
Robust and Efficient Processing Techniques for Static and Dynamic Geometric Data

Oliver Schall

**Max-Planck-Institut für Informatik
Saarbrücken, Germany**

Dissertation zur Erlangung des Grades
Doktor der Ingenieurwissenschaften (Dr.-Ing.)
der Naturwissenschaftlich-Technischen Fakultät I
der Universität des Saarlandes

Eingereicht am 18. März 2009 in Saarbrücken.

Betreuender Hochschullehrer – Supervisor

Dr. Alexander Belyaev,
Heriot-Watt University, Edinburgh, United Kingdom

Gutachter – Reviewers

Prof. Dr. Hans-Peter Seidel,
Max-Planck-Institut für Informatik, Saarbrücken, Germany
Dr. Alexander Belyaev,
Heriot-Watt University, Edinburgh, United Kingdom

Prüfungsvorsitzender – Head of the committee

Prof. Dr. Philipp Slusallek,
Universität des Saarlandes, Saarbrücken, Germany
Deutsches Forschungszentrum für Künstliche Intelligenz, Saarbrücken, Germany

Dekan – Dean

Prof. Dr. Joachim Weickert,
Universität des Saarlandes, Saarbrücken, Germany

**Promovierter akademischer Mitarbeiter –
Academic Member of the Faculty having a Doctorate**

Dr. Michael Wand,
Universität des Saarlandes, Saarbrücken, Germany
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Datum des Kolloquiums – Date of Defense

28. August 2009 – August 28th, 2009

Oliver Schall
MPI Informatik
Campus E1 4
66123 Saarbrücken
schall@mpi-inf.mpg.de

Abstract

Generating high quality geometric representations from real-world objects is a fundamental problem in computer graphics which is motivated by manifold applications. They comprise image synthesis for movie production or computer games but also industrial applications such as quality assurance in mechanical engineering, the preservation of cultural heritage and the medical adaptation of prostheses or orthoses. Common demands of these applications on their underlying algorithms are robustness and efficiency. In addition, technological improvements of scanning devices and cameras which allow for the acquisition of new data types such as dynamic geometric data, create novel requirements which rise new challenges for processing algorithms. This dissertation focuses on these aspects and presents different contributions for flexible, efficient and robust processing of static and time-varying geometric data. Two techniques focus on the problem of denoising. A statistical filtering algorithm for point cloud data building on non-parametric density estimation is introduced as well as a neighborhood filter for static and time-varying range data which is based on a novel non-local similarity measure. The third contribution unifies partition of unity decomposition and a global surface reconstruction algorithm based on the Fast Fourier Transform which results in a novel, robust and efficient reconstruction technique. Concluding, two flexible and versatile tools for designing scalar fields on meshes are presented which are useful to facilitate a controllable quadrangular remeshing.

Kurzfassung

Die Erzeugung von hochqualitativen geometrischen Darstellungen realer Objekte ist ein grundlegendes Problem der Computergraphik, welches von vielfältigen Anwendungen motiviert ist. Diese umfassen die Synthese von Bildern zur Herstellung von Filmen oder Computerspielen, aber auch industrielle Anwendungen, wie die Qualitätsprüfung im Maschinenbau, die Erhaltung von Kulturgütern und die medizinische Anpassung von Prothesen oder Orthesen. Gemeinsame Anforderungen dieser Anwendungen an die beteiligten Algorithmen sind Robustheit und Effizienz. Zusätzlich erzeugen technologische Fortentwicklungen von Scannern und Kameras, welche die Aufnahme von neuen Datentypen, wie dynamischen Geometriedaten, erlauben, Anforderungen, die neue Herausforderungen für Verarbeitungsalgorithmen darstellen. Diese Dissertation betrachtet diese Aspekte und präsentiert verschiedene neue Beiträge zur flexiblen, effizienten und robusten Verarbeitung statischer und zeitlich dynamischer geometrischer Daten.

Zwei Verfahren befassen sich mit dem Problem des Entrauschens. Ein statistischer Algorithmus zur Filterung von Punktwolken, welcher auf nicht-parametrischer Dichteschätzung beruht, wird eingeführt, sowie ein Nachbarschaftsfilter für statische und zeitlich variierende Entfernungsdaten, welcher auf einem neuen nicht lokalen Ähnlichkeitsmaß basiert. Der dritte Beitrag verbindet eine Zerlegung basierend auf der Partitionierung der Eins mit einem globalen Algorithmus zur Oberflächenrekonstruktion, welcher auf der schnellen Fouriertransformation beruht. Dies resultiert in einem neuen, robusten und effizienten Rekonstruktionsverfahren. Abschließend werden zwei flexible und vielseitige Werkzeuge zum Design von Skalarfeldern auf Netzen vorgestellt, welche eine steuerbare Vierecksneuvernetzung ermöglichen.

Summary

A central problem in computer graphics is the generation of high quality meshes from real-world objects as they build the foundation for many graphics algorithms. Their versatile applicability motivated intensive research which produced a broad range of processing techniques. These methods focus on different problems of the mesh generation process. Examples comprise the acquisition of geometry, data preprocessing, surface reconstruction and remeshing. The different techniques are applied in a sequence which is referred to as the geometry processing pipeline.

Demanding requirements on processing algorithms are robustness and efficiency. In addition, the continuous improvement of scanning devices and digital cameras drives the availability of new data types such as dynamic range data which entail new challenges for processing algorithms.

This dissertation focuses on these aspects and presents different contributions for the flexible, efficient and robust processing of static and time-varying geometric data which address different steps of the geometry processing pipeline.

Two techniques consider the problem of denoising. The first method is devoted to the robust filtering of point cloud surface data. It approaches the problem statistically. Building on non-parametric density estimation, a smooth probability function is defined to move noisy input points to their maximum likelihood positions. The derived technique is insensitive to parameter variations, handles high noise amplitudes properly and exhibits a good outlier resistance. The second technique chooses a different approach to the denoising problem by filtering range scans individually before they are merged in a point cloud in order to be able to utilize their inherent structure. The method focuses on denoising static and time-varying range data whose acquisition has increasingly become feasible due to improvements in scanning technology. The technique introduces a novel non-local similarity measure which determines the resemblance of two points on a range scan not only by their local properties but by incorporating context information and comparing the surface regions around them. The measure is utilized to develop a denoising algorithm which produces a more accurate denoising result on range data than previous state-of-the-art approaches while having a better feature-preservation. Additionally, an extension is presented that enables a robust and stable treatment of time-varying range data by augmenting the similarity measure into the time domain.

The first chapters of this thesis advocate that data cleaning improves the performance of subsequent processing steps such as surface reconstruction. In a new algorithm, we use partition of unity decomposition to develop a novel local variant

of a global surface reconstruction technique based on the Fast Fourier Transform. The new approach is characterized by the robustness of the global technique while being more memory efficient which allows it to achieve a higher reconstruction accuracy.

The fourth contribution considers controllable quadrilateral remeshing. We introduce two flexible and versatile tools for designing scalar fields on surfaces. The first can enforce directional constraints which is useful to avoid tedious post-processing, for instance, by aligning quads to important features while the second utilizes quasi-harmonic fields to consider the surface as an inhomogeneous domain to endow it with attraction or repulsion properties.

Zusammenfassung

Ein zentrales Problem der Computergraphik ist die Erzeugung hochqualitativer Netze realer Objekte, da sie die Grundlage vieler Graphikalgorithmien darstellen. Ihre vielseitige Anwendbarkeit motivierte intensive Forschung, was ein breites Angebot an Verarbeitungstechniken hervorbrachte. Die Methoden betrachten verschiedene Probleme des Netzgenerierungsprozesses. Beispiele sind die Akquise von Geometrie, die Vorverarbeitung von Daten, Oberflächenrekonstruktion und Neuvernetzung. Die verschiedenen Techniken werden sequentiell in einer Folge angewendet, die als geometrische Datenverarbeitungskette bezeichnet wird.

Anspruchsvolle Anforderungen an die Verarbeitungsalgorithmen sind Robustheit und Effizienz. Des Weiteren fördert die anhaltende Fortentwicklung von Scannern und Digitalkameras die Verfügbarkeit von neuen Datenformen, wie dynamischen Entfernungsdaten, welche neue Herausforderungen für Verarbeitungsalgorithmen mit sich bringen.

Diese Dissertation betrachtet diese Aspekte und stellt verschiedene Beiträge zur flexiblen, effizienten und robusten Verarbeitung statischer und zeitlich variierender geometrischer Daten vor, die sich mit verschiedenen Schritten der geometrischen Datenverarbeitungskette befassen.

Zwei Verfahren behandeln das Problem des Entrauschens. Die erste Methode ist der robusten Filterung von Punktwolken, welche Oberflächen darstellen, gewidmet und nähert sich dem Problem statistisch. Aufbauend auf nicht-parametrischer Dichteschätzung wird eine glatte Wahrscheinlichkeitsfunktion definiert, um die verrauschten Punkte der Eingabedaten zu deren Positionen maximaler Wahrscheinlichkeit zu verschieben. Das hergeleitete Verfahren ist unempfindlich gegenüber Parameteränderungen, fähig große Rauschamplituden zu verarbeiten und zeigt eine gute Beständigkeit gegenüber Ausreißern. Das zweite Verfahren wählt einen anderen Entrauschungsansatz und filtert die Entfernungsdatenscans einzeln bevor sie zu einer Punktwolke zusammengefasst werden, um deren inherente Struktur nutzen zu können. Die Methode beschäftigt sich mit dem Entrauschen statischer und zeitlich variierender Entfernungsdaten, deren Akquise zunehmend durch die Fortentwicklung von Scannertechnologien ermöglicht wurde. Der Ansatz führt ein neues nicht lokales Ähnlichkeitsmaß ein, welches die Ähnlichkeit zweier Punkte von Entfernungsdaten nicht nur durch deren lokale Eigenschaften bestimmt, sondern deren Kontext einbezieht und Regionen der Oberfläche um sie herum vergleicht. Das Maß wird zur Entwicklung eines Entrauschungsverfahrens verwendet, welches ein genaueres Filterresultat auf Entfernungsdaten als bisherige aktuelle Techniken erreicht und gleichzeitig Details besser erhält. Zusätzlich wird eine Erweiterung vorgestellt, die eine robuste und stabile Verarbeitung zeitlich

variierender Entfernungsdaten ermöglicht, durch eine Erweiterung des Ähnlichkeitsmaßes in die zeitliche Dimension.

Die ersten Kapitel dieser Dissertation befürworten das Entrauschen von Daten, da es die Performance nachfolgender Verarbeitungsschritte, wie Oberflächenrekonstruktion, verbessert. In einem neuen Verfahren wird eine Zerlegung basierend auf der Partitionierung der Eins verwendet, um eine neue lokale Variante einer globalen Oberflächenrekonstruktionstechnik zu entwickeln, welche auf der schnellen Fouriertransformation basiert. Das neue Verfahren zeichnet sich durch die Robustheit der globalen Methode aus, während es speichereffizienter ist und damit eine höhere Rekonstruktionsgenauigkeit erlaubt.

Der vierte Beitrag beschäftigt sich mit der steuerbaren Neuvernetzung in Vierecksoberflächen. Es werden zwei flexible und vielseitige Werkzeuge zum Design von Skalarfeldern auf Oberflächennetzen eingeführt. Das erste kann Richtungseinschränkungen festlegen, was nützlich ist, um umständliches Nachverarbeiten zu vermeiden, z.B. bei der Ausrichtung der Vierecke an wichtigen Details. Das zweite Werkzeug verwendet quasi-harmonische Felder, um die Oberfläche als ungleichförmige Umgebung aufzufassen und sie mit Anziehungs- oder Abstoßungseigenschaften zu versehen.

Acknowledgements

First of all, I would like to thank my supervisor Dr. Alexander Belyaev for his continuous support, his valuable comments, and for giving me the freedom to pursue my own ideas.

I am especially grateful to Prof. Dr. Hans-Peter Seidel for giving me the opportunity to perform my doctoral studies at the Max-Planck-Institute for Computer Science. With the Computer Graphics Group he shaped a unique working environment I enjoyed to be part of.

During my time in the group I had the opportunity to meet and work with many talented researchers who contribute to a great working atmosphere in which I found many new friends. In particular, I owe special thanks to Volker Scholz, Torsten Langer, Rhaleb Zayer, Christian Rössl, Waqar Saleem, Zachi Karni and Hitoshi Yamauchi for fruitful discussions.

I thank Conny Liegl and Sabine Budde for their help in administrative matters and Martin Fuchs for his technical support.

Furthermore, I am grateful to my parents and my complete family who encouraged and supported me during my whole life. Especially, I would like to thank Katrin Dedden for her patience, support, and for enriching my life in numerous ways.

Contents

1	Introduction	1
1.1	Main Contributions	2
1.2	Chapter Overview	3
2	The Geometry Processing Pipeline	5
2.1	Geometry Acquisition	5
2.1.1	Contact Approaches	6
2.1.2	Non-Contact Approaches	6
2.2	Registration	8
2.3	Denoising	9
2.4	Surface Reconstruction	9
2.5	Remeshing	9
3	Robust Filtering of Noisy Scattered Point Data	11
3.1	Introduction	11
3.2	Likelihood and Convergence	14
3.2.1	Likelihood Function	14
3.2.2	Convergence	16
3.2.3	Adaptive Kernel Size	18
3.3	Results and Applications	19
3.3.1	Filtering and Outlier Removal	20
3.3.2	Surface Reconstruction	22
3.4	Conclusion	23
4	Feature-preserving Non-local Denoising of Dynamic Range Data	25
4.1	Introduction	25
4.2	Related Work	28
4.2.1	PDE Approaches	28
4.2.2	Spectral Techniques	28
4.2.3	Neighborhood Filtering	28

4.2.4	Projection-based Approaches	29
4.2.5	Statistical Techniques	29
4.2.6	Dynamic Data Acquisition and Processing	30
4.3	Non-local Denoising	30
4.3.1	Non-local Image Filtering	31
4.3.2	Static Range Data	32
4.3.3	Time-varying Range Data	36
4.4	Results	37
4.5	Conclusion	41
5	Error-guided Adaptive Fourier-based Surface Reconstruction	43
5.1	Introduction	43
5.2	Related Work	45
5.2.1	Delaunay-based Techniques	45
5.2.2	Radial Basis Functions (RBFs)	47
5.2.3	Partition of Unity (PU)	49
5.2.4	Fourier-based Surface Reconstruction	50
5.2.5	Other Reconstruction Techniques	52
5.3	Adaptive FFT-based Surface Reconstruction	53
5.3.1	Adaptive Octree Subdivision	53
5.3.2	Integration	56
5.4	Results	58
5.5	Conclusion	62
6	Controlled Field Generation for Quad Remeshing	65
6.1	Introduction	65
6.2	Related Work	66
6.3	Constrained Fields on Surface Meshes	68
6.3.1	Gradient Constraints	68
6.3.2	Quasi-harmonic Fields	69
6.4	Quad Construction	71
6.5	Results and Discussion	74
6.6	Conclusion	79
7	Discussion and Conclusion	81
7.1	Summary	81
7.1.1	Statistical Denoising	81
7.1.2	Non-local Temporal Denoising	82
7.1.3	Surface Reconstruction	82
7.1.4	Controlled Quad Remeshing	83
7.2	Future Work	83

A Publications**85****Bibliography****87**

List of Figures

2.1	Schematic laser scanner setup.	7
3.1	Robust filtering of a structured light face scan corrupted by small-amplitude noise and outliers.	12
3.2	Illustration of the kernel density estimation technique for 1D scattered point data. Local maxima of the kernel estimation \hat{f} define cluster centers of the original data.	14
3.3	2D example of the weighted least-squares fitting plane and ellipsoid kernel computation.	16
3.4	A color coded slice of the likelihood function L defined for the noisy Dragon head model.	17
3.5	The effect of adaptive kernels on the filtering performance on the example of the noisy Dragon dataset.	18
3.6	Removal of a dense cloud of outliers added to the Stanford Bunny dataset.	20
3.7	Smoothing results for two range scans of the Ippolita Sforza Bust.	21
3.8	Denoising results for the Bimba model composed of raw registered range data obtained using a laser scanner.	22
3.9	Surface reconstruction results before and after our filtering procedure of the Dragon scans from the Stanford Scanning Repository.	24
4.1	Non-local denoising result of a raw laser scan of the Bimba model.	26
4.2	The definition of the similarity of neighborhoods in the image case is based on the pixel-wise difference of intensity values.	31
4.3	The effect of adaptive neighborhoods on the denoising performance on the example of a structured light scan.	33
4.4	Comparison of bilateral and non-local filtering on two acquired noisy range sequences.	35
4.5	Comparison of feature-preservation properties of bilateral and non-local filtering on a laser range scan of the Turbine Blade model.	38
4.6	Non-local filtering results for different noise models.	39

5.1	Illustration of a local surface approximation for points inside and in the vicinity of an octree leaf cell.	54
5.2	Example octree configuration for partition of unity blending. All local characteristic functions have a common global resolution. . .	55
5.3	Reconstructions of the XYZ RGB Dragon model using the FFT-based, the Poisson-based approach and our technique.	57
5.4	Reconstructions of the Thai Statuette created using the global FFT-based approach, the Poisson-based method and our technique. . .	59
5.5	Comparison of our reconstruction approach with other state-of-the-art techniques illustrated on the Dragon head scans.	60
5.6	Illustration of blending for a reconstruction of the Armadillo dataset composed of 114 registered range scans using our technique. . . .	61
6.1	Heat distribution on a homogeneous and inhomogeneous plate, modeled using the Laplacian operator and the quasi-harmonic operator respectively.	68
6.2	Contours of quasi-harmonic fields with different conductivity ratios on a rectangular plate.	70
6.3	Modeling contours of a quasi-harmonic field on a rectangular plate.	71
6.4	Comparison between the approach introduced by Dong et al. and our technique.	72
6.5	Treatment of the special case where a line constraint covers two edges of one triangle.	73
6.6	Remeshing of a reconstructed laser-scanned Turbine Blade model guided by gradient constraints.	75
6.7	Remeshing of a laser-scanned hand model guided by gradient constraints.	76
6.8	Robustness properties of the remeshing approach on data with irregular connectivity.	77
6.9	Remeshing of the mannequin model using areas with modified conductivity and gradient constraints.	78

List of Tables

3.1	Timings for the ellipsoid kernel computation and for filtering the models presented in this chapter.	19
4.1	Parameter settings and timings for the results presented in this chapter.	40
5.1	Timings and memory consumption for the reconstructions shown in Figure (5.5).	63
5.2	Reconstruction information for the results presented in this chapter and computed using our method.	64

Chapter 1

Introduction

Computer graphics is a research field which is characterized by fast progress and development making it a wide research area. Starting from early applications as, for instance, the computer drawing device Sketchpad which allowed to draw simple primitives on a computer screen under the use of a light pen at the beginning of the 1960s, it now spreads over diverse focus areas in geometry processing, rendering and computer animation which are motivated by manifold applications in industrial engineering and entertainment.

The foundation for many graphics algorithms is the availability of a digital model of a real-world object. One option to create high quality models is their manual design by skilled artists which usually takes extensive effort in terms of time and labor. An alternative is the acquisition of a real-world model and its transformation into a geometric representation suitable for subsequent applications. This process is performed in different steps which are combined in the geometry processing pipeline whose automation and development is one of the fundamental problems and a long standing goal in computer graphics research.

Geometric data is commonly represented in the form of tri- or quadrangular surface meshes which are ideally composed of regular facets. Other applied representations are, for instance, implicit surfaces which are described by the level sets of a function or point clouds which can be rendered as splats if no explicit connectivity is required.

The requirements on the resulting mesh are versatile and dependent upon the subsequent applications which comprise image synthesis for feature production or computer games or industrial applications as the acquisition of mechanical parts

for design or surface analysis purposes such as tolerance testing for quality assurance. Other interesting applications are the preservation of cultural heritage in a digital form to support their restoration process or the medical adaptation of prostheses or orthoses.

These applications make great demands on the processing algorithms in terms of robustness, accuracy, flexibility and efficiency. Furthermore, technological improvements of cameras and scanning devices drive the availability of new types of data which rise new challenges for processing algorithms.

This thesis focuses on these aspects and presents new techniques for the robust and efficient processing of static *and* time-varying range data which cover different steps of the geometry processing pipeline.

1.1 Main Contributions

Contributions are made in the area of point cloud and dynamic range scan filtering. By approaching the denoising problem statistically or by considering it as a non-local problem, improvements in robustness and filtering accuracy are achieved. Furthermore, the problems of surface reconstruction and quadrilateral remeshing are addressed. The developed reconstruction approach combines the advantages of an efficient local processing realized by partition of unity decomposition, which is yet sufficiently global, indicated by the size of the reconstructed primitives, to ensure a more robust and efficient reconstruction compared to recent state-of-the-art techniques. The quadrangular remeshing approach presents novel tools for the flexible design of scalar fields on meshes which are used for a controllable remeshing of irregular triangular meshes.

This dissertation builds upon scientific articles which have been published in various journals [Saleem07, Schall07a, Schall08a] and conference proceedings [Schall05b, Isgro05, Schall05a, Schall06a, Schall06b, Schall07b, Schall08b]. A compiled list of publications can be found in Appendix A. The main contributions of this dissertation can be summarized as follows:

- A robust statistical point cloud denoising and clustering algorithm derived from a non-parametric density estimation technique called Mean Shift.
- A non-local similarity measure to determine the resemblance of two points on a range scan depending on their geometric neighborhood. The measure was used to develop an efficient feature-preserving neighborhood filtering technique for static as well as dynamic range data which is an interesting

novel data format available due to important progress in the area of 3D video acquisition.

- A robust and efficient surface reconstruction approach which unifies partition of unity decomposition and a global reconstruction technique based on the Fast Fourier Transform.
- Flexible and versatile tools for designing scalar fields on surfaces. They can enforce directional constraints by operating on the gradient of the scalar field and area constraints by regarding the surface as an inhomogeneous domain to endow regions with attraction/repulsion properties. These techniques are used to develop a flexible algorithm for controllable quadrangular remeshing.

1.2 Chapter Overview

This thesis starts with an overview of the geometry processing pipeline in Chapter 2. Following the steps of the pipeline, the contributions of this dissertation are presented. Chapter 3 describes the novel statistical point cloud denoising approach. The filtering technique for static and time-varying range data (range video) is introduced in Chapter 4 followed by the surface reconstruction approach in Chapter 5. In Chapter 6, the quad-remeshing technique based on quasi-harmonic fields which allows for design constraints is presented before this thesis is concluded in Chapter 7 where all contributions as well as open questions are discussed.

Chapter 2

The Geometry Processing Pipeline

Creating a geometric representation for a real-world object is a process which consists of several steps. These steps are performed in a sequence combined in the geometry processing pipeline. This chapter presents an overview of a pipeline which focuses on reconstructing a well-behaved mesh representation from several range scans obtained from a real-world object. Extending the focus of the pipeline to subsequent processing, more steps than the ones introduced in this chapter are possible. Examples comprise hole filling, surface simplification, the estimation of differential quantities as for example curvature or surface parameterization and modeling.

2.1 Geometry Acquisition

The first step of the pipeline is to capture the geometry of a real-world object in form of points embedded in three-dimensional space. For this problem, different approaches have been proposed in the past which can be categorized into different groups according to their way of sampling the surface of the real-world object.

2.1.1 Contact Approaches

Contact-based devices sample the object by physical touch using a probe. This is performed either automatically or hand-driven. Application areas are for instance manufacturing since the approaches can be very precise and the computer animation industry where hand-crafted models are digitized. Since contact is required the probe has to be moved sequentially over the surface causing the methods to be rather slow. Depending on the object such as historical artifacts another limitation is that contact is not desired due to the risk of damage.

2.1.2 Non-Contact Approaches

Non-contact approaches sample the environment or object optically. For a given position of the camera their goal is to reconstruct a depth map of the recorded scene where each pixel of the image is associated with the distance from the image plane to the recorded part of the surface. By obtaining multiple depth maps from different viewing directions the geometry of a complete object can be acquired. The alignment of the individual range images into a common coordinate system is performed in a subsequent registration.

Two groups of techniques have been established. Active approaches emit light into the scene or onto the object in order to support the depth estimation with a priori knowledge. Passive methods, on the other hand, avoid light projection which might be disturbing in the scene and rely solely on the image information.

Time-of-flight scanners estimate the distance to an object by the time laser light emitted from the scanner requires to reach the object's surface. Depth from focus/defocus techniques reconstruct depth from a set of images with different focal planes. They can be passive as well as active by using camera or projector defocus to reconstruct the depth of a scene.

A large group of approaches reconstructs depth using triangulation. The principle of this type of depth estimation can be illustrated using simple trigonometric rules. Please see Figure (2.1) for notations.

$$d \cdot \sin \gamma = h \quad b \cdot \sin \alpha = h$$

$$\frac{d}{\sin \alpha} = \frac{b}{\sin \gamma} \quad (2.1)$$

As the angles form a triangle the following relation holds

$$\alpha + \beta = \pi - \gamma$$




Figure 2.1: Schematic laser scanner setup. The triangulation is built from a laser beam projected onto a surface point (top) which is observed by a calibrated camera (bottom). For more details on the depth computation please see the text.

$$\sin(\alpha + \beta) = \sin(\pi - \gamma) = -\sin(\gamma). \quad (2.2)$$

Combining equation (2.1) and (2.2) the depth can be calculated as

$$d = \frac{b \cdot \sin \alpha}{\sin \gamma} = -\frac{b \cdot \sin \alpha}{\sin(\alpha + \beta)}. \quad (2.3)$$

In order to determine the unknowns, correspondences between the views need to be established and the observers need to be calibrated to each other. Different variations of this method exist. Well-known techniques comprise laser stripe scanning, stereo vision, and depth from motion. In laser scanning the triangulation is built from a laser emitter, a camera and the surface point to be measured. Stereo vision utilizes two cameras whereas depth from motion applies only one camera which is in motion relative to the scene to create a triangle by recording a surface point at least at two different time instances. Furthermore, other methods for multi-camera systems have been developed.

The emission of light is used by active techniques in order to simplify the detection of correspondences. In structured light scanning a light stripe is projected onto a surface and observed by a camera analog to the laser scanning principle. In order to be able to acquire more geometric information of the object simultaneously, more than one stripe is projected at once. In order to identify the individual stripes a coded pattern is used. Variations of this method with color coded patterns have been proposed to increase the robustness of the system regarding environment illumination and object color. While traditional laser scanning and passive stereo

detect correspondences only in the spatial domain, novel structured light methods make use of features in the spatio-temporal domain which allows for a more stable and robust acquisition over time. Furthermore, capturing a time-varying geometric sequence of non-rigid objects becomes feasible.

Other techniques do not estimate the depth to an object directly but aim to obtain the surface normals of the object first. In a subsequent step the surface itself can be derived as a height map by integration of the normal maps. This group comprises techniques such as photometric stereo and shape from shading. For photometric stereo the object is illuminated from different directions and captured from a constant viewpoint. The normal map can be computed from the reflectance maps of a Lambertian surface as their intensities are determined by the lighting direction and the surface orientation. Shape from shading aims to recover the normal map as well as lighting direction from the variation of shading in one image.

2.2 Registration

As discussed earlier in the previous section the acquisition of a real-world object is mostly conducted in multiple iterations from different directions in order to obtain a complete scan covering the whole object. The relative orientation between the scanner views is typically unknown. In order to align them, a subsequent registration into a common coordinate system is necessary.

A famous technique for registration is the Iterative Closest Point algorithm (ICP). For two raw scans it computes an affine transformation in order to merge them. It finds for all points of one scan the closest point in the second scan. The position of the first scan is then transformed to minimize the root mean squared error of the distances between the point pairs. This process is iterated until the error falls under a user prescribed threshold. The method is used for precise alignment of scans as it requires an initial configuration close to the solution in order to converge. Therefore, user interaction is often required in case the scans are highly disaligned. Other approaches using shape descriptors alleviate this problem by searching for similar feature points on the scans. Once potential corresponding points are detected they are used to coarsely align the set of range images.

2.3 Denoising

Data denoising or cleaning is an important part of the geometry processing pipeline. Since data acquisition is a physical process, it is inevitably imprecise which is observable as noise present in the scanned data. Furthermore, larger errors in depth estimation at grazing angles of the camera view to the surface of the object create the problem that range scan pairs cannot be perfectly aligned. Novel registration techniques aim to handle this problem by performing a non-rigid alignment. Data cleaning can be performed at different steps of the pipeline. Early algorithms focused on filtering already reconstructed surface meshes. In this dissertation it will be demonstrated that denoising is recommendable also early in the geometry processing pipeline as it simplifies and increases the robustness of subsequent processing steps. It furthermore allows for the utilization of the structure given in the scanned data which is typically lost after further processing but beneficial for denoising static but especially time-varying geometric data.

2.4 Surface Reconstruction

At this stage of the geometry processing pipeline range scans have been combined into a point cloud which is often associated with oriented normals. Depending on the acquisition approach they have been already obtained during scanning or they can be easily computed and oriented using the neighborhood information given in the scan and the viewing direction of the camera. Then the goal is to recover the connectivity between the samples such that the resulting mesh is a discrete digital image of the scanned object.

2.5 Remeshing

Despite a lot of advances in surface reconstruction, it is often the case that resulting meshes are not readily suitable for numerical computations required by subsequent geometry processing algorithms. For instance, irregular changes in the normal derivative due to acute triangles complicate the estimation of curvature of the mesh. Therefore, one goal is to transfer the reconstructed surface into a well-behaved mesh representation which consists mostly out of equilateral triangles or quadrangles.

In the following chapters of this dissertation, the contributions will be presented

which cover different parts of the geometry processing pipeline. Chapters 3 and 4 introduce denoising approaches for the robust processing of point cloud and temporally varying geometric data. They are followed by the presentation of a robust and efficient surface reconstruction approach in Chapter 5 and a flexible quad-remeshing technique based on quasi-harmonic fields which allows for different design constraints in Chapter 6.

Chapter 3

Robust Filtering of Noisy Scattered Point Data

In this chapter, we develop a method for the robust filtering of a noisy set of points sampled from a smooth surface. The main idea of the method consists of using a kernel density estimation technique for point clustering. Specifically, we use a mean-shift based clustering procedure. With every point of the input data we associate a local likelihood measure capturing the probability that a 3D point is located on the sampled surface. The likelihood measure takes into account the normal directions estimated at the scattered points. Our filtering procedure suppresses noise of different amplitudes and allows for an easy detection of outliers which are then automatically removed by simple thresholding. The remaining set of maximum likelihood points delivers an accurate point-based approximation of the surface. We also show that while some established meshing techniques often fail to reconstruct the surface from original noisy point scattered data, they work well in conjunction with our filtering method.

3.1 Introduction

Point clouds have become increasingly popular in modeling and rendering applications [Pfister00, Rusinkiewicz00, Alexa01, Zwicker01, Botsch03, Pauly03]




Figure 3.1: Filtering of a face scan acquired using a structured light scanner. Initial scattered point data contains small-amplitude noise and outliers (left image). Our method automatically removes the outliers and effectively suppresses small-amplitude noise (right image).

due to improved graphics hardware and technologies for the acquisition of point geometry. Specifically, robust processing of scattered point data is a subject of intensive research [Stewart99, Mederos03, Pauly04, Fenn05, Steinke05, Jenke06, Suessmuth07, Lipman07b]. Here robustness means that an estimation/ filtering technique works well on noisy data with a small fraction of gross errors (“outliers”). A concise introduction into the field of robust filtering and estimation is available in [Press93]. In this chapter, we develop a technique for robust filtering of noisy sets of points scattered over surfaces and containing outliers. Such point datasets are routinely generated by optical and photometric range finders. The left image of Figure (3.1) shows a typical output of a structured light scanner. Beside usual small-amplitude noise the dataset contains many outliers. The right image presents the result of our filtering technique. The outliers are automatically removed and the small-amplitude noise is effectively suppressed.

While low-pass filtering [Linsen01], MLS fitting [Alexa01, Amenta04, Mederos04, Weyrich04, Dey04b] and PDE-based [Lange05] approaches remove small-amplitude noise well, eliminating outliers remains mostly a manual procedure. A concept for the removal of distant outliers can be found in [Xie04]. Recently, Lipman et al. [Lipman07b] presented an interesting projection-based geometry reconstruction approach which is parameterization-free in a sense that it does not rely on the computation of a local normal or on local plane fitting and thus performs well

in cases of ambiguous orientation. It provides a second order approximation to the given data points. Other works focusing on the accurate reconstruction of sharp features are introduced in [Fleishman05, Jenke06, Lipman07a]. Despite of continuous progress in extending robust statistics, statistical learning techniques and probabilistic approaches for processing scattered point data [Pauly04, Ivriissimtzis04, Schölkopf05, Fenn05, Steinke05, Jenke06, Suessmuth07], the problem of automatic outlier identification and removal from scattered point data remains a challenging task.

Our method can be considered as a non-parametric kernel density estimation scheme [Rosenblatt56, Parzen62]. Given 3D scattered point data $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$, we want to estimate an unknown density function $f(\mathbf{x})$ of the data. A simple kernel estimation $\hat{f}(\mathbf{x})$ of $f(\mathbf{x})$ is given by

$$\hat{f}(\mathbf{x}) = \frac{1}{Nh^3} \sum_{i=1}^N \Phi\left(\frac{\mathbf{x} - \mathbf{p}_i}{h}\right). \quad (3.1)$$

The smoothing parameter h is called the kernel size and Φ is the kernel function which is usually chosen to be a Gaussian function. Figure (3.2) illustrates the kernel-based density estimation approach. Local maxima of the kernel estimation $\hat{f}(\mathbf{x})$ naturally define centers of clusters in the scattered data \mathcal{P} .

The main idea behind our filtering approach consists of defining a kernel estimation $\hat{f}(\mathbf{x})$ to determine those cluster centers which deliver an accurate approximation of the sampled surface. To detect the local maxima of an appropriately constructed kernel estimation \hat{f} the Mean Shift technique [Fukunaga75, Cheng95, Comaniciu02] is used. Then clusters corresponding to the outliers are easily detected by using a simple thresholding scheme. Although the basic concept of our approach is well illustrated by (3.1) and Figure (3.2), the practical implementation of our kernel-based clustering and filtering procedure introduces and deals with much more complex kernels.

Recently, [Suessmuth07] presented a related approach with a different focus being reconstructing curves and surfaces from unstructured point clouds. Similar to our technique it is based on the observation that the curve or surface is located where the point cloud locally has highest density. More precisely, they detect relevant ridges of the density function. In practice, they evaluate first and second order differential geometric quantities on a regular grid surrounding the point cloud, perform a local analysis on the grid points and recover the curve or surface using an isoline extraction algorithm.

In Section (3.3), we demonstrate that our robust filtering method is useful to pre-process noisy data before it is used as input for surface reconstruction techniques. We show that this largely improves their quality of reconstruction.




Figure 3.2: Illustration of the kernel density estimation technique for 1D scattered point data. Local maxima of the kernel estimation \hat{f} define cluster centers of the original data.

3.2 Likelihood and Convergence

In this section, we present our statistical method to filter noisy point cloud surface data. We approach this problem by defining a smooth likelihood function L reflecting the probability that a point $\mathbf{x} \in \mathbf{R}^3$ is a point on the surface S sampled by a noisy point cloud \mathcal{P} . In order to filter the noisy samples we use an iterative scheme inspired by Mean Shift [Cheng95, Comaniciu02, Fukunaga75] to move the points along the gradient of L to maximum likelihood positions.

3.2.1 Likelihood Function

In order to define the likelihood function L we accumulate local likelihood functions L_i defined for every sample point $\mathbf{p}_i \in \mathcal{P}$. We measure the likelihood $L_i(\mathbf{x})$ for a certain \mathbf{x} considering the squared distance of \mathbf{x} to the least-squares plane fitted to a spatial neighborhood of \mathbf{p}_i . Being more specific, we determine the fitting

plane by computing the weighted covariance matrix

$$\mathbf{C}_i = \sum_{j=1}^N (\mathbf{p}_j - \mathbf{c}_i)(\mathbf{p}_j - \mathbf{c}_i)^T \chi \left(\frac{\|\mathbf{p}_j - \mathbf{p}_i\|}{h} \right), \quad (3.2)$$

where h is the kernel size, χ is a monotonically decreasing weighting function and \mathbf{c}_i is the weighted average of all samples inside the kernel. Since \mathbf{C}_i is symmetric and positive semi-definite, its eigenvalues λ_i^l , $l = 1, 2, 3$, are real-valued and non-negative: $0 \leq \lambda_i^3 \leq \lambda_i^2 \leq \lambda_i^1$. Furthermore, the corresponding eigenvectors \mathbf{v}_i^l form an orthonormal basis. Thus the covariance matrix (3.2) defines the ellipsoid

$$E_i(\mathbf{x}) = \{\mathbf{x} : (\mathbf{x} - \mathbf{c}_i)^T \mathbf{C}_i^{-1} (\mathbf{x} - \mathbf{c}_i) \leq 1\}, \quad (3.3)$$

where the least-squares fitting plane is spanned by the two main principal axes \mathbf{v}_i^1 and \mathbf{v}_i^2 of E_i and has the normal $\mathbf{v}_i^3 = \mathbf{n}_i$. A 2D example is illustrated in Figure (3.3).

If normals are provided by the scanning device we use them instead of the estimated normals. Using the squared distance of \mathbf{x} to the least-squares plane we measure the likelihood $L_i(\mathbf{x})$ as

$$L_i(\mathbf{x}) = \Phi_i(\mathbf{x} - \mathbf{c}_i) \left[h^2 - [(\mathbf{x} - \mathbf{c}_i) \cdot \mathbf{n}_i]^2 \right]. \quad (3.4)$$

Thus positions \mathbf{x} closer to the least-squares plane will be assigned a higher probability than positions being more distant. Additionally, we assume that the influence of a point \mathbf{p}_i on the likelihood of a position \mathbf{x} diminishes with increasing distance. To consider this behavior we use monotonically decreasing weighting functions Φ_i to reduce the influence of each L_i . In contrast to radial functions in [Pauly04, Ohtake05] we use a trivariate anisotropic Gaussian function Φ_i which is adapted to the shape of the ellipsoid E_i . This has the advantage that the weighting function is also adapted to the point distribution in a spatial neighborhood of \mathbf{p}_i .

To define the likelihood function L modeling the probability that a certain point \mathbf{x} is a point on the sampled surface S , we accumulate the local likelihoods $L_i(\mathbf{x})$ contributed by all points \mathbf{p}_i .

$$L(\mathbf{x}) = \sum_{i=1}^N w_i L_i(\mathbf{x}) \quad (3.5)$$

Note that we can easily incorporate scanning confidence measures $w_i \in [0, 1]$ associated with each point \mathbf{p}_i by scaling the amplitudes of the likelihood functions. If no scanning confidences are provided we use $w_i = 1$. Figure (3.4) shows an example of a slice of the likelihood function L .




Figure 3.3: 2D example of the weighted least-squares fitting plane and ellipsoid kernel computation.

3.2.2 Convergence

After determining the likelihood function $L(\mathbf{x})$ we use it to smooth the point cloud by moving all samples to positions of high probability. This means we move the samples to positions which are most likely locations on the sampled surface. To find the local maxima of $L(\mathbf{x})$ we use a procedure similar to a gradient-ascent maximization. We freeze the weighting functions Φ_j since they change slowly and approximate $\nabla L(\mathbf{x})$ by

$$-2 \sum_{j=1}^N w_j \Phi_j(\mathbf{x} - \mathbf{c}_j) \cdot [(\mathbf{x} - \mathbf{c}_j) \cdot \mathbf{n}_j] \cdot \mathbf{n}_j. \quad (3.6)$$

To allow a fast convergence of the samples to probability maxima we choose the step-size adaptive as

$$\tau = \frac{1}{2 \sum_{j=1}^N w_j \Phi_j(\mathbf{x} - \mathbf{c}_j)}. \quad (3.7)$$

This means that the step size is small near to the probability maximum and increases to the border of each kernel. This provides a fast and stable convergence of all sample points.

Combining equations (3.6) and (3.7) we get the resulting iterative scheme

$$\mathbf{p}_i^0 = \mathbf{p}_i \quad , \quad \mathbf{p}_i^{k+1} = \mathbf{p}_i^k - \mathbf{m}_i^k \quad (3.8)$$

$$\mathbf{m}_i^k = \frac{\sum_{j=1}^N w_j \Phi_j(\mathbf{p}_i^k - \mathbf{c}_j) \cdot [(\mathbf{p}_i^k - \mathbf{c}_j) \cdot \mathbf{n}_j] \cdot \mathbf{n}_j}{\sum_{j=1}^N w_j \Phi_j(\mathbf{p}_i^k - \mathbf{c}_j)}. \quad (3.9)$$




Figure 3.4: A slice of the likelihood function L of the noisy Dragon head model (bottom left) and a zoom of the framed region (bottom right). The function values are represented by colors increasing from deep blue to purple. Note that L is a smooth function.

In order to filter the point cloud \mathcal{P} we apply the iterative scheme individually to every sample. We stop the iterative process if

$$\|\mathbf{p}_i^{k+1} - \mathbf{p}_i^k\| < 10^{-4}h. \quad (3.10)$$

Each sample usually converges in less than 10 iterations.

A feature of our filtering method is the inherent clustering property. As the number of kernels is larger than the number of maxima in the likelihood function L (see Figure 3.2), several sample points converge to the same probability maximum. We cluster those samples and place one representative point at the local maximum of L . See Table (3.1) for details on the point reduction rate.




Figure 3.5: The effect of adaptive kernels. **Left:** The Dragon dataset is smoothed using a fixed kernel size. Large amplitude noise at the right foot of the dragon cannot be filtered due to maxima of L distant to the most likely surface. **Right:** Filtering result of the same dataset using adaptive kernels. Outlying maxima are well damped. Beside very few points, the noisy samples in the rectangular region are filtered completely.

3.2.3 Adaptive Kernel Size

So far we only used a fixed radius h to compute the local neighborhoods for the ellipsoidal weighting function and least-squares fitting plane computation. However, invariant kernels might not be suitable for datasets with varying sampling density. To overcome this problem we use the k -neighborhood of each sample \mathbf{p}_i for the PCA analysis to compute the ellipsoidal kernel E_i . In this manner we not only adapt the kernel shape to the point sample distribution in a neighborhood

Dataset	N	M	kernels	filtering	h
Face	180K	114K	1.38s	18.45s	0.8
Bunny	362K+25K	324K	3.2s	52s	0.001
Sforza (front)	123K	81K	1.33s	21.32s	2
Sforza (side)	143K	95K	1.49s	24.94s	2
Bimba	1.9M	1.2M	16s	80s	1
Dragon head	485K	170K	23.22s	10m 53s	0.0015
Dragon	2.1M	796K	1m 43s	36m 26s	0.0015
Dragon	2.1M	795K	6m 40s	38m 05s	$k = 250$

Table 3.1: Timings for the ellipsoid kernel computation and for filtering the models presented in this chapter. The kernel size h is chosen in the interval of one to ten times the average sampling density of the input data. The character N denotes the number of input samples and M the number of filtered points. The parameter k indicates the number of nearest neighbors used for the adaptive kernel computation. All results were computed on a 2.66 GHz Pentium 4 with 1.5 GB of RAM.

of \mathbf{p}_i but also the kernel size to the spatial sampling density. The motivation behind this choice can be observed in Figure (3.5). If a fixed radius h is used local maxima of L are created distant to the most likely surface in regions of the point cloud with large-amplitude noise. Those maxima also attract points during the iterative filtering process creating a second layer of points around the most likely surface (left). The usage of adaptive kernels leads to larger kernel sizes in these regions due to the lower sampling density of large amplitude noise. Therefore, kernels of both layers intersect which dampens the effect of local maxima. This results in an improved filtering of large scale noise (right).

3.3 Results and Applications

This section shows results of our point cloud data filtering technique. We apply our method to real-world datasets from laser and structured light scanners. Furthermore, we present the application of our algorithm to surface reconstruction. In general, surface reconstruction is performed on noisy data reducing the efficiency of surface reconstruction algorithms. We show that the results of well-known surface reconstruction methods can be improved in conjunction with our filtering method.




Figure 3.6: Left: Raw registered range scans of the Stanford Bunny dataset expanded by 25K random “salt and pepper” outliers. Right: Our method accurately denoises the given point set and removes the dense cloud of outliers properly.

3.3.1 Filtering and Outlier Removal

We demonstrate results of our filtering technique in Figure (3.1) and Figures (3.5) - (3.9). All images are rendered using shaded points. Normals for illustrating the results are computed using PCA analysis with small neighborhoods to avoid blurring effects. Meshes in Figure (3.9) are displayed using flat shading in order to show faceting. Table (3.1) summarizes our results and the parameters used to generate them.

Figure (3.1) shows a point cloud face dataset acquired by a structured light scanner before and after filtering using our method. The raw point cloud suffers from several outliers and ridges which are typical artifacts caused by the structured light scanning process. We show this comparison to illustrate the effectiveness of our method for removing outliers and smoothing of difficult datasets. Due to the clustering property of our method, groups of outliers usually converge to a set of single points sparsely distributed around the surface samples. These points can be characterized by a very low spatial sampling density compared to the surface samples. We use this criterion for the detection of outliers and remove them using simple thresholding. Figure (3.6) shows an additional example with a large amount of randomly generated points which can be interpreted as 3D “salt and pepper” outliers. In the case of images, “salt and pepper” noise corrupts random



Figure 3.7: Smoothing of two range scans of the Ippolita Sforza Bust. Note that details in hair, mouth and eye regions are accurately preserved.

image pixels with intensity spikes. This means that a number of pixels in the image have a very large intensity difference to neighboring pixels. For point clouds, we can model this kind of noise by displacing points of the dataset far from the smooth surface. In our example, we move points inside the bounding box of the dataset. Additionally, we add noise to the normals by perturbing them with random angles. Although the outlier density is high as shown in Figure (3.6), our algorithm is able to remove the noise and the outliers properly.

In Figure (3.7) we demonstrate the filtering efficiency of our algorithm on laser scan data. We show this comparison as laser scans are usually affected by different types of noise compared to structured light scans. Due to the different acquisition technique, laser scans are usually not corrupted by ridges and pits caused by structured light. Instead, they are affected by dense small-amplitude noise. Figure (3.7)



Figure 3.8: Left: Raw registered range data of the Bimba model obtained using a laser scanner. The data is corrupted by dense small-amplitude noise and scanning artifacts close to the mouth and the right eye of the model. Right: The artifacts are well dampened and noise is removed after filtering with our method.

illustrates that high-frequency noise is removed by our method while lower frequency details like hair, mouth and eyes of the Ippolita Sforza Bust are preserved. Figure (3.8) illustrates a second example consisting of a set of raw registered laser scans. The Bimba model suffers from the same small-amplitude noise and a few scanning artifacts close to the right eye and mouth region of the dataset. After denoising the artifacts are well dampened and the noise is properly removed while important details are preserved.

As noted previously, our method uses adaptive kernels to handle large scale noise. Figure (3.5) shows that while the dragon scan cannot be accurately filtered using a fixed kernel size, adaptive kernels provide a proper filtering of large amplitude noise.

3.3.2 Surface Reconstruction

Surface reconstruction is one of the most fundamental problems in geometry processing. One important group of reconstruction algorithms are Delaunay-based

methods. Those techniques are supported by rigorous mathematical results and provide correct reconstructions under specific sampling conditions. Furthermore, they have the great advantage to be able to reconstruct surfaces from points without normals. On the other hand, these methods are sensitive to data with noise and outliers which cannot be avoided in physical acquisition processes. Most Delaunay-based methods are therefore not practical to be applied to raw data. Therefore, making these methods robust to noisy data is currently a field of intensive research.

Our method can be used to filter real-world data before using it for surface reconstruction by computational geometry methods. We apply our filtering technique to noisy point clouds and reconstruct a surface from the preprocessed data. For comparison we also reconstruct a surface from the same noisy point cloud without cleaning it using our method. For surface reconstruction we use two well-known Delaunay-based reconstruction algorithms, namely Power Crust [Amenta01] and Tight Cocone [Dey03], which are available for scientific purposes. Figure (3.9) shows results of the comparison. The direct reconstruction of the noisy scattered data does not produce usable results. In contrast, results of both algorithms using the filtered scattered data show significantly improved reconstructions.

3.4 Conclusion

In this chapter, a kernel based clustering approach for the robust filtering of point cloud surface data has been introduced. It adapts to the density and distribution of the given input points which allows the technique to remove noise of different amplitudes accurately. Due to its inherent clustering property it permits an easy detection of outliers. It was demonstrated that the algorithm is able to handle even a large amount of “salt and pepper” outliers. The effectiveness of the approach was shown on real-world datasets acquired using structured light and laser scanners. In addition, the method can be used in combination with surface reconstruction algorithms which significantly improves their results on noisy data. Furthermore, a parallelization of the presented approach is imaginable as every input point converges independently to a maximum of the likelihood function.

While the presented method operates on already registered scans and does not require neighborhood information, the technique introduced in the next chapter follows a different approach. It denoises range scans separately before they are registered and is thus able to exploit the connectivity information between the data points. The next chapter furthermore illustrates that the method extends to the treatment of time-varying range data whose acquisition has recently become an active research direction.

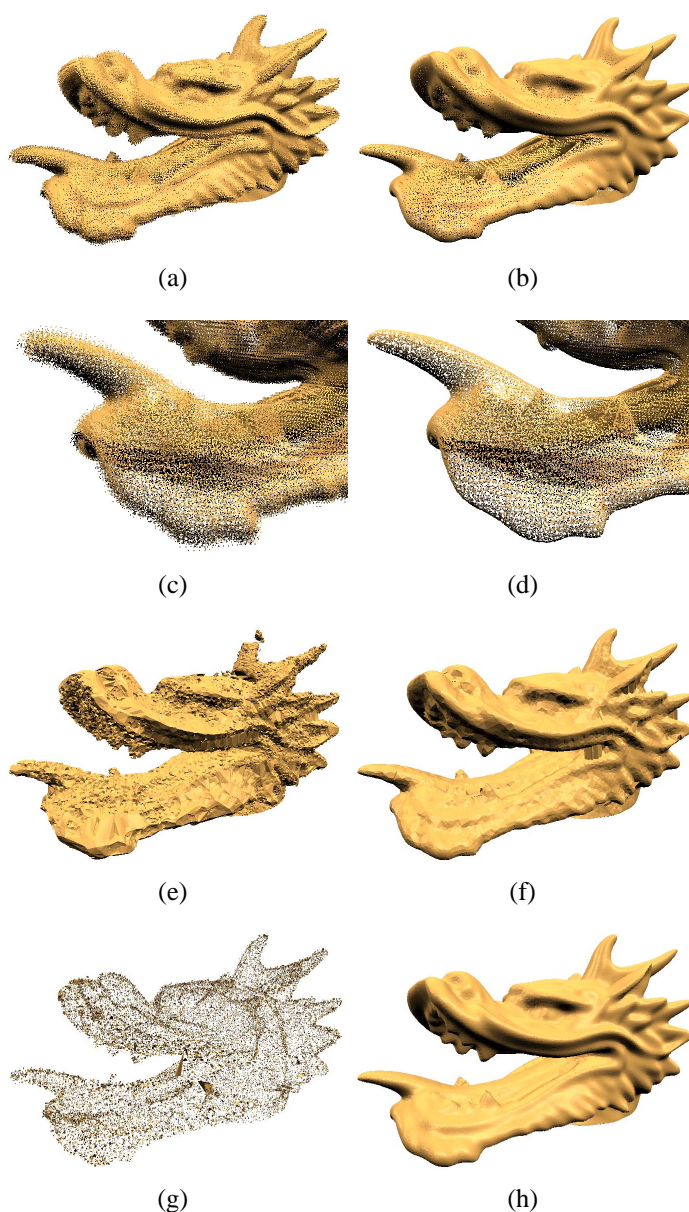


Figure 3.9: Figures (a) and (b) present the head of the Dragon scans from the Stanford Scanning Repository before and after our filtering procedure. Figures (c) and (d) show zooms of the images (a) and (b) close to the tongue region. Notice that noise is removed and that the filtered samples indicate a surface. Figures (e) and (g) illustrate Power Crust and Tight Cocone reconstructions from the noisy samples shown in (a). Figures (f) and (h) show reconstruction results from the filtered data shown in (b). While the Power Crust algorithm shows noticeably improved results with small defects (f), the Tight Cocone algorithm reconstructs a smooth mesh (h).

Chapter 4

Feature-preserving Non-local Denoising of Dynamic Range Data

We present a new method for noise removal on static and time-varying range data. Our approach predicts the restored position of a perturbed vertex using similar vertices in its neighborhood. It defines the required similarity measure in a new non-local fashion which compares regions of the surface instead of point pairs. This allows our algorithm to obtain a more accurate denoising result than previous state-of-the-art approaches and, at the same time, to better preserve fine features of the surface. Furthermore, our approach is easy to implement, effective, and flexibly applicable to different types of scanned data. We demonstrate this on several static and interesting new time-varying datasets obtained using laser and structured light scanners.

4.1 Introduction

With the increased usage of scanning devices, denoising of digitized models has become one of the most fundamental problems in computer graphics. It remains a challenging task to remove the inevitable noise created in every acquisition




Figure 4.1: A raw laser scan of the Bimba model (top row), the denoised result obtained using our algorithm (middle row), and a closer view of the noisy and denoised ear of the model (bottom row). For all images, the corresponding mean curvature visualizations are shown. Notice that high-frequency noise is properly removed after only one iteration of our method, while fine details in hair, ear and eye regions are accurately preserved.

process while preserving the details of the underlying image or shape. Especially, fine features are often lost if no special treatment is provided. The continuous progress achieved by a variety of denoising approaches which have been introduced in recent years in the fields of image processing and computer vision [Perona90, Rudin92, Tomasi98] as well as in computer graphics [Taubin95, Desbrun99, Tasdizen02, Fleishman03, Hildebrandt04, Lipman07a], clarifies that the development of reliable, accurate and versatile denoising techniques is a lively area of research and the foundation for a wide range of applications.

In this chapter, we introduce a new method for range scan denoising called *non-local* neighborhood filtering inspired by a image processing technique [Buades05] which presents remarkable results. The main idea of neighborhood filtering in general is to determine the denoised position of a vertex as a weighted average of similar vertices in its vicinity. In particular, the choice of the similarity measure has a strong influence on the efficiency of the denoising approach. Unlike previous neighborhood filters which determine the similarity of two points locally using only their positions and sometimes normals, our approach is *non-local* and defines the similarity by comparing regions of the surface around the vertices instead. This yields a more accurate denoising of the surface and improves the removal of higher noise levels compared to previous state-of-the-art filtering approaches. At the same time, fine shape features are better preserved. Another extension of our method is that the size of the regions which are used to determine the similarity measure and the restored point position are not constant over the whole dataset but adapted close to the boundaries which allows it to obtain a more accurate denoising result in these regions. In addition, our method is easy to implement and can be flexibly applied to different types of noisy data obtained using laser and structured light scanners. An example for the effectiveness of our method is presented in Figure (4.1).

Our method addresses the denoising problem differently compared to most previous approaches since we denoise range scans before they are combined within the scanning pipeline. This is more efficient since the given structure of the data can be utilized in a simpler similarity measure which allows for a faster evaluation. Furthermore, our approach extends more naturally to dynamic range data whose acquisition has become feasible thanks to interesting improvements in scanning technology [Zhang04, Davis05]. It is the first approach which is designed to filter time-varying geometric data, and we believe that proper denoising is the first step to open this new type of data up to various kinds of applications.

4.2 Related Work

Denoising is a field of intensive research in image processing, computer vision, and computer graphics which comprises different classes of algorithms. The main approaches can be categorized into the following groups.

4.2.1 PDE Approaches

Early works in image processing which introduced PDEs for denoising are anisotropic diffusion [Perona90] and total variation minimization [Rudin92]. Desbrun et al. [Desbrun99] proposed a geometric diffusion flow algorithm for irregular meshes and introduced the use of an implicit integration method which stabilizes the flow to allow larger time steps. Later, a new variant of anisotropic mean curvature flow which preserves non-linear features was introduced [Hildebrandt04]. This approach was extended in [Lange05] in order to denoise point clouds. Hildebrandt and Polthier [Hildebrandt07] presented a constrained-based fairing approach where the final smoothed surface remains in a user-prescribed distance to the input mesh.

4.2.2 Spectral Techniques

Taubin [Taubin95] first introduced signal processing on meshes based on the definition of the Laplacian operator on surfaces. In [Alexa02], Wiener filtering was applied to meshes. Pauly and Gross [Pauly01] created a spectral decomposition of a point cloud and denoised it by manipulating the spectral coefficients.

4.2.3 Neighborhood Filtering

Neighborhood filters were early addressed in [Yaroslavsky85]. In a later work, Tomasi and Manduchi [Tomasi98] introduced the well known bilateral filter for images which was extended to non-local neighborhoods in [Buades05]. Paris and Durand [Paris06] proposed an interesting acceleration technique for the bilateral filter which allowed it to be mostly expressed as simple linear convolutions. Weiss [Weiss06] introduced a very fast vectorizable algorithm for median and bilateral filtering. Other works adapted neighborhood filters for video enhancement [Bennett05] and processing [Mahmoudi05]. Further interesting research areas used the information contained in local neighborhoods for other tasks, such

as texture synthesis [Efros99, Wei00] and mesh completion [Sharf04]. Fleishman et al. [Fleishman03] proposed an anisotropic mesh denoising algorithm derived from the bilateral filter for images. Concurrently, Jones et al. [Jones03] introduced a similar method based on robust statistics which uses local first-order predictors of a surface. Yoshizawa et al. [Yoshizawa06] extended the non-local image filter to meshes by computing a local RBF approximation to define the similarity measure. Our work approaches the problem in a different way by introducing a new similarity measure for range scans, which avoids the computation of a local approximation and thus allows a much faster evaluation. Additionally, our approach considers the denoising of time-varying geometric data.

4.2.4 Projection-based Approaches

Algorithms that attracted the interest of many researchers are moving-least squares (MLS) approaches. They were first proposed by Levin [Levin98] and introduced to computer graphics by Alexa et al. [Alexa01]. The main idea of MLS is the definition of a projection operator which takes points scattered in the vicinity of a surface onto the surface itself. More precisely, the MLS surface is defined by the fixpoints of the given projection operator. Mederos et al. [Mederos03] applied the MLS projection for point cloud denoising. Amenta and Kil [Amenta04] analyzed different MLS operators by separating them into two components being a vector field and an energy function. They used this representation to introduce a new variant of MLS with a better behavior near sharp features. In a recent work, Fleishman et al. [Fleishman05] chose a different approach and represented sharp features by defining piecewise smooth moving least-squares surfaces using robust statistics. Dey and Sun [Dey05] proposed the AMLS operator which provides reconstruction guarantees for the underlying surface of a point set with a non-uniform sampling density. Lipman et al. [Lipman07a] introduced a data-dependent variant of MLS which does not use a fixed approximation space as for instance polynomials of a certain degree. For each projected point, it finds a proper local approximation space of piecewise polynomials. The locally constructed spline accounts for singularities which may be present in the data.

4.2.5 Statistical Techniques

Pauly et al. [Pauly04] introduced a framework for measuring uncertainty in point-sampled geometry based on statistical data analysis which can be used, for instance, to merge range scans. Schall et al. [Schall05a] defined a global probability

distribution function for a given noisy point set using locally defined kernels. Positions on a smooth surface were then found by moving every sample of the noisy dataset to maximum likelihood positions. Jenke et al. [Jenke06] proposed how to produce a smooth point cloud from a given noisy one using Bayesian statistics. For image restoration, Awate and Whitaker [Awate06] obtained remarkable results with an unsupervised, information-theoretic, adaptive filter which improved the predictability of a pixel by reducing the joint entropy between its neighborhood and other pixel neighborhoods.

4.2.6 Dynamic Data Acquisition and Processing

The development of new scanning technology with the goal of acquiring dynamic geometric data opened a new interesting research direction. Zhang et al. [Zhang03] and Davis et al. [Davis03] concurrently developed novel spacetime stereo algorithms. While Zhang et al. focused on capturing dynamic scenes, the main interest of Davis et al. was to develop a unifying framework for existing more specialized depth from triangulation techniques. Building on this work, Zhang et al. [Zhang04] presented an approach for the acquisition of facial geometry. After obtaining the geometric data using a spacetime stereo algorithm they fit a template mesh to the data in order to avoid denoising and to create a face sequence with vertex correspondences. The resulting sequence can be manipulated using a data-driven inverse kinematics technique. Bickel et al. [Bickel07] presented a different approach which does not use a generic face template but a high-resolution scan of the face for which motion data is obtained using a traditional marker-based motion capture system. The system is augmented by two synchronized video cameras to track facial expressions. In order to transfer the motion to the static face scan a shell-based mesh deformation method is used. Wand et al. [Wand07] recently proposed a method for the reconstruction of deforming geometry in the form of a sequence of noisy point clouds. The method is based on a statistical framework and reconstructs the geometry as well as dense correspondences over time.

4.3 Non-local Denoising

We begin this section by describing the idea of non-local filtering for images in more detail before we introduce our extension of this approach for denoising static range data. Building on this extension, we then show how to apply our algorithm to filter time-varying range data.

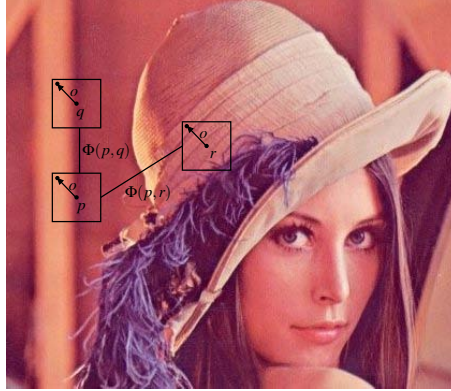


Figure 4.2: The similarity of neighborhoods is computed based on the pixel-wise difference of intensity values. Similar neighborhoods of p and q have a large weight $\Phi(p,q)$, while different neighborhoods of p and r have a small weight $\Phi(p,r)$.

4.3.1 Non-local Image Filtering

The non-local image filter [Buades05] belongs to the group of neighborhood filtering schemes, which define the intensity value of a restored pixel of an image as the weighted average of neighboring pixels with similar intensity values.

More precisely, if an image $\mathcal{I} = \{I(\mathbf{u}) | \mathbf{u} \in P\}$ is given, where $\mathbf{u} = (x,y)$ is a pixel and $I(\mathbf{u})$ is the intensity value at \mathbf{u} , the smoothed pixel intensity $I'(\mathbf{u})$ can be computed as the average of all pixel intensities in the image

$$I'(\mathbf{u}) = \frac{\sum_{\mathbf{v} \in P} \Phi(\mathbf{u}, \mathbf{v}) I(\mathbf{v})}{\sum_{\mathbf{v} \in P} \Phi(\mathbf{u}, \mathbf{v})} \quad (4.1)$$

weighted by a similarity factor which measures the similarity between \mathbf{u} and \mathbf{v} as

$$\Phi(\mathbf{u}, \mathbf{v}) = \exp \left(- \frac{\sum_{\mathbf{o}} G_a(\|\mathbf{o}\|) |I(\mathbf{u} + \mathbf{o}) - I(\mathbf{v} + \mathbf{o})|^2}{h^2} \right). \quad (4.2)$$

Figure (4.2) illustrates the computation of the region-based similarity measure. It depends on the pixel-wise intensity difference of two square neighborhoods centered at the pixels \mathbf{u} and \mathbf{v} . The vector \mathbf{o} denotes the offset between the center pixel and an arbitrary neighborhood pixel. The influence of a pixel pair on the similarity falls with increasing Euclidean distance to the center of the neighborhoods. For the distance weighting a Gaussian kernel $G_a(\cdot)$ with a user-defined

standard deviation a is used. Additionally, the method depends on the parameter h which controls the degree of smoothing.

4.3.2 Static Range Data

We want to adapt this approach from the 2-dimensional plane to range data. This is not a straight-forward task due to two fundamental differences between intensity images and range scans. Firstly, image pixels are usually aligned on a regular and equispaced grid which is in general not true for range images. Secondly, removing noise from a range scan is more complex than from intensity images since noise on a range scan is not necessarily additive to the surface but can, for instance, be dependent on the view of the scanner camera. In the case of images, noise is usually only additive to the intensity values. Therefore, we have to consider different displacement directions for vertices from range images which makes range scan denoising different and more difficult than image denoising.

We assume that the data is given in the form of points \mathbf{p}_i which are arranged on a grid structure. In this way, the neighborhood information for all points is known, but they are not required to be equispaced. Since this data representation can be easily computed from the output of different scanning devices, our algorithm is easily applicable to filter different types of range data.

We find counterparts for the intensity values of an image by the distance of points \mathbf{p}_j from a defined plane computed at a given point \mathbf{p}_i . Thereby, the orientation of the plane is orthogonal to the displacement direction \mathbf{d}_i of a vertex \mathbf{p}_i . By determining the weighted average of the offsets, we find the displacement for \mathbf{p}_i , to remove the noise component from the range scan surface. In order to adapt to different types of noise such as view-dependent or additive noise, we choose different displacement directions \mathbf{d}_i . In the case of additive noise, we first estimate normals \mathbf{n}_i for all points \mathbf{p}_i by least-squares fitting to their local neighborhoods and choose $\mathbf{d}_i = \mathbf{n}_i$. For view-dependent noise, we select \mathbf{d}_i to be the line-of-sight of the camera of the scanning device. In general, we determine the filtered points \mathbf{p}'_i by computing

$$\mathbf{p}'_i = \mathbf{p}_i - \frac{\sum_{\mathbf{p}_j \in \mathcal{N}(\mathbf{p}_i)} \Phi_d \Phi_s [(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{d}_i] \mathbf{d}_i}{\sum_{\mathbf{p}_j \in \mathcal{N}(\mathbf{p}_i)} \Phi_d \Phi_s} \quad (4.3)$$

where Φ_d represents the distance and Φ_s the similarity weight.

Unlike the non-local image filtering algorithm, we do not sum over all point positions to filter a point but over a local square neighborhood $\mathcal{N}(\mathbf{p}_i)$ surrounding \mathbf{p}_i .

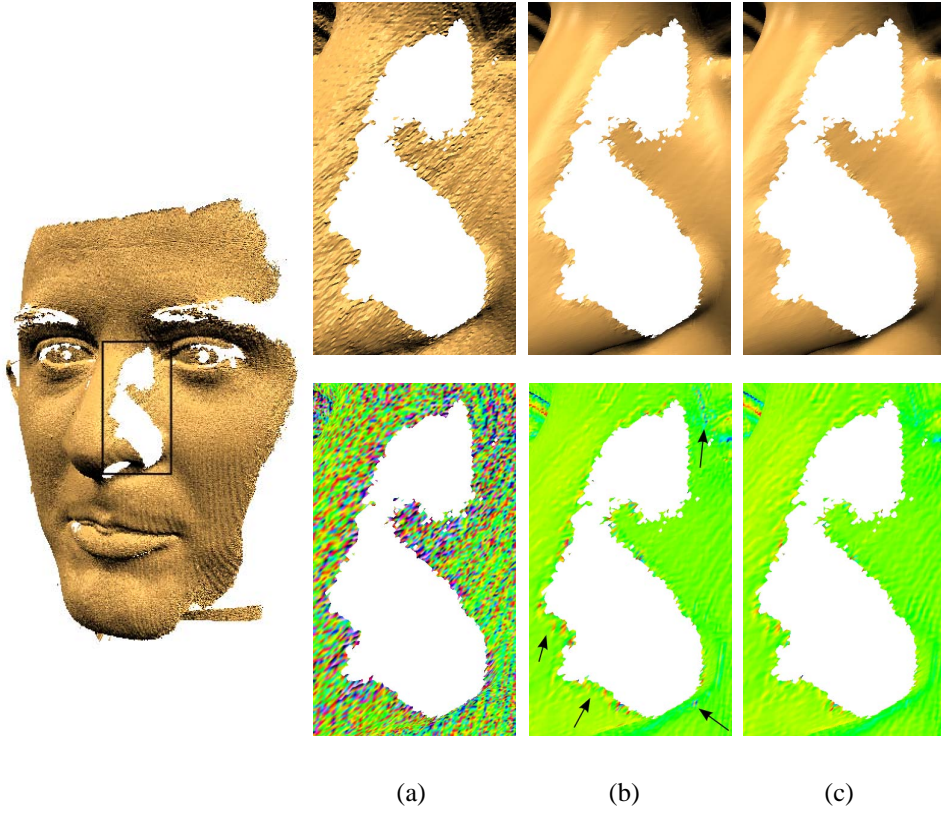


Figure 4.3: A face scan obtained using a structured light scanner (a) is denoised using uniform neighborhood sizes (for $\mathcal{N}(\cdot)$ and Sim) (b) and using adaptive neighborhoods (c). Figure (c) shows that the denoising result close to the boundary is improved and that artifacts are removed.

Additionally, we separate the distance weighting factor from the similarity measure. This allows us a more efficient computation of the similarity weight and the denoised point position \mathbf{p}'_i .

The fundamental difference of our method from previous neighborhood filtering approaches for meshes is the selection of the similarity weight Φ_s . Unlike the bilateral filtering algorithm, where Φ_s only weights the similarity between the two points \mathbf{p}_i and \mathbf{p}_j , our approach considers the similarity of their geometric neighborhoods:

$$\Phi_d(\mathbf{p}_i, \mathbf{p}_j) = e^{-\frac{\|\mathbf{p}_i - \mathbf{p}_j\|^2}{d^2}} \quad \Phi_s(\mathbf{p}_i, \mathbf{p}_j) = e^{-\frac{Sim(\mathbf{p}_i, \mathbf{p}_j)^2}{s^2}} \quad (4.4)$$

$$Sim(\mathbf{p}_i, \mathbf{p}_j) = \frac{\sum_{o \in O} |(\mathbf{p}_{i+o} - \mathbf{p}_{j+o}) \cdot \mathbf{d}_i|^2}{\|O\|}. \quad (4.5)$$

This results in a more accurate filtering performance. We determine the point-wise difference of two square neighborhoods centered at \mathbf{p}_i and \mathbf{p}_j and project the distances onto the displacement direction \mathbf{d}_i . In the process, invalid points in the neighborhoods $\mathcal{N}(\mathbf{p}_i)$ and $\mathcal{N}(\mathbf{p}_j)$ are ignored. The result is the point-wise height difference of both neighborhoods which is averaged by the number of points inside the neighborhood (equal to the number of offsets $\|O\|$) to compute the non-local similarity $Sim(\mathbf{p}_i, \mathbf{p}_j)$. Our similarity measure automatically compensates for translational movements of the neighborhoods along the defined plane. We also conducted experiments with a similarity measure that additionally compensates for rotations by aligning the directions \mathbf{d}_i and \mathbf{d}_j with each other before the similarity computation but it was more expensive to compute and did not improve the results significantly. We use Gaussian weighting functions for Φ_d and Φ_s and an automatic procedure to determine their bandwidths d and s . For this, we first choose random points \mathbf{p}_k of the range scan. We then determine the maximal distance of the points of $\mathcal{N}(\mathbf{p}_k)$ to \mathbf{p}_k and the standard deviation of all offsets to the plane defined at \mathbf{p}_k . The average maximal distance and standard deviation over all random samples are then assigned μ_d and σ_s . We set $d = 0.75\mu_d$ and $s = \sigma_s$.

The user-defined parameters of our algorithm are thus only the size of the neighborhood $\mathcal{N}(\mathbf{p}_i)$ which controls the degree of smoothing and the size of the neighborhood used to determine $Sim(\mathbf{p}_i, \mathbf{p}_j)$ which regulates the homogeneity of the filtering result.

The size of the neighborhoods is usually chosen to be uniform over the dataset. According to our experiments the results of our approach can be improved if adaptive neighborhoods are used. If the neighborhood square $\mathcal{N}(\mathbf{p})$ of a point \mathbf{p} contains insufficient information due to too many invalid points which is likely to happen mostly at the boundary of the dataset, our experiments showed that it is better to reduce the size of $\mathcal{N}(\mathbf{p})$ stepwise in order to optimize the ratio of its valid to its complete number of points. The effect of this enhancement is shown in Figure (4.3). It is important to mention that the adaptation of the neighborhood size is optional and that our method does not require special boundary treatment which is important as scanned data is not closed and often has holes.




Figure 4.4: Denoising results for two acquired noisy range sequences. The raw input from the structured light scanner (top row) is denoised using bilateral filtering (middle row) and our technique (bottom row). Coloring by mean curvature is used to illustrate the smoothness of the range data. Note that our algorithm is able to remove noise more properly while accurately preserving surface features.

4.3.3 Time-varying Range Data

Building on the previous section, we extend our approach to process time-varying range data. The data is given as a sequence of frames each of which is a static range scan. When we apply our algorithm to each frame independently, we obtain a result that is satisfying for each frame but which is not temporally stable.

Therefore, we extend $\mathcal{N}(\mathbf{p}_i)$ which is usually only defined as a spatial neighborhood by the temporal domain. This means that we choose sample points for $\mathcal{N}(\mathbf{p}_i)$ not only inside the current frame but also in neighboring frames. We usually consider one frame before and after the current frame for $\mathcal{N}(\mathbf{p}_i)$. In the following, we use the notation $\mathcal{N}_k(\mathbf{p}_i)$ for the slice of the neighborhood $\mathcal{N}(\mathbf{p}_i)$ which is contributed by the frame k . Consequently, we have to adapt Φ_d and Φ_s to weight the distance and the similarity between \mathbf{p}_i and \mathbf{p}_j which can be points in different frames. We adapt the bandwidths of Φ_d and Φ_s depending on the frame with which \mathbf{p}_j is associated. We detect the parameters automatically as described in Section (4.3.2) for each frame k and identify the weighting functions as Φ_{dk} and Φ_{sk} .

Similar to the spatial domain, we want that neighborhoods from distant frames contribute less to the new point position. We therefore introduce the temporal distance factor Ψ_{dk} which weights the contribution of the frame k . If c is the index of the current frame, we select $\Psi_{dk} = (1/2)^{|k-c|}$. Additionally, we can weight a frame based on its noise level. Neighborhoods from frames with a higher amount of noise can contribute less to a smooth solution and should thus have a lower weight. We obtain an estimate for the noise-level from the bandwidth s_k of Φ_{sk} for each frame k . We use these values to set the weighting factor $\Psi_{sk} = \exp(-s_k^2 / \max_k \{s_k\}^2)$. By combining all elements, we determine the denoised point position as

$$\mathbf{p}'_i = \mathbf{p}_i - \frac{\sum_k \Psi_{dk} \Psi_{sk} \sum_{\mathbf{p}_j \in \mathcal{N}_k(\mathbf{p}_i)} \Phi_{dk} \Phi_{sk} [(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{d}_i] \mathbf{d}_i}{\sum_k \Psi_{dk} \Psi_{sk} \sum_{\mathbf{p}_j \in \mathcal{N}_k(\mathbf{p}_i)} \Phi_{dk} \Phi_{sk}}. \quad (4.6)$$

One advantage of our approach is that we do not necessarily need to compensate for motion between frames as the similarity of the whole temporal neighborhood is evaluated. If the motion is high, the similarity of the whole neighborhood will be low and it will only marginally contribute to the new point position. In this way, our approach also automatically accounts for scene changes.

4.4 Results

We demonstrate results of our denoising approach in Figure (4.1) and Figures (4.4) - (4.6). We test our method on scanned data from various sources. It is applied to laser scanned models Figures (4.1) + (4.5) as well as to a face and hand sequence Figure (4.4) and to the Bust model Figure (4.6) which were acquired using two different structured light scanners. We compare our result to the bilateral filtering algorithm. Table (4.1) summarizes the timings for our results and the parameter settings used to generate them.

In Figure (4.1), we show the filtering efficiency of our approach on real-world laser scanned data. The images show that high-frequency noise on the Bimba model is removed after only one iteration of our algorithm while lower-frequency details such as hair, ear and eye are accurately preserved.

Figure (4.4) illustrates results of the bilateral filter and our algorithm on three frames of two acquired structured light sequences. To filter the scans, we apply two different kernel settings for each algorithm. First, we filter with a larger kernel size to remove the stripe artifacts created due to the projection of regular line patterns onto the scanned object during the acquisition process. As the stripe pattern varies over time, we let both methods filter across frames to increase the temporal stability of the smoothed sequence. We consider one frame before and after the current frame while filtering both sequences. High-frequency noise distributed over the whole scan does not show any temporal coherence. Therefore, we subsequently filter every frame separately with a smaller kernel size. The parameters are chosen such that each algorithm has its optimal performance. Our experiments show that using a larger kernel size for the bilateral filter than chosen in this comparison does not yield a smoother result. Instead, we allow the bilateral filtering algorithm to iterate twice over the sequences with both kernel sizes to make the runtimes of the bilateral filter and our algorithm match (see Table 4.1). Figure (4.4) shows that our method removes the stripe artifacts and high-frequency noise properly and achieves a more accurate result than the bilateral filter due to our region-based definition of the similarity measure which adds more geometric information into the filtering process. Furthermore, our algorithm preserves high-curvature regions accurately, for instance, the eyes of the face scan and the wrinkles of the hand model.

Figure (4.5) shows a comparison of the bilateral filter and our approach concerning feature preservation. Note that our algorithm creates a smoother result of the Turbine Blade model than the bilateral filter while preserving sharp features more accurately.




Figure 4.5: Comparison of feature-preservation properties of bilateral filtering and our approach on a laser range scan of the Turbine Blade model. The zoomed mean curvature visualizations show that our approach preserves sharp features more accurately than bilateral filtering while simultaneously producing a smoother result.

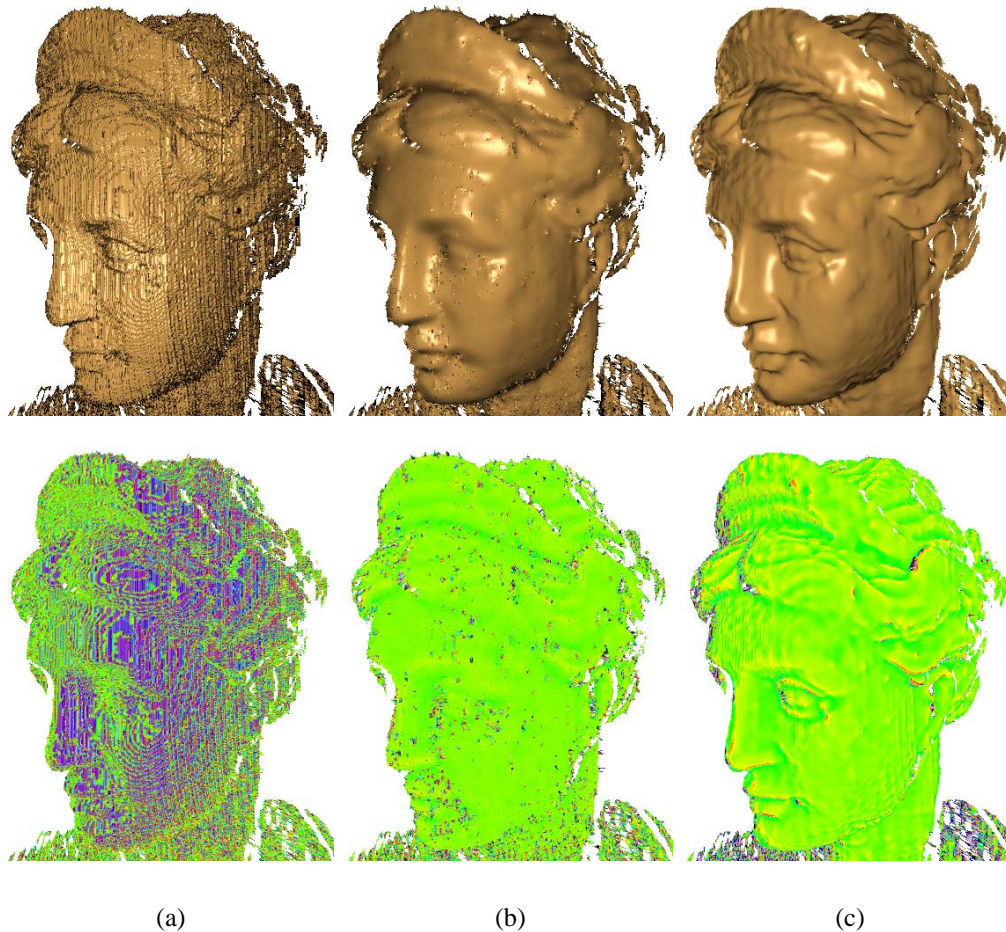


Figure 4.6: The raw structured light scan of a Bust model (a) is corrupted by view-dependent noise. Simple denoising using an additive noise model, which is usually assumed or required by previous approaches, does not deliver satisfactory results (b). Due to the generality of our method, it is able to utilize additional information as, for instance, the camera viewing direction. This allows it to obtain a more accurate and detailed denoising result (c). For the computation of the results in (b)+(c) the same number of iterations and parameters are used for our method (see Table 4.1).

bilateral filtering				
model	P	Sim	$\mathcal{N}(\cdot)$	time/frame
Blade	59K	–	11x11	2.8s
Face Sequence (2 iterations)	192K	–	23x23	94s
		–	11x11	16s
Hand Sequence (2 iterations)	131K	–	23x23	68s
		–	11x11	13s
our approach				
model	P	Sim	$\mathcal{N}(\cdot)$	time/frame
Blade	59K	5x5	7x7	5.6s
Bimba	212K	5x5	7x7	7.4s
Bust (2 iterations)	358K	9x9	13x13	240s
		5x5	11x11	48s
Face Sequence (1 iteration)	192K	9x9	13x13	95s
		5x5	11x11	18s
Hand Sequence (1 iteration)	131K	9x9	13x13	69s
		5x5	11x11	13s

Table 4.1: Parameter settings and timings for the results presented in this chapter. The parameter P labels the average number of input points per frame. Sim denotes the size of the neighborhood considered to compute the similarity measure of our algorithm. All results were computed on a 2.66 GHz Pentium 4.

In Figure (4.6), we illustrate the adaptability of our method to different noise models. This is an important property since many previous approaches are bound to the case of additive noise where it is assumed that vertices are mostly displaced into the normal direction of the underlying smooth surface of the range scan. This yields a bad denoising performance if the noise present in the data differs largely from the imposed model. Figure (4.6 a) shows an example with an unprocessed structured light scan which is corrupted by view-dependent noise. By simply denoising the dataset with our method using the additive noise model, it is clearly noticeable that surface features are not well preserved and that smoothing artifacts are introduced (see Figure 4.6 b). In Figure (4.6 c), we use the identical parameter settings for our denoising algorithm but use the viewing direction of the scanner, which is usually available from the scanner calibration data, for the displacement directions \mathbf{d}_i . This yields a significantly more accurate denoising result and a better preservation of high-curvature regions.

4.5 Conclusion

In this chapter, we presented a new similarity-based neighborhood filtering technique for static and dynamic range data which is the standard output of scanning devices and, in particular, of recently developed 3D video cameras. We introduced a new non-local similarity measure which determines the resemblance of two points on the surface not only by utilizing their local properties like position or normal but by also comparing the region of the surface surrounding the vertices instead. We demonstrated on several different types of scanned data that the idea of adding context information to the similarity definition allows our method to produce a more accurate denoising result than previous state-of-the-art approaches while having a better feature preservation. Additionally, we showed that the usage of adaptive neighborhoods improves the denoising result in the vicinity of boundaries of the given input. Furthermore, our method is easy to implement and flexibly adaptable to scans with different noise properties. It thus delivers a practical, versatile and powerful tool for filtering range data. In this way, our approach naturally fits into the scanning pipeline by denoising range scans before they are combined for further processing, which is more efficient since the given structure of the data can be utilized. Furthermore, we showed an interesting extension of our approach for filtering time-varying geometric data which is important since we expect a wider use of 3D video cameras in the future. It exploits the temporal coherence of the sequence in order to guarantee smoothness along the time domain.

The last two chapters introduced two novel denoising algorithms for static and dynamic geometric data. While both methods approach the denoising problem in different ways by either denoising the given geometric data before or after the registration step, the output of both methods is a clean point cloud with associated oriented normal information. The next chapter, introduces an adaptive surface reconstruction method which takes this type of data as input and is based on the Fast Fourier Transform (FFT) as well as on partition of unity composition.

Chapter 5

Error-guided Adaptive Fourier-based Surface Reconstruction

In this chapter, we propose to combine a FFT-based approach to surface reconstruction from oriented points with adaptive subdivision and partition of unity composition. This removes the main drawback of the FFT-based approach which is a high memory consumption for geometrically complex datasets. This allows us to achieve a higher reconstruction accuracy compared to the original global approach. Furthermore, our reconstruction process is guided by a global error control determined by the Hausdorff distance of selected input samples to intermediate reconstructions. The advantages of our surface reconstruction method also include a more robust surface restoration in regions where the surface folds back to itself.

5.1 Introduction

Many of today's applications make use of 3D models reconstructed from digitized real-world objects such as machine parts, terrain data, and cultural heritage. In spite of progress in developing speedy and reliable methods for surface reconstruction from scattered data [Dey05, Fleishman05, Kazhdan05, Mederos05,

Nehab05, Ohtake05] (see also references therein), the quest for fast, accurate, and adaptive surface reconstruction techniques capable of processing large and noisy datasets remains a major research issue in computer graphics and geometric modeling areas [Hornung06, Jenke06, Kazhdan06, Sharf06, Walder06, Schall06a].

Kazhdan [Kazhdan05] introduced a novel and elegant FFT-based reconstruction technique. His approach reconstructs a solid, watertight model from an oriented point set by determining a characteristic function which has the value one inside and the value zero outside of the resulting solid. The method allows for a robust and fast reconstruction from noisy samples. On the other hand, the approach has a high memory requirement due to its global nature. The integral of the characteristic function has to be sampled on a uniform grid for the whole volume in order to be able to apply the inverse FFT. This limits the maximal reconstruction resolution of the approach on today's computers to a level where the reconstruction of fine details of the input data is not possible. Furthermore, the approach has no global error control and its globality prevents the accurate reconstruction of regions which are close to each other but represent disconnected regions of the surface. Our work proposes a simple solution to overcome these limitations while preserving the advantages of the global approach.

The general idea of our technique is to employ an error-guided subdivision of the input data. For this, we compute the bounding box of the input and apply an octree subdivision. In order to decide whether an octree leaf cell needs to be subdivided, we compute a local characteristic function for the points inside the cell using the global FFT-based approach. This is a non-trivial task since the points inside a cell do in general not represent a solid. We propose a solution to that problem to avoid that surface parts are created which are not represented by points. If the resulting local approximation inside the cell is not accurate enough, the cell needs to be subdivided. By iterating this procedure, we compute overlapping local characteristic functions at the octree leaves for each part of the input with a user-defined accuracy. We obtain the final reconstruction by combining the local approximations using partition of unity and by extracting the surface using a polygonization algorithm. One advantage of our adaptive approach is that the characteristic function is only determined close to the surface and not for the whole volume. As the reconstruction accuracy is mainly limited by memory requirements, this allows us to obtain higher reconstruction resolutions. Additionally, the adaptiveness allows us a more accurate reconstruction of strongly bended regions of the input.

The rest of this chapter is organized as follows. Section (5.2) reviews related work on surface reconstruction. In Section (5.3) we present details on the data partitioning step as well as of the local characteristic functions of our approach. After this, their integration into a global function and the extraction of the final

surface is explained. We show results of our technique in Section (5.4) before we conclude in Section (5.5).

5.2 Related Work

Surface reconstruction is a research area with a long tradition and a large variety of techniques.

5.2.1 Delaunay-based Techniques

One important group of surface reconstruction methods which were already mentioned in Chapter 3 in the context of statistical point cloud denoising for improving reconstruction performance are Delaunay- or Voronoi-based approaches. They are mainly investigated in the area of computational geometry and exhibit interesting properties. One of them is that they are supported by rigorous mathematical results and provide reconstruction guarantees under certain sampling conditions. These conditions are typically formulated in terms of local feature size which is the distance of a surface point to its medial axis. The medial axis of a closed surface again is defined as the set of centers of empty balls which touch the surface at more than one point.

The first Delaunay-based method introduced to computer graphics is the crust algorithm by Amenta, Bern and Kamvyselis [Amenta98b]. It was developed from the crust algorithm for curve reconstruction by Amenta, Bern and Eppstein [Amenta98a]. The idea of the algorithm is to compute a Delaunay complex of the input points and their Voronoi poles which are an approximation to the medial axis of the sampled surface. From this complex all triangles crossing the medial axis are removed. The resulting candidate triangles form in general not the sampled surface but they contain a surface that is homeomorphic to it if the sampling is dense enough. This property was shown by Amenta and Bern using a result of Edelsbrunner and Shah [Edelsbrunner97]. Since the pole vector which points from the center of the Voronoi cell to its farthest Voronoi vertex is a good approximation for the surface normal, the algorithm is able to reconstruct a surface from a given point cloud without associated normals which is another important feature.

The cocone algorithm [Amenta02] was developed by Amenta, Choi, Dey and Leekha. It has the advantage that it requires the computation of only one Delaunay triangulation. Similar to the crust algorithm it provides the same mathematical

guarantee that the set of candidate triangles contains the resulting surface as a restricted Delaunay triangulation if the given sampling is dense enough.

Important extensions of both methods are algorithms that produce watertight surfaces. These techniques usually belong to the class of inside/outside labeling algorithms. They first compute a Delaunay tetrahedralization of the scattered data. Then the resulting tetrahedra are labeled inside or outside depending whether a tetrahedron is inside the solid bounded by the scattered data or outside. Two important methods generating watertight reconstructions are Power Crust developed by Amenta, Choi and Kolluri [Amenta01] and Tight Cocone by Dey and Goswami [Dey03].

The sampling preconditions under which the Delaunay complex contains the sampled surface are usually strict and do not often hold in practice. Therefore, computational geometry methods face difficulties while dealing with noisy data and undersampling. To alleviate these problems preparatory denoising is one option as illustrated in Chapter 3.

Dey and Goswami [Dey04a] presented the Robust Cocone algorithm which applies Delaunay prefiltering. It uses the observation that polar balls approximate the solid bounded by the input point set. A polar ball is defined as the maximal empty ball around a pole of the Voronoi diagram of a point set. If balls of adjacent tetrahedra intersect deeply both tetrahedra belong to the same component being either inside or outside. On the other hand if polar balls intersect shallowly their tetrahedra belong to different components. Dey and Goswami showed that this might not be true for noisy regions. They also showed that those polar balls have a small radius and can thus be detected. Robust Cocone only preserves the samples on the outer polar balls and reconstructs a watertight surface based on the Tight Cocone methodology.

Kolluri et al. [Kolluri04] introduced the Eigen Crust algorithm. The method computes a watertight surface reconstruction from noisy scattered data with outliers using spectral graph partitioning to label each tetrahedron of the Delaunay tetrahedralization. A pole graph is constructed whose edges are weighted according to the likelihood that pairs of dual tetrahedra are on the same side of the surface. This weighted graph is represented by a pole matrix whose smallest eigenvector is used to partition the graph in two subgraphs containing inside and outside poles. Tetrahedra which are not duals of poles are classified in order to produce a smooth surface with low genus.

Alliez et al. [Alliez07] reconstructed a watertight surface from an unoriented point set by using its Voronoi diagram to determine a tensor field whose principal axes locally represent the most likely direction of the normal to the surface as well

as a confidence for this estimation. Subsequently, an implicit function is computed by solving a generalized eigenvalue problem such that its gradient is most aligned with the principal axes of the tensor field. In particular, the implicit function optimization provides resilience to noise, adjustable fitting to the data, and controllable smoothness of the reconstructed surface.

Samozino et al. [Samozino06] presented a hybrid approach which combined Voronoi diagrams and Radial Basis Functions (RBFs) which are reviewed in the next section. The sampled shape is approximated as the zero-level set of a function. This function is defined as a linear combination of compactly supported radial basis functions. Different from previous work, centers of the basis functions are located on an estimate of the medial axis, instead of the input data points. Those centers are selected among the vertices of the Voronoi diagram of the given samples. Being a Voronoi vertex, each center is associated with a maximal empty ball. The radius of this ball is used to adapt the support of each radial basis function. The method can fit a user-defined amount of centers which are selected from the Voronoi vertices using the lambda medial axis. The main advantage of the approach is that it can achieve the same reconstruction accuracy as previous RBF techniques with a smaller amount of centers.

5.2.2 Radial Basis Functions (RBFs)

Opposed to Delaunay-based approaches which directly construct a triangulation or tetrahedralization respectively, the main idea behind most implicit surface interpolation techniques consists of building a function f whose zero-level set $\mathcal{Z}(f) = \{\mathbf{x} \mid f(\mathbf{x}) = 0\}$ approximates or interpolates a given input point set $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$. Usually f is constructed as a composition (weighted sum) of simple primitives.

Radial Basis Function (RBF) techniques are now standard tools for geometric data analysis [Franke80, Lodha99], in pattern recognition [Kirby01], statistical learning [Hastie01], and neural networks [Haykin99]. Properties of RBFs are widely studied in mathematical literature [Duchon77, Buhman03, Iske04, Wendland04].

Given a scattered point dataset, it is interpolated or approximated by the zero-level set of a composite function $f : \mathbf{R}^3 \rightarrow \mathbf{R}$ defined as a linear combination of relatively simple primitives

$$f(\mathbf{x}) = \sum_{i=1}^m \alpha_i \Phi(\mathbf{x}, \mathbf{c}_i) \quad (5.1)$$

where $\Phi(\cdot, \mathbf{c}_i) : \mathbf{R}^3 \rightarrow \mathbf{R}$ are functions centered at \mathbf{c}_i and α_i are the unknown weights [Girosi93].

It is desirable to constrain the solution to be stable to translation and rotation of the point set. The functions Φ are thus given by:

$$\Phi(\mathbf{x}, \mathbf{c}_i) = \phi(\|\mathbf{x} - \mathbf{c}_i\|), \quad (5.2)$$

where $\|\cdot\|$ denotes the Euclidian distance and $\phi : \mathbf{R}^+ \rightarrow \mathbf{R}$.

Relations between the RBF and variational approaches to scattered data interpolation / approximation are analyzed in [Duchon77, Haykin99].

Reconstruction using Radial Basis Functions gives a smooth implicit interpolating or approximating surface, since both the implicit solution and its zero-level set have the same continuity properties as the ones of the basis functions Φ .

Let $F^T = (f_1, \dots, f_n)$ be a vector of values of a function f at some scattered distinct points $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathbf{R}^3$. We want to find a function $f : \mathbf{R}^3 \rightarrow \mathbf{R}$ such that

$$f(\mathbf{x}_j) = \sum_{i=1}^n \alpha_i \phi(\|\mathbf{x}_j - \mathbf{x}_i\|) = f_j \quad (5.3)$$

The reconstruction problem thus boils down to determining the vector $\alpha^T = (\alpha_1, \dots, \alpha_n)$ by solving a linear system of equations given by the constraints (5.3). Since all constraints are located on the surface, all f_j s are valued zero. In order to avoid the trivial solution $\alpha = \mathbf{0}$, we add interior and exterior constraints where the function is non-zero and assign them the values $-d$ and d , respectively. In order to compute the weights α , we denote $[\phi(\|\mathbf{x}_i - \mathbf{x}_j\|)] = A_{X,\Phi}$ and have to solve the following linear system:

$$A_{X,\Phi} \cdot \alpha = F. \quad (5.4)$$

Since this linear system is under-determined the orthogonality conditions

$$\sum_{i=1}^n \alpha_i p(\mathbf{x}_i) = 0 \quad \forall p \in \mathbf{P}_k \quad (5.5)$$

where \mathbf{P}_k is the set of polynomials of order up to k are imposed on the coefficients α . Furthermore, a polynomial $p \in \mathbf{P}_k$ is added to (5.1) leading to the equation

$$f(\mathbf{x}) = p(\mathbf{x}) + \sum_{i=1}^m \alpha_i \Phi(\mathbf{x}, \mathbf{c}_i). \quad (5.6)$$

Some conventional radial basis functions are: $\phi : \mathbf{R}^+ \rightarrow \mathbf{R}$

biharmonic RBF $\phi(r) = r$ with a linear polynomial

pseudo-cubic RBF $\phi(r) = r^3$ with a linear polynomial

triharmonic RBF $\phi(r) = r^3$ with a quadratic polynomial

thin plate RBF $\phi(r) = r^2 \log(r)$ with a linear polynomial

All functions listed above have an unbounded support. The corresponding equations lead to a dense linear system. Therefore, recovering a solution is tractable only for small data sets. To overcome this problem, Morse et al. [Morse01] used Gaussians as Compactly Supported RBFs (CSRBFs) to obtain a sparse interpolation matrix. Other CSRBFs are proposed in [Wendland95, Wu95]. The main disadvantage of radial basis functions with compact support is that they are not well-suited for reconstruction from incomplete data. In order to handle large and incomplete datasets two strategies have been proposed. One approach uses polyharmonic RBFs (non-compactly supported functions) [Carr97], reduces the number of centers by a greedy selection procedure and performs fast evaluation using the so-called Fast Multipole Method (FMM). Another approach consists of using locally supported functions [Tobor04], where partition of unity is used for blending. The function support is computed locally for all centers, as described in [Ohtake04b]. A multiresolution version of this approach has been proposed in [Ohtake03b].

We can notice that radially symmetric functions are not suited for piecewise smooth surface reconstruction. Dinh et al. [Dinh01] have presented a method using anisotropic basis functions to overcome this issue.

5.2.3 Partition of Unity (PU)

Divide and conquer is the main idea behind the partition of unity approach. The main idea consists of breaking the domain into smaller subdomains where the problem can be solved locally. The data is first approximated on each subdomain separately, and the local solutions are blended together using a weighted sum of local subdomain approximations. The weights are smooth functions and sum up to one everywhere on the domain.

Tobor et al. [Tobor04] combined the partition of unity method and radial basis functions. Ohtake et al. [Ohtake03a] used weighted sums of different kinds of piecewise quadratic functions in order to capture the local shape of the surface. This way implicit surfaces from very large scattered point sets can be reconstructed.

Consider a global bounded domain Ω in an Euclidian space. Divide Ω into M overlapping subdomains $\{\Omega_i\}_{i=1,\dots,M}$ with $\Omega \subseteq \cup_i \Omega_i$. Together with this covering, we construct a partition of unity, i.e. a family of non-negative continuous com-

pactly supported functions $\{w_i\}_{i=1,\dots,M}$ such that $Supp(w_i) \subseteq \Omega_i$ and $\sum_{i=1}^M w_i = 1$ everywhere. Let $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_N\} \in \mathbf{R}^3$ be a set of N points on the surface. For each domain Ω_i , a set $\mathcal{P}_i = \{\mathbf{p} \in \mathcal{P} | \mathbf{p} \in \Omega_i\}$ is built, and the surface is approximated on each subdomain by a local approximant f_i . The global function is then defined as a combination of the local functions as:

$$f(\mathbf{x}) = \sum_{i=1}^N w_i(\mathbf{x}) f_i(\mathbf{x}). \quad (5.7)$$

The condition $\sum_{i=1}^M w_i = 1$ can be obtained from any other set of smooth functions W_i by a normalization process:

$$w_i(\mathbf{x}) = \frac{W_i(\mathbf{x})}{\sum_{j=1}^N W_j(\mathbf{x})}. \quad (5.8)$$

The weighting function W_i determines the continuity of the global reconstruction function f . We can generate these functions using local geometry of the corresponding cell as for instance the distance function or the center and radius of the cell.

For domain decomposition usually space subdivision trees are used. The cells can have different shapes such as axis-aligned bounding boxes, balls or axis-aligned ellipsoids. The degree of subdivision can be adapted to both local sampling density, desired smoothness and precision. The choice of the local fitting method is another degree of freedom. Radial basis functions are used in [Tobor04, Wendland02], while quadrics are used in [Ohtake03a]. Furthermore, one can adapt the fitting strategy for each cell according to the number of points and distribution of associated normals. Reconstructing sharp features is this way possible as described in [Ohtake03a].

An interesting combination of RBF and PU was proposed in [Ohtake04b] where a partition of unity is used to obtain an initial rough approximation of given scattered data and then RBFs are used to refine the PU approximation.

5.2.4 Fourier-based Surface Reconstruction

As we combine the FFT-based method introduced in [Kazhdan05] with partition of unity composition, we present an overview of the technique in this section.

The goal of the approach is to compute for an oriented point set $(\mathbf{p}_i, \mathbf{n}_i)$, $1 \leq i \leq N$ which bounds a solid V its characteristic function

$$\chi_V(x, y, z) = \begin{cases} 1 & (x, y, z) \in V \\ 0 & \text{otherwise} \end{cases} \quad (5.9)$$

which has the value one inside the solid and is zero otherwise. Instead of determining the characteristic function directly, the approach aims to find its Fourier coefficients which are defined as

$$\hat{\chi}_V(l, m, n) = \int_{\mathbf{R}^3} \chi_V(x, y, z) e^{-i(lx+my+nz)} dx dy dz \quad (5.10)$$

$$= \int_V e^{-i(lx+my+nz)} dx dy dz \quad (5.11)$$

In order to transfer this volume integral into a surface integral Stoke's theorem is used. The divergence theorem states that the volume integral over the divergence of a function F on its domain is equal to the surface integral of $\langle F(\mathbf{p}), \mathbf{n}(\mathbf{p}) \rangle$ over the boundary of the domain of the function F . More specifically the theorem can be written as

$$\int_V (\nabla \cdot F)(\mathbf{p}) d\mathbf{p} = \int_{\partial V} \langle F(\mathbf{p}), \mathbf{n}(\mathbf{p}) \rangle d\mathbf{p}. \quad (5.12)$$

Assuming $F_{lmn}(x, y, z)$ is a function whose divergence is equal to the (l, m, n) complex exponential

$$(\nabla \cdot F_{lmn})(x, y, z) = e^{-i(lx+my+nz)} \quad (5.13)$$

equation (5.11) can be rewritten as

$$\hat{\chi}_V(l, m, n) = \int_V (\nabla \cdot F_{lmn})(x, y, z) dx dy dz \quad (5.14)$$

$$= \int_{\partial V} \langle F(x, y, z), \mathbf{n}(x, y, z) \rangle dx dy dz. \quad (5.15)$$

As the given oriented point set is a Monte Carlo sampling of the boundary of the solid V , the last equation can be approximated by

$$\hat{\chi}(l, m, n) \approx \sum_{i=1}^N \langle F_{lmn}(\mathbf{p}_i), \mathbf{n}_i \rangle. \quad (5.16)$$

There are different choices for the function F . The one that is used in the paper has the advantage that it does not depend on the alignment of the coordinate axes.

$$F_{lmn}(x, y, z) = \begin{pmatrix} \frac{il}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \\ \frac{im}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \\ \frac{in}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \end{pmatrix} \quad (5.17)$$

The method reconstructs the characteristic function from the coefficients $\hat{\chi}(l, m, n)$ by using the inverse FFT on a regular grid. Subsequently, the surface is obtained by using a standard iso-surface extraction technique.

In a later work, Kazhdan et al. [Kazhdan06] reformulated the reconstruction problem into a Poisson problem. The goal of the approach is to find a characteristic function χ whose gradient approximates the vector field W given by the oriented point set as good as possible, i.e. $\min_{\chi} \|\nabla\chi - W\|$. Applying the divergence operator, the optimization problem is transferred into a Poisson problem

$$\Delta\chi = \nabla \cdot \nabla\chi = \nabla \cdot W. \quad (5.18)$$

The resulting linear system is sparse and solved using an adaptive Poisson solver. In contrast to the FFT-based approach the memory footprint is not cubic but proportional to the size of the reconstructed surface which allows the approach to achieve higher reconstruction resolutions. Bolitho et al. [Bolitho07] extended this method for out-of-core processing of very large meshes in a streaming framework. Manson et al. [Manson08] used wavelets for surface reconstruction since they provide a localized, multiresolution representation of functions suitable for stream processing.

5.2.5 Other Reconstruction Techniques

Another interesting approach not related to the presented categories of reconstruction techniques was presented by Sharf et al. [Sharf06] who determined a surface from a point cloud using an explicit mesh representation with multiple competing evolving fronts. These fronts adapt to the input data in a coarse-to-fine manner. Their use of an explicit surface representation guarantees watertightness. Furthermore, the approach allows for adaptive handling of regions with a non-homogeneous sampling density and reconstructs missing data in defected areas.

A more recent topology-aware approach of Sharf et al. [Sharf07] takes user input to correct reconstruction decisions at regions where the topology of the model cannot be automatically induced with a reasonable degree of confidence. First a continuous function over a three-dimensional domain is constructed. The zero-level set of this function is a first approximation of the reconstructed surface. At complex undersampled regions they analyze the local topological stability of the zero-level set which are, if necessary, suggested to the user who removes the ambiguities by adding local inside/outside constraints.

Gal et al. [Gal07] introduced an example-based surface reconstruction method. The approach uses a database of local shape priors built from a set of given context models that are chosen specifically to match a specific scan. Local neighborhoods of the input scan are matched with enriched patches of these models at multiple scales. Using the additional information given for the prior models,

the reconstruction can be enriched by normal data or with a feature classification. Depending on the quality of the prior models, noisy and under-sampled data can be reconstructed and sharp features can be recovered.

The interesting problem of building surfaces from non-parallel cross-section curves lying on planes with arbitrary orientations which has applications in bio-medical modeling was considered by Liu et al. [Liu08]. The desired output is a surface network that models both the exterior surface and the internal partitioning of the object. The method is guaranteed to produce a closed surface mesh that interpolates the curve network on each cross-section.

Hornung et al. [Hornung06] reconstructed a watertight triangle mesh from an un-oriented point cloud using an unsigned distance function. Therefore, no information about the local surface orientation is required. The approach estimates local surface confidence values within a region around the input samples. The surface which maximizes the global confidence is then extracted by computing the minimum cut of a weighted spatial graph structure in which edges with low weights pass regions with high confidence values.

5.3 Adaptive FFT-based Surface Reconstruction

In this section, we present our adaptive FFT-based surface reconstruction technique (in the following denoted as AdFFT) in detail. We first describe the error-controlled subdivision of the adaptive octree structure and the computation of overlapping local surface approximations for the input points associated with the octree leaves. We then integrate the local approximations using partition of unity composition to reconstruct the final model.

5.3.1 Adaptive Octree Subdivision

The general idea of the partition of unity approach is to divide the data domain into several pieces and to approximate the data in these domains separately. The resulting local approximations are then blended together using smooth and local weighting functions which sum up to one over the whole domain.

In order to find local characteristic functions of the domain bounding the input point cloud, we first compute the axis-aligned bounding box of the input data. We then apply an adaptive octree subdivision of this bounding box. In order to decide




Figure 5.1: Left: Local curve approximation for points inside and in the vicinity of a leaf cell (inner rectangle). The dashed line indicates the irrelevant region of the reconstructed solid. Right: Real 3D example of the sketch in the left image after pruning meaningless regions of the solid.

whether a cell needs to be subdivided, we compute the characteristic function of this cell and its vicinity with a fixed accuracy. If the surface extracted from this characteristic function approximates the points in the cell sufficiently close according to a user-defined accuracy, the cell has not to be subdivided further.

How to compute the characteristic function for a cell of the octree is not obvious, as a straightforward application of the global FFT-based method always determines a characteristic function representing a solid, whereas the points in a cell form in general non-closed surface patches. To avoid that irrelevant surface parts occur in the local characteristic function, we use the construction shown in Figure (5.1). We embed the octree cell including its oriented input samples at the center of a larger cell with doubled edge lengths. In order to allow a smooth transition between adjacent local characteristic functions later in the integration step, we add points in the vicinity of the original octree cell to the construction. In our implementation, we choose all points in the octree leaf cell scaled by a constant factor c around its center for the computation of the local approximation. If the parameter c is small, few samples in the neighborhood of the octree cell are considered to compute the local approximation. This might cause that reconstructions of adjacent octree cells have no smooth transition across their common boundary. Therefore, it is important that the parameter c is sufficiently large so that enough neighboring samples are considered. According to our experiments a constant factor of $c = 1.8$ works well for all performed tests.

By using the global FFT-based method with a fixed resolution (2^5 in our implementation) on the larger volume, we then compute its characteristic function at regular grid positions. As the shape of the octree cells is usually not cubical, we




Figure 5.2: Left: All local characteristic functions of octree cells containing the final iso-surface have a common resolution. This allows an easy interpolation between adjacent cells. Right: Example octree configuration for partition of unity blending. A corner point p of a Marching Cubes cell and radial kernels of octree cells with centers c_i and c_j are shown.

transform all candidate data points and normals to fit into a cube to enable the use of the FFT. Figure (5.1) sketches the idea behind this construction. The surface patch inside and in the vicinity of the octree cell is correctly reconstructed and the irrelevant surface part of the solid is outside of the inner cell. This works in the majority of the cases as the irrelevant surface part has a curved shape (see Figure 5.1). Additionally, adding sufficient samples in the vicinity of the octree cell increases the diameter of the shape that the unwanted part does not cross the smaller cell. In rare cases, the crossing cannot be avoided due to very different alignment of octree cell and local surface approximation. But since the resulting unwanted surface parts are small and distant to the real surface they can be pruned easily during the polygonization. The right image of Figure (5.1) shows a real example of a local surface approximation for an octree cell and its vicinity.

To measure the accuracy of the resulting local approximation, we construct a mesh from the computed characteristic function using the Marching Cubes algorithm [Lorensen87] and compute the Hausdorff distance of selected samples inside the cell to the mesh. If the average computed Hausdorff distance is above the user-defined error, the cell needs to be subdivided further. If a cell is empty, no approximation needs to be computed and we leave it untreated. In order to

guarantee an efficient computation of the Hausdorff error, we use only a subset of points inside the octree cell. In our implementation we select 10% of the cell points to obtain a stable estimation.

In the presence of noise it might happen that the error criterion cannot be reached everywhere on the dataset. This leads to an oversubdivision of octree cells in very noisy regions until subcells contain not enough samples to allow a robust local reconstruction. To avoid this, we introduce a stopping criterion to ensure a minimum number of samples in non-empty cells. We fix this lower bound to be 0.5% of the number of input points.

5.3.2 Integration


After the subdivision step, we obtain an octree with leaves on different depths which are either empty or contain a sampling of a local characteristic function which has the same resolution for all leaf cells. In order to obtain a common global resolution for all local characteristic functions, we reconstruct leaves with lower tree depths, which are larger octree cells, using a higher resolution inside each cell (see left illustration of Figure 5.2). This allows us to blend and to interpolate between adjacent cells and to apply the Marching Cubes algorithm on the resulting grid. To obtain the final reconstruction, we interleave the extraction of the iso-surface and the combination of the local characteristic functions. In order to be able to extract an iso-surface of a characteristic function which has a value of one inside the surface and zero outside of the surface, we need to choose an appropriate iso-value. We follow the global approach and choose it as the average value of the characteristic function values obtained at the input samples.

Our Marching Cubes implementation processes all octree cells for which local characteristic functions have been computed. As the local characteristic functions overlap each other, cubes close to the boundary of octree cells have more than one characteristic function value associated with its corners. To merge them into one value, we use partition of unity blending. More precisely, if we denote the corner position \mathbf{p} and use our octree data structure to find all local function values $\{f_0, \dots, f_N\}$ at this position which are associated with the cells $\{c_0, \dots, c_N\}$, we determine the global characteristic function value as

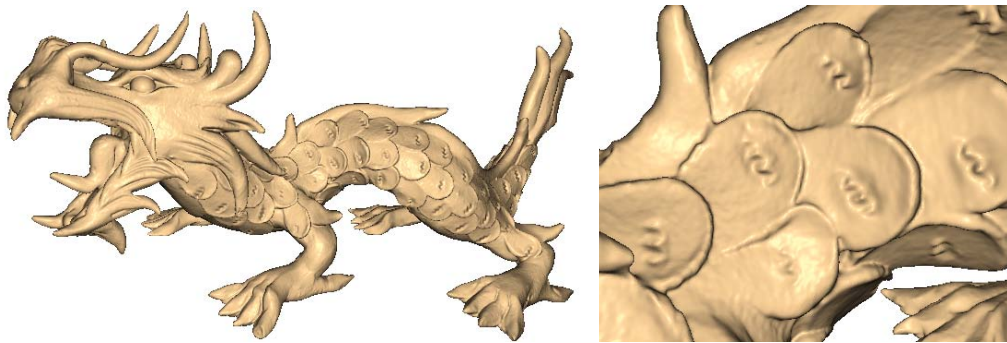
$$f_g = \frac{\sum_{i=0}^N w_i f_i}{\sum_{i=0}^N w_i} \quad (5.19)$$

where


$$w_i = G_i(\|\mathbf{c}_i - \mathbf{p}\|_2), \quad (5.20)$$



FFT [Kazhdan05], reconstruction time: 1m15s, peak memory: 1.4 GB



Poisson [Kazhdan06], reconstruction time: 65m15s, peak memory: 2.1 GB



our approach (AdFFT), reconstruction time: 21m10s, peak memory: 1.6 GB

Figure 5.3: The XYZ RGB Dragon model is reconstructed using the FFT-based, the Poisson-based approach and our technique. The meshes are computed using the maximal feasible resolution for all methods. The zooms show that the FFT approach is limited in its reconstruction accuracy due to its high memory consumption while the Poisson-based method and our technique achieve an almost identical high reconstruction accuracy. Note that our method needs less memory and reconstruction time to obtain these results.

and \mathbf{c}_i is the center of the cell c_i . The center of the radial Gaussian weighting function $G_i(\cdot)$ is fixed at \mathbf{c}_i . The bandwidth is chosen such that grid positions, which are farther from \mathbf{c}_i than the radius including the overlap of the cell c_i , have weights close to zero. For illustration see the right image of Figure (5.2).

After determining the global characteristic function values for the corners of the cubes, we can interpolate them across the cube edges to compute the position of the chosen iso-value. Our Marching Cubes implementation interpolates the resulting global function quadratically.

5.4 Results

In this section, we present results of our reconstruction algorithm. We compare our method with several state-of-the-art reconstruction techniques. Furthermore, we apply our method to real-world laser scanner data as well as large and complex point cloud data and discuss computation times and memory consumption.

Results of our reconstruction algorithm are shown in Figures (5.3) - (5.6). The reconstructions in Figures (5.5) and (5.6) are shown in flat shading to illustrate faceting. The meshes in Figures (5.4) and (5.3) are rendered in Phong shading to bring out small details on the surface as single triangles are not visible. Tables (5.1) and (5.2) in the end of this chapter summarize details for the presented reconstructions.

In Figure (5.3) we compare our technique with the global FFT-based and the Poisson-based approach [Kazhdan06]. The latter method reformulates the surface reconstruction problem into a sparse Poisson problem to overcome the limited reconstruction accuracy of the FFT-based technique. To compare the three algorithms, we reconstructed the XYZ RGB Dragon model as detailed as possible using the maximal feasible resolution for all approaches. Due to the lower memory consumption of the Poisson-based approach and our method (see Table 5.1), both techniques are able to reach higher reconstruction resolutions than the FFT-based method. This allows them to faithfully capture fine details for instance on the scales of the model. While both algorithms produce almost identical reconstruction results, our approach has a significantly better time and peak memory performance than the Poisson-based technique. Figure (5.4) shows another comparison with reconstructions of the Thai Statuette which are not created using the maximal but, if feasible, the same grid resolution. The result obtained using our algorithm shown in Figure (5.4 c) is more detailed than the Poisson-based in Figure (5.4 b). This is visible for instance on the trunk of the elephant and the necklace of the


(a) FFT [Kazhdan05] 512^3 (b) Poisson [Kazhdan06] 1024^3 (c) AdFFT 1024^3

Figure 5.4: Reconstructions of the Thai Statuette created using the global FFT-based approach (a), the Poisson-based approach (b) and our technique (c). Although the reconstructions (b) and (c) were created using the same grid resolution, our approach (c) preserves fine details more accurately. This can be observed for instance on the woman’s necklace or the elephant’s trunk.

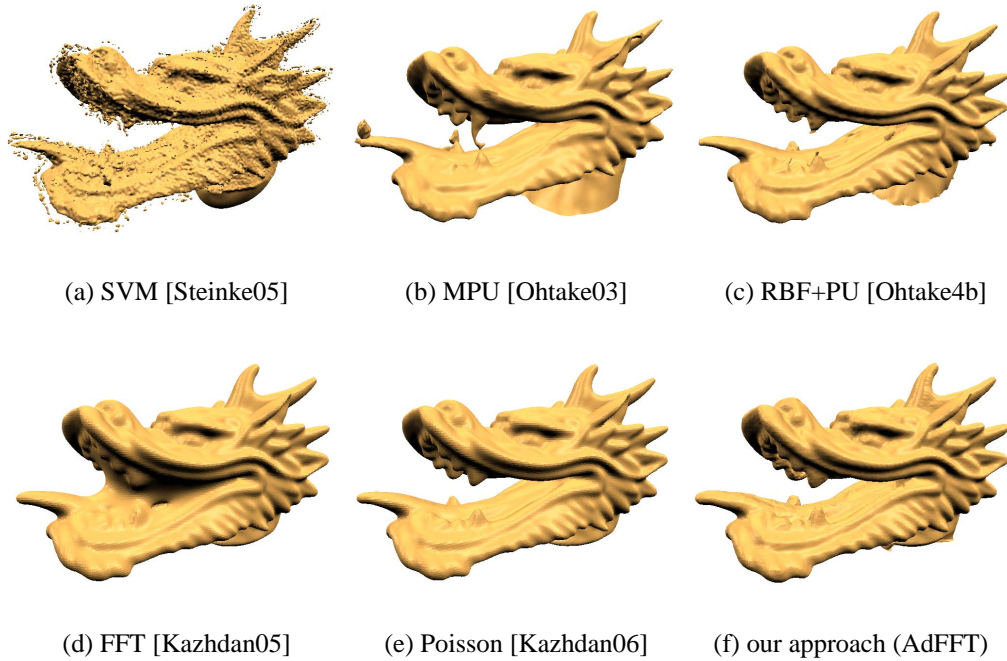


Figure 5.5: Comparison of our reconstruction approach (f) with other state-of-the-art techniques. We illustrate results on the unprocessed Dragon head composed of registered range scans which was already used in Chapter 3.9. Notice that our technique is more robust on noisy data than previous approaches (a)-(c) and generates a more faithful reconstruction in highly bended regions than the global FFT-based method (d). Its robustness is comparable to the recently introduced Poisson-based technique (e). Corresponding timings are reported in Table (5.1).

woman model. On the other hand, a close observation of our results on the belly of the woman statuette shows that our approach might reconstruct small bumps on the surface. These pose no problem as they can be removed, for instance, by smoothing the surface slightly subsequently to reconstruction. Note that although the Thai Statuette and the XYZ RGB Dragon model were decomposed into thousands of patches (see Table 5.2), our reconstructions show no blending artifacts. For all models in this chapter, we choose an overlap of 5 cells to smoothly blend adjacent reconstructions.

Figure (5.5) shows a comparison of recent state-of-the-art surface reconstruction techniques with our approach. As input data we choose the head of the original Dragon range scans from the Stanford 3D Scanning Repository. We compare our approach with the learning-based reconstruction technique using Support




Figure 5.6: The original Armadillo dataset composed of 114 registered range scans from the Stanford Scanning Repository (a) and a reconstruction from the noisy data using our method (c). Image (b) illustrates the patches without overlap used to reconstruct the Armadillo model. Figure (d) shows the mean curvature of our reconstruction (red represents negative and blue positive mean curvature values). Although the final reconstruction is composed of many patches, the mean curvature plot does not show blending artifacts.

Vector Machines (SVMs) [Steinke05], MPU [Ohtake03a], RBF+PU [Ohtake04b], the global FFT-based method [Kazhdan05] and the Poisson-based reconstruction technique [Kazhdan06]. The figure shows that SVM, MPU and RBF+PU create noisy reconstructions of the Dragon head scans and produce additional zero-level sets around the surface. Due to the global nature of the FFT approach, it robustly reconstructs noisy real-world data but has problems capturing regions where the surface folds back to itself. By localizing the global approach using adaptive decomposition and partition of unity blending, our algorithm accurately reconstructs those regions while retaining the robustness of the global approach. Results of the Poisson-based approach show that its robustness is comparable to our technique. Note that some methods, as for instance MPU and RBF+PU, obtain a better performance on real-world data by utilizing scanning confidence values, while our approach is robust on noisy data without using additional scanning information. Similar to computational geometry approaches, SVM is mainly suited to reconstruct clean data. Its performance in the presence of noise can be improved, for instance, by preprocessing the point cloud [Schall05a] before reconstructing a surface.

Figure (5.6) analyzes the effect of blending on the results of our surface reconstruction algorithm in more detail. For this, we computed a reconstruction of the original Armadillo range scans using our method. Figure (5.6 b) illustrates the patches without overlap used to create the integrated reconstruction shown in (c).

Figure (5.6 d) shows a mean curvature plot of (c). The results indicate that no visible blending artifacts close to the cell boundaries are introduced by our approach.

5.5 Conclusion

In this chapter, we have suggested to localize a global FFT-based reconstruction approach [Kazhdan05] by using adaptive subdivision and partition of unity blending. We have shown that our method preserves the resilience of the global approach and is more robust against noise than previous state-of-the-art surface reconstruction techniques. Furthermore, our reconstruction process is error-controlled, is capable of delivering a reliable surface reconstruction from noisy real-world data, and allows an accurate restoration of highly bended regions. The lower memory consumption of the method allows us to achieve a higher reconstruction accuracy and enables to capture fine and small details in large and complex point clouds. Another attractive feature of our method consists of its readiness for an out-of-core implementation.

The following chapter of this thesis describes the next step in the geometry processing pipeline which is surface remeshing. Often surface reconstruction algorithms do not produce meshes with geometry and connectivity suitable for efficient and stable further processing, i.e. curvature estimation. Therefore, the transition into a well-behaved mesh representation is essential. The next chapter introduces a technique which transforms irregular triangle meshes into regular quadrangular meshes while allowing for guidance of the remeshing process with design constraints.

method	reconstruction	polygonization	memory	user-defined error (global resolution)
SVM [Steinke05]	42s	5m15s	304MB	$1 \cdot 10^{-4}$
MPU [Ohtake03a]	1m40s		164MB	$7 \cdot 10^{-3}$
RBF+PU [Ohtake04b]	2m33s	34s	98MB	$1 \cdot 10^{-5}$
FFT [Kazhdan05]	8.9s	1.4s	179MB	(256^3)
FFT [Kazhdan05]	1m42s	9.2s	1.1GB	(512^3)
Poisson [Kazhdan06]	1m56s		68MB	(256^3)
Poisson [Kazhdan06]	7m02s		247MB	(512^3)
AdFFT	42s	1m03s	119MB	$1.7 \cdot 10^{-3} (256^3)$
AdFFT	3m08s	4m58s	462MB	$1.2 \cdot 10^{-3} (512^3)$

Table 5.1: Timings and memory consumption for the reconstructions shown in Figure (5.5). All results were computed on a 2.66 GHz Pentium 4 with 1.5 GB of RAM.

model	N	reconstruction	polygonization	memory	error	global resolution	M
Thai Statuette	5M	6m46s	11m15s	2.2GB	$2.4 \cdot 10^{-4}$	1024^3	2260
XYZ RGB Dragon	3.6M	11m15s	9m54s	1.6GB	$1.9 \cdot 10^{-4}$	1024^3	2614
Dragon head scans	485K	3m08s	4m58s	462MB	$1.2 \cdot 10^{-3}$	512^3	874
Dragon head scans	485K	42s	1m03s	119MB	$1.7 \cdot 10^{-3}$	256^3	626
Armadillo scans	2.4M	1m22s	55s	273MB	$1.1 \cdot 10^{-3}$	256^3	565

Table 5.2: Reconstruction information for the results presented in this chapter and computed using our method. The character N denotes the number of input samples and M the number of used patches. The results were computed on a 2.66 GHz Pentium 4 with 1.5 GB of RAM (only the Statuette and the XYZ RGB Dragon were computed on a 2.4 GHz AMD Opteron with 3 GB of RAM).

Chapter 6

Controlled Field Generation for Quad Remeshing

Quadrangular remeshing of triangulated surfaces has received an increasing attention in recent years. A particularly elegant approach is the extraction of quads from the streamlines of a harmonic field. While the construction of such fields is by now a standard technique in geometry processing, enforcing design constraints is still not fully investigated. This work presents a technique for handling directional constraints by directly controlling the gradient of the field. In this way, line constraints sketched by the user or automatically obtained as feature lines can be fulfilled efficiently. Furthermore, we show the potential of quasi-harmonic fields as a flexible tool for controlling the behavior of the field over the surface. Treating the surface as an inhomogeneous domain we can endow specific surface regions with field attraction/repulsion properties.

6.1 Introduction

Surface meshes delivered by laser scanning technology or iso-surface extraction as illustrated in the last chapter are in general irregularly sampled which reduces the efficiency of subsequent mesh processing applications. Therefore, conversion into regular triangular or quadrilateral meshes is a common requirement.

While triangular meshes are a widespread surface representation, quadrangular meshes are preferable for a considerable number of applications. Their tensor-product nature makes them particularly suited for serving as the parameter domain for spline representations [Li06]. Besides computer graphics, other indispensable applications comprise simulations using finite elements or architectural design [Liu06, Pottmann07]. This stimulated lively research and continuous progress in the areas of quad generation and remeshing.

In this chapter, we focus on the design aspects of quad remeshing using vector fields defined over triangular meshes. While the construction of such fields is by now a standard technique in geometry processing, enforcing design constraints is still not fully investigated. This work features the following contributions. First, we present a technique which allows control over the gradient of a harmonic field by aligning it to a set of line constraints. The constraints can be sketched by the user or automatically obtained using a feature line detection algorithm. Furthermore, inspired by the problem of modeling heat flow on inhomogeneous surfaces, we investigate the potential of quasi-harmonic fields as a tool for controlling the behavior of the field over the surface. We demonstrate that it can be used for allowing certain regions on the surface to attract or repulse field contour lines. Both techniques can be used separately or together without affecting the computational cost since the Laplacian is a special case of the quasi-harmonic operator. In all cases the runtime is dominated by solving a single linear system. Additionally, we address issues related to quad construction from the resulting vector fields.

Our approach offers the advantage that no post-processing is needed for resolving clipped primitives as proper alignment is addressed during the field construction stage. Additionally, the tools presented can be seamlessly used in combination with many of the existing quad remeshing techniques [Kälberer07, Ray06, Tong06].

The rest of this chapter is organized as follows. An overview of related literature is given in Section (6.2). Section (6.3) addresses the construction of gradient constrained harmonic fields and illustrates how quasi-harmonic fields can be used as a tool for field design. The construction of quad meshes is covered in Section (6.4). Section (6.5) presents and discusses the results of this work.

6.2 Related Work

In order to fulfill the ever increasing need of quad representations in a wide range of disciplines, research has a productive tradition in the closely related fields of

quad remeshing, parameterization, and vector field design bringing forth a large variety of approaches.

The work of Alliez et al. [Alliez03] on quadrangulation of triangle meshes uses principle curvature directions to guide the remeshing process. This approach was later extended by [Marinov04] to arbitrary meshes by applying curvature line integration on the underlying surface. Dong et al. [Dong05] compute a harmonic scalar field on the surface and determine the quadrangular facets by tracing integral lines of its gradient and orthogonal co-gradient vector field.

Boier-Martin et al. [Boier-Martin04] employ spatial- and normal-based clustering in order to segment the given triangular mesh into patches from which polygons are computed. Those are subsequently quadrangulated and subdivided resulting in the final quad-mesh. Kharevych et al. [Kharevych06] generate a patch layout using circle patterns while [Dong06] obtain a segmentation from the Morse-Smale complex of the eigenfunctions of the Laplacian. Marinov and Kobbelt [Marinov06] propose a two-step approach which first segments the mesh using a variant of variational shape approximation [Cohen-Steiner04] and then quadrangulates each patch independently using curves with minimum bending energy.

Tong et al. [Tong06] design quadrangulations by specifying a singularity graph on the triangular mesh. It allows for the representation of line singularities as well as singularities with fractional indices. Based on a modified discrete Laplacian operator, two scalar fields whose iso-contours form the quadrangular mesh are computed. Ray et al. [Ray06] determine a parameterization of a surface with arbitrary topology by defining two piecewise linear periodic functions which are aligned with two orthogonal vector fields defined on the surface. The quadrilaterals are subsequently extracted from the bivariate parameterization function. Inspired by this approach, [Kälberer07] propose frame fields based on branched covering spaces on a surface. In their context branch points are conceptually similar to singular points with fractional indices. They consider locally integrable fields which are not divergence free to improve the alignment of parameter lines with the given vector field.

Closely related to our approach, Fisher et al. [Fisher07] introduce a technique for the design of tangent vector fields based on discrete exterior calculus which is targeted at designing textures on surfaces. They constrain line integrals of a given vector field over the mesh edges according to a sparse set of user-provided constraints. Unlike their approach which operates on line integrals, we enforce vector constraints directly on the gradient of the field.




Figure 6.1: Heat distribution on a homogeneous (left) and inhomogeneous (right) plate, modeled using the Laplacian operator and the quasi-harmonic operator respectively.

6.3 Constrained Fields on Surface Meshes

In this section, we describe the techniques for controlling the construction of constrained scalar fields on triangular surface meshes. We start from the simple observation that the contours (co-gradient streamlines) of a scalar field are generally easier to construct. Therefore, it seems natural to apply the constraints directly to the contours. Thus the aim is to build harmonic fields whose contours satisfy the design constraints. In order to get a well behaved field on the surface mesh, we require the field to be harmonic. In the present discrete setup this translates to the construction of a piecewise linear function f within each triangle. The construction of such a function amounts to determining its values at the mesh vertices.

In order to control the behavior of the contours, we present two scenarios. The first imposes constraints on the gradient of the harmonic field while the second relies on quasi-harmonic fields which allow for more flexibility in vector field design in comparison to harmonic fields. In both scenarios, the field computation reduces to the solution of a linear system, a task that can be performed efficiently using standard direct or iterative solvers [Davis04, Chen06].

6.3.1 Gradient Constraints

On a triangular surface mesh, the gradient of a piecewise linear scalar field f is a piecewise constant vector field which exhibits discontinuity on the triangle sides. Analytically, on a triangle T described by its vertices $\{v_1, v_2, v_3\}$ and of area A

and normal n the gradient can be derived as

$$\nabla f = f_1 \frac{n \times (v_3 - v_2)}{2A} + f_2 \frac{n \times (v_1 - v_3)}{2A} + f_3 \frac{n \times (v_2 - v_1)}{2A}. \quad (6.1)$$

With this definition in mind, the alignment of contours to line constraints described by the triangle edges they traverse $\{e_i, i = 1..n\}$ amounts to imposing the following set of equations on the scalar field f

$$\nabla^2 f = 0 \quad (6.2)$$

$$\langle \nabla f, e_i \rangle = 0, i = 1..n, \quad (6.3)$$

where $\langle \cdot, \cdot \rangle$ stands for the dot product. In matrix form, this leads to an augmented matrix consisting of the Laplacian matrix and additional rows representing the gradient constraints. For this linear system to have a unique solution the value of the scalar field needs to be prescribed for at least one vertex.

We note that this approach is independent of the way the line constraints are determined. They can be obtained as feature lines automatically detected using methods such as [Ohtake04a, Yoshizawa05], or defined by the user using a sketching interface. In a preprocessing step, an intermediate mesh which contains new triangles along the feature lines is constructed and all the calculations are performed on it. As the re-triangulation is adaptive w.r.t. the feature lines, the mesh size does not increase significantly. This way our approach is more cost effective in comparison to the construction of higher order fields on the whole surface since elaborate interpolation schemes within the triangulation are avoided.

6.3.2 Quasi-harmonic Fields

While gradient constraints allow to directly enforce field directions by adding additional constraints to the harmonic equation, the approach described in this section allows certain regions to attract or repulse contour lines without introducing additional constraints. For this purpose, we rely on the notion of quasi-harmonic maps [Zayer05] to control the behavior of the field contours. The inspiring idea behind the approach stems from simple physical considerations. Let us consider the steady state heat equation on a quadrangular plate. In a first stage, we treat the surface as a homogeneous domain in the sense that the heat conductance is constant over the whole surface. The heat distribution can be obtained by solving a Laplace equation with prescribed conditions. In this example, we apply Dirichlet




Figure 6.2: Contours of quasi-harmonic fields on a rectangular plate. The ratio of conductivity (inversely related to diffusion) between the large and small rectangle was set to 1, 0, 1e-3, 1e3 from left to right respectively. Dirichlet boundary conditions were applied to the top and bottom sides and Neumann boundary conditions to the right and left sides of the plate.

conditions to the top and bottom sides and Neumann boundary conditions to the right and left sides as illustrated in Figure (6.1-left).

On the other hand, if we impose specific conductance values for the circular and rectangular sub-domains inside the plate, the standard Laplace equation is not suitable anymore for modeling the heat distribution and we have to rely on the so called quasi-harmonic equation which incorporates the conductance terms, and is therefore sensitive to the inhomogeneous nature of the plate. The heat distribution in this setup is depicted in Figure (6.1-right).

This example illustrates how simple scalar conductance values can alter the heat distribution on a simple domain. We capitalize on this observation for controlling the behavior of the contour lines.

Our approach proceeds by altering the scalar conductance values C (which are inversely related to diffusion) at the regions of interest and minimizes the following energy functional over the whole surface domain Ω

$$\int_{\Omega} (C \nabla f) \cdot (\nabla f), \quad (6.4)$$

with prescribed Dirichlet or Neumann boundary conditions. Figures (6.2) and (6.3) illustrate the effect of conductance values on the field contour lines. When the conductance is set to 0 on a certain region the field behaves as if Neumann boundary condition were applied at the region boundary. Higher diffusion values make the region repulse the contour lines while lower values make the region attract them.





Figure 6.3: Contours of a quasi-harmonic field on a rectangular plate. The conductance inside the circular and rectangular regions (red) is several orders of magnitude higher than the plate conductance.


6.4 Quad Construction

Once the constrained harmonic field has been computed, we use its contours to determine the orientation of the quads. The bottom row of Figure (6.4) illustrates the work-flow of our quadrangulation algorithm consisting out of two main steps. Firstly, we trace streamlines along and orthogonal to the contours of the harmonic field. In a second step, we use the set of connected streamlines to obtain the final quadrangular mesh.

More precisely, we compute given an harmonic scalar value at every vertex of the triangular mesh, the piecewise constant gradient on each triangle using the gradient discretization provided in equation (6.1). As it is usually desirable to have quads as rectangular as possible, we also determine the vector field orthogonal to the gradient vectors on each triangle which will be denoted in the following as co-gradient field. The co-gradient field is computed as the vector product of the gradient vector and the triangle normal on each face of the mesh. In specific situations the placement of line constraints may lead to a configuration where two edges of a triangle are constrained. A special treatment consisting of targeted subdivisions of the affected faces yields a correct behavior as illustrated in Figure (6.5).



(a)



(b)

Figure 6.4: Comparison between the approach introduced by Dong et al. [Dong05] (a) and our technique (b). The textured models (left column) illustrate the behavior of the contours of the harmonic fields for both techniques. For our approach the gradient is constrained orthogonal to the ascending edges and tangential to the base edges of the tetrahedron. This way, we avoid clipped primitives and obtain a proper alignment of the quads (right column) to the features of the surface.




Figure 6.5: In the special case where a line constraint (bold blue lines) covers two edges of one triangle, the gradient vector on the triangle (green) cannot be aligned to both vector constraint (orange) simultaneously. This is remedied by half-splitting the triangle as shown in the image to the right.

We regard a streamline as a piecewise linear curve on the surface which integrates one of the tangential vector fields and whose vertices are located on the edges of the triangular mesh. Starting from a given seed point, the streamline is integrated in the positive and negative field direction until it either creates a loop, approaches another streamline too closely or meets a singularity.

The singularities in the gradient vector field are detected based on the definition of the index of a critical point. Consider a continuous vector field V and a closed curve γ . Suppose that there are no critical points of V on γ . Let us move a point P along the curve in the counterclockwise direction. The vector $V(P)$ will rotate during the motion. When P returns to its starting place after one revolution along the curve, $V(P)$ also returns to its original position. During the journey $V(P)$ will make some whole number of revolutions. Counting these revolutions positively if they are counterclockwise, negatively if they are clockwise, the resulting algebraic sum of the number of revolutions is called the winding number of V on γ . The index of a point in the vector field V is then defined as the winding number of a small counterclockwise oriented circle centered at that point. A discretization of this definition for piecewise linear vector fields on surfaces is introduced in [Ray07]. We use a simplified version of their formulations which allows us to compute the index at a vertex v of our gradient vector field as

$$I(v) = \frac{1}{2\pi} \sum_{e \in \mathcal{N}(v)} \Theta(e) + A_d(v) \quad (6.5)$$

where $A_d(v)$ is the angle deficit at v and $\Theta(e)$ is the angle between the gradient vectors $\vec{g}(t_0)$ and $\vec{g}(t_1)$ after flattening the pair of triangles t_0 and t_1 adjacent to the edge e . We thus determine the extremal vertices as well as the saddle points as the points with the index 1 and -1 , respectively.

In order to trace streamlines on the surface, we sample the line constraints regularly and choose the selected points as seeds for the gradient streamlines. To cover the whole mesh we also place seed points on both sides of the streamline while it is traced according to a user-defined distance measure which controls the quad size. The co-gradient streamlines are traced accordingly by propagating them over the surface starting from the line constraints.

In the second step, we reconstruct quads from the set of streamlines. For this, we first determine the intersection points of gradient and co-gradient streamlines. In order to perform this efficiently, we associate the line segments of all streamlines with the triangles they are integrated on and compute the intersections on all triangles. This way, we obtain a graph of streamlines which are interconnected at their intersection points. Starting at these intersections we traverse the graph to create the quadrangular faces.

6.5 Results and Discussion

We tested our approach on a benchmark of triangular meshes covering synthetic and reconstructed laser-scanned data. Typical results are featured in Figure (6.4) and Figures (6.6) - (6.9).

A comparison to the clipping approach used in [Dong05], on a tetrahedron model reveals that our approach handles sharp features in an accurate manner as the alignment to line constraints is performed during the field computation (see Figure 6.4).

Figures (6.6) + (6.7) illustrate remeshing results of our approach on irregularly sampled data reconstructed from laser scans. While line constraints are useful to align quads to important features of the surface as illustrated on the Turbine Blade model, they can be further used to guide streamlines along user-specified constraints. The resulting remesh of the hand model reflects that the quads intuitively align to the shape of the surface.

Figure (6.8) shows the robustness of our technique on datasets with very irregular connectivity. While small triangles and faces with acute angles are usually challenging, the computation of the constrained harmonic field as well as the quad




Figure 6.6: Remeshing of a reconstructed laser-scanned Turbine Blade model with irregular triangulation. The remeshing process is guided by gradient constraints indicated by the blue lines (b). The remeshed result (c) as well as the zooms (a) show that the quads are properly aligned along the blade as well as the selected prominent features.

extraction step of our method remain numerically stable on these problematic datasets.

In Figure (6.9) we apply quasi-harmonic fields for driving the remeshing process. Using higher conductivity, contours around the eye region are repelled which allows the generated quads to follow the natural shape of the eyes.

Our approach is fast as the runtime is dominated by the solution of a linear system which can be solved efficiently using direct or iterative solvers. Even for large meshes such as the hand model consisting out of nearly 400K triangles runtimes are in the order of seconds. The subsequent streamline tracing as well as the quad generation is performed in only a few seconds.

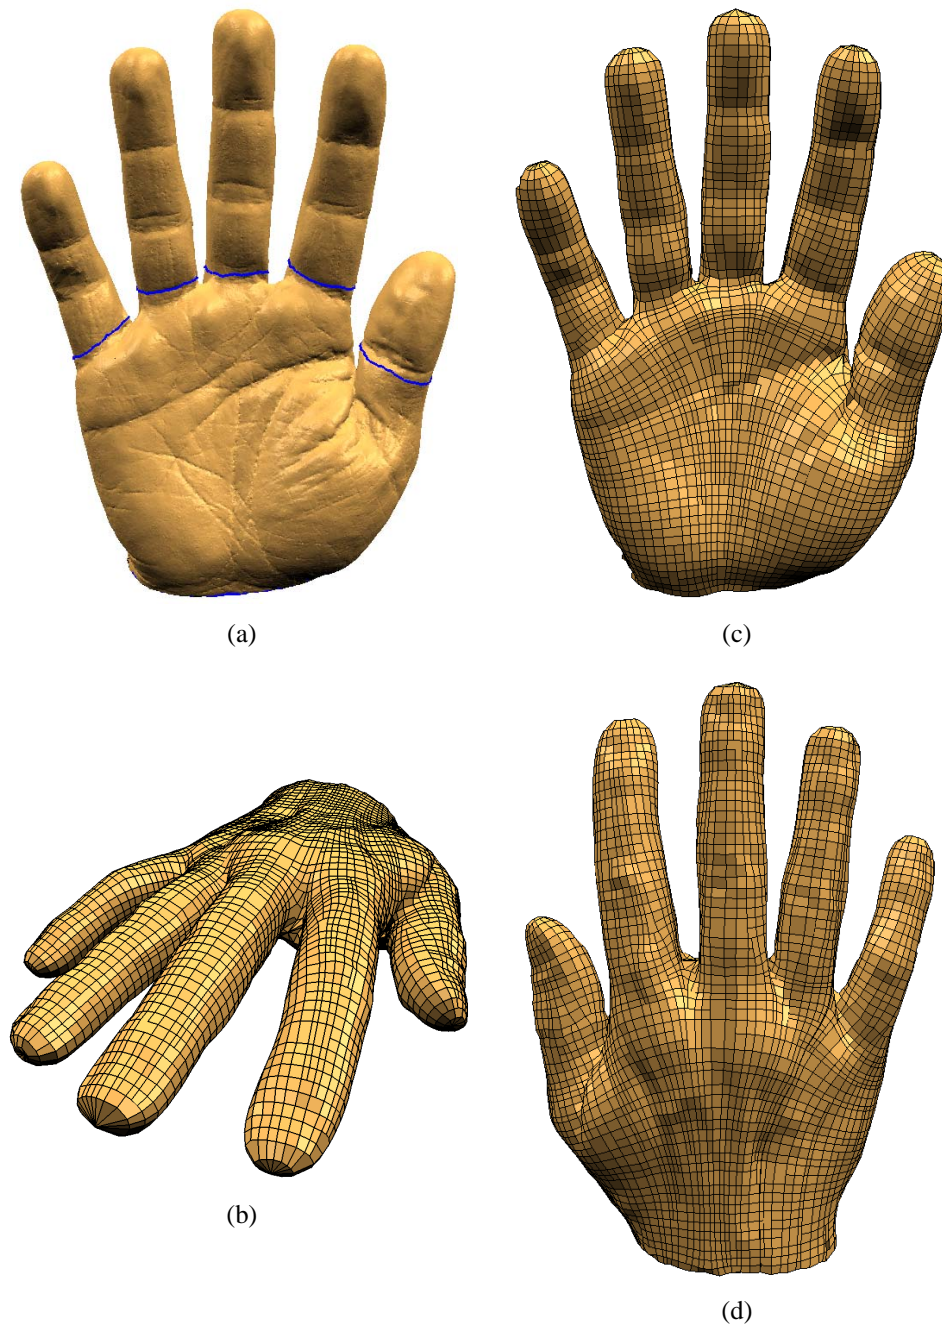
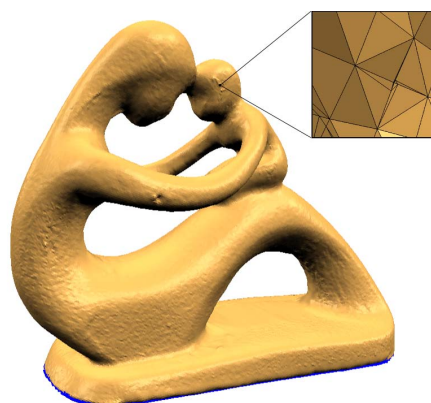



Figure 6.7: The laser-scanned hand model is remeshed guided by gradient constraints (a) which permits the resulting quads to follow the shape of the hand in an intuitive manner (c)+(d). Furthermore, our approach automatically places the singularities at the fingertips properly although the constraints are chosen very distant to them (b).



(a)



(b)



(c)

Figure 6.8: The fertility model is a reconstruction with extremely irregular connectivity containing tiny triangles and faces with very acute angles (a). Our method is robust on this type of datasets and creates a proper remeshing with alignment of the quads to the bottom of the pedestal of the figure (c).




Figure 6.9: Remeshing of the mannequin model using areas with modified conductivity (red regions) and gradient constraints (blue lines). The variation in the conductivity creates a repulsive effect such that the contour lines of the harmonic field (b) and thus the streamlines (c) bend around the eye region resulting in a special consideration of this surface part in the final quadrangular remeshing (d). Note that the gradient constraints specified at the lips as well as the neck ensure a proper quad alignment.

As our approach enforces directional constraints it may introduce additional singularities especially when certain line constraints form a closed curve. Those can be found efficiently during streamline tracing using the singularity detection technique described earlier. We do not see this as a limitation of our approach as it can be used to place new singularities at desired locations for design purposes such as the fingertips of the hand model illustrated in Figure (6.7).

6.6 Conclusion

In this chapter, we presented a set of flexible and versatile tools for designing scalar fields on surfaces. Two scenarios for controlling the field behavior on the surface were demonstrated. By operating directly on the gradient of the scalar field our technique can enforce directional constraints which makes it suitable for avoiding tedious post-processing, generally needed for aligning quads to important features. Regarding a surface as an inhomogeneous domain, we introduced quasi-harmonic fields as a design tool which endows surface regions with attraction/repulsion properties. This makes it a more general and flexible design tool in comparison to standard harmonic fields. These techniques can be used independently or on top of existing field-based quad remeshing methods. The substantiated results demonstrate the quality and robustness of our quad remeshing approach and confirm the flexibility of our field construction techniques.

In the next chapter, the presented topics are summarized and discussed which concludes this thesis.

Chapter 7

Discussion and Conclusion

This thesis presents new approaches towards a robust and efficient processing of data along the geometry processing pipeline. The focus areas in which new contributions are introduced comprise the denoising of static point clouds and dynamic range data, FFT-based surface reconstruction with partition of unity decomposition and guided quadrangular remeshing using gradient and area constraints.

7.1 Summary

In the following, we briefly summarize and discuss our techniques and show their advantages over existing approaches.

7.1.1 Statistical Denoising

In Chapter 3, a kernel based clustering approach for the robust filtering of point cloud surface data has been introduced. For a given point cloud the method determines a smooth global likelihood function which is used to guide the noisy input points onto a smooth surface at positions of maximum probability. It adapts to the density and distribution of the given input points which allows the technique to remove noise of different amplitudes accurately. Due to its inherent clustering property the approach permits an easy detection of outliers. It was demonstrated that the algorithm is able to handle even a large amount of “salt and pepper” outliers. The effectiveness of the approach was shown on real-world datasets acquired

using structured light and laser scanners. In addition, the method can be used in combination with surface reconstruction algorithms which significantly improves their results on noisy data.

7.1.2 Non-local Temporal Denoising

In Chapter 4, we presented a new similarity-based neighborhood filtering technique for static and dynamic range data which is the standard output of scanning devices and, in particular, of recently developed 3D video cameras. We introduced a new non-local similarity measure which determines the resemblance of two points on the surface not only by utilizing their local properties like position or normal but by also comparing the region of the surface surrounding the vertices instead. We demonstrated on several different types of scanned data that the idea of adding context information to the similarity definition allows our method to produce a more accurate denoising result than previous state-of-the-art approaches while having a better feature preservation. Additionally, we showed that the usage of adaptive neighborhoods improves the denoising result in the vicinity of boundaries of the given input. Furthermore, our method is easy to implement and flexibly adaptable to scans with different noise properties. It thus delivers a practical, versatile and powerful tool for filtering range data. In this way, our approach naturally fits into the scanning pipeline by denoising range scans before they are combined for further processing, which is more efficient since the given structure of the data can be utilized. Furthermore, we showed an interesting extension of our approach for filtering time-varying geometric data which is important since we expect a wider use of 3D video cameras in the future. It exploits the temporal coherence of the sequence in order to guarantee smoothness along the time domain.

7.1.3 Surface Reconstruction

We have suggested to localize a global FFT-based reconstruction approach by using adaptive subdivision and partition of unity composition. In Chapter 5, we have shown that our method preserves the resilience of the global approach and is more robust against noise than previous state-of-the-art surface reconstruction techniques. Furthermore, our reconstruction process is error-controlled, is capable of delivering a reliable surface reconstruction from noisy real-world data, and allows an accurate restoration of highly bended regions. The lower memory consumption of the method allows us to achieve a higher reconstruction accuracy and enables to capture fine and small details in large and complex point clouds.

7.1.4 Controlled Quad Remeshing

We presented a set of flexible and versatile tools for designing scalar fields on surfaces. In Chapter 6 two scenarios for controlling the field behavior on the surface were demonstrated. By operating directly on the gradient of the scalar field our technique can enforce directional constraints which makes it suitable for avoiding tedious post-processing, generally needed for aligning quads to important features. Regarding a surface as an inhomogeneous domain, we introduced quasi-harmonic fields as a design tool which endows surface regions with attraction/repulsion properties. This makes it a more general and flexible design tool for quadrangular remeshing in comparison to standard harmonic fields.

7.2 Future Work

While the presented techniques introduce novel solutions, their investigation reveals new interesting problems and extensions.

In Chapter 4, we presented that using context information in contrast to a point-wise defined similarity comparison, significantly improves the denoising performance. While we defined similarity in a spatio-temporal manner, it is interesting to explore the augmentation of the measure by additional attributes such as color which are usually acquired simultaneously with the geometric data. This could, for instance, support the detection of discontinuities to further improve the performance of neighborhood filtering schemes. The statistical point cloud denoising approach presented in Chapter 3 is realized as a research prototype and can be improved in terms of runtimes. As every input point converges independently to a maximum of the likelihood function, a massive parallelization of the presented approach is imaginable for a commercial implementation. Furthermore, it is interesting to explore how the idea of non-local neighborhoods can be transferred to our proposed statistical denoising approach. In the case of the mean-shift based denoiser each point in space is attracted by the other samples only depending on their local neighborhood and with an intensity diminishing with distance. Similar to the approach of non-local denoising, it is sensible to grant samples with a neighborhood similar to the vicinity of the current position in space a higher impact on the weighted mean which determines the next update step. Thus the problem amounts to determining the similarity of point configurations in three dimensional space.

In Chapter 5, we presented an approach which unifies partition of unity decomposition and surface reconstruction based on the Fast Fourier Transform. An attrac-

tive feature of our method consists of its readiness for an out-of-core implementation which is desirable for processing very large datasets especially in commercial applications. Furthermore, it is worth investigating how to incorporate confidence values, which can be determined for digitized point samples, into the reconstruction approach. While the technique is currently very robust, we are confident that this would enhance the performance of our approach.

The method introduced in Chapter 6 allows for a guided remeshing of triangular into quadrangular meshes using gradient as well as area constraints. An interesting future research direction is to investigate the automatic placement of constraints for design purposes or a perceptual appealing remeshing. Furthermore, the development of a unifying approach which combines all the individual advantages of different methods such as a well-behaved remeshing, accuracy, guidance, planar primitives and no T-junctions remains challenging.

Appendix A

Publications

The work presented in this thesis was published as articles in the following proceedings and journals:

- [1] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Adaptive Feature-preserving Non-local Denoising of Static and Time-varying Range Data. *Computer-Aided Design*, 40:701–707, 2008.
- [2] O. SCHALL, R. ZAYER, AND H.-P. SEIDEL. Controlled Field Generation for Quad-Remeshing. In *ACM Symposium on Solid and Physical Modeling 2008*, pages 295–300, Stony Brook, New York, USA, 2008. ACM.
- [3] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Feature-preserving Non-local Denoising of Static and Time-varying Range Data. In *ACM Symposium on Solid and Physical Modeling 2007*, pages 217–222, Beijing, China, 2007. ACM. **1st Best Paper Award.**
- [4] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Error-guided Adaptive Fourier-based Surface Reconstruction. *Computer-Aided Design*, 39(5):421–426, 2007.
- [5] W. SALEEM, O. SCHALL, G. PATANÈ, A. BELYAEV, AND H.-P. SEIDEL. On Stochastic Methods for Surface Reconstruction. *The Visual Computer*, 23(6):381–395, 2007. **AIM@SHAPE Best Paper of the Year 2006.**

- [6] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Adaptive Fourier-based Surface Reconstruction. In M.-S. Kim and K. Shimada, editors, *Geometric Modeling and Processing*, volume 4077 of *Lecture Notes in Computer Science*, pages 34–44, Pittsburgh, Pennsylvania, USA, 2006. Springer.
- [7] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Robust Filtering of Noisy Scattered Point Data. In M. Pauly and M. Zwicker, editors, *Eurographics Symposium on Point-Based Graphics 2005*, pages 71–77, Stony Brook, New York, USA, June 2005.
- [8] F. ISGRO, F. ODONE, W. SALEEM, AND O. SCHALL. Clustering for Surface Reconstruction. In *1st Int. Workshop on Semantic Virtual Environments*, pages 156–162, Villars sur Ollon, Switzerland, March 2005.
- [9] O. SCHALL, AND M. SAMOZINO. Surface from Scattered Points: A Brief Survey of Recent Developments. In *1st Int. Workshop on Semantic Virtual Environments*, pages 138–147, Villars sur Ollon, Switzerland, March 2005.

In addition, the temporal denoising approach from Chapter 4 was presented as a sketch at SIGGRAPH 2006:

- [10] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Feature-preserving Denoising of Time-varying Range Data. In H. Pfister, editor, *SIGGRAPH 2006 Sketches and Applications*, page 56, Boston, Massachusetts, USA, 2006. ACM.

Bibliography

- [Alexa01] M. ALEXA, J. BEHR, D. COHEN-OR, S. FLEISHMAN, AND C. T. SILVA. Point Set Surfaces. *IEEE Visualization 2001*, pages 21–28, October 2001.
- [Alexa02] M. ALEXA. Wiener Filtering of Meshes. In *Proceedings of Shape Modeling International*, pages 51–57, 2002.
- [Alliez03] P. ALLIEZ, D. COHEN-STEINER, O. DEVILLERS, B. LÉVY, AND M. DESBRUN. Anisotropic Polygonal Remeshing. *ACM Transactions on Graphics*, 22(3):485–493, 2003. Proceedings of ACM SIGGRAPH 2003.
- [Alliez07] P. ALLIEZ, D. COHEN-STEINER, Y. TONG, AND M. DESBRUN. Voronoi-based Variational Reconstruction of Unoriented Point Sets. In *Fifth Eurographics Symposium on Geometry Processing*, pages 39–48, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Amenta98a] N. AMENTA, M. BERN, AND D. EPPSTEIN. The Crust and the β -skeleton: Combinatorial Curve Reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998.
- [Amenta98b] N. AMENTA, M. BERN, AND M. KAMVYSSELIS. A new Voronoi-based Surface Reconstruction Algorithm. In *Proceedings of ACM SIGGRAPH 1998*, pages 415–421, 1998.
- [Amenta01] N. AMENTA, S. CHOI, AND R. KOLLURI. The Power Crust. In *Proceedings of 6th ACM Symposium on Solid Modeling*, pages 249–260, 2001.
- [Amenta02] N. AMENTA, S. CHOI, T. K. DEY, AND N. LEEKHA. A Simple Algorithm for Homeomorphic Surface Reconstruction. *Internat. J. Comput. Geom. & Appl.*, 12:125–141, 2002.

- [Amenta04] N. AMENTA AND Y. J. KIL. Defining Point-Set Surfaces. *ACM Transactions on Graphics*, 23(3):264–270, August 2004. Proceedings of SIGGRAPH 2004.
- [Awate06] S. P. AWATE AND R. T. WHITAKER. Unsupervised, Information-Theoretic, Adaptive Image Filtering with Applications to Image Restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 28(3):364–376, March 2006.
- [Bennett05] E. P. BENNETT AND L. MCMILLAN. Video Enhancement using per-pixel Virtual Exposures. *ACM Transactions on Graphics*, 24(3):845–852, 2005. Proceedings of ACM SIGGRAPH 2005.
- [Bickel07] B. BICKEL, M. BOTSCH, R. ANGST, W. MATUSIK, M. OTADUY, H. PFISTER, AND M. GROSS. Multi-scale Capture of Facial Geometry and Motion. *ACM Transactions on Graphics*, 26(3), 2007.
- [Boier-Martin04] I. BOIER-MARTIN, H. RUSHMEIER, AND J. JIN. Parameterization of Triangle Meshes over Quadrilateral Domains. In *Eurographics/ACM SIGGRAPH Symposium on Geometry Processing 2004*, pages 193–203, New York, NY, USA, 2004. ACM.
- [Bolitho07] M. BOLITHO, M. KAZHDAN, R. BURNS, AND H. HOPPE. Multilevel Streaming for Out-of-Core Surface Reconstruction. In *Fifth Eurographics Symposium on Geometry Processing (SGP '07)*, pages 69–78, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Botsch03] M. BOTSCH AND L. KOBELT. High-quality Point-based Rendering on Modern GPUs. In *Proceedings of Pacific Graphics '03*, 2003.
- [Buades05] A. BUADES, B. COLL, AND J. M. MOREL. A Non-local Algorithm for Image Denoising. In *Computer Vision and Pattern Recognition (CVPR) 2005*, volume 2, pages 60–65, 2005.
- [Buhman03] M. D. BUHMAN. *Radial Basis Functions: Theory and Implementations*, volume 12. Cambridge monographs on applied and computational mathematics edition, 2003.

- [Carr97] J. CARR, W. FRIGHT, AND R. BEATSON. Surface Interpolation with Radial Basis Functions for Medical Imaging. *IEEE Transactions Med. Imag.*, 16(1), 1997.
- [Chen06] Y. CHEN, T. A. DAVIS, W. W. HAGER, AND S. RAJAMANICKAM. Algorithm 8xx: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate. Technical Report TR-2006-005, 2006. submitted to ACM Transactions on Mathematical Software.
- [Cheng95] Y. CHENG. Mean Shift, Mode Seeking, and Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:790–799, 1995.
- [Cohen-Steiner04] D. COHEN-STEINER, P. ALLIEZ, AND M. DESBRUN. Variational Shape Approximation. *ACM Transactions on Graphics*, 23(3):905–914, 2004. Proceedings of ACM SIGGRAPH 2004.
- [Comaniciu02] D. COMANICIU AND P. MEER. Mean Shift: A Robust Approach Toward Feature Space Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, May 2002.
- [Davis03] J. DAVIS, R. RAMAMOOTHY, AND S. RUSINKIEWICZ. Spacetime Stereo: A Unifying Framework for Depth from Triangulation. In *Computer Vision and Pattern Recognition (CVPR) 2003*, pages 359–366, June 2003.
- [Davis04] T. A. DAVIS. Algorithm 832: UMFPACK, an Unsymmetric-pattern Multifrontal Method. *ACM Transactions on Mathematical Software*, 30(2):196–199, June 2004.
- [Davis05] J. DAVIS, D. NEHAB, R. RAMAMOOTHY, AND S. RUSINKIEWICZ. Spacetime Stereo: A Unifying Framework for Depth from Triangulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 27(2):296–302, February 2005.
- [Desbrun99] M. DESBRUN, M. MEYER, P. SCHRÖDER, AND A. H. BARR. Implicit Fairing of Irregular Meshes Using Diffusion and Curvature Flow. In *Proceedings of SIGGRAPH 99*, pages 317–324, 1999.

- [Dey03] T. K. DEY AND S. GOSWAMI. Tight Cocone: A Water-tight Surface Reconstructor. In *Proc. 8th ACM Sympos. Solid Modeling Applications*, pages 127–134, 2003.
- [Dey04a] T. K. DEY AND S. GOSWAMI. Provable Surface Reconstruction from Noisy Samples. In *Proc. 20th ACM Sympos. Comput. Geom.*, 2004.
- [Dey04b] T. K. DEY, S. GOSWAMI, AND J. SUN. Smoothing Noisy Point Clouds with Delaunay Preprocessing and MLS. Technical Report OSU-CISRC-3/04-TR17, The Ohio State University, 2004.
- [Dey05] T. K. DEY AND J. SUN. Adaptive MLS Surfaces for Reconstruction with Guarantees. In *Eurographics Symposium on Geometry Processing 2005*, pages 43–52, 2005.
- [Dinh01] H. Q. DINH, G. TURK, AND G. SLABAUGH. Reconstructing Surfaces Using Anisotropic Basis Functions. In *International Conference on Computer Vision (ICCV) 2001*, volume 2, pages 606–613, 2001.
- [Dong05] S. DONG, S. KIRCHER, AND M. GARLAND. Harmonic Functions for Quadrilateral Remeshing of Arbitrary Manifolds. *Computer-Aided Design*, 22(4):392–423, 2005.
- [Dong06] S. DONG, P.-T. BREMER, M. GARLAND, V. PASCUCCI, AND J. C. HART. Spectral Surface Quadrangulation. *ACM Transactions on Graphics*, 25(3):1057–1066, 2006. Proceedings of ACM SIGGRAPH 2006.
- [Duchon77] J. DUCHON. Spline minimizing rotation-invariant seminorms in Sobolev spaces. In W. Schempp and K. Zeller, editors, *Constructive Theory of Functions of Several Variables*, volume 571 of *Lecture Notes in Mathematics*, pages 85–100, 1977.
- [Edelsbrunner97] H. EDELSBRUNNER AND N. SHAH. Triangulating Topological Spaces. *Internat. J. Comput. Geom. & Appl.*, 7:365–378, 1997.
- [Efros99] A. A. EFROS AND T. K. LEUNG. Texture Synthesis by Non-Parametric Sampling. In *International Conference on Computer Vision (ICCV'99), Volume 2*, pages 1033–1038, 1999.

- [Fenn05] M. FENN AND G. STEIDL. Robust Local Approximation of Scattered Data. *Geometric Properties from Incomplete Data*, 2005.
- [Fisher07] M. FISHER, P. SCHRÖDER, M. DESBRUN, AND H. HOPPE. Design of Tangent Vector Fields. *ACM Transactions on Graphics*, 26(3):56, 2007. Proceedings of ACM SIGGRAPH 2007.
- [Fleishman03] S. FLEISHMAN, I. DRORI, AND D. COHEN-OR. Bilateral Mesh Denoising. *ACM Transactions on Graphics*, 22(3):950–953, 2003. Proceedings of ACM SIGGRAPH 2003.
- [Fleishman05] S. FLEISHMAN, D. COHEN-OR, AND C. T. SILVA. Robust Moving Least-squares Fitting with Sharp Features. *ACM Transactions on Graphics*, 24(3):544–552, 2005. Proceedings of ACM SIGGRAPH 2005.
- [Franke80] R. FRANKE AND G. NIELSON. Smooth Interpolation of Large Sets of Scattered Data. *International Journal for Numerical Methods in Engineering*, 15(11):1691–1704, 1980.
- [Fukunaga75] K. FUKUNAGA AND L. D. HOSTETLER. The Estimation of the Gradient of a Density Function with Applications in Pattern Recognition. *IEEE Transactions on Information Theory*, 21:32–40, 1975.
- [Gal07] R. GAL, A. SHAMIR, T. HASSNER, M. PAULY, AND D. COHEN-OR. Surface Reconstruction using Local Shape Priors. In *Fifth Eurographics Symposium on Geometry Processing (SGP '07)*, pages 253–262, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Girosi93] F. GIROSI, M. JONES, AND T. POGGIO. Priors Stabilizers and Basis Functions: From Regularization to Radial, Tensor and Additive Splines. Technical Report AIM-1430, 1993.
- [Hastie01] H. HASTIE, R. TIBSHIRANI, AND J. H. FRIEDMAN. *The Elements of Statistical Learning*. Springer, 2001.
- [Haykin99] S. HAYKIN. *Neural Networks: A Comprehensive Foundation*. Prentice Hall, New Jersey, USA, 1999.
- [Hildebrandt04] K. HILDEBRANDT AND K. POLTHIER. Anisotropic Filtering of Non-linear Surface Features. *Computer Graphics Fo-*

- rum*, 23(3):391–400, 2004. Proceedings of EUROGRAPHICS 2004.
- [Hildebrandt07] K. HILDEBRANDT AND K. POLTHIER. Constraint-based Fairing of Surface Meshes. In *Fifth Eurographics Symposium on Geometry Processing (SGP '07)*, pages 203–212, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Hornung06] A. HORNUNG AND L. KOBBELT. Robust Reconstruction of Watertight 3D Models from Non-uniformly Sampled Point Clouds Without Normal Information. In *Fourth Eurographics Symposium on Geometry Processing (SGP '06)*, pages 41–50. Eurographics Association, 2006.
- [Isgro05] F. ISGRO, F. ODONE, W. SALEEM, AND O. SCHALL. Clustering for Surface Reconstruction. In *1st Int. Workshop on Semantic Virtual Environments*, pages 156–162, Villars sur Ollon, Switzerland, March 2005.
- [Iske04] A. ISKE. *Multiresolution Methods in Scattered Data Modelling*. Springer-Verlag, Heidelberg, Germany, 2004.
- [Ivrissimtzi04] I. IVRISSIMTZIS, Y. LEE, S. LEE, W.-K. JEONG, AND H.-P. SEIDEL. Neural Mesh Ensembles. In Y. Aloimonos and G. Taubin, editors, *3DPTV 2004*, Los Alamitos, USA, 2004. IEEE.
- [Jenke06] P. JENKE, M. WAND, M. BOKELOH, A. SCHILLING, AND W. STRASSER. Bayesian Point Cloud Reconstruction. *Computer Graphics Forum*, 25(3), 2006. Proceedings of EUROGRAPHICS 2006.
- [Jones03] T. R. JONES, F. DURAND, AND M. DESBRUN. Non-iterative Feature-preserving Mesh Smoothing. *ACM Transactions on Graphics*, 22(3):943–949, July 2003. Proceedings of SIGGRAPH 2003.
- [Kälberer07] F. KÄLBERER, M. NIESER, AND K. POLTHIER. QuadCover - Surface Parameterization using Branched Coverings. *Computer Graphics Forum*, 26(3):375–384, 2007.
- [Kazhdan05] M. KAZHDAN. Reconstruction of Solid Models from Oriented Point Sets. In *Third Eurographics Symposium on Geometry Processing (SGP '05)*, pages 73–82, 2005.

- [Kazhdan06] M. KAZHDAN, M BOLITHO, AND H. HOPPE. Poisson Surface Reconstruction. In *Fourth Eurographics Symposium on Geometry Processing (SGP'06)*, pages 61–70, 2006.
- [Kharevych06] L. KHAREVYCH, B. SPRINGBORN, AND P. SCHRÖDER. Discrete Conformal Mappings via Circle Patterns. *ACM Transactions on Graphics*, 25(2):412–438, 2006.
- [Kirby01] M. KIRBY. *Geometric Data Analysis*. John Wiley & Sons, 2001.
- [Kolluri04] R. KOLLURI, J. R. SHEWCHUK, AND J. F. O'BRIEN. Spectral Surface Reconstruction from Noisy Point Clouds. In *Symposium on Geometry Processing*, pages 11–21. ACM Press, July 2004.
- [Lange05] C. LANGE AND K. POLTHIER. Anisotropic Fairing of Point Sets. *Special Issue of Computer Aided Geometric Design*, 22(7):680–692, 2005.
- [Levin98] D. LEVIN. The Approximation Power of Moving Least-squares. *Math. Comput.*, 67(224):1517–1531, 1998.
- [Li06] W.-C. LI, N. RAY, AND B. LÉVY. Automatic and Interactive Mesh to T-spline Conversion. In *Eurographics/ACM SIGGRAPH Symposium on Geometry Processing 2006*, pages 191–200, 2006.
- [Linsen01] L. LINSEN. Point Cloud Representation. Technical Report 2001-3, Fakultät für Informatik, Universität Karlsruhe, 2001.
- [Lipman07a] Y. LIPMAN, D. COHEN-OR, AND D. LEVIN. Data-dependent MLS for Faithful Surface Approximation. In *Fifth Eurographics Symposium on Geometry Processing (SGP '07)*, pages 59–67, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Lipman07b] Y. LIPMAN, D. COHEN-OR, D. LEVIN, AND H. TAL-EZER. Parameterization-free Projection for Geometry Reconstruction. *ACM Transactions on Graphics*, 26(3), 2007.
- [Liu06] Y. LIU, H. POTTMANN, J. WALLNER, Y.-L. YANG, AND W. WANG. Geometric Modeling with Conical Meshes and Developable Surfaces. *ACM Transactions on Graphics*, 25(3):681–689, 2006. Proceedings of ACM SIGGRAPH 2006.

- [Liu08] L. LIU, C. BAJAJ, J. O. DEASY, D. A. LOW, AND T. JU. Surface Reconstruction From Non-parallel Curve Networks. *Computer Graphics Forum*, 27(2):155–163, 2008. Proceedings of Eurographics 2008.
- [Lodha99] S. K. LODHA AND R. FRANKE. Scattered Data Techniques for Surfaces. In *Dagstuhl '97, Scientific Visualization*, pages 181–222. IEEE Computer Society, 1999.
- [Lorensen87] W. E. LORENSEN AND H. E. CLINE. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *Computer Graphics*, 21(3):163–169, 1987. Proceedings of ACM SIGGRAPH '87.
- [Mahmoudi05] M. MAHMOUDI AND G. SAPIRO. Fast Image and Video Denoising via Nonlocal Means of Similar Neighborhoods. *Signal Processing Letters*, 12(12):839–842, 2005.
- [Manson08] J. MANSON, G. PETROVA, AND S. SCHAEFER. Streaming Surface Reconstruction Using Wavelets. *Computer Graphics Forum*, 27(5):1411–1420, 2008. Proceedings of SGP 2008.
- [Marinov04] M. MARINOV AND L. KOBBELT. Direct Anisotropic Quad-dominant Remeshing. In *Proc. of the 12th Pacific Conf. on Comp. Graph. and Appl.*, pages 207–216, 2004.
- [Marinov06] M. MARINOV AND L. KOBBELT. A Robust Two-step Procedure for Quad-dominant Remeshing. *Computer Graphics Forum*, 25(3):207–216, 2006.
- [Mederos03] B. MEDEROS, L. VELHO, AND L. H. DE FIGUEIREDO. Robust Smoothing of Noisy Point Clouds. In *Proc. SIAM Conference on Geometric Design and Computing*, Seattle, USA, 2003. Nashboro Press.
- [Mederos04] B. MEDEROS, L. VELHO, AND L. H. DE FIGUEIREDO. Smooth Surface Reconstruction from Noisy Clouds. *Journal of the Brazilian Computing Society*, 2004.
- [Mederos05] B. MEDEROS, N. AMENTA, L. VELHO, AND L. H. DE FIGUEIREDO. Surface Reconstruction from Noisy Point Clouds. In *Eurographics Symposium on Geometry Processing 2005*, pages 53–62, 2005.
- [Morse01] B. S. MORSE, T. S. YOO, D. T. CHEN, P. RHEINGANS, AND K. R. SUBRAMANIAN. Interpolating Implicit Surfaces

- from Scattered Surface Data Using Compactly Supported Radial Basis Functions. In *SMI '01: Proceedings of the International Conference on Shape Modeling & Applications*, pages 89–98. IEEE Computer Society, 2001.
- [Nehab05] D. NEHAB, S. RUSINKIEWICZ, J. DAVIS, AND R. RAMAMOORTHI. Efficiently Combining Positions and Normals for Precise 3D Geometry. *ACM Transactions on Graphics*, 24(3):536–543, 2005. Proceedings of ACM SIGGRAPH 2005.
- [Ohtake03a] Y. OHTAKE, A. BELYAEV, M. ALEXA, G. TURK, AND H.-P. SEIDEL. Multi-level Partition of Unity Implicits. *ACM Transactions on Graphics*, 22(3):463–470, July 2003. Proceedings of SIGGRAPH 2003.
- [Ohtake03b] Y. OHTAKE, A. G. BELYAEV, AND H.-P. SEIDEL. A Multi-scale Approach to 3D Scattered Data Interpolation with Compactly Supported Basis Functions. In *Shape Modeling International 2003*, pages 153–161, Seoul, Korea, May 2003.
- [Ohtake04a] Y. OHTAKE, A. BELYAEV, AND H.-P. SEIDEL. Ridge-valley Lines on Meshes via Implicit Surface Fitting. *ACM Transactions on Graphics*, 23(3):609–612, 2004. Proceedings of ACM SIGGRAPH 2004.
- [Ohtake04b] Y. OHTAKE, A. G. BELYAEV, AND H.-P. SEIDEL. 3D Scattered Data Approximation with Adaptive Compactly Supported Radial Basis Functions. In *Shape Modeling International 2004*, pages 31–39, Genova, Italy, June 2004.
- [Ohtake05] Y. OHTAKE, A. G. BELYAEV, AND H.-P. SEIDEL. An Integrating Approach to Meshing Scattered Point Data. In *ACM Symposium on Solid and Physical Modeling*, pages 61–69, 2005.
- [Paris06] S. PARIS AND F. DURAND. A Fast Approximation of the Bilateral Filter using a Signal Processing Approach. In *European Conference on Computer Vision (ECCV)*, 2006.
- [Parzen62] E. PARZEN. On the Estimation of a Probability Density Function and the Mode. *Annals of Mathematical Statistics*, 33:1065–1076, 1962.

- [Pauly01] M. PAULY AND M. GROSS. Spectral Processing of Point-Sampled Geometry. *Proceedings of SIGGRAPH 2001*, pages 379–386, 2001.
- [Pauly03] M. PAULY, R. KEISER, L. P. KOBBELT, AND M. GROSS. Shape Modeling with Point-sampled Geometry. *Proceedings of SIGGRAPH 2003*, 22:641–650, 2003.
- [Pauly04] M. PAULY, N. J. MITRA, AND L. J. GUIBAS. Uncertainty and Variability in Point Cloud Surface Data. In *Eurographics Symposium on Point-Based Graphics*, pages 77–84, Zurich, Switzerland, June 2004.
- [Perona90] P. PERONA AND J. MALIK. Scale-Space and Edge Detection Using Anisotropic Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [Pfister00] H. PFISTER, M. ZWICKER, J. VAN BAAR, AND M. GROSS. Surfels: Surface Elements as Rendering Primitives. In *Proceedings of ACM SIGGRAPH 2000*, pages 335–342, July 2000.
- [Pottmann07] H. POTTMANN, Y. LIU, J. WALLNER, A. BOBENKO, AND W. WANG. Geometry of Multi-layer Freeform Structures for Architecture. *ACM Transactions on Graphics*, 26(3):65, 2007. Proceedings of ACM SIGGRAPH 2007.
- [Press93] W. H. PRESS, S. A. TEUKOLSKY, W. T. VETTERLING, AND B. P. FLANNERY. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1993.
- [Ray06] N. RAY, LI W.-C., B. LÉVY, A. SHEFFER, AND P. ALLIEZ. Periodic Global Parameterization. *ACM Transactions on Graphics*, 25(4):1460–1485, 2006.
- [Ray07] N. RAY, B. VALLET, W.-C. LI, AND B. LÉVY. N-Symmetry Direction Field Design. *ACM Transactions on Graphics*, 2007. accepted pending revisions, to appear.
- [Rosenblatt56] M. ROSENBLATT. Remarks on some Non-parametric Estimates of a Density Function. *Annals of Mathematical Statistics*, 27:832–837, 1956.
- [Rudin92] L. I. RUDIN, S. OSHER, AND E. FATEMI. Nonlinear Total Variation based Noise Removal Algorithms. In *Physica D 60*, pages 259–268. Elsevier North-Holland, Inc., 1992.

- [Rusinkiewicz00] S. RUSINKIEWICZ AND M. LEVOY. QSplat: A Multiresolution Point Rendering System for Large Meshes. In *Proceedings of ACM SIGGRAPH 2000*, pages 343–352, 2000.
- [Saleem07] W. SALEEM, O. SCHALL, G. PATANÈ, A. BELYAEV, AND H.-P. SEIDEL. On Stochastic Methods for Surface Reconstruction. *The Visual Computer*, 23(6):381–395, 2007. AIM@SHAPE Best Paper of the Year 2006.
- [Samozino06] M. SAMOZINO, M. ALEXA, P. ALLIEZ, AND M. YVINEC. Reconstruction with Voronoi centered Radial Basis Functions. In *Fourth Eurographics Symposium on Geometry Processing (SGP '06)*, pages 51–60. Eurographics Association, 2006.
- [Schall05a] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Robust Filtering of Noisy Scattered Point Data. In M. Pauly and M. Zwicker, editors, *Eurographics Symposium on Point-Based Graphics 2005*, pages 71–77, Stony Brook, New York, USA, June 2005.
- [Schall05b] O. SCHALL AND M. SAMOZINO. Surface from Scattered Points: A Brief Survey of Recent Developments. In *1st Int. Workshop on Semantic Virtual Environments*, pages 138–147, Villars sur Ollon, Switzerland, March 2005.
- [Schall06a] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Adaptive Fourier-Based Surface Reconstruction. In M.-S. Kim and K. Shimada, editors, *Geometric Modeling and Processing*, volume 4077 of *Lecture Notes in Computer Science*, pages 34–44, Pittsburgh, Pennsylvania, USA, 2006. Springer.
- [Schall06b] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Feature-preserving Denoising of Time-varying Range Data. In H. Pfister, editor, *SIGGRAPH 2006 Sketches and Applications*, page 56, Boston, Massachusetts, USA, 2006. ACM.
- [Schall07a] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Error-guided Adaptive Fourier-based Surface Reconstruction. *Computer-Aided Design*, 39(5):421–426, 2007.
- [Schall07b] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Feature-preserving Non-local Denoising of Static and Time-varying Range Data. In *ACM Symposium on Solid and Physical Modeling 2007*, pages 217–222, Beijing, China, 2007. ACM. 1st Best Paper Award.

- [Schall08a] O. SCHALL, A. BELYAEV, AND H.-P. SEIDEL. Adaptive Feature-preserving Non-local Denoising of Static and Time-varying Range Data. *Computer-Aided Design*, 40:701–707, 2008.
- [Schall08b] O. SCHALL, R. ZAYER, AND H.-P. SEIDEL. Controlled Field Generation for Quad-remeshing. In *ACM Symposium on Solid and Physical Modeling 2008*, pages 295–300, Stony Brook, New York, USA, 2008. ACM.
- [Schölkopf05] B. SCHÖLKOPF, J. GIESEN, AND S. SPALINGER. Kernel Methods for Implicit Surface Modeling. In *Advances in Neural Information Processing Systems 17*, pages 1193–1200. MIT Press, Cambridge, MA, 2005.
- [Sharf04] A. SHARF, M. ALEXA, AND D. COHEN-OR. Context-based Surface Completion. *ACM Transactions on Graphics*, 23(3):878–887, August 2004. Proceedings of SIGGRAPH 2004.
- [Sharf06] A. SHARF, T. LEWINER, A. SHAMIR, L. KOBBELT, AND D. COHEN-OR. Competing Fronts for Coarse-to-Fine Surface Reconstruction. *Computer Graphics Forum*, 25(3):389–398, 2006. Proceedings of Eurographics 2006.
- [Sharf07] A. SHARF, T. LEWINER, G. SHKLARSKI, S. TOLEDO, AND D. COHEN-OR. Interactive Topology-aware Surface Reconstruction. In *ACM Transactions on Graphics*, volume 26, New York, NY, USA, 2007. ACM.
- [Steinke05] F. STEINKE, B. SCHÖLKOPF, AND V. BLANZ. Support Vector Machines for 3D Shape Processing. *Computer Graphics Forum*, 24(3):285–294, 2005. Proceedings of EUROGRAPHICS 2005.
- [Stewart99] C. V. STEWART. Robust Parameter Estimation in Computer Vision. *SIAM Review*, 41:513–537, 1999.
- [Suessmuth07] J. SUESSMUTH AND G. GREINER. Ridge based Curve and Surface Reconstruction. In *Fifth Eurographics Symposium on Geometry Processing (SGP '07)*, pages 243–251, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Tasdizen02] T. TASDIZEN, R. WHITAKER, P. BURCHARD, AND S. OSHER. Geometric Surface Smoothing via Anisotropic Diffu-

- sion of Normals. In *Proceedings of IEEE Visualization 2002*, pages 125–132, Washington, DC, USA, 2002. IEEE Computer Society.
- [Taubin95] G. TAUBIN. A Signal Processing Approach To Fair Surface Design. In *Proceedings of SIGGRAPH 95*, pages 351–358, 1995.
- [Tobor04] I. TOBOR, P. REUTER, AND C. SCHILCK. Efficient Reconstruction of Large Scattered Geometric Datasets using the Partition of Unity and Radial Basis Functions. *Journal of WSCG 2004*, 12:467–474, 2004.
- [Tomasi98] C. TOMASI AND R. MANDUCHI. Bilateral Filtering for Gray and Color Images. In *Proceedings of the Sixth International Conference on Computer Vision (ICCV)*, pages 839–846, 1998.
- [Tong06] Y. TONG, P. ALLIEZ, D. COHEN-STEINER, AND M. DESBRUN. Designing Quadrangulations with Discrete Harmonic Forms. In *Eurographics/ACM SIGGRAPH Symposium on Geometry Processing 2006*, pages 201–210, 2006.
- [Walder06] B. SCHÖLKOPF WALDER, C. AND O. CHAPELLE. Implicit Surface Modelling with a Globally Regularised Basis of Compact Support. *Computer Graphics Forum*, 25(3), September 2006. Proceedings of EUROGRAPHICS 2006.
- [Wand07] M. WAND, P. JENKE, Q. HUANG, M. BOKELOH, L. GUIBAS, AND A. SCHILLING. Reconstruction of Deforming Geometry from Time-varying Point Clouds. In *Fifth Eurographics Symposium on Geometry Processing (SGP '07)*, pages 49–58, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [Wei00] L.-Y. WEI AND M. LEVOY. Fast Texture Synthesis using Tree-structured Vector Quantization. In *Proceedings of ACM SIGGRAPH 2000*, pages 479–488, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.
- [Weiss06] B. WEISS. Fast Median and Bilateral Filtering. *ACM Transactions on Graphics*, 25(3):519–526, 2006.

- [Wendland95] H. WENDLAND. Piecewise Polynomial, Positive Definite and Compactly Supported Radial Functions of Minimal Degree. *Advances in Computational Mathematics*, 4:389–396, 1995.
- [Wendland02] H. WENDLAND. Fast Evaluation of Radial Basis Functions: Methods based on Partition of Unity. *Approximation Theory X: Wavelets, Splines, and Applications*, pages 473–483, 2002.
- [Wendland04] H. WENDLAND. *Scattered Data Approximation*. Cambridge Monographs on Applied and Computational Mathematics (No. 17). Cambridge University Press, 2004.
- [Weyrich04] T. WEYRICH, M. PAULY, S. HEINZLE, R. KEISER, AND S. SCANDELLA. Post-processing of Scanned 3D Surface Data. In *Eurographics Symposium on Point-Based Graphics*, pages 85–94, Zurich, Switzerland, June 2004.
- [Wu95] Z. WU. Compactly Supported Positive Definite Radial Functions. *Advances in Computational Mathematics*, 4:283–292, 1995.
- [Xie04] H. XIE, K. T. MCDONNELL, AND H. QIN. Surface Reconstruction of Noisy and Defective Data Sets. *IEEE Visualization 2004*, 2004.
- [Yaroslavsky85] L. P. YAROSLAVSKY. *Digital Picture Processing. An Introduction*. Springer Verlag, Berlin, Heidelberg, 1985.
- [Yoshizawa05] S. YOSHIZAWA, A. BELYAEV, AND H.-P. SEIDEL. Fast and Robust Detection of Crest Lines on Meshes. In *ACM Symposium on Solid and Physical Modeling 2005*, pages 227–232, New York, NY, USA, 2005. ACM.
- [Yoshizawa06] S. YOSHIZAWA, A. BELYAEV, AND H.-P. SEIDEL. Smoothing by Example: Mesh Denoising by Averaging with Similarity-based Weights. In *Proceedings of Shape Modeling International*, pages 38–44, 2006.
- [Zayer05] R. ZAYER, C. RÖSSL, AND H.-P. SEIDEL. Discrete Tensorial Quasi-Harmonic Maps. In *Shape Modeling International 2005 (SMI 2005)*, pages 276–285, Cambridge, MA, U.S.A., 2005.
- [Zhang03] L. ZHANG, B. CURLESS, AND S. M. SEITZ. Spacetime Stereo: Shape Recovery for Dynamic Scenes. In *Computer*

Vision and Pattern Recognition (CVPR) 2003, pages 367–374, June 2003.

- [Zhang04] L. ZHANG, N. SNAVELY, B. CURLESS, AND S. M. SEITZ. Spacetime Faces: High Resolution Capture for Modeling and Animation. *ACM Transactions on Graphics*, 23(3):548–558, 2004. Proceedings of SIGGRAPH 2004.
- [Zwicker01] M. ZWICKER, H. PFISTER, J. VAN BAAR, AND M. GROSS. Surface Splatting. In *Proceedings of SIGGRAPH 2001*, pages 371–378, New York, NY, USA, 2001. ACM Press.