \,\Big]^{\mathrm{T}}. \tag{3.18}$$

In the case of a Lambertian reflectance model, assuming that the unit light direction vector $\mathbf{L} = [l_1,\ \ l_2,\ \ l_3]^{\mathrm{T}}$ is known, we have $R(p_i,q_i)$ in the form of (3.13) and,

$$\begin{aligned} \frac{\partial R(p_i,q_i)}{\partial p_i} &= -\frac{l_1\,(1+{q_i}^2) + p_i\,(l_3 - l_2\,q_i)}{(1+{p_i}^2+{q_i}^2)^{\frac{3}{2}}}, \\ \frac{\partial R(p_i,q_i)}{\partial q_i} &= -\frac{l_1\,(1+{p_i}^2) + q_i\,(l_3 - l_2\,p_i)}{(1+{p_i}^2+{q_i}^2)^{\frac{3}{2}}}. \end{aligned} \tag{3.19}$$

The whole procedure for recovering local orientation near the surface vertices is summarized in Algorithm 3 on the next page. The surface need not necessarily be polyhedral.

Algorithm 3 (Surface Orientation near Vertices)

requirements

Illuminant position and reflectance function are known.

1. *Detect junctions in image where at least 3 edges meet so that there is no pair of parallel edges among them. For each junction perform steps 2–9.*
2. *Index all adjacent edges with an index i. Determine their orientation angles ϕ_i and compute the matrices $\mathbf{D}_{pL}$ and $\mathbf{D}_{qL}$ from* (3.17)*.*
3. *Index each face bounded by edges i and $i+1$ with index i.*
4. *For each face i determine the brightness value I_i near the vertex.*
5. *Set a $2m$-vector variable $\mathbf{x}$ to an initial value satisfying* (3.12)*.*
6. *Perform a maximum of M steps of the process* (3.14) *and* (3.15)*. Update the matrices $\mathbf{D}_{pN}$ and $\mathbf{D}_{qN}$ according to* (3.16) *and the vector $\mathbf{f}$ according to* (3.18)*.*
7. *Split the $2m$-vector $\mathbf{x}$ into two m-vectors and assign the first and the second parts to m-vector variables $\mathbf{P} := [p_i]$ and $\mathbf{Q} := [q_i]$, respectively. Compute the unit normals $\mathbf{N}_i$ as*

$$\mathbf{N}_i := \frac{[-p_i,\ \ -q_i,\ \ 1]^{\mathrm{T}}}{\sqrt{1 + {p_i}^2 + {q_i}^2}}, \qquad i = 1, 2, \ldots, m\,.$$

8. *If the solution satisfies* (3.12) *add it to the list of acceptable solutions.*
9. *If N_s different solutions have been found, exit with the list of acceptable solutions. Go to step 5 otherwise.*

When the surface is Lambertian, use (3.13) *and* (3.19) *for updating $\mathbf{D}_{pN}$, $\mathbf{D}_{qN}$, and $\mathbf{f}$. In this case, the albedo α need not be known: the image can be re-scaled by I_{max} which is the maximum image value. For details see Section 5.3. After this step, the albedo can be considered unit.*

3.2 Image Singularities

Singular image points are such loci in the image plane where the image function's directional derivative $D_{\mathbf{u}} I$ vanishes in every direction $\mathbf{u}$. This also means that the gradient magnitude $\|\nabla I\|$ vanishes. At such a point, no first-order directional information is available and this case must usually be treated separately in local shading analysis. If the singular image points form a line or a region, we speak of *singular image lines* and *singular image regions*, respectively.

After presenting an image singularity classification in this section, a detailed analysis of the shading in the neighbourhood of the most important point-singularity class will be given. We will show that the local surface orientation as well as the absolute value of the principal curvatures can be determined uniquely there.

3.2.1 Classification

Let us suppose the viewed surface $\mathbf{x}$ is mapped onto a Gaussian sphere so that the *Gaussian image* of a point $\mathbf{x}$ is the surface normal $\mathbf{N} = \mathbf{N}(\sigma, \tau)$. The Gaussian sphere is parameterized so that the latitude circles correspond to the τ-curves and the longitude circles correspond to the σ-curves, where σ is the slant of the normal $\mathbf{N}$ from the light direction vector $\mathbf{L}$ and τ is the normal tilt in the $\mathcal{L}$-plane. For each regular surface $\mathbf{x}$, there is a corresponding Gaussian image $\mathbf{N}$. The pre-images of the σ- and the τ-curves on the surface always exist and they are two independent systems of curves if the surface is locally non-developable.[3] The τ-curves on a non-developable surface project onto image isophotes as will be seen later in Section 4.2.

We will study the cause for image singularities under the Lambertian reflectance assumption. This simplifies the analysis without a significant loss of generality. The specular reflectance functions behave similarly, except for specular points that will be mentioned as a separate case in the following text. When assuming unit albedo, the image irradiance equation (2.4) reduces to the form

$$I = \mathbf{L} \cdot \mathbf{N}.$$

The differential of I is

$$\mathrm{d}I = \mathbf{L} \cdot \mathrm{d}\mathbf{N} = \mathbf{L} \cdot \mathbf{N}_\sigma \,\mathrm{d}\sigma + \mathbf{L} \cdot \mathbf{N}_\tau \,\mathrm{d}\tau$$

because the light direction vector $\mathbf{L}$ is supposed to be constant. The τ-direction is chosen so that the vector $\mathbf{N}_\tau$ is always perpendicular to $\mathbf{L}$ (or it has a vanishing magnitude). Therefore,

$$\mathrm{d}I = \mathbf{L} \cdot \mathbf{N}_\sigma \,\mathrm{d}\sigma \,.$$

Let us first look which events can give rise to a singular point in the image:

1. $\mathbf{N} = \pm\mathbf{L}\,.$

   According to the classification given in [Kv80], such a point is called *apex* when $\mathbf{N} = \mathbf{L}$ and *antapex* when $\mathbf{N} = -\mathbf{L}$. On a closed $\mathcal{C}^2$ surface, there is at least one apex and at least one antapex.

   Isolated apices and antapices give rise to local extrema of the image function in the $\mathcal{V}$-plane. On developable surfaces, those rulings[4] where the normal satisfies

[3] On a developable surface, the Gaussian curvature vanishes everywhere and the Gaussian image of such a surface is a single curve on the sphere.

[4] Rulings are straight lines on a developable parallel to the direction in which the principal curvature vanishes.

$\mathbf{N} = \pm\mathbf{L}$ give rise to singular image lines.

2. $\mathbf{N}_\tau \neq 0$ and $\mathbf{N}_\sigma = 0$.

   This event can only occur at a point on a parabolic line in the Gaussian image (and, consequently, on its corresponding pre-image on the surface), since there is a direction (the σ-direction) in which the normal curvature vanishes (because of $\mathbf{N}_\sigma = 0$). Since there are two independent directions in which the change of brightness I vanish in the Gaussian image—one along the τ-curve, because $\mathbf{L} \cdot \mathbf{N}_\tau = 0$, and the other along the σ-curve, because $\mathbf{N}_\sigma = 0$—this event gives rise to a singular point of the image function I. The event occurs at such a point T in Gaussian image, where the parabolic line just touches a τ-curve, i.e., the latitude circle on the sphere [Kv80] (see Fig. 3.8).

   In the image projected on the $\mathcal{V}$-plane, isolated saddle points and extrema, or singular lines can occur within this class, depending on the qualitative type of the surface:

   The extremum singularity occurs when the Gaussian image fold becomes convex with respect to the illuminant. In this case, the slant σ changes monotonically in *any* direction from the point of touch on the Gaussian sphere (see Fig. 3.8).

   The saddle singularity occurs when the fold becomes concave with respect to the illuminant. In this case, the slant increases in one direction and it decreases in another direction.

   The singular lines occur on developable surfaces along those rulings where the slant of the normal increases or decreases in both directions that are perpendicular to the ruling.

3. $\mathrm{d}\mathbf{N} = 0$.

   Such a surface point is called *planar*—both principal curvatures vanish there. There can exist isolated planar points, lines of planar points, or planar regions on a general, a developable, or a planar surface, respectively.

4. $\mathbf{N} = \mathbf{H}$.

   Singularities of this type occur on the illuminated surface only if the reflectance function contains a specular component. $\mathbf{H}$ is the bisectrix of the viewing and light directions. Singularities can be isolated, or they can form lines or regions according to the surface type.

3.2.2 Properties Related to Surface Geometry

The singularities of the second and the third type that were introduced in the previous subsection are important because of their invariant relationship to shape. The first and the fourth types are important because of their relation to the imaging geometry.

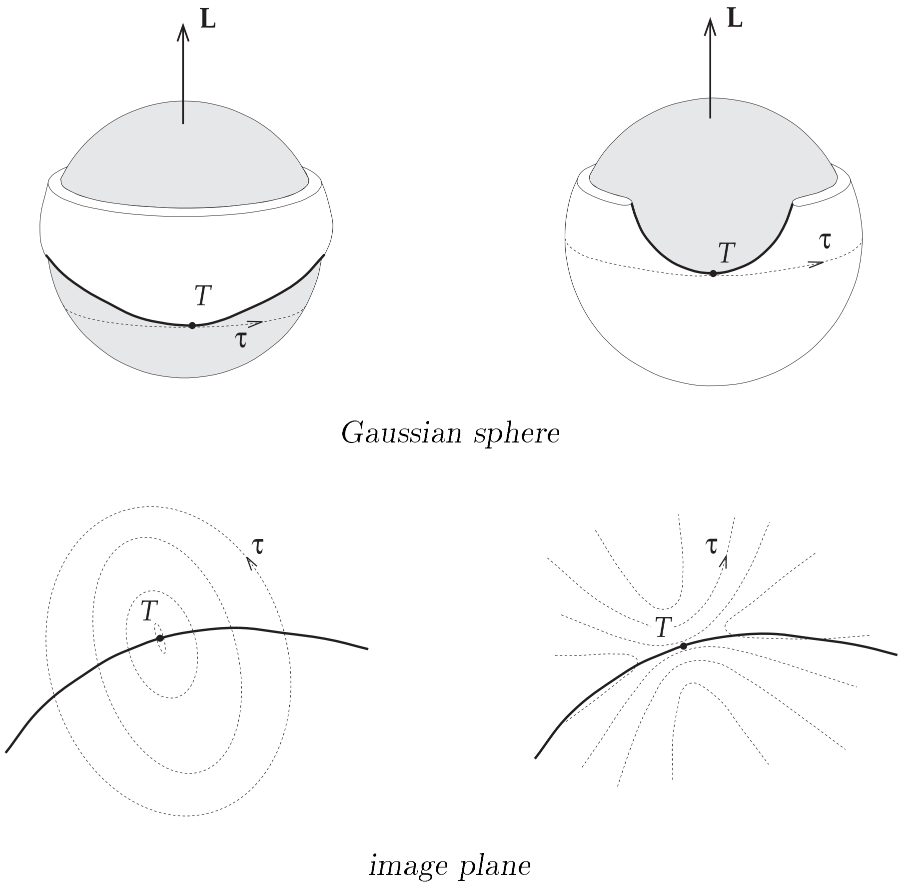

Figure 3.8: *Two types of image singularities due to a τ-parameter curve (dashed) touching the parabolic line (thick) on the Gaussian sphere. The Gaussian image shows folds due to parabolic lines. A fold that is convex with respect to the light direction vector* $\mathbf{L}$ *(upper left) gives rise to an extremum in the image (lower left). A concave fold gives rise to a saddle singularity in the image (right).* T *is the point of touch. When* $\mathbf{L}$ *moves,* T *will travel along the parabolic curve, since the τ-curves tilt from their original position. The 'thickness' of the Gaussian image serves only to give a clear perception of the fold's structure.*

To see why, let us look how these singularities move when the imaging geometry changes, how and where they are created and annihilated.

Let us suppose there is a photometric event $E(\mathbf{N}; \mathbf{V}, \mathbf{L})$ that is related to the local surface orientation $\mathbf{N}$ taken relative to the global imaging geometry $(\mathbf{V}, \mathbf{L})$. All the singularities listed in the previous subsection are of this type. The event is located at a surface point T: $\mathbf{x}(u_0, v_0)$. If the global geometry changes, the event must move so that the normal $\mathbf{N}$ grows again into a correct position relative to $\mathbf{V}$ and $\mathbf{L}$. The velocity of the movement depends on the speed of the global geometry change, of course, but also

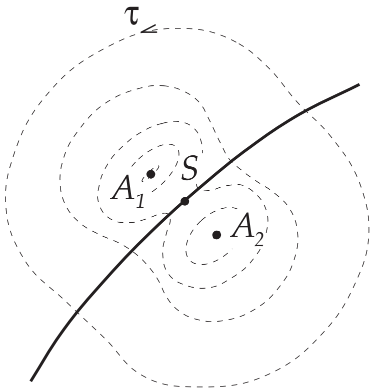

Figure 3.9: *The moment of creation of a couple of apices A_1, A_2 on a parabolic line in the image (thick). At the same moment, there necessarily appears a saddle point S, since the density of τ-curves (dashed) cannot be infinite. In an extreme case, only a single apex and a saddle occur. The saddle remains at the parabolic line.*

on the length of the arc along which the event must move on the surface. In other words, the turn of the surface normal $\mathbf{N}$ from its original position that is necessary to perform to restore the event after the global geometry change requires a movement of certain length in an appropriate direction on the surface. The amount of the turn after a movement by a unit distance in a particular direction is related to the surface 'curvature'. Thus, the velocity of the photometric event movement is inversely proportional to the 'curvature' in the direction in which the event moves.

The motion velocity on differently curved surfaces is also related to the density of their likely occurrence: the more the surface is curved, the slower the movement will be. The events, therefore, cling to strongly curved parts on the surface. This can also be explained so that the distribution of surface normals on a more strongly curved surface patch 'sweeps' a larger portion of space where the illuminant may occur and the patch thus attracts the singularities with higher probability.

What happens when the point of interest T is located just at a parabolic point on the surface? Since the normal $\mathbf{N}$ *must change,* the movement is not possible in the zero-curvature principal direction dv, since $\mathbf{N}$ does not change along such a direction: the change of $\mathbf{N}$ in a direction (du, dv) is, according to Rodrigues's formula,

$$d\mathbf{N} = \kappa_2 \, \mathbf{x}_u \, du \, .$$

Therefore, the trajectory leaving T must be parallel to the non-zero-curvature principal direction at T. Note that this direction is generally neither parallel nor perpendicular to the parabolic line *on the surface.*

Let us look how the different types of singularities move and where they are created and annihilated [Kv80]. The following description can also be interpreted as a phenomenologic equivalent of the classification given above. The different behaviour of the singularities under a change of the imaging geometry is a suitable cue to their type.

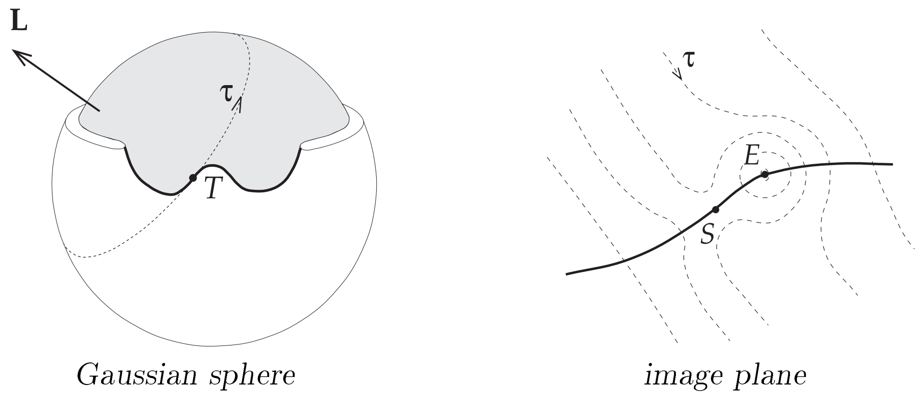

Figure 3.10: *The moment when a saddle S and an extremum E are created on a parabolic line (thick in both cases). The τ-curve just touches the inflection point T on the parabolic curve on the Gaussian sphere (left). The extremum E and the saddle S move apart along the parabolic line in image (right). Dashed are the image isocontours (τ-curves).*

Type 1. (apex, antapex)

Because of the Lambertian assumption, apices and antapices do not move if the vantage point changes but they move with the change of the illuminant position.

They are created or annihilated in pairs on parabolic curves and they move transversely to these lines at the moment of creation or annihilation. Together with the pair of apices, a single saddle singularity occurs. The saddle remains on the parabolic curve when the apices move apart (see Fig. 3.9 on the previous page) and thus belongs to the class 2. If the length of the fold's cross-section on the Gaussian sphere is infinitely short, only a single apex and a single saddle occur.

Type 2. (singularities stuck to parabolic lines)

The point singularities of this type move along parabolic lines if the illuminant moves and they are created or annihilated in pairs 'extremum – saddle,' with the exception of a pair consisting of the 'type–1–extremum' and the 'type–2–saddle' singularities.

In the regular case, they are created or annihilated at those loci on Gaussian images of parabolic lines, where an inflection occurs (see Fig. 3.10).

The line singularities on developable surfaces move perpendicularly to the rulings if either the viewer or the illuminant position changes, and they are created or annihilated along the lines of planar points.

Type 3. (images of planar points)

The images of planar points do not move with the change of the viewer or the illuminant positions, since the local planarity is inherent to the surface.

Type 4. (specular points)

Specular points occur on surfaces with a specular reflectance function. They are such loci on the surface, where the normal vector becomes parallel to the vector $\mathbf{H}$ (see Fig. 2.1 on page 20).

Specular points move if either the viewer or the illuminant moves. They are created and annihilated at the parabolic points. They move transversely to the parabolic lines in the moment of their creation or destruction. They cling to strongly curves parts for the same reason as apices and antapices do.

Specular lines on developable surfaces move in the same manner. They are created or annihilated along the lines of planar points.

The velocity of the specularity movement on the surface is generally lower than the velocity of the first two singularity types. This is caused by the fact that the event is generated by the bisectrix of the viewing and illumination directions:

$$\mathbf{H} = \frac{\mathbf{L}+\mathbf{V}}{\|\mathbf{L}+\mathbf{V}\|} = \frac{\mathbf{L}+\mathbf{V}}{\sqrt{2\,(1+\mathbf{L}\cdot\mathbf{V})}}\,.$$

Let us assume that $\mathbf{L}$ changes infinitesimally (the change of $\mathbf{V}$ has the same effect). After some algebra, we can find that the magnitude of the $\mathbf{H}$ change is

$$\|\mathrm{d}\mathbf{H}\|^2 = \ldots = \frac{\|\mathrm{d}\mathbf{L}\|^2}{2\,(1+\mathbf{L}\cdot\mathbf{V})}\left[1 - \frac{\left(\mathbf{V}\cdot\frac{\mathrm{d}\mathbf{L}}{\|\mathrm{d}\mathbf{L}\|}\right)^2}{2\,(1+\mathbf{L}\cdot\mathbf{V})}\right] < \frac{\|\mathrm{d}\mathbf{L}\|^2}{2},$$

since $\mathbf{L}\cdot\mathbf{V} > 0$ (the illuminated point is visible) and $\left(\mathbf{V}\cdot\frac{\mathrm{d}\mathbf{L}}{\|\mathrm{d}\mathbf{L}\|}\right)^2 < 1$. Thus the velocity of the specularities is lower than the velocity of photometric events generated by $\mathbf{L}$ or $\mathbf{N}$. This has the consequence that the specularities accumulate on strongly curved parts *with higher density* than other singularities.

3.2.3 Shading and Shape Near Isolated Singular Points

In this subsection, we will restrict our attention to isolated apices. We will study what kind of surface-related information can be recovered provided that the type of the singularity is known.

The possibility of recovering local geometric surface properties from image function properties in the neighbourhood of a singular point of the second type remains an open question that needs further investigation. In this case, the directionality of the image Hessian probably plays a role.

The image brightness is often saturated at locations of specularities. If not, the local surface orientation can be specified, since $\mathbf{N} = \mathbf{H}$ must hold there. If the imaging geometry is known, the vector $\mathbf{H}$ is known as well.

Shading near an Apex[5]

We will look for answers to the following two questions:

Q 1: What kind of information related to local shape is available within $\mathcal{F}_I^2$?

Q 2: What is the ambiguity of local shading near an apex?

In order to simplify the expressions, a Lambertian reflectance function and unit constant albedo α are assumed. The image function can be considered normalized by the albedo value: $I = \mathbf{N} \cdot \mathbf{L}$.

The second partial derivatives of the normalized image function $I(x, y)$ can be arranged into the Hessian matrix and evaluated:

$$\mathbf{H}_I = \begin{bmatrix} I_{xx} & I_{xy} \\ I_{xy} & I_{yy} \end{bmatrix} = \begin{bmatrix} \mathbf{L}\cdot\mathbf{N}_{xx} & \mathbf{L}\cdot\mathbf{N}_{xy} \\ \mathbf{L}\cdot\mathbf{N}_{xy} & \mathbf{L}\cdot\mathbf{N}_{yy} \end{bmatrix} = \begin{bmatrix} \mathbf{N}\cdot\mathbf{N}_{xx} & \mathbf{N}\cdot\mathbf{N}_{xy} \\ \mathbf{N}\cdot\mathbf{N}_{xy} & \mathbf{N}\cdot\mathbf{N}_{yy} \end{bmatrix},$$

since $\mathbf{L} = \mathbf{N}$ at the apex. Image coordinates x, y can be used as surface parameters, because each surface patch can be expressed in the Monge form

$$\mathbf{x} = [x,\ \ y,\ \ f(x, y)]^{\mathrm{T}}, \tag{3.20}$$

where f is a function whose 'plot' represents the surface [Lip69].

The normal derivatives $\mathbf{N}_i$ are related to surface tangent vectors $\mathbf{x}_x$ and $\mathbf{x}_y$ by the Weingarten equations

$$-\mathbf{N}_i = \beta_i^x\,\mathbf{x}_x + \beta_i^y\,\mathbf{x}_y\,, \qquad i = x,\, y \tag{3.21}$$

where β_i^j are the Weingarten coefficients. Let the matrix $\mathbf{S} = [\beta_i^j]$ represent the shape operator.[6] It is related to the first and the second fundamental forms by

$$\mathbf{S}^{\mathrm{T}} = \mathbf{G}^{-1}\,\mathbf{B}\,,$$

where $\mathbf{G}$ is a symmetric, positive definite matrix of the first fundamental form and $\mathbf{B}$ is a symmetric matrix of the second fundamental form.

[5]The antapex will not be considered, since it can not be illuminated by the assumed single-point illuminant.

[6]More details on the Weingarten mapping and fundamental forms can be found in Section 2.1.

Differentiating (3.21), we get

$$-\mathbf{N}_{ij} = \beta_i^x\, \mathbf{x}_{xj} + \beta_i^y\, \mathbf{x}_{yj} + (\beta_i^x)_j\, \mathbf{x}_x + (\beta_i^y)_j\, \mathbf{x}_y\,.$$

Because $\mathbf{N} \cdot \mathbf{x}_i = 0$, we have

$$-\mathbf{H}_I = \left[\begin{array}{cc} \beta_x^x\, \mathbf{N} \cdot \mathbf{x}_{xx} + \beta_x^y\, \mathbf{N} \cdot \mathbf{x}_{yx} & \beta_x^x\, \mathbf{N} \cdot \mathbf{x}_{xy} + \beta_x^y\, \mathbf{N} \cdot \mathbf{x}_{yy} \\ \beta_y^x\, \mathbf{N} \cdot \mathbf{x}_{xx} + \beta_y^y\, \mathbf{N} \cdot \mathbf{x}_{yx} & \beta_y^x\, \mathbf{N} \cdot \mathbf{x}_{xy} + \beta_y^y\, \mathbf{N} \cdot \mathbf{x}_{yy} \end{array}\right] = \mathbf{S}\,\mathbf{B} = \mathbf{B}\,\mathbf{G}^{-1}\,\mathbf{B}\,,$$

since all the involved matrices are symmetric and thus equivalent to their transposes. Let us multiply both sides of the equation by $\mathbf{G}^{-1}$ from the right side. Then,

$$-\mathbf{H}_I\,\mathbf{G}^{-1} = \left(\mathbf{B}\,\mathbf{G}^{-1}\right)^2 = \left[\left(\mathbf{G}^{-1}\,\mathbf{B}\right)^{\mathrm{T}}\right]^2 = \mathbf{S}^2,$$

since, again, all the involved matrices are symmetric. Since for any matrix $\mathbf{A_1}$

$$\det({\mathbf{A_1}}^2) = \left(\det \mathbf{A_1}\right)^2,$$

and for any symmetric matrix $\mathbf{A_2}$

$$\mathrm{tr}({\mathbf{A_2}}^2) = \left(\mathrm{tr}\, \mathbf{A_2}\right)^2 - 2\, \det \mathbf{A_2}\,,$$

the Gaussian and the mean curvatures of the surface near the apex are related to local image properties by

$$K^2 = \det(\mathbf{H}_I\,\mathbf{G}^{-1}), \tag{3.22}$$

$$2H^2 - K = -\frac{1}{2}\,\mathrm{tr}(\mathbf{H}_I\,\mathbf{G}^{-1}). \tag{3.23}$$

This is a system of two second-order polynomial equations with at most four real solutions for K and H:

a single solution exists when $K = H = 0$, which happens at a planar point,

two solutions exist when $K = 0$, i.e., at a parabolic point,

four solutions exist in the case of a generally curved surface.

The matrix $\mathbf{G}$ can be evaluated at the apex. The surface normal $\mathbf{N}$ for a surface of the form (3.20) becomes

$$\mathbf{N}(x, y) = \frac{\mathbf{x}_x \times \mathbf{x}_y}{\|\mathbf{x}_x \times \mathbf{x}_y\|} = \frac{\left[-f_x,\ \ -f_y,\ \ 1\right]^{\mathrm{T}}}{\sqrt{1 + {f_x}^2 + {f_y}^2}} = \frac{\left[-l_1,\ \ -l_2,\ \ 1\right]^{\mathrm{T}}}{\sqrt{1 + {l_1}^2 + {l_2}^2}}\,, \tag{3.24}$$

where $\mathbf{N} = \mathbf{L} = \left[-l_1,\ \ -l_2,\ \ 1\right]^{\mathrm{T}}$. Using (3.20) and (3.24), the matrix $\mathbf{G}$ becomes

$$\mathbf{G} = \left[\begin{array}{cc} 1 + {l_1}^2 & l_1\, l_2 \\ l_1\, l_2 & 1 + {l_2}^2 \end{array}\right]. \tag{3.25}$$

The determinant and the trace of $\mathbf{H}_I\,\mathbf{G}^{-1}$ in (3.22) and (3.23) can therefore be computed from the image data and the known imaging geometry.

Algorithm 4 (Local Orientation and Surface Curvature at an Apex)

requirements

The surface is Lambertian, the illuminant position and the surface albedo α are known (see comments on Alg. 3).

1. *Locate apices as those image points where the noise less image function reaches its maximum.*
2. *Compute the image Hessian matrix* $\mathbf{H}_I$.
3. *Construct the* 2×2 *matrix* $\mathbf{G}$ *from the known light direction* $\mathbf{L} = [-l_1,\ -l_2,\ 1]^{\mathrm{T}}$. *Use* (3.25).
4. *Compute the determinant and the trace of the matrix* $\mathbf{H}_I\, \mathbf{G}^{-1}$ *and solve the system of two quadratic equations* (3.22) *and* (3.23) *for the Gaussian and the mean curvatures. Koenderink's curvature measures can be computed using* (3.26) *and* (3.27). *The local surface orientation is found as* $\mathbf{N} = \mathbf{L}$.

According to the combination of K and H signs, four different local surface interpretations are possible:

sign of K	*sign of H*	*surface*
$+$	$+$	*convex elliptic*
$+$	$-$	*concave elliptic*
$-$	*any*	*hyperbolic*

The concavity/convexity is considered relative to the illuminant.

Now, we can answer the questions that we have raised at the beginning of this subsection:

A 1: The principal curvatures are recoverable from $\mathcal{F}_I^2$ at an apex.

A 2: There are up to four possible solutions. The sign of the Gaussian curvature cannot be determined, thus giving *one elliptic* and *one hyperbolic* solution. Also, the sign of the mean curvature cannot be determined, giving *one convex* and *one concave* solution.

This confirms Saxberg's observation that, at a singular image point, there are at least two and at most four solutions for the principal curvatures of a non-planar surface [Sax89].

Note that Koenderink's curvedness is defined *uniquely* at an apex:

$$C = \sqrt{\frac{{\kappa_1}^2 + {\kappa_2}^2}{2}} = \sqrt{2\, H^2 - K} = \sqrt{-\frac{1}{2}\, \mathrm{tr}(\mathbf{H}_I\, \mathbf{G}^{-1})}. \tag{3.26}$$

The *absolute value* of Koenderink's shape index can have up to two solutions:

$$S^2 = \left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right)^2 = \frac{\mathrm{tr}(\mathbf{H}_I\,\mathbf{G}^{-1}) \pm 2\sqrt{\det(\mathbf{H}_I\,\mathbf{G}^{-1})}}{\mathrm{tr}(\mathbf{H}_I\,\mathbf{G}^{-1}) \mp 2\sqrt{\det(\mathbf{H}_I\,\mathbf{G}^{-1})}}\,. \tag{3.27}$$

Algorithm 4 on the previous page summarizes the results of the analysis in a form suitable for implementation.

Chapter 4

Local Analysis of the Isophote Field

In this chapter, image isophotes are studied in order to establish a relation between the image properties and differential properties of the viewed surface. We will show how the image properties can be transformed into geometric properties of isophotes. The change of those properties under invertible image transformations, as well as under the change of reflectance function, will be discussed. A principal equation that relates the image gradient direction to the local tilt of the viewed surface will be derived. In Chapter 5, this equation will be used in two modifications of a *LSA* algorithm.

Here we choose a different strategy of presenting information compared to Chapter 3. To keep the line of argument, no "algorithms" will be presented inside the text. Instead, Chapter 4 will be rather theoretical. The following chapter deals with the algorithm for local shading analysis in full detail, for which the brief form of "algorithm description" becomes insufficient. There will be enough space to discuss the computational aspects more deeply. Thus, the material contained here constitutes only a theoretical prerequisite.

4.1 Properties of Image Isophotes

As has been mentioned in Chapter 2, image isophotes are curves of constant brightness in an intensity image. For the purpose of this chapter and in order to avoid problems with the definition of a curve in a discrete image, we will assume that the image is a *continuous*, scalar, two-parametric function $I(x, y)$.[1]

The aim of this section is to study local properties of image flow-lines and isophotes

[1] This does not cause a serious problem with the implementation. The ill-posedness of the discrete image function derivatives can be overcome by computing them at a certain scale [Koe90, RFK91, TP86].

out of singular elements of the image gradient field, in order to find an isomorphic set for a "reasonable" subset of the set of all partial image function derivatives

$$\mathcal{F}_I{}^n = \{I,\, I_x,\, I_y,\, I_{xx},\, I_{xy},\, I_{yy},\, \ldots\}\ ,$$

where n denotes the order of the highest derivative. The isomorphic subset $\mathcal{F}_I^{*n}$ should have a *geometric interpretation.* This is because we want to relate the geometric properties of the image isophotes and flow-lines to the differential properties of the viewed surface as has been suggested in Chapter 2.

We restrict n to the order of two because a higher-order image characteristics detector becomes very sensitive to noise in image data, since its frequency band includes the higher frequencies the higher the order n is.

For the sake of simplicity of the subsequent formulas, we introduce two perpendicular directions in the image, $\mathbf{u}$ and $\mathbf{v}$, such that[2]

$$I_{\mathbf{u}} = D_{\mathbf{u}}\, I = \nabla I \cdot \frac{\mathbf{u}}{\|\mathbf{u}\|} = -\|\nabla I(x,y)\|$$

and, consequently,

$$I_{\mathbf{v}} = D_{\mathbf{v}}\, I = \nabla I \cdot \frac{\mathbf{u}}{\|\mathbf{u}\|} = 0\,.$$

Suppose each isophote has a *label* I, which is the brightness it represents. Two natural local descriptors of a set of isocontours (like those in a topographical chart) are their *local orientation* γ, giving the "steepest descent" direction, and their *density* ρ, giving the amount of the descent. The density may be thought of as the inverse of the perpendicular distance between two neighbouring isophotes. If we move along an infinitesimal arclength of the image flow-line, the number of isophotes that are "pierced" during that move divided by the arclength gives the density of isophotes:

$$\rho = -I_{\mathbf{u}} = \|\nabla I\|\,.$$

Thus γ and ρ are the *first-order* natural image descriptors with a geometric interpretation. For reasons that will become clear later, we will work with a modification of this density that is defined as

$$\rho^* = \frac{\rho}{I}\,. \tag{4.1}$$

Two of the three *second-order* natural descriptors are the curvatures of the image isophote and the image flow-line. Let us assume the flow-line has its local orientation defined uniquely by the negative image gradient direction and the isophote orientation

[2] We hope the directional derivatives $D_{\mathbf{u}}\, f = f_{\mathbf{u}}$ used rather frequently in this section will not be confused with partial derivatives. We denote the directional derivatives by subscripts to save space. The subscripts are typographically different from those corresponding to partial derivatives, which should help the reader.

is chosen so that its tangent vector is parallel to the flow-line tangent vector rotated counterclockwise by $\pi/2$. See Fig. 4.3 on page 68, where the *unit* vectors $\mathbf{u}$ and $\mathbf{v}$ represent the flow-line and the isophote tangent vectors, respectively. Since $\mathbf{u}$ and $\mathbf{v}$ are *always* perpendicular, both curvatures are defined by an elementary turn of $\mathbf{u}$ after a movement about a unit arclength of either curve, and they can be written in the form:

$$\begin{aligned} \kappa_{\mathrm{F}} &= \mathbf{v} \cdot \frac{d\mathbf{u}}{ds_{\mathrm{F}}}, \\ \kappa_{\mathrm{I}} &= \mathbf{v} \cdot \frac{d\mathbf{u}}{ds_{\mathrm{I}}}, \end{aligned} \tag{4.2}$$

where s_{F} and s_{I} are the natural parameters of the flow-line and the isophote, respectively. We choose them so that the curvatures become

$$\begin{aligned} \kappa_{\mathrm{F}} &= \mathbf{v} \cdot D_{\mathbf{u}} \left(\frac{-\nabla I}{\|\nabla I\|} \right), \\ \kappa_{\mathrm{I}} &= \mathbf{v} \cdot D_{\mathbf{v}} \left(\frac{-\nabla I}{\|\nabla I\|} \right), \end{aligned}$$

where κ_{F} and κ_{I} are the flow-line curvature and the isophote curvature, respectively. Since $I_{\mathbf{v}} = 0$, they evaluate to

$$\kappa_{\mathrm{F}} = -\mathbf{v} \cdot \left(\frac{I_{\mathbf{u}}\,\mathbf{u} + I_{\mathbf{v}}\,\mathbf{v}}{-I_{\mathbf{u}}} \right)_{\mathbf{u}} = \frac{I_{\mathbf{uv}}}{I_{\mathbf{u}}}, \tag{4.3}$$

$$\kappa_{\mathrm{I}} = -\mathbf{v} \cdot \left(\frac{I_{\mathbf{u}}\,\mathbf{u} + I_{\mathbf{v}}\,\mathbf{v}}{-I_{\mathbf{u}}} \right)_{\mathbf{v}} = \frac{I_{\mathbf{vv}}}{I_{\mathbf{u}}}. \tag{4.4}$$

There is one more second-order directional derivative in ${\mathcal{F}_I}^2$, which is $I_{\mathbf{uu}}$. It has been shown in [FRKV91] that the *relative variation of isophote density along the flow-line scaled by the image gradient magnitude is*

$$\nu = D_{\mathbf{u}}(\ln \rho) = \frac{I_{\mathbf{uu}}}{I_{\mathbf{u}}}. \tag{4.5}$$

We have therefore found a set that is isomorphic to the set of second-order image function properties ${\mathcal{F}_I}^2$; it is

$$\mathcal{F}_I^{*2} = \{I,\, \rho^*,\, \gamma,\, \kappa_{\mathrm{F}},\, \kappa_{\mathrm{I}},\, \nu\}, \tag{4.6}$$

where,

- I is the image function value,
- ρ^* is the isophote density scaled by I,
- γ is the negative image gradient direction,

κ_{F} is the flow-line curvature,

κ_{I} is the isophote curvature,

ν is the relative variation of the isophote density in the negative image gradient direction.

4.1.1 T-invariance and R-invariance

Here we discuss certain important properties of some of the members of $\mathcal{F}_I^{*2}$ that do not have the members of $\mathcal{F}_I{}^2$.

T-invariants from $\mathcal{F}_I^{*2}$

First, it turns out that *the gradient direction, the isophote curvature, and the flow-line curvature are invariant to any invertible image value transformation.* This may be viewed as follows. Let us consider a continuous invertible map T ($T' \neq 0$) and let us study its influence upon the $\mathcal{F}_I^{*2}$ members (see also [FRKV91]). From Leibnitz's product rule and the chain rule for differentiation follows that

$$\begin{aligned} \Big(T(I)\Big)_{\mathbf{i}} &= I_{\mathbf{i}}\, T' \\ \Big(T(I)\Big)_{\mathbf{ij}} &= I_{\mathbf{ij}}\, T' + I_{\mathbf{i}}\, I_{\mathbf{j}}\, T'' \end{aligned} \tag{4.7}$$

for $\mathbf{i}, \mathbf{j} \in \{\mathbf{u}, \mathbf{v}\}$. Substituting into (4.1), (4.3), (4.4), and (4.5), we can see that the members of $\mathcal{F}_I^{*2}$ transform as follows:

$$\begin{aligned} I &\mapsto T(I) \\ \rho^* = -\frac{I_{\mathbf{u}}}{I} &\mapsto -\frac{I_{\mathbf{u}}}{T(I)}\, T' \\ \gamma = \operatorname{tilt}(-\nabla I) &\mapsto \operatorname{tilt}(-\nabla I\, T') = \gamma \\ \kappa_{\text{F}} = \frac{I_{\mathbf{uv}}}{I_{\mathbf{u}}} &\mapsto \frac{I_{\mathbf{uv}}\, T' + I_{\mathbf{u}}\, I_{\mathbf{v}}\, T''}{I_{\mathbf{u}}\, T'} = \frac{I_{\mathbf{uv}}}{I_{\mathbf{u}}} = \kappa_{\text{F}} \\ \kappa_{\text{I}} = \frac{I_{\mathbf{vv}}}{I_{\mathbf{u}}} &\mapsto \frac{I_{\mathbf{vv}}\, T' + I_{\mathbf{v}}{}^2\, T''}{I_{\mathbf{u}}\, T'} = \frac{I_{\mathbf{vv}}}{I_{\mathbf{u}}} = \kappa_{\text{I}} \\ \nu = \frac{I_{\mathbf{uu}}}{I_{\mathbf{u}}} &\mapsto \frac{I_{\mathbf{uu}}\, T' + I_{\mathbf{u}}{}^2\, T''}{I_{\mathbf{u}}\, T'} = \nu - \rho\, \frac{T''}{T'} \end{aligned} \tag{4.8}$$

Thus γ, κ_{F}, and κ_{I} are invariant to any invertible image transform T. We shortly call them T-invariants. If T is linear, T'' vanishes, and ν becomes also invariant. If T is an offset, all the members of $\mathcal{F}_I^{*2}$ but the first two become invariant.

R-invariants

A member of $\mathcal{F}_I^{*2}$ is an R-invariant if it is independent of a change of the reflectance function R from the class of admissible reflectance functions.[3] According to the image irradiance equation (2.4), assuming $\mathbf{V}$ = const. and $\mathbf{L}$ = const., and constant unit albedo α, we have[4]

$$I = R(\mathbf{N}\cdot\mathbf{L},\ \mathbf{N}\cdot\mathbf{V},\ \mathbf{V}\cdot\mathbf{L}).$$

Then,

$$\begin{aligned} I_{\mathbf{i}} &= \mathbf{N_i}\cdot\mathbf{L}\ R_1() + \mathbf{N_i}\cdot\mathbf{V}\ R_2()\,, \\ I_{\mathbf{ij}} &= \mathbf{N_{ij}}\cdot\mathbf{L}\ R_1() + \mathbf{N_{ij}}\cdot\mathbf{V}\ R_2() \\ &+ \mathbf{N_i}\cdot\mathbf{L}\ \mathbf{N_j}\cdot\mathbf{L}\ R_{11}() + \mathbf{N_i}\cdot\mathbf{L}\ \mathbf{N_j}\cdot\mathbf{V}\ R_{12}() \\ &+ \mathbf{N_j}\cdot\mathbf{L}\ \mathbf{N_i}\cdot\mathbf{V}\ R_{21}() + \mathbf{N_i}\cdot\mathbf{V}\ \mathbf{N_j}\cdot\mathbf{V}\ R_{22}()\,. \end{aligned}$$

$R_i()$ denotes a partial derivative of the reflectance function $R()$ with respect to the i-th argument.

Note that the isophote direction is defined so that

$$I_{\mathbf{v}} = \mathbf{N_v}\cdot\mathbf{L}\ R_1() + \mathbf{N_v}\cdot\mathbf{V}\ R_2() = 0.$$

The local vectors $\mathbf{u}$ and $\mathbf{v}$ therefore *change* under a change of R—because of the changing isophote direction. This causes that none of the $\mathcal{F}_I^{*2}$ members remains R-invariant under general conditions. If we restrict the imaging geometry or the class of admissible reflectance functions, the R-invariance is achieved when

1. $\mathbf{V} = \mathbf{L}$ (illuminant at vantage point). Then,

$$\begin{aligned} I_{\mathbf{i}} &= \mathbf{N_i}\cdot\mathbf{L}\ R_k\,, \\ I_{\mathbf{ij}} &= \mathbf{N_i}\cdot\mathbf{L}\ \mathbf{N_j}\cdot\mathbf{L}\ R_{lm} + \mathbf{N_{ij}}\cdot\mathbf{L}\ R_k\,, \end{aligned}$$

where we used Einstein's notation.[5] From the first equation follows that $\mathbf{N_v}\cdot\mathbf{L} = 0$. Then, we get the same invariants as in (4.8) because

$$\begin{aligned} I &= R() \\ \rho^* &= -\mathbf{N_u}\cdot\mathbf{L}\ \frac{R_k}{R()} \end{aligned}$$

[3]For our purposes, this requirement states nothing but the function is in the form of (2.3). We may also require that $0 \le R \le 1$.

[4]The albedo is assumed as unity to simplify the subsequent formulas. A general albedo does not affect the conclusions that will be drawn later.

[5]R_k stands for $\sum_k R_k()$, and R_{lm} stands for $\sum_l \sum_m R_{lm}()$.

$$\begin{aligned}
\gamma &= \mathrm{tilt}(-\nabla I\, R_k) \\
\kappa_{\mathrm{F}} &= \frac{\mathbf{N_{uv}} \cdot \mathbf{L}\, R_k}{\mathbf{N_u} \cdot \mathbf{L}\, R_k} \\
\kappa_{\mathrm{I}} &= \frac{\mathbf{N_{vv}} \cdot \mathbf{L}\, R_k}{\mathbf{N_u} \cdot \mathbf{L}\, R_k} \\
\nu &= \frac{(\mathbf{N_u} \cdot \mathbf{L})^2\, R_{lm} + \mathbf{N_{uu}} \cdot \mathbf{L}\, R_k}{\mathbf{N_u} \cdot \mathbf{L}\, R_k} = \mathbf{N_u}\,\frac{R_{lm}}{R_k} + \frac{\mathbf{N_{uu}} \cdot \mathbf{L}}{\mathbf{N_u} \cdot \mathbf{L}}
\end{aligned}$$

2. R is space-invariant. Then,

$$\begin{aligned}
I_{\mathbf{i}} &= \mathbf{N_i} \cdot \mathbf{L}\, R_1()\,, \\
I_{\mathbf{ij}} &= \mathbf{N_i} \cdot \mathbf{L}\, \mathbf{N_j} \cdot \mathbf{L}\, R_{11}() + \mathbf{N_{ij}} \cdot \mathbf{L}\, R_1()\,,
\end{aligned}$$

and we therefore get the same invariants as in (4.8) if we evaluate the members of $\mathcal{F}_I^{*2}$ as above.

We can conclude:

Theorem 4.1 (R-Invariants from $\mathcal{F}_I^{*2}$)

The image gradient direction γ, the image flow-line curvature κ_{F}, and the image isophote curvature κ_{I} are R-invariants if

1. $\mathbf{L} = \mathbf{V}$, *or,*
2. *the class of R is restricted to space-invariant reflectance functions.*

This fact will play an important role for the behaviour of the local tilt estimator proposed in Section 5.4.

4.2 Gaussian Parameterization of a Non-Developable Surface

From the fundamental theorem of differential geometry of surfaces it is known that the six coefficients of the first and second fundamental forms determine a surface of class at least $\mathcal{C}^3$ *uniquely*, up to a proper Euclidean (rigid) motion. Those six coefficients (in fact, two-parametric scalar functions) must also be somehow related to the intensity image of the illuminated surface, since the surface radiance in a particular direction depends on the local surface geometry, namely on the surface tangent plane orientation at the radiating point. On the other hand, the surface reflectance in a particular direction does not depend on the position of the rigid surface in space, assuming the conditions given

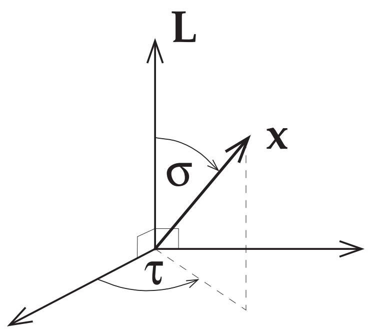

Figure 4.1: *Slant σ and tilt τ of a vector $\mathbf{x}$ with respect to a vector $\mathbf{L}$.*

in Section 1.1 on page 4. The relation between the local properties of the image and the six fundamental coefficients is not straightforward because the coefficients have no direct physical interpretation. We therefore look for different surface descriptors that are *explicitly* related to the surface tangent plane orientation. The set of local surface descriptors, necessary to model the surface uniquely and such that they are well-related to the intensity image of the viewed surface, will be called a *local surface representation.* We prove that this representation exists and it determines the surface uniquely, as do the six fundamental coefficients.

A surface normal defining the surface local tangent plane orientation can be thought as a two-parametric vector

$$\mathbf{N}(\sigma,\tau) = [\cos\tau\ \sin\sigma,\ \ \sin\tau\ \sin\sigma,\ \ \cos\sigma]^{\mathrm{T}}\,,\ \ \sigma \neq k\,\pi\,, \tag{4.9}$$

where σ and τ are the slant and the tilt of the vector from the light direction $\mathbf{L}$, respectively (see Fig. 4.1). The tilt and the slant define polar coordinates on a unit Gaussian sphere located at the coordinate origin. This representation has several local properties, the advantage of which becomes clear later. For the surface, whose surface normal is $\mathbf{N}(\sigma,\tau)$, it holds that

$$\begin{aligned} \mathbf{N}_\sigma \cdot \mathbf{N}_\tau &= 0\,, \\ \mathbf{N}_\sigma \cdot \mathbf{N}_\sigma &= 1\,, \\ \mathbf{N}_\tau \cdot \mathbf{N}_\tau &= \sin^2\sigma\,. \end{aligned} \tag{4.10}$$

A surface can be parameterized by any two independent parameters under very weak conditions. We choose those two that are closely related to the normal vector: the slant σ and tilt τ of the normal $\mathbf{N}(\sigma,\tau)$ with respect to the (constant) light direction vector $\mathbf{L}$. We call this choice the *Gaussian parameterization.* It parameterizes the surface "by its normal vector." The parametric curves corresponding to constant slant and constant tilt will be called *τ-curves* and *σ-curves,* respectively. The parameterization does not make any sense for a developable surface, since its Gaussian image is a single one-parametric object (curve) on the Gaussian sphere. The local surface tilt and slant are therefore

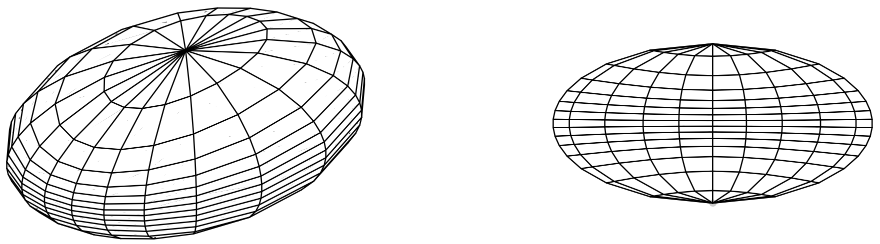

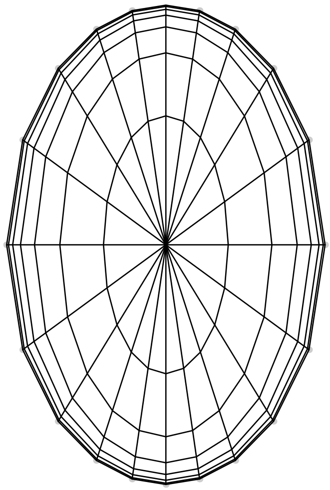

Figure 4.2: *Gaussian parameterization of an ellipsoid. Note that the curves of constant tilt, i.e., the σ-curves are planar and intersect at two singular points of the parameterization: $\sigma = 0, \pi$. The non-planar curves of constant slant are perpendicular to σ-curves in the case when the vectors tangent to both the σ- and the τ-curves lie in principal directions. All the three projections are orthogonal.*

not independent characteristics in the case of a developable surface. See Fig. 4.2 for an example of the Gaussian parameterization.

The change of surface radiance in a particular direction on the surface is proportional to the change of the surface normal $\mathbf{N}$ in that direction. The first derivatives of the image function are therefore related to the surface curvature. More specifically, the image gradient is related to the shape operator. We should therefore include the shape operator into our local surface model.

The change of the normal vector $d\mathbf{N} = \mathbf{N}_\sigma d\sigma + \mathbf{N}_\tau d\tau$ lies in the surface tangent plane and it is related to surface parametric curves tangent vectors by means of the *shape operator* $\mathbf{S} = [\beta_i^j]$. In our case,

$$
\begin{aligned}
\mathbf{N}_\sigma &= -\beta_\sigma^\sigma \mathbf{x}_\sigma - \beta_\sigma^\tau \mathbf{x}_\tau \, , \\
\mathbf{N}_\tau &= -\beta_\tau^\sigma \mathbf{x}_\sigma - \beta_\tau^\tau \mathbf{x}_\tau \, .
\end{aligned}
\tag{4.11}
$$

It seems that the explicit vector function (4.9) together with the four Weingarten coefficients in (4.11) is a surface description explicitly related to the intensity image.

Proposition 4.1

> *A local surface representation which is equivalent to the representation based on the first and the second fundamental forms and which is well-related to the intensity image can be based on the shape operator together with an explicit model of the surface normal in the form of* (4.9).

There is another physical motivation for this parameter choice: the image isophotes are a natural projection of τ-curves onto the image plane under certain conditions. To see this, consider a surface isophote arc locally aligned with a direction $\mathbf{w}$ on a surface. The derivative of the reflectance function in the $\mathbf{w}$ direction must vanish. From (2.3) it follows that

$$D_{\mathbf{w}}\, R(\mathbf{N}\cdot\mathbf{L},\, \mathbf{N}\cdot\mathbf{V},\, \mathbf{V}\cdot\mathbf{L}) = D_{\mathbf{w}}\, R(\mathbf{N}\cdot\mathbf{L}) = D_{\mathbf{w}}\, R(\cos\sigma) = -R'\sin\sigma\, D_{\mathbf{w}}\,\sigma \qquad (4.12)$$

when $\mathbf{V} = \mathbf{L}$ or when the reflectance function is space-invariant. The expression vanishes for $D_{\mathbf{w}}\,\sigma = 0$, i.e., along the τ-curve.

Corollary 4.1.1

> *The surface isophote and the image isophote in the $\mathcal{V}$-plane are therefore curves of constant slant of the normal $\mathbf{N}$ from the light direction $\mathbf{L}$, when $\mathbf{L} = \mathbf{V}$ or when R is space-invariant.*

In the next subsection, some of the properties of the representation based on Proposition 4.1 will be demonstrated, and the uniqueness of the representation will be proved.

4.2.1 Existence Theorem

Given an explicit Gaussian map (4.9) of a surface of a sufficiently high class, we want to investigate whether or not the coefficients of the Weingarten mapping (4.11) are a unique representation of the surface, up to a proper Euclidean motion. We also look for an existence condition regarding the underlying surface in the Gaussian parameterization $\mathbf{x} = \mathbf{x}(\sigma, \tau)$.

If $\mathbf{x}(\sigma,\tau)$ is a surface of at least $\mathcal{C}^2$, then the second-order mixed partial derivatives of $\mathbf{x}$ must be independent of the order of differentiation; i.e., $(\mathbf{x}_\sigma)_\tau = (\mathbf{x}_\tau)_\sigma$. Differentiating (4.11) and multiplying by (4.9), we have

$$\begin{array}{lcrcrcl}
\beta_\sigma^\sigma\, \mathbf{x}_{\sigma\sigma}\cdot\mathbf{N} & + & \beta_\sigma^\tau\, \mathbf{x}_{\tau\sigma}\cdot\mathbf{N} & = & -\mathbf{N}_{\sigma\sigma}\cdot\mathbf{N} & = & 1\,, \\
\beta_\sigma^\sigma\, \mathbf{x}_{\sigma\tau}\cdot\mathbf{N} & + & \beta_\sigma^\tau\, \mathbf{x}_{\tau\tau}\cdot\mathbf{N} & = & -\mathbf{N}_{\sigma\tau}\cdot\mathbf{N} & = & 0\,, \\
\beta_\tau^\sigma\, \mathbf{x}_{\sigma\sigma}\cdot\mathbf{N} & + & \beta_\tau^\tau\, \mathbf{x}_{\tau\sigma}\cdot\mathbf{N} & = & -\mathbf{N}_{\tau\sigma}\cdot\mathbf{N} & = & 0\,, \\
\beta_\tau^\sigma\, \mathbf{x}_{\sigma\tau}\cdot\mathbf{N} & + & \beta_\tau^\tau\, \mathbf{x}_{\tau\tau}\cdot\mathbf{N} & = & -\mathbf{N}_{\tau\tau}\cdot\mathbf{N} & = & \sin^2\sigma\,.
\end{array}$$

Rearranging the equations, solving for β_τ^σ and β_σ^τ, and substituting for $\beta_\sigma^\sigma\beta_\tau^\tau - \beta_\sigma^\tau\beta_\tau^\sigma = K$, where K be the Gaussian curvature, we obtain

$$\begin{aligned} \beta_\tau^\sigma &= -K\,\mathbf{x}_{\tau\sigma}\cdot\mathbf{N} = -K\,M\,, \\ \beta_\sigma^\tau\,\sin^2\sigma &= -K\,\mathbf{x}_{\sigma\tau}\cdot\mathbf{N} = -K\,M\,. \end{aligned} \tag{4.13}$$

The case $\beta_\tau^\sigma = \beta_\sigma^\tau = 0$ occurs when either $K = 0$, or

$$M = \mathbf{x}_{\sigma\tau}\cdot\mathbf{N} = \mathbf{x}_{\tau\sigma}\cdot\mathbf{N} = -\mathbf{x}_\sigma\cdot\mathbf{N}_\tau = -\mathbf{x}_\tau\cdot\mathbf{N}_\sigma = 0\,.$$

See also (2.1) on page 15.

Lemma 4.1

$M = 0$ is a sufficient condition for the σ- and τ-curves of the Gaussian parameterization to be (at least locally) tangent to the lines of curvature.

Proof. It is known from the differential geometry of surfaces that a necessary and sufficient condition for parametric curves to be tangent to lines of curvature is $M = 0$ *and* $F = \mathbf{x}_\sigma\cdot\mathbf{x}_\tau = 0$. In other words, the pairs of vectors $(\mathbf{x}_\sigma, \mathbf{x}_\tau)$, $(\mathbf{x}_\sigma, \mathbf{N}_\tau)$, and $(\mathbf{x}_\tau, \mathbf{N}_\sigma)$ must all be pairwise perpendicular. Since it follows from the conditions $M = 0$ and $K \neq 0$ that the $(\mathbf{x}_\sigma, \mathbf{N}_\tau)$ and $(\mathbf{x}_\tau, \mathbf{N}_\sigma)$ vectors are perpendicular, and, according to (4.10), the $\mathbf{N}_\sigma$ and $\mathbf{N}_\tau$ vectors are perpendicular as well, then $\mathbf{x}_\sigma$ and $\mathbf{x}_\tau$ must also be perpendicular, which completes the proof. □

This means that

$$\begin{cases} \beta_\tau^\sigma = \beta_\sigma^\tau = 0 & \text{iff } \sigma\text{- and } \tau\text{-curves are locally} \\ & \text{tangent to the lines of curvature;} \\ \beta_\tau^\sigma = \beta_\sigma^\tau\,\sin^2\sigma & \text{otherwise} \end{cases} \tag{4.14}$$

for a surface with Gaussian parameterization.

We can finally state the following theorem that is a modification of the fundamental theorem of surfaces for the case of a Gaussian surface parameterization.

Theorem 4.2 (Existence and Uniqueness)

Let $\mathbf{N}(\sigma, \tau)$ be a smooth two-parametric vector function in the form of (4.9) and β_i^j be four two-parametric scalar functions of class at least $\mathcal{C}^1$ satisfying the condition

$$\beta_\tau^\sigma = \beta_\sigma^\tau \, \sin^2 \sigma \,,$$

such that $\det[\beta_i^j] \not\equiv 0$. Let the partial derivatives of $\mathbf{N}$ satisfy the Weingarten equations (4.11) for a two-parametric vector function $\mathbf{x}(\sigma, \tau)$.

Then $\mathbf{N}(\sigma, \tau)$ is a Gaussian image of a non-developable surface $\mathbf{x}(\sigma, \tau)$ of class at least $\mathcal{C}^3$, and $[\beta_i^j]$ is the shape operator of the surface. The surface is unique up to a proper Euclidean motion and it can be parameterized in the same manner as $\mathbf{N}(\sigma, \tau)$.

The condition of $K = \det[\beta_i^j] \not\equiv 0$ follows from the fact that, on each point on the surface where $K \equiv 0$, there is a direction $\mathbf{w}$ in which $D_{\mathbf{w}} \mathbf{N} = 0$. Such a case can not be described by (4.9).

Proof. Let us have three scalar two-parametric $\mathcal{C}^2$ functions E, F, G and another three scalar two-parametric $\mathcal{C}^1$ functions L, M, N defined on the surface, such that $E = \mathbf{x}_\sigma \cdot \mathbf{x}_\sigma$, $F = \mathbf{x}_\sigma \cdot \mathbf{x}_\tau$, $G = \mathbf{x}_\tau \cdot \mathbf{x}_\tau$, $L = -\mathbf{x}_\sigma \cdot \mathbf{N}_\sigma$, $M = -\mathbf{x}_\sigma \cdot \mathbf{N}_\tau = -\mathbf{x}_\tau \cdot \mathbf{N}_\sigma$, $N = \mathbf{x}_\tau \cdot \mathbf{N}_\tau$. We will show that these functions are unique first and second fundamental coefficients of the surface $\mathbf{x}$ with normal vector (4.9), assuming the existence of β_i^j satisfying (4.11) and (4.14).

If we subtract the second equation in (4.11) multiplied by $\mathbf{x}_\sigma$ from the first equation multiplied by $\mathbf{x}_\tau$, and then multiply both equations in (4.11) by $\mathbf{N}_\sigma$ and $\mathbf{N}_\tau$, we get four independent equations

$$\begin{aligned}
0 &= \beta_\tau^\sigma E + (\beta_\tau^\tau - \beta_\sigma^\sigma) F - \beta_\sigma^\tau G \,, \\
1 &= \beta_\sigma^{\sigma 2} E + 2 \beta_\sigma^\sigma \beta_\sigma^\tau F + \beta_\sigma^{\tau 2} G \,, \\
0 &= \beta_\sigma^\sigma \beta_\tau^\sigma E + (\beta_\tau^\sigma \beta_\sigma^\tau + \beta_\sigma^\sigma \beta_\tau^\tau) F + \beta_\sigma^\tau \beta_\tau^\tau G \,, \\
\sin^2 \sigma &= \beta_\tau^{\sigma 2} E + 2 \beta_\tau^\sigma \beta_\tau^\tau F + \beta_\tau^{\tau 2} G \,.
\end{aligned}$$

These equations are consistent with (4.14). Solving for E, F, G and rearranging the result, we obtain

$$\begin{aligned}
E &= \frac{1}{K^2} \left[(\beta_\sigma^\tau \sin \sigma)^2 + \beta_\tau^{\tau 2} \right] > 0 \,, \\
G &= \frac{\sin^2 \sigma}{K^2} \left[(\beta_\sigma^\tau \sin \sigma)^2 + \beta_\sigma^{\sigma 2} \right] > 0 \,, \\
EG - F^2 &= \frac{\sin^2 \sigma}{K^2} > 0 \,.
\end{aligned} \tag{4.15}$$

These three functions are positive definite in a non-singular point of the Gaussian parameterization and they satisfy the compatibility equations together with L, M, N if $\beta_i^j \in \mathcal{C}^1$ [Sto69]. L, M, N can be obtained from (4.11) and (4.15) in a similar way as the E, F, G. Thus the necessary and sufficient conditions for the existence of a surface $\mathbf{x} = \mathbf{x}(\sigma, \tau)$ of class $\mathcal{C}^3$, defined in a neighborhood of (σ_0, τ_0), for which E, F, G, L, M, N are the first and the second fundamental coefficients, are satisfied. The surface represented is unique except for position in space. □

Note: Additional necessary conditions for derivatives of β_i^j to define a surface $\mathbf{x} \in \mathcal{C}^3$ follow from the conditions $(\mathbf{x}_\sigma)_\tau = (\mathbf{x}_\tau)_\sigma$, $(\mathbf{x}_\sigma)_{\sigma\tau} = (\mathbf{x}_\sigma)_{\tau\sigma}$, and $(\mathbf{x}_\tau)_{\sigma\tau} = (\mathbf{x}_\tau)_{\tau\sigma}$. We do not discuss them here because they are not important for the algorithms presented in Chapter 5.

4.3 The Difference Angle between Normal Projection and Image Gradient in $\mathcal{L}$-Plane

When analyzing a single monocular intensity image with the aim to find a relation to the underlying surface, we need to consider properties of the image function I at a neighborhood of an image point (x, y). We have already seen that suitable local descriptors with a geometric interpretation are the elements of the class $\mathcal{F}_I^{*2}$. We will now look at the relation between the normal $\mathbf{N}$ projection onto the $\mathcal{L}$-plane and the elements of $\mathcal{F}_I^{*2}$.

Suppose first that there is an intensity image I of a *spherical surface* $\mathbf{x}(\sigma, \tau)$ illuminated from the vantage point. We are interested in the difference angle ϑ between the image gradient $-\nabla I$ and the vector $\widetilde{\mathbf{N}}$, which is the orthogonal projection of the surface normal vector $\mathbf{N}$ onto the $\mathcal{L}$-plane. The Gaussian mapping now becomes equivalent to space-scaling by the factor of $1/R$, where R is the radius of the sphere. The set of surface τ-curves (surface isophotes) is projected onto the set of concentric circles in the $\mathcal{L}$-plane. Therefore, the image gradient curves are straight in the $\mathcal{L}$-plane and coincide with the projection of the surface σ-curves. The projection of vector $\mathbf{N}_\sigma$ has the same tilt as $\mathbf{N}$, which follows from (4.9). Since any direction on an umbilical surface is a principal direction, it holds by Rodrigues's formula [Sto69] that $\mathbf{N}_\sigma = -\beta_\sigma^\sigma \mathbf{x}_\sigma$, which finally implies that the projection of the vector $\mathbf{x}_\sigma$ (tangent to the σ-curve) onto image plane has the same tilt as the projection of $\mathbf{N}$. Therefore the image gradient has the same tilt as the surface normal and $\vartheta \equiv 0$ on a sphere illuminated from vantage point. This property of spherical surfaces has already been observed by Lee & Rosenfeld in [LR85].

The difference ϑ vanishes identically for a broader class of all *surfaces of revolution* viewed and illuminated from their axis of rotation. The isophotes become parallel to latitude circles in this case and they project to concentric circles in the image plane. Another class of surfaces satisfying $\vartheta \equiv 0$ is the class of all cylinders oriented so that their axis is parallel to the image plane. In this case, the isophotes form a system of parallel straight lines running along the surface rulings independent of the position of an infinitely distant illuminant.[6]

On a general surface, the tilt difference ϑ does not necessarily vanish. We will now derive a general expression that relates ϑ to the image properties from the set $\mathcal{F}_I^{*2}$, assuming *no a priori restriction* on the surface form in the sense of restricting the surface type like Pentland in [Pen84], Lee & Rosenfeld in [LR85], or Ferrie & Levine in [FL89], or the surface order like Arnspang in [Arn91], or assuming an explicit local surface approximation like Oliensis in [Oli90]. Note again that it is assumed that the image is projected on the $\mathcal{L}$-plane rather than on a general $\mathcal{V}$-plane. The generalization to an arbitrary case of *known* $\mathbf{L} \neq \mathbf{V}$ is possible using a similar approach.

Assuming an infinitely distant observer (and orthogonal projection), a derivative of a vector projection $(\tilde{\mathbf{x}})_i$ is equivalent to a projection of the vector derivative $\widetilde{\mathbf{x}_i}$,

$$(\tilde{\mathbf{x}})_i = [\mathbf{x} - (\mathbf{x}\cdot\mathbf{V})\mathbf{V}]_i = \mathbf{x}_i - (\mathbf{x}_i\cdot\mathbf{V})\mathbf{V} = \widetilde{\mathbf{x}_i}\,.$$

Then, if $\mathbf{u}$ and $\mathbf{v}$ are two independent directions (see Fig. 4.3), the derivatives of a vector $\tilde{\mathbf{x}}$ can be written in the form

$$\begin{array}{rclcl} \tilde{\mathbf{x}}_\mathbf{u} & = & (\tilde{\mathbf{x}})_\sigma\,\sigma_\mathbf{u} + (\tilde{\mathbf{x}})_\tau\,\tau_\mathbf{u} & = & \widetilde{\mathbf{x}_\sigma}\,\sigma_\mathbf{u} + \widetilde{\mathbf{x}_\tau}\,\tau_\mathbf{u}\,, \\ \tilde{\mathbf{x}}_\mathbf{v} & = & \cdots & = & \widetilde{\mathbf{x}_\sigma}\,\sigma_\mathbf{v} + \widetilde{\mathbf{x}_\tau}\,\tau_\mathbf{v}\,. \end{array} \tag{4.16}$$

Before we continue, the relation of $\tau_\mathbf{u}$, $\tau_\mathbf{v}$, $\sigma_\mathbf{u}$, $\sigma_\mathbf{v}$ to the members of $\mathcal{F}_I^{*2}$ for a *Lambertian surface* will be established. It allows us to simplify the analysis. The general case leads to similar results.

We know that the slant is constant along an isophote if the reflectance function becomes space-invariant or when $\mathbf{L} = \mathbf{V}$, i.e.,

$$\sigma_\mathbf{v} = 0\,.$$

From (4.1) and from $I = \alpha\,\mathbf{N}\cdot\mathbf{L} = \alpha\,\cos\sigma$ follows that

$$\sigma_\mathbf{u} = -\frac{D_\mathbf{u}\cos\sigma}{\sin\sigma} = \frac{1}{\alpha}\,\frac{-I_\mathbf{u}}{\sin\sigma} = \frac{1}{\alpha}\,\frac{\rho}{\sin\sigma} = \frac{I}{\alpha}\,\frac{\rho}{I}\,\frac{1}{\sin\sigma} = \frac{\alpha\,\cos\sigma}{\alpha}\,\rho^*\,\frac{1}{\sin\sigma} = \rho^*\,\cot\sigma\,,$$

where ρ is the isophote density and ρ^* is the isophote density scaled by the image value. The advantage of ρ^* over ρ is in its invariance to the Lambertian surface albedo α.

Now, we will show how the $\tau_\mathbf{u}$ and $\tau_\mathbf{v}$ are related to the isophote and the flow-line curvatures, respectively. The curvature of a *planar* curve $\mathcal{C}$ can be defined as an angular

[6] Also rulings of a slanted cylinder project on straight isophotes. Nevertheless, in this case $\vartheta \neq 0$.

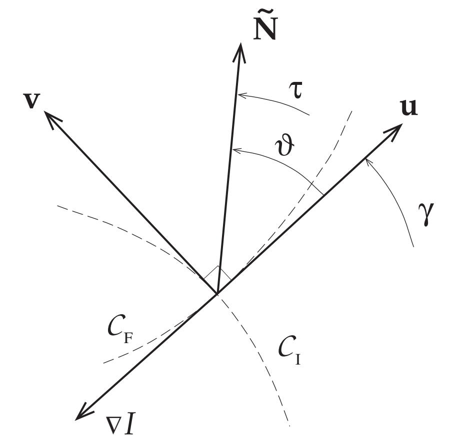

Figure 4.3: *The relation between the surface normal vector $\tilde{\mathbf{N}}$ projected on $\mathcal{L}$-plane and the vectors $\mathbf{u}$ and $\mathbf{v}$ that are tangent to image flow-line $\mathcal{C}_{\mathrm{F}}$ and isophote $\mathcal{C}_{\mathrm{I}}$, respectively. The angles τ and ϑ represent the tilt of the surface normal in the $\mathcal{L}$-plane and the deviation of the image gradient $-\nabla I$ from $\tilde{\mathbf{N}}$, respectively.*

turn $d\theta$ of its tangent vector after a move by a unit arc dt along the curve. More specifically,

$$\kappa(t) = \frac{d\theta}{dt}\,,$$

where t is the natural parameter of $\mathcal{C}$, and θ is the angle between the vector tangent to the curve at a point located at $t = t_0$ and a reference vector in an outer coordinate system. The curvature can have assigned a sign if the direction of the curve, i.e., the direction in which the parameter t increases, remains well-defined.

If τ is the tilt of the surface normal projected onto the image plane and γ is the orientation of the image flow-line (see Fig. 4.3), we can write

$$\begin{aligned} s\,\kappa_{\mathrm{F}} &= D_{\mathbf{u}}\gamma = D_{\mathbf{u}}(\tau - \vartheta) = \tau_{\mathbf{u}} - \vartheta_{\mathbf{u}}\,, \\ s\,\kappa_{\mathrm{I}} &= D_{\mathbf{v}}(\gamma \pm \frac{\pi}{2}) = D_{\mathbf{v}}(\tau - \vartheta) = \tau_{\mathbf{v}} - \vartheta_{\mathbf{v}}\,, \end{aligned}$$

where s is an extrinsically given sign related to the sign of the Gaussian curvature that is locally unknown except for special cases (points on the occlusion contour). This relation will be explained later.[7] From (4.2) it can be seen that both the signs of κ_{I} and κ_{F} change simultaneously when the orientation of the vector $\mathbf{u}$ reverses.

[7] Loosely speaking, while the direction of increasing slant σ along an image flow-line is well-defined in an intensity image, the direction of increasing tilt along an image isophote is not—the natural direction of an isophote remains undefined.

To sum up, the unknown scalar functions in (4.16) are related to the members of $\mathcal{F}_I^{*2}$ as follows:

$$\begin{aligned}
\sigma_{\mathbf{v}} &= 0\,, && (4.17)\\
\sigma_{\mathbf{u}} &= \rho^*\,\cot\sigma\,, && (4.18)\\
\tau_{\mathbf{u}} &= s\,\kappa_{\mathrm{F}} + \vartheta_{\mathbf{u}}\,, && (4.19)\\
\tau_{\mathbf{v}} &= s\,\kappa_{\mathrm{I}} + \vartheta_{\mathbf{v}}\,. && (4.20)
\end{aligned}$$

For the sake of shortening the subsequent formulas, we will use the unevaluated quantities, except for $\sigma_{\mathbf{v}} = 0$, in the following analysis.

The vectors $\widetilde{\mathbf{x}}_{\mathbf{u}}$ and $\widetilde{\mathbf{x}}_{\mathbf{v}}$ in (4.16) can be chosen such that they are equivalent to the unit tangent vectors to the image flow-line $\mathcal{C}_{\mathrm{F}}$ and the image isophote $\mathcal{C}_{\mathrm{I}}$, respectively:

$$\begin{aligned}
\widetilde{\mathbf{x}}_{\mathbf{u}} &= \mathbf{u}\,,\\
\widetilde{\mathbf{x}}_{\mathbf{v}} &= \mathbf{v}\,.
\end{aligned}$$

From (4.16) follows that

$$\begin{aligned}
\widetilde{\mathbf{x}_\tau} &= \frac{1}{\tau_{\mathbf{v}}}\,\mathbf{v}\,,\\
\widetilde{\mathbf{x}_\sigma} &= \frac{1}{\sigma_{\mathbf{u}}}\left(\mathbf{u} - \frac{\tau_{\mathbf{u}}}{\tau_{\mathbf{v}}}\,\mathbf{v}\right)\,.
\end{aligned} \tag{4.21}$$

Weingarten equations must hold for the projections of surface normal vector derivatives:

$$\widetilde{\mathbf{N}_i} = -\beta_i^\sigma\,\widetilde{\mathbf{x}_\sigma} - \beta_i^\tau\,\widetilde{\mathbf{x}_\tau}\,,\quad i = \sigma,\,\tau\,.$$

Substituting derivatives of $\mathbf{N}$ and (4.21) into these two equations and solving for β_i^j, we obtain

$$\begin{aligned}
\beta_\sigma^\sigma &= -\sigma_{\mathbf{u}}\,\cos\sigma\,\cos\vartheta\,,\\
\beta_\sigma^\tau &= -\cos\sigma\,(\tau_{\mathbf{v}}\,\sin\vartheta + \tau_{\mathbf{u}}\,\cos\vartheta)\,,\\
\beta_\tau^\sigma &= \sigma_{\mathbf{u}}\,\sin\sigma\,\sin\vartheta\,,\\
\beta_\tau^\tau &= -\sin\sigma\,(\tau_{\mathbf{v}}\,\cos\vartheta - \tau_{\mathbf{u}}\,\sin\vartheta)\,.
\end{aligned} \tag{4.22}$$

Finally, substituting for β_σ^τ and β_τ^σ into the existence condition (4.14), we arrive at the desired equation for ϑ:

$$2\,\sigma_{\mathbf{u}}\,\sin\vartheta + \sin 2\sigma\,(\tau_{\mathbf{v}}\,\sin\vartheta + \tau_{\mathbf{u}}\,\cos\vartheta) = 0\,, \tag{4.23}$$

which evaluates to

$$\rho^*\,\sin\vartheta + \sin^2\sigma\,\Big[(\vartheta_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\sin\vartheta + (\vartheta_{\mathbf{u}} + s\,\kappa_{\mathrm{F}})\cos\vartheta\Big] = 0\,. \tag{4.24}$$

This is a differential equation that expresses the desired relation between the isophote properties and the angle between the image gradient vector and the projection of the surface normal vector on the image plane. It is satisfied for every non-singular point of the image of a non-developable surface[8] $\mathbf{x}(\sigma, \tau)$ of class at least $\mathcal{C}^3$.

It can be seen that four different solutions for ϑ comply with (4.23) because the common sign s of the image isophote and the image flow-line remains undetermined, and the equation holds both for ϑ and $\vartheta + \pi$ within the basic interval $\langle 0, 2\pi)$ for ϑ.

For the purposes of the following text, we express the Gaussian and mean curvatures from (4.22):

$$K = \frac{1}{2}\,\sigma_{\mathbf{u}}\,\tau_{\mathbf{v}}\,\sin 2\sigma\,, \tag{4.25}$$

$$H = -\frac{1}{2}\left[\sigma_{\mathbf{u}}\,\cos\sigma\,\cos\vartheta + \sin\sigma\,(\tau_{\mathbf{v}}\,\cos\vartheta - \tau_{\mathbf{u}}\,\sin\vartheta)\right]. \tag{4.26}$$

Let us look more closely at the case $\vartheta = 0$ now. We already know that ϑ vanishes along the occlusion contour as well as in images of spheres and some other surfaces viewed and illuminated from a special position. The equation (4.23) defines the general condition for ϑ to vanish at a non-singular point. Rearranging it, we get

$$[2\,\sigma_{\mathbf{u}} + \tau_{\mathbf{v}}\,\sin 2\sigma]\sin\vartheta + \tau_{\mathbf{u}}\,\sin 2\sigma\,\cos\vartheta = 0\,. \tag{4.27}$$

The term at $\sin\vartheta$ can be expressed using (4.25):

$$[2\,\sigma_{\mathbf{u}} + \tau_{\mathbf{v}}\,\sin 2\sigma] = \frac{2}{\sigma_{\mathbf{u}}}\,(\sigma_{\mathbf{u}}^2 + K)$$

and it vanishes only on a *hyperbolic surface* out of a singular point of the Gaussian parameterization and out of the occlusion contour when

$$K = -\sigma_{\mathbf{u}}^2\,.$$

If, in addition, there is a rotational symmetry in the image of a hyperbolic surface (the image isophotes are concentric), it holds that $\tau_{\mathbf{u}} = 0$. In that case the equation (4.27) vanishes identically. But this is an isolated singular case in the sense that a small perturbation of the surface or small displacement of the illuminant breaks the symmetry. The equation (4.24) is therefore well-posed for a generic surface.

If the terms at $\sin\vartheta$ and $\cos\vartheta$ in (4.27) do not vanish simultaneously, the cause of vanishing ϑ is

$$\tau_{\mathbf{u}} = 0\,. \tag{4.28}$$

[8] A developable surface cannot be parameterized in the Gaussian manner, see also Section 4.5.

It is evident that this condition covers the special cases mentioned at the beginning of this section where we spoke about surfaces satisfying $\vartheta \equiv 0$. If we evaluate (4.22) under the conditions $\vartheta = 0$ and the (4.28) conjecture, we will get $\beta_\sigma^\tau = \beta_\tau^\sigma = 0$ in a non-singular image point that is not located on the occlusion contour. But this implies that the surface parametric curves become locally tangent to lines of curvature.

Therefore the general condition for ϑ to vanish out of the occlusion contour and singular points is that the parametric curves of Gaussian parameterization locally coincide with the lines of curvature on the surface. The converse is also true.

4.4 The Sign of the Gaussian Curvature on Occlusion Boundaries Revisited

The aim of this subsection is to explain the relation between the unknown sign s that appeared in (4.24) and the Gaussian curvature at a point on the occlusion boundary. It will be shown that the sign can be specified in certain regions in the image with the possibility to reduce the image interpretation ambiguity by propagating this information over most of the image.

We have shown in Section 3.1.1 that along the occlusion contour, it holds that $\vartheta \equiv 0$. Therefore also $\vartheta_{\mathbf{v}} = 0$ holds because isophotes, like any other curves approaching the occlusion contour, are parallel to it. It follows, according to (4.25), (4.18), (4.1) and (4.20), that

$$K \to s\,\kappa_{\mathrm{I}}\,\rho^*\,\cos^2\sigma = s\,\kappa_{\mathrm{I}}\,\rho\,\cos\sigma$$

when approaching the contour. Let us assume we approach it from the visible part of the surface. Then $\sigma < \pi/2$ and $\cos\sigma \to 0_+$. The isophote density tends to explode to infinity, $\rho \to +\infty$, because of projective foreshortening of the surface metric vector in the direction perpendicular to the contour. Since the isophote is parallel to the contour, it holds that $\kappa_{\mathrm{I}} = \widehat{\kappa}_t$, which is the curvature of the occluding contour introduced in Section 3.1.2. From the expression for Gaussian curvature given on page 35, stating that $K = \kappa_s\,\widehat{\kappa}_t$, where κ_s is the unknown *but positive* curvature of the surface normal section parallel to viewing direction at a point on the occlusion boundary, it follows that

$$0 \;\leq\; \kappa_s = \lim_{\sigma \to \left(\frac{\pi}{2}\right)_-} \rho\,\cos\sigma \;<\; \infty\,.$$

Therefore, at a point in the image, the unknown orientation of image isophotes and image flow-lines that is represented by the sign s can be specified, *knowing that the point in question is located on the occlusion contour.*

The information about the surface hyperbolicity can be propagated from the occlusion contour inside the visible surface patch. We hypothesize that, if there is a $\mathcal{C}^2$ pre-image of a path connecting an image point $\widehat{\mathbf{x}}$ with a point on the occlusion contour, a correct decision about the local Gaussian curvature sign can be made at $\widehat{\mathbf{x}}$. The problem of reducing the ambiguity of the local image interpretation by propagating this information needs further investigation. The work of Oliensis on a similar problem suggests some solutions [OD93]. From the practical point of view, a successful solution requires a computationally efficient representation of the surface structure [Šár92, PVM+93, Šár94].

Our experiments have shown that the expected ill-posedness of numerical computations of K near the occlusion contour (due to the term $\rho \cos\sigma$) does not play a significant role and the results are correct. See the discussion of experimental results given in Section 6.2.4.

4.5 Shading on Developable Surfaces

Surfaces that can be isometrically mapped onto a plane are called *developable surfaces*[9] and they are defined as regular surfaces whose Gaussian curvature K vanishes everywhere. As mentioned earlier, the Gaussian parameterization makes no sense on a developable, since it is not related to isophote properties: the σ- and τ-curves do not form two independent systems of curves because the Gaussian image of a developable is a single curve on the Gaussian sphere. The assumption made in the analysis given in the previous section is thus violated and (4.24) does not describe the properties of local shading on a developable surface. The problem might be avoided by defining the surface parameterization in a way that, instead of the tilt τ, one could choose a parameter t such that the t-curve always projects onto an image isophote. A possible choice is

$$t = \frac{\tau}{K}\,,$$

where K is the Gaussian curvature. This needs further investigation, however.

In the rest of this subsection, we will prove that the direction of isophotes on a developable surface is always parallel to the surface rulings, which means that a surface isophote projects to a single point on Gaussian sphere.

Theorem 4.3

Let $\mathbf{w}$ be a line element of a surface isophote. It is conjugate to the projection of a vector $\mathbf{W}$ on the surface tangent plane at the point in question. $\mathbf{W}$, lying in the principal plane of the imaging geometry, is determined by the geometry and by the reflectance function properties.

Proof. The theorem states that $\mathbf{W}^* \cdot D_{\mathbf{w}}\,\mathbf{N} = 0$, where

$$\mathbf{W}^* = \frac{f}{\mathbf{W}\cdot\mathbf{N}}[\mathbf{W} - (\mathbf{W}\cdot\mathbf{N})\mathbf{N}]$$

[9] We also call them *developables.*

is a (perspective) projection of $\mathbf{W}$ on the surface tangent plane. It is readily understood that, if we move along isophote, i.e., in the direction of $\mathbf{w}$, then the change of brightness must vanish:

$$0 = D_{\mathbf{w}}\, I = D_{\mathbf{w}}\, R(\mathbf{N}\cdot\mathbf{L},\, \mathbf{N}\cdot\mathbf{V},\, \mathbf{V}\cdot\mathbf{L}) = (R_1\mathbf{L} + R_2\mathbf{V})\cdot D_{\mathbf{w}}\,\mathbf{N}\,,$$

where $R_i,\ i = 1, 2$ are the derivatives of R with respect to i-th argument. The vector $R_1\,\mathbf{L} + R_2\,\mathbf{V}$ lies in the principal plane of the imaging geometry spanned by vectors $\mathbf{L}$ and $\mathbf{V}$. We set $\mathbf{W} = R_1\,\mathbf{L} + R_2\,\mathbf{V}$. Then

$$0 = D_{\mathbf{w}}\, I = \mathbf{W}\cdot D_{\mathbf{w}}\,\mathbf{N} = \frac{\mathbf{W}\cdot\mathbf{N}}{f}\,\mathbf{W}^*\cdot D_{\mathbf{w}}\,\mathbf{N}\,,$$

since $\mathbf{N}\cdot D_{\mathbf{w}}\,\mathbf{N} = 0$ and $\mathbf{W}\cdot\mathbf{N} \neq 0$ in general. This means that the dot-product $\mathbf{W}^*\cdot D_{\mathbf{w}}\,\mathbf{N}$ must vanish. Therefore, by the definition of conjugate directions [Lip69], the direction parallel to $\mathbf{w}$ in the surface tangent plane and the direction parallel to the vector

$$\mathbf{W}^* = (R_1\,\mathbf{L}^* + R_2\,\mathbf{V}^*)$$

are conjugate. In the case of a space-invariant reflectance function, or when $\mathbf{L} = \mathbf{V}$, the vector $\mathbf{W}^*$ is parallel to the projection of the light direction vector onto the surface tangent plane. □

Corollary 4.3.1

> *The isophote direction is parallel to the zero-curvature principal direction independently of the imaging geometry at a parabolic point on a surface* $\mathbf{x}(s,t)$.

Let the surface parameterization (s,t) be chosen so that the parametric curves' tangent vectors become parallel to principal directions at each point on the surface, and so that the principal curvature corresponding to t-direction is $\kappa_1 = 0$. Such a parameterization always exists and is unique for a given surface [Sto69]. Then, in a direction $d\mathbf{x} = \mathbf{x}_s\,ds + \mathbf{x}_t\,dt$ on the surface, the normal vector changes about

$$d\mathbf{N} = \mathbf{N}_s\,ds + \mathbf{N}_t\,dt = -\kappa_2\mathbf{x}_s\,ds$$

according to Rodrigues's formula [Sto69, Lip69]. We choose another direction on the surface $\delta\mathbf{x} = \mathbf{x}_s\,\delta s + \mathbf{x}_t\,\delta t$ that is conjugate to $d\mathbf{x}$. Then $0 = d\mathbf{N}\cdot\delta\mathbf{x} = -\kappa_2\,\mathbf{x}_s\cdot\mathbf{x}_s\,ds\,\delta s$, since $\mathbf{x}_t\cdot\mathbf{x}_s = 0$. It follows that $ds = 0$, since $\|\mathbf{x}_s\| > 0$, $\kappa_2 \neq 0$, and δs is generally non-zero. It means that *any* direction at a parabolic point, including the projection of the vector $\mathbf{W}$ on the surface, is conjugate to the zero-normal-curvature direction (i.e., the isophote direction), as follows from Theorem 4.3.

In other words, *the isophote direction at a parabolic point on the surface is invariant to the imaging geometry.* It also follows that the crossing of isophote and a parabolic line gives rise to a singular point of the isophote parameterization. This happens due to $d\tau = 0$ at the crossing when the isophote is parameterized by the tilt τ. If we admit these singular points on isophotes, the Gaussian parameterization can be defined there. This is the reason why we do not exclude points of vanishing Gaussian curvature from consideration in Theorem 4.2 as far as the curvature does not vanish identically.

Corollary 4.3.2

> *Isophotes on developable surfaces are always parallel to the rulings of the surface.*

This is true because the zero-principal-curvature direction is parallel to rulings everywhere, which is known from elementary differential geometry. When the light source moves, only the isophote density changes.

Chapter 5

Local Orientation and Local Shape Quality from Shading

The material presented here summarizes our results in *LSA* based on the theoretical analysis given in the previous chapter. The first several sections discuss proposed algorithms for local shading analysis that increase in power and complexity. The text structure will be similar to the structure of Chapter 3 but without explaining the theoretical background that is given in Chapter 4. Here we aim at giving more "technical details" related to the algorithmic implementation. All the algorithms presented here were implemented and tested. The results of computational experiments are discussed in Chapter 6.

5.1 The Existence of a Second-Order *LSA* Algorithm

In general, the members of the set ${\mathcal{F}_I}^n$—the set of local features of an intensity image of a properly illuminated surface—are related to the differential surface properties, global geometry parameters, and photometric surface properties.

Definition 5.1

Let $\mathcal{G} = \{\sigma,\, \tau,\, \beta^\sigma_\sigma,\, \beta^\tau_\sigma,\, \beta^\tau_\tau,\, \sigma_\mathcal{L},\, \tau_\mathcal{L},\, \alpha\}$, where

$\sigma,\, \tau$	*are the local surface slant and tilt, respectively,*
β^j_i	*are three of the four Weingarten coefficients (note that $\beta^\sigma_\tau = \beta^\tau_\sigma \sin^2\sigma$ for Gaussian parameterization),*
$\sigma_\mathcal{L},\, \tau_\mathcal{L}$	*are the light direction slant and tilt, respectively,*
α	*is the surface albedo.*

A theoretical LSA algorithm is a system of relations between the elements of $\mathcal{G}$ and computationally detectable *image properties ${\mathcal{F}_I}^n$, where n is the order of the algorithm.*

In the previous chapter, the relation to the first two types of properties has been discussed under the assumption that *the illuminant is located at the vantage point.* We will keep this assumption. Suppose further that the reflectance function R is known. In this case there are six quantities of ${\mathcal{F}_I}^2$ related to six unknown parameters: σ, τ, β^σ_σ, β^τ_σ, β^τ_τ, α. The corresponding six independent equations can theoretically be solved for these unknowns.

Consider the set $\mathcal{F}^{*2}_I$ (see 4.6). The curvatures κ_I and κ_F and the relative isophote density change ν are related *to partial derivatives of the Weingarten coefficients* as well. (Or, equally, to the second derivatives of the coefficients of the second fundamental form of the surface.) In the previous chapter, we proposed an approach that allows to "hide" these derivatives into derivatives of the difference angle ϑ and let them possibly recover by means of a relaxation algorithm based on a relatively well-posed relation (4.23). There is probably nothing like that for the slant σ. Some information about the slant is contained in the density change ν. But the relation of this quantity to the underlying surface shape is rather complicated and, moreover, ν is heavily dependent on the surface reflectance function: if the image irradiance equation is in the form $I = \alpha\, R(\mathbf{N}\cdot\mathbf{L})$, assuming either constant $\mathbf{L} = \mathbf{V}$ or space invariant R, we have, according to (4.5):

$$\nu = \frac{I_\mathbf{uu}}{I_\mathbf{u}} = \frac{R''\,(\mathbf{N_u}\cdot\mathbf{L})^2 + R'\,\mathbf{N_{uu}}\cdot\mathbf{L}}{R'\,\mathbf{N_u}\cdot\mathbf{L}}.$$

Even if the reflectance function becomes linear (i.e., $R'' = 0$) and ν thus becomes invariant to R', the term $\mathbf{N_{uu}}\cdot\mathbf{L}$ has no direct relation to the local surface properties discussed above. In

$$\mathbf{N_{uu}}\cdot\mathbf{L} = -\left(\frac{\cos\sigma}{\sin^2\sigma}\,\rho^2 + \sigma_\mathbf{uu}\,\sin\sigma\right), \tag{5.1}$$

namely in the unknown parameter $\sigma_\mathbf{uu}$, there occur derivatives of Weingarten coefficients. Therefore, except for trivial surfaces with constant curvature (like spheres), the relative isophote density change ν can probably not be used for *LSA*. Experiments with disregarding $\sigma_\mathbf{uu}$ and using only the first part of (5.1) did not lead to acceptable results. A reliable relaxation algorithm has not been found.

Hypothesis 2 (The Fundamental Ambiguity of *LSA*)

> *There is probably no theoretical LSA algorithm of second order giving results with finite ambiguity even under maximum restrictions on surface reflectance function and imaging geometry, supposing the viewed surface is general.*

Definition 5.2

> *A heuristic LSA algorithm is a system of relations between the ${\mathcal{F}_I}^n$ image properties and the parameters introduced above under heuristic restricting assumptions about the viewed surface, imaging geometry, or surface reflectance that remove the infinite ambiguity of theoretical LSA.*

Goal 1

> *The aim of this work is to find the most realistic heuristics for LSA to work with minimum ambiguity, without loosing any possible solutions.*

This is a very vague notion, however. Suitable heuristics can depend on the circumstances under which the algorithm is expected to work. Since the aim of *LSA* is to recover local geometry-related information about an unknown general surface, we will prefer those heuristics that restrict the *photometric properties* of the surface or the *global imaging geometry* over the heuristics restricting the *geometric properties* of the surface. In practice, the global geometry as well as the photometric properties can be more or less controlled.[1] What cannot usually be controlled is the *shape* of the surface that is to be recovered.

Some acceptable solutions to the problem will be given later in this chapter.

5.2 Recovering Local Orientation and Surface Curvature from Images of Spheres

Equation (4.23) expresses the necessary condition for ϑ to be a difference angle between the image gradient direction and the projection of normal vector on $\mathcal{L}$-plane. If ϑ, the

[1] The surface can be painted and the light source can be designed so that it satisfies the requirements. If either of those requirements cannot be guaranteed, *LSA* is the misleading approach to the problem and a different technique has to be chosen.

normal slant σ, and the normalized image gradient

$$\frac{-\nabla I}{||\nabla I||} = [\cos\gamma,\ \sin\gamma]^{\mathrm{T}}$$

are known at a pixel, the surface normal vector is given by

$$\mathbf{N} = [\cos(\gamma+\vartheta)\ \sin\sigma,\ \ \cos(\gamma+\vartheta)\ \sin\sigma,\ \ \cos\sigma]^{\mathrm{T}}\,.$$

Suppose first that

$$\vartheta = k\,\pi, \quad k = 0, 1 \tag{5.2}$$

everywhere on the viewed surface *patch* (k is locally unknown). This is true for the class of surfaces mentioned at the beginning of Section 4.3. On all the other surfaces, zero represents an estimator for ϑ.

From (5.2), (4.25), (4.26), and (4.17) – (4.20) it follows that

$$\begin{aligned} K &= s\,\kappa_{\mathrm{I}}\,\rho^*\ \cos^2\sigma\,, \\ H &= \frac{(-1)^{k+1}}{2}(\rho^*\cot\sigma\ \cos\sigma + s\,\kappa_{\mathrm{I}}\ \sin\sigma), \quad k = 0, 1\,. \end{aligned} \tag{5.3}$$

The undetermined sign for the mean curvature is due to $\cos\vartheta$ term that evaluates to ± 1.

A necessary condition for a surface to be spherical is $H^2 - K = 0$ because the principal curvatures κ_1, κ_2 are equal. From (5.3) follows that this condition is satisfied when

$$\tan\sigma = \pm\frac{\rho^*}{\kappa_{\mathrm{I}}}\cot\sigma. \tag{5.4}$$

Since the scaled density ρ^* is always positive, the only allowed sign for κ_{I} is $\kappa_{\mathrm{I}} \geq 0$. The Gaussian curvature in (5.3) will therefore always be evaluated as positive, which is true for a spherical surface but it is not true for a hyperbolic surface. Notice that (5.4) represents a slant estimator invariant to (locally constant) Lambertian albedo.

The corresponding Algorithm 5 on the next page gives therefore correct results on Lambertian spherical surfaces illuminated from a distant vantage point by a point light source. The performance worsens as the true surface deviates from being spherical and the light source moves from the assumed position. If the surface deviates too far from a sphere and becomes hyperbolic, two of the four possible solutions (the two corresponding to the hyperbolic solution) will be completely lost when using the image gradient as a tilt-estimator. The advantage of the corresponding slant-estimator is that it cannot fail since the arc-tangent function has an infinite domain and thus the estimator always has a single real solution independently of the violation of the surface sphericity assumption.

Algorithm 5 (Local Shading Analysis for Spheres)

requirements

Illuminant is located at the vantage point, $\mathbf{L} = \mathbf{V}$. The surface is Lambertian and spherical.

For each pixel (x, y) that is not located at a singular point of the input image $I(x, y)$ do

1. *Compute the scaled density ρ^*, the image isophote curvature κ_{I}, and the image gradient direction γ (see (4.1) and (4.4)).*
2. *The surface normal slant is*

$$\sigma := \arctan\left(\sqrt{\frac{\rho^*}{|\kappa_{\mathrm{I}}|}}\right),$$

and the surface normal tilt and curvatures are

$$\begin{aligned} \tau(k) &:= \gamma + k\pi\,, \\ K &:= \frac{\kappa_{\mathrm{I}}^2 \rho^*}{|\kappa_{\mathrm{I}}| + \rho^*}\,, \\ H(k) &:= (2k-1)\sqrt{\frac{\kappa_{\mathrm{I}}^2 \rho^*}{|\kappa_{\mathrm{I}}| + \rho^*}}\,. \end{aligned}$$

The sign of the recovered K will always be positive because of (5.4) and (5.3), and the value of k gives two local solutions for the surface according to the following table:

k	surface at (x, y)
0	convex elliptic
1	concave elliptic

5.3 A Reasonable Estimator of Local Surface Slant

Violation of the sphericity assumption in the previous slant and tilt estimators causes a loss of a suitable slant estimator within ${\mathcal{F}_I}^2$ when the albedo becomes unknown and the surface becomes non-spherical. We have seen in Section 5.1 that, without an explicit assumption about the surface, there exists probably no *LSA* algorithm. The reason is that there exists no slant-estimator. The purpose of this section is to show that, under some weaker assumptions, a reasonable slant estimator can exist.

Let us suppose that the reflectance function R is known and the albedo α is constant but unknown. These are realistic conditions for *LSA* to work. The local image function singularities fall into one of the categories mentioned in Section 3.2. If the category of the singularity *is known* and the local orientation $\mathbf{N}$ at such a point can be specified (see Section 3.2.3), the irradiance equation $I = \alpha\, R$ can be solved for α, because R is then a known function of the global surface geometry. The simplest case mentioned in the literature on *LSA* is an apex on Lambertian surface [LR85, FL89, FL92]. At the apex $\mathbf{L} = \mathbf{N}$ and the image value will be $I = \alpha$ if the reflectance function is space-invariant.

The requirement of $\alpha =$ const. need not always be satisfied. Distinguishing the kind of singular image point in single monocular images without any additional information (motion, stereopsis, photometric stereo) may be impossible as well. Therefore, a reasonable choice for image point candidates (x, y) satisfying $I(x, y) = \alpha$ are those singularities, where the image function attains its maximum value I_{max} *under the assumption of a space-invariant reflectance function or the illuminant located at the vantage point.* The singularities on parabolic lines have a corresponding image value less than I_{max} because there $\mathbf{N} \neq \mathbf{L}$, in general. Then, the only image singularities where $I = I_{max} = \alpha$ are those where $\mathbf{N} = \mathbf{L}$. This requires that such a point is present in the image.

The corresponding slant estimator is

$$\cos\sigma \approx \frac{I}{I_{max}}\,.$$

This estimator has already been suggested by several authors [LR85, FL89, FL92] and we therefore accepted it for Algorithm 6 on the next page. The influence of the bad slant estimation on the resulting tilt estimate depends on the type of the tilt estimator; this is discussed in Subsection 6.2.6 of the next chapter.

5.4 Estimator of Local Tilt for a General Surface

We would like to broaden the narrow class of surfaces where Algorithm 5 yields correct results. Although there is no single-step algorithm giving *exact results,* a good estimator (at least within the ${\mathcal{F}_I}^2$ class) can be provided or a relaxation algorithm can be designed. The estimator will be described in this section and the relaxation algorithm will be dealt with in the next section.

If the condition (5.2) is relaxed, we need an independent image feature related to the difference angle ϑ. As may be seen from the analysis in Section 4.3, a suitable candidate is the flow-line curvature. By rearranging (4.24) we get

$$\tan\vartheta = -\frac{(\vartheta_{\mathbf{u}} + s\,\kappa_{\mathrm{F}})\sin^2\sigma}{\rho^* + (\vartheta_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\sin^2\sigma}\,.$$

If we disregard the derivatives of the difference angle ϑ, we get an estimator:

$$\tan\vartheta \approx -\frac{s\,\kappa_{\mathrm{F}}\,\sin^2\sigma}{\rho^* + s\,\kappa_{\mathrm{I}}\,\sin^2\sigma}\,. \tag{5.5}$$

Algorithm 6 (Single-Step Local Shading Analysis for General Surfaces)

requirements

Illuminant is located at the vantage point, $\mathbf{L} = \mathbf{V}$. Surface reflectance is Lambertian.

For each pixel (x, y) that is not located at a singular point of the input image $I(x, y)$ do

1. *Compute the scaled density ρ^*, the isophote curvature κ_{I}, the flow-line curvature κ_{F}, and the image gradient direction γ (see* (4.1), (4.4), *and* (4.3)*).*
2. *Compute the local slant from the re-scaled value of $I(x, y)$*

$$\sigma \quad := \quad \arccos \frac{I(x, y)}{I_{max}} \,,$$

and surface normal tilt and curvatures from

$$\begin{aligned}
\vartheta(k, s) &:= -\arctan \frac{s\, \kappa_{\mathrm{F}} \sin^2 \sigma}{\rho^* + s\, \kappa_{\mathrm{I}} \sin^2 \sigma} + k\, \pi, \\
\tau(k, s) &:= \gamma + \vartheta(k, s)\,, \\
K(s) &:= s\, \rho^*\, \kappa_{\mathrm{I}} \cos^2 \sigma\,, \\
H(k, s) &:= -\frac{1}{2 \sin \sigma} \Big\{ \rho^* \cos^2 \sigma \cos \vartheta(k, s) + \\
&+ \quad s \sin^2 \sigma \left[\kappa_{\mathrm{I}} \cos \vartheta(k, s) - \kappa_{\mathrm{F}} \sin \vartheta(k, s) \right] \Big\}\,.
\end{aligned}$$

The sign $s = +1, -1$ and the value of $k = 0, 1$ give four local interpretations of the surface.

This estimator is used in Algorithm 6 and it gives very good results. The behaviour of this algorithm in comparison with the algorithms published in literature will be discussed in Chapter 6.

5.5 A Relaxation Algorithm for Normal Tilt Recovery

The Algorithm 6 from the previous section can yield a good first estimate of ϑ. We will use the expression (4.24) to construct a *local* relaxation algorithm. The motivation for the maximum algorithm locality is to preserve the locality of the process of surface feature recovery that is required for local surface interpretation. This requirement also implies that the algorithm must converge fast to reduce the extent of image data around a point in question that is necessary to support the final interpretation.

We can develop the expression (4.24)

$$J(\vartheta) = \left[\rho^* + (\vartheta_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\sin^2\sigma\right]\,\sin\vartheta + (\vartheta_{\mathbf{u}} + s\,\kappa_{\mathrm{F}})\sin^2\sigma\,\cos\vartheta = 0 \tag{5.6}$$

into a Taylor series:

$$J(\vartheta + \delta\vartheta) = J(\vartheta_0) + \frac{dJ(\vartheta)}{d\vartheta}\delta\vartheta + O(\delta\vartheta^2) \approx 0, \tag{5.7}$$

where $O(\delta\vartheta^2) \to 0$ as $\delta\vartheta \to 0$, and

$$\frac{dJ(\vartheta)}{d\vartheta} = J'(\vartheta) = \left[\rho^* + (\vartheta_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\sin^2\sigma\right]\,\cos\vartheta - (\vartheta_{\mathbf{u}} + s\,\kappa_{\mathrm{F}})\sin^2\sigma\,\sin\vartheta\,. \tag{5.8}$$

The new estimate for ϑ can then be obtained from the old one after a single Newton-Raphson step:

$$\vartheta^n = \vartheta^{n-1} - \xi\,\frac{J(\vartheta^{n-1})}{J'(\vartheta^{n-1})}\,, \tag{5.9}$$

where ξ is a (constant) convergence factor, usually $\xi \leq 1$. The Newton-Raphson process is known to converge very fast. A problem with the convergence stability can occur when both $J' \to 0$ and $J \to 0$, however. To overcome this difficulty we can use the method proposed in [TS92] that showed a very good behaviour. The authors propose substituting $J'(\vartheta^{n-1})^{-1}$ in (5.9) by a filter K^n that has the following properties:

1. $K^n \approx J'(\vartheta^{n-1})^{-1}$,
2. $K^n \to 0$ when $J'(\vartheta^{n-1}) \to 0$,
3. $K^n \to 0$ when $\vartheta^{n-1} \to \vartheta$.

A suitable choice is

$$K^n = \frac{P^{n-1}J'(\vartheta^{n-1})}{R + P^{n-1}J'(\vartheta^{n-1})^2}\,,$$

where R is a small non-negative constant guaranteeing the first requirement, and the positive definite variable P^{n-1} is an "expectation" of $\vartheta^{n-1} = \vartheta$ guaranteeing the last requirement. In [TS92], Tsai & Shah propose updating P according to the formula

$$P^n = \begin{cases} P_0 & \text{when } n = 0\,, \\ \left[1 - K^n\,J'(\vartheta^{n-1})\right]\,P^{n-1} & \text{when } n > 0. \end{cases}$$

This procedure has a formal relation to Kalman filtering [May79]. We may rewrite (5.7) into the form of

$$- v^n = J(\vartheta^{n-1}) + J'(\vartheta^{n-1})\,(\vartheta^n - \vartheta^{n-1}), \tag{5.10}$$

where v is a random variable whose properties are related to the properties of the remainder $O(\delta\vartheta^2)$ in (5.7). The upper index n is related to "discrete time" of the filtering process. We will suppose that v has the properties of zero-mean uncorrelated Gaussian noise with statistics

$$\begin{aligned} E[v^i] &= 0 \\ E[v^i\, v^j] &= \begin{cases} R^i & i = j \\ 0 & i \neq j \end{cases} \end{aligned} \tag{5.11}$$

where R is the noise covariance. These requirements, namely the whiteness of the noise, are unrealistic, however. Our experiments showed that violating this condition is not critical in practice.

To establish the relation to Kalman filtering, we can treat ϑ^n in (5.10) as an unknown system variable. The estimation is based on the set of measurements of a linear system output z^n, more specifically,

$$z^n = H^n\,\vartheta^n + v^n,$$

where z^n is a *discrete-time measurement process,* H^n is a *measurement scalar* transforming the system state on the observed system output, ϑ^n is a *state variable,* and v^n is discrete-time white Gaussian measurement noise with statistics (5.11).

We assume that the unknown system is *static, autonomous,* and *time-varying.* Then an estimate $\widehat{\vartheta}^n$ of the system variable ϑ^n is updated by defining the Kalman filter gain K^n according to:

$$\begin{aligned} K^n &= \frac{P^{n-1}\,H^n}{R^n + P^{n-1}\,(H^n)^2}\,, \\ \widehat{\vartheta}^n &= \widehat{\vartheta}^{n-1} + K^n\left[z^n - H^n\,\widehat{\vartheta}^{n-1}\right], \\ P^n &= (1 - K^n H^n)\,P^{n-1}\,, \end{aligned} \tag{5.12}$$

where

$$\begin{aligned} \widehat{\vartheta}^0 &= \widehat{\vartheta}_0\,, \\ P^0 &= P_0 \end{aligned}$$

are the initial estimates for the state and for the state conditional covariance, respectively [May79].

Returning back to our analogy, it follows from (5.10) that $H^n \approx J'(\widehat{\vartheta}^{n-1})$ and the measured output is

$$z^n = -J(\widehat{\vartheta}^{n-1}) + J'(\widehat{\vartheta}^{n-1})\, \widehat{\vartheta}^{n-1}.$$

The expression $\left[z^n - H^n\, \widehat{\vartheta}^{n-1}\right]$ in (5.12) then becomes

$$\left[z^n - H^n\, \widehat{\vartheta}^{n-1}\right] = -J(\widehat{\vartheta}^{n-1}),$$

therefore the update step in our case is

$$\begin{aligned} K^n &= \frac{P^{n-1}\, J'(\widehat{\vartheta}^{n-1})}{R^n + P^{n-1}\left(J'(\widehat{\vartheta}^{n-1})\right)^2}, \\ \widehat{\vartheta}^n &= \widehat{\vartheta}^{n-1} - K^n\, J(\widehat{\vartheta}^{n-1}), \\ P^n &= (1 - K^n J'(\widehat{\vartheta}^{n-1}))\, P^{n-1}, \end{aligned} \tag{5.13}$$

which is formally identical with (5.9), but possesses the desired properties. Algorithm 7 on the next page incorporates this scheme.

The discrete-space formulation of the process is as follows. The continuous function $\vartheta(x, y)$ can be substituted by a two-parametric discrete function, where x, y are image coordinates. The $\vartheta_{\mathbf{u}}$ and $\vartheta_{\mathbf{v}}$ functions in (5.6) and (5.8) are computed as central differences

$$\begin{aligned} \vartheta_{\mathbf{u}} = D_{\mathbf{u}}\, \vartheta \approx \Delta^{\epsilon}_{\mathbf{u}}\, \vartheta = \frac{\vartheta(u+\epsilon, v) - \vartheta(u-\epsilon, v)}{2\,\epsilon}, \\ \vartheta_{\mathbf{v}} = D_{\mathbf{v}}\, \vartheta \approx \Delta^{\epsilon}_{\mathbf{v}}\, \vartheta = \frac{\vartheta(u, v+\epsilon) - \vartheta(u, v-\epsilon)}{2\,\epsilon}, \end{aligned} \tag{5.14}$$

and the discrete equivalents of ρ^*, κ_{F}, and κ_{I} are evaluated from the image data according to (4.1), (4.3), and (4.4), respectively. The update process (5.13) remains unchanged.

5.6 When Does the Tilt Estimator Based on (4.24) Become Ill-Posed?

Both the Algorithms 6 and 7 are based on equation (4.24) given on page 69. We have seen in Section 4.3 that the equation becomes ill-posed when

$$K = -{\sigma_{\mathbf{u}}}^2 \quad \text{and} \quad \tau_{\mathbf{u}} = 0\,.$$

In this case the equation (4.24) vanishes identically. This may happen in the image of a hyperbolic surface.

Another cause of ill-posedness of (4.24) is when the point in question lies just at an image of a parabolic point. Let us suppose that the surface is non-developable.

Algorithm 7 (Relaxation Local Shading Analysis for General Surfaces)

requirements

Same as for Algorithm 6.

1. *Initialize images $\mathbf{\Theta}^0(x,y)$ to the values of $\vartheta_0(x,y)$ computed by Algorithm 6, $\mathbf{P}^0(x,y)$ to P_0, and a scalar variable R to R_0.*
2. *For $n = 0, \ldots, N$ and $s = +1, -1$ do steps 3 and 4:*
3. *At each pixel (x,y) compute the directional differences in the image gradient direction ∇I and in the perpendicular direction $\nabla_\perp I$:*

$$\vartheta^n_{\mathbf{u}} := -\,\Delta^{\epsilon}_{\nabla I}\,\mathbf{\Theta}^n(x,y), \qquad \vartheta^n_{\mathbf{v}} := -\,\Delta^{\epsilon}_{\nabla_\perp I}\,\mathbf{\Theta}^n(x,y)\,,$$

and compute J^n and $(J^n)' = dJ^n/d\vartheta$ from

$$J^n = \left[\frac{\rho^*}{\sin^2\sigma} + (\vartheta^n_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\right]\,\sin\vartheta^n + (\vartheta^n_{\mathbf{u}} + s\,\kappa_{\mathrm{F}})\,\cos\vartheta^n.$$

4. *Perform the following updates*

$$\begin{aligned} K^{n+1} &:= \frac{\mathbf{P}^n(x,y)\,(J^n)'}{R + \mathbf{P}^n(x,y)\,(J^n)'^2}\,, \\ \mathbf{\Theta}^{n+1}(x,y) &:= \mathbf{\Theta}^n(x,y) - \xi\,K^{n+1}\,J^n\,, \\ \mathbf{P}^{n+1}(x,y) &:= (1 - K^{n+1}\,(J^n)')\,\mathbf{P}^n(x,y)\,. \end{aligned}$$

5. *For the combinations of $k = 0, 1$ and $s = +1, -1$ compute the final results at each pixel (s, k have the same meaning as in Alg. 6):*

$$\begin{aligned} \vartheta(k,s) &:= -\arctan\frac{\vartheta^N_{\mathbf{u}} + s\,\kappa_{\mathrm{F}}}{\dfrac{\rho^*}{\sin^2\sigma} + \vartheta^N_{\mathbf{v}} + s\,\kappa_{\mathrm{I}}} + k\,\pi\,, \\ \tau(k,s) &:= \gamma + \vartheta(k,s)\,, \\ K(s) &:= s\,\rho^*\,(\vartheta^N_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\,\cos^2\sigma\,, \\ H(k,s) &:= \frac{-1}{2\,\sin\sigma}\,\Big\{\rho^*\cos^2\sigma\cos\vartheta(k,s) + \\ &+ \sin^2\sigma\left[(\vartheta^N_{\mathbf{v}} + s\,\kappa_{\mathrm{I}})\cos\vartheta(k,s) - (\vartheta^N_{\mathbf{u}} + s\,\kappa_{\mathrm{F}})\sin\vartheta(k,s)\right]\Big\}\,. \end{aligned}$$

The convergence constant $\xi = 0.25$ proved to be a suitable choice between the convergence speed and stability. $R_0 = 10^{-3}$ was found suitable for noiseless synthetic images. In practise, there is no need to exceed 10 iterations. The positive value $P_0 = 1$ is not critical.

Then the τ-curve touches the zero-curvature principal direction on the surface and the σ-curve generally crosses the τ-curve so that the Gaussian parameterization in the vicinity of the parabolic point exists (at the parabolic point, there is a singular point of the parameterization due to $d\tau = 0$ along the τ-curve). The cause of the ill-posedness may be the exploding magnitude of the first fundamental form determinant, which can be seen from (4.15) on page 65. This may be avoided by choosing a modified Gaussian parameterization such that we substitute a parameter $K\,t$ for τ as has already been suggested in Section 4.5. The determinant of the first fundamental form becomes

$$E\,G - F^2 = \sin^2\sigma.$$

But such a parameter substitution does not substantially affect (4.24). The equation therefore remains well-posed at a parabolic point on a non-developable surface.

We have already seen that equation (4.24) assumes the point in question is not located at a singular image point. This is the last cause of possible ill-posedness of the equation.

In practice, the ill-posedness is not critical for a generic surface. The operators for the ${\mathcal{F}_I}^2$ features computation have a finite extent which guarantees that the behaviour of the *LSA* algorithm near the images of parabolic lines or singular points is never worse than the behaviour of standard estimators. On the other hand, the proposed estimators give *significantly better* results out of the ill-posed regions. This issue is discussed in the next chapter.

Chapter 6

Computational Experiments

This chapter deals with numerous experiments possibly with the aim of proposing and utilizing a suitable methodology for evaluating such algorithms that could generally be called "needle map from local shading." Because of the lack of such a methodology, we discuss it in more detail and present more material concerning accuracy, bias, and convergence (when relevant) of the various algorithms. The main attention will be paid to the algorithms presented in the previous chapter and to the comparison with some of the published algorithms.

The structure of this chapter differs from previous text—it stresses the experimental aspects over the theoretical ones. Basic facts about the test results are found mostly in the captions of figures and tables—this gives the reader an opportunity to study the results without turning the pages back and forth when skimming the text. Details and discussion are given in the main text.

Although the detection of ${\mathcal{F}_I}^n$ image properties represents a *necessary prerequisite* for a successful *LSA* algorithm implementation, we curiously give only little attention to it here. Some notes on robust detection of ${\mathcal{F}_I}^2$ image properties are given in Chapter 7.

All the algorithm tests presented here were run on *artificially generated data.* For the experiments presented in this chapter, we used only simple ${\mathcal{F}_I}^2$ image feature detectors based on (Gaussian) blurred image derivatives with small support fields. The experiments test mainly the influence of the violation of assumptions under which a particular algorithm is expected to work.

6.1 Standard Test

For the purpose of evaluating various surface geometric characteristic estimators from intensity images, we propose a test that is referred to as *the standard test.* The test defines *standard images, standard statistical evaluation* of the algorithm results, and *standard views* of 3-D mesh-plot data.

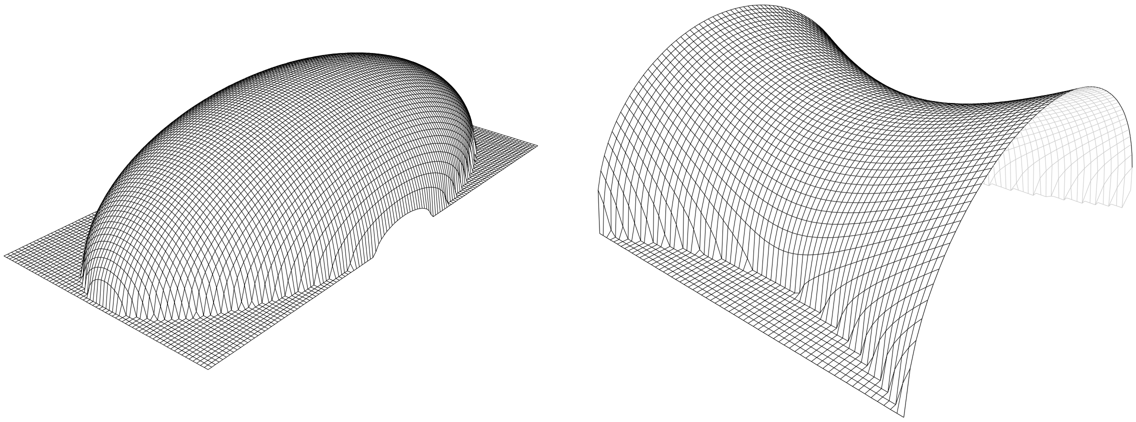

Figure 6.1: *Standard views of the standard test bodies. (a) ellipsoid with axes ratios* $x : y : z = 64 : 128 : 90$*, (b) hyperboloid with axes ratios* $x : y : z = -64 : 128 : 64$*. Both viewing vectors have a slant of* 65 *degrees from the* z*-axis and a tilt of* -50 *degrees from the* x*-axis of the global coordinate system.*

Standard Input Data

Images of *an ellipsoid* and *a hyperboloid of one sheet viewed from aside* have been chosen as standard *input grey-level data.* The three principal axes of both bodies are aligned with the axes of the image coordinate system and with the direction perpendicular to the image plane (i.e., with the axes of the $\mathcal{V}$-coordinate system, see Fig. 2.3 on page 24). Both bodies have Lambertian reflectance and they are *properly illuminated* from the vantage point (except for the test in Subsection 6.2.6). For the purpose of *LSA,* the standard test surfaces are sufficiently complex bodies representing the convex and concave classes.

The standard test ellipsoid has axes ratios $x : y : z = 64 : 128 : 90$ and the standard test hyperboloid has $x : y : z = -64 : 128 : 64$ (in pixels). The mesh-plots of both surfaces are shown in Fig. 6.1. The corresponding images are shown in Figs. 6.2a on page 91 and 6.3a on page 92, respectively. The image sizes were chosen appropriately: the grey-level 4-byte floating-point images have the size of 128×256 pixels and 128×128 pixels for the ellipsoid and for the hyperboloid, respectively.

The curvature of the bodies was chosen so that:

- the discretization of the grey-level image is not the main cause for estimators to behave incorrectly,
- the mesh-plots of surfaces and the results of the integration of the data obtained

from *LSA* algorithms are easily visually comparable,

- the error maps are easily visually comparable.

All these represent restrictions on the test surface curvature—it should not be too high. An exception is the test of the surface eccentricity influence upon the tilt estimate that is presented in Subsection 6.2.5.

Standard Mesh-Plot Views

For the purpose of visual comparison of the reconstructed 3-D data, standard views of the mesh-plots of the two standard bodies are defined. The viewing vector and the data scale are chosen so that the artifacts of various reconstruction methods are well-visible.

The viewing vector for the mesh-plots of both surfaces has a slant of 65 degrees from the z-axis of the body and a tilt of -50 degrees from the x-axis. The z-scale measured relative to the longer size of the image is chosen to be 0.25 for the ellipsoid and 0.49 for the hyperboloid. The projection of the mesh plot is perspective. The plots are pasted into the text so that the relative scale remains unchanged in all mesh-plots of both surfaces.

Standard Statistical Evaluation

Absolute-error maps and *relative-error maps* were computed for the purpose of the statistical evaluation. The *minimum and maximum values, the mean value, the standard deviation, the 5%-quantile, the median, and the 95%-quantile* were evaluated.

Since the minimum and the maximum errors often do not yield any relevant information because of their occurrence due to bad computational behaviour of the image derivative operators on occlusion contours, we do not include them into most of the statistics presented in this text. Instead, 5%-quantile and 95%-quantile errors are used as representative quantities related to the “maximum deviation” of the estimates from the true values. Where the difference between the 5%-quantile and the 95%-quantile errors becomes minimal (except for the sign), we denote them together as a 5/95%-quantile error and give the absolute value of it.

The standard deviation statistic was used as a measure of the “average deviation” of the estimates from the true data.

The mean error computed over the error map was not included in many statistics presented here for the same reasons as the minimum and the maximum errors. For the purpose of evaluation of the estimator’s “bias,” the median value was used instead. This is because the two statistics do usually not differ substantially in its values and the median value is not affected by several outliers in the input data that can occur due to possible misbehaviour of the operator on occlusion contours. The term “bias” differs slightly from the statistical definition of the quantity [Háj69]: In the context of this work, we use the term *bias* for an *expectation of the estimate error* over the visible part of a closed $\mathcal{C}^2$ surface.

Before the mentioned statistics were evaluated, the background data[1] and the data in a two-pixel wide strip around the occlusion contour had been masked out. This does not mean the test algorithms were run on masked data. All the computations except for the error analysis were performed on non-masked images. The reason for the masking is two-fold:

1. The statistics are influenced by vanishing errors in the portions of the image that correspond to the background;
2. Near the occlusion, the ${\mathcal{F}_I}^2$ image feature detector fails due to incorrect results of the second-order difference operator. This is because of its current simple implementation using a 3×3 convolution kernel (support field) for repeated first-order difference computation. Since the purpose of the experiment was to show the computational performance of the estimators *on noiseless synthetic data,* we were not interested in extending the support field of the detectors that would have improved the behaviour. We also wanted the influence of isophote characteristics detector misbehaviour on the results to be minimal. Moreover, it has been shown here that the error of tilt estimation decreases when approaching the occlusion contour. The statistical results have therefore not been much influenced by suppressing their evaluation near the occlusion contour.

6.2 Evaluation of *LSA* Algorithms

The aim of this section is to give the reader a good insight into what the reasonable heuristics demanded in Goal 1 on page 77 should be. We will deal with several algorithms including those that have been newly proposed in the previous chapter. Various aspects of the *LSA* algorithms' misbehaviour will be discussed here. Our experiments have shown that the following are the bottlenecks of local shading analysis that apply even under *ideal image formation process:*

- the tilt estimates are inaccurate for a general surface even when computed from artificially generated data under optimum and known conditions (proper illumination, known reflectance function, etc.),
- the surface curvatures computed from the recovered needle maps therefore *can rarely be used for any reasoning about the underlying surface,*
- all the local surface characteristic estimates are more or less dependent on the knowledge of the exact illuminant position relative to the viewer,

[1]The data out of the projection of the test surface onto the image plane.

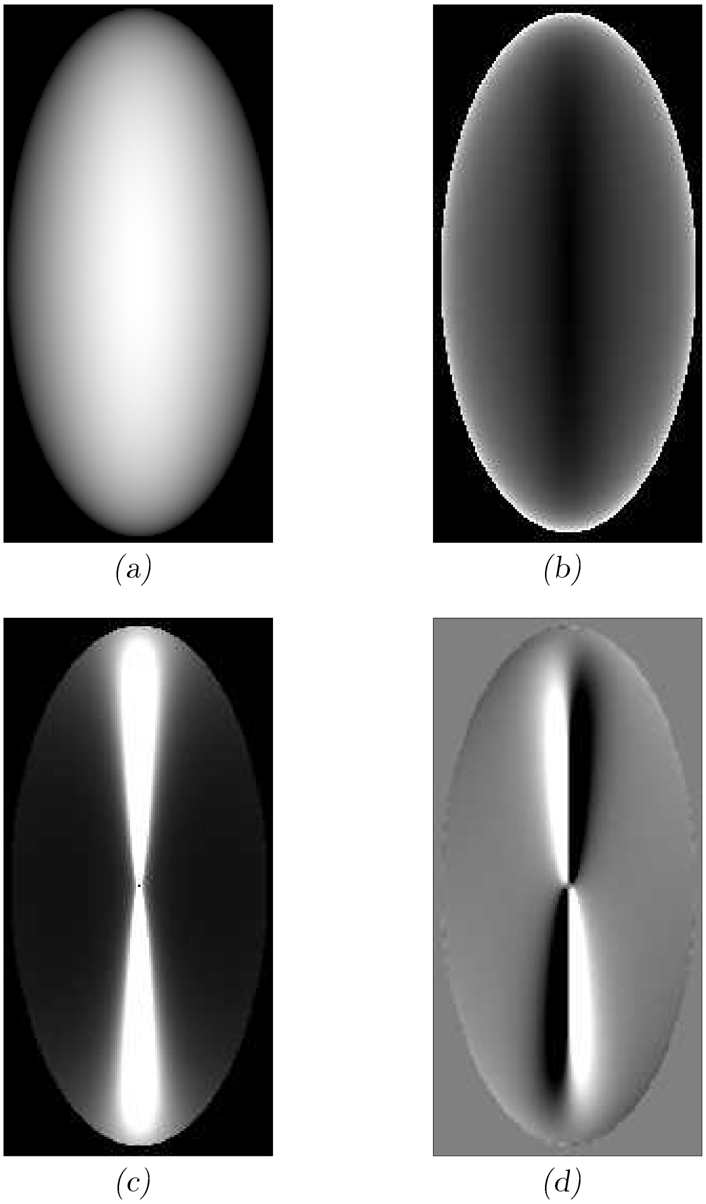

Figure 6.2: *Synthetic image of the test ellipsoid (a). The Lambertian surface is illuminated from the vantage point—the image corresponds to the* $\cos\sigma$ *map and the displayed range is* $\langle 0,\, 1\rangle$*; (b) is the grey-coded isophote density displayed in the range of* $\langle 0,\, 0.05\rangle$*; (c) is the isophote curvature displayed in the same range; (d) is the flow-line curvature displayed in the range of* $\langle -0.025,\, 0.025\rangle$*. The artifacts near the occlusion contour were masked out for the purpose of this display.*

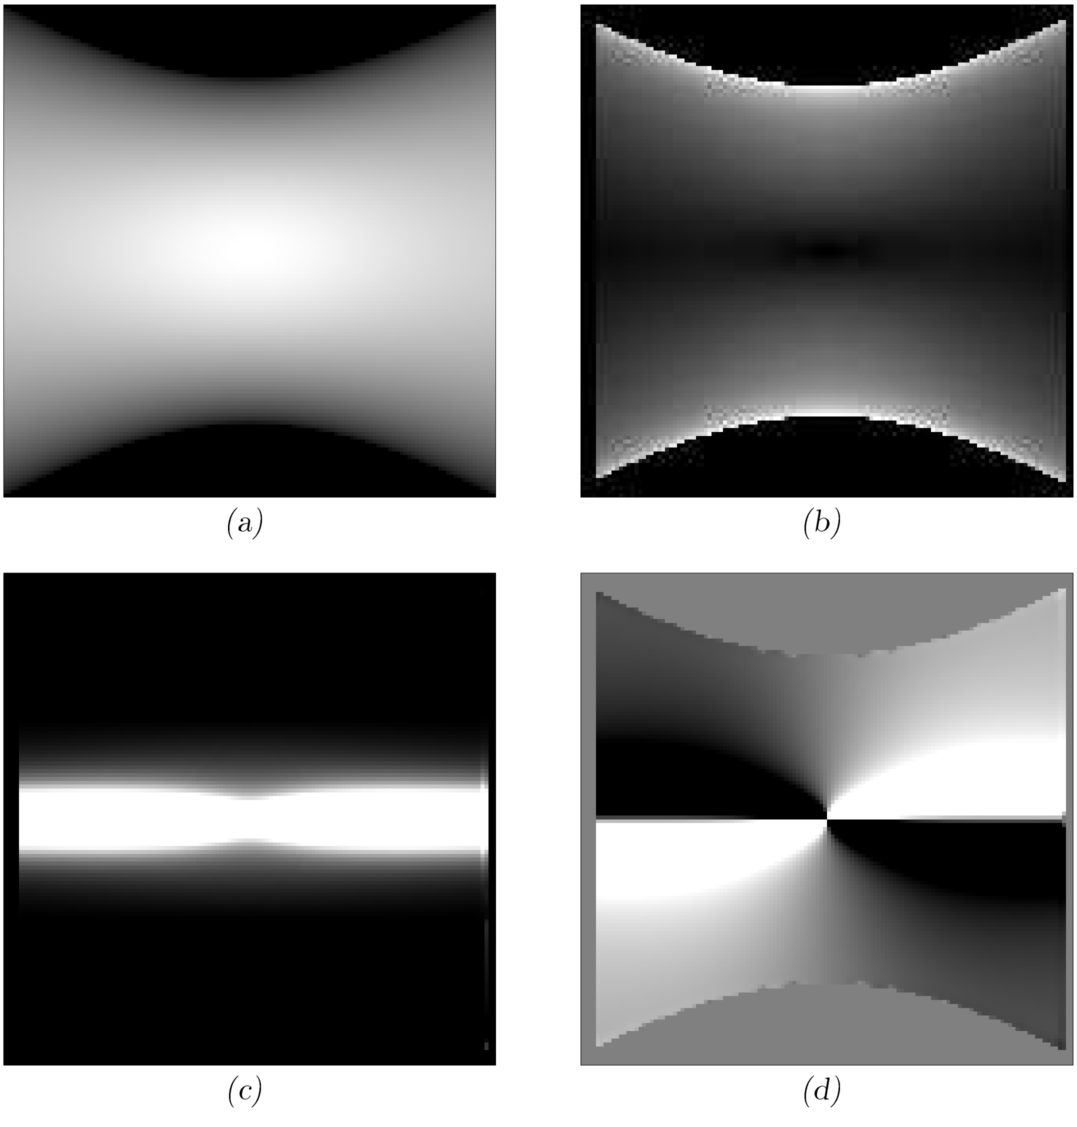

Figure 6.3: *Synthetic image of the test hyperboloid (a); (b) is the isophote density displayed in the range of $\langle 0, 0.05\rangle$; (c) is the isophote curvature in the range of $\langle -0.05, 0.05\rangle$; (d) is the flow-line curvature in the same range.*

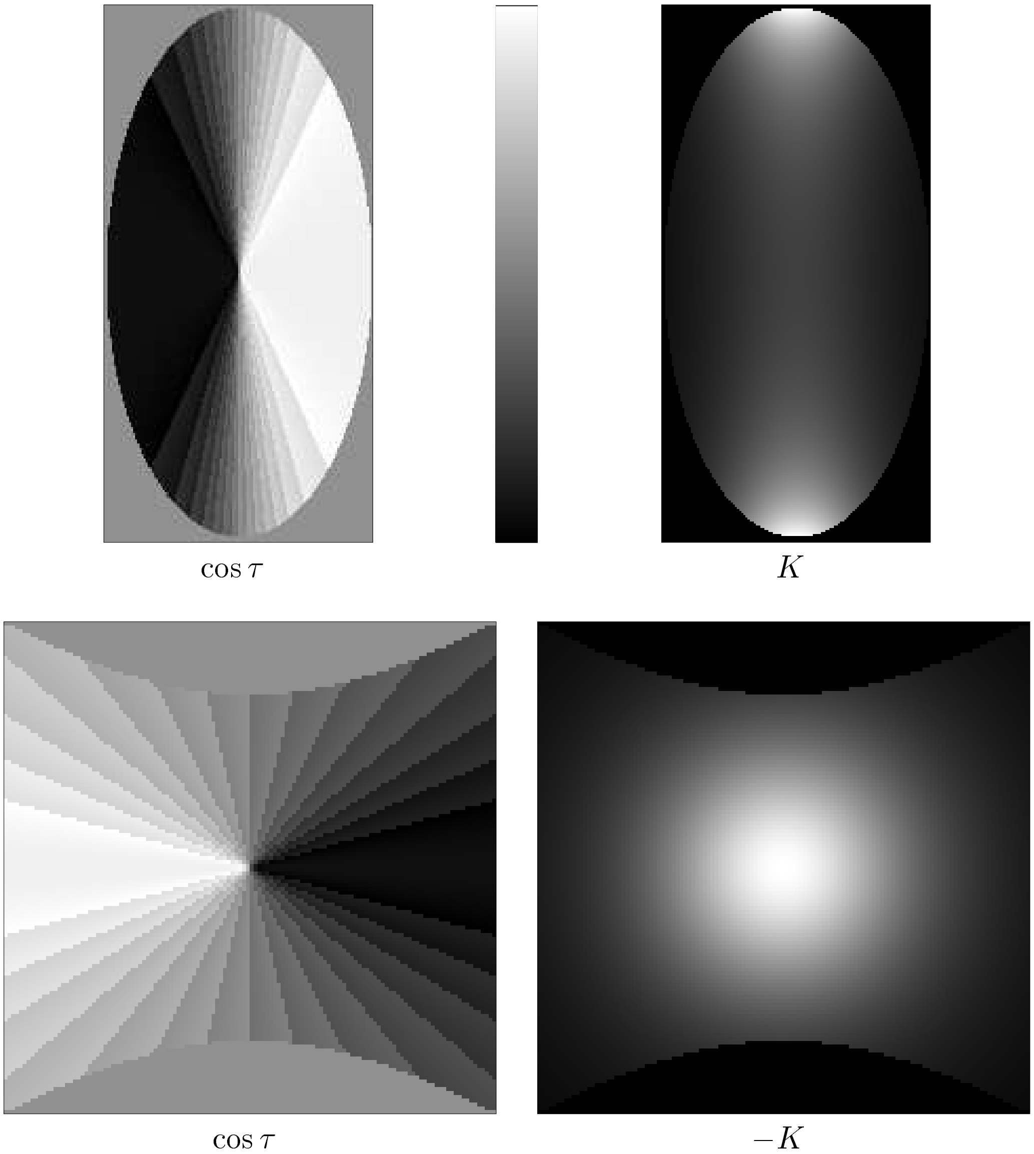

Figure 6.4: *Maps of the true local surface tilt ($\cos\tau$) and the true Gaussian curvature (K) of the test surfaces. All tilt maps in this chapter will be displayed in pseudo-grey-levels (using a saw-like lookup table) to see the structure of isogones. The displayed ranges are $\langle -1,\, 1\rangle$ in both cases. Notice that the isocontours of the tilt map are straight in both surfaces. The Gaussian curvature map is coded so that lighter parts correspond to higher absolute value of the curvature. The displayed value range is $\langle 0,\, 5\cdot 10^{-4}\rangle$ for the ellipsoid and $\langle -5\cdot 10^{-4},\, 0\rangle$ for the hyperboloid. The mapping into grey-levels is not fully monotonous due to imperfect grey-level data rendering in the printer. The PostScript mapping is shown in the linear vertical strip in between of the upper two maps.*

- all the estimates are influenced by the deviation of the surface reflectance function from the assumed one.

6.2.1 Tested Algorithms

We will compare the results of the proposed Algorithms 6 and 7 with two other *LSA* algorithms, both composed from a tilt and a slant estimator. The first one uses the image gradient tilt estimator (like in Alg. 5) and the second one uses a modification of Pentland tilt estimator [Pen84]. Both of them use the 'scaled image value slant estimator' discussed in Section 5.3. The tested algorithms will briefly be referred to as *the single-step algorithm, the relaxation algorithm, the image gradient algorithm, and the modified Pentland algorithm,* respectively.

The modification of Pentland tilt estimator is as follows. Pentland's original suggestion is to use the image Hessian eigenvectors as the estimator. There are up to four local solutions because it is not known in advance which of those vectors correspond to the true tilt. Assuming a spherical surface, the two solutions giving inconsistent normal slant estimates can be rejected. From the theory given in [Pen84] it follows that the slant is found as

$$\cos^2\sigma = \frac{\sin 2\tau\,(I_{xx}+I_{yy}) - 2\,I_{xy}}{\sin 2\tau\,(I_{xx}+I_{yy}) + 2\,I_{xy}}, \tag{6.1}$$

where τ is the corresponding tilt solution and I_{ij} are the image function derivatives. The rejection condition can therefore be easily evaluated. The two remaining solutions correspond to a convex and a concave local interpretations, respectively.

This rejection method cannot be used for a general surface because of the violation of the sphericity assumption—in the general case Equation (6.1) has no meaningful relation to the surface. If we want to use Pentland tilt estimator, we have to cope with all four solutions. Unfortunately, there are generally no loci where the number of local solutions decreases to less than four because there are always two perpendicular eigenvectors of the image Hessian, which is a symmetric matrix. This has the consequence that the solution which is locally "best" can suddenly "switch" to another best local solution when moving in some direction in the image. *This event in the image has no geometric relation to the underlying surface.* We have found that the switching of the tilt estimate to another local solution can even happen on the occlusion contour of the viewed surface if the slant inconsistency test is used to reject two of the four solutions. Moreover, when this rejection procedure is used, it can happen that all of the local solutions will be rejected in an image of a hyperbolic surface at a location sufficiently far from the singular image point.[2]

We have therefore modified Pentland's original tilt estimator so that it rejects the two solutions that differ most from the image gradient. Nevertheless, this modification

[2] In the image of the test hyperboloid, for instance, there is only a small region around the centre of the image where at least one solution for the tilt remains unrejected.

need not generally give correct results! For the purpose of the following evaluation on the two test surfaces, this modification gives the best solution.

6.2.2 Slant Estimators

Before we start experimenting with the algorithms mentioned in the previous section, we will present an experiment that should support our claims about the existence of a reasonable estimator of the local surface slant that has been discussed in Section 5.3 of the previous chapter.

This section compares the results of the following slant estimators: *the newly proposed spherical estimator (from Alg. 5), Pentland's spherical estimator* (6.1)*, and Pentland's heuristic biased estimator [Pen84].* All of these assume the surface is *locally spherical.* The purpose of this experiment is to show how they behave on images of non-spherical surfaces.[3]

Pentland's spherical slant estimator (6.1) has already been mentioned in Section 6.2.1. Theoretically, there are up to four possible solutions corresponding to the four possible solutions for the local surface tilt. In this experiment, the tilt estimate corresponding to the value closest to the image gradient was accepted as valid. The corresponding solution for the $\cos\sigma$ was computed as the square-root of the *absolute value* of the right side of (6.1).

Pentland's heuristic estimator confirms his results in texture analysis [Pen86]. The structure of the estimator is:

$$\sec\sigma \approx R\sqrt{\frac{\nabla^2 I}{I} - \frac{1}{R^2}}, \tag{6.2}$$

where R is an unknown constant referred to as the "radius of surface curvature" in Pentland's original text. We will call it *heuristic radius* since it can be determined only at a non-local scale. This information may not be available at the moment when the local analysis is performed. Our experiments showed that the minimum value of R satisfying the global constraints

$$\begin{aligned} 1 &\leq \sec\sigma\,, \\ R &\geq \frac{1}{\sqrt{\nabla^2 I}} \end{aligned}$$

over the *whole visible part* of the analyzed surface gives rise to slant estimates with minimum statistical error evaluated over the visible surface (see Fig. 6.5 on the next page).

[3] Pentland's heuristic estimator is the only one giving biased results on images of spheres. Pentland proposes a way to remove this bias assuming $\mathbf{L} = \mathbf{V}$ and a spherical surface. We do not present the results from the unbiased modification since its complexity and additional constraints on illuminant position do not guarantee a good behaviour for non-spherical surfaces.

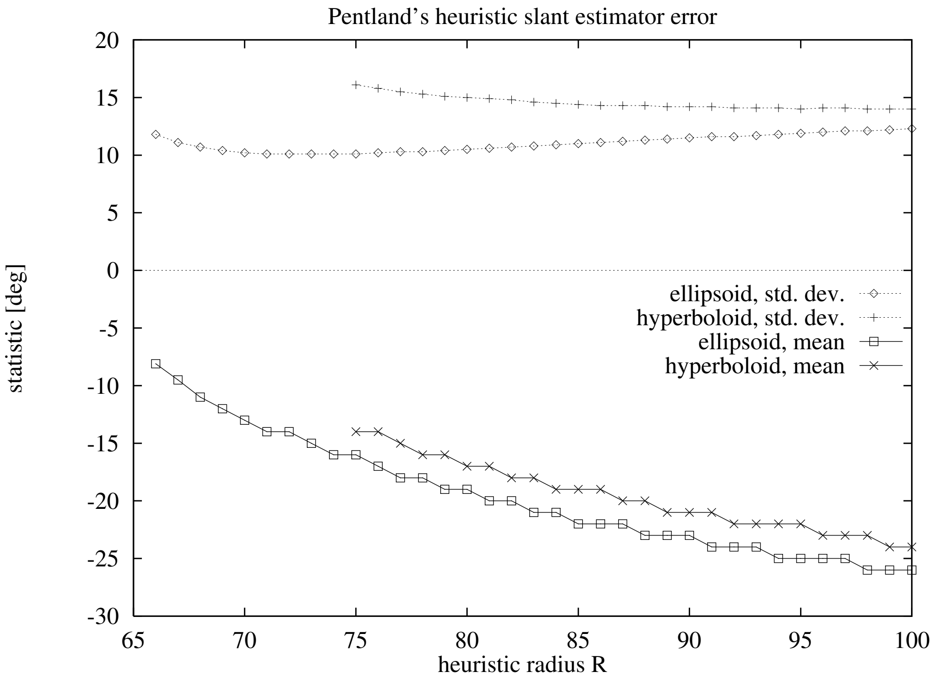

Figure 6.5: *The statistics of heuristic slant estimate error as functions of increasing Pentland's heuristic radius R. The solid line represents the mean error over the error map, the dashed line represents the standard deviation computed over the map. The mean error is minimal for the lowest acceptable radius R in both surfaces.*

The results for the test surfaces are presented in Figs. 6.6 on page 97 and 6.7 on page 98, respectively. The only estimator creating no "spurious extremal points" for the test ellipsoid is the proposed spherical estimator (Alg. 5), while all of the estimators create such points for the test hyperboloid.[4] Various statistics evaluated on error maps are also presented in Tab. 6.1 on page 99.

Evaluation

From the statistics follows that the proposed spherical estimator is slightly worse than Pentland's heuristic biased estimator. A great advantage of the proposed spherical estimator is that it is *unambiguous* and *works on a purely local basis.*

Neither of the tested estimators gives good results. On the ellipsoid, only the proposed spherical slant estimator and Pentland's heuristic estimator behave reason-

[4]We mean such loci in the recovered slant map where this map has a local maximum/minimum that does not corresponds to any geometric event on the surface.

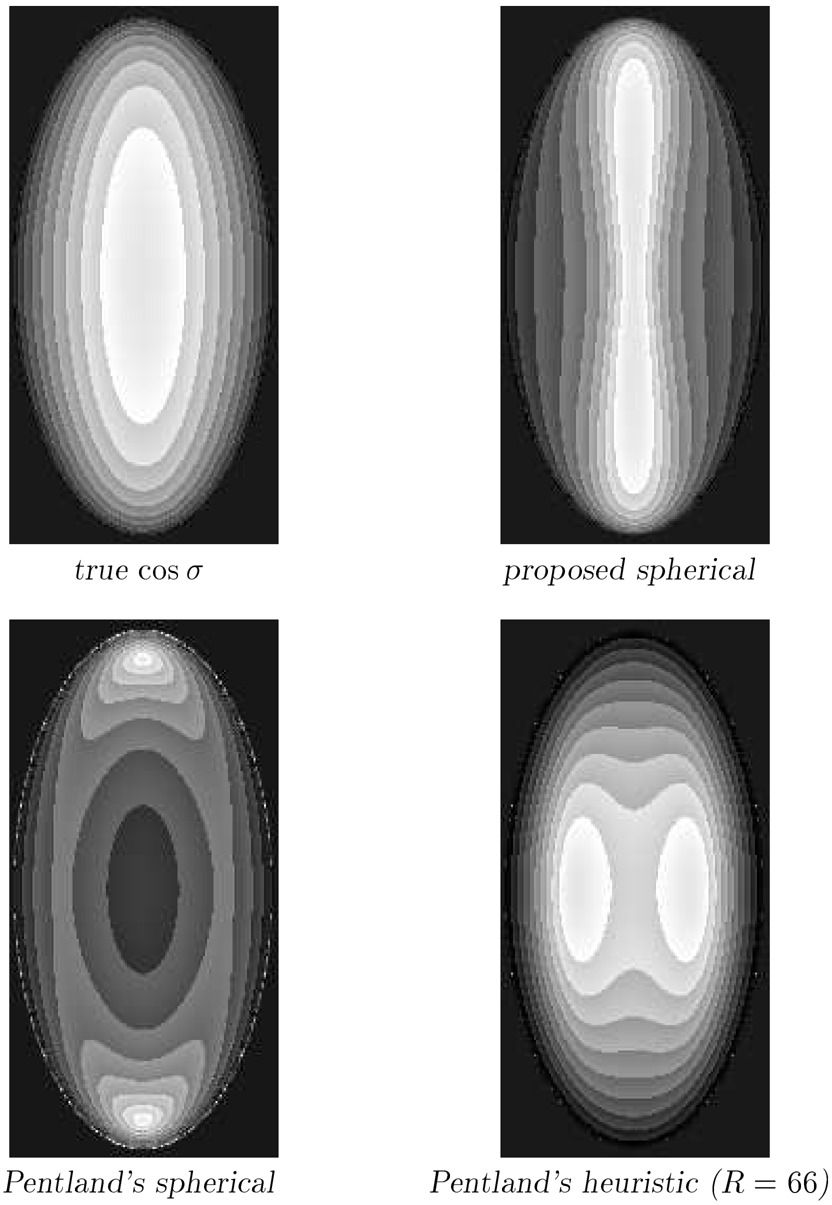

Figure 6.6: *Slant estimates for the test ellipsoid. To see the structure, the maps are displayed in pseudo grey-levels in the same way as in Fig. 6.4 (see also Fig. 6.2a). The displayed range is $\langle 0, 1\rangle$ in all cases.*

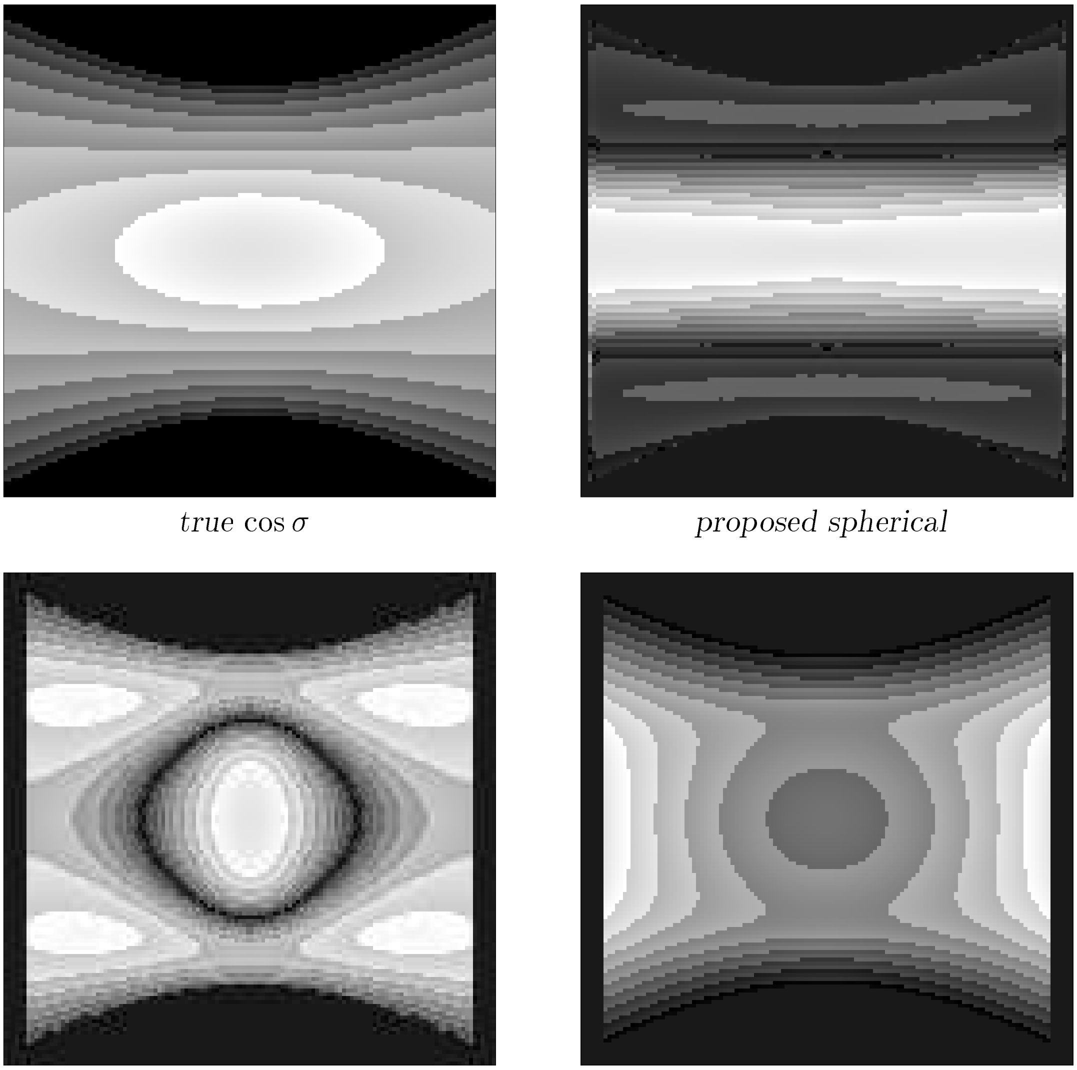

true $\cos\sigma$ — *proposed spherical*

Pentland's spherical — *Pentland's heuristic* ($R = 75$)

Figure 6.7: *Slant estimates for the test hyperboloid.*

ably. The results on the test hyperboloid are all unacceptable because of the odd structures having no relation to the underlying surface. This is an argument for preferring the scaled image function *because this slant estimator—although not necessarily accurate in real cases—does not create any new "spurious" surface structures.*

6.2.3 Tilt Estimators

The tilt estimation is the main bottleneck of *LSA* algorithms, causing the impossi-

					quantiles		
tilt estimator	min.	max.	mean	std. dev.	5%	median	95%
				ellipsoid			
Proposed spherical	−32.6	21.4	−9.4	13.2	−29.6	−10.0	15.5
Pentland's spherical	−75.5	43.4	−22.0	21.2	−61.5	−17.0	6.2
Pentland's heuristic ($R = 66$)	−30.9	27.0	−8.1	11.8	−24.3	−9.1	15.6
				hyperboloid			
Proposed spherical	−40.5	26.0	−12.0	14.7	−35.5	−13.0	16.0
Pentland's spherical	−66.1	3.7	−25.0	17.9	−54.9	−24.0	−2.9
Pentland's heuristic ($R = 75$)	−64.8	11.7	−15.0	14.7	−46.8	−10.0	4.6

Table 6.1: *Local surface slant recovery by different algorithms. The values represent statistics computed on absolute-error maps (in degrees).*

			quantiles		
tilt estimator	mean	standard deviation	5%	median	95%
			ellipsoid		
proposed single-step	0	8.2	−13.4	0	13.4
proposed relaxation	$-1.3 \cdot 10^{-4}$	7.0	−11.3	0	11.3
image gradient	0	14.5	−25.9	0	25.9
modified Pentland	0	31.0	−60.7	0	60.6
			hyperboloid		
proposed single-step	0	8.3	−13.3	−0.03	13.3
proposed relaxation	0	6.4	−10.8	−0.03	10.8
image gradient	0	62.0	−114	−0.09	114
modified Pentland	0	32.9	−64.9	−0.04	64.8

Table 6.2: *Local surface tilt recovery by different algorithms. The values represent statistics computed on absolute-error maps (in degrees). The convergence parameters of the relaxation algorithm were $\xi = 0.25$ and $R = 10^{-3}$ and five iterations were performed in both tests. The results for the ellipsoid correspond to the convex elliptic solution in the first two algorithms and to the convex solution in the others (a solution giving correct results at the occlusion contour). The results for the hyperboloid correspond to the hyperbolic solution in the first two algorithms and to the convex solution in the others.*

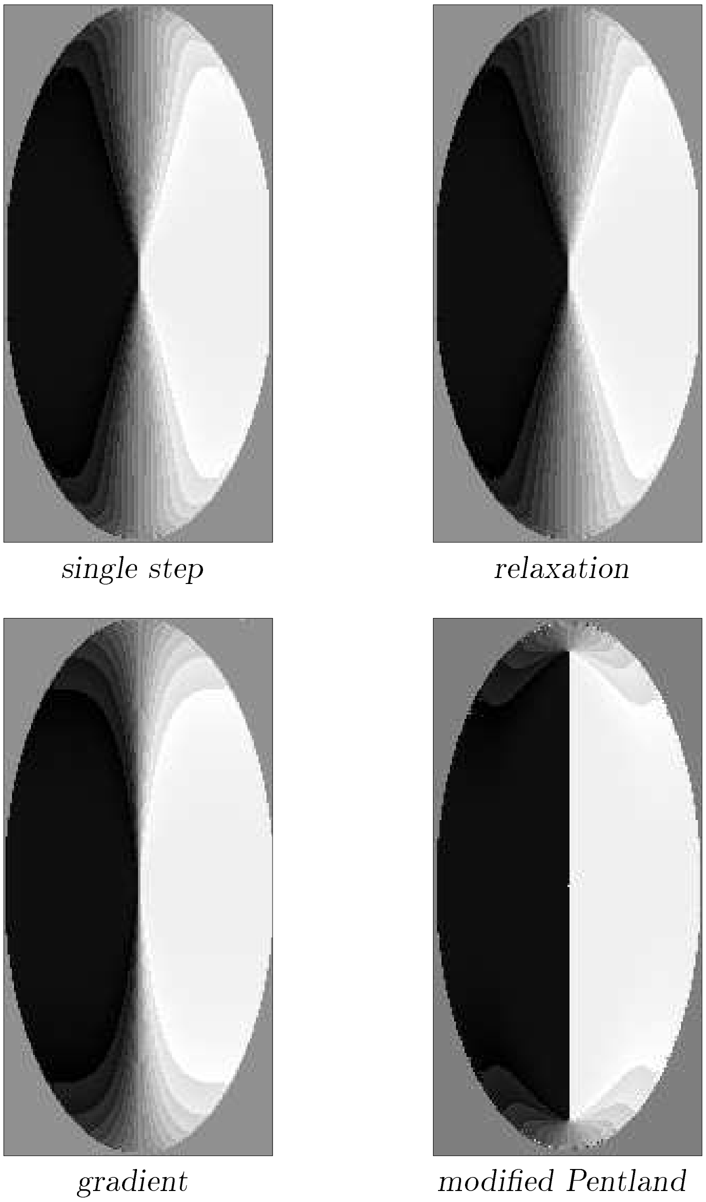

Figure 6.8: *Tilt maps ($\cos\tau$) of the test ellipsoid recovered by the tested algorithms. Compare with the corresponding maps of the true tilt in Fig. 6.4. All results on the occlusion contour are correct. The main difference between the algorithms is in the density of isogones along the longer ellipsoid axis. Note also the influence of the increased density on the recovered Gaussian curvature (cf. Fig. 6.12).*

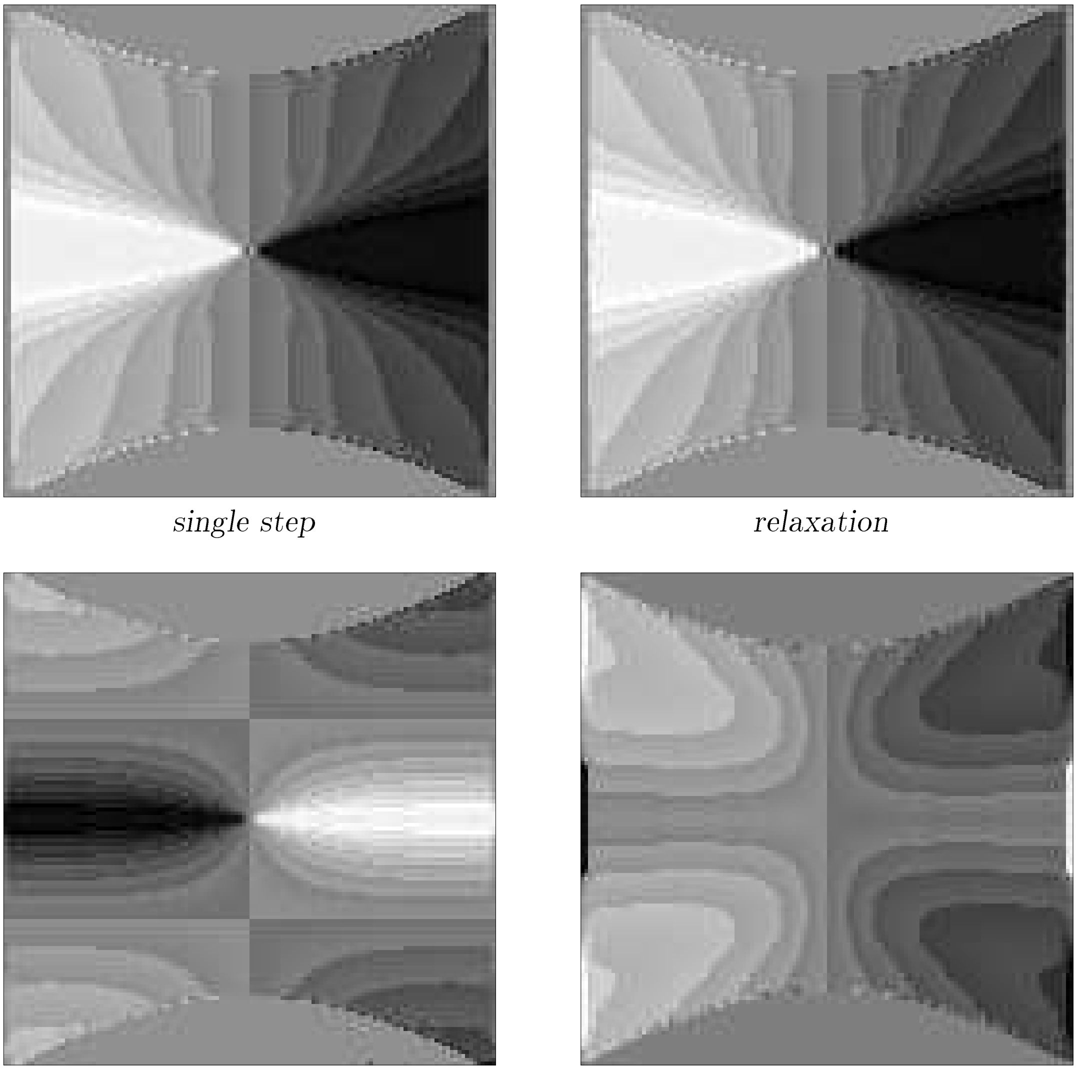

Figure 6.9: *Tilt maps ($\cos\tau$) of the test hyperboloid recovered by the tested algorithms. Compare with the corresponding maps of the true tilt in Fig. 6.4.*

bility of recovering (at least qualitative) surface curvature. The curvature recovery is discussed in Subsection 6.2.4. Here, a brief discussion of the differences between the tested tilt estimators will be given.

The reconstructed tilt maps are shown in Figs. 6.8 and 6.9. Notice the difference between those maps and the true map given in Fig. 6.4 on page 93. For both surfaces, the true tilt isogones[5] are *straight:* they all cross at a singular point in the image center and they end on the occlusion contour. All of the tested algorithms give correct results near the occlusion contour. *The isogones' deviation from being straight in the tilt map expresses therefore the degree of tilt misestimation.*

Only the proposed single-step and relaxation algorithms (6 and 7) can explicitly model the four solutions for the local surface quality. The solution corresponding to the true quality of the surface was chosen for evaluation. In the other algorithms, that solution giving correct results on occlusion boundary was chosen.

Since the reconstructed tilt in $\cos\tau$ maps is not very clearly visible, we also provide the corresponding needle maps in Figs. 6.10 on the next page and 6.11 on page 104.

From Tab. 6.2 on page 99 it follows that all the tested tilt estimators are *unbiased* under the conditions of the experiment because both the median and the mean values vanish. The proposed relaxation estimator shows the best overall behaviour. Standard convergence parameters $\xi = 0.25$ and $R = 10^{-3}$ were used—so that the convergence and stability within the first ten iterations is guaranteed. If the convergence parameters are tuned, the results are much better. Section 6.2.8 deals with the algorithm stability and accuracy in more detail.

It is worth mentioning that the relaxation algorithm does not become instable near the occlusion contour due to the discontinuity in image data. This is because the tilt difference angle ϑ vanishes near the contour and thus the directional derivatives in the gradient and the isophote directions do not influence the stability. The artifacts near the occlusion contour resulting from the operators for computing isophote properties are removed after a few iterations of the algorithm.

Evaluation

The disadvantage of the gradient algorithm is that the recovered isogones are too unevenly distributed and curved. Their high density along the longer axis of the ellipsoid causes the recovered surface curvature to be high there (see also Fig. 6.12 on page 106 and the corresponding discussion in Subsection 6.2.4).

The modified Pentland algorithm tends to recover a "chord" at the top of the ellipsoid that is also well visible in Fig. 6.10 on the next page and in the reconstructed surface view in Fig. 6.21 on page 119. Note that this is the only algorithm giving a consistent solution at the singular image point.

The proposed algorithms remove the disadvantage of the gradient algorithm. The improvement is maximal just along the critical ellipsoid axis. On the other hand, it is relatively small along the longer sides of the ellipsoid. This is due to a slow

[5] In this case, isogones are the isocontours in the tilt map.

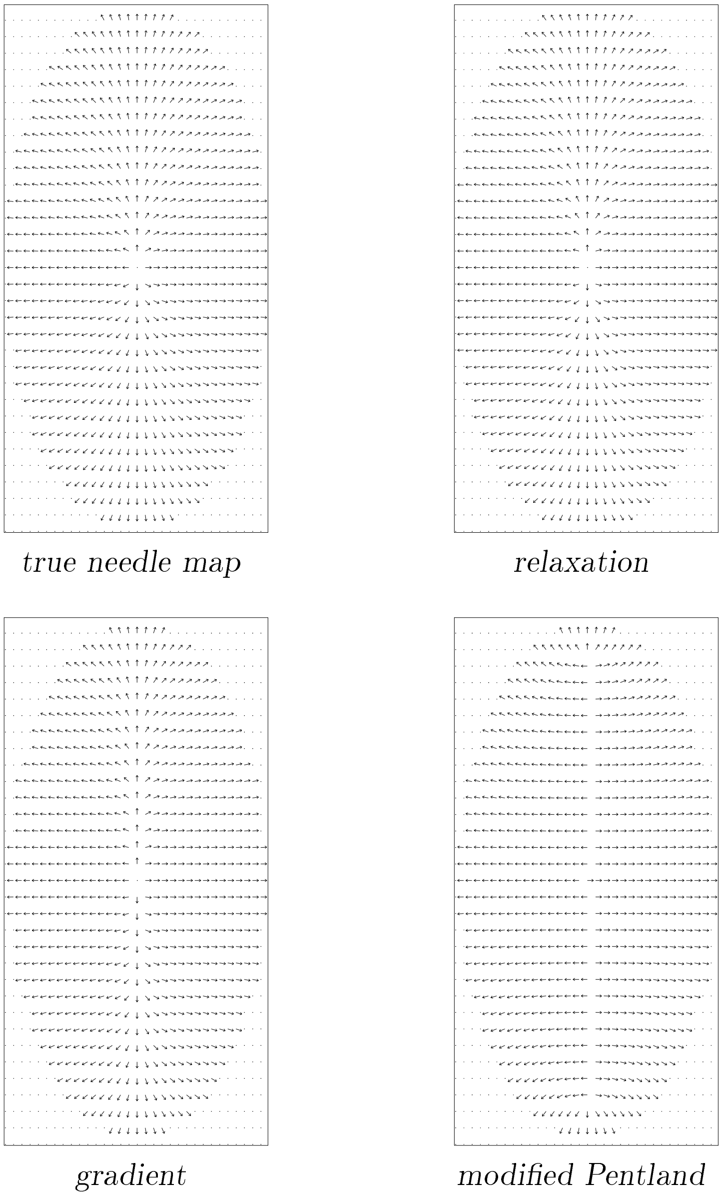

Figure 6.10: *Needle maps representing recovered local surface tilt of the test ellipsoid. The arrows show surface normals projected onto the image plane and normalized to unit-length vectors.*

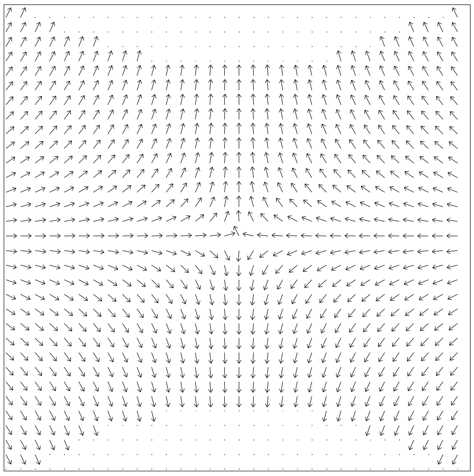

true needle map

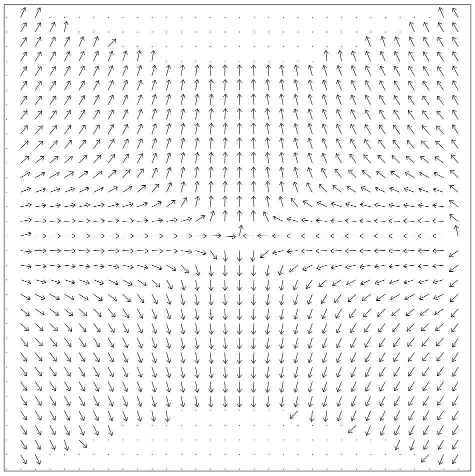

relaxation algorithm

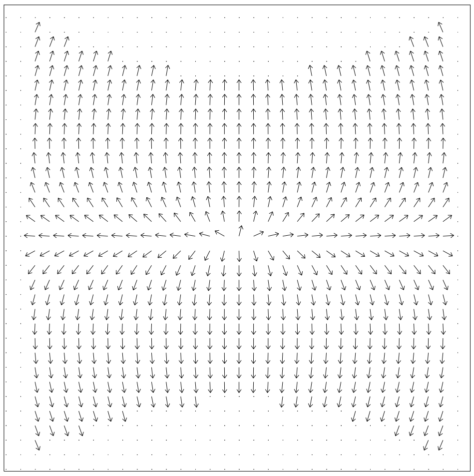

gradient estimator

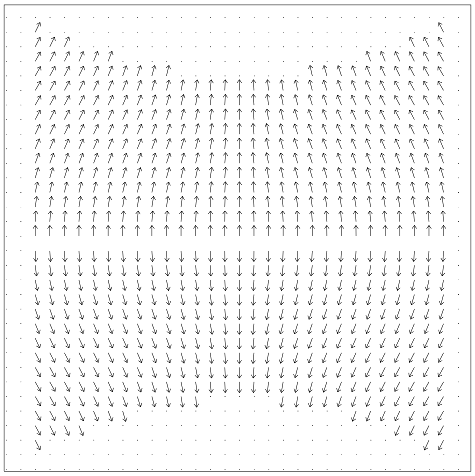

modified Pentland estimator

Figure 6.11: *Needle maps representing recovered local surface tilt of the test hyperboloid.*

convergence of the relaxation algorithm caused by the fact that the image flow-line curvature κ_{F} is relatively low there.

The improvement of the tilt estimates recovered by the proposed estimators over those recovered by the other estimators is quite substantial in the case of the hyperboloid. It is very well visible in the needle map in Fig. 6.11.

tilt estimator	mean	standard deviation	quantiles 5%	median	95%
ellipsoid					
proposed single-step	0.34	1.1	−0.32	0.28	1.2
proposed relaxation	0.28	0.77	−0.29	0.23	1
image gradient	0.59	0.74	−0.41	0.42	2.1
modified Pentland	3.5	540	−40	−1.6	24
hyperboloid					
proposed single-step	0.39	4.3	−0.60	0.065	2.1
proposed relaxation	0.39	4.2	−0.55	0.053	2.2
image gradient	−0.28	100	−13	−1.4	11
modified Pentland	6.9	29	−0.25	1.1	26

Table 6.3: *Gaussian curvature recovery from the needle maps computed by different algorithms. The values represent relative errors ($1 \widehat{=} 100\%$). The convergence parameter of the relaxation algorithm was $\xi = 0.25$ for the ellipsoid and $\xi = 0.1$ for the hyperboloid. $R = 10^{-3}$ and five iterations were performed in both tests. The result corresponding to the true elliptic/hyperbolic choice is presented for the first two cases and the result corresponding to the convex solution is presented for the others.*

6.2.4 Gaussian Curvature Recovery

In this subsection, the Gaussian-curvature maps recovered from needle maps computed by each of the tested algorithms are compared. Gaussian curvature is important since its sign distinguishes two qualitative types of surfaces: hyperbolic and elliptic. It has already been shown that the sign of this curvature *cannot be recovered locally:* there are at most two local solutions. Nevertheless, those two solutions must both come close to zero when the true K approaches zero. This happens in the vicinity of a parabolic line. We have also seen that the parabolic lines are important elements of the surface structure. These are the reasons we decided to compare the algorithms in their ability of recovering *qualitative Gaussian curvature.* We therefore consider it important that an algorithm does not produce "spurious parabolic lines," i.e., does not detect zero-crossings of K at locations out of the true parabolic lines. *Those algorithms that produce a minimum of "spurious parabolic lines"—particularly on qualitatively homogeneous surfaces—are generally preferred.*

The Gaussian-curvature maps were computed by differentiating the recovered needle maps. This simple method gives satisfactory results for the noiseless test data. We used this method also for the proposed Algorithms 6 and 7, although they can produce the

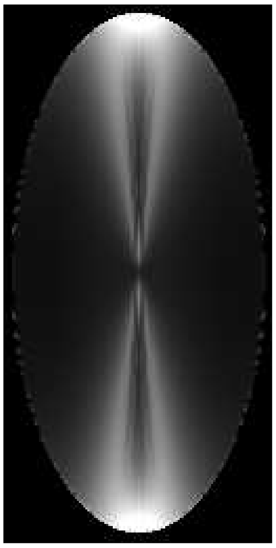

single step $\langle 0,\, 5 \cdot 10^{-4} \rangle$

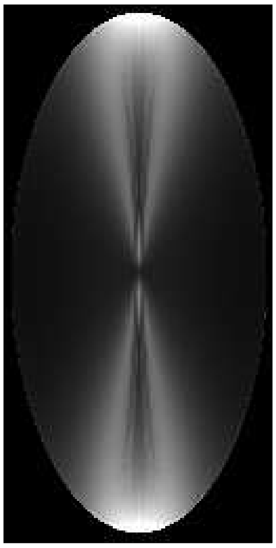

relaxation $\langle 0,\, 5 \cdot 10^{-4} \rangle$

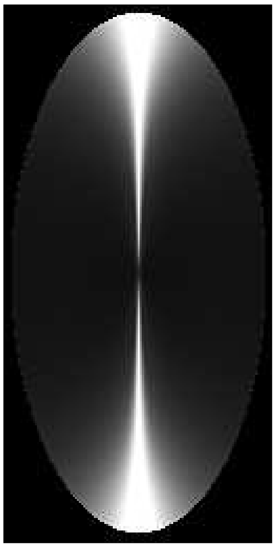

gradient $\langle 0,\, 5 \cdot 10^{-4} \rangle$

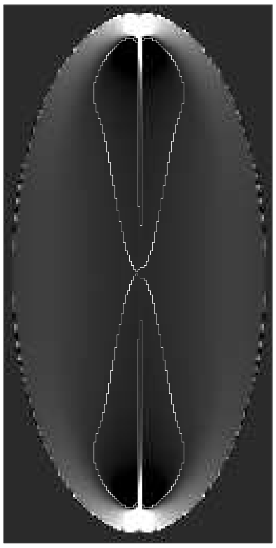

modified Pentland $\langle -10^{-4},\, 5 \cdot 10^{-4} \rangle$

Figure 6.12: *Gray-coded Gaussian-curvature maps recovered from the needle maps computed by the four tested algorithms from the image of the standard ellipsoid. The values in braces are the visible ranges. The brighter a point is in the map the more positive is K. The out-of-range values are displayed as black and white. Compare with the true map in Fig. 6.4. The white curve in the lower-right image is a "spurious parabolic line" recovered by the algorithm.*

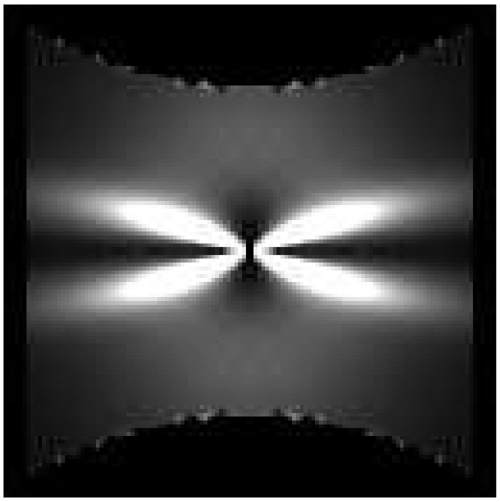

single step $\langle -5 \cdot 10^{-4},\, 0 \rangle$

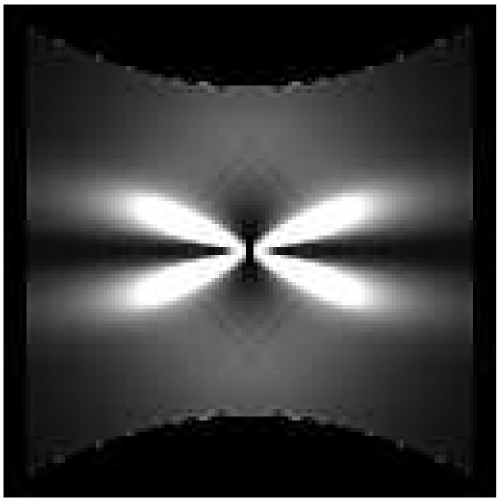

relaxation $\langle -5 \cdot 10^{-4},\, 0 \rangle$

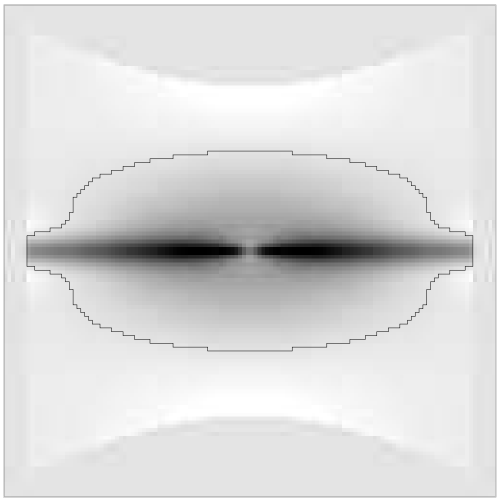

gradient $\langle -10^{-4},\, 5 \cdot 10^{-4} \rangle$

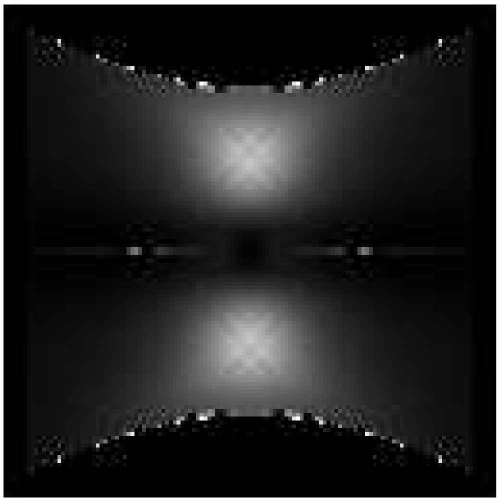

modified Pentland $\langle -5 \cdot 10^{-4},\, 0 \rangle$

Figure 6.13: *Gray-coded Gaussian-curvature maps recovered from the needle maps computed by the four tested algorithms from the image of the standard hyperboloid. The brighter point is in the map the more negative is K. Compare with the true map in Fig. 6.4. The black curve in the lower-left image is a "spurious parabolic line" recovered by the algorithm.*

maps directly. The reason is that the evaluation should compare the results recovered in a similar way. The direct recovery of K is discussed in Section 6.2.12.

If the surface is represented by a Monge patch

$$\mathbf{x}(x, y) = [x, \;\; y, \;\; f(x, y)]^{\mathrm{T}},$$

the normal vector can be written in the form

$$\mathbf{N}(x, y) = \frac{[-p, \;\; -q, \;\; 1]^{\mathrm{T}}}{\sqrt{1 + p^2 + q^2}},$$

and the needle map is represented by the tuple (p, q), where $p(x, y) = f_x(x, y)$ and $q(x, y) = f_y(x, y)$. The Gaussian curvature is then

$$K(x, y) = \frac{p_x\, q_y - p_y\, q_x}{(1 + p^2 + q^2)^2}, \qquad p_y = q_x. \tag{6.3}$$

For the sake of a better understanding the results, we prefer evaluating the standard statistics on *relative-error maps* over evaluating them on absolute-error maps as was done so far. This is mainly because of a wide range of the curvatures of the test surfaces and a rather wide range of the absolute errors.

Evaluation

From Fig. 6.12 on page 106 and 6.13 on the previous page it can be seen that the recovered curvature map shows various inhomogeneities—the amount of those is also expressed in the standard deviation entries in Tab. 6.3 on page 105. For both surfaces, the first three algorithms in the table give quite good estimates near the occlusion contour. But the algorithms' results significantly differ in the regions near the singular points.

LSA based on the image gradient tilt estimator running on the image of the hyperboloid as well as the modified Pentland *LSA* running on the image of the ellipsoid create "spurious parabolic lines" in the resulting curvature maps. *These methods therefore do not recover the qualitative Gaussian curvature correctly.*

It can be seen by comparing the results in Figs. 6.12, 6.13, and Tab. 6.3 that the lowest 5/95%-quantile error as well as the lowest bias (entries "median" in the table) are exhibited by the single-step and the relaxation algorithms (Alg. 6 and Alg. 7, respectively). The main difference to the true curvature map occur in the regions along the longer sides of the ellipsoid and along the principal axes of the hyperboloid.

If we compare the results of the single-step and the relaxation algorithms with the others, their ability to smooth out the high-curvature structures recovered by the image gradient algorithm in the vicinity of the longer principal axis of the ellipsoid is evident. Nevertheless, there still remain radial structures of higher curvature (visible as lighter strips leaving the map's center towards the occlusion boundary in the upper two maps in Fig. 6.12) that do not appear in the original map (cf. Fig. 6.4 on page 93). These structures appear at locations where the image flow-line curvature is high relative to image isophote curvature, i.e., where the ratio $\kappa_{\mathrm{F}}/\kappa_{\mathrm{I}}$ becomes high. We hypothesize that these structures are due to numerical ill-posedness of the local isophote characteristics computation in discrete images. A computationally more robust method could probably result in a better needle map.

6.2.5 The Influence of Surface Curvature on the Tilt Estimate

This is a standard test often presented in the literature on *LSA*. It measures the influence of the test body's eccentricity on the accuracy of the tilt estimators [Pen84, FL89, FL92]. The test body is usually a quadric—an ellipsoid. The test demonstrates the behaviour of a tilt estimator when the eccentricity of the body increases. One could argue that, since there is a theoretical maximum of deviation from the true tilt when the eccentricity increases[6], the gradient estimator is accurate for surfaces that are "nearly spherical" and "nearly cylindrical" [FL92, FL89].[7] This is *not* fully true, however, since the error for the body that is slanted from its basic position (when the principal axes of the test quadric are not aligned with the axes of the viewer-centered coordinate system) does not drop to zero when the eccentricity is so high that the body becomes "nearly cylindrical" over most of its surface.

The plot of the tilt error computed as a function of the test ellipsoid eccentricity is shown in Fig. 6.14 on the next page. The ellipsoid axes were parallel to the axes of the $\mathcal{V}$-system.

6.2.6 Sensitivity of Tilt Estimators to Explicit Slant Knowledge

This experiment tests the influence of non-Lambertianity as well as the wrong slant estimation influence on the resulting normal tilt estimate. The purpose of this test is to demonstrate that although the estimator (5.5) requires knowledge of the slant, it is relatively insensitive to its misestimation. This misestimation can either occur due to

[6]It is ca. 60 degrees for the eccentricity $x : y \approx 5$ according to [FL92, FL89].

[7]The authors do not give any details about how they derived their formula for the theoretical maximum error of the gradient estimator. They even do not present it. Our own results show that there is no local maximum of the error function in the real domain. The *theoretical maximum value* of the error function evaluated for an increasing eccentricity ratio does not show any local maximum either—it approaches ca. 60 degrees asymptotically for the infinite eccentricity of the ellipsoid.

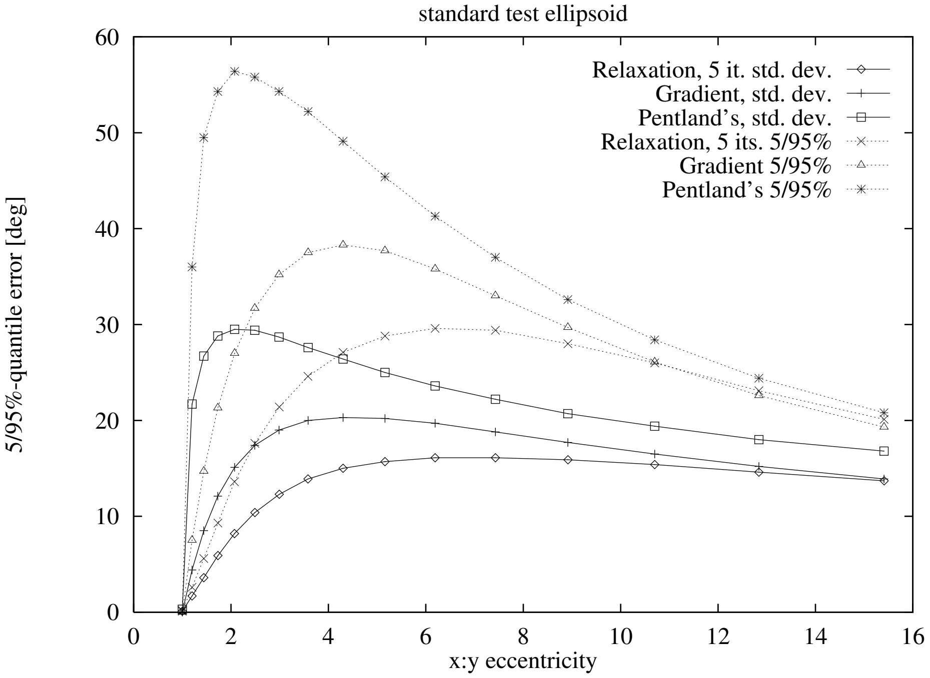

Figure 6.14: *Tilt estimation error as a function of the test ellipsoid $x : y$ eccentricity for the tested tilt estimators. The solid lines show the standard deviation error, the dashed lines show the 5/95%-quantile absolute error. The relaxation algorithm were run with standard parameters. The input image size varied from 128×128 to 128×1972 pixels. When the eccentricity increases, all the estimators show similar statistics. This is because of the smoothing effect of the operator at the locations with high surface curvature.*

a bad slant estimator or due to a deviation of the true reflectance function from the assumed one. We therefore designed an experiment that models both influences on the slant estimate.

It is supposed that the slant σ is estimated as the scaled Lambertian image value $I(x, y)/I_{max}$ (see Section 5.3). The scaled image of a test surface illuminated from the vantage point is a function $I(x, y) = \cos \sigma(x, y)$ as if the albedo and the light source intensity were unit.

Under these conditions, the slant misestimation as well as the degree of non-Lambertianity can be modeled by transforming the true image value $I \mapsto T(I)$. The transform T has to satisfy these requirements:

1. T is monotonic (so that no new extrema in transformed data occur),
2. $T(0) = 0$ and $T(1) = 1$ (so that the range of $T(\cos \sigma)$ values is preserved),

3. T should model the non-Lambertianity.

We propose $T(x) = x^n, \quad n > 0$. It is a simple function satisfying the first two requirements. According to Phong's reflectance model [Pho75], the third requirement is satisfied as well because the specularity strength can be modeled by the power of the term $\mathbf{N} \cdot \mathbf{H}$, where the specular direction $\mathbf{H}$ is defined as the bisectrix of the light direction $\mathbf{L}$ and the viewing direction $\mathbf{V}$. In our experimental setup, $\mathbf{L} = \mathbf{V} = \mathbf{H}$, therefore the exponent n in $T(I(x, y)) = \cos^n \sigma(x, y)$ is also related to the specularity strength. We call it the *non-Lambertianity index* and it has a value of $n = 1$ when the surface reflectance is perfectly Lambertian. The exponent n can also express the degree of slant misestimation due to a bad slant estimator.

Evaluation

The expected estimation error of the *image gradient tilt estimator* should be constant for any non-Lambertianity index $n > 0$ since this estimator is R-invariant if $\mathbf{L} = \mathbf{V}$ or when the reflectance function is space-invariant. Both these conditions are fulfilled. The experiment proved this expectation. The results are summarized in the following table:

valid for any non-Lambertianity index ranging from 0.1 to 10

					quantiles		
surface	min	max	mean	standard deviation	5%	median	95%
ellipsoid	−36.8	36.8	0	14.5	−25.9	0	25.9
hyperboloid	−175	175	0	62.0	−114	−0.085	113

The *modified Pentland estimator* is very sensitive to non-Lambertianity of the reflectance function—although the slant estimate is not required in an explicit form, transforming the image by a monotonic function T will considerably influence the estimate. The influence of non-Lambertianity of the illuminated surface on the tilt estimates is summarized in Fig. 6.15 on page 113.

The result of the *single-step tilt estimator* (Alg. 6) was taken as the first estimate for the *relaxation tilt estimator* (Alg. 7). The influence of slant misestimation is small, as can be seen from the test results that are summarized in Fig. 6.16 on page 114. It seems that the estimator is more sensitive to non-Lambertianity when the surface becomes hyperbolic.

In the range of the non-Lambertianity index $n = \langle 1/10, 10 \rangle$, the results of the proposed estimators are better than the gradient tilt estimator and the modified Pentland tilt estimator according to the following tables:

SINGLE-STEP ESTIMATE RELATIVE TO OTHER ESTIMATES

improvement over estimate	in standard deviation error		in 5/95%-quantile error	
	ellipsoid	hyperboloid	ellipsoid	hyperboloid
gradient	1.5–1.8×	2.4–9.4×	1.5–2.0×	2.6–11.0×
modified Pentland	3.4–5.9×	1.8–5.2×	4.3–6.0×	1.8–6.5×

RELAXATION ESTIMATE RELATIVE TO SINGLE-STEP ESTIMATE

improvement over estimate	in standard deviation error		in 5/95%-quantile error	
	ellipsoid	hyperboloid	ellipsoid	hyperboloid
single-step	0.96–1.20×	0.85–1.55×	1.0–1.23×	0.91–1.83×

6.2.7 Sensitivity to Illuminant Position

An important question is how the results of *LSA* are influenced by the displacement of the illuminant from its presupposed position. This experiment evaluates the influence of the violation of the condition $\mathbf{L} = \mathbf{V}$ required for the Algorithms 5, 6, and 7.

A spherical surface was illuminated by an infinitely distant illuminant with gradually increasing slant $\sigma_{\mathcal{L}}$ and constant tilt. The statistics were computed on the absolute-error map of the illuminated area. The influence on the tilt estimator error is summarized in Fig. 6.17 on page 115 and the influence on the overall local orientation error in Fig. 6.18 on page 116.

Evaluation

Pentland algorithm—which is designed to behave correctly on spheres—proved to be invariant under the illuminant position change. The advantage of this estimator is that it remains unbiased under the change of illuminant vector **L** *if the surface is spherical.* See the formula for the estimator (6.1) on page 94. This advantage is lost for non-spherical surfaces.

The overall local orientation error does not exceed the value of the illuminant slant $\sigma_{\mathcal{L}}$. This result is very similar for all *LSA* algorithms based on the image gradient.

The tilt error statistics are quite similar for all the gradient-based algorithms. The error increases linearly with the increasing illuminant slant.

All the tested algorithms give unbiased tilt estimates under the conditions of this experiment.

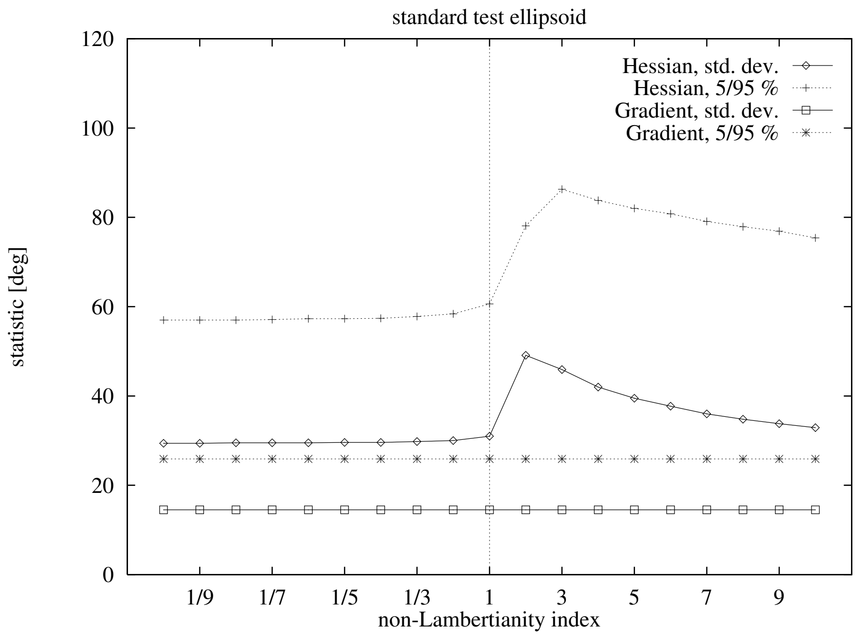

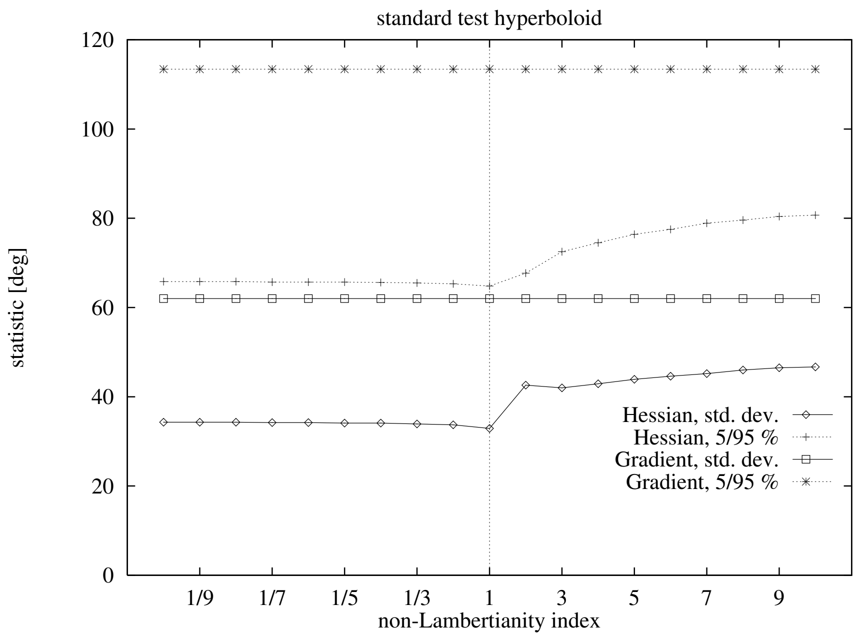

Figure 6.15: *The influence of non-Lambertianity on the* modified Pentland tilt *estimates and a comparison with the gradient tilt estimates. The solid lines show the standard deviation error, the dashed lines show the 5/95%-quantile error. The absolute errors are in degrees. The fast increase is caused by the "switching" between the locally best estimates when the Hessian eigenvectors turn due to the image transform.*

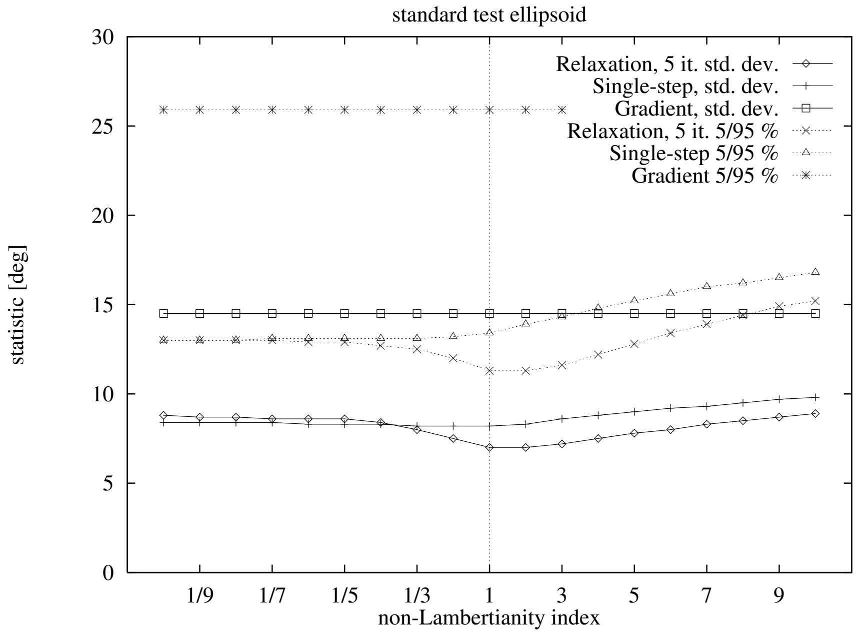

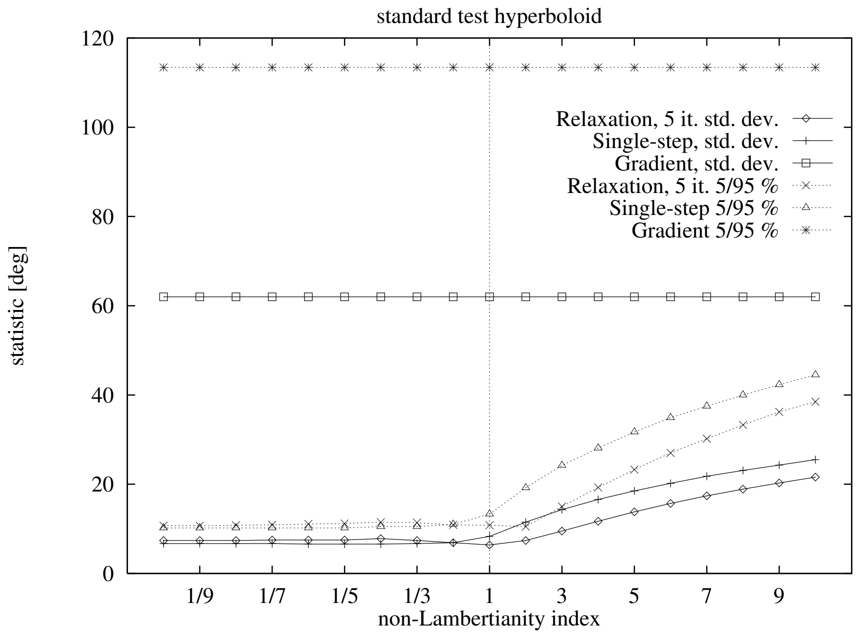

Figure 6.16: *The influence of non-Lambertianity on the* proposed tilt estimator *results and a comparison with the gradient tilt estimates. The solid lines show the standard deviation error, the dashed lines show the 5/95%-quantile error.* $\xi = 0.25$ *and* $R = 10^{-3}$ *in both test runs. The errors are in degrees. Note that the scale is different in both plots.*

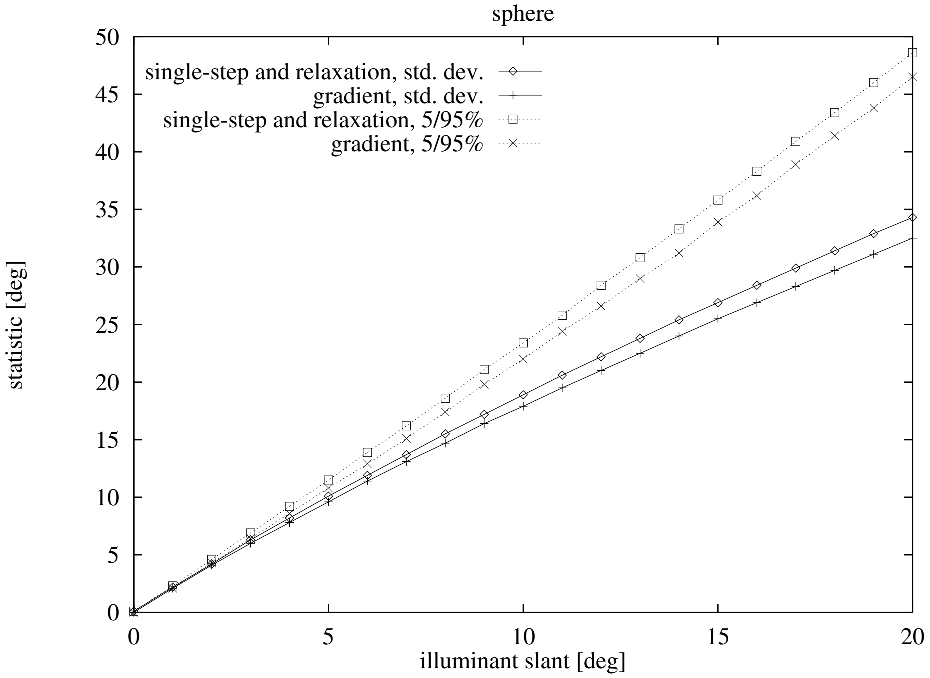

Figure 6.17: *The influence of the illuminant slant from the presupposed position on the* local tilt estimates. *The result is similar for all the image-gradient-based LSA algorithms. Pentland estimator is invariant under the illuminant slant since the surface is spherical.*

6.2.8 The Influence of ξ and R on the Iterative Algorithm Convergence

The aim of this experiment was to compare the relaxation Algorithm 7 results under the change of the convergence constant ξ and the covariance R. We performed the standard ellipsoid and hyperboloid tests. The results are summarized in Figs. 6.19 on page 117 and 6.20 on page 118.

The cause of the global statistical instability is the local instability that occurs in regions with higher slant (but not on the occlusion contour) when the convergence constant becomes improper. There can also occur a small instability around the singular points of the images of hyperbolic surfaces that causes the slow increase of global statistics observable in Fig. 6.20.

6.2.9 3-D Surface Reconstruction from Needle Maps

For completeness and for the sake of clear discernibility of the *LSA* algorithms properties, we included an experiment with the reconstruction of the depth map from the needle

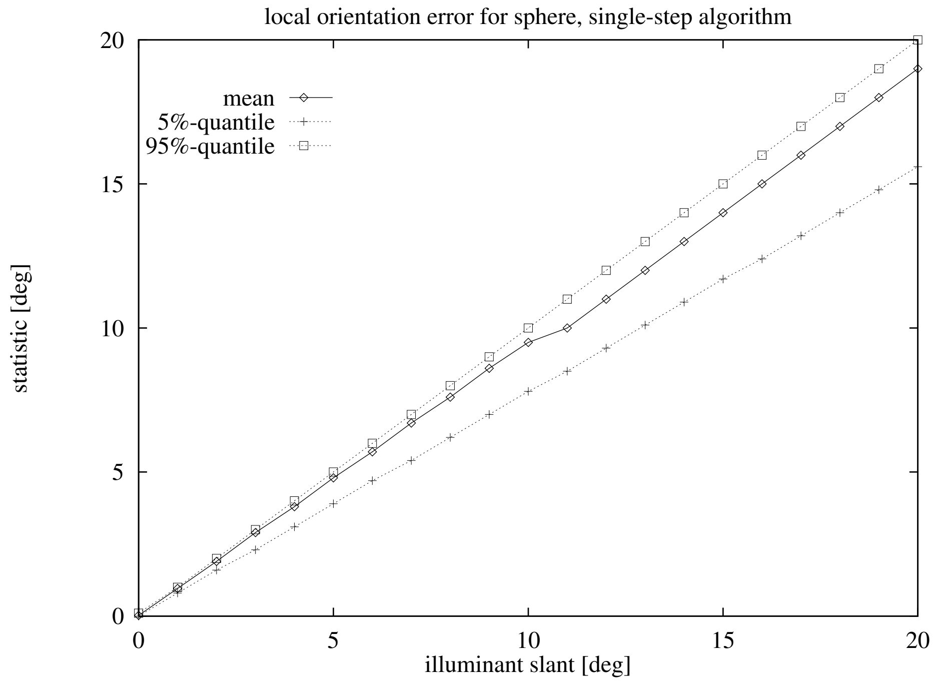

Figure 6.18: *The influence of the illuminant slant from the presupposed position on the* local orientation estimates. *The error is the angular difference between the true and the recovered local normal. The result is similar for all image-gradient-based LSA algorithms.*

map recovered by *LSA* algorithms. The results are shown in Figs. 6.21 on page 119 and 6.22 on page 120.

The reconstruction algorithm was adapted from [Ter88]. We will briefly describe it here. Details can be found in Terzopoulos's article. The variational method is based on deforming a thin plate $v(x, y)$ under tension so long as the penalty for the discrepancy between the local orientations of the deformed plate and the constraints represented by the needle map data becomes minimum. The distance is expressed as an energy functional

$$\mathcal{E}(v) = \mathcal{S}(v) + \mathcal{P}(v),$$

where $\mathcal{P}$ is a *penalty functional* measuring the discrepancy between v and given constraints and $\mathcal{S}$ is a *stabilizing functional* measuring the lack of smoothness of v. The function v is supposed to be a member of the space of admissible functions $\mathcal{H}$. The regularized reconstruction problem is then formulated according to the following variational principle:

Find $u \in \mathcal{H}$ such that

$$\mathcal{E}(u) = \inf_{v \in \mathcal{H}} \mathcal{E}(v).$$

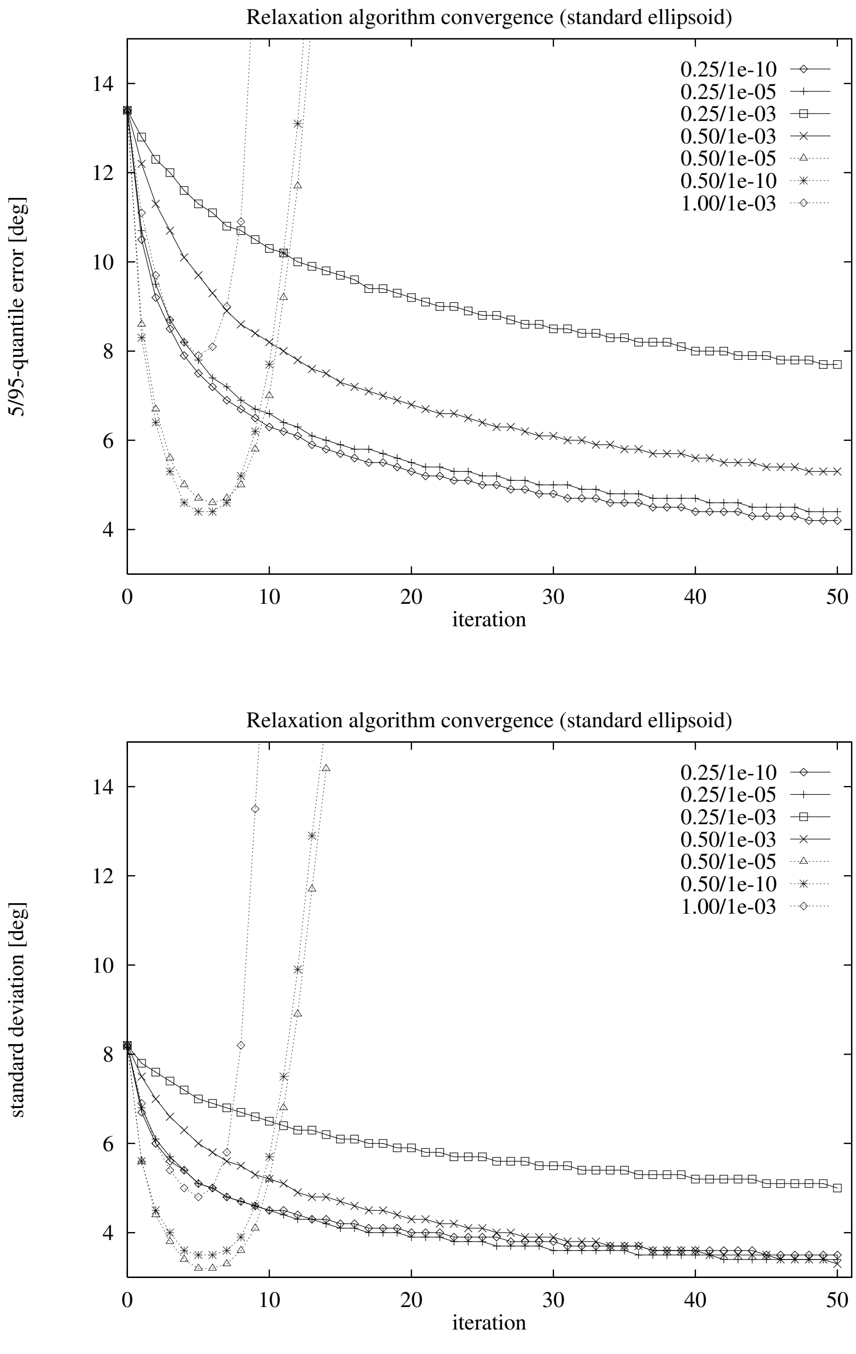

Figure 6.19: *Convergence of the relaxation algorithm run on the ellipsoid image. The numbers in the key denote the value of ξ and R for each algorithm run. The algorithm diverges in the observed global statistics for $\xi > 0.5$. Upper plot: 5/95%-quantile error. Lower plot: standard deviation error.*

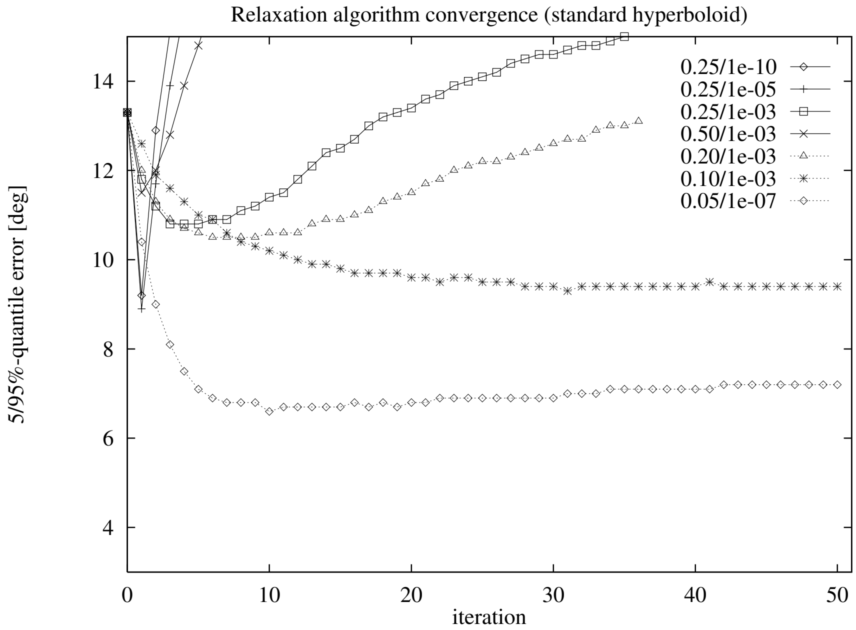

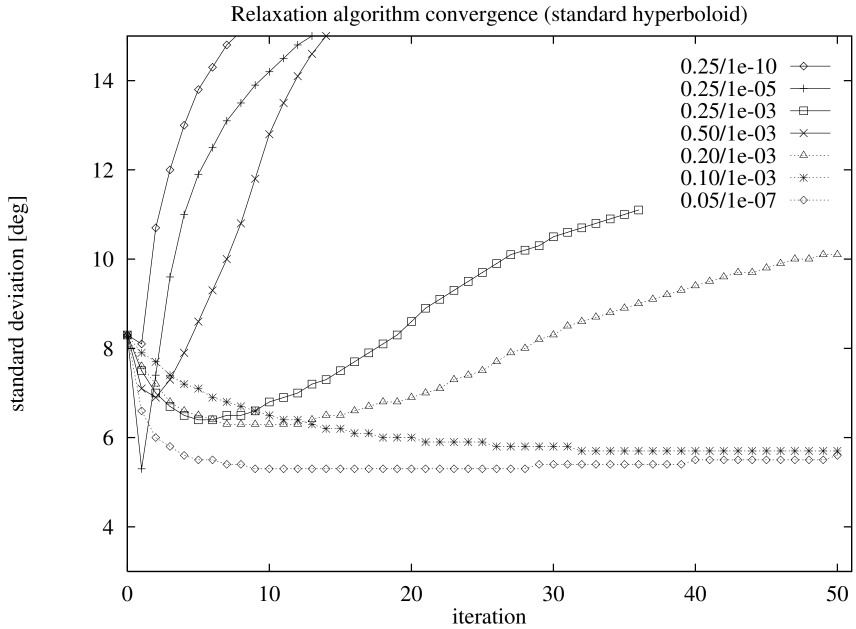

Figure 6.20: *Convergence of the relaxation algorithm run on the hyperboloid image. The numbers in the key denote the value of ξ and R for each algorithm run. The algorithm diverges in the observed global statistics for $\xi > 0.1$. Upper plot: 5/95%-quantile error. Lower plot: standard deviation error. The solid lines correspond to the solid lines in Fig. 6.19.*

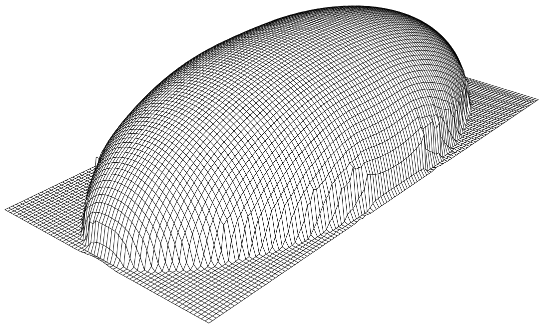

reconstructed from the true map

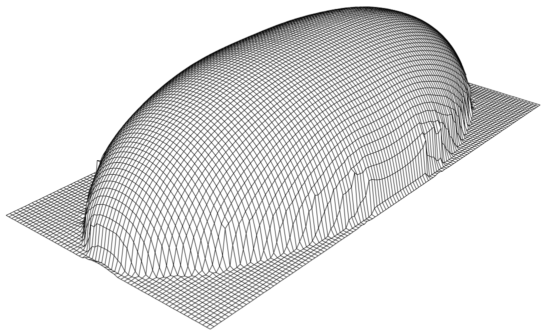

relaxation algorithm

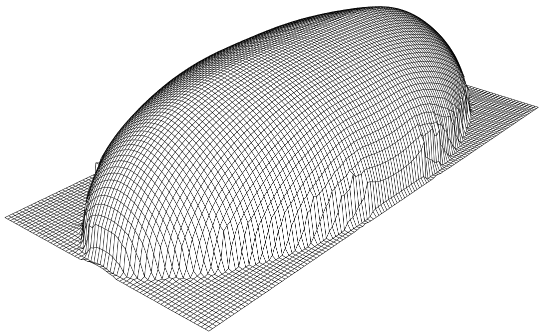

image gradient algorithm

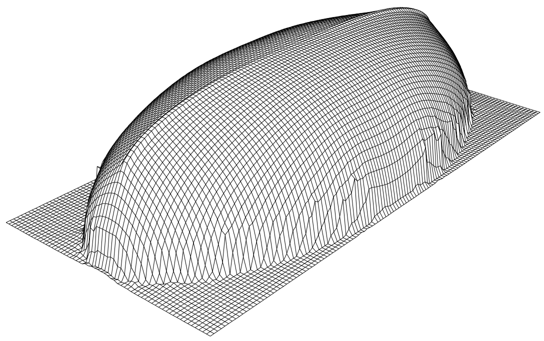

modified Pentland algorithm

Figure 6.21: *Standard views of the test ellipsoid reconstructed from the needle maps that were recovered by various algorithms. The upper left plot is the surface reconstructed from the true needle map—it is not the true surface (for that see Fig. 6.1). The reconstruction algorithm ran 2000 Jacobi relaxations on the input data, reduced 2× in size. The initial reconstruction error evaluated as* $\sum \sin^2 \epsilon_i$ *reduced 12×* (ϵ_i *is the local angular difference between the local orientation constraint and the normal vector of the reconstructed surface). The initial estimate was a plane. The good looking reconstruction results in all cases are due to the slant estimator giving exact results under the experiment conditions.*

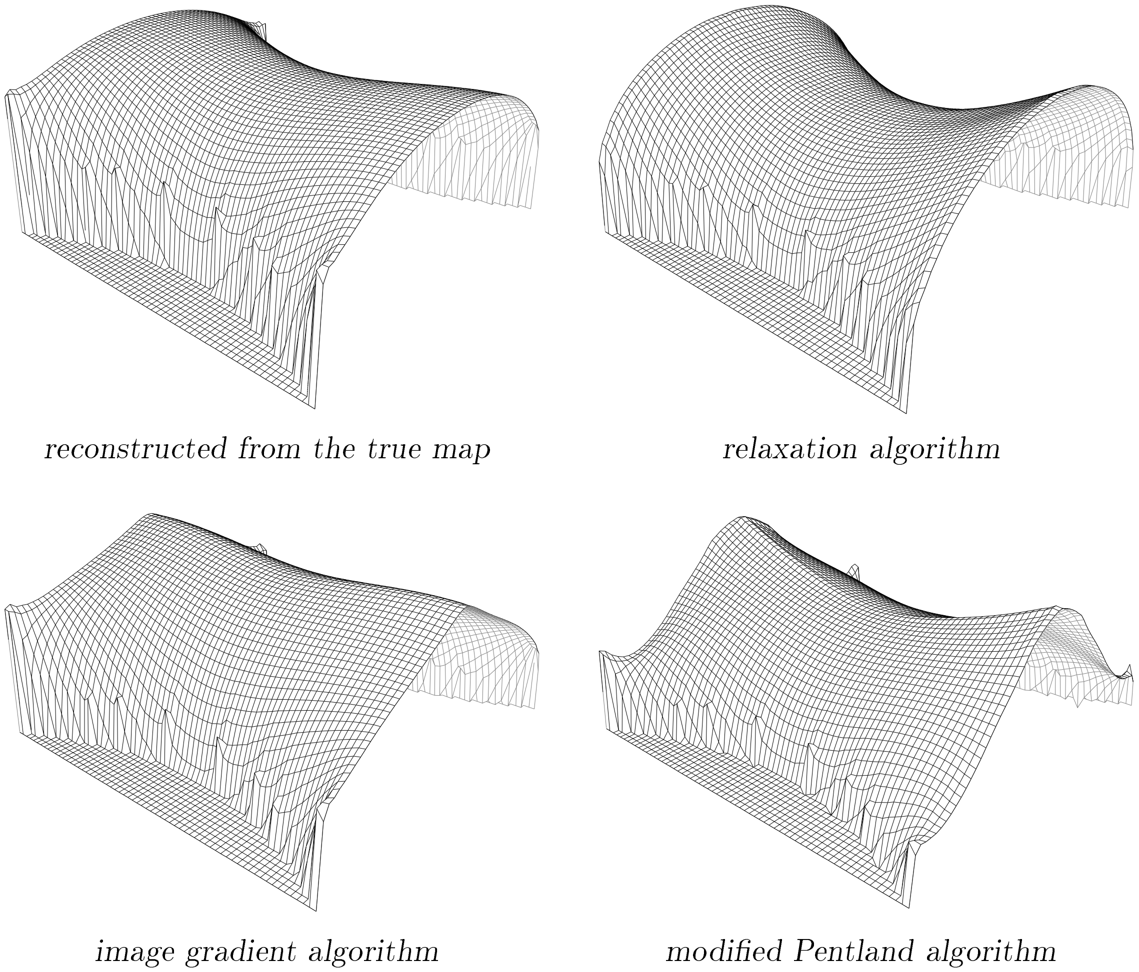

Figure 6.22: *Standard views of the test hyperboloid reconstructed from the needle maps that were recovered by various algorithms. The upper-left plot is the surface reconstructed from the true needle map. The initial reconstruction algorithm ran 2000 Jacobi relaxations on input data reduced $2\times$. The reconstruction error reduced from $2.8\times$ to $4.5\times$ (for the result from the relaxation algorithm).*

Terzopoulos proposed the following *controlled-continuity* stabilizer

$$\mathcal{S}_{\rho\tau}(v) = \frac{1}{2} \iint_{\Omega} \rho \, \left[\tau \, (v_{xx}^2 + 2\, v_{xy}^2 + v_{yy}^2) + (1-\tau)\,(v_x^2 + v_y^2) \right] \, d\Omega \,,$$

where ρ and τ are the continuity control functions modelling the depth and orientation constraints over the image domain Ω.

The penalty functional can be written in the form

$$\begin{aligned}\mathcal{P}(v) = &\frac{1}{2}\sum_{i\in D} a_{d_i}\left[v(x_i,y_i) - d_{(x_i,y_i)}\right]^2\\ &+ \frac{1}{2}\sum_{j\in P} a_{p_j}\left[v_x(x_j,y_j) - p_{(x_j,y_j)}\right]^2\\ &+ \frac{1}{2}\sum_{i\in Q} a_{q_k}\left[v_y(x_k,y_k) - q_{(x_k,y_k)}\right]^2,\end{aligned}$$

where $d_{(x_i,y_i)}$ is the set of local depth constraints, $p_{(x_j,y_j)}$ and $q_{(x_k,y_k)}$ are the sets of local orientation constraints for $i \in D$, $j \in P$, and $k \in Q$, respectively, and a are the constraints weights.

If the solution to the variational task exists, it satisfies the Euler-Lagrange equations, which express the necessary conditions for the minimum. Details about the discrete formulation of the problem can be found in [Ter88].

We have simplified the task so that

- the continuity control functions are assumed constant, since it has been observed that the reconstruction results are good without the necessity to model the discontinuity at the occlusion contour,
- no depth constraints were introduced,
- the domains $P = Q = \Omega$.

Evaluation

The results for the standard ellipsoid look very similar. Notice the "chord" in the surface reconstructed from the needle map obtained from the gradient algorithm and, particularly, from the Pentland algorithm.

The standard hyperboloid test needs some explanation. The result in Fig. 6.22 on the previous page, which is the most similar to the true surface in Fig. 6.1 on page 88, is the surface reconstructed from the needle map recovered by the relaxation algorithm. Paradoxically, this result seems to be better than the reconstruction from the true needle map. It is caused by the reconstruction algorithm that does not have proper boundary constraints. Therefore the reconstructed surface tends to be locally perpendicular to the plane defined by the image boundary and the image plane's normal vector. This is the reason why the surface reconstructed from the true needle map does not seem to be sufficiently "hyperbolic." Another reason is that the needle map produced by the relaxation algorithm is better integrable because the relaxation improves the integrability of the needle map compared to the 'true' needle map obtained by a simple sampling of the analytic vector function representing the true surface orientation.

6.2.10 Recovering Needle Maps of Developable Surfaces

Developables are surfaces that cannot be parameterized in the Gaussian manner. Nevertheless, all the tested algorithms do not fail in the sense that they could not produce any output. Even the proposed Algorithms 6 and 7 give some estimates, although they are supposed to work only on images of non-developable surfaces.

This test was performed on two qualitatively different representatives of developables: a right circular cylinder and a right circular cone (see Fig. 6.23 on the next page). Both surfaces were gradually slanted from the initial position (a position with the axis parallel to image plane). The standard error statistics computed on absolute error maps were evaluated for the increasing slant of the bodies. The results are shown in Figs. 6.24 on page 124 and 6.25 on page 125.

All tilt estimators based on the directionality of the image gradient or the image Hessian eigenvectors give identical (incorrect) results on images of slanted cylinders *regardless of the illuminant position.* The isophotes will always follow the cylinder's rulings (that are parallel), so they project onto parallel straight image isophotes. If the cylinder slants from its initial position, only the density of isophotes and the image value change. This is not sufficient information for the *LSA* algorithm to recover local surface tilt and curvature. The only thing that can be recovered is the *local surface quality:* the isophotes are straight and parallel, while the isophote density is non-zero, which happens only in an image of a locally cylindrical surface.

A cone is a slightly more complicated surface. Unfortunately, there is *no* view where any of the tilt estimators give exact results.[8] There is a view, however, when the cone is slanted so that only the ruling that is closest to the observer is just parallel to the image plane (see Fig. 6.23, where the cone is in this position). When the surface slant is close to this value, the tilt estimates show a minimum deviation from the true values (see Fig. 6.25). The proposed Algorithm 6 showed the best behaviour in this case.

The quality of the surface can also be recognized: in this case, the isophotes remain straight but they are not parallel any more.

Evaluation

From the experiment it can be concluded that the proposed tilt estimators are not worse than any other estimator. For a cylinder, the proposed tilt estimates are equivalent to the image gradient tilt estimates—the tilt difference angle ϑ vanishes because of vanishing flow-line curvature (the isophotes are parallel). In the case of a cone, the

[8]This is not fully true. There is an instable viewer and illuminant position that allows *LSA* to work properly. The case occurs when both the viewing and the light vectors become parallel to the cone's axis. An infinitely small movement of the infinitely distant point light source from this position causes the isophotes "switch" from the position perpendicular to the surface rulings to the position parallel to the rulings. We therefore exclude this case from consideration because it is not likely to occur.

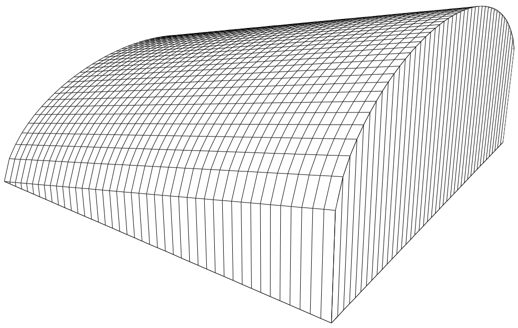

The test cylinder slanted by 26 degrees.

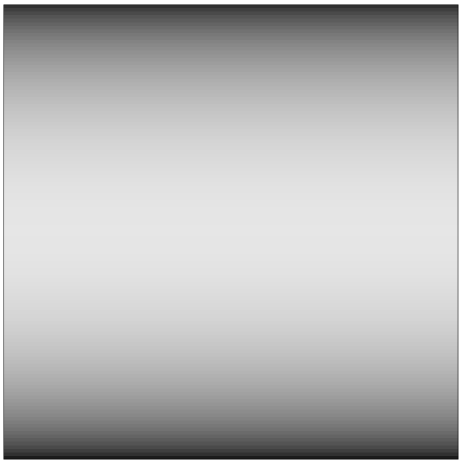

The corresponding image.

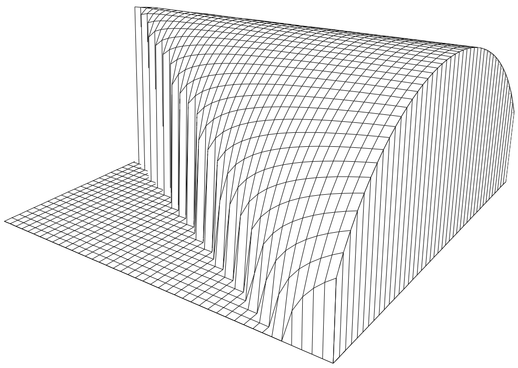

An equally slanted test cone.

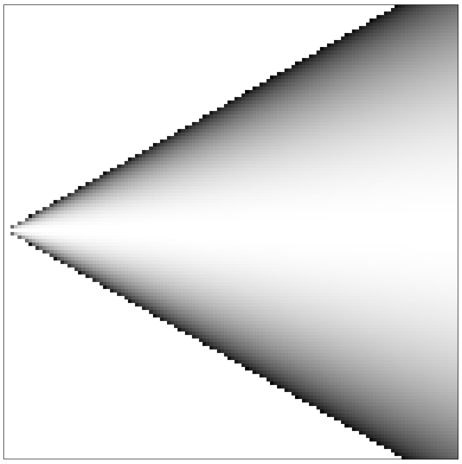

The corresponding image.

Figure 6.23: *Two test surfaces representing developables and their images in a slanted position. The full range of grey levels from black to white corresponds to the image value range of* $\langle 0,\, 1\rangle$. *The cone is shown in a position when the ruling nearest to the viewer (located above the surface) becomes parallel to the image plane.*

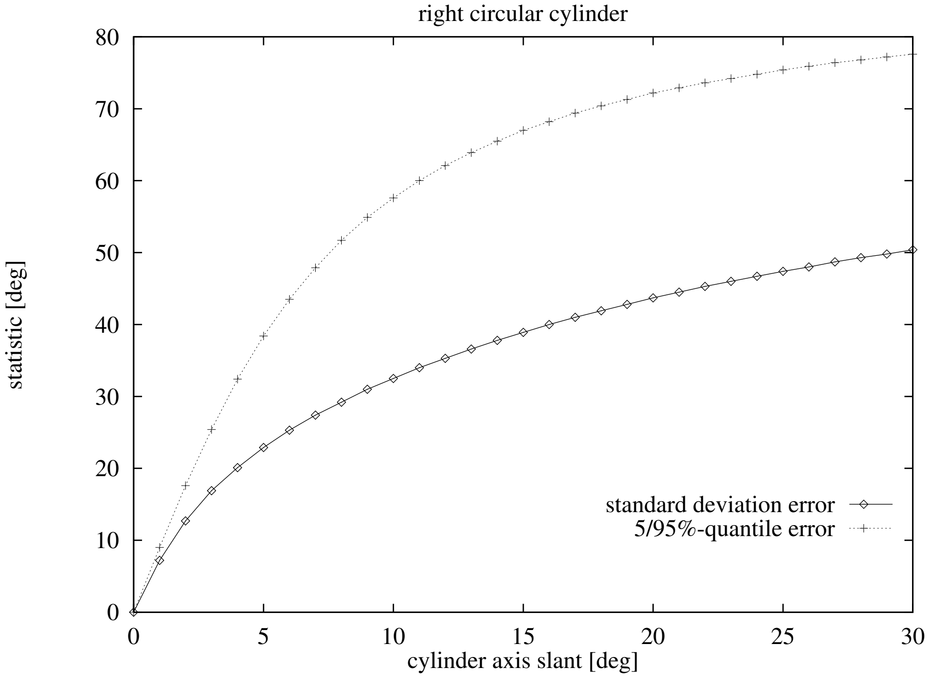

Figure 6.24: *Absolute error of the tilt map recovered from the image of a slanted cylinder. These statistics give only some information about the error spread over the visible surface—the maximum error is 90 degrees except for the case of a vanishing slant. But the region of maximum error is limited to a very narrow strip around the cylinder's axis.*

flow-line curvature is non-zero, which causes a difference from the gradient tilt estimate. But this difference has no relation to the local surface tilt—as can be seen from Fig. 6.25. All estimators remain unbiased for any slant of the tested bodies.

6.2.11 Summary of the Proposed Algorithms Properties

Provided that the surface is properly illuminated and the illuminant is located at the vantage point, it follows from the tests presented in this chapter that:

1. The proposed algorithms work on generic images of general non-developable surfaces.

2. The proposed algorithms do not create any spurious structures in the resulting maps.

3. The proposed algorithms can model all four solutions for the principal curvatures, i.e., two solutions for the Gaussian curvature and four solutions for the mean

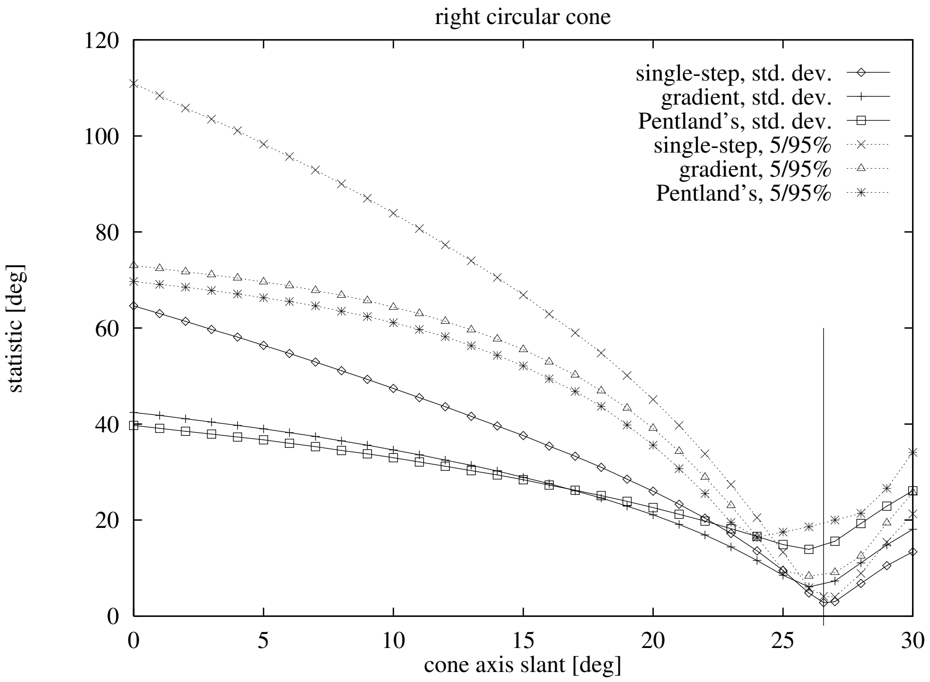

Figure 6.25: *Absolute error of the tilt map recovered from the image of a slanted cone. The vertical bar marks the slant of 26.565 degrees when the cone touches the image plane along a single ruling. This position is the only one where the maximum tilt error is lower than 90 degrees.*

curvature.

4. Both the proposed single-step tilt estimator 6 and the relaxation Algorithm 7 proved to be unbiased.

5. The relaxation algorithm generally improves the estimate from the single-step algorithm. This improvement is better if the underlying surface is convex.

If the constraints on the illuminant position and the Lambertianity of the surface are violated, then

1. The average angular error of the local surface orientation does not exceed the value of the illuminant slant from the position when $\mathbf{L} = \mathbf{V}$.

2. The tilt estimates of the single-step estimator due to bad slant estimate are not worse than $1.25\times$ for the test ellipsoid and $3.4\times$ for the test hyperboloid. Paradoxically, when $n < 1$ the estimates can improve. The "improvement" is less than 3% for the ellipsoid and less than 30% for the hyperboloid.

			quantiles		
surface / *method*	mean	standard deviation	5%	median	95%
ellipsoid / *direct*	0.34	3.3	−0.38	0.17	1
ellipsoid / *needle map*	0.28	0.77	−0.29	0.23	1
hyperboloid / *direct*	−1.4	12	−12	−1.4	9.1
hyperboloid / *needle map*	0.57	13	−0.55	−0.072	2.3

Table 6.4: *Gaussian curvature recovered directly using the formulas in Alg. 7 and from the needle map recovered by the same algorithm. The table shows statistics computed on the relative-error map ($1 \mathrel{\widehat{=}} 100\%$).*

6.2.12 Notes on Gaussian Curvature Recovery

In Subsection 6.2.4 on page 105, where the results of different algorithms in recovering Gaussian curvature K have been presented, we explained why we used the curvature computed from the needle map instead of using the direct results from Algorithm 6 (7). The reason given there was briefly 'to compare *comparable* results.'

Here we will compare the results obtained by direct evaluation of the formula for K in Algorithm 7 with K computed from the needle map recovered by the same algorithm.

Standard statistics are evaluated in Tab. 6.4. Algorithm 7 ran five iterations with $\xi = 0.25$ and $R = 1 \cdot 10^{-3}$. The results of both methods for the standard ellipsoid does not differ significantly. It can be seen, however, that the results for the standard hyperboloid are much closer to reality when computed from the needle map. The results are also visualized in Fig. 6.26 on the next page.

The possibility of recovering Gaussian and mean curvatures from the needle maps could be used for evaluation of $\tau_{\mathbf{u}}$ and $\tau_{\mathbf{v}}$ appearing in (4.25) and (4.26). There is a possibility that a computationally robust method for recovering K and H from needle maps could be based on statistical analysis of normal directions around the pixel in question by utilizing the two following heuristics [Koe90, p. 215ff]:

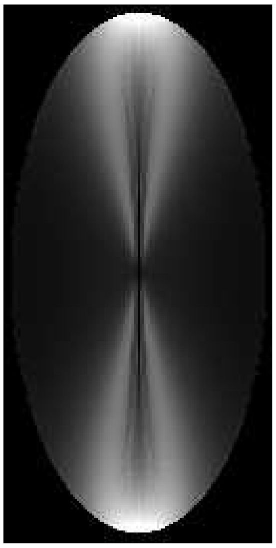
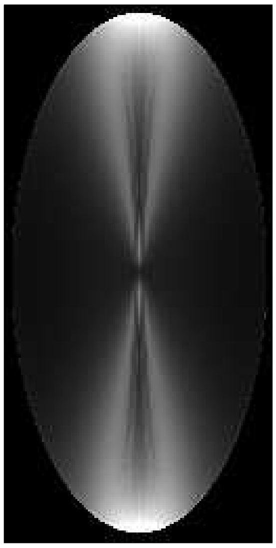

(a) *(b)*

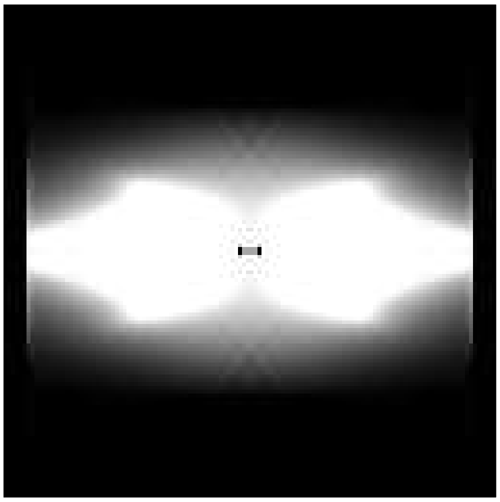
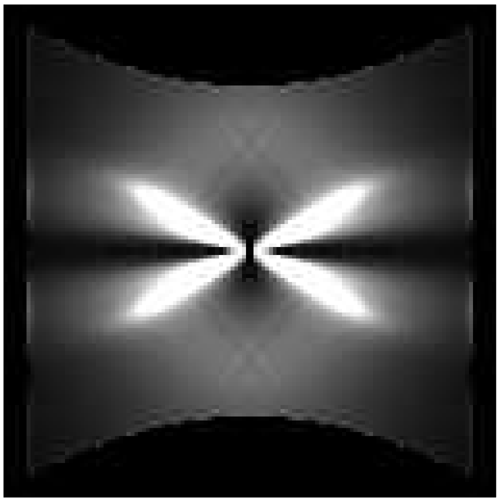

(a) *(b)*

Figure 6.26: *Gaussian curvature map recovered (a) directly, i.e., by evaluating the corresponding formulas in Alg. 7, (b) from the recovered needle map by evaluating (6.3).*

Heuristic 1

> *K measures the total spread of normal directions per unit surface area.*

Heuristic 2

> *H measures the turn of the normal direction for a unit step along the surface averaged over all directions.*

We suggest that the normal direction spread and the average turn could correspond to values computed from suitable statistics computed on a pixel neighbourhood. The question remains whether suitable (computationally stable) statistics exist and whether this approach would be more stable than the one we have used.

Chapter 7

Conclusions and Pointers

7.1 Open Questions

Robust Image Features Detectors

All *LSA* algorithms rely on the robust detection of image function features. High-order detectors are of particular importance since there is no general *LSA* algorithm for which the features ${\mathcal{F}_I}^1$ are sufficient. The computational behaviour of an n-th order detector deteriorates quickly with increasing n.

A reason why we do not present 3-D reconstruction results on natural images is that the link between the image and the image properties required for the algorithm is missing. (Another, more important reason is that we have not implemented a module for global interpretation based on linking multiple local interpretations to a consistent entity.) We currently do not have suitable robust estimators of ${\mathcal{F}_I}^2$ image features to be able to cope with the image noise in ill-posed areas. Our approach has been to use "blurred image derivatives" computed over Gaussian-weighted support field [TP86, Koe90]. This technique did not prove to be robust enough. Pentland in [Pen84] suggests how to estimate a second-order directional derivative by computing a Gaussian-weighted sum of blurred image Laplacians along a straight line in the direction in question. This approach has some support from the psychology of vision. There is a kind of cortical neurons in the primary visual cortex summing the responses from the center-surround receptive fields in a very similar way as Pentland has suggested [Ros79, Gau80, BG85]. Although we did not perform our own experiments, we expect the results will probably not show much difference from computing blurred image directional derivatives because the direction must be known in advance in both cases and it remains ill-posed at some loci.

A direct relationship between second-order image function partial derivatives and the geometric properties of the image isophotes and flow-lines was proved in [RFK91] and has been discussed in Section 4.1. The geometric interpretation of ${\mathcal{F}_I}^2$ image features suggest there could exist second-order image derivative estimators based on geometric principles. This is a motivation for research towards a robust detection method. A suitable approach could be motivated by the method used in early stages of computing *saliency maps* for perceptual grouping [GM92, SU88]. This method collects the support for hypothesis about a curve of certain curvature and orientation from a local neighbourhood of an image pixel *(extension field).* Such an approach could overcome computing any kind of "image derivatives" and could probably be implemented in a neural network very efficiently. This is a topic for future research that has to be done before the proposed algorithms can run on real images.

Local Analysis of a General Gradient Field

All the analysis given in Chapter 4 and the algorithms presented in Chapter 5 are based on local analysis of a *gradient field.* It was defined as a gradient field of a continuous $\mathcal{C}^2$ function corresponding to the radiance map of an illuminated surface without any texture.

But the field can be defined arbitrarily: it can, for instance, be the field of the texture gradient. If the surface texture satisfies some criteria (homogeneity, isotropy), the analysis given in Chapter 4 can be applied with only minor changes. We have not investigated this topic. This problem needs further research. A fact worth noting is that, in a texture image, there prevails the projection component of image formation, since the texture, i.e., the varying albedo, is almost an "intrinsic" property of the surface—its directional properties do not depend (much) on the actual illumination. The reflectance function component of the image formation plays only a minor role here. The consequence is that a surface texture image is not affected by mutual illumination among bodies in the scene.[1]

A very interesting question is the possible existence of a unified framework that could cope with both the texture and the radiance gradient fields.

Shading near Parabolic Lines

If we want to recover the qualitative structure of the viewed surface from a single monocular image based on local analysis, it is extremely important to have a Gaussian curvature zero-crossing detector, since the parabolic lines are very important for modelling the surface structure (also other singularities of Gaussian images occur on parabolic lines). Unfortunately, a reliable single-step detector of parabolic lines based on second-order image properties has not yet been found.

[1] There may be an apparent texture due to surface roughness. Such "texture" properties depend on illumination, of course.

Consistent Region Representation

An important problem that has not been touched here is that of building a consistent geometric interpretation of an image region. This requires linking local interpretations to larger regions. On a surface, the curves along which the interpretation may switch are the parabolic lines. But, since there is no parabolic line detector working on intensity images, there are no natural image curves along which the interpretation should switch. Linking must therefore be based on local curvature consistency: if we suppose there is a Darboux frame $\mathcal{D}(\mathbf{x} + \mathrm{d}\mathbf{x})$ attached to any neighbouring point of $\mathbf{x}$ in image, then a parallel transport of this frame to the point $\mathbf{x}$ on the surface must be consistent with the frame $\mathcal{D}(\mathbf{x})$. This approach is due to Ferrie & Lagarde [FL92].

7.2 Summary

In this work we have investigated how an image, together with a (partially) known reflectance function and imaging geometry, constrains the local geometric structure of the viewed surface. We have seen that it is a difficult problem. The difficulty is caused by the complex relation between the local surface shape and the intensity image of the viewed surface. Although the image irradiance depends on local surface orientation (normal), and the image gradient is related to the surface curvature, the complexity of the relation does not allow a simple formulation of the problem. This is the reason why all the previous attempts at accurate or at least qualitatively correct *LSA* failed.

We have examined the system of image isophotes near $\mathcal{C}^1$ surface discontinuities and vertices, near singular image points, and near occlusion and self-shadow contours, with the aim to get the least ambiguous constraints on the local surface orientation and curvature. The generic surface discontinuity contours can be distinguished from the occlusion contours on a purely local basis without high-level knowledge: the occlusion contour is always an envelope of images of all non-perpendicular surface curves approaching it from the interior of the surface patch, while a $\mathcal{C}^1$ surface discontinuity is generally not. At a point on an occlusion contour, the local surface orientation and quality can be determined uniquely: the normal vector lies in the image plane oriented perpendicularly to the contour and the sign of the contour curvature is equivalent to the sign of the Gaussian curvature of the surface there.

We have shown that, besides the occlusion contour, the self-shadow contour is also a rich source of unambiguous information about the surface. As a new result, we hypothesize that the local surface orientation can be determined uniquely at these locations if the illuminant position is known and the reflectance function is space-invariant. Also, the surface quality (Gaussian curvature) could probably be specified there in the general case. This problem needs further investigation.

The shading near a general $\mathcal{C}^1$ discontinuity does not constrain the shape sufficiently—there is an infinite (one-parametric) ambiguity. We have seen that, when several discontinuities meet at a vertex, the structure of the corresponding image edges, together with the shading near the edges junction, constrain the shape up to a finite ambiguity when there is no pair of parallel edges among those forming the junction in the image.

The shape is constrained up to a low degree of ambiguity at image singularities of certain type: if the imaging geometry and the type of the singularity are known, the local surface orientation can be specified uniquely at specular points and apices. Moreover, at an apex, principal surface curvatures can be recovered up to their sign, which gives four possible local surface interpretations. We have shown that Koenderink's 'curvedness,' which is a surface curvature measure, can be specified uniquely at an apex. Unfortunately, there is no monocular cue to the extremal singularity type. On the other hand, the saddle singularities are always of the same type (they occur only on parabolic lines). But we have found neither a local orientation detector nor a curvature detector for that case. This also needs further research.

We have studied shading out of singular image elements as well. In previous work on *LSA,* either a heuristic approach was chosen (assuming mostly non-realistic restrictions regarding the viewed surface) [Pen84, LR85], or an approach based on the geometry of the surface itself (Darboux frame chosen as a local surface-related coordinate system) [Oli90]. The drawback of the former is the lack of generality that causes non-integrability of the results and a low correspondence with the true surface [FC88, FL92]; the latter approach suffers from a bad tractability of the complex relation between the image function properties and the geometric properties of the viewed surface. To enforce the tractability, heuristic assumptions about the surface are made again, which, after all, leads to similar drawbacks as in the former method.

We have therefore chosen a physically-motivated approach that expresses the surface properties in terms of structures generated by the reflectance physics of the surface and whose generality allows further analysis of image events. The advantage of utilizing this concept is that the *LSA* problem becomes *well-tractable* with no a priori restriction on the surface form in the sense of restricting the surface type like Pentland in [Pen84], Lee & Rosenfeld in [LR85], or Ferrie & Levine in [FL89], or the surface order, like Arnspang in [Arn91], or assuming an explicit local surface approximation like Oliensis in [Oli90]. The drawback of the proposed method (like any other method based on the image gradient) is its ill-posedness at image extremal points. The method (like any of the known local methods) cannot be used for recovering surface normals and the mean curvature of developable (and planar) surfaces.

We show that isophotes possess an important property related to the viewed surface geometry: they are projections of curves of constant slant from the light direction if the surface reflectance is space-invariant or the illuminant is located at the vantage point. We also show that the isophote and the flow-line curvatures are invariant to the reflectance function under either of the two conditions. This results in a low sensitivity of the proposed local surface tilt estimator to the true reflectance function. The tilt estimator is based on the equation that relates the image gradient vector to the projection of the surface normal onto the image plane. This equation is derived provided that the illuminant is located at the vantage point, but a generalization to an arbitrarily positioned illuminant is possible.

We also present expressions for Gaussian and mean curvatures valid at a non-singular image point. We show that there are up to four local interpretations of the surface at any point of the image. The number of interpretations decreases to two near the occlusion contour.

Based on the analysis of the isophote field, a good tilt estimator and a relaxation algorithm for *LSA* were presented.

We have proposed a method of local shading analysis that:

1. recovers the local surface orientation (needle map) as well as the curvature (K map and H map),
2. maintains all local solutions,
3. does not assume any a priori restrictions on surface shape or its order,
4. the accuracy of the results increases with the computational effort,
5. the recovered needle map is integrable,
6. the tilt estimates are significantly more accurate than in the standard methods.

The approach we have used leads to a sufficiently simple problem description that allows further theoretical analysis of local shading.

The drawbacks of the approach are as follows:

1. it relies on known illuminant position, surface reflectance function, and albedo[2]
   (this is a general *LSA* feature),

2. it relies on the non-existence of inter-reflections between surfaces
   (this is a general assumption for any *LSA*),

3. it relies on second-order image function properties
   (this is a general drawback of *LSA*),

4. it does not work at general singular image points because the image gradient is ill-defined there
   (this is a feature of any *LSA* based on the image gradient),

5. it does not work on developable and planar surfaces
   (we believe this is a general drawback of any *LSA*).

[2]Although the sensitivity to the violation of those conditions is relatively low.

It can be seen that the drawbacks of the proposed approach are due to the fundamental limitations of local shading analysis. These limitations cannot be overcome within the scope of *LSA*. A global approach can solve the problems of 1 (see [LR85, Pen82, NB93]), 2 (see [NIK91a]), and 5 (see [BT81, IH81, Woo81]).

The requirement 3 implies that a good detector of second-order image features must exist. The need for such a detector for local shading analysis is not contradictory to the facts known from the neuropsychology of vision [BG85].

We have shown that the local orientation and surface curvatures at a singular image point of certain type can be determined with the same degree of ambiguity as at a general image point. The problem 4 can therefore be solved as well, in most cases.

The local cues to 3-D surface quality and structure are very important for low-level data acquisition for the bottom-up processing of the visual information without specific world knowledge. This is the reason why they are widely studied. Differential geometry is a suitable tool that allows to examine a visual problem both locally and globally. The theoretical and computational research is necessary for the understanding of human and machine vision.

Bibliography

[Arn91] J. Arnspang. Direct shape measures from shading and impossible shaded images. In P. Johansen and S. Olsen, editors, *Proceedings of the 7nd Scandinavian Conf. on Image Analysis*, pages 40–47, Aalborg University, Denmark, August 1991. Pattern Recognition Society of Denmark.

[Baj90] R. Bajcsy, 1990. Personal communication.

[BG85] V. Bruce and P. R. Green. *Visual Perception, Physiology, Psychology, and Ecology*. Lawrence Erlbaum Associated Ltd., 1985.

[BJ86] P. J. Besl and R. C. Jain. Invariant surface characteristics for 3-D object recognition in range images. *Computer Vision, Graphics, and Image Processing*, 33:33–80, January 1986.

[BPYA85] M. Brady, J. Ponce, A. Yuille, and H. Asada. Describing surfaces. *Computer Vision, Graphics, and Image Processing*, 32:1–28, 1985.

[BS90] L. G. Brown and H. Shvaytser. Surface orientation from projective foreshortening of isotropic autocorrelation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(6):584–588, June 1990.

[BT81] H. G. Barrow and J. M. Tenenbaum. Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, 17:75–116, 1981.

[BY84] M. Brady and A. Yuille. An extremum principle for shape from contour. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(3):288–301, May 1984.

[FC88] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451, July 1988.

[FL92] F. P. Ferrie and J. Lagarde. Curvature consistency improves local shading analysis. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 55(1):95–105, January 1992.

[FL89] F. P. Ferrie and M. D. Levine. Where and why local shading analysis works. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):198–206, February 1989.

[FRKV91] L. M. J. Florack, B. M. H. Romeny, J. J. Koenderink, and M. A. Viergever. General intensity transformations and second order invariants. In P. Johansen and S. Olsen, editors, *Proceedings of the 7nd Scandinavian Conf. on Image Analysis*, pages 338–345, Aalborg University, Denmark, August 1991. Pattern Recognition Society of Denmark.

[FZ91] D. Forsyth and A. Zisserman. Reflections on shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):671–679, July 1991.

[Gau80] J. G. Gaugman. Two-dimensional spectral analysis of cortical receptive fields profiles. *Vision Res.*, 20:847–856, 1980.

[GM92] G. Guy and G. Medioni. Perceptual grouping using global saliency-enhacing operators. In *Proc. ICPR Conf.*, pages 99–103, 1992.

[Háj69] J. Hájek. *Survey of Applicable Mathematics*, chapter Mathematical Statistics, pages 1264–1284. Iliffe Books, 1969.

[Hav56] B. Havelka. *Tři práce o kvalitě optického obrazu (Three Works on Optical Image Quality)*, volume 4 of *Řada matematických a přírodních věd*. ČSAV, Praha, 1956. In Czech.

[Hor75] B. K. P. Horn. Obtaining shape from shading information. In P. H. Winston, editor, *Psychology of Computer Vision*, pages 115–155. McGraw-Hill, New York, 1975.

[Hor77] B. K. P. Horn. Understanding image intensities. *Artificial Intelligence*, 8:201–231, 1977.

[HS79] B. K. P. Horn and R. W. Sjoberg. Calculating the reflectance maps. *Applied optics*, 18(11):1770–1779, 1979.

[HWL83] R. M. Haralick, L. T. Watson, and T. J. Laffey. The topographic primal sketch. *International Journal of Robotics Research*, 2:50–72, 1983.

[IH81] K. Ikeuchi and B. K. P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17:141–184, 1981.

[JH88] A. K. Jain and R. Hoffman. Evidence-based recognition of 3-D objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(6):783–801, November 1988.

[Ken79] J. R. Kender. Shape from texture: an aggregation transform that maps a class of textures into surface orientation. In *Proc. 6th Int. Joint Conf. on Artificial Intell.*, Tokyo, Japan, 1979.

[Koe90] J. J. Koenderink. *Solid Shape*. The M.I.T. Press, 1990.

[Koe84] J. J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13, 1984.

[Kv80] J. J. Koenderink and A. J. van Doorn. Photometric invariants related to solid shape. *Optica Acta*, 27(7):981–996, 1980.

[Len91] R. Lenz. Group theoretical feature extraction: Weighted invariance and texture analysis. In P. Johansen and S. Olsen, editors, *Proceedings of the 7nd Scandinavian Conf. on Image Analysis*, pages 127–134, Aalborg University, Denmark, August 1991. Pattern Recognition Society of Denmark.

[LF84] W. C. Lin and K. S. Fu. A syntactic approach to 3-D object representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(3):351, May 1984.

[LF86] W. Ch. Lin and K. S. Fu. A syntactic approach to three-dimensional object recognition. *IEEE Transactions on Systems, Man, and Cybernetics*, 16(3):405–422, May/June 1986.

[Lip69] M. M. Lipschutz. *Differential Geometry*. Shaum's Outline Series. Mc Graw-Hill, 1st edition, 1969.

[LR85] Ch.-H. Lee and A. Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, 26:125–143, 1985.

[LZ87] Y. G. Leclerc and S. W. Zucker. The local structure of image discontinuities in one dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(3):341–355, May 1987.

[Mac73] A. K. Mackworth. Interpreting pictures of polyhedral scenes. *Artificial Intelligence*, 4:121–137, 1973.

[Mar82] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman & Co., San Francisco, 1982.

[May79] P. S. Maybeck. *Stochastic Models, Estimation, and Control*, volume 1 of *Mathematics in Science and Engineering*. Academic Press, 1979.

[MN87] M. Magee and M. Nathan. A viewpoint independent modelling approach to object recognition. *IEEE Journal of Robotics and Automation*, 3(4):351–356, August 1987.

[NB93] S. K. Nayar and R. M. Bolle. Computing reflectance ratios from an image. *Pattern Recognition*, 26(10):1529–1542, 1993.

[NIK91a] S. K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *Int. Journal of Computer Vision*, 6(3):173–195, 1991.

[NIK91b] S. K. Nayar, K. Ikeuchi, and T. Kanade. Surface reflection: Physical and geometrical perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):611–634, July 1991.

[NRH+77] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. NBS Monograph 160, U. S. Department of Commerce, National Bureau of Standards, 1977.

[OD93] J. Oliensis and P. Dupuis. A global algorithm for shape from shading. In *Proc. 4th Int. Conf. on Computer Vision*, pages 692–701, Berlin, Germany, May 1993. IEEE Computer Society Press.

[Oli90] J. Oliensis. New results in shape from shading. In *procs. Image Understanding Workshop*, pages 145–153, September 1990.

[Oli91] J. Oliensis. Shape from shading as a partially well-constrained problem. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 54(2):163–183, September 1991.

[Pen82] A. P. Pentland. Finding the illuminant direction. *J. Optical Soc. Amer.*, 72(4):448–455, April 1982.

[Pen84] A. P. Pentland. Local shading analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(2):170–187, March 1984.

[Pen86] A. P. Pentland. Shading into texture. *Artificial Intelligence*, 29:147–170, 1986.

[Pho75] B. T. Phong. Illumination for computer generated pictures. *Comm. ACM*, 18(6):311–317, June 1975.

[PK90] J. Ponce and D. J. Kriegman. On recognizing and positioning curved 3-D objects from image contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(12):1127–1137, December 1990.

[PSH85] T. Ch. Pong, L. G. Shapiro, and R. M. Haralick. Shape estimation from topographic primal sketch. *Pattern Recognition*, 18(5):333–347, 1985.

[PVM+93] L. Přeučil, J. Vašek, J. Martínek, R. Mařík, and R. Šára. Vision-based robot navigation in the 3-D world. In *Proc. 4th Conf. AI, Simulation, and Planning in High Autonomy Systems*, pages 200–204, Tucson, Arizona, USA, September 1993. IEEE Computer Society Press.

[RFK91] B. M. H. Romeny, L. M. J. Florack, and J. J. Koenderink. Invariant third order properties of isophotes: T-junction detection. In P. Johansen and S. Olsen, editors, *Proceedings of the 7nd Scandinavian Conf. on Image Analysis*, pages 346–353, Aalborg University, Denmark, August 1991. Pattern Recognition Society of Denmark.

[Ros79] D. Rose. Mechanisms underlying the receptive field properties in cat visual cortex. *Vision Res.*, 19:533–544, 1979.

[Šár94] R. Šára. The impact of smooth surface perception capability on the structure of a vision system. *Cybernetics and Systems*, 25(1):163–176, January 1994. Presented at 3rd Int. Workshop on CAST, Las Palmas, Spain, February 21-26, 1993.

[Šár92] R. Šára. The surface structure of 3-D curved objects. In H. Bischof and W. Kropatch, editors, *Proceedings of 16th ÖAGM Workshop*, pages 129–138, Vienna, May 1992.

[Sax89] B. V. H. Saxberg. *A Modern Differential Geometric Approach to Shape from Shading.* PhD thesis, MIT Artificial Intelligence Laboratory, 1989. (TR 1117).

[Ste81] K. A. Stevens. The visual interpretation of surface contours. *Artificial Intelligence*, 17:47–73, 1981.

[Sto69] J. J. Stoker. *Differential Geometry.* John Wiley and Sons, 1969.

[SU88] A. Sha'ashua and S. Ullman. Structural saliency: the detection of globally salient structures using a locally connected network. In *Proceedings of the ICCV*, pages 321–327, 1988.

[Td88] H. D. Tagare and R. J. P. de Figueriedo. Simultaneous estimation of shape and reflectance maps from photometric stereo. In *Proc. 3rd Int. Conf. Comput. Vision*, pages 340–343, December 1988.

[Ter88] D. Terzopoulos. The computation of visible-surface representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):417–437, July 1988.

[TP81] H. D. Tagare and de Figueiredo R. J. P. A theory of photometric stereo for a class of diffuse non-Lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(2):133–152, February 1981.

[TP86] V. Torre and T. A. Poggio. On edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(2):147–163, 1986.

[TS92] P. S. Tsai and M. Shah. A simple shape from shading algorithm. Technical report, University of Central Florida, Orlando, FL., 1992.

[Wit81] A. P. Witkin. Recovering surface shape and orientation from texture. *Artificial Intelligence*, 17:17–45, 1981.

[Woo81] R. J. Woodham. Analysing images of curved surfaces. *Artificial Intelligence*, 17:117–140, 1981.

[Woo80] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19:139–144, 1980.

CURRICULUM VITAE

Personal data: Ing. Radim Šára, born on July 11th, 1963 in Prague, Czech Republic, single.

Office Address: Department of Control Engineering, Faculty of Electrical Engineering, Czech Technical University, Karlovo náměstí 13, CZ–121 35 Prague 2, Czech Republic.

Telephone: (+422) 295664, fax: (+422) 290159, e-mail: sara@vision.felk.cvut.cz

Employment: Research fellow at the Department of Control Engineering, Faculty of Electrical Engineering, Czech Technical University in Prague (since 1990).

Qualifications:

1982–1987 Study at TU Prague, Electrotechnical Faculty, stream Technical Cybernetics, specialization Robotics. Master thesis: *Classification and Image Description Using Relational Structures* (in Czech).

1987–1990 Post-graduate study at the Dept. of Electrotechnology, at the same school.

1990–present Doctoral study at the Dept. of Control Engineering, at the same school.

1992–1993 A two-year doctoral-student research place at the Institut für Systemwissenschaften, Johannes Kepler Universität Linz, Austria.

Research Projects:

1991–1992 Co-investigator of the research project "Model-Based 3D Computer Vision". Grant C-04 of the I.B.M. Corporation Initiative in Czechoslovakia.

1991–1993 Joint-applicant and co-investigator of the research project "Digitalizovaná rekonstrukce chrupu a její aplikace ve formě artikulačního přístroje". Applicant: the Institute of Dental Research, Prague. Grant Z-038 of the Ministry of the Economic Policy and the Czech Republic Development.

1992–1993 Co-investigator of the research project "Theory and Applications of Digital Image Processing and Pattern recognition". Applicant: PRIP TU-Wien. Grant P-8785 FWF Österreich.

1992 Co-investigator of the research project "New Image Analysis Algorithms". Internal grant 8066 of TU-Prague.

1992 Co-investigator of the research project "Robot Orientation in a 3D World" Internal grant 8093 of TU-Prague.

1993–1995 Co-investigator of the research project "3D Shape Reconstruction and Modelling Using Active Vision". Grant 102/93/0954 of the Grant Agency of the Czech Republic.

Research Interests: Computer vision; structural description, reconstruction, and recognition of three-dimensional objects; local shading analysis; applied mathematical methods.

Selected Publications:

R. Šára. Graph structural description of two-dimensional patterns. In *Proceedings of Modern Trends in Cybernetics and Systems*, Temešvár, 1990.

T. Dostálová, J. Dušková, and R. Šára. Digitalizace vyhodnocování plošných měření ve stomatologii. *Lékař a technika*, 22(5-6):73–75, 1991. In Czech (A digital method of planimetric measurements evaluation for stomatology).

T. Dostálová, R. Šára, and V. Hlaváč. A digital reconstruction of occlusion. In *Proceedings of Biosignal 92 Conf.*, pages 28–30, Brno, June 1992.

R. Šára. The surface structure of 3-D curved objects. In H. Bischof and W. Kropatch, editors, *Proceedings of 16th ÖAGM Workshop*, pages 129–138, Vienna, May 1992.

R. Šára. Projection of smooth surface invariants into intensity image. In Pölzleitner W. and Wenger E., editors, *Image Analysis and Synthesis (Proc. 17th ÖAGM Workshop)*, pages 83–92, Graz, Austria, June 1993. Österreichische Computer Gesellshaft, Wien; München: Oldenbourg.

R. Šára. The impact of smooth surface perception capability on the structure of a vision system. *Cybernetics and Systems*, 25(1):163–176, January 1994. Presented at 3rd Int. Workshop on CAST, Las Palmas, Spain, February 21-26, 1993.

L. Přeučil, J. Vašek, J. Martínek, R. Mařík, and R. Šára. Vision-based robot navigation in the 3-D world. In *Proc. 4th Conf. AI, Simulation, and Planning in High Autonomy Systems*, pages 200–204, Tucson, Arizona, USA, September 1993. IEEE Computer Society Press.

R. Šára. Local shading analysis revisited. In V. Hlaváč and T. Pajdla, editors, *Proceedings of the 1st Czech Pattern Recognition Workshop*, pages 67–74, Temešvár u Písku, Czech Republic, November 1993. Czechoslovak Pattern Recognition Society.

T. Dostálová, R. Šára, and J. Procházková. The digital determination of static occlusion type based on Angle classification. *International Journal on Adult Orthodontics and Orthodonthic Surgery*. (to appear).