

UNIVERSITAT POLITÈCNICA DE CATALUNYA

Ph.D. Dissertation

Surface Reconstruction for Multi-View Video

Jordi Salvador Marcos

Thesis supervisor:

Dr. Josep R. Casas Pla

Department of Signal Theory and Communications Universitat Politècnica de Catalunya

Barcelona, April 2011

ACTA DE QUALIFICACIÓ DE LA TESI DOCTORAL

Reunit el tribunal integrat pels sota signants per jutjar la tesi doctoral:

Autor de la tesi:	

Acorda atorgar la qualificació de:

No apte
Aprovat
Notable
Excel·lent
Excel·lent Cum Laude

Barcelona, de/d'..... de

El President

El Secretari

(nom i cognoms)

(nom i cognoms)

El vocal

El vocal

El vocal

(nom i cognoms)

(nom i cognoms)

(nom i cognoms)

A la meva estimada Sibylle

Resum

questa tesi presenta diferents tècniques per a la definició d'una metodologia per obtenir una representació alternativa de les seqüències de vídeo capturades per sistemes multi-càmera calibrats en entorns controlats, amb fons de l'escena conegut.

Com el títol de la tesi suggereix, aquesta representació consisteix en una descripció tridimensional de les superfícies dels objectes de primer pla. Aquesta aproximació per la representació de les dades multi-vista permet recuperar part de la informació tridimensional de l'escena original perduda en el procés de projecció que fa cada càmera.

L'elecció del tipus de representació i el disseny de les tècniques per la reconstrucció de l'escena responen a tres requeriments que apareixen en entorns controlats del tipus *smart room* o estudis de gravació, en què les seqüències capturades pel sistema multi-càmera són utilitzades tant per aplicacions d'anàlisi com per diferents mètodes de visualització interactius.

El primer requeriment és que el mètode de reconstrucció ha de ser *ràpid*, per tal de poder-ho utilitzar en aplicacions interactives. El segon és que la representació de les superfícies sigui *eficient*, de manera que en resulti una compressió de les dades multi-vista. El tercer requeriment és que aquesta representació sigui *efectiva*, és a dir, que pugui ser utilitzada en aplicacions d'anàlisi, així com per visualitació.

Un cop separats els continguts de primer pla i de fons de cada vista –possible en entorns controlats amb fons conegut–, l'estratègia que es segueix en el desenvolupament de la tesi és la de dividir el procés de reconstrucció en dues etapes. La primera consisteix en obtenir un *mostreig* de les superfícies (incloent orientació i textura). La segona etapa proporciona superfícies tancades, contínues, a partir del conjunt de mostres, mitjançant un procés d'*interpolació*.

El resultat de la primera etapa és un conjunt de punts orientats a l'espai 3D que representen localment la posició, orientació i textura de les superfícies visibles pel conjunt de càmeres. El procés de mostreig s'interpreta com un procés de cerca de posicions 3D que resulten en correspondències de característiques de la imatge entre diferents vistes. Aquest procés de cerca pot ser conduït mitjançant diferents mecanismes, els quals es presenten a la primera part d'aquesta tesi.

La primera proposta és fer servir un mètode basat en les imatges que busca mostres de superfície al llarg de la semi-recta que comença al centre de projeccions de cada càmera i passa per un determinat punt de la imatge corresponent. Aquest mètode s'adapta correctament al cas de voler explotar foto-consistència en un escenari estàtic i presenta característiques favorables per la seva utilizació en *GPUs* –desitjable–, però no està orientat a explotar les redundàncies temporals existents en seqüències multi-vista ni proporciona superfícies tancades.

El segon mètode efectua la cerca a partir d'una superfície inicial mostrejada que tanca l'espai on es troben els objectes a reconstruir. La cerca en direcció inversa a les normals –apuntant a l'interior– permet obtenir superfícies tancades amb un algorisme que explota la correlació temporal de l'escena per a l'evolució de reconstruccions 3D successives al llarg del temps. Un inconvenient d'aquest mètode és el conjunt d'operacions topològiques sobre la superfície inicial, que en general no són aplicables eficientment en GPUs.

La tercera estratègia de mostreig està orientada a la paral·lelització -GPUi l'explotació de correlacions temporals i espacials en la cerca de mostres de superfície. Definint un espai inicial de cerca que inclou els objectes a reconstruir, es busquen aleatòriament unes quantes mostres llavor sobre la superfície dels objectes. A continuació, es continuen buscant noves mostres de superfície al voltant de cada llavor –procés d'expansió– fins que s'aconsegueix una densitat suficient. Per tal de millorar l'eficiència de la cerca inicial de llavors, es proposa reduir l'espai de cerca, explotant d'una banda correlacions temporals en seqüències multi-vista i de l'altra aplicant multi-resolució. A continuació es procedeix amb l'expansió, que explota la correlació espacial en la distribució de les mostres de superfície.

A la segona part de la tesi es presenta un algorisme de mallat que permet interpolar la superfície entre les mostres. A partir d'un triangle inicial, que connecta tres punts coherentment orientats, es procedeix a una expansió iterativa de la superfície sobre el conjunt complet de mostres. En relació amb l'estat de l'art, el mètode proposat presenta una reconstrucció molt precisa (no modifica la posició de les mostres) i resulta en una topologia correcta. A més, és prou ràpid com per ser utilitzable en aplicacions interactives, a diferència de la majoria de mètodes disponibles.

Els resultats finals, aplicant ambdues etapes –mostreig i interpolació–, demostren la validesa de la proposta. Les dades experimentals mostren com la metodologia presentada permet obtenir una representació ràpida, eficient –compressió– i efectiva –completa– dels elements de primer pla de l'escena.

ii

Summary

his thesis presents different techniques for the definition of a methodology for obtaining an alternative representation of video sequences captured by calibrated multi-camera systems in controlled environments, with known scene background.

As suggested by the title of the thesis, this representation consists in a threedimensional description of the surfaces of foreground objects. This approach for the representation of multi-view data allows for the recovering of part of the threedimensional information of the original scene lost int the projection process happening in each camera.

The choice of a type of representation and the design of the techniques for 3D scene reconstruction are driven by three requirements that appear in controlled environments such as a *smart room* or recording studios. In these scenarios, video sequences captured by a multi-camera rig are used both for analysis applications and interactive visualization methods.

The first requirement is that the reconstruction method must be *fast* in order to be usable in interactive applications. The second one is that the surface representation must be *efficient*, providing a compression of the multi-view data. The third requirement is that this representation must be *effective* or, in other words, that it provides the relevant information to be used for analysis applications as well as for visualization.

Once separated foreground and background elements from each view –possible in controlled environments with known background–, the strategy that is followed in the development of the thesis is that of dividing the reconstruction process in two stages. The first one consists in obtaining a *sampling* of the foreground surfaces (including orientation and texture). The second stage provides closed, continuous surfaces from the set of samples, through an interpolation process.

The result of the first stage is a set of oriented points in 3D space that locally represent the position, orientation and texture of surfaces that are visible from multi-camera settings. The sampling process is interpreted as a search for 3D positions that result in feature matchings between different views. This search process can be driven by different mechanisms, which are presented in the first part of this thesis.

The first proposal is to use an image-based method that searches surface samples along the half-line which starts at each camera's center of projections and goes through a certain point in the corresponding image. This method is well-suited to exploiting photo-consistency in a static scenario and presents favorable features for its usage in GPUs –desirable–, but it is not oriented towards the exploitation of temporal redundancies in multi-view sequences, nor provides closed surfaces.

The second method performs the search starting from an initial sampled surface which encloses the space where the objects to be reconstructed lie. The search, in the inverse direction of the normals –pointing towards its interior–, allows to obtain closed surfaces with a algorithm which exploits the correlation in the temporal sequence of reconstructions of the scene. A drawback of this method is the set of topological operations over the initial surface, which are in general not efficiently applicable in GPUs.

The third sampling strategy is oriented towards parallelization -GPU- and the exploitation of spatial and temporal correlations in the search of surface samples. Defining an initial search space that includes the objects to be reconstructed, some seed samples are randomly searched, which lie on the surfaces of the objects. Next, new surface samples are searched for around each seed –expansion process– until a sufficient density is achieved. In order to improve the search efficiency for seed samples, a reduction of the search space is introduced, either by exploiting temporal correlations in multi-view sequences or by applying multi-resolution. Next, the algorithm proceeds with the expansion, which exploits the spatial correlation in the distribution of surface samples.

In the second part of the thesis, a meshing algorithm is presented, which allows for the interpolation of the surfaces between its samples. Starting by an initial triangle, which connects the points coherently oriented, an iterative expansion of the surface over the complete set of samples takes place. Compared to the *state-ofthe-art*, the proposed method presents a very accurate reconstruction (it does not modify the position of samples) and results in a correct topology. Furthermore, it is fast enough as to be usable in interactive applications, as a difference with the majority of available methods.

The final results, applying both stages –sampling and interpolation–, reflect the validity of the proposal. Experimental data show how the presented methodology permits obtaining a fast, efficient –compression– and effective –complete– representation of the foreground elements of a scene.

Contents

Li	List of Figures xi			
1	Intr	oducti	on	1
	1.1	Contex	κt	2
		1.1.1	Challenges	3
	1.2	Object	zives	5
	1.3	Struct	ure	5
2	Stat	te of tł	ne Art	7
	2.1	Multi-	View 3D Reconstruction	8
		2.1.1	3D reconstruction taxonomy	8
		2.1.2	Visual hull	11
		2.1.3	Photo hull	14
		2.1.4	Multi-view stereo	14
	2.2	Multi-	View Video Reconstruction	15
		2.2.1	Scene flow	16
		2.2.2	Surface tracking	17
	2.3	Meshi	ng Algorithms	19
		2.3.1	Propagation-based methods	19
		2.3.2	Marching cubes-based methods $\ldots \ldots \ldots \ldots \ldots \ldots \ldots$	20
	2.4	Conclu	sions	21
3	App	oroach		25
	3.1	Surfac	e Description	26
		3.1.1	Sample-based surface description	27
		3.1.2	Polygon mesh surface description	29
	3.2	Metho	dology	31
		3.2.1	Surface sampling	32
		3.2.2	Surface interpolation	34
	3.3	Thesis	Organization	35
		3.3.1	Part I	35
		3.3.2	Part II	36

		3.3.3	Part III	36
Ι	Su	rface	Sampling	37
4	Ima	ge-bas	sed Surface Sampling	41
	4.1	Motiv	ation	41
		4.1.1	Related work	42
		4.1.2	Problem statement	43
	4.2	Photo	-consistency Test	44
	4.3	Sampl	ling of Visible Photo-consistent Surfaces	46
		4.3.1	Automatic neighborhood selection	47
		4.3.2	Extraction of candidate surface points	47
		4.3.3	Global photo-consistency	48
	4.4	Orient	tation Estimation	49
		4.4.1	Orientation in image-based sampling	51
	4.5	Conclu	usions	51
5	Sur	face Sa	ampling by Deformation	53
	5.1	Motiv	ation	53
		5.1.1	Related work	54
		5.1.2	Strategy	55
	5.2	Initial	ization: Static Reconstruction	55
		5.2.1	Silhouette constraints	57
		5.2.2	Orientation estimation	59
	5.3	Tracki	ing: Dynamic Reconstruction	61
		5.3.1	Surface expansion	62
		5.3.2	Spatial regularization	62
	5.4	Photo	-consistency Constraints	64
		5.4.1	Visibility estimation	65
		5.4.2	Segment shrinking	66
	5.5	Conclu	usions	66
6	Stat	tistical	Surface Sampling	69
	6.1	Motiv	ation	69
		6.1.1	Random sampling in multi-view environments	70
		6.1.2	Silhouette-consistent surface reconstruction $\ldots \ldots \ldots$	71
	6.2	Initial	ization	72
		6.2.1	Surface point validation	73
		6.2.2	Local normal estimation $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$	73
	6.3	Propa	gation	74
		6.3.1	Propagation region	75
		6.3.2	Iterative propagation algorithm	77

CONTENTS

	6.4	Efficient Sampling Schemes	7
		6.4.1 Dynamic sampling	8
		6.4.2 Multi-resolution sampling	9
	6.5	Surface post-processing 8	1
		6.5.1 Fine normal estimation 8	1
		6.5.2 Anisotropic smoothing	2
		6.5.3 Surface coloring	3
	6.6	Conclusions	4
7	Surf	face Sampling Results 8	7
	7.1	Image-based Surface Sampling 8	7
		7.1.1 Qualitative validation	8
		7.1.2 Quantitative results $\ldots \ldots \ldots \ldots \ldots \ldots \ldots $	0
	7.2	Surface Sampling by Deformation	2
		7.2.1 Static vs. dynamic reconstruction	3
		7.2.2 Usage of resources	4
		7.2.3 Accuracy	5
	7.3	Statistical Surface Sampling	6
		7.3.1 Multi-resolution and dynamic sampling	8
		7.3.2 Efficiency in multi-camera environments 9	9
	7.4	Comparison of the Three Sampling Approaches 10	0
	7.5	Conclusions	4
	a		
11	Sı	urface Interpolation 10 ⁷	7
8	Surf	face Interpolation 11	1
	8.1	Motivation	1
		8.1.1 Interpolation	2
		8.1.2 Related work	4
	8.2	Methodology 11	8
		8.2.1 Spatial queries	9
		8.2.2 Half edges	9
	8.3	Algorithm	1
		8.3.1 Initial triangle	1
		8.3.2 Propagation	2
	8.4	Conclusions	7
9	Surf	face Interpolation Results 12	9
			9
	9.1	Methods	
	9.1	9.1.1 Propagation methods 13	0
	9.1	Methods 12 9.1.1 Propagation methods 13 9.1.2 Marching cubes-based methods 13	0 0

			101
		9.2.1 Datasets	131
		9.2.2 Metrics	132
	9.3	Results	132
		9.3.1 Quantitative comparison	133
		9.3.2 Qualitative comparison	134
	9.4	Conclusions	142
II	IC	Overall Results and Conclusions 1	.43
10	Ove	rall Method: Sampling and Interpolation Results	145
	10.1	Evaluation	145
		10.1.1 Data and methodology	146
		10.1.2 Experiments	146
	10.2	Validation	149
		10.2.1 Data and methodology	149
		10.2.2 Experiments	151
	10.3	Conclusions	174
11	Cor	nelusions .	177
11	11 1	Summary of Contributions	177
	11.1	11.1.1. Surface sampling	178
		11.1.2 Surface interpolation	170
		11.1.2 Surface interpolation	180
	11.9	Future Work	181
	11.2	11.2.1 Surface compling	101
		11.2.1 Surface intermoletion	101
		11.2.2 Surface interpolation	102
	11 2	Perspectives	102
	11.5	reispectives	100
	-		~ ~
Iν	/	Appendices 1	.85
A	Can	nera Model	187
	A.1	Extrinsic Parameters	187
		A.1.1 Camera position	188
		A.1.2 Camera orientation	188
	A.2	Intrinsic Parameters	189
		A.2.1 Pinhole camera model	189
		A.2.2 Lens distortion	189
	A.3	Back-Projection	190
		A.3.1 Known depth	191

CONTENTS

	A.4	Epipolar Geometry)1
		A.4.1 Derivation of the fundamental matrix 19)2
в	kd-	Tree 19	13
	B.1	Construction)3
		B.1.1 Dynamic addition of elements)5
	B.2	Queries	95
		B.2.1 Nearest-neighbors search	96
		B.2.2 Range search)7
	B.3	Complexity	97
		B.3.1 High-dimensional data)7
С	Hau	Isdorff Distance 19	9
	C.1	Definition	99
		C.1.1 Properties)0
	C.2	Motivation)0
	C.3	Applications)1
		C.3.1 Polygon mesh dissimilarity 20)1
D	Rela	ated Publications 20)3
	D.1	Surface Sampling)3
	D.2	Surface Interpolation and Applications)4
Bi	bliog	graphy 20)7

$\mathbf{i}\mathbf{x}$

List of Figures

1.1	3D reconstruction after projection losses	3
2.1	$Visual Hull$ as a result of a shape-from-silhouette algorithm $\ldots \ldots$	8
2.2	Comparison between Visual Hull and Photo Hull	9
2.3	Level sets of a voxelized solid cube	11
2.4	Scene flow as a 3D dense motion field	16
2.5	Meshing of a set of points	19
3.1	Surface samples, describing position, orientation and texture	27
3.2	Triangle mesh containing faces and vertices	29
3.3	Right-hand rule	30
3.4	3D surface reconstruction as a two-stage approach $\hfill \ldots \ldots \ldots \ldots$	32
4.1	NCC in small-baseline setups	44
4.2	Distortion measurements in RGB space	46
4.3	Photo-consistent candidates with increasing depth	47
4.4	Space carving – rejection	48
4.5	Space carving – acceptance	49
5.1	Dilation of a surface and its enclosed space	54
5.2	Surface sampling by deformation. Initialization and tracking	56
5.3	Topological operations for deformation of a closed surface	58
5.4	Events related to the frustum and silhouette from a viewpoint	59
5.5	Partial silhouette and frustum constraints for one view	60
5.6	Multi-view merging of partial frustum and silhouette constraints	60
5.7	Surface points with geometrically estimated orientation	61
5.8	Dilation of the samples of a surface	62
5.9	Need for spatial regularization	63
5.10	Dynamic surface sampling with color-encoded velocity field $\ .\ .\ .$.	65
6.1	Motivation to the use of statistical sampling	71
6.2	Rectangular propagation region for statistical sampling	76
6.3	Propagation of surface samples in statistical sampling	78

6.4	Efficiency improvement in dynamic statistical sampling	79
6.5	Multi-resolution as a reduction of the search space	81
6.6	Normal estimation from a local set of surface points	82
6.7	Anisotropic smoothing of surface samples	83
6.8	Choice of the best oriented cameras for coloring	84
7.1	Image-based sampling overview	88
7.2	Image-based sampling, qualitative results	89
7.3	Image-based sampling vs. voxelized space carving	92
7.4	Dynamic sampling by surface deformation in five real sequences	94
7.5	Sampling by deformation vs. voxelized shape-from-silhouette $\ . \ . \ .$	95
7.6	Statistical sampling with multi-resolution scouting comp. time $\ . \ . \ .$	98
7.7	Statistical sampling for different levels of image decimation $\ . \ . \ .$	99
7.8	Statistical sampling efficiency in large multi-view settings $\ldots \ldots$	100
7.9	Statistical sampling with increasing number of views: samples	101
7.10	Statistical sampling with increasing number of views: meshes $\ . \ . \ .$	102
7.11	Surface obtained with each of the sampling strategies	104
8.1	Range and nearest-neighbor searches	118
8.2	Creation of a half-edge structure	120
8.3	Rules for initial triangle creation	122
8.4	Additional rules for triangle propagation	123
8.5	Surface propagation over a set of oriented points	124
8.6	Non-manifold edge detection and correction	126
8.7	Topological correctness of the output mesh	127
9.1	Surface interpolation: <i>drill</i> dataset	135
9.2	Surface interpolation: <i>bunny</i> dataset	137
9.3	Surface interpolation: armadillo dataset	138
9.4	Surface interpolation: hand dataset	139
9.5	Surface interpolation: dragon dataset	140
9.6	Surface interpolation: <i>happy</i> dataset	141
10.1	Surfaces obtained by the baseline system and the proposed one	148
10.2	Topological correctness of the mesh from the proposed method $\ . \ .$.	149
10.3	Free-viewpoint video application from the proposed methodology $\ . \ .$	150
10.4	Results from sequence <i>dancer</i>	153
10.5	Results from sequence <i>antoine</i>	155
10.6	Results from sequence antoine assis ballon	156
10.7	Results from sequence antoine edmond ballon	157
10.8	Results from sequence <i>antoine lucie</i>	158
10.9	Results from sequence <i>antoine roue</i>	159
10.1	0 Results from sequence <i>lucie corde</i>	160

10.11 Results from sequence benjamin baton
10.12 Results from sequence coup de pied 1 $\ldots \ldots $
10.13 Results from sequence coup de pied 2 $\ldots \ldots $
10.14 Results from sequence parade coup de point 1 $\ldots \ldots \ldots \ldots \ldots 164$
10.15 Results from sequence parade coup de point 2 $\ldots \ldots \ldots \ldots \ldots 165$
10.16 Results from sequence saisie col \ldots
10.17 Results from sequence saisie poigne
10.18 PSNR vs. number of surface samples, $Dancer$ dataset \hdots
10.19 PSNR vs. number of surface samples, Children dataset $\ .$ 169
10.20 PSNR vs. number of surface samples, Martial dataset 170
10.21 Surfaces reconstructed with growing number of surface samples 171
10.22 Criterion for deciding an optimal number of surface samples $\ . \ . \ . \ 172$
B.1 Construction of a kd -tree with maximum-variance splitting axis 194

Chapter 1

Introduction

ver the last years, analysis applications exploiting visual information have evolved from environments where scenes are captured by a single camera to setups where several cameras provide multi-view image sets. Newer imaging technologies provide also more advanced features, such as depth, but they still lack the interference-free, placing flexibility of standard video cameras.

Thus, the main advantage of multi-view scenarios is the availability of richer information about the actual 3D scene, when compared to imaging a scene from a single viewpoint. What is lost in the imaging process is, precisely, the 3D structure of the scene. This is due to the projection onto a 2D image on the camera sensor offering the perspective of a single view-point. Multi-view systems are able to recover 3D information from 2D projections in multiple viewpoints. This information allows for a richer representation of the scene from which both visualization and analysis applications can benefit. For example, objects affected by occlusions when being observed from a certain viewpoint can be observed without ambiguity by using data from other viewpoints. Furthermore, redundancies among different views contribute positively by enhancing the robustness of the scene representation, either by resolving inconsistencies in detected features or by adding 3D location information.

A practical disadvantage in a multi-view scenario is the management cost of the richer available data. An obvious problem in multi-view scenarios is that the storage –or transmission, in a streaming application– of a large number of video sequences that contain the projections of a certain scene from each of the available viewpoints is not negligible. Indeed, if a very large number of views is considered, the situation may become intractable. An important observation is that this cost should be considered unjustified, because the multiple sequences captured by each camera present redundancies –the larger number of available views, the more redundant the multi-view data is–.

The target of this thesis is to present an efficient method for representing this multi-view information, at the expense of a certain, reasonable computational cost, providing an alternative representation for visualization, analysis, transmission and storage. The computational resources required by such an efficient representation should not increase linearly with the number of views, regarding the information redundancy of multi-view data. Another requisite of such a representation is that it allows to go back to the original views, given the fact that it conveys all the information contained in the latter.

This representation is the result of the 3D reconstruction methodology that is proposed in this thesis. The main idea is not only to reduce the storage and transmission costs related to the multi-view data, but also to do that in an efficient way, such that it is adequate to be used by both visualization and analysis applications. Regarding this last point, the 3D representation of the multi-view data should enrich the possibilities of such applications. Some features -e.g. surface normals– are not easily computable from the multiple views, whereas their computation can be a relatively simple task given a suitable 3D representation. To sum up, using a 3D representation we obtain efficiency (compression) and effectiveness (adaptation to visualization and analysis) in order to represent the rich information available in multi-view setups.

1.1 Context

The target scenarios of this thesis are controlled environments. Examples of these are *smart rooms* –or *smart spaces*– and recording studios. In the former, a varying number of people or objects can interact with each other as well as with computer systems. This sets the adequate scenario for the development of services related to advanced human-computer interfaces (HCI) not requiring special attention by the user, as a the typical HCI display-keyboard-mouse would require. In the latter, recording studios, multi-view techniques can be exploited also for video production.

In these environments, video cameras -as well as microphone arrays and other capture equipment– are distributed around the scene in an almost non-invasive manner, usually in rather wide-baseline configurations. In many cases, the contents of the background of the scene can be assumed to be static, or are even removed $-chroma \ keying$ –. We consider this case, and focus on the foreground -i.e. the dynamic part– of the scene.

The projection process that takes place in each camera in order to generate a 2D image out of the actual 3D scene is an inherently lossy process, for most of the 3D information about the scene is lost. An image-based representation of the multi-view data provides cues that can be used to recover part of this 3D



Figure 1.1: 3D reconstruction as a recovering mechanism for the 3D cues lost in the imaging process. The result of the reconstruction process is an approximation to the actual geometry and texture of the 3D scene, from which only some surfaces are visible

information. However, it is the process of multi-view reconstruction the one that explicitly attempts to recover the information lost in the projection process in order to provide a richer representation, at least of some parts of the scene. This concept is illustrated in Fig. 1.1.

Hence, the adopted approach is that of obtaining a unique, 3D representation of the foreground contents of the scene captured by such a multi-camera rig. In more detail, the main subject that this thesis addresses is the *multi-view reconstruction* of the foreground elements of scenes in the aforementioned scenarios with known background. The purpose of such 3D representation is two-fold, since it is to be used by different algorithms exploiting the enriched information obtained by multi-camera rigs for enriched visualization -free-viewpoint video [Lipski et al., 2010]– and also for analysis of the contents of the scene.

1.1.1 Challenges

Solutions to several problems in multi-view reconstruction have already been proposed, which provide different levels of accuracy depending on the amount and nature of the available information. Thus, whereas there exist techniques able to compete with range scanners in terms of accuracy [Seitz et al., 2006], they require a small enough baseline in order to determine the 3D structure of the scene being reconstructed. On the other extreme, there also exist techniques which, although being able to obtain 3D representations from a small number of viewpoints in a much wider baseline configuration, they are limited by the accuracy or the robustness with respect to inconsistencies in the input data they can provide.

In this thesis, we discuss the suitability of a surface-based representation of the 3D data as opposed to a volumetric one in terms of its efficiency at describing the

visible features in multi-view scenes. Continuous surface representations have been widely used in computer graphics, resulting in powerful graphics hardware able to handle triangle mesh representations of 3D surfaces. Also sample-based descriptions of surfaces can be considered when the sampling density is high enough and an explicit description of the topology is not required.

Few methods are available -EPVH in [Franco and Boyer, 2003] being perhaps the clearest example– which are capable of obtaining a triangle mesh directly from image-based features, without inputting explicit 3D estimates to the location of some surface points. However, many techniques exist –*e.g. Poisson reconstruction* [Kazhdan et al., 2006]–, which are capable of robustly obtaining triangle meshes from sets of oriented surface points. One drawback of the approach consisting in computing the 3D position of surface points and obtaining a triangle mesh by using such a technique is that the computation time for the mesh might be too large for being used on-line.

Whereas the former approach has already been used for real-time operation on adequate hardware configurations, the latter seems more suitable in order to provide a surface description useful for both analysis and visualization. This is due to the more uniform distribution of the size of the triangle faces composing the resulting surface description and the higher density of vertices, which might introduce an additional rendering cost but provides richer information about local features of the surface -e.g. orientation or curvature–.

It is also interesting to note that most works in 3D reconstruction target static scenes, *i.e.*, scenes where no more than one single frame is captured –for multi-view settings using multi-camera rigs– or scenes where a static object is captured at different time instants from different viewpoints. Few authors, [Vedula et al., 1999, Starck and Hilton, 2005] among others, have addressed the problem of exploiting the redundancies between consecutive time instants in video sequences to either improve the quality or reduce the computational cost of 3D reconstruction.

There are three main challenges for the 3D reconstruction methodology proposed in this thesis. The first one is that its result must be able to be used as a replacement of the multi-view data without introducing a significant loss of information. The second one is that it has to be efficient in terms of its representation, resulting in a compression of the multi-view data. Finally, it has to be effective in order to facilitate the posterior use of the data in visualization and analysis applications.

Going slightly more in detail, some challenges regarding the design of new algorithms for surface-based 3D reconstruction can also be detected. These algorithms must provide a sufficient level of accuracy –projection of the 3D shape vs. original image– without restrictions to the baseline. In second place, it is necessary to obtain a throughput such that the reconstructed surface can be used for real-time

 $\mathbf{4}$

visualization and analysis applications and, in third place, it is also necessary, in terms of efficiency and effectiveness of the approach, to adapt the 3D reconstruction to the case of multi-view video sequences in order to exploit temporal correlations.

1.2 Objectives

Regarding the previously described challenges in the context of 3D reconstruction of the surfaces of foreground objects in multi-view video sequences, the main objective of this thesis is to provide a unique surface-based representation of the relevant multi-view data under the aforementioned conditions:

- The resulting surface representation must be efficient –compressing the redundancies in multi-view information–, effective –suitable for visualization and analysis– and accurate –minimizing losses with respect to the original data–.
- It must be able to retrieve a close approximation to the original images –with minimal losses– by re-projection of the unique surface-based representation and addition of the background information.
- It should exploit both spatial and temporal correlation in multi-view video sequences in order to reduce the computation time and make it suitable for real-time operation.
- Finally, another goal is to obtain an efficient, continuous representation of the surface –a triangular mesh– in a usable computation time. Such a representation can contribute to further compress the multi-view data and can be used for analysis as well as for rendering using standard graphics pipelines.

As a summary, the ultimate goal of this thesis is to obtain a fast -towards real-time-, efficient -within a compressed/limited support-, effective -exploitable by forthcoming applications- and unique surface-based representation of each frame in multi-view video sequences, which provides all the relevant information about foreground objects contained in the original streams and can be used for both interactive visualization and analysis applications.

1.3 Structure

Following this brief introduction to the challenges that are tackled in this thesis and the goals that are to be achieved, the reader will find a review of the related state of the art in Chapter 2. Next, Chapter 3 presents the approach that is followed throughout this thesis, which is divided in two main concepts.

The first main concept is the dense sampling of surfaces from multi-view data –3D reconstruction–. This is presented in Part I as an evolutionary collection of different techniques, each with its advantages and its drawbacks.

The second main concept is the interpolation of surfaces from a dense set of samples –meshing–. The proposed algorithm for this purpose and a comparison to other methods in the literature are presented in Part II.

In Part III, a validation of the objectives of the thesis is presented. These results are followed by the final discussion about the achievement of the proposed goals and the resulting contributions.

Chapter 2

State of the Art

pplications derived from research in the exploitation of the rich data available in multi-view scenarios range from 3D reconstruction of monuments and buildings [Mueller et al., 2004] to volumetric computer tomography [Sardarescu et al., 2008]. Computer vision algorithms benefit of the enriched information available in such scenarios, as well. In general, visual analysis from multiple views exploits the similarity and disparity features present in the available projections of the scene.

Similarity features (redundancies) that can be identified from different viewpoints allow for the establishment of correspondences for the registration of the images in a common geometric framework. This allows for an easier detection and analysis of these redundancies in the several views. On the other hand, the disparity features that can be detected in the multiple views contribute to augment the visual information available through the inclusion of visual elements belonging to areas that are hidden in one of the views, thus disambiguating occlusions.

One of the possibilities when working in multi-view scenarios consists in reconstructing the 3D shapes of elements in a scene that has been captured by a multi-camera setting. This approach appears in opposition to working in an imagebased manner, using geometric relationships between pairs of images described by epipolar constraints. Therefore, multi-view approaches can be classified in:

- **3D** reconstruction. Three-dimensional shapes are computed from the multiview data, with several possibilities for their representation, each providing different advantages and drawbacks, as seen below.
- Image-based multi-view. Analysis is performed directly on the captured images and multi-view cues are exploited by considering epipolar constraints that can be computed from *fundamental matrices* (Appendix A).



Figure 2.1: Visual Hull of a set of silhouettes. Only foreground/background classification information is used to reconstruct the 3D shape

Thus, a detailed description of the visual data in multi-view settings, for applications ranging from visualization to analysis -e.g. object manipulation or visual recognition–, might be based on the reconstruction of 3D models of the environment and of the objects present in it. Such reconstruction will be more precise as more cameras observe the space where the scene is located.

We focus this dissertation on 3D reconstruction to obtain a compact and meaningful representation of multi-view data. This can be done efficiently and contribute to the state of the art. Besides, this representation can be used to freely choose the viewpoint for visualization applications. Real-time analysis applications can also benefit of the enriched information available in settings with an increasing number of viewpoints, without necessarily increasing the computational load as would result in image-based scenarios, involving comparisons or consistency checks. Instead, a unique, compact representation of the visual cues provides a generic framework that isolates visualization applications or high-level analysis from the low-level details of the multi-view capture process.

2.1 Multi-View 3D Reconstruction

In contrast with image-based multi-view approaches, 3D reconstruction methods provide a representation with a 3D support that can be used to explain in a single, compact structure the multi-view cues from the input images with an arbitrary level of detail, limited by the baseline configuration of the multi-view setup.

2.1.1 3D reconstruction taxonomy

There are several techniques for computing 3D reconstructions from a set of images captured in multi-camera settings. These techniques can be classified accordingly



Figure 2.2: Tightness of the photo hull compared to the visual hull. In (a), the bold line represents the visual hull of a 2D scene (in this case, the central occupied square) reconstructed from a set of silhouettes (segments) available in a set of 1D cameras. In (b), the photo hull is computed and represented by the bold line, resulting in a tighter reconstruction of the actual shape, which is the colored square

to several differentiating properties. In the following, two possible classifications, each regarding a certain differentiating property, are listed.

Classification by type of input data. Depending on the type of input data, we can find different types of 3D reconstruction methods in the literature:

- Visual hull. In [Laurentini, 1994], the visual hull was introduced as the maximal 3D volume which is consistent with the silhouettes of an object projected onto a set of images. In this case, input data consist solely of binary silhouettes. This type of 3D reconstruction is depicted in Fig. 2.1.
- Photo hull. In [Kutulakos and Seitz, 1999], the *photo hull* was introduced as the maximal 3D volume which is photo-consistent with each image. In this case, input data are the captured color images for a given time instant. Due to the richer information available in the input data, the photo hull constitutes a tighter estimate of the actual volume than the visual hull, being its lower robustness, in comparison with that of the visual hull, a drawback to consider when attempting to use it for analysis applications. The tightness of the photo hull, compared to that of the visual hull, is illustrated in Fig. 2.2.
- Multi-view stereo. Similarly to the photo hull computation, in multi-view stereo [Goesele et al., 2006] the input images captured by the cameras are considered for reconstruction. However, in this case a small baseline camera setup is required in order to precisely compute the position of surface points that match a photo-consistency test between pairs of images. Typically, normalized cross-correlation is used as a photo-consistency test, either in grayscale or in the RGB color space.

Classification by representation support. The majority of 3D reconstruction algorithms in the related literature rely on one of the following representations:

• Volumetric. In many algorithms, for example in [Szeliski, 1993], a discrete occupancy function over a regularly sampled 3D grid -voxelization- is used to describe the 3D structure of the reconstructed scene. Whereas this support is useful to represent three-dimensional dense volumes, it appears as a poor descriptor of the visible data in multi-camera settings, which are the projections of object surfaces. In the cited reference, a progressive algorithm allows to obtain the occupancy function in homogeneous areas without having to visit every voxel of the highest-resolution regular grid individually (these are exclusively visited when needed).

In [Faugeras and Keriven, 2002], **level sets** were introduced. They are functions encoding distances of elements in a regularly sampled 3D grid to their closest surfaces. The advantage of methods based on the level set representation is that numerical computations involving curves and surfaces can be computed without having to parameterize these objects. Furthermore, methods based on level sets provide an effective manner of following shapes that change their topology, making level sets a convenient representation for modeling time-varying objects. An illustration of such representation is shown in Fig. 2.3.

• Surface-based. Due to their efficiency for tessellation of surfaces, polygon meshes are a popular output format for multi-view algorithms. They have also been used as the central representation by some authors, for example in [Isidro and Sclaroff, 2003, Franco and Boyer, 2003]. As a difference with the previous representation supports, polygon meshes describe only the geometry of surfaces, and not their interior. Thus, it is a representation better suited to the data available in multi-camera settings. An advantage with respect to the previous representations is that 3D space does not need to be discretized in order to obtain it.

Surface patches, also called surfels or oriented points in some references in the literature [Jia et al., 2006, Carceroni and Kutulakos, 2002], are sampling elements designed to describe 3D surfaces. As such, they basically consist in a 3D location and a 3D orientation descriptor. They constitute a suitable support for describing surfaces of the objects in a multi-view scene, since they efficiently describe the only features visible by the cameras.

• Image-based. A depth map representation consists in setting a depth value (distance to image) for each pixel in each available view. It has been used, for example, in [Kolmogorov and Zabih, 2002, Kauff et al., 2007]. This representation, compared to the others, lacks a 3D support, which makes its



Figure 2.3: Level sets of a voxelized solid cube. In this example, where a cut of a solid cube is shown, gray levels would encode the minimal distance of each voxel to the surface of the reconstructed cube

usage in analysis applications potentially more costly, due to the increasing computational demands of multi-view matching with elevate numbers of available input views. However, depth maps can be an efficient representation for classical stereo or tri-focal with narrow-baseline [Waizenegger et al., 2011].

Additionally, in [Casas and Salvador, 2006], a discrete occupancy function over an irregular 3D grid, sampled in a specific way such that it appears as adapted to the calibration geometry of the input images, is also shown as an intermediate step between volumetric and image-based representations because, like the latter, the image-adapted sampling scheme does not introduce an additional, arbitrary sampling of 3D space that would result in sampling artifacts like those appearing in volumetric techniques, whereas, like the former, it keeps a sampled representation that can potentially expand current algorithms for processing voxelized volumes.

2.1.2 Visual hull

As introduced above, the *visual hull* is the maximal volume consistent with a set of input silhouettes. Depending on the formation model of the input silhouettes, different techniques allow the estimation of the visual hull, each of them using one of the representation supports introduced above.

Shape from silhouette. Shape from silhouette (SfS) techniques require the previous computation of binary silhouettes of foreground objects in each of the available views. In this case, the computed silhouettes are considered noiseless, thus not needing for special care beyond that of geometrical constraints when checking consistencies between views. Reviewing the literature, we report the following approaches:

- Voxelized Shape from Silhouette. A multi-resolution technique based in *octrees* [Meagher, 1982, Szeliski, 1993] is a widely extended approach. In it, a coarse representation of the visual hull is obtained and, at subsequent iterations, the resolution of the representation is increased in parts of the volume that require higher sampling frequencies for obtaining a suitable representation without excessive aliasing (surfaces). Another technique, with fixed resolution for all voxels but computationally efficient in the definition of its projection test, consists in randomly sampling the actual projection of a voxel on the images, rather than using the whole *splat* onto which a voxel projects for a given point of view. This projection test is called SPOT (*Sparse Pixel Occupancy Test*), and was introduced in [Cheung et al., 2000].
- Implicit Surface Visual Hull. This technique attempts to achieve a higher reconstruction quality by the adoption of a specific adaptive sampling scheme [Erol et al., 2005] that, similarly to the octree-based voxelization, uses a finer resolution in those regions where it is needed. As a second step, it is accompanied by a post-processing aiming at obtaining crack-free polygonal surfaces, using *marching cubes* [Lorensen and Cline, 1987].
- Polyhedral Visual Hull. This technique directly computes the 3D surface of the visual hull and describes it as a polygon mesh. Different views of this approach can be found in [Matusik et al., 2001] and [Lazebnik et al., 2007], [Franco and Boyer, 2003] or [Franco and Boyer, 2009]. Furthermore, [Franco et al., 2006] introduces the visual shape concept. Visual shapes are a class of silhouette-consistent shapes including, but not limited to, the visual hull. They are useful when observing shapes with known properties, such as smoothness, because this prior knowledge of the geometrical structure can be used for retrieving shapes with an appearance closer to the actual ones.
- Compressed Sensing Reconstruction. The application of the *compressed* sensing theory [Donoho, 2006] to the field of 3D reconstruction allows obtaining the volumetric occupancy from a smaller number of random projections, thanks to the exploitation of the sparse structure of the resulting occupancy function. This approach allows for the fast reconstruction of the visual hull, rendering the computation time practically independent of the number of input views, as shown in [Reddy et al., 2008]. The reconstruction of the actual scene occupancy from the sparse set of samples requires the efficient computation of the so-called *Basis Pursuit* [Chen and Donoho, 1994], a widely studied linear programming problem, which is considered tractable under the data structure assumptions in the compressed sensing theory.

Shape from silhouette with enhanced robustness. In a more general case, silhouette formation must be considered noisy, due to limitations in silhouette ex-

traction algorithms. Hence, consistency tests between views and further processing must address this problem. There are several approaches in the literature that can be grouped as follows:

- Minimization of energy functions. An energy function is defined, including functionals based on the local neighborhood structures of three-dimensional elements and smoothing factors, like in [Kolmogorov and Zabih, 2004] and [Alcoverro and Pardàs, 2009]. Algorithms based on *graph cuts* allow to obtain a global minimum of the defined energy function ([Kolmogorov and Zabih, 2002]) with great computational efficiency. The optimal graph cuts are obtained by first constructing a directed graph with a reversed edge for each edge between two nodes. The maximum flow algorithm is executed on the resulting graph, with each node representing a binary variable, and the resulting labeling of the nodes represents the minimum cut of the graph derived from the defined energy function.
- Space occupancy grids and other Bayesian frameworks. With space occupancy grids [Franco and Boyer, 2005], each pixel is considered as an occupancy sensor, and the visual hull computation is formulated as a problem of fusion of sensors with Bayesian networks. This method allows to avoid obtaining error-prone binary silhouettes by, instead, computing an occupancy probability for each pixel that is fused with the probabilities of other pixels in other views in order to set the voxel occupancy probability. This voxel occupancy probability can be thresholded at a final step in order to set the binary occupancy values.
- Shape from Inconsistent Silhouette. For cases when silhouettes contain systematic errors, [Landabaso et al., 2006] proposes a framework to minimize the classification error of a voxel when inconsistent views are a minority. To achieve that, the *inconsistent hull* is obtained, representing those voxels that are potentially occupied although they do not belong to the visual hull due to inconsistencies between silhouettes. The *unbiased hull* computed from the inconsistent hull, which represents voxels that should be part of the reconstructed volume, is added to the visual hull, resulting in a reconstruction robust to systematic errors in the silhouette extraction process.

Additionally, in [Shin and Tjahjadi, 2008], the *local convex hull* is introduced. This technique addresses some problems related to discontinuities, derived from the surface triangulation via *marching cubes* of octree-based voxelized representations of visual hulls. It proposes slicing an initial reconstructed volume for, in subsequent steps, generating local convex surfaces that are combined to obtain a more robust estimate of the surfaces of foreground objects.

2.1.3 Photo hull

As introduced above, when, instead of binary silhouettes, color images captured by multi-camera settings are used for reconstructing a scene, the *photo hull* can be obtained, representing the largest photo-consistent volume. In the literature we can find two different approaches for obtaining the photo hull of a scene in arbitrarily wide-baseline setups:

- Voxel coloring. In the algorithm proposed in [Seitz and Dyer, 1997], starting with an empty volume, voxels are visited in a specific order that provides consistent visibility for all the cameras and added to the reconstructed volume if they pass a photo-consistency test. One of its drawbacks is the need for a certain order for visiting the voxels, in order to ensure visibility. Furthermore, it assumes surfaces to be Lambertian, *i.e.*, they only reflect the diffuse component of the ambient light, making the appearance of surfaces independent of the point of view. Unfortunately, this assumption, made by many algorithms exploiting photo-consistency, is in general not true in real cases, which makes such methods much less robust than silhouette-based ones.
- **Space carving.** As a difference with the previous one, this algorithm [Kutulakos and Seitz, 1999] consists in setting as occupied a whole initial volume, which includes the actual scene, and iteratively removing those voxels that, visited in a specific order (typically using *sweep planes*), do not pass a photo-consistency test. Again, the Lambertian surface assumption appears in this approach, which keeps this method from achieving the robustness of silhouette-based ones.

2.1.4 Multi-view stereo

When small-baseline setups are available, finer photo-consistency tests can be used that take into account data in a vicinity of each pixel, assuming that perspectives will be similar between views. This set of techniques can be considered as an extension of classical stereo vision to a multi-camera environment, and introduces the problem of the appearance of inconsistent estimates between different pairs of views.

Thus, different multi-view stereo techniques introducing different types of refinements and solutions to the consistency problem can be found in the literature:

• Depth-map extraction and fusion. In [Goesele et al., 2006], depth-maps are extracted from each view by finding correspondences with, at least, two neighbor views, and computing its correspondent 3D position by triangula-

tion. Then, a volumetric combination of the depth-maps and posterior surface extraction is applied in order to obtain the photo-consistent reconstructed surface.

- Energy minimization. Volumetric min-cut methods (that are obtained from the maximum flow algorithm for graphs with reversed edges) combine photo-consistency with shape smoothing factors, as in [Sinha and Pollefeys, 2005]. More recently, [Campbell et al., 2008] considers several hypothesis for the placement of the actual surface, and an optimization step using Markov Random Fields gives impressive results for highly textured scenes. Its main drawback is that the method is unable to infer the surface position for objects that are not highly textured. This algorithm can be viewed as an improvement over [Goesele et al., 2006] using the same working principle.
- Fusion with silhouettes. In this case, the robustness of shape from silhouette approaches is combined with the availability of richer features in multi-view stereo by minimizing a convex functional with the exact silhouette consistency imposed as a convex constraint that restricts the domain of admissible surfaces to those that fit in the binary silhouettes once projected onto the original viewpoints, as in [Kolev and Cremers, 2008]. The authors claim high-accurate, silhouette-consistent reconstructions for real-world data in a computationally efficient manner.
- Surface reconstruction from dense stereopsis. In [Furukawa and Ponce, 2007], an initial sparse set of surface elements (surfels, surface patches or oriented points) are obtained from matching image-based features of several input images. From this initial estimation, an iterative algorithm augments the number of surface elements and rejects outliers, achieving results with high accuracy, even for small image sets. At a post-processing step, a closed polygonal mesh is morphed until it fits the set of reconstructed surface elements with an energy minimization algorithm. This method has the drawback of a high computational cost, with its bottleneck in the iterative stage that increases the density of surface elements.

2.2 Multi-View Video Reconstruction

One of the fundamental properties of a scene is its motion. Exploiting the temporal correlations of a scene is a task in computer vision that can benefit of the enriched information available in multi-view scenarios. Depending on the methodology applied for exploiting this valuable feature of a scene, which complements and enriches the geometric description that can be obtained by some of the methods above, *scene*



Figure 2.4: Scene flow. Each vector represents the motion of a feature of the reconstructed shape from one time instant to the next one. Note that the reconstructed shape has rotated 90 degrees around its center and also translated. The scene flow implicitly describes both types of motion

flow and *surface tracking* approaches can be found in the literature, which present important differences on the nature of the data used for estimation.

2.2.1 Scene flow

Earlier approaches have separated the shape and motion estimation tasks, in a *scene flow* formulation of the problem. The scene flow concept was introduced as a 3D equivalent to the *optical flow* in monocular video sequences. In this case, input data consist in sequences of images captured by a multi-camera setting.

The techniques in the literature first pre-compute shape estimates at each time instant. Then, velocity fields are obtained by feature matching between reconstructions at different instants. The scene flow typically complements photo hulls or multi-view stereo surface reconstructions. These techniques are reviewed next. Fig. 2.4 shows how rigid motion can be represented accompanying an existing shape at different time instants in a toy example.

As seen in [Vedula et al., 1999], when many cameras are used to image a scene, it is necessary to apply some regularization or smoothing to compute the scene flow describing the non-rigid motion of the scene. Reviewing the literature, we find two different ways to compute the scene flow for non-rigid motion, regarding the way regularization is applied:

- Image-based regularization, as in [Vedula et al., 2002]. This approach is used with voxel coloring for obtaining the voxelized photo hull of the scene and its scene flow. In this case, smoothing is performed in each image plane and the smoothed 2D optical flows are combined for obtaining the 3D scene flow.
- **3D** regularization on 3D surfaces. In [Carceroni and Kutulakos, 2001], this approach is used for recovering the 3D shape, reflectance and non-rigid mo-
tion of scenes in a surfel-based representation, with the scene flow smoothing directly computed in 3D space. One of the most interesting points of this approach is the usage of a model for the appearance of the projections of surfaces in the scene not assuming Lambertian surfaces, although it requires the knowledge of the physical position of light sources.

Some works in the literature obtaining dense non-rigid scene motion estimates using a polygon mesh representation of surfaces, in a scene flow formulation of the motion estimation problem, can be found in [Ahmed et al., 2008] and [Klie et al., 2007]. The former receives as input a sequence of visual hull surfaces obtained from static surface reconstructions in each time instant. Then, sparse correspondences from SIFT features [Lowe, 2004] between adjacent frames are obtained. Finally, dense correspondences are inferred by finding matches at neighborhoods of each found correspondence. The latter [Klie et al., 2007] receives the same type of input and obtain a sparse 3D motion field from combining 2D optical flows from different views. Similarly to the previous case, sparse motion vectors are propagated to their vicinities in order to obtain a dense motion field. Dues to the noisy motion vectors at the output of the algorithm, temporal smoothing is applied in the form of low pass filters.

In [Starck and Hilton, 2005], a spherical model is constructed by mapping the surface obtained at each time instant to a sphere. In this case, surfaces are estimated by first concatenating voxel coloring and marching cubes [Lorensen and Cline, 1987], and then applying projective texturing [Mueller et al., 2004] in order to provide texture to the surfaces. As its authors claim, the spherical domain gives a consistent 2D parametrization of non-rigid surfaces for matching, which can be used in order to find the 2D bijection that guarantees a continuous one-to-one surface correspondence.

2.2.2 Surface tracking

More recent approaches exploit both spatial and temporal redundancies in sequences with motion by tracking surfaces from multi-view captures. Some of the existing techniques exploit multi-view data in a manner similar to multi-view stereo reconstruction, while others aim at obtaining shapes equivalent to photo hulls with wider-baseline scenarios in mind. Also for the latter scenario, silhouette-based approaches can be used to introduce precise deformations to an existing model in order to fit it to the silhouettes at each time instant. Other approaches aim at tracking objects in crowded scenes with very sparse cameras without color calibration. All these approaches are viewed in more detail next. Multi-view stereo-based. Impressive time-coherent reconstructions can be obtained by tracking surfaces reconstructed by multi-view stereo, at the cost of large computation times. In [Furukawa and Ponce, 2008, Furukawa and Ponce, 2009b], normalized cross-correlation is used as a photo-consistency measurement between local neighborhoods of vertices in the mesh surface being tracked and images captured at the next time instant with a detailed treatment of visibility. First, smoothing constraints are applied at a local scale, in order to capture both radial and tangential surface motion. Then, at a global scale, a second level of refinement is achieved by introducing both smoothing and rigidity constraints.

In [Courchay et al., 2009], a global variational approach is applied in order to jointly estimate shape and motion. Their method fixes the topology of the surface, represented as a mesh, at the initial instant and then recovers both shape and motion by optimizing the positions of vertices at each instant. The optimization is driven by an energy function that incorporates photo-consistency and smoothness terms for both the velocity field and the surface across time.

In [Goldluecke and Magnor, 2004], spatio-temporal 3D hyper-surfaces are introduced in an elegant notation able to combine data from several time instants in a sequence in order to provide smooth surface estimates over time. The dynamic geometry is volumetrically modeled as a 3D iso-surface in space-time. To create the model, the photo-consistency of the whole scene is optimized with temporal smoothness constraints, which are intrinsic to the *PDE*-based evolution method used for energy minimization. The intersection of the 3D hyper-surface with planes of constant time gives the 2D surface at each time instant.

Silhouette-based. In [de Aguiar et al., 2007], range data is used for creating a high resolution mesh model of an object to be tracked. Then, *Laplacian* deformations are applied in order to fit the surface of the model being tracked to the set of silhouettes in each time instant, obtaining impressive offline results. In its follow-up [de Aguiar et al., 2008], photo-consistency constraints from multi-view stereo are introduced in the framework in order to provide accurate mesh fitting to small scale shape detail.

In [Guan et al., 2008], color information is introduced in a way that can be used in very wide-baseline scenarios and without requiring color calibration between cameras. A *Gaussian Mixture Model* (GMM) model of each subject in the scene is automatically learned in each view when a new shape is found in the scene. A *Bayesian* method is used to infer the voxel occupancy, combining the silhouette information with the matching score to each multi-view GMM model. This approach does not provide dense motion fields, but it can be useful for tracking multiple objects in crowded scenes with sparse multi-view settings.



Figure 2.5: Meshing of a set of points. From left to right, an initial set of points and different surfaces represented as polygon meshes, obtained with increasing spatial resolution. Courtesy of Xavier Suau, extracted from [Suau et al., 2010]

2.3 Meshing Algorithms

Meshing algorithms provide a means to compute a closed, continuous surface out of a set of input points. Typically, the output surface is provided as a set of connected triangular facets. In Fig. 2.5, the data points from the *dragon* dataset in [The Stanford 3D Scanning Repository, 2010] are used to generate three different surfaces with increasing spatial resolution.

Regarding the literature in meshing algorithms from a set of oriented points, two families of methods can be considered. The first one comprises methods based on the propagation of an existing surface, whereas the second family comprises methods based on a volumetric reconstruction followed by marching cubes [Lorensen and Cline, 1987].

2.3.1 Propagation-based methods

Propagation methods for surface reconstruction have in common the initial definition of one or more regions that are used as seeds for a propagation algorithm that attempts to cover the whole surface between the input 3D points. Some of the methods available in the literature that fall in this category are:

- Ball pivoting [Bernardini et al., 1999]. It starts with a seed triangle and iteratively pivots a ball around an edge (*i.e.* it revolves around the edge while keeping in contact with edgepoints) until it touches another point, forming another triangle.
- Restricted and oriented propagation [Suau et al., 2010]. It starts by voxelizing the space where the surface lies and fills each voxel with an original point. Then, a specifically-designed propagation scheme connects close points in order to create new faces of the resulting surface.

Ball pivoting

The ball pivoting algorithm is an advancing-front algorithm to incrementally build an interpolating triangulation of a given point cloud. The method is conceptually simple. Starting with a seed triangle, it pivots a ball around each edge on the current mesh boundary until a new point is hit by the ball. The edge and point define a new triangle, which is added to the mesh, and the algorithm iteratively considers a new boundary edge for pivoting.

Restricted and oriented propagation

Restricted and oriented propagation is a method to reconstruct meshed surfaces from point sets based on a propagation through a voxelized space performed in order to find the closest neighbors of every data point. Propagation is done following a specific propagation pattern which exploits reciprocity in neighbor finding. Every obtained pair of neighbors defines a new edge of the output mesh.

2.3.2 Marching cubes-based methods

Marching cubes (MC) is a well-known method for extracting a polygonal mesh of an iso-surface from a 3D scalar field –regular voxelization–. A set of oriented points belonging to the surface of the objects –explicit surface representation– is a very common output of 3D reconstruction methods (the first part of [Furukawa and Ponce, 2007] constitutes an impressing multi-view stereo example.) It is possible to translate it to a volumetric representation, usually under certain conditions of spatial uniformity, in order to apply marching cubes at the last stage.

Therefore, some approaches exist which, first, determine voxel occupancy by means of a certain type of regularization and, then, extract the surface by means of MC. Some methods in the literature are:

- Robust implicit moving least squares-marching cubes [Oztireli et al., 2009]. It extracts the iso-surface of a *Moving Least Squares* surface, defined by a point set as a robust implicit extension of Moving Least Squares that preserves sharp features, by using non-linear regression and finally applies the marching cubes algorithm.
- *Poisson reconstruction* [Kazhdan et al., 2006]. It obtains volumetric occupancy by solving a Poisson equation and then performs marching cubes in order to extract the surface.

Robust implicit moving least squares-marching cubes

Moving least squares is a very attractive tool to design effective meshless surface representations. However, as long as approximations are performed in a least square sense, the resulting definitions remain sensitive to outliers, and smooth-out small or sharp features. In [Oztireli et al., 2009], these major issues are considered and a novel point-based surface definition combining the simplicity of implicit MLS surfaces [Kolluri, 2005] with the strength of robust statistics is presented.

To reach this new definition, MLS surfaces are reviewed in terms of local kernel regression, opening the doors to the utilization of a robust kernel regression. This representation, which can handle sparse sampling, generates a continuous surface, which better preserves fine details, and can naturally handle any kind of sharp features with controllable sharpness. It combines ease of implementation with performance competing with other non-robust approaches. At the latest stage, a triangular mesh is extracted by using the marching cubes algorithm.

Poisson Reconstruction

Poisson reconstruction proposes a formulation that considers all the points at once, without resorting to heuristic spatial partitioning or blending, and is therefore highly resilient to data noise. Unlike radial basis function-based schemes [Turk and O'Brien, 2005], the Poisson approach presented in [Kazhdan et al., 2006] allows a hierarchy of locally supported basis functions, and therefore the solution reduces to a well conditioned sparse linear system. A spatially adaptive multi-scale algorithm is used, whose time and space complexities are proportional to the size of the reconstructed model.

Using this approach, surface reconstruction is expressed as a Poisson problem, which seeks the indicator function that best agrees with a set of possibly noisy, nonuniform observations. This approach can, for example, robustly recover fine detail from noisy real-world scans. At its last stage, an octree-based marching cubes implementation extracts a triangular mesh from the volumetric representation.

2.4 Conclusions

The main techniques and representation formats for 3D reconstruction from a set of calibrated input views found in the literature have been introduced. The state of the art provides a number of methods, based on different approaches, that are suitable for different types of scenarios that can be considered when working with multi-view camera settings. A taxonomy of the methods and types of representation that translate relevant features from the multi-view input data to 3D features (occupancy, surface, color...) has been presented.

When focusing on the 3D reconstruction of the objects in a scene captured by a set of multiple cameras, which are useful for both representation and further processing purposes, different methods yield results with dramatic differences, regarding the nature of the data used for reconstruction and the requirements of the specific application. Thus, methods offering high detail at its output are costly and oriented to small-baseline scenarios, such as the multi-view stereo approaches, whereas methods which provide lesser detail are suitable to interactive applications demanding small processing times and can have an arbitrarily wide-baseline configuration, which influences the reachable level of accuracy on the reconstruction.

We can observe that the only visible 3D features in multi-view data are the projections of the surfaces -i.e. the interfaces between the interior of objects and air- of the objects contained in the actual 3D scene. Therefore, we propose to exploit surface representations, with the use of surface patches or polygon meshes, in order to efficiently represent the reconstructed 3D features, instead of using volumetric representations, which in general spend valuable computational resources in computing and storing the occupancy of empty (air) or invisible (interior of objects) regions.

In order to efficiently compute surfaces in multi-view video sequences, the application of a tracking strategy seems a suitable option. In contrast with techniques for motion estimation by computation of the scene flow, tracking strategies can be used to exploit the temporal correlation of the scene in order to either reduce the computation time or improve the accuracy and consistency of the surface. Since reconstruction techniques are in general costly, it will be important to provide tracking techniques for the reconstruction of sequences, which allow for a combined exploitation of spatial and temporal correlations in the position of the surfaces.

Polygon meshes appear as good candidates for representing the visible features in 3D scenes, due to their powerful post-processing capabilities, being easy to handle in standard hardware-accelerated graphics pipelines and providing an explicit description of the surface topology. Meshes can be obtained from a discrete representation, based on surface samples, by different techniques, as shown in Section 2.3. However, the reviewed techniques are either computationally demanding or inaccurate in terms of topological correctness. Therefore, it is also proposed to obtain a method able to obtain a topologically correct continuous surface with a low computational cost.

In order to summarize the challenges of this thesis already mentioned in the previous chapter related to the state of the art, we attempt to provide a framework for the efficient processing of the available data in multi-camera scenarios by obtaining a single instance representation of the objects in the scene that:

- 1. Has little computational cost when compared to other methods delivering a similar accuracy in their results.
- 2. It is suitable for the reconstruction of multi-view sequences, exploiting both spatial and temporal correlations.
- 3. Contains all the relevant information contained in the original images.

Finally, it is important to remark that, although this thesis is not focused on specific higher-level computer vision or computer graphics problems such as pose estimation or techniques for scene visualization –for free-viewpoint video–, the ultimate goal of the obtained scene representation is to be used by these applications, without the expense of operating directly on the data captured by a multi-camera setting that would arise in image-based approaches.

Chapter 3

Approach

he previous chapter shows that the availability of multi-view data is beneficial in order to provide state-of-the-art performance in several disciplines in computer vision. In addition, the rich multi-view information can be more easily exploited when presented as a *single instance* of a 3D data structure properly describing all features of interest. Imposing the condition that the reconstruction is accurate, with a minimal information loss, it is always possible to go back to a good estimate of the original images by re-projecting the 3D models onto the original viewpoints. The methodology for obtaining this representation must solve early vision concepts –such as visibility, occlusion or image distortion– and leave for a posterior stage the feature analysis for specific applications.

One of the main goals of this approach, which can be interpreted as a low-level data fusion from multiple sensors, is to avoid using a *multi-instance* representation composed by the raw images from each available viewpoint and their corresponding calibration data at an application level. Such a multi-instance representation would further impact an analysis application attempting to exploit not only multi-view data, but also motion information by processing multi-view video sequences. Such scenario might require costly processing, transmission and storage of multi-view data from several time instants. Furthermore, algorithms based on the processing of a single instance 3D representation of a multi-view scene can be more general, because they will never depend on the number of available cameras and their interactions. Otherwise, the cost of processing multi-view data could exponentially grow with the number of input views.

This thesis attempts to demonstrate that surface representations of the elements of interest in multi-view video sequences can be obtained at a reasonable computational cost, providing complete and compact descriptions of the relevant visual features available in the original multi-view streams. In other words, this type of description can provide enough visual information to analysis applications as to render image-based approaches unpractical, due to the additional management cost associated to using raw multi-view data in the latter.

This chapter describes the approach presented throughout this thesis, designed to provide a complete and compact multi-view data representation of foreground elements in dynamic scenes, which can be used by both analysis and visualization applications. First, the set of visual features included in the proposed surface representations are presented and classified regarding their nature. Then, the necessary methods to obtain each of these representations are introduced. Finally, the organization of the rest of the dissertation is presented.

3.1 Surface Description

Relevant visual features in multi-view scenes are always found on the surfaces of objects of interest. Indeed, the interior space of objects of interest is typically hidden by an opaque surface, whereas empty space occupied by air does not provide any visual information. Therefore, it is this interface between the interior of objects and empty space the one that contains the relevant visual data.

Surfaces in 3D space can be represented in a number of ways, either explicit or implicit –a volumetric field determining spatial occupancy implicitly contains the location of surfaces–. In our approach, surfaces are explicitly described by means of two different, although related, types of primitives. The first representation consists in a sample-based description of the surface, whereas the second one is a polygon mesh.

One important advantage of considering a surface sampling scheme instead of a volumetric approach is that it is possible to augment the level of detail exactly on the surface of interest, by increasing the sampling density in that region. In contrast, a volumetric approach would require an increase of the sampling density in a region which would comprise the surface, part of the interior of the object and part of the empty space around the object, due to the impossibility of exactly determining the position of the surface.

The benefit of a polygon mesh compared to a sample based description is that it provides a mechanism –linear interpolation– to determine an estimate of the surface between existing samples. In comparison with other possible interpolation schemes, this one is supported by current mainstream graphics hardware, which makes it a good choice for interactive applications.



Figure 3.1: Three surface samples of the proposed 3D representation characterized by their geometric (3D position and orientation) and texture descriptors. The size of each surface sample in the proposed 3D representation is considered to be infinitesimal

3.1.1 Sample-based surface description

The first type of support for the description of 3D objects of interest in a multi-view scene is, as mentioned, a collection of surface samples. In its most elementary form, such representation consists in a point cloud with the 3D positions of a certain number of surface points. However, with the help of additional descriptors, higher degrees of information detail can be obtained.

In order to describe the shape of a 3D surface as estimated from data captured by a calibrated multi-camera setting, geometry and texture are complementary features which, combined, provide a rather complete set of the available visual features. Whereas the former describes the position and pose of a 3D surface element, as referred to a reference coordinate system, the latter describes its color at a local scale. An illustration of the concept of a surface sample is shown in Fig. 3.1.

Geometric descriptors

Geometric information can be expressed in different manners, depending on the requirements of the application, ranging from very low level descriptions, which require less computational resources for their extraction and representation, up to higher level descriptors, the extraction and representation of which require more resources.

In general, three different types of geometric descriptors can be considered for each surface sample:

• 3D position \mathbf{x}_i , with respect to an arbitrary 3D reference system.

- Orientation $\hat{\mathbf{n}}_i$, a unitary outward-pointing vector in a direction normal to the surface.
- *Curvature*, described by its two principal curvatures $\mathbf{k}_{1,i}$ and $\mathbf{k}_{2,i}$ [Guggenheimer, 1963].

At this point, it is interesting to discuss the properties of these three descriptors. The 3D position of a surface sample constitutes the most basic geometric description. By considering the variation of the position of the points included in a conveniently sized neighborhood of the surface sample, a higher level descriptor –its orientation– can be estimated. At the highest level, the variation of the orientation of the points belonging to the neighborhood of the surface sample can be used to estimate its curvature parameters. However, this third type of descriptor will not be used in our approach, for curvature can be determined at any later stage by exploiting the extracted location and orientation of a set of infinitesimal or *elementary* surface samples at close distances.

Alternative approaches consider larger surface samples [Carceroni and Kutulakos, 2002], which require more detail in the represented features (curvature, orientation and also texture). In this case, even the computation of the orientation of each surface sample would require more than the knowledge of the location of a set of close surface samples and a hint about the interior and exterior of the object, because samples would be too distant to describe local variations. Thus, an interesting property of elementary surface samples is that, given a sufficient sampling density, a small amount of information at every point –location and orientation– suffices to geometrically determine the surface.

Texture descriptors

Texture can also be encoded as an attribute of each surface sample. Assigning an RGB color that is consistent with each one of the viewpoints that have a line of sight with the surface point described by the sample already constitutes a low-level texture description. Such a method for describing texture can have an arbitrary level of detail, by varying the density of the surface samples.

Higher-level features, such as *Gray-Level Co-occurrence Matrix* (GLCM)-derived figures [Haralick et al., 1973] or edge histograms can also be used when a low density of surface points is needed for implementation constraints. The use of this type of features allows using a sparser surface description while keeping valuable information about the texture of the object. Therefore, they can constitute a useful complement to the low-level RGB color descriptor when the sampling density over the surface cannot provide enough level of detail of its texture variations, which are in general subtler than their geometric counterparts.



Figure 3.2: Local detail of a polygon mesh describing a surface. Each triangular face, which might contain texture information, shares its vertices with other neighbor faces, thus describing the surface topology

In this thesis, only the RGB color has been used as a texture descriptor, assuming a sufficient surface sampling density. Again, the adoption of an approach based on elementary surface samples, instead of larger surface samples, results in a benefit in terms of the required per-sample information. For both types of features, geometric and texture, a rather complete description of the actual surfaces can be obtained by encoding a small amount of data (only position, orientation and color) for every sample when a sufficient sampling density is considered.

3.1.2 Polygon mesh surface description

In order to describe the reconstructed surface even in regions lying between elementary surface samples, a polygon mesh is used. Polygon meshes can be considered a first-order polynomial interpolant – or *linear* interpolant– of the actual surface, given a set of elementary surface samples. It constitutes, therefore, a suitable continuous representation of the information contained in a set of discrete surface samples.

There are different ways to describe a polygon mesh, such as lists of vertices and faces or a half-edge structure. In our case, we consider a the former type of description, composed by triangular faces that share a set of vertices. An illustration of such description is shown in Fig. 3.2. However, the latter is also used at an intermediate step in order to check the topological correctness of the polygon mesh.

Some properties regarding the topology of a mesh must be fulfilled if the mesh is used for analysis. A two-manifold mesh can be properly oriented, whereas a mesh not fulfilling this topological property cannot be, for example, unfolded onto a 2D plane [Floater and Hormann, 2001]. Since surfaces reconstructed from multiview video sequences using the proposed approach are going to be used for both



Figure 3.3: Ordering of the three vertices of a triangular face following the right-hand rule, given the orientation of the face

visualization and analysis, special care will be taken in order to provide surfaces with correct topology.

Geometric descriptors

As pointed out above, surface geometry described by polygon meshes is considered to be consisting of triangular faces that share their vertices with neighbor faces.

Vertices. Similar to the case of surface samples, vertices in a mesh can contain different types of geometric descriptors. Apart from the 3D position of the vertex in an arbitrary reference system, a commonly included feature is an orientation, which can be used to compute surface illumination. However, unless otherwise stated, only the 3D position of each vertex will be considered as a per-vertex feature throughout this dissertation.

Faces. Typically, triangular faces are geometrically described as a set of three indices to their corresponding vertices. This is due to the fact that their orientation is commonly assumed to follow the right-hand rule of ordering of their vertices, as shown in Fig. 3.3.

Texture descriptors

Without considering more sophisticated effects such as *mip-mapping* –used in computer graphics in order to provide superficial detail beyond the geometric resolution of the polygon mesh–, two types of texture descriptors can be considered, regarding a polygon mesh (geometrically) described by means of lists of vertices and faces.

- Face texture, containing high resolution detail between vertices.
- Per-vertex RGB color, used to interpolate the surface color between vertices.

Due to the minimalistic texture representation chosen in surface sampling, the most natural choice for the mesh representation is a per-vertex color description. Such a texture description will be accurate as long as the vertex density of the mesh is sufficiently high. Again, the choice of an elementary surface sampling approach with a sufficient density allows encoding the continuous surface with a small amount of data when compared to that of providing a texture to each face composing the mesh. Indeed, in a closed surface with N_f faces, approximately $N_v = 2N_f$ vertices exist. Whereas encoding the texture coordinates in a parameterizable two-manifold mesh would require N_v 2D coordinates $-3N_f$ in non-manifold surfaces– and a texture image for each surface only N_v RGB colors (3D coordinates) are required for pervertex color information.

In this section, two types of surface representation have been presented, which will be used at different stages throughout the methodology presented in this thesis. Such methodology, resulting in the extraction of these two types of surface descriptions, is introduced in the next section.

3.2 Methodology

Several factors strongly influence the design decisions in this thesis towards methods able to work with arbitrarily wide baselines in real-time scenarios. Examples of these factors are the computational cost of processing an arbitrarily high number of cameras or the impossibility of placing a desired number of cameras due to spatial and usability constraints.

Whereas multi-view stereo methods [Campbell et al., 2008, Esteban and Schmitt, 2003] provide state-of-the-art quality with properly conditioned small-baseline setups, the surface reconstruction methods proposed in this thesis might fall one step behind those in terms of quality in properly-conditioned scenarios. However, this strategy allows us to provide surface reconstruction in challenging wide-baseline scenarios, where more accurate methods fail due to the absence of the required availability of redundant data.

The two different types of surface representations that are considered in the following chapters are used in two different stages of the processing chain depicted in Fig. 3.4. Therefore, sample-based representations of the reconstructed surfaces are obtained at the end of the *surface sampling* stage, whereas continuous polygon mesh representations are available after the *surface interpolation* stage. Due to the different nature of these two methods and type of representation, this thesis has been divided in two parts, surface sampling and surface interpolation, tackling the specific challenges at each of these stages. These are described next.



Figure 3.4: 3D surface reconstruction as a two-stage approach. After the first stage, a samplebased representation of surfaces is obtained, whereas, after the second one, a polygon mesh is obtained as the interpolation between the surface samples.

3.2.1 Surface sampling

The surface sampling methods presented in this thesis can be classified by two related criteria:

- The type of *available data*, leading to different possible reconstruction methods, some of them faster and others more accurate and offering richer information of the surfaces.
- The type of *search strategy* for each surface sample, leading to different types of algorithms, some better adapted to exploit photo-consistency and others to exploit temporal correlation.

In the following, the differences that appear from the choice of one of the types of available data and search strategies are briefly introduced.

Reconstruction methods and available data

As depicted in Fig. 3.4 Part I, a number of possible pipelines has been considered in order to extract surface samples from multi-view data. Thus, in case of considering only the silhouettes of objects of interest for a single time instant, a first type of surface reconstruction (*silhouette-consistent*) appears. This initial pipeline can be slightly modified to exploit temporal correlations (*dynamic surface*) or color information (*photo-consistency*).

Silhouette-consistent reconstruction. The reconstructed surfaces obtained by this approach are those of a *visual hull* [Laurentini, 1995, Franco and Boyer, 2009], corresponding to the techniques in Section 2.1.2 in the previous chapter.

Photo-consistent reconstruction. When attempting to obtain a single-instance representation containing all the relevant features of the input images, it is necessary to also consider color or texture information. As seen in Section 2.1.4, it is useful to exploit both binary silhouettes of the objects of interest and color images [Kolev and Cremers, 2008]. The former data provides a strong robustness and the ability to reconstruct shapes in multi-view scenarios with an arbitrarily wide baseline by limiting the spatial search area. The latter, color information, allows refining shapes by enforcing photo-consistency between images, using techniques similar to those in Section 2.1.3.

With this approach, apart from obtaining tighter surface estimates than those obtained with the silhouette-consistent approach, additional shape features are also directly available at the output: a texture descriptor –such as a RGB color– is included at each surface sample when photo-consistency is also enforced. Please note that a texture descriptor could also be attached to silhouette-consistent reconstructions, although photometric information would not be used to improve the accuracy of the reconstructed surface. In turn, this latter approach involves a smaller computational cost.

Dynamic surface reconstruction. Inter-frame correlation is exploited in this thesis for reducing the computational cost involved in the reconstruction of surfaces in consecutive time instants of multi-view video sequences. Interestingly, as presented in Section 2.2.2, the exploitation of the correlation of a scene between consecutive time instants can also provide an estimate of the motion of each surface sample.

Reconstruction methods and search strategies

Surface sampling can be intuitively interpreted as the distribution of samples, at a constant or varying density, all over the surfaces of the objects that are being reconstructed. One desirable property related to the distribution of samples would be that their density is increased where required, while keeping a lower density in areas without relevant information.

Whereas this would be difficult to obtain with classical volumetric techniques, where 3D space is arbitrarily divided in order to determine the occupancy of each subdivision –with some possible refinements such as *octrees* [Szeliski, 1993]–, the surface reconstruction methods present here an advantage, for they are free of these arbitrary divisions and sampling is more suitable to be adapted to the contents of the scene.

Thus, a *search* strategy specifically designed to exploit the features of the available data are expected to result in the placement of surface samples in valid positions.

As it will be seen in subsequent chapters, the choice of a strategy for search of surface samples leads to different algorithms. Some of these will present convenient behaviors to exploit temporal correlation, whereas others will adapt better to the exploitation of photo-consistency. A possible classification of the methods presented in this thesis is:

- *Image-based* search, with search regions defined along the back-projected rays of camera pixels, where surface samples are associated to individual pixels.
- *Surface deformation-based* search, with search regions along the normals of each surface sample, assuming a certain level of correlation between the surfaces in consecutive time instants.
- *Statistical* search, with search regions distributed around other valid surface samples, where the likelihood of finding new surface samples is higher.

The advantages and disadvantages of these methods will be presented and discussed throughout the chapters dedicated to surface sampling.

3.2.2 Surface interpolation

Surface interpolation corresponds to the *Part II* in Fig. 3.4. The input data consists of surface samples with or without color information. Although there exist several methods in the literature for extracting a polygon mesh from a set of oriented surface samples [Bernardini et al., 1999, Kazhdan et al., 2006], the target of the technique proposed here is to obtain a high degree of precision while keeping a reasonably reduced computational cost.

Some methods in the literature [Kazhdan et al., 2006] consist in first robustly determining the volumetric occupancy of 3D space and then extracting the surface as the interface between occupied and empty space (Section 2.3.2). The method

proposed in this thesis shares similarities with another family of mesh extraction methods, which can be referred to as propagation-based schemes (Section 2.3.1), with [Bernardini et al., 1999] constituting a well-known example. The main advantage of our method with respect to other propagation-based schemes is the usage of fewer computational resources, accomplished by the use of suitable data structures and a quick propagation algorithm based on a set of rules for guaranteeing the desired topological properties and a high accuracy.

At the output of this stage, an interpolating surface is obtained from the set of input surface samples. This surface can be used for both visualization and analysis, thanks to its topological correctness.

3.3 Thesis Organization

The main body of this thesis dissertation is divided in three parts. The first one presents a set of methods suitable for surface sampling from different types of input data. The second part describes the proposed meshing algorithm for surface interpolation. The third part presents the global results combining methods from Parts I and II, as well as some concluding remarks and proposed lines of future work.

3.3.1 Part I

The first part of the document contains chapters dedicated to the different surface sampling schemes, providing some advantages and disadvantages for each data configuration considered for reconstruction. In the final one, the techniques are validated and compared. The information presented in each of these chapters is the following:

- In Chapter 4, an *image-based surface sampling* approach with search regions consisting in the back-projected rays of each pixel, exploiting photoconsistency.
- In Chapter 5, a surface sampling based on the *deformation* of an existing surface, with search regions along the normal of each surface sample, exploiting temporal correlation.
- In Chapter 6, a surface sampling scheme based on the *statistical propagation* of surface samples in regions with higher likelihood of containing surfaces, exploiting spatial correlation.
- In Chapter 7, experiments with each one of the proposed techniques and the corresponding results, followed by a comparison of the three methods.

3.3.2 Part II

The second part of the dissertation contains two chapters dedicated to the extraction of a mesh interpolating a continuous surface between surface samples. First, the methodology is introduced and then, in Chapter 9, results are presented. The contents of these chapters are:

- In Chapter 8, a *surface interpolation* approach –meshing algorithm– is presented, which provides accurate polygon meshes with a low associated computational cost.
- In Chapter 9, some results obtained with the interpolation technique are presented.

3.3.3 Part III

The two previous parts of the thesis include partial results corresponding to the sampling and interpolation techniques. The third part of the thesis dissertation contains overall results, validating the complete approach –sampling and interpolation– and the conclusions. The chapters conforming this last part of the document contain:

- Chapter 10 presents quantitative and qualitative evidence of the suitability of the complete approach –surface sampling and interpolation– for the problem of surface reconstruction.
- Chapter 11 lists an overview of the accomplished goals and proposes some lines of future work.

Finally, Appendices A, B, C and D introduce, respectively, concepts related to the geometry of calibrated cameras, the *kd*-tree structure for efficient spatial queries, the *Hausdorff*-distance as a metric for surface quality and a list of publications obtained as a result of the research conducted in the elaboration of this thesis.

Part I

Surface Sampling

Surface Sampling

his part of the dissertation introduces a methodology for providing a complete, discrete representation of the surfaces of moving objects –scene foreground– that are observed by a rig of calibrated cameras. This is accomplished by a set of procedures oriented to the finding of the location and orientation of surface samples by means of specifically designed search schemes aiming at efficiently exploiting the spatial or temporal redundancies available in the input multi-view data.

Surface sampling, when compared to other 3D reconstruction methods, provides a higher degree of freedom in the desired output resolution —in this case, sampling density—. Although this might difficult the process of interpolating between surface samples, the technique presented in Part II correctly handles this possible limitation. In volumetric analysis, computational costs grow when an increase of the sampling density at the visible features of the 3D scene —surfaces of objects— is desired, due to the uniform distribution of samples in the entire working volume. However, in surface-oriented approaches, sampling density can be more freely increased, because the additional computational cost will only be used for samples representing objects surfaces.

In the following chapters, three different techniques are presented, which develop the principle of finding the position and orientation of surface samples in different manners. Regarding different types of input data –binary silhouettes or color images, multi-view video sequences or static multi-view scenes–, each of the proposed techniques shows some advantages and some limitations. These techniques are the result of an evolutionary approach in the search for the best possible technique that efficiently exploits the available information and improves the limitations of each previous one.

Chapter 4

Image-based Surface Sampling

s presented in Section 2.1.1, some approaches for scene reconstruction from a set of calibrated views show a high level of robustness, such as the silhouette-based methods, whereas others attempt to obtain the finest possible detail, such as current multi-view stereo techniques. Photo-consistent techniques fill in the gap between these two types of approaches, providing photohulls. In certain environments, it is required to obtain a complete and robust representation of a multi-view scene under tight constraints on the smallest achievable baseline of the camera setting. In such a scenario, photo-consistency can provide better accuracy with respect to the real shapes than silhouette-based methods, while being more flexible in terms of baseline than multi-view stereo approaches.

In order to exploit photo-consistency, an image-based approach appears as a good candidate to reconstruct surfaces from multi-view data [Furukawa and Ponce, 2007]. Therefore, in this chapter an image-based sampling strategy is presented, which provides a photo-consistent surface reconstruction of foreground objects in scenes with known background.

4.1 Motivation

Multi-view stereo techniques aim at reconstructing the 3D shapes of objects from images captured from small baseline multi-view setups. This setup permits reconstructing surfaces of objects in high detail, as long as the surfaces are visible from a large enough number of viewpoints. When the volume enclosing the scene to be reconstructed is large, more viewpoints must be placed around the scene in order to keep a small baseline. However, it is not always possible to arbitrarily increase the number of viewpoints, due to either physical –camera placement– or computational constraints.

The goal of the presented method is an alternative approach to tackle the problem of improving accuracy without greatly increasing the number of views. The target scenario consists in a scene in an indoor environment with known background whose images are captured from a highly sparse set of viewpoints, as few as *e.g.* 4 or 5 –as the cardinal points and a top view– in a wide baseline setup. These viewpoints are placed around the whole scene in a way that they offer a wide spatial coverage. In contrast with the usual *sparse* scenarios that are considered in multi-view stereo, which typically include between 10 and 20 viewpoints, the discrepancies due to perspective changes between images in such setups are dramatic and influence the reachable reconstruction accuracy.

4.1.1 Related work

The visual hull was introduced in [Laurentini, 1994] as the maximal 3D shape consistent with the silhouettes of an object projected onto a set of viewpoints. Several methods in the literature, usually referred to as *shape-from-silhouette*, estimate the visual hull. In these methods, input data consists solely of binary silhouettes. In [Franco and Boyer, 2009], a polygon mesh is directly reconstructed, which represents the surfaces of the visual hull. This technique provides a highly precise surface description when compared to most common volumetric counterparts.

A pre-processing stage, which may introduce segmentation errors, must provide accurate silhouettes of the objects of interest. In order to make visual hulls more robust against errors in silhouettes, [Franco and Boyer, 2005] and [Landabaso et al., 2006] provide algorithms for obtaining visual hulls from silhouettes presenting false detections and misses, both using a voxel grid representation.

The photo hull was introduced in [Kutulakos and Seitz, 1999] as the maximal 3D shape which is photo-consistent with the set of input images. In this case, the captured images are used as input data, instead of binary silhouettes. In their work, a voxel grid is again used as 3D support. The algorithm proposed for extracting the 3D shapes, called *space carving*, consists in setting all voxels of an initial volume that contains the actual scene as occupied and iteratively carving away those visible voxels that, visited using *sweep planes* –in order to avoid computing the visibility of each voxel from each camera–, do not pass a photo-consistency test. Due to the richer information available, the photo hull constitutes an estimate of the actual shape tighter than the visual hull, allowing concavities on the reconstructed shapes.

Voxels are useful to represent dense volumes, but they do not accurately describe

their surfaces. Alternatively, *surface samples* are designed for that purpose. Several examples in the multi-view stereo literature, like [Carceroni and Kutulakos, 2001] and [Furukawa and Ponce, 2007], have already used this type of representation, either to provide the final result or as an intermediate representation. In [Furukawa and Ponce, 2007] an initial sparse set of surface samples is obtained from matching image-based features. From this initial estimates, an iterative algorithm augments the density of samples and rejects outliers, achieving results with high accuracy. Then, a polygon mesh is morphed with an energy minimization algorithm until it fits the set of reconstructed surface samples.

In [Campbell et al., 2008], another common strategy is utilized, which consists in, first, estimating the depth map of each available view by using local groups of input images and then combining them in a voxel grid, in order to obtain the global surface estimate. Usually, combining methods use some regularization technique. Minimization via graph cuts offers fast performance when minimizing a certain domain of energy functions with binary variables, as defined in [Kolmogorov and Zabih, 2004].

In some other works, such as [Kolev and Cremers, 2008], the robustness of the visual hull is combined with color information, allowing the reconstruction of concavities. A convex functional is minimized with the exact silhouette consistency imposed as a constraint that enforces the resulting shape to fit in the binary silhouettes once re-projected onto the original viewpoints. However, multi-view stereo techniques are designed for scenarios where the robustness against the dramatic discrepancies between views present in wide-baseline setups is not a determining factor and, therefore, do not suit the requirements imposed by our target scenario.

4.1.2 Problem statement

The proposed method aims at reconstructing the photo-consistent surfaces of objects, the images of which have been captured by a very sparse set of cameras with known pose and intrinsic parameters. Similarly to [Furukawa and Ponce, 2007], the presented method exploits photo-consistency by performing an image-based search for surface samples. From [Kolev and Cremers, 2008], the proposed method keeps a resemblance in the sense of exploiting binary silhouettes in addition to photoconsistency measurements in order to constraint the position of the surface samples. The complete method consists in the concatenation of three stages:

- 1. First, candidate depths for the image points in each view are obtained, based on photo-consistency measurements between pairs of pixels from neighbor cameras.
- 2. Then, the surface points are extracted by globally imposing photo-consistency



Figure 4.1: Normalized Cross-Correlation (NCC) exploits perspective similarity in order to find accurate matches between neighbor views when used in typical multi-view stereo scenarios. A match for a given image point is found along its epipolar line on the neighbor view

[Kutulakos and Seitz, 1999], considering the complete set of available views.

3. Finally, a normal for each surface point is estimated by processing the local geometry of its neighborhood.

4.2 Photo-consistency Test

Due to the severe perspective discrepancies between different views –wide-baseline setups–, the properties of photo-consistency measurements used in state-of-the-art multi-view stereo approaches do not match our specific requirements, despite of their great level of precision.

A widely utilized form of photo-consistency measurement in multi-view stereo, the so-called *normalized cross correlation* (NCC), assumes a certain perspective similarity between viewpoints, which actually exists in small-baseline image sets. This assumption permits measuring the NCC in a straight-forward, fast manner by just computing it over 2D windows at each image, as illustrated in Fig. 4.1. In scenarios with a slightly wider baseline, a planar grid can be defined in 3D space, with sample points on the grid being projected onto the images. This grid must vary its orientation depending on the actual 3D geometry of the surfaces corresponding to the observed projections, as shown in [Furukawa and Ponce, 2007]. This alternative is computationally expensive, since both photo-consistency and orientation have to be jointly extracted and there remains the assumption of a maximal baseline that is too small for our target scenario.

In dynamic scenes where background images are available, foreground silhouettes of objects of interest in each image can be automatically extracted with the technique proposed in [Stauffer and Grimson, 1999]. Silhouettes are useful for photo-consistency for two reasons: on the one hand, they reduce the search space in the 3D domain; on the other hand, they can be used to reject many photoconsistent matches that may lie far from the actual object surfaces. Therefore, a simpler photo-consistency test with little impact from the perspective discrepancies can be considered if it is applied in combination with the constraints imposed by the silhouettes. In order to test the photo-consistency of a 3D point that belongs to the visual hull, we choose a less computationally expensive photo-consistency test that measures the squared Euclidean RGB distance between the pair of image points where it projects¹. In order to cope with non-Lambertian surfaces, the photo-consistency measurement admits an additional tolerance to brightness differences when the chromaticity shows a variation below an adjustable threshold. A more detailed description of the photo-consistency test applied to pixels belonging to the projections of objects of interest follows.

A robust photo-consistency measure for non-Lambertian surfaces, which results in different projections of the surface showing different brightness levels, should account for the possible discrepancies in both their brightness and chromaticity components. Based in the foreground/background color-based segmentation scheme presented in [Horprasert et al., 1999, Xu et al., 2005], brightness distortion (BD)can be defined as a scalar value that brings expected photo-consistent pixels close to the observed chromaticity line. Color distortion (CD) can be defined as the orthogonal distance between the expected photo-consistent color and the observed chromaticity line. Both measures are shown in Fig. 4.2 and formulated in Eqs. 4.1.

$$BD = \arg\min_{\alpha} \left(\mathbf{RGB}_1 - \alpha \cdot \mathbf{RGB}_2 \right)^2$$

$$CD = \|\mathbf{RGB}_2 - BD \cdot \mathbf{RGB}_1\|$$

$$(4.1)$$

¹In our experiments, using other color spaces, such as YUV or LabCIE, did not provide clear improvements with respect to RGB.



Figure 4.2: Distortion measurements in RGB space. \mathbf{RGB}_1 and \mathbf{RGB}_2 are the RGB values of two pixels from different images. CD is the chromaticity distortion between both colors and BD is the corresponding brightness distortion

In practice, BD can be easily computed as $BD = \mathbf{RGB}_2 \cdot \mathbf{RGB}_1 / \|\mathbf{RGB}_1\|^2$. A set of thresholds and rules are necessary to define the final photo-consistency test. This test, which takes a maximum RGB distance (RGB_{max}) , has the following form:

- 1. If $\|\mathbf{RGB}_1 \mathbf{RGB}_2\| < RGB_{max}$, then pixels are photo-consistent.
- 2. Else, if $\|\mathbf{RGB}_1 \mathbf{RGB}_2\| \ge RGB_{max}$, but CD < 10 and 0.5 < BD < 1.25, then pixels are also photo-consistent.
- 3. Otherwise, pixels are not photo-consistent.

This photo-consistency test, the thresholds of which have been empirically determined in order to obtain the desired behavior, provides a robustness improvement for photo-consistency measurements when considering non-Lambertian surfaces. In fact, when chromaticity remains similar enough between views, the range of allowed brightness distortion will include a wide range of shadowing and highlighting effects.

4.3 Sampling of Visible Photo-consistent Surfaces

Although the technique presented next could also be used to reconstruct the surfaces of a visual hull by using only binary silhouettes as input data, in the following we consider the more general configuration of using both color images and binary silhouettes of objects of interest. The method consists in generating lists of candidate locations for photo-consistent surface samples along the back-projected ray of each pixel and then imposing global photo-consistency by means of a carving procedure.



Figure 4.3: A list of photo-consistent candidates for a given pixel (marked in yellow in the left view) is obtained from testing pixels along its epipolar line in a neighbor view and stored in a list, in order of increasing depth with respect to the current viewpoint

4.3.1 Automatic neighborhood selection

In order to generate lists of photo-consistent candidates –that will contain the resulting surface points– from pixels in every view, a set of neighbor views for the image-based search must be defined.

The neighborhood of each viewpoint is automatically obtained by taking its closest viewpoints, defining the proximity of two viewpoints as

$$\pi_{ij} = \pi_{ji} = \left(1 - \frac{1}{2}\cos(\alpha_{ij})\right)\delta_{ij},\tag{4.2}$$

where α_{ij} is the angle between the two unitary vectors orthogonal to the image planes of viewpoints *i* and *j* and δ_{ij} is the Euclidean distance between the centers of projections –Appendix A– of the two viewpoints.

4.3.2 Extraction of candidate surface points

For each viewpoint v_c of the complete set V, we define two sets of neighbor, N_c , and complementary, C_c , viewpoints, which have empty intersection and such that $N_c \cup C_c \cup v_c = V$.

Any image point back-projects to the 3D world as a ray with its origin at the center of projections of the given viewpoint, as seen in Appendix A. Thus, if we want to know the exact position of the 3D point which originated the captured



Figure 4.4: A candidate surface point visible but not photo-consistent in a third viewpoint is rejected. The colored square inside of each camera represents the color of the corresponding pixel when the 3D point is visible.

2D image point, we must search for it along the mentioned back-projected ray. In order to do so, we utilize the information available in the other views by means of a valuable geometric tool, the *epipolar line*, also introduced in Appendix A.

As illustrated in Fig. 4.3, for a given 2D point x_c that lies on the extracted foreground silhouette of an image i_c corresponding to the viewpoint v_c , we obtain the epipolar line $l_n(x_c)$ at each neighbor viewpoint $n \in N_c$. For each foreground pixel p traversed by $l_n(x_c)$, we measure its photo-consistency with the pixel at the point x_c . If the photo-consistency test defined above is positive, we triangulate the 3D position of the new candidate surface point. If the triangulated point projects to a foreground pixel for each view in $C_c \cup N_c$, we store it in a list of candidate 3D points $\{z_{x_c}\}$ with an increasing value of depth with respect to the center of projections of the view v_c .

At the end of this stage, a list of candidate surface points for each image point in each view is obtained. Furthermore, since the lists of candidate points are ordered in increasing depth from each viewpoint, estimates of the depth maps are implicitly obtained. The surface described by the points at the end of this stage can be interpreted as a shape which is already tighter to the actual shape than the visual hull provided by a *shape-from-silhouette* method, although it does not yet constitute a photo hull, since it lacks global photo-consistency.

4.3.3 Global photo-consistency

The next stage consists in discarding candidate surface points which are not photoconsistent with the complete set of viewpoints V. In order to do so, we will carve away those 3D points which do not pass the photo-consistency test defined in Section 4.2 for at least one of the viewpoints.

The algorithm works as follows: for each viewpoint v_c , we take the first point at every $\{z_{x_c}\}$. We consider the case that the 3D point obtained from two neighbor



Figure 4.5: A candidate surface point visible and also photo-consistent in a third viewpoint is taken as the new closest candidate in this viewpoint. The colored square inside of each camera represents the color of the corresponding pixel when the 3D point is visible.

views is visible from another viewpoint, meaning that its distance to the center of projections is smaller than the current depth estimate for the corresponding image point. If the candidate point is photo-inconsistent with the new viewpoint, it is removed and the next closest candidate for each view where it was visible is taken. This situation is illustrated in Figure 4.4. Otherwise, if the 3D point is photoconsistent with the new viewpoint, then we set it as the new closest candidate for the corresponding image point. This second situation is illustrated in Figure 4.5.

This filtering strategy is similar to space carving, with the difference that, instead of carving visible voxels, we carve candidate surface points along back-projected rays. As an advantage, the proposed method does not introduce additional spatial sampling artifacts due to the arbitrary division of the 3D space that occurs when using a voxel grid. All of the candidate surface points are obtained from triangulating image samples with the available geometric constraints from camera calibration. Furthermore, the search space is restricted, because it is bounded by the visual hull as a result of the inclusion of silhouette information. As in space carving, every time some point has been carved away, the algorithm iterates on the new set of surface points. This is repeated until no more changes are detected, which means that the set of surface samples is globally photo-consistent.

After running this algorithm, we obtain the surface corresponding to a photo hull, represented by surface samples which are tighter to the actual surface than the initial candidates and can show concavities. In order to obtain a complete description of the surface, the orientation of these points must be estimated.

4.4 Orientation Estimation

In order to estimate the orientation of each surface point, a least squares method is applied for fitting a plane to a set of points in a neighborhood of it, assuming local flatness. It is possible to fit a plane to a set of 3D points using least squares with errors measured orthogonally to the proposed plane. Let the plane be $\mathbf{\hat{n}} \cdot (\mathbf{x} - \mathbf{x}_a) = 0$, where $\mathbf{\hat{n}}$ is a unit length normal to the plane and \mathbf{x}_a is a point on the plane. Define \mathbf{x}_j to be the sample points; then

$$\mathbf{x}_j = \mathbf{x}_a + \lambda_j \hat{\mathbf{n}} + p_j \hat{\mathbf{n}}_j^* \tag{4.3}$$

where $\lambda_j = \hat{\mathbf{n}} \cdot (\mathbf{x}_j - \mathbf{x}_a)$ and $\hat{\mathbf{n}}_j^*$ is some unit length vector perpendicular to $\hat{\mathbf{n}}$ with appropriate coefficient p_j . Define $\mathbf{y}_j = \mathbf{x}_j - \mathbf{x}_a$. The vector from \mathbf{x}_j to its projection onto the plane is $\lambda_j \hat{\mathbf{n}}$. The squared length of this vector is $\lambda_j^2 = (\hat{\mathbf{n}} \cdot \mathbf{y}_j)^2$. The energy function for the least squares minimization is $\mathbf{E}(\mathbf{x}_a, \hat{\mathbf{n}}) = \sum_{j=1}^m \lambda_j^2$. Two alternate forms for this function are

$$E\left(\mathbf{x}_{a}, \hat{\mathbf{n}}\right) = \sum_{j=1}^{m} \mathbf{y}_{j}^{\top} \left(\hat{\mathbf{n}} \hat{\mathbf{n}}^{\top}\right) \mathbf{y}_{j}$$

$$(4.4)$$

and

$$\mathbf{E}\left(\mathbf{x}_{a},\hat{\mathbf{n}}\right) = \hat{\mathbf{n}}^{\top} \left(\sum_{j=1}^{m} \mathbf{y}_{j} \mathbf{y}_{j}^{\top}\right) \hat{\mathbf{n}} = \hat{\mathbf{n}}^{\top} \mathbf{M}\left(\mathbf{x}_{a}\right) \hat{\mathbf{n}}$$
(4.5)

Using the first form of E in the previous equation, take the derivative with respect to \mathbf{x}_a to get

$$\frac{\partial \mathbf{E}}{\partial \mathbf{x}_a} = -2 \left(\mathbf{\hat{n}} \mathbf{\hat{n}}^\top \right) \sum_{j=1}^m \mathbf{y}_j \tag{4.6}$$

This partial derivative is zero whenever $\sum_{j=1}^{m} \mathbf{y}_j = 0$ in which case $\mathbf{x}_a = \frac{1}{m} \sum_{j=1}^{m} \mathbf{x}_j$ (the average of the sample points).

Given \mathbf{x}_a , the matrix $\mathbf{M}(\mathbf{x}_a)$ is determined in the second form of the energy function. The quantity $\hat{\mathbf{n}}^{\top}\mathbf{M}(\mathbf{x}_a)\hat{\mathbf{n}}$ is a quadratic form whose minimum is the smallest eigenvalue of $\mathbf{M}(\mathbf{x}_a)$. This can be found by standard eigensystem solvers. The corresponding unit length eigenvector $\hat{\mathbf{n}}$ is the result of the extraction of the normal direction to a set of points by least squares fitting of a plane.

If
$$\mathbf{x}_a = (a, b, c)^{\top}$$
 and $\mathbf{x}_j = (x_j, y_j, z_j)^{\top}$, then matrix $\mathbf{M}(\mathbf{x}_a)$ is given by

$$\begin{pmatrix} \sum_{j=1}^{m} (x_j - a)^2 & \sum_{j=1}^{m} (x_j - a)(y_j - b) & \sum_{j=1}^{m} (x_j - a)(z_j - c) \\ \sum_{j=1}^{m} (x_j - a)(y_j - b) & \sum_{j=1}^{m} (y_j - b)^2 & \sum_{j=1}^{m} (y_j - b)(z_j - c) \\ \sum_{j=1}^{m} (x_j - a)(z_j - c) & \sum_{j=1}^{m} (y_j - b)(z_j - c) & \sum_{j=1}^{m} (z_j - c)^2 \end{pmatrix}.$$
(4.7)

The normal estimated by using this approach still shows one degree of freedom. The sign of the unit length eigenvector –the normal, $\hat{\mathbf{n}}$ – corresponding to the smallest eigenvalue is not defined. Therefore, additional features have to be exploited to completely determine surface orientation. In the following, a mechanism for this image-based sampling strategy is presented.

4.4.1 Orientation in image-based sampling

Let r_n be an arbitrary maximum distance for neighbor points. The algorithm applied to each surface point \mathbf{x}_i works as follows:

- 1. Define an extrinsic neighborhood ν_i including all the surface points at an Euclidean distance from \mathbf{x}_i smaller or equal to r_n .
- 2. Estimate the mean position μ_i of the neighborhood ν_i and modify the position of each neighbor point \mathbf{x}_j by subtracting it:

$$\mathbf{x}_j := \mathbf{x}_j - \boldsymbol{\mu}_i, \quad \forall \mathbf{x}_j \in \nu_i$$

3. Create a matrix $\mathbf{M}_i = \sum_{\mathbf{x}_j \in \nu_i} \mathbf{x}_j \mathbf{x}_j^{\top}$. A vector $\hat{\mathbf{n}}_i$, orthogonal to the plane fitted by least squares to the neighborhood ν_i , is obtained as the first eigenvector of \mathbf{M}_i .

This method does not control whether the resulting orthogonal vector $\hat{\mathbf{n}}_i$ is inwards or outwards pointing, but the knowledge of the camera viewing the surface point helps setting the sign of the orientation by enforcing a negative dot product between the point normal and the director vector of the corresponding back-projected ray, which has a direction inverse to that from which the surface is visible.

4.5 Conclusions

The presented image-based approach to surface sampling obtains a set of colored and oriented surface samples, which describe the visible photo-consistent surfaces of the objects of interest in images –foreground objects– using an arbitrarily sparse set of views. With this technique, surface samples are found by a searching mechanism over regions defined as back-projected rays of pixels in every image.

As introduced in Chapter 3, it would be wise to exploit the information about the contents in a certain time instant to improve the quality or the efficiency of the reconstruction at the next time instant, due to the strong temporal correlation on the position of the surface samples, when working with multi-view video sequences. However, due to the selection of search regions for surface samples, an arbitrary motion of the surfaces in 3D space cannot be naturally modeled and incorporated in the proposed technique in order to enhance the efficiency or the accuracy of surface sampling.

As a summary, the presented technique makes a search of visible surface samples along the back-projected rays of each pixel in each view, in an image-based procedure exploiting epipolar constraints and photo-consistency between pairs of neighbor views in arbitrarily wide baselines. Some valuable properties of this approach are:

- The visibility computation is implicit. Whenever a surface candidate is carved away, its next candidate becomes visible. This feature shows the method is suitable for the efficient computation of photo-consistent surfaces.
- It can provide a pixel-perfect projection of the reconstructed surface under accurate photometric and geometric camera calibration. This corresponds to the goal of being able to retrieve the original images by the projection of the reconstructed surfaces.

However, this approach presents some drawbacks that limit its efficient applicability for real-time processing of multi-view video sequences, which is the main target of this thesis:

- The reconstruction mechanism makes it difficult to exploit temporal correlations when reconstructing surfaces in multi-view video sequences, due to the strict definition of search regions along back-projected rays of pixels. This results in a limited throughput.
- Reconstructed surfaces are open, because only surfaces samples visible from at least one view are part of the final surface –surfaces are not extrapolated–. Therefore, algorithms exploiting topological properties of the objects surfaces will be affected by this limitation.

The next chapter proposes an alternative surface sampling approach that tackles these drawbacks by searching the positions of surface samples in every time instant around the deformation of an existing, closed surface from a previous time instant.
Chapter 5

Surface Sampling by Deformation

he previous chapter has introduced an image-based sampling scheme to obtain photo-consistent samples of visible surfaces in arbitrarily wide baseline setups. Such technique is able to extract a cloud of oriented points that can also describe the silhouette-consistent surface of a static scene, when replacing the color information by binary silhouettes. However, the design of the search strategy for surface samples is not suitable for exploiting temporal correlations in multi-view video sequences.

In order to cope with this limitation, a new sampling scheme based on the deformation of an existing surface is proposed, which naturally exploits temporal redundancies in addition to spatial ones.

5.1 Motivation

When multi-view video sequences are available, a large amount of information about the scene at each new time instant is known from the reconstruction at the previous time instant. Therefore, it seems straightforward to provide algorithms able to tackle the problem of 3D reconstruction as the exploitation of both spatial and temporal redundancies in multi-view video sequences simultaneously.

In contrast with the approach presented in the previous chapter, where the search regions for surface samples were back-projected rays of pixels, this time a different search strategy is considered, which is more oriented towards the actual structure of the surfaces to be reconstructed. The new search regions are defined



Figure 5.1: The volume enclosed by the dilation (dark thick line) of the existing surfaces (light thick line) of objects in certain time instant very likely contains the surfaces of these moving objects (in orange) in the next time instant, thus reducing the search space

along the (inward) normals of samples of an existing surface. This arbitrary choice exploits the fact that the volume enclosed by a sufficient dilation of an existing surface in a time instant most likely contains the object in the next time instant, under the assumption of a limited amount of motion between frames. This is illustrated in Fig. 5.1. The main advantage of this approach is that the search space for every time instant is greatly reduced, once conditioned to the knowledge of the previous position of surface samples.

5.1.1 Related work

A number of methods exists for extracting the *visual hull* (VH) from a set of binary silhouettes. Some provide a voxelized volumetric representation [Cheung et al., 2000] and others provide surfaces described by polygonal meshes [Franco and Boyer, 2009]. The VH is fast to estimate and offers an initial description of the volumes in the scene for applications with camera rigs in arbitrarily wide-baseline setups. However, these algorithms are not designed to exploit temporal redundancies available in video sequences.

Some achievements in motion estimation have been made by separating the shape and motion estimation tasks, in a *scene flow* estimation approach of the problem that can use color [Vedula et al., 2005] or spherical models [Starck and Hilton, 2005] as matching features, although the shape estimation problem is treated statically, without exploiting temporal redundancies. There also exists a growing number of algorithms that jointly estimate the 3D motion and photo-consistent shape of scenes by tracking surfaces in multi-view settings [Furukawa and Ponce, 2008, Courchay et al., 2009, Goldluecke and Magnor, 2004] using local normalized cross-correlation and smoothing constraints, a global variational approach or

spatio-temporal hypersurfaces, respectively. As a result, impressive time-coherent reconstructions are obtained at the cost of large computation times. [Guan et al., 2008] tracks people in crowded scenes, exploiting image-based appearance models, which help disambiguating shape modeling without requiring color calibration between cameras. This approach is designed to work even in wide-baseline setups.

Using range data for initializing a high-quality mesh model of the surface of a person, [de Aguiar et al., 2007] applies Laplacian deformations in order to fit the high-quality surface to the set of silhouettes at each instant, obtaining impressive offline results. However, in *smart room* environments it is desirable to only consider the data obtained from calibrated cameras, with known or static background, in order to make a better user experience. In such a scenario, no range data is available in order to estimate precise models of the dynamic objects in a scene.

5.1.2 Strategy

The methodology for efficiently extracting silhouette-consistent surface samples by the exploitation of temporal redundancies that is presented next corresponds to a tracking strategy as opposed to *scene flow* estimation strategies. The proposed strategy is divided in two parts:

- Initialization: Static reconstruction. First, samples of a silhouetteconsistent surface are obtained for an initial time instant without having any knowledge of the shape at previous time instants by exploiting only spatial redundancies.
- **Tracking: Dynamic reconstruction**. Then, the knowledge of the shape in the previous time instant allows for a more efficient sampling of the surfaces in the next time instant, by exploiting both spatial and temporal redundancies.

The different stages composing this technique for surface sampling, based on the deformation of an existing surface, are illustrated in Fig. 5.2. In both stages of the complete tracking strategy –initialization and dynamic surface–, the resulting output is a sampled representation of the surfaces of the VH at every time instant.

5.2 Initialization: Static Reconstruction

The starting point of the initialization procedure is the definition of an arbitrary, closed initial surface that encloses a large volume containing the scene to be reconstructed. The initial surface is described as a point cloud of arbitrary density with the property that each point p_i has a unitary displacement vector \mathbf{m}_i , orthogonal to



Figure 5.2: The tracking strategy is divided in an initialization stage, at which an arbitrary sampled surface is deformed to represent the surface of the foreground objects of the scene, and a tracking stage (dynamic surface), at which the surface in the previous time instant is used in order to estimate more efficiently the position of surface samples in the new time instant

the initial surface, pointing inwards to another point p_i^* , also lying on the surface. This property defines a segment s_i for each pair of surface points p_i and p_i^* in which their displacement range is constrained.

This arbitrarily defined constraint on the displacement range of surface points is suitable for two reasons. On the one hand, it ensures that the processing of a new view keeps each surface point inside of the limits imposed by any other previously processed viewpoint. On the other, it allows to keep topological information about interior and exterior of the surface.

The main concept behind the proposed algorithm for the initialization stage is the iterative introduction of modifications of the closed surface in order to fit its projection to the foreground silhouette of each viewpoint in the multi-camera setting. This concept in spatial domain resembles that of adapting an existing surface to a new set of silhouette constraints in another instant in time domain, which is the principle of the tracking part of the presented surface sampling approach. At the output of the static reconstruction stage, oriented silhouette-consistent surface samples, corresponding to a Visual Hull, are obtained. First, silhouetteconsistency is imposed by a per-view application of a set of topological operations on each segment connecting surface points and a global merging of the resulting segments. Then, orientation is estimated by exploiting local distributions of surface samples and knowledge about the interior and exterior of the objects, given by the segments.

5.2.1 Silhouette constraints

Silhouette constraints are imposed in a two-staged approach that copes with the problem of camera settings not offering complete views of the scene to be reconstructed. First, in a per-view stage, every segment is conceptually shrunk and cut into smaller pieces by imposing the silhouette constraints corresponding to each view. Then, a global combination of these partial results delivers the final set of surface points, which are the extremes of such segments.

Per-view constraints

The first stage when adapting the arbitrary initial surface –described as a point cloud with connected points forming segments– to a new silhouette is the *shrinkage* of each segment s_i . Both extrema p_i and p_i^* are displaced along the segment towards each other until they either intersect the back-projected silhouette or meet each other. In the former case, the segment is just shrunk in order to fit the silhouette constraint. In the latter case, the segment –and therefore its extreme points– can be discarded.

However, after shrinking a segment, holes in silhouettes and empty regions between silhouettes still have to be accounted for. In this case, the extremes are further moved towards each other until they meet, introducing additional segments by *cutting* the current segment every time the point displacement along the segment makes them enter and leave the interior of the back-projection of the silhouette. These two topological operations on the segments are illustrated in Fig. 5.3.

After these two operations, each segment is replaced by a set of new, shorter segments that fit the constraints of the silhouette corresponding to the view under consideration. When daisy-chaining this stage between all views, these new constraints are automatically merged to those resulting from the silhouettes of previously processed views. Throughout the entire process the surface remains closed.

However, as seen in the next section, it is better to consider this daisy-chained version of the method only as a concept, for it would be faster to execute –each



Figure 5.3: Topological operations on a segment connecting two surface points by the imposition of the silhouette constraints from one view. (a) Initial segment connecting two points on an arbitrary initial surface. (b) Shrinkage in order to fit the visible extents corresponding to the silhouettes in that view. (c) Segment is cut in two pieces due to the silhouette configuration. (d) Resulting set of segments for this view

view reduces the search space of the subsequent views in the chain– but would not provide mechanisms to reconstruct parts of the foreground objects that lie out of the *frustum* –or viewing cone– of one or more views.

Global combination

Rather than directly performing the topological operations –in other words, cutting and shrinking segments– during the processing of each input silhouette, an ordered list of events for each initial segment accumulates both the partial inclusion in the camera's frustum and the *occupancy* from each view. The types of events are illustrated in Fig. 5.4. A final stage merges the per-view constraints and decides the topological operations applied to each segment. The resulting algorithm processes each initial segment as follows:

1. For each view, the events corresponding to the partial frustum inclusion and occupancy extremes –shown in Fig. 5.5– are obtained.



Figure 5.4: 2D representation of a set of events detected along a segment from a certain viewpoint

- 2. The events from all views are stored in a list, ordered by the distance of the position of each event to a pivot point on the line where the segment lies.
- 3. The extremes of the resulting segments are those points where either the occupancy count becomes equal to the number of inclusions in frusta (entering shape) or the occupancy becomes smaller than the number of inclusions in frusta (exiting shape), as shown in Fig. 5.6.

The resulting surface points are the extremes of the segments created during the global combination stage. Using this mechanism, the topological operations on the segments joining points from the initial closed surface are applied at the last stage. This way, it is not necessary that the object is included in the frustum of each view in order to be reconstructed.

Furthermore, conservative estimates of the surface of the visual hull are possible, by assuming the existence of τ silhouette under-segmentation errors of a point included in the volume of a foreground object –in other words, in τ views, the projection of the same part of an object is classified as background–. Then, in the third step of the previous algorithm, a tolerance τ can be used as a sufficiently small difference between the occupancy count and the number of inclusions in frusta that produces a surface point. $\tau = 0$ provides the non-conservative, visual hull surface.

In order to characterize the surface by a set of points, without keeping topological information included in the extents of the segments, the orientation of each surface sample must be defined as shown next.

5.2.2 Orientation estimation

Normals are geometrically estimated in a way similar to that in the image-based method in Section 4.4. However, in order to estimate whether the resulting vector



Figure 5.5: Partial frustum intersection and occupancy of an initial segment from a viewpoint. The orange segment represents the projection of a 3D segment onto the image plane.

 $\hat{\mathbf{n}}_i$ is inwards or outwards pointing, the knowledge of the unitary displacement vector of each surface point –pointing towards where the segment lies– is the tool used to determine the sign of the estimated normal. Using the same notation as in the image-based case, let λ_i be the average projection norm of the inward-pointing unitary director vector \mathbf{m}_j tied to each $\mathbf{x}_j \in \nu_i$ onto $\hat{\mathbf{n}}_i$. If $\lambda_i < 0$, $\hat{\mathbf{n}}_i$ is considered to point outwards and therefore its sign remains unchanged, whereas, if $\lambda_i > 0$, it is forced to point outwards by changing its sign: $\hat{\mathbf{n}}_i := -\hat{\mathbf{n}}_i$. In Fig. 5.7, a set of surface points lit by a virtual light at the viewing position is shown. At the left side normals are not available and, therefore, all points are lit equally, whereas at the right side the use of normal information provides a better representation of the surface.

The neighborhood ν_i used for estimating the normal direction also provides an approximate description of the local topology of the reconstructed surface around each sample. This information proves useful for the dynamic estimation of the position of surface samples, which will require the application of spatial regularization.



Figure 5.6: Occupancy along an initial segment for 3 input views with inclusion in a minimum of 2 frusta. The resulting surface points are the extremes of the occupied segments



Figure 5.7: Surface points lit considering normals (right), which are geometrically obtained from the positions of sample points on the surface of the visual hull (left)

5.3 Tracking: Dynamic Reconstruction

As introduced in the previous chapter, the image-based reconstruction technique defines an arbitrary search region for each surface element, which is the back-projected ray corresponding to each pixel of the objects of interest –foreground objects–. This search criterion keeps the method from naturally handling shape priors in sequences with motion, due to the fact that search areas cannot be modified in order to exploit directions of motion of a surface sample between two consecutive time instants different from the arbitrarily defined ones.

However, a surface deformation-based method allows for the choice of a better adapted search region for each surface sample. Hence, the design for the search region is exploited in this section in order to accept shape priors from a previous time instant and track them at each new time instant. The proposed method is the first one in this thesis exploiting spatial correlation. It is capable of obtaining a surface sampling at each time instant equivalent to that obtained with static approaches (either the initialization stage of this technique or the image-based approach in the previous chapter), but with the benefit of a smaller computational cost.

In the initialization stage of this technique, parts of the initial arbitrary surface -represented as a point cloud with pairs of points defining segments- were allowed to vanish in cases where the strictly guided displacement of the surface points along their arbitrarily defined search regions makes the extremes of a segment collapse without having crossed a surface along their path. However, when attempting to estimate shape in a tracking strategy, it is not permissible to discard surface points, since that would result in a growing under-sampling of the surfaces as time passes.



Figure 5.8: A sampled surface and its dilation following the estimated outward normals for each surface sample. The size of the surface samples has been enlarged for clarity

Therefore, the proposed method introduces a regularization strategy for reducing the side-effects of arbitrarily defining the search regions for each surface sample.

5.3.1 Surface expansion

Assuming that the displacement of surface samples between consecutive frames is bounded, r_t is defined as an arbitrary radius of the 3D search area from one time instant to the next. The search space is defined in spatial units of the 3D reference system, instead of pixel units that would be used in classical image processing. This is a clear advantage of working with 3D reconstructions of multi-view data, which is fundamental in the proposal of this thesis.

The dilation of a sampled surface can be defined as the displacement of each surface sample by a parametric distance r_t in the direction of the normal of each surface sample. Fig. 5.8 depicts the *dilation* of a surface obtained by translating the surface points a distance r_t following their estimated outward normals. It is interesting to note that the volume enclosed by the resulting surface very likely contains the object at the next time instant. Thus, r_t controls the range of distances a surface sample can displace between from frame to frame such that it can be tracked by the proposed algorithm. However, a mechanism for recovering samples displaced a distance larger than r_t is also considered by means of spatial regularization.

5.3.2 Spatial regularization

Similarly to the initialization procedure resulting in a static reconstruction, the search direction for each surface sample is again set as an inner pointing vector –in



Figure 5.9: In (a) a set of oriented surface samples (green) represents the surface of the orange shape. In (b) the shape has moved with respect to its initial position. In (c) the initial surface samples have been dilated along their normals in order to generate a new search volume. In (d) some samples reach the surface following the inward normal direction. In (e) local rigidity is imposed in the remaining sample by adding a tangential displacement (blue arrow). In (f) the whole set of samples lies on the new surface thanks to the use of regularization

this case, the inverted normal of a surface sample–. The main difference is that in this case the search space is greatly reduced with respect to that of the initialization. Instead of an arbitrary surface enclosing the volume where foreground objects lie, the enclosing surface is the result of dilating the surface in the previous time instant.

As shown in Fig. 5.9, the arbitrarily defined search direction results in some surface points not being able to meet the surface in the previous time instants, whereas others slide over the surface. In order to cope with the former problem and reduce the latter, regularization is applied in order to enforce local rigidity, as seen below.

Let \mathbf{x}_i^e be a surface point expanded along its normal $\hat{\mathbf{n}}_i$ by r_t units. Its inward normal direction $-\hat{\mathbf{n}}_i$ can be considered as a direction to look for an initial estimate of its position at the new time instant. However, some sort of regularization is needed in order to keep local rigidity on the set of points. Therefore, an iterative method with motion feed-back is applied on the whole list of points. In outline, this iterative method modifies the search range of each point \mathbf{x}_i^{t-1} by adding a set of tangential displacements d_i obtained from its neighbors ν_i . In more detail, the initial set of displacements for each surface point is initialized $\mathbf{d}_j := (0, 0, 0)^{\top}$, $\forall \mathbf{d}_j \in d_i$. Then, the following method is applied to each surface point \mathbf{x}_i iteratively:

1. Let N_i be the number of \mathbf{x}_i 's assigned displacements. Define an initial *displaced* position as the addition of its position at the previous time instant and

its current set of displacements $\mathbf{x}_i^d := \mathbf{x}_i^{t-1} + \frac{1}{N_i} \sum_{\mathbf{d}_i \in d_i} \mathbf{d}_j$.

- 2. Define a virtual segment, joining its displaced expanded position $\mathbf{x}_i^e := \mathbf{x}_i^d + r_t \hat{\mathbf{n}}_i$ and its displaced *contracted* position $\mathbf{x}_i^c := \mathbf{x}_i^d r_t \hat{\mathbf{n}}_i$.
- 3. Shrink the virtual segment by displacing the *outer* surface point \mathbf{x}_i^e towards \mathbf{x}_i^c . Along this path, store the closest position to \mathbf{x}_i^{t-1} , namely \mathbf{x}_i^s , at which the outer surface point crosses the surface and enters the intersection of the back-projection of all the available silhouettes corresponding to cameras with frusta containing that region.
- 4. If the \mathbf{x}_i^s is found, keep it as the current candidate for its new position $\mathbf{x}_i^t := \mathbf{x}_i^s$ and obtain the corresponding displacement $\boldsymbol{\delta}_i := \mathbf{x}_i^t - \mathbf{x}_i^{t-1}$.
- 5. If δ_i is defined, set a displacement term for each neighbor $\mathbf{x}_j \in \nu_i$ orthogonal to its displacement direction $\tau_{ij} = \delta_i \delta_i \cdot \hat{\mathbf{n}}_j$. The motion component along the inward normal is directly set by the shrinking mechanism.

The regularization procedure described above is executed until no more samples can be fit to the new silhouettes. Some samples, lying in regions such as limb extremes, might not be attached to the surface in the new time instant, despite of the regularization. However, they can be integrated in the reconstructed surface by applying the same rigid transform (in this case a translation) as the resulting one of its closest neighbor that fits the new surface. This method allows these surface samples to be recovered at subsequent iterations of the algorithm in another time instant, without the need to destroy them and create new ones in order to prevent an otherwise progressively growing uncovered area of the surface.

Finally, the normal estimation method introduced in static reconstruction (Section 5.2.2) is applied to each of the resulting surface points. A dense velocity field can be obtained as a first order estimate by simply taking the displacement of each surface point between two consecutive time instants. This velocity field is not intended to be accurate, but it illustrates the direction and magnitude of the displacements that result after the application of the dynamic method. In Fig. 5.10, it can be observed how the color-encoded velocity field estimated in this manner actually provides a plausible approximation to the actual motion of the scene.

5.4 Photo-consistency Constraints

So far we have seen how silhouette constraints can be used in order to deform an initial surface and make it represent the surface enclosing the visual hull of the set of silhouettes. However, color information contained in the images captured by the multi-camera rig has not yet been introduced to the surface deformation mechanism.



Figure 5.10: Shape and motion jointly estimated from sequences of silhouettes extracted from 8 cameras placed around a person. The estimated velocity is color encoded following the axes at the bottom-left side. White points correspond to the initial surface.

The approach we propose in order to introduce the photo-consistency constraints imposed by the contents of the color images is that of further deforming the resulting surfaces until they are photo-consistent with the complete set of available views. The approach is divided in two steps that are repeated iteratively until no more segments are shrunk in order to fulfill the photo-consistent constraints imposed by the available images:

- 1. Obtain the visibility of each surface point.
- 2. Shrink each segment until its extremes (the surface points) are either nonvisible or photo-consistent.

5.4.1 Visibility estimation

Computing the visibility of a surface point from a given viewpoint is not straightforward, due to the fact that the topology of the surface is unknown. Furthermore, the density of the point cloud describing the surface can be arbitrarily small. Even if the exact topology of the surfaces in the scene remains unknown, the neighborhood of each surface point provides useful information for obtaining its visibility.

The condition for a surface point to be visible from a certain viewpoint is that it has a direct line of sight with the center of projections of the camera. Let z (p = (x, y)) be the depth map of a certain viewpoint. If a surface point is visible, its depth measured from the center of projections must be equal to the depth of the pixel onto which it projects for this cam.

The problem that arises is that of computing a proper depth map from a cloud of 3D points projected onto a given viewpoint without an exact knowledge of its topology. In practice, this means that an area around the projection of each surface point has to be drawn in order to close the depth map and avoid the possibility that a surface point in a hidden surface becomes visible through the holes. The chosen approach consists in drawing the convex hull of the neighborhood of each surface point, which ensures filling in the gaps between projected surface points. The resulting algorithm for each view works as follows. First, initialize the depth map to a value further than any surface point. Then, for each surface point:

- 1. Compute the depth of the surface point with respect to the viewpoint z_p .
- 2. Compute the convex hull of its neighborhood.
- 3. For each pixel in the convex hull, if its depth is smaller than the value stored in the depth map, set its depth to z_p .

Once the depth map is computed, each surface point whose depth with respect to a viewpoint is smaller than $z(p) + \delta$, with $p = (x_0, y_0)$ and δ a threshold for balancing the inaccuracy of the convex hull representation empirically set to $1/2r_n$, is considered to be visible from that viewpoint.

The depth map can be improved by computing the exact topology of the surface and drawing faces instead of convex hulls. Doing so, the δ threshold would be unnecessary and the visibility computation would be, as a result, more precise.

5.4.2 Segment shrinking

When an extreme of a segment is visible from more than two views, its photoconsistency has to be measured, using the same method as in the image-based method (Section 4.2). When the photo-consistency test is not passed, the segment starts a shrinking process until it either becomes hidden, turns photo-consistent or vanishes by meeting its other extreme.

When all visible segments have been processed in this manner, a new carving iteration starts. First, the current visibility of each view is computed and, then, the segments are shrunk accordingly to their visibility and photo-consistency. The algorithm ends when all the extrema of the segments, *i.e.* the surface points, are photo-consistent.

5.5 Conclusions

A new surface sampling strategy has been presented, which exploits temporal and spatial correlations in order to reduce the computational cost associated to the reconstruction of each frame in scenarios with static, known background. This technique is based on the deformation of a closed initial surface in order to obtain silhouette-consistency.

The design of the technique has been driven by the proposal of solutions to the problems found in the sampling strategy presented in the previous chapter. However, the algorithm presented in this chapter shows some drawbacks:

- Obtaining pixel-perfect projections requires a high number of initial surface points due to the random initial placement of surface points, many of which are discarded as a result of the deformation strategy for initialization.
- Photo-consistency constraints require costly visibility computations when compared to the technique presented in the previous chapter.
- The initialization part of the technique is not well-suited to massively parallel platforms, such as GPUs, due allocation and memory alignment issues derived from the division of the initial segments connecting two surface points into an undetermined quantity of smaller segments resulting from the two applied topological operations.
- There exists a long-term drift of surface samples towards areas further apart from the motion direction, due to the action of rigidity constraints –which are, in turn, necessary in order to recover surface samples close to extremes, such as limbs, and avoid excessive sliding of points on the surface–.

On the other hand, this method shows important advantages with respect to the technique presented in the previous chapter:

- The reconstruction mechanism allows exploiting spatial and temporal correlations on the position of surface samples in multi-view video sequences, thus improving the efficiency of the sampling stage.
- The search space is not determined by the images, but rather by the 3D structure of the scene.
- The reconstructed surfaces are closed, with extrapolated non-visible parts, which provides a more effective representation for visualization *-stencil buffer* shadowing [Segal et al., 1992]- and analysis applications.
- The strategy might be used to obtain an initial estimate of the dense motion of the scene. Although this is not the main target of the sampling approach, it could be used as an additional feature in multi-modal analysis applications for motion detection.

Chapter 6

Statistical Surface Sampling

he previous chapter has presented a sampling approach intended for the exploitation of temporal correlations in multi-view video sequences. Although spatial correlations are implicitly exploited in the regularization process taking place in the tracking stage of that algorithm, in this chapter we present a methodology for taking advantage of spatial correlations in a more principled manner.

Samples of 2D surfaces in 3D space show a high degree of spatial compactness. Indeed, given a certain volume in 3D space which dimensions are in the order of $\propto r^3$, its corresponding 2D enclosing surface has dimension $\propto r^2$. This result implies that, given a known surface sample, new surface samples can be efficiently found by inspecting a compact region surrounding it.

In this chapter, this property is exploited in order to improve the process of surface sampling from a set of calibrated views. First, we motivate the exploitation of spatial correlation. Then, we present an initialization or *scouting* stage, in which a small number of *seed* surface samples is obtained, and a propagation stage, where a dense set of surface samples is efficiently obtained by exploiting spatial correlation. In order to improve the efficiency of the initialization stage, two techniques are proposed, which exploit temporal correlations in multi-view video sequences and multi-resolution. Finally, some techniques for the post-processing of the resulting set of surface samples are presented.

6.1 Motivation

The working principle of the technique presented here can be justified by an example with lower dimensionality. Let p_j denote a pixel and $i(p_j)$ a binary 2D image with $i(p_j) = 0$ for pixels representing an empty area and $i(p_j) = 1$ for pixels representing an occupied area, with an image size equal to $a_t = \sigma_x \times \sigma_y \propto r^2$.

Let $a_b \propto r$ be the size of the border of such region. Then, if we define a procedure where p_j takes a random pixel in the rectangle $[1..\sigma_x] \times [1..\sigma_y]$, the probability of one such pixel lying on the border of the occupied region, π_b , is the ratio between the size of the border and the size of the complete image, or $\pi_b = a_b/a_t \propto r/r^2 \propto 1/r$. This procedure is illustrated in Fig. 6.1 (a).

However, after a large number of attempts $O(1/\pi_b)$, the probability that the result of at least one of these pixels lies on the border of the occupied region tends to one. Although the number of attempts could be prohibitively expensive if the procedure was to be applied for each border point, it could be affordable at an initialization step if then a better suited strategy follows.

Let δ_n be a representative distance between neighbor points, and $\rho_n \propto \delta_n$ a search radius. A second procedure is defined as follows: let p_b be a border pixel obtained after a large number of tries of the first procedure and generate a new p_j as the sum of p_b and a random value in the ball $p'_j : ||p'_j|| \leq \rho_n$. Since the border in a local region is approximately a straight line, the number of border points is $\propto \rho_n$, whereas the area of the search region is $\propto \rho_n^2$. This means that the probability that the picked up pixel lies on the border of the occupied region, π'_b , is the ratio between the size of the intersection of the border with the new search region and the size of this new search region, or $\pi'_b \propto \rho_n / \rho_n^2 \propto 1 / \rho_n$. Since $\rho_n \ll r$, the resulting success probability for the new procedure accomplishes $\pi'_b \gg \pi_b$. This second procedure is illustrated in Fig. 6.1 (b).

With this example we have shown how the search in local regions can provide an efficiency improvement with respect to global searches. The same idea can be translated to a higher dimensional space, which results of great interest for the surface sampling problem. In the following, the details of the application of this principle in the extraction of surface samples is explained in detail: the strategy for obtaining the random surface sampling consists in an initial scouting procedure. As a result, a collection of *seed* surface points uniformly distributed over the surface is obtained. Finally, a finer local search provides a dense sampling of the object surfaces. The challenge of this method is to guarantee that surface parts are not missing due to the non-exhaustive search of surface points.

6.1.1 Random sampling in multi-view environments

Similarly to the case of finding a 1D contour in a 2D space, in multi-view environments it is necessary to find 2D contours (surfaces) in a 3D space, with the difference that, in this case, the only data from which we have actual measurements



Figure 6.1: Random sampling of a 1D contour in a 2D space. (a) It takes a large number of random tries to find a contour sample. (b) Once a contour sample has been found, a proper neighborhood (pixels delimited by orange lines) around a given contour point (light green pixel) helps in reducing the number of tries to find another contour sample.

are the 2D projections of the 3D scene onto each available view. Thus, 3D points under study have to be evaluated from their 2D projections onto each view.

6.1.2 Silhouette-consistent surface reconstruction

Given a set of silhouettes corresponding to different views of a scene, a silhouetteconsistent surface point necessarily projects onto a silhouette contour pixel in at least one of the images. In order to validate this statement, please keep in mind that the occupied region of the visual hull [Laurentini, 1994] is uniquely defined by a set of silhouettes. A property supporting this statement is the fact that the back-projection of a silhouette point contacts the actual surface of the object in at least one point.

Therefore, the problem of reconstructing the silhouette-consistent surface by random sampling can be tackled in the following manner. Let **P** be the outcome of a random experiment that generates 3D points inside of a predefined bounding volume $v_t = \sigma_x \times \sigma_y \times \sigma_z$. Let now \mathbf{p}_c be the 2D projection of **P** on a view c. If \mathbf{p}_c belongs to the contour of the silhouette in a camera c, then **P** is part of the surface. Naturally, the efficiency of this experiment, which is the base for the scouting or initialization stage, is rather low.

However, the actual improvement introduced by random sampling is not shown by this experiment. The search efficiency can be improved by reducing the search range. If we define a proper neighborhood around a surface sample, a second experiment can generate random 3D points \mathbf{P}' lying inside. This corresponds to the second stage, *propagation*, of the proposed method. In this stage, the likelihood that a randomly generated sample \mathbf{P}' belongs to the surface will be much greater than in the initialization stage.

6.2 Initialization

The algorithm for silhouette-consistent reconstruction consists of two stages. The first stage, *scouting*, is responsible for generating a collection of *seed* surface samples by random sampling the space enclosed by a bounding box where foreground objects are to be found. The second stage extracts new surface samples by iteratively testing new points in suitable neighborhoods around already obtained surface samples.

The concept of propagation can be understood as that of assigning a property of a certain sample to others lying in its neighborhood. As a result, a sparse initial set of samples of a certain class can be expanded to a dense set of samples. In our method, points lying around seed surface samples can be selected as new surface samples after a local propagation, as introduced in the next section.

The decision of generating more than one seed surface sample is due to the chosen type of propagation scheme, which is introduced below. Propagation can be done in two different manners:

- *Omni-directional.* Each surface sample propagates with equal likelihood in every direction inside a suitable neighborhood.
- *Directional.* Each surface sample is given a propagation direction such that the propagation process does not result in new samples in regions that have already been visited.

In practice, since the location of the surfaces is unknown, the only possibility is to apply the former. In order to apply the omni-directional propagation scheme without introducing local oversampling, it is necessary to define a uniformly distributed collection of surface samples that will propagate to a dense set of surface samples after an adequate number of iterations. The details will be seen in the next section, but for now we keep in mind the idea of obtaining a sparse set of uniformly distributed seed surface samples.

Let $[m_x..M_x] \times [m_y..M_y] \times [m_z..M_z]$ be a bounding box enclosing the scene to be reconstructed. An initial random experiment is defined, which produces random 3D points inside of this bounding box. Then, in order to accept a generated 3D random point **P** as belonging to the surface of one of the objects to be reconstructed, three conditions have to be fulfilled:

- 1. The 3D point projects onto a foreground pixel for each available view.
- 2. The 3D point projects to a contour pixel for at least one of the views.
- 3. A normal (the point's orientation) can be estimated by operations on a close neighborhood of the 3D point.

When these three tests are passed by a random 3D point, this is accepted as a seed surface sample after attaching its orientation estimation.

6.2.1 Surface point validation

Let N_c be the number of cameras with *frustum* including the 3D point under evaluation or, in other words, the number of cameras for which the projection of the 3D point **P** lies inside of the corresponding image. If for one of these N_c cameras, the pixel where **P** projects to belongs to the background of the scene, then the point is discarded.

The projection test described above checks the first of the three conditions for accepting the 3D point as belonging to the surface. In order to ensure that \mathbf{P} is an actual surface point, a second test is also performed on the point. Given a pixel \mathbf{p}_c where the 3D point \mathbf{P} projects onto the image corresponding to camera c, we test if at least one of the pixels in its 4-neighborhood belongs to the background. If the test is positive for at least one of the cameras, the point is accepted for testing the third acceptance condition.

Conservative estimation. This surface point validation method is correct in absence of segmentation errors. In case that there are some pixels wrongly marked as belonging to the foreground, the consistency checks between a large enough number of cameras will most likely reject false surface points projecting onto one of such pixels. The case of false background detection, however, will introduce severe reconstruction errors.

Assuming segmentation errors are uncorrelated between views, we can derive that the presence of one such an error in one of the N_c cameras onto which a 3D point projects is most surely accompanied by $N_c - 1$ correct decisions in the remaining cameras. Thus, by setting a tolerance to one segmentation error, a better estimate of the actual object can be obtained from noisy foreground detections, at the cost of slightly wider reconstructed shapes, providing a conservative estimate of the silhouette-consistent surface.

6.2.2 Local normal estimation

The orientation of each seed surface sample is important for two reasons. On the one hand, it offers necessary information about the topology of the surface which will be useful in order to interpolate the surface, as presented in Part II of this thesis. On the other hand, and more importantly for the approach presented in this chapter, it is necessary in order to define suitable search regions for the propagation stage. In order to estimate the orientation of a surface point, a local cluster of surface points is estimated in a close neighborhood of a valid surface point. Given a valid surface point \mathbf{P} , a new random search is performed in a ball of radius ρ_n centered at \mathbf{P} . ρ_n is an adjustable parameter that defines a small distance such that the local surface remains practically flat (constant normal). In general, it is defined as a small fraction of the main diagonal of the bounding box of the working volume.

Three new points are generated in this neighborhood, which pass the surface point validation test described above. Then, a plane is fitted to the set of 4 points by *least squares* and the director vector of the plane is taken. As introduced in Section 4.4, doing so leaves us with a degree of freedom: the sign of the normal vector. In order to determine its sign, an iterative method is applied, which consists in the following steps:

- 1. Add the orientation of the surface point multiplied by a small scaling value to the position of the surface point $\mathbf{x}_+ := \mathbf{x} + \alpha \mathbf{n}$.
- 2. Subtract the orientation of the surface point multiplied by a small scaling value to the position from the surface point $\mathbf{x}_{-} := \mathbf{x} \alpha \mathbf{n}$.
- 3. Obtain the results of the projection tests for both \mathbf{x}_+ and \mathbf{x}_- .
- 4. If one of the projection tests results in empty space and the other results in occupied space, the point can be accepted. If \mathbf{x}_+ belongs to empty space, the normal's sign is already correct. Otherwise, we change its sign: $\mathbf{n} := -\mathbf{n}$.
- 5. If both projection tests have the same result and $\alpha > \alpha_{\min}$, repeat the process with $\alpha := 4/5\alpha$. If $\alpha \le \alpha_{\min}$, the point cannot be successfully oriented.

If at the output of this test the result is that the 3D point \mathbf{P} could not be successfully oriented, it is discarded, together with its local neighbors.

6.3 Propagation

From the collection of seed surface points, which are sparse but uniformly distributed over the surface, a propagation procedure is defined in order to efficiently provide a dense sampling of the surface by exploiting spatial correlation. Thus, the following method aims at obtaining new surface samples at smaller distances from each other, resulting in a denser surface sampling. The benefit of introducing new samples taking into account the position of already obtained samples is the statistical reduction of the number of random tries until we find new valid surface samples.

6.3.1 Propagation region

Given a known surface sample, we define a local 3D region around it where new surface samples are very likely to be found, given that surfaces are continuous. The design rules of this propagation region are the following:

- Assuming surfaces are locally flat, or present a large curvature radius, preference will be given for new surface samples to be placed close to the plane defined by the surface sample's position and orientation.
- Given the fact that a certain density of surface samples is already placed around each surface sample, in order to estimate its orientation, new surface samples will be placed at a distance which will be larger or equal than the distance at which the furthest neighbor sample is placed.

One possible shape for the propagation region is a rectangle, although others could be used with similar result. Given a seed surface sample, we define two normal vectors $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ which are orthogonal to each other and also to the normal $\hat{\mathbf{n}}$. By construction, these two vectors $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ lie on the plane defined by the seed surface sample's position and orientation.

Rectangular propagation region. In order to generate this region, the following three random values are required for each sample:

- 1. Normal offset ω_n , or distance between the seed sample and the new random sample in the axis of the seed sample's normal, $\hat{\mathbf{n}}$.
- 2. Axial offset ω_u , or distance between the seed sample and the new random sample in the $\hat{\mathbf{u}}$ axis.
- 3. Axial offset ω_v , or distance between the seed sample and the new random sample in the $\hat{\mathbf{v}}$ axis.

Let ρ_n be a small maximum distance between surface samples such that these are close enough to assume they lie on the same plane. Let now ρ_r be a larger distance about one order of magnitude larger than ρ_n that is still small enough to assume a large number of surface samples are almost coplanar. Let n, u and v be uniform random variables in the range [0, 1]. Then, the previous random values ω_n ,



Figure 6.2: The rectangular propagation region around a surface sample \mathbf{x}_s –and its neighbor samples, used for estimating the local orientation–. It favors propagation on the plane orthogonal to the normal vector while allowing small curvatures with variations of up $2\rho_n$ along the normal

 ω_u and ω_v can be generated as follows:

$$\begin{aligned}
\omega_n &:= \rho_n \left(2n - 1 \right) \\
\omega_u &:= \begin{cases} \rho_r \left(u + \frac{1}{2} \right), u \ge 0 \\
\rho_r \left(u - \frac{1}{2} \right), u < 0 \end{cases} \\
\omega_v &:= \begin{cases} \rho_r \left(v + \frac{1}{2} \right), v \ge 0 \\
\rho_r \left(v - \frac{1}{2} \right), v < 0 \end{cases}
\end{aligned} \tag{6.1}$$

A gap between a seed surface sample and the ones obtained from propagation is left empty in order to apply the propagation algorithm iteratively, at increasingly smaller scales, until the desired number of surface samples is reached. With this, the omni-directional search for new surface samples using the same propagation region at different scales can be performed without creating clusters of surface samples at small distances. Also, since any surface sample is surrounded by three additional surface samples at a small distance –used for estimating its orientation–, we leave an extra space $\rho_r/2$ in order to avoid the generation of larger clusters of samples in the vicinities of seed surface samples.

The final random sample is computed by combining the contributions of the three random values:

$$\mathbf{x}_r = \mathbf{x}_s + \omega_n \hat{\mathbf{n}} + \omega_u \hat{\mathbf{u}} + \omega_v \hat{\mathbf{v}},\tag{6.2}$$

with \mathbf{x}_r and \mathbf{x}_s the random sample in a rectangular propagation region and the seed surface sample, respectively. The appearance of the rectangular propagation region is shown in Fig. 6.2. In the figure, a seed sample \mathbf{x}_s , surrounded by three neighbor samples used for estimating the local surface orientation, is the center of the oriented propagation region.

6.3.2 Iterative propagation algorithm

Let $q := {\mathbf{x}_s}$ be a queue of seed surface samples and N_s the target number of surface samples. Initially, q contains N_q seed surface samples obtained in the *scouting* stage (q does not contain the additional samples obtained for the estimation of the orientation, since propagating around them would generate local clusters of samples). ρ_n and ρ_r are initialized to valid values, according to the scene scale.

Then, while the required number of surface samples N_s is larger than the current number N_c –which is initially equal to $4N_q$, *i.e.* the total number of samples adding the seed surface samples and their neighbors–, the following algorithm is iteratively applied at monotonically decreasing scales (note step 6):

- 1. Create an empty queue of seed surface samples q'.
- 2. Pop the first seed sample \mathbf{x}_s from q.
- 3. Obtain a test surface sample \mathbf{x}_t lying in the rectangular propagation region and perform the foreground, contour and orientation tests introduced in the initialization stage. If \mathbf{x}_t is valid (it belongs to the surface of the visual hull and can be correctly oriented), add it to q', store its neighbors in a separate list and update N_c .
- 4. Repeat the previous step until 4 new surface samples are obtained by propagation of \mathbf{x}_s .
- 5. Repeat from 2 to 4 until q is empty.
- 6. Set q := q' and update the propagation region dimensions $\rho_n := \rho_n$ and $\rho_r := 0.5\rho_r$.

At the end, a total of N_s surface samples are obtained at a very small computational cost when compared to the initialization or *scouting* stage, as we will see in the experiments in Chapter 7. This results from the fact that the propagation region is designed to exploit the local distribution of surface samples.

Note that the additional gap introduced in the propagation region allows applying an omni-directional search for new surface samples at different scales of the propagation region without creating clusters of surface samples at small distances. This is illustrated in Fig. 6.3.

6.4 Efficient Sampling Schemes

Analyzing the strategy presented in this chapter, it is clear that the bottleneck resides in the initial search for seed surface samples, which is driven at random in



Figure 6.3: Propagation of a surface sample –and three neighbor samples– in order to generate 4 new surface samples –with three neighbor samples each–. Such propagation scheme allows to multiply by 5 the number of samples in every iteration, thus densely covering the surface

usually large working volumes –the space where people interact while being captures by the multi-camera rig–. Even though most of the surfaces samples are obtained during the propagation stage of the method, a large part of the computational resources are dedicated to this initial search.

In this section, two different techniques are proposed, which reduce the amount of computational resources spent in the initial search of surface samples. The first one is based on an idea already presented in Chapter 5, the existence of temporal correlations between the position of surfaces samples in a given a time instant and the next one, whereas the second one is based on considerations about the spatial sampling density implicit to the resolution of the input images.

6.4.1 Dynamic sampling

Similar to the dynamic surface sampling strategy presented in Section 5.3, the statistical approach for surface sampling can also benefit from temporal correlations in multi-view sequences. Indeed, the information about the position of surface samples in previous time instants can reduce the time required to find the seeds in the current time instant by limiting the volume in which new surface samples are to be searched for.

The approach taken in this chapter is not as tight as the one presented in the previous one. Indeed, only the seed samples are transferred from one time instant to the next one, leaving the complete covering of the surfaces to the efficient propagation scheme presented in Section 6.3. Thus, in this case an approximation of the dense velocity field of the surface samples cannot be obtained as a sample-wise first order estimate.

Assuming that the seed surface samples are evenly distributed over the surfaces of objects –which actually occurs, due to to the initial random search for surface points–, new positions for these samples can be safely searched for by independently setting a symmetric search region with a parametric radius for dynamic search ρ_d . This way, the placement of new seed samples for the current time instant does not favor any specific direction. Given three random values u, v, w in the range [-1, 1],



Figure 6.4: By exploiting temporal correlations between consecutive frames, the search space for seed surface samples can be reduced from the whole volume of $(M_x - m_x) \times (M_y - m_y) \times (M_z - m_z)$ to a smaller one $(4/3\pi\rho_d^3)$, thus increasing the search efficiency

such that $u^2 + v^2 + w^2 \leq 1$, the candidate to seed sample is computed as follows:

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \rho_d(u, v, w)^\top.$$
(6.3)

When the surface candidate and its three neighbor samples obtained for orientation estimation fulfill the three conditions presented in Section 6.2, they are added to the list of seed surface samples and the list of neighbor samples for the current time instant, respectively. Then, the next seed sample from the previous time instant is used for searching a new seed for the current time instant. This dynamic alternative to the initialization algorithm proceeds until finishing processing the whole list of seed surface samples from the previous time instant.

Fig. 6.4 illustrates the gain in search efficiency that can be achieved by exploiting the temporal correlation of the scene in this manner. Instead of finding a correct seed surface sample in the complete volume $V = (M_x - m_x) \times (M_y - m_y) \times (M_z - m_z)$, the search volume is reduced to $V' = 4/3\pi\rho_d^3$, which on its hand increases the probability of succeeding at finding a surface sample $\propto V/V'$.

After this alternative initialization, the propagation algorithm proceeds as in Section 6.3, resulting in the final sampling of the surface with the desired number of surface samples.

6.4.2 Multi-resolution sampling

The main purpose of this technique is to reduce the amount of random searches needed to find the initial seed samples by increasing their success probability. Please note that, given a set of silhouettes corresponding to different views of a scene, the likelihood of finding a contour pixel in random searches decreases when the resolution of the images increases. As a reminder, the projection onto a contour pixel for at least one silhouette is a requirement for a 3D point to be part of the silhouette-consistent surface –see Section 6.2–.

Let $N_x \times N_y$ and l_c be the resolution of the input images –measured in pixels– and corresponding contour length –also measured in pixels–, respectively. If the resolution of an image is reduced by decimating in each dimension by a factor δ , the corresponding contour length is also divided by δ .

Now let a randomly chosen 3D point project onto one of the input views. The success probability at full resolution π_f can be considered proportional to $\frac{l_c}{N_x \times N_y}$, assuming that random points project with equal likelihood to any point in the image. Then, for the success probability after decimation, $\pi_d \propto \frac{l_c/\delta}{N_x/\delta \times N_y/\delta} = \delta \pi_f$, the result is that it increases with respect to that at full resolution. This translates in a reduced number of random attempts to find a surface point.

In order to accept a generated 3D random point \mathbf{x}_l as a low-resolution seed to the silhouette-consistent surface, two conditions have to be fulfilled:

- 1. The 3D point projects onto a foreground pixel for all low-resolution views.
- 2. The 3D point projects to a contour pixel for at least one of these views.

The low-resolution seeds obtained with this method cannot be used as surface points because they are not accurate with respect to the high-resolution silhouettes. However, similar to the case of dynamic sampling, an actual surface sample can be found in a local neighborhood of low-resolution seed. This situation is illustrated in Fig. 6.5. The complete algorithm for multi-resolution sampling is as follows:

- 1. A list of low-resolution seeds is obtained from the decimated multi-view set.
- 2. As shown in Fig. 6.5, for each low-resolution seed, we apply an approach similar to that in the previous section, but using a search region around each seed with a multi-resolution search radius ρ_m , which should be as small as possible (in order to be efficient), but large enough to be able to find a valid surface sample in the region delimited by the back-projection of the pixels of the low-resolution views. As a result, a list of oriented seed samples and their corresponding neighbors is obtained.
- 3. Finally, the propagation procedure covers the entire surface with the desired number of surface samples.



Figure 6.5: The search for a low-resolution seed –gray dot– needs a much smaller number of random searches to obtain than the high-resolution counterpart –black dot–. In order to obtain the high-resolution seed from the low-resolution one, a local search a region around the latter must be performed. The area between the gray and black half-lines starting at each camera correspond to the back-projection of a low-resolution contour pixel and a high-resolution one, respectively

6.5 Surface post-processing

In order to improve the visual appearance of the reconstructed surface, some refinements and addition of complementary information can be applied to the resulting set of surface samples. In the development of this thesis, three post-processing operations have been considered.

The first one is a finer estimation of the surface normals by using the distribution of surface samples in a wider area that comprises more samples than the used to obtain the surface orientation in the initialization stage, thus providing more robust estimates for the surface orientation. The second one is an anisotropic smoothing to obtain a smooth surface at a sup-pixel scale. The third and final one assigns a color to each surface sample without imposing photo-consistency, targeting realtime applications.

6.5.1 Fine normal estimation

Assuming the surface to have a large local curvature radius, the same least squares method presented in Section 4.4 can be applied for fitting a plane to a set of points in a neighborhood of each point on the surface offering better robustness than the one of the method for local estimation introduced above.

Let r_n be an arbitrary maximum distance for neighbor points satisfying the previous assumption. The algorithm applied to each surface point \mathbf{x}_i works as follows:



Figure 6.6: Normal estimated from a set of points by fitting a plane passing through the point resulting from averaging the positions of the complete set (μ_i)

- 1. Define an extrinsic neighborhood $\nu_i := {\mathbf{x}_j}$ including each surface point \mathbf{x}_j at an Euclidean distance from \mathbf{x}_i smaller or equal to r_n .
- 2. Estimate the mean position μ_i of ν_i and modify the position of each neighbor point \mathbf{x}_j by subtracting it $\mathbf{x}_j := \mathbf{x}_j - \mu_i$, $\forall \mathbf{x}_j \in \nu_i$.
- 3. Create a matrix $\mathbf{M}_i = \sum_{\mathbf{x}_i \in \{\nu_i\}} \mathbf{x}_j \mathbf{x}_j^\top$.

A vector $\hat{\mathbf{n}}_i$, orthogonal to the plane fitted by least squares to the neighborhood ν_i , is obtained as the eigenvector of largest eigenvalue of \mathbf{M}_i . Finally, the sign of $\hat{\mathbf{n}}_i$ has to be determined. Let λ_i be the average projection norm of the estimated normals $\hat{\mathbf{n}}_j$ of each $\mathbf{x}_j \in \nu_i$ onto $\hat{\mathbf{n}}_i$. If $\lambda_i > 0$, $\hat{\mathbf{n}}_i$ is considered to point outwards and therefore its sign remains unchanged whereas, if $\lambda_i < 0$, it is forced to point outwards by changing its sign: $\hat{\mathbf{n}}_i := -\hat{\mathbf{n}}_i$. The method is illustrated in Fig. 6.6.

6.5.2 Anisotropic smoothing

Once normals have been robustly obtained with the application of the previous method, the surface orientation can be used to eliminate the high-frequency noise in the placement of the surface samples due to the finite resolution of the input silhouettes. We plan to smooth surfaces along the direction of the normal of each sample. This technique allows neutralizing the effect of local skewness in the distribution of the samples in a neighborhood, while providing a suitable smoothing of the detected surface samples' positions.

Thus, in order to obtain a more visually appealing distribution of surface points, a spatial smoothness constraint is imposed to a local neighborhood of each surface sample. Let the superscript o denote an output sample and r_n a neighborhood radius, equal to the one used for normal estimation; then, the algorithm applied to each surface sample \mathbf{x}_i works as follows:

1. Define a neighborhood $\nu_i := {\mathbf{x}_j}$ including each surface sample \mathbf{x}_j at an Euclidean distance from \mathbf{x}_i smaller or equal to r_n with a normal at an angle



Figure 6.7: Each surface sample is displaced in the direction of its normal until intersecting the plane orthogonal to its normal, producing an anisotropic smoothing of the sampled surface

- smaller than θ_{ν} with respect to the normal of the sample under consideration: $\hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_j > \theta_{\nu}$.
- 2. Estimate the mean position $\boldsymbol{\mu}_i$ of ν_i . Set the new value $\mathbf{x}_i^o := \mathbf{x}_i + \alpha \hat{\mathbf{n}}_i$, with $\alpha = (\boldsymbol{\mu}_i \mathbf{x}_i) \cdot \hat{\mathbf{n}}_i$.
- 3. If \mathbf{x}_i^o does not project onto any silhouette contour, repeat the previous step after setting $\alpha := \frac{4}{5}\alpha$.

Had we chosen to apply an isotropic smoothing technique, clusters in regions with higher sampling density and gaps in their emptier counterparts would have appeared. An illustration of an ideal case where the displacement results in a perfect placement of the surface sample –while projecting onto at least one silhouette contour– can be found in Fig. 6.7.

6.5.3 Surface coloring

Exploiting photo-consistency in this sampling approach would present the same type of difficulties as in the deformation-based approach from the previous chapter. The explicit computation of the visibility of each surface sample in order to impose global photo-consistency might be too costly for real-time applications, although a similar method to that in Section 5.4 could be applied to that respect.

However, in order to add color information to the reconstructed surface samples, a two-pass approach can be adopted, which will not provide photo-consistent surfaces but will contain important texture information that can be used in real-time applications. Such approach is based on the photo-consistency method in Section 5.4, but removing the possibility of changing the position of any surface point, thus exiting after the first iteration. Therefore, in first place a depth map for each view is obtained by projecting each surface point onto the corresponding viewpoint –with a sufficient area, in order to avoid hidden surface points to be visible through gaps between visible ones– and storing the closest depth value of each pixel.

In the second pass, we can determine the visibility of each surface point from



Figure 6.8: Given the normal (orange arrow) of a surface sample \mathbf{x}_i , its visibility from each view (computed using depth maps) and a set of cameras with orientation vectors $\hat{\mathbf{z}}_c$ (blue arrows), the three best oriented cameras viewing the sample can be selected

each available view by projecting it one more time and comparing its depth to that of the corresponding pixel in the depth map. Using this information, we collect the colors of the 3 best oriented cameras where a surface point is visible. The choice of these 3 cameras is based on the angles between the surface point normal and the $-\hat{\mathbf{z}}_c$ vectors of each camera c's coordinate system expressed in world coordinates. For each camera, $-\hat{\mathbf{z}}_c$ is a vector pointing towards its center of projections –Appendix A– in a direction orthogonal to the image plane. Let $\hat{\mathbf{n}}_i$ be the normal of an oriented surface point. Then, $\omega_{i,c} = -\hat{\mathbf{z}}_c \cdot \hat{\mathbf{n}}_i$, is taken as orientation score of a camera c for the point \mathbf{x}_i . The selection of the three best oriented cameras for a given point is illustrated in Fig. 6.8. The final color of the surface point is assigned as a weighted average of the colors sampled from the three best oriented views, with the weight of camera c defined as $w_{i,c} = \omega_{i,c} / \sum_{c'=1...3} \omega_{i,c'}$.

6.6 Conclusions

In this chapter we have presented an alternative sampling strategy to obtain the position and orientation of samples of the surface of a visual hull that outperforms the other two sampling approaches presented in the previous chapters in terms of computational load and ability to scale to a large number of input views, as it will be seen in Chapter 7, where all the presented sampling strategies are validated.

This is accomplished thanks to an alternative search for surface samples driven by statistical principles that aim at exploiting the spatial correlation on the location of surface samples. The main advantage of this approach is that, although it initially generates a sparse surface sampling with a low search efficiency –random search–, an efficient search strategy in regions very likely containing surface samples around existing ones provides a fast dense sampling of surfaces. Among its advantages, this technique allows to define exactly the number of samples in the reconstructed surface. This feature improves the ability of this sampling strategy –compared to the one presented in the previous chapter– of providing surfaces that, once re-projected onto the original viewpoints, accurately describe the original images. Furthermore, the algorithm can also exploit correlations in the temporal axis. The inefficient initial search for a sparse sampling of the surface is not required from frame to frame. Instead, a search over a suitably smaller range around each sample can be used. In order to improve the efficiency of the initial scouting stage when no information about previous time instants is available, a multi-resolution technique is also presented. The gains achieved with these two enhancements for the scouting stage (dynamic and multi-resolution) will also be presented in Chapter 7.

This technique still presents a limitation when compared to the image-based approach presented in Chapter 4: the application of photo-consistency constraints still requires costly global visibility computations, like in the approach presented in Chapter 5. However, in comparison to the technique presented in Chapter 5, this new algorithm is easy to parallelize, which makes it suitable to its application in GPU contexts. Furthermore, this sampling strategy is the fastest among the three presented in this thesis, showing the smallest computation time for a given number of surface samples, especially when considering larger multi-view settings, where the efficiency at exploiting spatial correlations of the proposed method shows up more clearly.

Chapter 7

Surface Sampling Results

n this chapter, the sampling methodology proposed in this first part of the thesis is validated by evaluating the results obtained when applying each of the techniques presented in the previous three chapters to different sets of multi-view static scenes and video sequences. Some of the results are purely qualitative, depicting the apparent improvement in quality of the proposed methodology with respect to some classical methods. A quantitative evaluation of the new techniques is also included, regarding different aspects relevant to their performance, such as the computational load and geometric or photometric accuracy.

This chapter is divided in four sections. The first one presents photo-consistent surfaces obtained by using the image-based sampling strategy introduced in Chapter 4. The next one presents silhouette-consistent surfaces obtained from both static multi-view scenes and dynamic multi-view sequences, using the surface deformation technique from Chapter 5. The third section presents colored, silhouette-consistent surfaces from multi-view sequences by using the methodology described in Chapter 6. Finally, a comparison of these three sampling strategies is presented.

7.1 Image-based Surface Sampling

The technique presented in Chapter 4 is able to reconstruct the visible surfaces of foreground objects from a very sparse set of images. An illustration of some of its stages is shown in Figure 7.1. It can be summarized as:

1. First, obtaining candidate depths for the image points in each view, based on photo-consistency measurements between pairs of pixels from neighbor cameras, using epipolar constraints.



Figure 7.1: Image-based surface sampling. From left to right: two sample input images out of six available; initial depth estimates for one of the input viewpoints; reconstructed surface points; reconstructed surface rendered with colored colored surface samples from two novel viewpoints

- 2. Then, surface points are extracted by iteratively carving away those points which are not globally photo-consistent.
- 3. Finally, a normal for each surface point is estimated by fitting a plane to a local neighborhood.

In the following, two different types of results are presented. The first one is a qualitative comparison between the results obtained with both typical small baselines that are considered in multi-view stereo scenarios and more challenging ones obtained by dramatically increasing the baseline –reducing the number of views– in the same scenes. These results are accompanied by examples of surface samples from static scenes in real-world scenarios with known –and removable– background in challenging wide-baseline configurations.

The second type of results, focusing on the target scenarios, present a quantitative measure of the quality improvement that can be achieved by a surface sampling strategy, compared to a volumetric sampling strategy, under equivalent conditions.

7.1.1 Qualitative validation

The datasets listed in Table 7.1 have been used as input in our experiments. In the table, apart from the number of available input images and their resolution, we have also included the maximum allowed squared RGB distance for photo-consistency – manually adjusted– that has been used for each dataset, ρ .

The first two datasets, *temple* and *dino*, are part of the benchmark promoted by [Seitz et al., 2006]. In the *temple* dataset, one of the 16 available images has been removed, because it was repeated. The datasets *dino-s* and *temple-s* are obtained by just picking 6 of the available views. This produces a very sparse scenario that helps us demonstrating how the sampling strategy is actually able to extract photoconsistent surfaces under such adverse conditions.

The *persons* dataset corresponds to 5 views of two persons captured at the *smart* room facility of the Signal Processing Group of the UPC, whereas the *shirt* dataset


Figure 7.2: Sample results on the used datasets. (a) *temple*, (b) *dino*, (c) *temple-s*, (d) *dino-s*, (e) *persons* and (f) *shirt*. In each case, one of the input images is shown, along with the initial estimate of one of the depth maps corresponding to one of the input views and the reconstructed surface samples seen from a novel viewpoint

consists of 6 views of a person in a smart room environment in the *Telefonica* R & D facilities in Barcelona.

The only relevant parameter for obtaining correct reconstructions in all the scenarios is the aforementioned photo-consistency threshold ρ . For example, in datasets where the color captured by different cameras has not been calibrated, as in *persons*, this threshold must be loose enough to allow the reconstruction of almost complete surfaces. Please note that the higher the threshold, the looser, since our photo-consistency test (Section 4.2) measures Euclidean color distances. Whereas a high threshold implies the acceptance of some outliers as surface elements, for us it is more important to obtain an almost complete surface, with both visualization and analysis applications in mind.

The results obtained from testing our algorithm with the datasets in Table 7.1 are shown in Fig. 7.2. For each dataset, we have included one of the input views, an initial depth estimate for one of the viewpoints and the resulting surface rendered with colored surface samples viewed from a novel viewpoint.

The method shows a relatively fast performance, mainly when compared to

Name	Views	Image Size	ρ
temple	15	640×480	100
dino	16	640×480	30
temple-s	6	640×480	100
dino-s	6	640×480	30
persons	5	768×576	3000
shirt	6	640×480	1000

Table 7.1: Datasets for the qualitative validation of the image-based sampling approach. The large *temple* and *dino* datasets are included for reference. The only parameter that needs tuning is the photo-consistency threshold ρ over the squared Euclidean RGB distance

multi-view stereo techniques, with computation times ranging from more than 3 minutes for the complete *dino* dataset (16 views) down to a few seconds for the *persons* dataset. These times are obtained on an Intel Xeon CPU at 3.0 GHz.

The reconstructions with the *temple* and *dino* datasets are useful for comparative purposes. The very sparse datasets (those with less than 10 input viewpoints placed all around the scene), which are prohibitive for multi-view stereo techniques due to the wide baseline, can be handled with the proposed sampling strategy. Indeed, it trades-off geometric or photometric accuracy with spatial coverage of the scene.

7.1.2 Quantitative results

The neighbors of each view are automatically obtained by taking the 3 closest viewpoints. The proximity of two viewpoints is defined as $\pi_{ij} = \pi_{ji} = \left(1 - \frac{1}{2}\cos(\alpha_{ij})\right)\delta_{ij}$, where α_{ij} is the angle between the two unitary vectors orthogonal to the image planes *i* and *j* and δ_{ij} is the Euclidean distance between the centers of projections of the two views. Photo-consistency thresholds have been adjusted empirically.

The dataset used for the evaluation is composed by 33 static multi-view scenes captured in a controlled environment by a set of 18 cameras placed around the scene of which we will make subsets for the experiments. Eight of these cameras offer partial views of the objects of interest in the scene, whereas two other cameras are reserved as ground-truth for indirect evaluation of the reconstruction quality. The datasets for evaluation are obtained by picking some of the views for reconstruction. As a result, three datasets with 6, 8 and 16 cameras respectively are available for our experiments. The corresponding foreground silhouettes are automatically extracted and show misclassified pixels. The foreground silhouettes of the two ground-truth views are manually extracted.

The evaluation procedure consists in using the 6, 8 or 16 available views for reconstructing the photo-consistent shapes of the foreground elements of the scene,

	F-measure		PSNR	
# Views	Surface patch	Voxel	Surface patch	Voxel
6	0.85	0.84	40.64 dB	$34.77~\mathrm{dB}$
8	0.87	0.86	$41.67~\mathrm{dB}$	$37.08~\mathrm{dB}$
16	0.88	0.88	$43.62~\mathrm{dB}$	$39.94~\mathrm{dB}$

 Table 7.2:
 Evaluation of image-based sampling and volumetric photo-consistent reconstruction

 schemes considering a growing number of input views for the experiments illustrated if Fig. 7.3

projecting them onto the two unused viewpoints and extracting the performance measurements from comparison with the ground-truth. The measured magnitudes are:

- *F-measure*, or the harmonic mean of *precision* (rejection of false foreground pixels) and *recall* (rejection of missing foreground pixels), measured on the silhouettes,
- *PSNR* of the pixels that are part of the foreground on both the ground-truth and testing silhouettes.

A volumetric equivalent of the proposed method has been implemented, which also uses the proposed photo-consistency test and the visual hull-bounding of the final photo-consistent reconstructed shapes. In order to compare the results, we follow the equivalency condition between voxels and image sampling density proposed in a previous work [Casas and Salvador, 2006]. Thus, given a sampling distance in each major direction x and y used on the available input images, $\sigma_{\{x,y\}} = 4$ pixels in our experiments, we define its equivalent voxel size, measured in units of the coordinate system of the scene.

The average results obtained by comparing the projections of the reconstructed shapes with the two ground-truth images for the surface patch-based and voxelized methods for each of the three considered datasets are shown in Table 7.2. As it can be observed, the accuracy of the proposed method is higher than that obtained by a voxelized approach, both in terms of shape (F-measure) and color (PSNR). Regarding the latter, our measurements confirm that surface color is not properly represented by the inaccurate surface description provided by a volumetric reconstruction. When analyzing the results, please keep in mind that the ground-truth images have not been used for reconstruction. Some results for the evaluated scenes are shown in Fig. 7.3, including geometric and textured renderings of the reconstructed shapes from novel viewpoints. The computation time of the proposed surface patch-based method is within one order of magnitude larger than that of its volumetric counterpart (ranging from fractions of a second for the datasets with 6 views up to more than ten seconds for the datasets with 16 views).



Figure 7.3: Reconstruction results for 8 cameras. The groundtruth image has not been used for reconstruction. (a) Groundtruth image masked with its manually extracted silhouette. (b) Reconstructed geometry represented by surface patches. (c) Colored surface patches. (d) Voxelized colored surfaces with equivalent resolution

7.2 Surface Sampling by Deformation

The methodology for extracting silhouette-consistent surface samples by the deformation of an existing surface is presented in Chapter 5. As an advantage with respect to the image-based approach, it allows the exploitation of temporal redundancies, following a tracking strategy. As usual among tracking strategies, two separate stages provide initialization and tracking:

- *Initialization: Static reconstruction.* First, samples of a silhouette-consistent surface are obtained for an initial time instant –without having any knowledge of the shape at previous time instants– by exploiting only spatial redundancies.
- *Tracking: Dynamic reconstruction.* Then, the knowledge of the shape in the previous time instant allows for a more efficient sampling of the surfaces in the next time instant, by exploiting both spatial and temporal redundancies.

Three experiments are used to validate the proposed method. The first experiment checks the consistency of the shapes resulting from dynamic reconstruction, compared to their equivalents obtained by static reconstruction, in order to validate the faster former method. The second experiment evaluates the usage of computational resources of dynamic reconstruction when compared to a classical volumetric

Sequence	Time static	Time dynamic	RMS distance
walking	$3.43 \mathrm{\ s}$	$1.07 \mathrm{~s}$	0.000102
open arms	$3.77 \mathrm{\ s}$	$1.17 \mathrm{~s}$	0.000134
move arm	$3.45 \mathrm{\ s}$	$0.93 \mathrm{\ s}$	0.000100
kick	$3.45 \mathrm{\ s}$	$1.23 \mathrm{~s}$	0.000094
jump	$3.41 \mathrm{~s}$	$1.16 \mathrm{~s}$	0.000119

 Table 7.3: Average RMS Hausdorff distance of surfaces obtained by dynamic reconstruction vs.

 static reconstruction. The maximal edge length of the associated bounding box is 0.101 in all cases

shape-from-silhouette algorithm. The third experiment quantitatively evaluates the accuracy of the method, using a synthetic scene for which we have 3D ground-truth data.

The synthetic scene is obtained as the projection of an animated human-like 3D model onto a number of virtual viewpoints. Therefore, the reconstructed surfaces can be compared to those of the actual model in every time instant.

For the other experiments, five real sequences with known background have been used, with 50 frames each. The silhouettes of a person captured by 8 cameras placed all around are extracted using [Stauffer and Grimson, 1999]. These sequences *-walking, open arms, move arm, kick* and *jump*-, show different types of motion, ranging from fast motion of limbs to slower displacement of the person.

The initial surface for the tracking strategy is defined by 10^5 surface points randomly distributed over a box-shaped surface with dimensions slightly larger than that of the bounding box of each scene; the diameter for normal estimation $2 \times r_n$ is set to 1/100 of its main diagonal and the expansion distance is set to $r_t := 5 \times r_n$. In order to reconstruct shapes when parts of the scene are not included in the frusta of all views, a minimal inclusion in 4 frusta is required. All measurements have been obtained from an Intel Xeon 3.0 GHz.

7.2.1 Static vs. dynamic reconstruction

Fig. 7.4 shows the reconstructed surfaces with color-encoded estimated velocities for their samples. The sequences *walking*, *open arms* and *jump* show displacement of large surfaces in scenes with displacement of large surfaces, whereas *move arm* and *kick* show fast motion of limbs in scenes with displacement of small surfaces. The velocity estimated by the dynamic reconstruction method shows outliers in areas where surfaces collide, but, overall, it plausibly describes the apparent surface motion.

In order to quantitatively validate the dynamic method as an efficient alternative



Figure 7.4: Surface samples from sampling by surface deformation and color-encoded velocity estimates for (a) *walking*, (b) *open arms*, (c) *move arm*, (d) *kick* and (e) *jump*. The RGB color code for velocity follows the included axes, although the normal information included for better visualization of shape under illumination slightly distorts its actual intensity. White colored points correspond to surfaces at initial time instants (static reconstruction)

to static reconstruction, each of the scenes have been processed with both methods, and the average RMS Hausdorff distance [Aspert et al., 2002] is measured between two meshes obtained by applying *Poisson reconstruction* [Kazhdan et al., 2006] on the sets of oriented surface points. The results in Table 7.3 reflect that shapes are practically equivalent and the processing times are clearly and consistently reduced.

7.2.2 Usage of resources

In Table 7.4, the memory usage for two configurations with approximately equal processing times shows how the reconstruction of surfaces with the proposed method needs less computational resources than the voxelized alternative and provides a

Method	Initial pts	Surface pts	Memory
Voxel	$121\times 300\times 311$	53706	$700 \mathrm{MB}$
Dynamic	$5 imes 10^5$	84962	$495~\mathrm{MB}$

Table 7.4: Memory usage of surface sampling by deformation compared to that of volumetric shape-from-silhouette with approximate equal processing times of 2.2 s and 2.3 s respectively

higher level of detail. This can be visually confirmed in Fig. 7.5, where triangular meshes have been extracted from both voxelized shape-from-silhouette and the proposed method with *marching cubes* [Lorensen and Cline, 1987] and *Poisson reconstruction*, respectively. The memory usage and execution times in the voxelized approach are influenced by the implementation strategy of pre-computing the projections of all voxel centers and storing the pixel coordinates in *Look-Up Tables*, which greatly shortens the processing time.

7.2.3 Accuracy

Synthetic sets of silhouette sequences are generated by the projections onto a set of views of an animated mesh model. We utilize the average RMS Hausdorff distance between a *target mesh*, which is obtained either from voxelized shape-from-silhouette followed by marching cubes or from dynamic reconstruction followed by Poisson reconstruction; and a *sampled mesh*, which is the initial mesh model. The results are summarized in Table 7.5. As expected, dynamic reconstruction outperforms the voxelized shape-from-silhouette method in terms of accuracy, because the 3D sampling at regular positions that takes place in the voxelized method cannot accurately describe surfaces, and it also confirms that the dynamic reconstruction of shape and motion achieves a correct estimate of the visual hull from frame to frame, as suggested by the first experiment.



Figure 7.5: Comparison of the reconstructed surface with voxelized shape-from-silhouette (left) and the proposed method (right) for approximately equal processing times of 2.3 s and 2.2 s, respectively

Sequence	RMS dist. volumetric	RMS dist. dynamic
walking	0.001286	0.000664
open arms	0.000902	0.000447
move arm	0.000873	0.000587
kick	0.001100	0.000745
jump	0.000910	0.000517

 Table 7.5:
 Average RMS Hausdorff distance for approximately equal processing times using volumetric shape-from-silhouette and dynamic sampling by deformation. The maximal edge length of the associated bounding box is 0.101 in all cases

7.3 Statistical Surface Sampling

The statistical surface sampling strategy presented in Chapter 6 is designed to efficiently obtain silhouette-consistent surface samples by directing the search to regions where a high likelihood of finding surfaces exists. This is achieved by a two-staged approach:

- First, an initial *scouting* procedure in the working volume, obtains a small amount of *seed* surface samples, which are costly to find but constitute a small percentage of the final set of surface samples.
- Then, a propagation strategy redirects the search to suitably defined regions around each seed sample, and quickly covers the complete surface until reaching the desired number of surface samples.

The bottleneck of the algorithm is the initial scouting, which takes the most of the computation time, due to the blind search that takes place in a large volume. Therefore, two strategies for reducing the computation time required to obtain the seed surface samples have been proposed:

- The first one, *dynamic* statistical sampling, exploits temporal correlations in multi-view sequences, which translates in a reduced search region for seed samples from frame to frame and posterior propagation.
- The second strategy introduces a *multi-resolution* technique, which performs a fast greedy search for surface points from heavily decimated input that is followed by a search for seed surface samples in small regions around each low-resolution surface point in a manner similar to that in the dynamic strategy.

The algorithm is parameterized by the number of output surface samples –kept constant as 100000 in our experiments– and the size of the different search regions used at every stage. All of these parameters, including those only used for the

Parameter	Value
Bounding box volume	$s_x \times s_y \times s_z$
Fine normal estimation radius ρ_n	$\frac{1}{200}\sqrt{s_x^2+s_y^2+s_z^2}$
Initial propagation radius ρ_p	$2\rho_n$
Local normal estimation radius ρ_l	$\frac{1}{2} ho_n$
Dynamic scouting radius ρ_d	$8 ho_n$
Multi-resolution scouting radius ρ_m	$2\rho_n$
Number of samples N_s	100000

 Table 7.6:
 Parameters for the statistical surface sampling, including those for dynamic and multi-resolution sampling. Note that the number of surface samples is restricted to 100000

improvement of the *scouting* stage using one of the two proposed alternatives, have been empirically determined in order to provide the expected behavior and are listed in Table 7.6.

For the experiments, three multi-view real sequences from datasets available in [4D Repository, 2010] and a synthetic one -kung-fu girl [Kung-Fu Girl, 2005]– have been used, which are listed in Table 7.7. As it is shown, two of the real sequences are captured by 16 views placed all around the scene, whereas the remaining one contains only captures from 8 views. The synthetic sequence is captured by 25 virtual cameras placed on a hemisphere around the animated model.

Three experiments have been performed. The first two experiments demonstrate the efficiency increase that can be achieved by applying any of the two types of proposed improvements to the initial scouting for seed surface samples, e.g., the dynamic sampling and the multi-resolution sampling. The third experiment shows a property about the scaling of the proposed surface sampling strategy to large multi-view settings that makes it ideal to be used in such conditions. All the results are obtained from a single-threaded implementation running on an Intel Core 2 Duo 2.8 GHz processor.

Sequence	# Views	# Frames	Resolution
dancer	8	201	780×582
children	16	339	1624×1224
martial	16	210	1624×1224
kung-fu girl	25	200	320×240

 Table 7.7:
 Three multi-view real sequences from the 4D repository and a synthetic sequence (kung-fu girl) used for the experiments related to statistical sampling

Sequence	Default	Dynamic
dancer	$1.02 \mathrm{~s}$	$0.46~{\rm s}$
children	$2.73 \mathrm{~s}$	$0.92~{\rm s}$
martial	$2.35 \mathrm{~s}$	$0.88 \mathrm{\ s}$
kung-fu girl	$1.4 \mathrm{~s}$	$0.56~{\rm s}$

Table 7.8: Per-frame computation time with default and dynamic scouting for statistical sampling

7.3.1 Multi-resolution and dynamic sampling

In this section we present two experiments that demonstrate the increase in efficiency that can be achieved when modifying the initial scouting stage with the two proposed improvement types.

In Section 6.4.1, a dynamic scouting strategy has been presented, which reduces the computation time in multi-view video sequences by exploiting temporal correlation. This technique has been used for the four sequences. The average computation times for each sequence with and without dynamic scouting are listed in Table 7.8. Inter-frame similarity of the position of the surfaces of foreground objects can be used for reducing the search region for seed surface samples, which results in a clear increase in frame-rate.

In Section 6.4.2, a multi-resolution approach has been proposed to improve the efficiency of the scouting stage. The most important adjustable parameter is the decimation level in the input images. In Fig. 7.6, the execution times for sampling with different levels of initial decimation for the input silhouettes are shown. Note that the case with decimation set to one corresponds to the default scouting strategy.



Figure 7.6: Reconstruction time vs. initial level of image decimation in multi-resolution scouting. A decimation of 1 is equivalent to the default scouting with full-resolution images



Figure 7.7: From left to right, statistical sampling results for the *kung-fu girl* sequence for initial decimation of 1, 2, 8 and 32 in each image dimension. Decimation of 1 is equivalent to default scouting. Incorrect values for the search radius lead to uneven distribution of seed surface samples, which cannot be correctly balanced during the propagation stage

The resulting efficiency gain obtained by both dynamic and multi-resolution scouting is similar. The extra cost from initially obtaining seed samples at low resolution in multi-resolution scouting is partly balanced with the fact that the lowresolution seeds are in average closer to the actual surface than the high-resolution seeds from a previous time instant in dynamic scouting, which translates in the choice of a multi-resolution scouting radius smaller than its dynamic scouting counterpart.

In Fig. 7.7, surface samples corresponding to a frame from the *kung-fu girl* sequence are shown, as obtained when using the multi-resolution scouting strategy with different levels of initial decimation. Due to an incorrect choice of the search radius for high-resolution seed surface samples, most of the low-resolution samples are discarded by the impossibility to find a high-resolution matching sample. The propagation mechanism is not designed to correct such situation, resulting in a highly uneven distribution of samples.

7.3.2 Efficiency in multi-camera environments

A limiting problem when working in multi-camera environments is the cost associated to the addition of each new camera. Typically, algorithms exploiting multi-view data can grow linearly –i.e., in volumetric reconstruction techniques– or even exponentially –i.e., in image-based approaches where multiple cross-validations between images take place– with the number of available cameras.

However, the proposed strategy for sampling benefits from the fact that the ratio between contour pixels and total number of pixels remains practically constant with respect to the number of available cameras when these are placed at similar distances



Figure 7.8: Reconstruction time *vs.* number of cameras for three real sequences from [4D Repository, 2010] and a continuous sequence from [Kung-Fu Girl, 2005] (all in continuous lines). The dash lines show the computation time that would be required for sampling if the computational cost of adding a view was constant (as in volumetric shape-from-silhouette). The details of the sequences are in Table 7.7

from the objects of interest. Even better, the ratio between silhouette contour pixels and occupied volume increases with the addition of each new camera. Thus, even if the cost of a projection test grows linearly with the number of views, the probability that a 3D point is a surface point also grows, partly balancing this effect. As shown by the following experiment, the result is that the proposed sampling strategy is well suited to large multi-camera settings.

Fig. 7.8 shows, in dashed lines, the extrapolated computation time from the first five up to the whole range of available views. As shown, the actual computation time grows at a much smaller rate for subsequent views, due to the introduction of a larger number of silhouette contour pixels. In Figs. 7.9 and 7.10, the reconstructed surfaces for one time instant of each sequence with increasing numbers of input views are shown as a set of surface samples or its interpolating mesh obtained with the algorithm presented in Part II, respectively. The marginal increase of detail in the resulting surface becomes smaller, and so does the marginal cost of adding each additional view. This shows the efficiency of the proposed sampling strategy.

7.4 Comparison of the Three Sampling Approaches

The three sampling strategies presented in this part of the thesis are now compared in terms of computation time and visual appearance of the resulting surface. The video sequence used for this comparison is *dancer*, from [4D Repository, 2010], listed in Table 7.7.



Figure 7.9: Statistical sampling with increasing number of views. The leftmost column shows surface samples reconstructed from only 5 viewpoints of each sequence. Towards the right, surface samples obtained with increasing number of views. The rightmost column shows the surface samples reconstructed from the whole multi-view set for each sequence



Figure 7.10: Statistical sampling with increasing number of views. The continuous surface is obtained with the algorithm presented in Part II. The leftmost column shows surfaces reconstructed from only 5 viewpoints of each sequence. Towards the right, surfaces obtained with increasing number of views. The rightmost column shows the surfaces reconstructed from the whole multiview set for each sequence

Method	No. of samples	Time	Initialization time
Image-based	56502	$167.34~\mathrm{s}$	-
Deformation	56610	$2.5 \mathrm{~s}$	$8.62 \mathrm{~s}$
Statistical	56610	$0.25 \ s$	$0.63 \mathrm{\ s}$

 Table 7.9: Average computation time per frame (in seconds) for sequence dancer using the three

 sampling strategies with approximately equal number of reconstructed surface samples

In Table 7.9, the average computation time per frame using each of the proposed techniques is listed. For the sampling strategy based on the deformation of an existing surface, the obtained time corresponds to that of the tracking strategy (initialization 8.62 s). For the statistical sampling strategy, the time in the table corresponds to the application of the dynamic improvement in the *scouting* stage (initialization without multi-resolution decimation 0.63 s). The number of surface samples has been determined by the image-based strategy, with the deformation-based strategy configured in order to obtain a value as close as possible to the former and the statistical sampling configured in order to retrieve exactly the same number of samples as the highest one among the other two strategies.

The average computations times confirm the expected behavior of the three strategies. Whereas the image-based approach does not exploit any sort of spatial or temporal correlation on the distribution of surface samples and the strategy based on the deformation of an existing surface exploits mainly the temporal correlation (and indirectly also spatial during the regularization algorithm), the statistical sampling strategy is the one that best exploits these two types of correlations outperforming the other two methods in terms of computation time for a certain sampling density.

Fig. 7.11 shows the continuous surface corresponding to one time instant of the *dancer* sequence, obtained using *ball pivoting* [Bernardini et al., 1999] over the set of surface samples reconstructed with each of the three sampling strategies. Whereas the surface obtained with the image-based strategy is the closest to the actual surface –see the narrower head, achieved by imposing photo-consistency–, it also presents holes –like those on the chin of the dancer–. The surface obtained from the deformation-based strategy is practically equivalent to that obtained with statistical sampling, with a clearly higher presence of high-frequency detail and, most prominently, noise on the distribution of surface samples in the former, avoided by the application of anisotropic smoothing (Section 6.5.2 in Chapter 6) as postprocessing in the latter.

In general, considering the processing time and the resulting topological correctness of the statistical sampling strategy, the presented results assess this strategy as the best of the three proposed strategies in terms of the goals of this thesis. Indeed, it is the approach that takes a smaller processing time –close to real-time opera-



Figure 7.11: Continuous surface for a time instant in the video sequence *dancer*, obtained by applying *ball pivoting* on the set of approximately 56000 surface samples reconstructed by the image-based, the surface deformation and the statistical sampling strategies

tion using a single core of the CPU, while being easy to parallelize– and provides a closed surface, free of holes. Whereas the posterior application of *ball pivoting*, in order to provide a continuous triangulated surface, renders this approach useful for off-line processing only, the surface interpolation algorithm presented in Part II of this thesis tackles the problem of tessellating a continuous surface between samples targeting speed and topological correctness.

7.5 Conclusions

The experiments presented in this chapter provide evidence of the high accuracy level that can be achieved by obtaining surface reconstructions using the techniques introduced in the previous chapters of this first part of the thesis. Different types of experiments have been conducted using each of the three presented surface sampling approaches, as well as a final comparison of the three methods.

Using the image-based surface sampling approach, experimental results show that, in order to obtain an equivalent resolution level in a volumetric methodology, a much higher amount of computational resources would be necessary. Compared to multi-view stereo methods, the photo-consistency test proposed in Chapter 4 proves useful in arbitrarily wide-baseline scenarios. This technique is the slowest among the presented ones, but it can be used to provide photo-consistent reconstructions when a small number of views is available.

With the surface sampling strategy by surface deformation, experiments show that the method allows taking advantage of temporal redundancies in dynamic scenes, which are the source of a great reduction of the required computation time and can be used to provide an explicit description of motion that can be used by higher level algorithms as an additional set of features detected from the raw data, as introduced in Chapter 5. This technique is much faster than the previous one, due to the exploitation of temporal correlation, and can be, therefore, used with a larger number of views if the latency due to the slower initialization stage can be allowed in a specific real-time application.

Finally, using the statistical sampling strategy presented in Chapter 6, experiments show that it is highly efficient with respect to large number of input views, yet capable of obtaining reconstructions in very wide-baseline configurations, thanks to the proper exploitation of spatial correlations. Furthermore, the exploitation of temporal correlations is also introduced as a mechanism to improve the initial scouting stage for seed surface samples. Similarly, a multi-resolution technique can also be designed, which produces similar gains in processing time when no information from the current position of surfaces is available. When a large number of views is to be processed for real-time applications, this would be the algorithm that best fits the goals of the thesis presented in Chapter 1, thanks to the more principled exploitation of both spatial and temporal correlations.

As the final comparison between the three strategies shows, the statistical sampling approach is the fastest, resulting from the better exploitation of spatial correlation, and provides a topologically correct, closed surface. Due to the high degree of parallelization of the latter algorithm, it is also the best suited for GPU implementations. Indeed, a Master's Thesis developed by Marc Maceira and co-directed by my advisor and myself has recently resulted in a real-time implementation with speed-ups of around $5\times-7\times$, with room for further improvements.

Part II

Surface Interpolation

Surface Interpolation

hereas the previous part of the thesis copes with the definition of strategies for the reconstruction of sets of surface samples corresponding to foreground objects in multi-view video sequences, this second part tackles the problem of retrieving a closed, continuous surface – presented as a triangle mesh– in order to interpolate the information contained in such surface samples.

Again, the design of the corresponding methodology is driven by the goal of obtaining fast methods applicable in interactive applications, while offering topological guarantees in order to keep the resulting continuous surfaces usable for forthcoming visualization and multi-view analysis applications.

In the two chapters composing this second part of the thesis, we first present the details of the proposed method for surface interpolation –meshing algorithm– and, then, we compare its performance to that of other meshing methods in the literature, by using well-known datasets of surface samples obtained from range data. The application of the method proposed in this second part of the thesis on the sets of surface samples obtained by the best sampling approach in the first part is left for Part III, where a validation for the complete surface reconstruction approach –including both sampling and interpolation– is presented.

Chapter 8

Surface Interpolation

set of surface samples constitutes a minimal discrete representation of a surface. In order to obtain a complete reconstruction of the surface, an interpolation scheme must be provided, which tessellates a continuous surface between its samples. Due to the possible irregularity resulting from the sampling scheme, interpolating a surface between these samples is a nontrivial problem that requires special attention, as it will be explained below. In this chapter, a novel meshing algorithm is introduced, which produces a two-manifold mesh from a set of surface samples. Such a mesh is useful for visualization purposes, but also for *analysis-by-synthesis* methods that might exploit convenient topological properties of 2D manifolds in 3D space.

8.1 Motivation

Surface samples can be used for visualization, without the need for an explicit computation of a continuous surface out of the rendering pipeline, as it has been demonstrated with impressing results in [Pfister et al., 2000] or [Guennebaud et al., 2004]. However, a discrete representation of a surface does not provide a complete description of its topology, which is usually required when processing 3D data (in some applications through surface parameterization [Hormann et al., 2007]). Furthermore, a suitable continuous description of the surface can also simplify visualization by using the standard rendering pipelines in current graphics hardware.

As mentioned above, the problem of interpolating a continuous surface given a set of discrete surface samples presents several challenges. In first place, the problem can be computationally expensive when the number of surface samples is large. The problem is also ill-conditioned: in general, some assumptions about the topology of the surface have to be made in order to obtain a solution, which will in general not be unique. The combination of these problems translate to the main challenge, which is the definition of the ordering of the surface samples for surface interpolation. Furthermore, in order to be convenient for analysis and visualization, the interpolated continuous surface must guarantee certain topological properties, such as being closed and orientable.

8.1.1 Interpolation

Mathematically, interpolation can be described as a method of constructing new data points within the spatial range of a discrete set of known data points. It can also be interpreted as a specific case of curve fitting, in which the resulting curve goes exactly through the data points or samples.

In datasets where the independent variable –sample points x_i – is unidimensional, the order of the samples is uniquely defined. Thus, the problem in such case is just that of defining the *interpolant* function.

However, in datasets like the ones we deal with, the independent variable is a position in 3D space. Therefore, the order of samples, which will be used in order to generate the interpolant function, cannot be uniquely chosen, since a well-defined ordering of the samples does not exist. Indeed, the most challenging part of surface interpolation is precisely the definition of the neighborhood of each sample point, which in turn results in the explicit description of the surface topology. Whereas for analysis applications it is convenient to obtain a topologically correct continuous surface, obtaining such a surface in general introduces more complexity to the procedure.

The problem of defining the topology of the independent variable is therefore the main topic of this chapter. In the following, we justify the selection of an interpolant function by studying the different available interpolation types:

- Polynomial. The interpolant is a polynomial of a certain degree.
- Spline. The interpolant is a special type of picewise polynomial, or spline.

Polynomial interpolation

Given a set of n + 1 data points (x_i, y_i) where no two x_i are the same –in case of a surface, x_i would be the coordinate of vertex i and y_i some attribute related to such vertex, i.e., its color–, one is looking for a polynomial p of degree at most nwith the property

$$p(x_i) = y_i, \forall i \in \{0, 1, ..., n\}$$
(8.1)

The polynomial p is called the *interpolant*. Generally, if we have n data points, there is exactly one polynomial of degree at most n-1 going through all the data points in a certain order. The interpolation error is proportional to the distance between the data points to the power n. One advantage of this method is that the interpolant is a polynomial and thus infinitely differentiable in the range (x_0, x_n) , with $x_0 < x_1 < ... < x_n$.

However, polynomial interpolation has some disadvantages. Calculating the interpolating polynomial can be computationally expensive for large n. Furthermore, polynomial interpolation may exhibit oscillatory artifacts, especially at the end points, the so-called Runge's phenomenon. This disadvantage can be avoided by using spline interpolation, which is presented later.

In the following, two particular cases of polynomial interpolation are considered, which solve the computation cost problem that appears for large n at the cost of increasing the interpolation error. The first, with n = 0, is the so-called *nearestneighbor* or *piecewise constant* interpolation. The second, with n = 1, is the socalled *linear* interpolation.

Nearest-neighbor interpolation. The nearest-neighbor algorithm simply selects the value of the nearest sample, and does not consider the values of other neighboring samples at all, yielding a piecewise-constant interpolant. The algorithm is very simple to implement, and has been widely used in real-time 3D rendering to select color values for a textured surface. In this case, the differentiability property is lost.

Linear interpolation. Linear interpolation on a set of data points is defined as the concatenation of linear interpolants between each pair of data points. This results in a continuous curve, with a discontinuous derivative, thus of differentiability class C^0 .

This particular case of polynomial interpolation is widely used in surface reconstruction, since it provides a sufficient level of accuracy for visualization and also for analysis, given a dense enough sampling of the surface, as it is shown in Chapter 10 in Part III. Furthermore, the resulting surface representation (a polygon mesh), can be processed by the standard rendering pipelines in graphics hardware. This type of interpolation is therefore the most suitable for interactive applications.

Spline interpolation

In the specific case of linear interpolation, a linear function for each of intervals $[x_k, x_{k+1}]$ is used as interpolant, whereas in the more general polynomial interpola-

tion a polynomial of high-degree n is used as the interpolant between a large number of data points or even the whole set of points. Spline interpolation uses low-degree polynomials in each of the intervals that would be used in linear interpolation, with the polynomial pieces chosen such that they smoothly fit in the resulting function.

Spline interpolation incurs a smaller error than linear interpolation and the resulting interpolant is smoother. It is also easier to evaluate than the high-degree polynomials used in the more general polynomial interpolation and does not suffer from Runge's phenomenon. However, the resulting interpolated surface, although showing a high level of accuracy for analysis, cannot be processed by the standard rendering pipelines in current graphics hardware with interactive visualization in mind.

Chosen methodology

A possible drawback in linear interpolation is the discontinuous derivative. However, given the fact that the input surface samples already contain information about the local orientation, this does not generate any practical limitation.

In the following, linear interpolation is taken as the most suitable option for its usage in analysis and visualization. Even though the interpolating error might be larger than that in spline-based interpolation, the possibility of using current graphics hardware for interactive visualization and the smaller computational cost associated to the extraction of the interpolant (a polygon mesh) makes us decide for this approach. Furthermore, a *spline* interpolation of a surface could also be obtained from a polygon mesh [Krishnamurthy and Levoy, 1996] at a post-processing stage if required.

About the ordering of the independent variable -i.e. the positions in 3D space of the surface samples–, different techniques exist in the literature, which produce linear interpolants which either pass exactly through the existing sample points or through close points defined by smoothing constraints on local sets of input samples. Some of these techniques are introduced in the next section, which are used to contextualize the necessity for the proposed method.

8.1.2 Related work

Many surface meshing techniques have been developed over the past years. The *marching cubes* (MC) algorithm was proposed long ago in [Lorensen and Cline, 1987]. This algorithm is a well-known method for extracting iso-surfaces from three-dimensional scalar fields. In practice, MC allows to extract a triangle mesh from a volumetric representation of a 3D scene.

Marching cubes-based methods

Many reconstruction techniques, like the ones presented in the first part of this thesis and others in the 3D reconstruction literature [Furukawa and Ponce, 2008], as well as range data from laser scans, do not provide a volumetric representation of 3D shapes, but rather a set of oriented surface samples. Hence, some approaches exist which, first, determine voxel occupancy by means of a certain type of regularization and, then, extract the surface by means of MC. Two state-of-the-art methods falling in this category are *RIMLS-MC* [Oztireli et al., 2009] and *Poisson-Rec* [Kazhdan et al., 2006].

RIMLS-MC. Moving least squares (MLS) is an attractive tool to design effective meshless surface representations. However, as long as approximations are performed in a least square sense, the resulting definitions remain sensitive to outliers, and smooth-out small or sharp features. In [Oztireli et al., 2009] these issues are considered and they present a novel point-based surface definition combining the simplicity of implicit MLS surfaces [Kolluri, 2005] with the strength of robust statistics, resulting in a Robust Implicit Moving Least Squares (RIMLS).

To reach this new definition, MLS surfaces are reviewed in terms of local kernel regression, opening the doors to the introduction of robustness. The authors claim that this representation can handle sparse sampling, it generates a continuous surface better preserving fine details, and can naturally handle any kind of sharp features with controllable sharpness. It also combines ease of implementation with performance competing with other non-robust approaches. Finally, a triangular mesh is extracted by using the Marching Cubes algorithm. RIMLS-MC has some desirable properties:

- It provides an accurate description of the surface, even in presence of noisy data.
- It shows better detail in edges, when compared to other marching cubes-based approaches.

However, it also presents some drawbacks:

- The method is relatively slow compared to other methods in the literature.
- The output surface, although topologically correct, can present holes, due to insufficient volumetric density.

Poisson-Rec. *Poisson Reconstruction* proposes a formulation that considers all the points at once, without resorting to heuristic spatial partitioning or blending,

and is therefore highly robust against data noise. As claimed by its authors, the Poisson approach presented in [Kazhdan et al., 2006] allows a hierarchy of locally supported basis functions, and therefore the solution reduces to a well conditioned sparse linear system. A spatially adaptive multi-scale algorithm is used, whose time and space complexities are proportional to the size of the reconstructed model.

In other words, surface reconstruction is expressed as a Poisson problem, which seeks the indicator function that best agrees with a set of possibly noisy, non-uniform observations. This approach can, for example, robustly recover fine detail from noisy real-world scans. At its last stage, an octree-based Marching cubes implementation extracts a triangular mesh from the volumetric representation. Therefore, Poisson-Rec has several desirable properties:

- It provides an accurate description of the surface, even in presence of noisy data.
- It allows choosing the level of detail in a trade-off with computational complexity.

However, it also presents some drawbacks:

- The precision around edges is not controlled.
- The output surface shows a limited amount of detail for similar computation times when compared to propagation methods.

Propagation methods

Propagation methods for surface reconstruction have in common the initial definition of one or more regions that are used as seeds for a propagation algorithm that attempts to cover the whole surface between the input 3D points. Two methods in the literature falling in this category are *Ball Pivoting* [Bernardini et al., 1999] and *ReOP* [Suau et al., 2010].

Ball Pivoting. The Ball Pivoting Algorithm (BPA), is an advancing-front algorithm to incrementally build an interpolating triangulation of a given point cloud. The method is conceptually simple. Starting with a seed triangle, it pivots a ball around each edge on the current mesh boundary until a new point is hit by the ball. The edge and point define a new triangle, which is added to the mesh, and the algorithm considers a new boundary edge for pivoting. BPA has several desirable properties:

- It is intuitive, since it triangulates a set of points by *rolling* a ball on the point cloud. The user chooses only a single parameter, its radius.
- It has a theoretical foundation, since it is related to alpha-shapes [Edelsbrunner and Mücke, 1994], and given sufficiently dense sampling, it is guaranteed to reconstruct a surface within a bounded distance from the original manifold.
- It keeps the positions of the original surface samples in the output mesh.

However, it also presents some drawbacks:

- It can produce surfaces with holes when sampling is not dense enough. In case of iterating with a larger radius for the pivoting ball, wrong topology might result from some local configurations.
- It is slow when compared to other methods, both based on a propagation scheme and on marching cubes.

ReOP. Restricted and Oriented Propagation is a method to reconstruct polygon meshes from oriented point sets based on a propagation through a voxelized space performed in order to find the closest neighbors of every data point. Propagation is done following a specific pattern which exploits reciprocity in neighbor finding. Every obtained pair of neighbors defines a new edge of the output mesh. ReOP has several desirable properties:

- It provides a quick reconstruction of the surface, which makes real-time applications conceivable.
- It keeps the positions of the original surface samples (or an average of 3D point cliques lying on the same voxel).

However, it also presents some drawbacks:

- The resulting mesh is composed by tetrahedrons rather than surface triangles, due to the 3D search strategy for neighbor points.
- The use of memory might be very high at high resolution, due to the necessity of regularly sampling space.

Goals

The goals of the method proposed in this chapter, compared to the features of those available in the literature, are:



Figure 8.1: In green color, a range search with specified radius r around a given point –marked in red– returns a single result. In blue, a nearest-neighbors search with k := 5 returns a set of neighbor points around the red point

- To preserve the position of the existing surface samples, like the propagationbased methods, in order to maximize the accuracy of the reconstructed surface when error-free data points are available. ReOP and BPA are therefore fulfilling this requirement.
- To obtain fast interpolation of surfaces with the goal of being usable in realtime applications. ReOP is fast, but does not guarantee topological correctness.
- To obtain a correct topology of the interpolated surface, making it usable for analysis applications. BPA, Poisson-Rec or RIMLS-MC can provide topologically correct surfaces, but are slow to be applied in real-time.

8.2 Methodology

The proposed algorithm follows the same principle as the ball pivoting algorithm, *i.e.* it is based in a surface propagation scheme, with the imposition of a set of rules to guarantee the topological correctness of the resulting surface.

First, the algorithm randomly selects a starting face and, then, it propagates on its contour, resulting in a growing region of triangular faces. The objective of the propagation phase is to find adjacent faces to, ideally, cover the whole surface with topologically correct triangles. As mentioned, a set of rules is checked whenever a face is to be added to the region, to ensure the resulting mesh is complete and manifold.

8.2.1 Spatial queries

Knowledge of local neighborhoods of each point and fast access to this information is a key factor of our proposal. A *kd-tree* decomposition of the 3D point cloud helps out in finding this information. As proposed in [Friedman et al., 1977], a kd-tree decomposition is an efficient structure –in terms of memory footprint and search speed– onto which *nearest-neighbors* and *range* queries are efficiently performed. The main concepts about this structure are presented in Appendix B, but here we focus on the two types of search –nearest-neighbors and range search– that are of interest for our proposed algorithm and are illustrated in Fig. 8.1. We recommend readers unfamiliar with this technique to look at Appendix B, in order to better understand the benefits of using such representation for efficient spatial queries.

In order to create the tree, instead of splitting the kd-tree in the axis of maximum spread as in [Friedman et al., 1977], we choose to split in that of maximum variance. This makes our the tree more robust against outliers in the sense that the efficiency of the search process is not affected by their presence. Other splitting rules might be applied, as suggested in [Maneewongvatana and Mount, 1999].

Nearest-neighbors search. In order to determine the composition of the neighborhood of a given surface point it is often necessary to obtain a list of the k nearest-neighbors to the given point. This is efficiently provided by a nearest-neighbors search in the kd-tree, as shown in Section B.2.1.

Range search. In some other cases, it is required to obtain all the points lying inside of a sphere of a determined radius centered at a given point. This search, which retrieves a list of points at a distance smaller or equal to r from the specified point, can also be efficiently implemented as a range search in a kd-tree, as shown in Section B.2.2.

8.2.2 Half edges

In order to obtain a topologically correct polygon mesh, it is necessary to check whether a newly created face is compatible with the set of existing ones. Considering triangular faces, with three edges forming the contour of each face, it is necessary to ensure that the two new edges created after each propagation step are compatible with the already existing edges that may connect the involved vertices.

One useful representation of the polygon mesh is the so-called *half-edge* structure [Kettner, 1999]. In this structure, only edges are represented, which connect the vertices of the mesh in a certain orientation. In Fig. 8.2, the creation of a half-edge



Figure 8.2: Evolution of a set of half-edges after the addition of a new triangular face at every step. The current contour half edges at each step are named as \mathbf{h}_{ij} , where the subscript indicates the source and sink sample points that are connected by the half edge. The corresponding sparse matrix H^t after every step t shows the set of current half edges. The ones that continue the propagation are those left in bold in the matrices

representation and the corresponding contour at each propagation stage is shown.

The half-edge representation of meshes used by our method is stored as a sparse matrix H where all elements are empty, except for those that correspond to an active half edge, as shown in Fig. 8.2. These occupied elements are the only ones that are actually stored in memory.

The labeling of each non-zero element of the matrix is set to the identifier of the triangular face that generated the half-edge, which is used for correcting the topology of the mesh, as presented later in Section 8.3.2.

8.3 Algorithm

The algorithm presented in the following returns a polygon mesh that can be used as a linear interpolant of the surface between a set of input surface samples. In fact, the goal is to define the connections between close samples, or the ordering of the independent variable -i.e. the 3D position of the surface samples–.

The description of the resulting polygon mesh is such that it suffices at providing the required information for both visualization and analysis, without requiring additional information. As presented in Section 3.1.2, the polygon mesh will be represented as a list of vertices, each with its 3D position –and its RGB color, when available–, and a list of faces, each represented by three vertex identifiers.

The problem of finding the set of faces that correctly connect the input points is tackled in a region propagation fashion, as introduced above. First, a robust initial surface region –a triangle– will be defined and, then, the propagation on the current surface contour will start from the triangle's contour and iteratively end up in the covering of the whole surface.

8.3.1 Initial triangle

The choice of a starting face for propagation is of great importance, since it sets where the surface will start growing from. Thus, a set of conditions are required to be fulfilled in the form of a set of rules, in order to ensure a correct subsequent propagation.

A random unused starting point A and its k nearest-neighbors C_i are queried to the kd-tree, where C_1 and C_k are the closest and the farthest points in the returned set, respectively. The objective is to find a triangle between $A, B = C_j$ in the range j = 1..k and a third valid C_i in the range i = j..k. In order to do so, points C_i must verify the following conditions (also depicted in Fig. 8.3):



Figure 8.3: Rules for initial triangle creation. Left, situation where the *shape rule* (S) is not fulfilled. Right, situation where the *orientation rule* (O) is not fulfilled. Red color means *rejected*, green means *new* and grey means *already there*

- (S) Shape rule: Angle α_i at C_i must verify that $\alpha_{min} < \alpha_i < \alpha_{max}$.
- (O) Orientation rule: The normals of the three points, i.e. $\hat{\mathbf{n}}_A, \hat{\mathbf{n}}_B$ and $\hat{\mathbf{n}}_{C_i}$, must be coherently oriented.

If it is not possible to find a correct configuration from the randomly chosen point A and its neighbors, the method is restarted with a new randomly chosen point A. The first $C = C_i$ point which verifies the shape rule (S) and the orientation rule (O) with respect to two close points A and B is kept, obtaining an $F_1 = \widehat{ABC}$ starting triangle. The half-edge structure presented above is also updated with the three new half-edges just created, which are also added to a queue of contour edges that will be used throughout the propagation. In an abuse of notation, an *edge* in the queue actually corresponds to a *half-edge* in the former structure.

8.3.2 Propagation

Propagation starts at the initial triangle F_1 , its three edges being the first contour edges to be processed. Contours to be propagated are extracted from a FIFO queue of contour edges. Then, it is checked whether they are active contours -i.e. the half-edge structure does not contain the opposite-oriented half-edge-. The idea is to expand contours with faces, while preserving topological properties at the same time. We define the following features related to a contour edge:

- F: Face the edge belongs to, with normal $\hat{\mathbf{n}}_f$.
- V_1, V_2 : Edge's start (source) and end (sink) vertices, respectively.
- V_3 : Opposite vertex which completes the existing F.
- M_e : Edge's middle point, or $M_e = \frac{1}{2}(V_1 + V_2)$.
- $\hat{\mathbf{d}}_e$: Propagation direction from M_e , defined as $\hat{\mathbf{d}}_e = \hat{\mathbf{n}}_f \times (V_2 V_1)$.



Figure 8.4: Rules for triangle propagation, added to those in Fig. 8.3. Top left, situation where the *used vertex inside rule* (U) is not fulfilled. Top right, situation where the *unused vertex inside rule* (UN) produces the rejection of such vertex. Bottom left, situation where the *spatial position rule* (P) is not fulfilled. Bottom right, a valid face is created, which fulfills the complete set of propagation rules. Red color means *rejected*, green means *new* and gray means *already there*

The k-nearest neighbors of M_e are queried to the kd-tree. Such vertices, called candidates or C_i , are processed in increasing distance order, beginning with the closest one C_1 . In addition to rules for correct shape (S) and orientation (O), a candidate C_i (and the new candidate face $\tilde{F} = \widehat{V_1 V_2 C_i}$) must satisfy four more rules to be considered valid (see Fig. 8.4). The first two of these additional rules are the following:

- (P) Spatial position : The candidate vertex must be placed in a $\beta = 2\pi$ sr solid angle centered at the edge's middle point M_e and oriented as the propagation direction $\hat{\mathbf{d}}_e$.
- (C) Contour or unused : The candidate vertex must be either a non-used vertex or a contour one.

For the two last additional rules, we make use of the second type of search implemented for the kd-tree structure. Since it deals with a local neighborhood around each created face, a range search with a suitable radius is performed, in order to find all the existing points regardless of the sampling density. The idea behind these rules is to prevent the creation of two close layers of triangle faces coherently oriented. This might occur due to the existence of points distributed in two close layers when working with noisy surface samples.

• (U) Used vertex inside : Let bc_f be the baricenter of the new face \tilde{F} , and



Figure 8.5: From left to right and top to bottom, propagation of a contour until completely covering the set of samples for a surface composed as samples of a bi-dimensional Gaussian function

 \mathcal{B} a ball centered at bc_f of radius $\rho = \max\{dist(p, bc_f)\}$ with $p = \{V_1, V_2, C_i\}$. If there exists an already used vertex $v_u \in \mathcal{B}$, which can be perpendicularly projected onto \tilde{F} , the candidate point C_i is rejected.

• (UN) Unused vertex inside : In the (U) context, if there exists an unused vertex $v_{un} \in \mathcal{B}$ which can be perpendicularly projected onto \tilde{F} , then v_{un} is rejected.

When a candidate point C_i , and its associated candidate face \tilde{F} , fulfill all of the preceding rules, this new triangle face is created. The queue of contour edges is updated with the newly added edges of \tilde{F} and the half-edge matrix also reflects the introduced changes. In case that no candidate satisfies the imposed rules, the current edge is marked as *contour* edge and it is not further processed.

The propagation process is illustrated in Fig. 8.5, where the contour –colored in red– is iteratively propagated until covering all the samples of a surface corresponding to a bi-dimensional Gaussian surface.
Non-manifold edge detection

Despite of the set of rules for topological correctness described above, in some cases a wrong set of triangular faces might be created in some regions where different fronts of the surface propagation meet. This situation is illustrated in Fig. 8.6 (a). The gray region corresponds to a hole that is to be tessellated by propagating any of the contours around it. In Fig. 8.6 (c), two contour edges have propagated to new faces fulfilling all of the correctness rules presented above. However, this results in an incorrect configuration of overlapping surfaces.

From Fig. 8.6 (f) on, the mechanism for detection of a non-manifold edge and its posterior correction is illustrated. When a new candidate face \tilde{F} is to be created, it is checked whether each of its contour half-edges can be correctly created –meaning that it does not exist in the half-edge structure–.

If one of the half-edges already exists, it is detected as a non-manifold edge. Consequently, the candidate face \tilde{F} and the faces sharing the colliding half-edge –face identifiers are stored in the half-edge structure– are deleted as shown in Fig. 8.6 (g).

The half-edges that are set free by the removal of these faces adjacent to the non-manifold edge are added to the contour queue. Then, the normal propagation proceeds, which eventually produces a correct configuration by considering a reduced set of candidate points.

Iterative propagation

Propagation stops when no more contour edges are left to be processed in the queue. However, there may still be many unused points left. Iterating the initial triangle and propagation steps presents two advantages. First, it helps to achieve a complete meshing of the surface by filling zones which, following the rules for topological correctness, were not reachable from the previous starting triangles. Second, the sampled scene may contain many unconnected surfaces to be interpolated. Letting the system find new starting triangles will result in an individual meshing of each object.

Thus, when several disconnected surfaces compose the sampled scene, a new starting triangle is searched among the remaining unused points, which will launch a new propagation phase. Such cycle may be repeated iteratively while the amount of used points N_u is below a threshold. The propagation implementation used through this thesis has been stopped in all cases when $N_u > 90\%$.

In case that the propagation to a certain part of a surface has not taken place



Figure 8.6: Incorrect surface interpolation and posterior correction. In every stage, the blue line signals the next contour edge to be processed in a queue. (a) An uncovered region (light gray) that is to be covered by the triangular mesh; (b) A half-edge propagates the surface following the set of rules for topological correctness; (c) another half-edge propagates the surface incorrectly, despite of following the set of rules for topological correctness; (d) and (e) two other contour half-edges propagate; (f) another contour half-edge propagates, in this case producing an overlap with an existing half-edge; (g) the three facets adjacent to the non-manifold edge are deleted; (h) and (i) the propagation proceeds without topological errors



Figure 8.7: Topologically correct output mesh and detail of a challenging region around a saddle point

due to limitations imposed by the spatial position rule (P), the propagation would also proceed correctly, since every propagation phase takes into account the connectivities found during the previous propagation steps.

8.4 Conclusions

In this chapter, a novel meshing algorithm has been presented, which is useful to interpolate a continuous 2D surface in 3D space between samples corresponding to a discrete representation. It is driven by the main goal of defining an ordering for the input surface samples which is useful for interpolating the values of a continuous surface between these samples within a limited computation time.

The proposed algorithm, which has some points in common with the ball pivoting algorithm [Bernardini et al., 1999], being the main one its propagation-based principle, has accomplished its goals:

- It preserves the position of the existing surface samples, and is therefore capable of maximizing the accuracy of the reconstructed surface when error-free data points are available.
- It provides fast interpolations of surfaces with the use of an efficient set of data structures for spatial queries.
- It is capable to obtain a correct orientable surface, through the application of several rules for topological correctness and the possibility of automatically recovering from errors beyond the scope of these rules.

When compared to the ball pivoting algorithm, the proposed method does not require an iterative increment of the search radius in order to cope with scenes with irregular distribution of surface samples. Instead, the use of kd-trees as a tool for spatial queries automatically isolates the meshing strategy from this problem. As illustrated in Fig. 8.7, the proposed method can reconstruct topologically correct surfaces even in cases where the input set of surface samples are unevenly distributed and represent a challenging geometric region such as a saddle point. More details about the performance of the meshing algorithm are presented in the next chapter.

As a difference with respect to volumetric algorithms like those based on the final extraction of the surface by marching cubes -e.g. [Kazhdan et al., 2006], a robust, state-of-the-art approach–, the chosen propagation scheme keeps unmodified the position of existing surface samples. This feature allows, for example, to easily interpolate the values contained in the surface samples to any position over the resulting continuous surface, using classical algorithms such as [Gouraud, 1971]. This can be done quickly thanks to the one-to-one correspondence between discrete surface samples and vertices in the mesh, which translates in a fulfillment of the main goal of the proposed method.

Chapter 9

Surface Interpolation Results

he algorithm for linear interpolation of a surface presented in Chapter 8 has been designed in order to provide a topologically correct triangular mesh that connects an input set of oriented points. The natural idea is to use this method for oriented points extracted with any of the surface sampling techniques proposed in the first part of this dissertation. However, in order to characterize its performance, in this chapter we present an evaluation based on a comparison of several surface reconstructions with range data, which have been obtained from [The Stanford 3D Scanning Repository, 2010].

First, the different techniques used for comparison are listed, which also provide a triangular mesh from sets of oriented 3D points. Next, the datasets and the metrics used for comparing them to our proposed technique are presented and quantitative results are provided for the six used datasets, for which a reference surface represented as a triangular mesh is available and used as ground-truth. Finally, a qualitative evaluation of the quality of the surfaces obtained with all methods for each used dataset is also introduced.

9.1 Methods

In this section, we present the techniques for triangulation of a surface from a set of oriented points that we have used in the evaluation of the technique presented in Chapter 8. The goal is to compare the results obtained with several approaches falling in the two categories presented in the Chapter 2 (Section 2.3) and in more detail in Section 8.1.2 in the previous chapter: propagation-based and marching cubes-based methods. All the following algorithms, except for the proposed method and ReOP, are implemented in the *MeshLab* software [Cignoni et al., 2008].

9.1.1 Propagation methods

From this category, the *ball pivoting* algorithm, *restricted and oriented propagation* and the *proposed* method have been used.

BPA. The *ball pivoting* algorithm [Bernardini et al., 1999] is related to the socalled *alpha-shapes*, which can be seen as a generalization of a convex hull [Edelsbrunner and Mücke, 1994]. In this method, a virtual ball of a user-specified radius scans a cloud of oriented points and obtains triangular faces connecting them in a surface propagation approach.

In our experiments, the initial value for the ball radius is automatically adjusted, and three iterations are then applied with increasing radius in steps of 0.5% of the bounding box diagonal.

ReOP. Restricted and oriented propagation [Suau et al., 2010] is a fast method based on voxel propagation, with the main drawback of a non-manifold output mesh, which renders this approach useful when targeting visualization, but not when the reconstructed surface has to be further processed.

In our experiments, only the parameter adjusting the voxel size is adjusted, leaving the 6-OCT pattern constant throughout all the executions. The rule for choosing the voxel size is such that at least 96% of the points are not sharing a voxel with other points.

Proposed. The *proposed* method presented in Chapter 8 creates an initial triangular face and iteratively propagates its edges to generate new faces under a set of rules for topological correctness

The only parameter that has been adjusted is the number of results k for nearestneighbor search, setting it to a value such that no holes appear in the resulting surfaces (ranging from 5 to 20 in our experiments).

9.1.2 Marching cubes-based methods

From this category, *RIMLS-MC* and *Poisson reconstruction* have been used, two methods that have in common the use of marching cubes [Lorensen and Cline, 1987] for extracting the final surface.

RIMLS-MC. Robust Implicit Moving Least Squares-Marching Cubes [Oztireli et al., 2009] obtains a Moving Least Squares surface defined by the point set as

a robust implicit extension of Moving Least Squares that preserves sharp features, by using non-linear regression. The surface is finally extracted with marching cubes.

In our experiments, the scale of the spatial low-pass filter, relative to the radius –local point spacing– of the vertices is set to 4. The grid resolution, as in ReOP, is manually adjusted until the number of vertices at the output is practically equal to that at the input. It has also been visually checked that the presence of holes due to volumetric under-sampling remains under reasonable limits.

Poisson-Rec. Poisson reconstruction [Kazhdan et al., 2006] solves a Poisson equation for volumetric occupancy trying to best approximate the vector field defined by the samples. Then, it obtains a smooth surface after applying an octree-based implementation of marching cubes to the resulting occupancy function.

In our experiments, only one parameter has been tuned. This is the depth of the octree used for extracting the final surface. Again, the criterion has been that the number of output vertices was approximately equal to the number of surface samples at the input.

9.2 Datasets and Metrics

In this section, we first present the six datasets used for both the quantitative and the qualitative comparisons of the polygon meshes obtained with each of the techniques listed in the previous section. Then, the metrics used for comparing the performance of each method are also presented.

9.2.1 Datasets

The experiments presented below are carried out with six datasets containing surfaces represented as oriented points, most of them provided by the Stanford 3D Scanning Repository [The Stanford 3D Scanning Repository, 2010]. These are listed in Table 9.1.

For all these datasets, a reference meshing is provided along with the oriented points, which is used as ground-truth for the quantitative comparison below. These datasets include both irregular and uniform distributions of surface points. Thus, they provide a sufficient measure of the performance of the proposed method under different scenarios.

Name	# Points	Points distribution
drill	1,960	Uniform
bunny	34,834	Uniform
arm a dillo	$172,\!974$	Uniform
hand	327,323	Irregular
dragon	437,645	Very irregular
happy	$543,\!652$	Very irregular

 Table 9.1: Datasets used for quantitative evaluation of our proposed surface interpolation algorithm

9.2.2 Metrics

The RMS Hausdorff distance metric is used to evaluate the similarity between the obtained surface and a ground-truth surface, giving an idea of the accuracy of the reconstruction. More details about this metric can be found in Appendix C.

A second parameter taken into account is the overall computation time, which does not include the time required to read the oriented points from the input file nor the time required to write the polygon mesh to the output file.

The third and last figure that has been taken for comparison is the memory footprint of each method. The memory measurements are only valid as an approximation to the actual values, since they have been obtained as a multiple of one thousandth part of the total amount of available memory. However, they reflect the trends with respect to memory efficiency of each compared method.

Thus, the proposed surface interpolation algorithm is quantitatively compared to the methods introduced above using these three figures. For reference, all the experiments are executed on a 64-bit Intel Xeon CPU 3.0 GHz with 32 GB RAM.

For the qualitative comparison, visual inspection of both the overall mesh and some details in challenging areas are considered in order to evaluate the performance of each method.

9.3 Results

In this section, we present the quantitative and qualitative results obtained with the proposed technique and each of the remaining five methods used for comparison. Details of the datasets presented above, which have also been used for the qualitative comparison, are introduced in the qualitative comparison.

Method		drill	bunny	arm a dillo	hand	dragon	happy
RIMLS-MC	Time	1.66	36.06	170.42	211.92	180.32	215.24
	Mem.	-	-	96	192	256	288
	RMS	1747	526	249	234	279	344
	Time	0.66	13.00	35.64	38.81	43.61	49.55
Poisson-Rec	Mem.	-	-	224	480	640	704
	RMS	2532	1367	591	543	874	1232
BPA	Time	0.15	48.62	347.00	935.36	1846.49	2988.64
	Mem.	-	-	64	128	${\bf 192}$	224
	RMS	506	99	61	92	105	144
ReOP	Time	0.70	0.81	4.10	8.73	14.36	14.89
	Mem.	-	-	288	1088	1696	1920
	RMS	464	94	61	163	56	116
Proposed	Time	0.03	0.92	3.06	8.12	10.11	15.44
	Mem.	-	-	96	192	288	320
	RMS	410	90	40	43	103	113

Table 9.2: Computation time, memory usage and accuracy for meshing each of the six sets of oriented points. The execution time is shown in seconds and the memory usage in megabytes, whereas the RMS Hasudorff distance is shown in units of a millionth part of the main diagonal of the bounding box of each dataset. In bold text, the shortest execution time, the least memory usage and the best accuracy for each dataset

9.3.1 Quantitative comparison

In Table 9.2, the execution time, approximate memory usage and RMS Hausdorff distance –in all cases, the smaller, the better– are shown for each of the datasets and the techniques used for comparison. About the units, execution time is measured in seconds and memory usage in megabytes, whereas the RMS Hausdorff distance is measured in units of a millionth part (10^{-6}) of the diagonal of the corresponding bounding box. The shortest computation time, smallest memory footprint and best accuracy are written in bold font for each dataset. Due to the small resolution of the memory usage measurements, the memory footprint of the datasets *drill* and *bunny* could not be obtained.

As shown in Table 9.2, the proposed method is always showing either the best computation time or the best accuracy among the tested methods. In terms of computation time, the properly created kd-tree proves as a suitable data structure for spatial queries, when compared to other propagation-based schemes. In terms of quality, when compared to marching cubes-based methods, the utilization of the exact position of surface samples makes propagation-based techniques much more accurate, whereas among these latter techniques, the proposed one provides in most cases the surface closest to the reference one due to the topologically correct

configuration of the resulting triangular faces.

With respect to the memory footprint, the comparison might be biased: the methods executed from *Meshlab* –BPA, RIMLS-MC and Poisson-Rec– have been measured by subtracting the initial memory usage of the program, before loading the datasets and starting the meshing algorithm. It might be possible that some memory is released by the program during the execution of potentially memory-demanding methods, or that some data structures actually used by the meshing methods are pre-loaded when the program starts. Thus, we cannot take solid conclusions about the memory usage beyond that, in general, methods requiring a global voxelization of the working volume –Poisson-Rec and ReOP– are more memory demanding. In general, the propose method offers a good computation time vs. memory footprint trade-off.

To sum up, the proposed method is the best in terms of accuracy and execution time, in general, being closely followed by ReOP in most cases and superseded by this method in the few remaining ones. In terms of memory usage, it was not possible to offer an accurate comparison with all of the methods, but it compares favorably to ReOP and Poisson-Rec, the closest contenders in terms of computation time.

9.3.2 Qualitative comparison

For the following experiments, a short introduction to each dataset is presented. Then, some qualitative valuations are drawn from the results depicted in the corresponding figures, presenting both general views and close-up details for the surfaces reconstructed using each of the presented methods as well as for the ground-truth surface.

Drill dataset. The *drill* dataset is the smallest one in the comparison. However, the relatively low sampling density of the surface makes this dataset challenging. In Fig. 9.1, it is shown how the marching cubes methods –RIMLS-MC and Poisson-Rec– fail to reconstruct the original, due to over-smoothing introduced by the regularization mechanism. BPA clearly shows topological errors, whereas ReOP creates an excessive number of faces, which corresponds to its condition of not guaranteeing topological correctness. The proposed method provides an almost-exact replica of the original surface.

Bunny dataset. This is a classical, common dataset in the literature that presents a mostly regular sampling with some holes on its base and irregular sampling on the seams between different scans. In Fig. 9.2, the detail view corresponding to the tail



Figure 9.1: First row: general view of the surface of the Drill dataset, from ground-truth – first column– and each of the compared techniques. All of them output the same number of vertices in the reconstructed surface. Second row: detailed view of the surface, with visible edges, corresponding to the top of the drill

shows how the proposed method and BPA provide a meshing very close to that of the ground-truth. ReOP, once again, generates an excessive amount of triangular faces, whereas RIMLS-MC presents some holes and Poisson-Rec oversmoothes the resulting surface.

Armadillo dataset. The armadillo dataset introduces a greater level of detail than the previous two datasets. It is also a common dataset in the literature that presents a very regular sampling with a somewhat more complicated topology on the surface corresponding to the teeth of the figurine. Fig. 9.3 shows the reconstructed surfaces obtained with each of the compared methods. The detail views of the proposed method, BPA and RIMLS-MC reveal different polygon meshes, which closely resemble the one corresponding to the ground-truth. ReOP presents an excessive number of triangular faces, whereas Poisson-Rec over-smooths the reconstructed surface.

Hand dataset. The *hand* dataset shows a challenging scenario with lots of joints between different pieces of the complete surface and hollow structures enclosed by two layers of opposite-oriented surfaces. However, the surface lacks fine detail, which makes it a good dataset for marching cubes-based techniques. In Fig. 9.4, general views already show some errors in BPA and ReOP, due to the small distance between the mentioned two-layers of opposite-oriented surfaces. The proposed method and RIMLS-MC provide the most accurate surfaces in topological terms, whereas Poisson-Rec tends to over-connect different pieces due to the excessive volumetric smoothing.

Dragon dataset. The *dragon* dataset shows a challenging scenario with a very irregular sampling and a great level of detail on its shape. In Fig. 9.5, the general views reveal that all methods provide apparently hole-free surfaces. However, a closer inspection of the back of the head shows that the proposed method and BPA are the ones that most closely resemble the ground-truth mesh. ReOP generates an excessive amount of faces and RIMLS-MC presents some small-scale holes, whereas Poisson-Rec presents an over-smoothed, yet watertight, output mesh.

Happy dataset. The *happy* dataset shows a very challenging scenario with the most irregular sampling and a great level of detail on its shape. In Fig. 9.6, the general views of the reconstructed surface already reveal a loss of detail with RIMLS-MC and Poisson-Rec, due to excessive regularization. BPA presents some cracks in the base, whereas ReOP produces an excessive amount of faces due to the lack of topological guarantees provided by the method. The surface provided by the proposed technique closely resembles that corresponding to the ground-truth.



Figure 9.2: First column: general view of the surface of the *Bunny* dataset, from ground-truth –first row– and each of the compared techniques. All of them output the same number of vertices in the reconstructed surface. Second column: detailed view of the surface, with visible edges, corresponding to the tail of the bunny



Figure 9.3: First column: general view of the surface of the *Armadillo* dataset, from ground-truth –first row– and each of the compared techniques. All of them output the same number of vertices in the reconstructed surface. Second column: detailed view of the surface, with visible edges, corresponding to a close-up of the face of the figurine



Figure 9.4: First column: general view of the surface of the *Hand* dataset, from ground-truth –first row– and each of the compared techniques. All of them output the same number of vertices in the reconstructed surface. Second column: detailed view of the surface, with visible edges, corresponding to the joins between three pieces close to the fifth finger



Figure 9.5: First column: general view of the surface of the *Dragon* dataset, from ground-truth –first row– and each of the compared techniques. All of them output the same number of vertices in the reconstructed surface. Second column: detailed view of the surface, with visible edges, corresponding to the back of the head



Figure 9.6: General and detailed –with visible edges– views of the surface of the *Happy* dataset, for ground-truth and each of the compared techniques. All of them output the same number of vertices in the reconstructed surface. The detailed views show a close-up of the face, where marching cubes-based methods lack fine detail and ReOP clearly introduces unnecessary redundant faces, as well as topological errors

In general, it can be observed that the proposed method adapts well to all of the datasets and presents a reconstructed mesh which, topologically, is very close to the ground-truth one, apparently delivering better results than the rest of proposed methods. Considering the results from the quantitative comparison above and these qualitative results about the topological correctness of the surface, the proposed technique compares favorably to all of the other tested methods.

9.4 Conclusions

In this chapter, both quantitative –accuracy, computation time and memory footprint– and qualitative results –visual similarity of the reconstructed mesh to that of the ground-truth– have been presented. For obtaining them, the surface interpolation technique proposed in the previous chapter as well as other state-of-the-art approaches, both propagation-based and marching cubes-based have been compared using the same datasets. In general, the main advantages of the proposed technique are the following:

- The proposed method is fast when compared to both propagation-based and marching cubes-based techniques. This is due to the efficient query for surface samples using the *kd*-tree structure.
- The proposed method, being part of the propagation-based techniques, offers a better accuracy than marching cubes-based counterparts, since the position of existing surface samples directly translates into the position of the vertices of the resulting mesh.
- When compared to the other propagation-based techniques, the topological correctness of the reconstructed surface obtained with the proposed method makes it a better alternative than the fast, accurate *Restricted and Oriented Propagation* method, the hardest contester in terms of computation time.
- By visual inspection, the mesh resulting from the proposed method resembles the most that of the ground-truth.

About its drawbacks, the most relevant one is the fact that, as a difference with marching cubes-based methods, in which a volumetric, global regularization takes place, when the input data is noisy, so is the output surface. However, this problem might be tackled by a pre-processing stage of the input surface samples, such as the surface smoothing that has been presented in Section 6.5.2, in the chapter devoted to statistical surface sampling.

Part III

Overall Results and Conclusions

Chapter 10

Overall Method: Sampling and Interpolation Results

n the two main parts in which the dissertation is divided, several (progressively improved) approaches for surface sampling and one for surface interpolation have been presented. In this chapter, the best surface sampling approach (the one presented in Chapter 6) and the surface interpolation technique presented in Chapter 8 are concatenated in order to generate 3D reconstructions out of several multi-view video sequences.

The chapter is divided in two parts. The first one presents an evaluation of the proposed approach by comparing it to a baseline system for scene reconstruction composed of classical techniques. The second part presents results obtained from multi-view video sequences publicly available from [4D Repository, 2010], with datasets considering different numbers of input views. These experiments have been used in order to demonstrate the achievement of the main goals of this thesis. As a remainder, the surface reconstruction was expected to be fast –towards realtime–, effective –usable for both analysis and visualization– and efficient –able to compress the multi-view data redundancies–, and the experiments corroborate the accomplishment of these requirements.

10.1 Evaluation

In this first section we present experimental evidence showing that the system resulting from the concatenation of surface sampling and interpolation obtains fast, effective and accurate descriptions of surfaces when compared to a baseline system built from existing techniques.

Dataset	# Views	# Frames	Color	Image size
Real	18	584	yes	640x480
Synthetic	8	977	no	640x480

Table 10.1: Data-sets used for evaluation of the proposed approach

10.1.1 Data and methodology

Two types of multi-view sequences are used for this first set of experiments. The first one consists on the projected silhouettes of an animated 3D model of a moving person onto a set of 8 cameras. Using such a synthetic sequence is useful for our experiments, since the 3D model used to generate it is available and can be used as 3D ground-truth, which in general is not available in real sequences. The second sequence consists in real data of a moving person captured by 18 cameras in the facilities of *Telefonica* $R \mathcal{C} D$ and the corresponding automatically obtained foreground silhouettes, which present some segmentation errors. The configuration details of these two data-sets are summarized in Table 10.1.

Baseline system. Both the baseline system and the proposed concatenation of surface sampling and interpolation (Table 10.2) use the same silhouette extraction method [Gallego and Pardàs, 2010] for generating their input silhouettes for the real sequences –for the synthetic sequence, silhouettes are noise-free and generated through OpenGL–. Then, in the baseline system, a voxelized *shape-from-silhouette* implementation that uses projection *look-up tables* provides a volumetric reconstruction. This implementation technique reduces computing time at the expense of an increased memory usage. At the last stage, a polygon mesh is extracted from the volumetric reconstruction by applying the *marching cubes* (MC) algorithm. The MC implementation also exploits precomputed tables to improve its throughput. The use of precomputed tables in each stage in order to improve its speed draws a challenging scenario for the comparison.

10.1.2 Experiments

For this evaluation, three experiments have been run on a system equipped with an Intel Xeon 3.0 GHz processor and 32 GB of RAM. These three experiments, which

System	Reconstruction	Mesh extraction
Baseline	Voxelized SfS	Marching cubes
Proposed	Statistical Sampling	Proposed meshing

Table 10.2: Features of the baseline system vs. those of the proposed approach

	Bas	seline	Pro	posed
# Points	Time (s)	$\rm RMS{\times}10^{6}$	Time (s)	$\rm RMS{\times}10^{6}$
5000	0.5	1006	0.4	676
10000	1.2	836	0.9	623
20000	3.0	752	1.9	575
40000	12.9	673	4.1	534

 Table 10.3:
 Quantitative comparison between the baseline system and the proposed sampling-interpolation concatenation

provide both quantitative and qualitative assessments about the performance of the methodology presented in this thesis, are the following:

- The first one is a performance comparison between both systems, using multiview sequences with 3D ground-truth data.
- Then, in the second experiment, surface reconstruction results obtained with both systems are qualitatively compared by visual inspection.
- Finally, in the third experiment, a possible *Free-Viewpoint Video* [Miller et al., 2005] application with the proposed technique is presented, in which a 3D reconstructed subject is placed in a virtual environment with a virtual camera that can be manually placed.

Quantitative performance comparison

The first experiment compares the performance of both systems, baseline and sampling-interpolation concatenation. This comparison consists in the reconstruction of the surface of the animated 3D model and its posterior comparison to the original mesh using the RMS Hausdorff distance [Aspert et al., 2002] between the reconstructed mesh and the real one, averaged along the whole sequence, as a quality measurement. As shown in Table 10.1, no color information is available in the synthetic sequence.

Table 10.3 summarizes the average times employed in the concatenation of 3D reconstruction –statistical surface sampling or volumetric shape-from-silhouette– and meshing –proposed surface interpolation method or marching cubes– and the average RMS Hausdorff distance for both the baseline system and the proposed sampling-interpolation scheme with growing sampling densities. In the table, lower values of RMS mean higher quality. The proposed system offers more accurate results at every resolution level, due to the adopted sampling scheme, better adapted to surfaces. At higher resolution levels, much smaller processing times and better accuracy reflect the suitability of the approach for dense surface reconstruction.



Figure 10.1: Mesh from the baseline system (left) and from the sampling-interpolation concatenation (right) with approx. processing times of about 5 s. Color is added for better discrimination

Qualitative comparison

The second experiment consists in a qualitative comparison of both systems in a real scenario (using the *real* sequence). In Fig. 10.1, two meshes corresponding to the surface of a person in a controlled environment, reconstructed by the baseline system and by the sampling-interpolation concatenation, are shown. Both for this experiment and the next one, a conservative estimate of the visual hull with $\tau = 2$, the tolerance to foreground segmentation errors, has been obtained with both methods, due to the notorious presence of foreground misses.

In this case, an equal processing time for both methods (\sim 5s) is the comparison criterion, which results in a larger number of reconstructed points for the proposed technique (30,000 versus 12,000) that, in addition, are more accurate in terms of their location with respect to the ground-truth. Visual inspection reveals that the mesh obtained by the sampling-interpolation concatenation is more accurate and realistic than the one obtained by the baseline system. Furthermore, the meshing algorithm proposed in Part II of this thesis delivers a manifold surface even in challenging regions (Fig. 10.2).

Free-viewpoint video application

In the third experiment, the complete surface reconstruction system –including surface coloring, not used in the previous experiments– has been applied to the *real* sequence. With this, we can demonstrate a possible free-viewpoint video application, consisting in a replacement of the real background by a virtual one and a free selection of the viewpoint. Please note that the mesh is not textured. Instead, we apply Gouraud shading [Gouraud, 1971] of the triangle faces of the mesh after set-



Figure 10.2: Manifold output mesh and detail of a challenging region around a saddle point

ting vertex colors equal to those of the corresponding reconstructed surface points. In Fig. 10.3, two instants of the sequence are viewed from three novel viewpoints.

This possible application reflects the fact that the reconstructed surface contains a suitable amount of information about the foreground elements of the scene. Such information can be employed, for example, for visualizing –from a freely chosen viewpoint– the foreground elements of the scene without the need of using the original images. In other words, the approach proves to be *effective* for its usage in one of the target applications.

10.2 Validation

In the previous section, it has been shown how the proposed methodology performs faster and with better accuracy than an equivalent system for surface reconstruction built from classic algorithms. In this section, we aim at illustrating how the proposed methodology can obtain high-quality reconstructions of scenes of diverse types. As a result of these experiments, both the *effectiveness* and the *efficiency* of the proposed surface reconstruction method for multi-view sequences are showcased.

10.2.1 Data and methodology

For this second set of experiments, fourteen sequences from three datasets *-dancer*, *children* and *martial-* available in [4D Repository, 2010] have been used, which are listed in Table 10.4. As it is shown, the two latter are captured by 16 cameras placed all around the scene, whereas the *dancer* sequence only contains 8 views. The complete set of multi-view sequences are used for the first experiment, whereas



Figure 10.3: Three novel viewpoints (rows) from two different time instants (columns) depicting how a *Free-Viewpoint Video* application, in which the original background has been replaced by a virtual one, would look like using the methodology presented in this thesis. The scene is captured by 18 cameras and the automatic foreground extraction introduces a notorious presence of segmentation errors. The original multi-view video sequences are courtesy of *Telefonica R&D*

Dataset	Resolution	# cams	Sequence	$\#\ {\rm frames}$
Dancer	780×582	8	dancer	201
Children	1624×1224	16	antoine	529
			antoine assis ballon	516
			antoine edmond ballon	339
			antoine lucie	485
			antoine roue	507
			lucie corde	410
Martial	1624×1224	16	benjamin baton	233
			coup de pied 1	210
			$coup \ de \ pied \ 2$	240
			parade coup de point 1	241
			parade coup de point 2	256
			saisie col	310
			saisie poigne	321

Table 10.4: Datasets from the 4D repository

only one sequence of each dataset has been used for the second experiment.

The video sequences are provided of foreground silhouettes of great quality, with minor errors in few frames. This means that for these experiments it will not be necessary to obtain a conservative estimate of the visual hull, which translates in a more accurate surface reconstruction. Another nice feature of this data is that color is calibrated across views, which is useful for both the application of photoconsistency or surface coloring.

As in the previous section, the sampling-interpolation concatenation presented throughout this thesis has been used for reconstructing the scene at each time instant. The surface sampling stage consists in the statistical sampling (Chapter 6), whereas the surface interpolation stage consists in the application of the method presented in Chapter 8. The number of surface samples in the first stage, unless otherwise stated, has been set to 100000 and the different distances parameterizing the method are set to the same values as those in Table 7.6. For the meshing stage –surface interpolation–, a number of 20 nearest-neighbors is chosen for surface propagation, which results in a closed surface without an excessive computational cost ($\sim 2 \text{ s/frame}$).

10.2.2 Experiments

The fourteen videos have been processed by the sampling-interpolation concatenation, and the resulting surfaces have been re-projected onto the original viewpoints. The experiments in this section show how the proposed methodology is *effective* –suitable for both analysis and visualization–:

- It is able to retrieve with a suitable level of precision the original images by composing the projection of the reconstructed surfaces with the original background or, in other words, it contains most of the information contained in the foreground of the original images.
- The reconstructed surface can also be used for visualization, since it is closed and contains the texture information –as already seen in the first part of this chapter–.

Using the same data, it is also shown that the presented methodology is efficient, since under acceptable level of losses, it is able to compress the multi-view data for video sequences.

Effectiveness

The effectiveness of the methodology presented in this thesis is showcased by presenting two types of qualitative results: the resemblance between the re-projection of the reconstructed surfaces and the original views and the comparison of the reconstructed surface with that obtained with a state-of-the-art method, the *Exact Polyhedral Visual Hull* (EPVH) [Franco and Boyer, 2003]. For this experiment, the complete set of multi-view sequences of each dataset has been used.

Dancer dataset. In Fig. 10.4, some sampled time instants of the results obtained using the *dancer* sequence are shown. In this figure, some of the original input images are shown next to novel images generated by composing the re-projection of the reconstructed surface with the original image in order to add the missing background information in 3D –separate background was not available–. The surfaces reconstructed with the proposed method –with and without surface colors– and with EPVH are also shown at the bottom of the figure, seen from a novel virtual viewpoint.

As it can be observed at the top of the figure, the resulting surface, once projected onto the original views, presents some artifacts, due to the conservative method to determine surface sample visibility in the coloring stage and, most importantly, the type of reconstruction –silhouette-consistent–. However, the reconstructed surface still accurately describes the surface of the actual dancer up to a point where it closely resembles the original views after re-projection.

As it is shown at the bottom, the sampling density is high enough to represent with an acceptable level of precision the texture of the face of the dancer and some



Figure 10.4: Dancer (8 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). Some artifacts are visible, (*e.g.* under the arm in the top right view. Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color. The proposed method is better oriented to represent per-vertex features thanks to the dense surface sampling

details in her dress. This contrasts with the reconstruction provided by the EPVH, in which most of the triangular faces are of great dimensions, because it is not oriented to represent per-vertex color information.

Children dataset. At the top of Figures 10.5, 10.6, 10.7, 10.8, 10.9 and 10.10, some sampled time instants of the sequences of the *Children* dataset are shown. As with the previous dataset, some of the original input images are shown next to novel images generated by composing the projection of the reconstructed surface onto the original image in order to add the background information. At the bottom of these figures, both the surfaces reconstructed with the proposed method and EPVH are also shown.

As a comment about the reconstructed surfaces, the specular reflections at the top of the red ball in Figs. 10.5 or 10.8 are lost due to the color average that takes place in the coloring stage (Section 6.5.3 from the chapter devoted to *statistical surface sampling*). One limitation of the proposed method is shown in Fig. 10.10. Due to the specific geometry of the string, which does not present a clear orientation for its surface –it constitutes an almost one-dimensional body immersed in 3D space–, it is unlikely that its surface can be reconstructed by the statistical surface sampling method.

Martial dataset. As with the two previous datasets, at the top of Figs. 10.11, 10.12, 10.13, 10.14, 10.15, 10.16 and 10.17, some results for sampled time instants of the sequences in the *Martial* dataset are shown. At the bottom of each figure, the surfaces corresponding to some of these sampled time instants are seen for novel virtual viewpoints.

Again, a problem related to the geometry of the surface arises in Fig. 10.11. Although this time most of the surface of the stick held by the performer is reconstructed, it clearly shows that the sampling strategy does not adapt properly to these structures. In the other sequences, the behavior of the method is such that the re-projected surfaces closely resemble the original images, although they lack specular reflections and present some coloring artifacts due to the high resolution of the input images and the relatively sparse surface sampling.

Overall, the experiments presented in this section are useful as a validation of the *effectiveness* of the proposed approach. Indeed, at its output, the proposed methodology is able to accurately represent the surfaces of the foreground objects in multi-view scenes, including color information. This representation is useful for visualization, as demonstrated by the projection of the reconstructed surfaces onto novel viewpoints shown at the bottom of each figure.

Furthermore, since the information about these surfaces is presented as a single



Figure 10.5: Antoine, from the Children dataset (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). There are some visible artifacts (*e.g.* the missing specular reflections on the top of the ball in the middle right image). Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.6: Antoine assis ballon (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). The most visible artifact is in the transition between the ball and the left leg in the bottom right image. Bottom: composition of reconstructed surfaces in the three time instants for (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.7: Antoine Edmond ballon (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). Some artifacts are visible, *e.g.* the missing color at the bottom of the ball in the bottom right image, or the missing specular reflection on the top of the ball. Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.8: Antoine Lucie (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). The most noticeable artifacts are the missing specular reflections at the top of the ball in the right images and the loss of geometric detail around the chest of Antoine. Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.9: Antoine roue (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). The most visible artifact is the loss of detail on the orange hood in the bottom right image. Bottom: composition of reconstructed surfaces in the three time instants (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.10: *Lucie corde* (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). A clear artifact is the loss of the string, which is masked in the top views by the composition with the original images. Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color. The sampling strategy does not adapt well to the geometry of the string


Figure 10.11: Benjamin baton, from the Martial dataset (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). Some artifacts are visible: *e.g.* the missing colors on the chin in the bottom right image. Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color. As in Fig. 10.10, the sampling strategy does not adapt well to the geometry of the stick



Figure 10.12: Coup de pied 1 (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). The most visible artifact is the wrong coloring of the chest of the standing man in the bottom right image. Bottom: composition of reconstructed surfaces in two time instants (a) EPVH; (b) proposed and (c) proposed with color. As seen below, 100000 samples are in the sparse side with this dataset



Figure 10.13: Coup de pied 2 (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). A visible artifact is the lack of detail around the belt of the man at the left side of the bottom right image. Bottom: composition of reconstructed surfaces in two time instants (a) EPVH; (b) proposed and (c) proposed with color. Again, 100000 samples are sparse for representing all the detail



Figure 10.14: Parade coup de point 1 (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). Some artifacts are visible, *e.g.* the missing detail on the hair of the man getting up in the bottom right image. Bottom: composition of reconstructed surfaces in two time instants (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.15: Parade coup de point 2 (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). The most visible artifact is the lack of detail around the belt of the man facing the camera in the top right image. Bottom: reconstructed surfaces (a) EPVH; (b) proposed and (c) proposed with color. The face of the man looking down has not been completely colored due to its orientation



Figure 10.16: Saisie col (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). A visible artifact is the lack of resolution on the face of the turning man in the bottom right image. Bottom: composition of reconstructed surfaces in two time instants (a) EPVH; (b) proposed and (c) proposed with color. The reconstruction from EPVH cuts the head of the man in one of the time instants



Figure 10.17: Saisie poigne (16 views). Top: comparison of three original images (left) and the overlay of the reconstructed surface over the original background (right). Some artifacts are visible: *e.g.* the missing detail on the hair of the man facing the camera in the bottom right image. Bottom: composition of reconstructed surfaces in two time instants (a) EPVH; (b) proposed and (c) proposed with color



Figure 10.18: Average PSNR of the foreground pixels of the re-projection of the *dancer* sequence on the view with largest silhouette for different amounts of surface samples. From 20000 samples on, the method performs at its best level of quality

instance, containing both geometry –topology and orientation– and color information, it can also be used for further analysis. The correctness of the surface topology obtained with the meshing approach presented in Part II of this thesis has proven useful in order to correctly interpolate the surface color between samples.

Efficiency

The efficiency of the methodology presented in this thesis is also based on the suitable representation of highly redundant data from multi-view settings as a unique instance, containing all the relevant information about the foreground objects. This is possible when the background is known and static, as occurs in the sequences presented for the experiments in this section.

The main idea behind this experiment is to compare the transmission or storage resources that are required to represent the relevant information of multi-view sequences when either an image-based representation of the scene or a 3D model of the foreground are chosen for representing the multi-view data.

The first question that arises when trying to compare the resources required by either approach is how to choose a reasonable surface sampling resolution given a set of multi-view video sequences. A possible answer to such question, which is the one we apply here, is to choose a surface sampling density such that an increase of such density does not produce a significant gain in accuracy.



Figure 10.19: Average PSNR of the foreground pixels of the re-projection of the *antoine edmond* sequence on the view with largest silhouette for different amounts of surface samples. From 50000 samples on, the method performs at its best level of quality

Accuracy. In order to quantify the accuracy, we use the PSNR between the reprojected surface onto the original viewpoints and the foreground pixels of the original images, which are the only ones used for reconstruction. In the sequences used for our experiments, we have found that the most limiting view in terms of PSNR is that in which the foreground objects occupy most space in the image. This seems reasonable, since more samples are expected to be required in order to accurately represent the contents of a larger number of pixels.

In Fig. 10.18, the average PSNR in the most limiting view of the *dancer* sequence, with increasing number of surface samples, is shown. Equivalently, in Fig. 10.19 and Fig. 10.20, the average PSNR of the *antoine edmond ballon* sequence from the *children* dataset and that of the *saisie col* sequence from the *martial* dataset are shown, respectively.

In all cases we observe that the surface reconstruction system reaches a limit of PSNR for a certain sampling density (~ 50 dB for the sequences from the *dancer* and *martial* datasets and ~ 70 dB for the sequence from the *children* dataset), beyond which little variation in the average PSNR occurs. In some cases it can even decrease. We attribute this to the reconstruction of regions corresponding to shadows for certain sampling densities.

In Fig. 10.21, the re-projection on the most limiting view of the surfaces obtained from a frame in each sequence are shown for increasing number of surface samples. As suggested by the charts, given a sufficient surface sampling, further increasing the sampling density does not produce a significant gain in re-projection accuracy.

For the sequence from the *dancer* dataset, a number of 20000 surface samples



Figure 10.20: Average PSNR of the foreground pixels of the re-projection of the *saisie col* sequence on the view with largest silhouette for different amounts of surface samples. From 200000 samples on, the method performs almost at its best level of quality

(between the second and third row in Fig. 10.21) already reaches a reasonable level of accuracy for the most limiting image in the multi-view set. For the sequence from the *children* dataset, a number of 50000 surface samples (third row in the same figure) is already approaching the limiting accuracy, whereas for the sequence from the *martial* dataset, 200000 surface samples (fourth row) retrieves the right level of detail.

Sampling density. In order to put these figures in a common framework, we normalize both the average PSNR by its maximum value and the number of surface samples by the maximal number of pixels in the most limiting view of each of the processed sequences and show the resulting evolutions in Fig. 10.22.

As shown in the chart, a reasonable choice for the number of surface samples in the sequences used for these experiments would be approximately equal to the number of foreground pixels in the most limiting view. However, in order to be on the safe side, a more conservative, yet still economical, choice would be to take a number of surface samples of approximately twice as many as foreground pixels in the most limiting view –that with the largest silhouette–.

This choice seems also reasonable when we think about an approximately symmetric foreground object –like a cylinder– with a very small scale texture covering its surface. In such a scenario, approximately half of the surface samples would project onto the most limiting view, since half the complete surface would be hidden by the visible half. In that case, in order to retrieve the small scale texture by re-projection of the samples, a sampling density of at least one visible surface sample per pixel would be required, assuming a uniform distribution of the projection



Figure 10.21: From top to bottom, surfaces reconstructed using 2000, 10000, 50000, 200000 and 1000000 output surface samples. 10% of these samples are obtained as seed samples and the rest through the propagation mechanism in statistical surface sampling



Figure 10.22: Relative PSNR with respect to the maximum in sequences from the *dancer*, *children* and *martial* datasets as a function of the ratio between the number of surface samples and the number of foreground pixels in the view where they take most image pixels. As a consequence, a number of surface samples of around twice the number of pixels occupied by the foreground objects in the most limiting view seems to be a safe, yet economical, choice

of the samples. Since most of the surfaces in natural scenes are expected to have less detail than the assumed in this example, a total number of surface samples of around twice as many as foreground pixels in the most limiting view is a plausible choice. With this result, we can assume to have an efficient, accurate sampling, with controlled losses, when the rule of thumb presented above is considered.

Resources. From the charts in Figs. 10.18, 10.19 and 10.20, a number of surface samples larger than 20000, 50000 and 200000 suffice in order to accurately represent the information contained in the views of the sequences from the *dancer*, *children* and *martial* datasets, respectively. In order to stay in the safe side, 50000 surface samples are used for *dancer*, 100000 for *children* and 500000 for martial.

Using these values, which provide the best accuracy in terms of re-projection while using a slightly over-dimensioned amount of resources, a fair comparison can be made about the resources needed to store or transmit the multi-view sequences with a classical image-based manner and with the representation proposed in this thesis, which still introduces losses. This comparison highlights the efficiency of the proposed method, but it does not pretend to introduce the methodology presented in this thesis as a coding mechanism -e.g. [Mueller et al., 2009]–.

In case of representing the multi-view streams of the sequence from the *dancer* dataset as sets of lossless PNG images, a total of 808 MB need to be stored (or trans-

Representation		dancer	children	martial
	Compression	8 views	16 views	16 views
		780×582	1624×1224	1624×1224
Image-based	PNG frames	808 MB	12 GB	$27.2~\mathrm{GB}$
	H.264 lossless	$352 \mathrm{MB}$	$4.6~\mathrm{GB}$	$4.3~\mathrm{GB}$
Surface-based	FG Uncompressed	$756 \mathrm{MB}$	$2.4~\mathrm{GB}$	$12.1~\mathrm{GB}$
	FG OpenCTM	$76 \mathrm{MB}$	$242~\mathrm{MB}$	$1.2~\mathrm{GB}$
	(PSNR)	$50.72 \mathrm{~dB}$	$70.85~\mathrm{dB}$	$48.40~\mathrm{dB}$

Table 10.5: Comparison of the data required to represent multi-view sequences in a classical image-based manner –either as individual frames or exploiting temporal redundancies as lossless video sequences– and using the approach presented in this thesis, keeping the background as a set of PNG images and exploiting spatial redundancies in order to represent the foreground

mitted), without considering the additional capacity required for the foreground silhouettes. In contrast, a surface-based representation where a set of images containing the static background is stored (transmitted) only once as a set of PNG images, and the complete geometry and color of the foreground elements is represented in a compact form for each frame –without compression–, requires 756 MB.

The gain in this case is not dramatic, due to the relatively sparse set of views (only 8) and the low resolution of the input images. However, lossless compressing the geometry using the open source mesh file format *OpenCTM* [Geelnard, 2010], the latter figure shrinks to just 76 MB. This does not only introduce a clear gain with respect to the former figure, but also when compared to the introduction of lossless video coding -H.264 lossless-, where the initial 808 MB are reduced to just 352 MB by the exploitation of temporal correlation.

In the other two sequences, where the number of views (16) and the image resolution are both larger, further shrinkage of storage or transmission resources are achieved. All these figures are summarized in Table 10.5. The main conclusion that can be extracted from the presented figures is that the introduction of losses with respect to the original images due to the reconstruction process –silhouetteconsistent in this case– is counter-balanced by the efficiency of the representation in terms of required resources.

This last experiment showcases the suitability of the proposed approach in terms of its efficiency. It clearly provides a suitable representation of the multi-view data by exploiting the spatial redundancies in the original views, which could still be enhanced by exploiting temporal redundancies using a dense motion estimation as the one presented in Chapter 5. This efficiency justifies its potential for real-time analysis and visualization applications.

10.3 Conclusions

In this chapter, two different types of experiments have been presented, which validate the multi-view reconstruction approach followed in this thesis, based on the division of the surface reconstruction task in two stages: the first one devoted to the sampling of the surfaces of foreground objects and the second one devoted to their interpolation in order to retrieve continuous surfaces. These two types of experiments, which are grouped in two sets of results, are:

- A comparison to a baseline system, in order to characterize the speediness and effectiveness –usability of the reconstructed surfaces– of the proposed approach, using both synthetic and real data.
- A validation of the effectiveness retrieval of the original images by re-projection, quality of the surfaces- and the efficiency –compression of the multi-view redundancies– of the proposed methodology, using high-quality multi-view data from [4D Repository, 2010].

The first set of experiments has shown how the proposed methodology permits obtaining more accurate reconstructions faster, with smaller computing times when compared to a system composed by classical techniques for volumetric reconstruction –voxelized shape-from-silhouette– and surface extraction –marching cubes–. This reflects the suitability of the design decision of directly reconstructing the surfaces with a set of conveniently designed algorithms in contrast with volumetric techniques, which spend a large part of their computational resources in regions not containing surfaces.

The last part of this first set of experiments also presents, qualitatively, the idea that the proposed representation is effective, for it can be used for a potential free-viewpoint video application. This idea is followed by the first part of the second set of experiments, where the proposed methodology shows its capability for recovering the contents of the original images with a high degree of fidelity by re-projecting the surface reconstructions of the foreground objects and adding background information. It also shows that the quality of the reconstructed surface can rival with that of a state-of-the-art technique (EPVH).

The last experiment in the second set demonstrates the efficiency of the resulting data representation for multi-view video sequences with known and static background. Accepting a reasonable level of losses, due to the reconstruction process -silhouette-consistent-, the resulting surface-based representation greatly reduces the required amount of computational resources for storing or transmitting the multi-view data. This could still be improved by exploiting temporal redundancies in order to represent the foreground -moving- elements in multi-view sequences. As a result of the presented experiments, we have also derived a rule of thumb for the required surface sampling density in order to represent the information contained in the original images with minimal losses: a reasonable number of surface samples is approximately twice the number of foreground pixels in the view with the largest foreground silhouette.

Chapter 11

Conclusions

he main goal of this thesis, summarized in a single sentence, was to obtain a fast –towards real-time–, efficient –within a compressed/limited support–, effective –exploitable by forthcoming applications– and unique surface-based representation of each frame in multi-view video sequences, which provides all the relevant information contained in the original streams and can be used for both interactive visualization and analysis applications.

In order to obtain such surface-based representation, a strategy consisting in dividing the surface reconstruction task in two stages has been chosen. First, one of the surface sampling strategies proposed in the first main part of the dissertation manages to find a dense cloud of surface points with additional information, such as orientation or color. Next, a fast meshing algorithm is responsible for obtaining the closed, continuous surface out of the samples, which can be used to interpolate the sampled surface features, let them be color, texture, orientation, velocity or any other features of interest.

In this closing chapter, we will first summarize the contributions that resulted from the research on effective, efficient and fast strategies for obtaining surface representations suitable for both analysis and visualization. Then, a number of future lines of work are proposed as either extensions or alternatives to the proposed framework.

11.1 Summary of Contributions

Checking the contents of the two parts of this thesis, different achievements have been made in each of the chapters, effectively contributing to the ultimate goals presented in the introduction in Chapter 1. These contributions basically consist in the design of methods to obtain the required output representation in every stage –surface sampling and surface interpolation– which are oriented towards real-time applications.

11.1.1 Surface sampling

In Part I, different strategies for sampling the surfaces of objects of interest (foreground objects in scenes with known, static, background) in multi-view settings have been presented, each with its advantages and drawbacks. However, they all fulfill one of the conditions of the proposed thesis, which is that they should be able to work in scenarios with arbitrarily wide baselines. Three different sampling strategies have been presented, comprising one chapter each.

Image-based surface sampling

In the first chapter devoted to surface sampling, an image-based sampling strategy has been presented, which succeeds at obtaining a pixel-wise photo-consistent surface from a set of views in an arbitrarily-wide baseline setup. This approach, which provides an accurate sampling of the visible features and an implicit handling of visibility, has two important drawbacks: on the one hand, it does not provide of closed surfaces; on the other, the strategy for search of surface samples does not provide a natural mechanism to exploit temporal redundancies in multi-view video sequences.

Surface sampling by deformation

In the second presented strategy for surface sampling, the design is driven by the search of a solution to the two main drawbacks of the previous approach. The technique, which consists in a continuous deformation of an existing surface from frame to frame in order to fit the multi-view silhouette constraints, succeeds at providing a closed surface. Not only the reconstructed surfaces are closed, but the proposed technique is also able to track their evolution in sequences, reducing the search area for surfaces in every instant and, therefore, also reducing the computational cost of reconstruction by exploiting temporal correlations.

However, this technique looses the efficiency at imposing photo-consistency of the previous approach and presents a feature that makes it difficult to use in a platform for massive parallelization, such as a GPU, due to the design of its initialization algorithm. Furthermore, it is costly to obtain a high sampling density in its initialization stage when the working space around objects of interest is large.

Statistical surface sampling

In the third chapter devoted to surface sampling, an alternative approach that drives the search for surface samples in a statistical manner is presented. The main advantage of this approach is that, although it initially generates a sparse surface sampling with a low search efficiency in large volumes -scouting-, an efficient search strategy in regions very likely containing surface samples around existing ones provides a fast sampling of surfaces –propagation–. This feature improves this sampling strategy's ability –compared to that of the one based in the deformation of an existing closed surface- of providing surfaces with a sampling density such that, once re-projected onto the original viewpoints, accurately describe the contents of the original images. In this sense, the algorithm does not only exploit multi-view redundancies, but also spatial correlation on the distribution of surface samples. Furthermore, the algorithm can also benefit from the correlations in the temporal axis. The inefficient initial search for a sparse sampling of the surface is not required from frame to frame. Instead, a search over a suitably smaller range around each sample can be used for scouting. The exploitation of multi-resolution also provides a comparable improvement of the efficiency of the scouting stage.

Even though this technique still presents a limitation when compared to the image-based approach (the application of photo-consistency constraints still requires costly global visibility computations), the algorithm on which it is based is easy to parallelize, which makes it suitable to its application in GPU contexts. In this sense, recent results in a Master's Thesis, developed by the undergraduate student Marc Maceira and co-directed with my advisor, have shown the advantages of using a surface-oriented, massively-parallel algorithm as the statistical surface sampling in a GPU, resulting in a real-time implementation. Furthermore, this sampling strategy is the fastest among the three presented in this thesis, showing the smallest computation time for a given number of surface samples, especially with a large number of input views, where the proposed approach best shows its potential by the efficient exploitation of spatial correlation.

11.1.2 Surface interpolation

In Part II, a fast strategy for generating a closed, continuous surface out of a dense enough set of surface samples is presented, which results in a triangle mesh that can be used to interpolate the information contained in the surface samples. Compared to other meshing techniques in the state-of-the-art, the proposed algorithm is generally faster and provides better accuracy –in terms of Hausdorff distance to ground-truth surfaces–. It is based on a propagation-based scheme which robustly generates a starting triangular oriented face and iteratively propagates through the current contour on the set of available surface samples until completely covering the surface.

The technique is efficient, using suitable data structures both for spatial queries -kd-tree– and for topology control –half-edge structure–. The algorithm also shows robustness to uneven distributions of surface samples. The spatial queries provided by the kd-tree, in contrast to the use of volumetric, fixed-size bins with close surface samples, are such that even samples unevenly distributed can be retrieved at the expense of an increase of computational cost by varying the number of nearest-neighbors query results. Thus, the algorithm does not require additional iterations at different scales in order to provide a closed surface.

A limitation of the proposed, propagation-based meshing algorithm exists with respect to another family of existing techniques, based on the definition of a volumetric representation by the application of a certain spatial regularization method and final surface extraction via *marching cubes*. This limitation is that the location of noisy samples is not corrected by imposing global constraints. However, the reverse of this limitation is that every vertex in the resulting mesh will exactly coincide with an input surface sample. Given a noise-free set of surface samples, this is the most accurate and fastest method to linearly interpolate a continuous surface with the information contained in the samples.

11.1.3 Accomplishments

Regarding the ultimate goals of this thesis as presented in the introduction, the contributions of this thesis can be summarized as follows:

- A unique surface-based representation of the multi-view information data, which is efficient –it compresses the multi-view information –, effective –suitable for analysis and visualization – and accurate –minimizing losses – has been obtained.
- The original images can be approximated –with minimal losses– by re-projection of the unique surface-based representation, which means all the relevant information regarding the original views is kept, with the addition of the 3D structure of the scene.
- Spatial and temporal correlations in multi-view video sequences have been exploited in order to reduce the computation time and make the surface reconstruction strategy suitable for real-time operation for both visualization and analysis applications.
- A triangular mesh that represents a closed surface has been obtained in a usable computation time, as a contrast with many of the techniques in the

state-of-the-art. Such a surface produces a compression of the multi-view data when the static background can be separately treated.

Several publications have resulted, which directly reflect, or are derived from, the achievements related to the methods and the results obtained during the development of this thesis. These are listed in Appendix D.

11.2 Future Work

As stated above, this thesis has succeeded in providing a solution to the problems of defining a methodology to compactly represent multi-view data by approximating the 3D information lost during the imaging process taking part in each camera.

Indeed, the followed approach, consisting in a division of the reconstruction process of the surfaces of foreground objects in multi-view sequences in two stages –surface sampling and surface interpolation– has successfully provided a means to determine the location and orientation of surfaces using a sample-based representation –Part I– and provide topologically correct, closed, continuous surfaces that approximate those of the actual 3D scene –Part II–.

However, some extensions to the methodology proposed in this thesis could be introduced in order to increase the suitability or applicability of the provided 3D data representation to different types of applications, which are out of the scope of this thesis.

11.2.1 Surface sampling

Common to all the surface sampling strategies presented in this thesis, and specially indicated for the image-based surface sampling strategy, it would be interesting to consider a tighter scenario, such as the typical one in *Multi-View Stereo* (MVS). With small-baseline setups, photo-consistency measurements of higher quality, such as *normalized cross-correlation*, can be applied when surfaces contain a suitable amount of texture. Combining high-quality photometric measurements with good quality silhouettes by imposing spatial regularization in a variational formulation could provide a higher level of accuracy than the one obtained under the consideration of arbitrarily wide-baseline setups.

Regarding the strategy for surface sampling by deformation, it would be wise to introduce spatial regularization in the tracking part of the algorithm in a more principled manner. This could be done by introducing the meshing of the surface samples in the sampling strategy with the use of MVS in small-baseline scenarios and without the strong limitation of close-to-real-time performance: once the initial set of high-quality surface samples have been obtained, the surface can be assumed to have the correct topology –which cannot be said from silhouette-consistent surfaces– and subsequent deformations can be applied on the resulting surface by using linear programming [Botsch and Sorkine, 2008], or even non-linear approaches [Botsch et al., 2006] in order to fit the new data constraints imposed by each new frame.

In the statistical strategy for surface sampling, it would be interesting to drive the sampling procedure by the scene contents of interest for a given application. For example, if gesture recognition was to be applied on multi-view data, an enhanced level of detail in the representation of limbs and head could be used to improve the fitting of an existing model –with approximately known pose– to the reconstructed data, using algorithms such as *Iterative Closest Point* or other variants [Rusinkiewicz and Levoy, 2001]. Such improvement in accuracy could be obtained by favoring the propagation in those areas of interest at the cost of a smaller sampling density over the rest of the surface. It would furthermore conform an interesting showcase of the suitability of the proposed statistical sampling strategy, driven by the exploitation of spatial correlation on the location of surface samples.

11.2.2 Surface interpolation

With respect to the meshing algorithm presented in this thesis, which is used in the surface interpolation stage, it could be sped-up by dividing the meshing of the complete surface into the meshing of several parts and posterior stitching. This approach would better exploit the multi-core architecture of current CPUs, although it would also introduce the problem of correctly stitching the different surface patches reconstructed in each thread without producing topological errors.

11.2.3 Background reconstruction

Last but not least, in the development of this thesis it has not been discussed how to proceed with the background. Indeed, assuming all the information of interest is contained in the foreground of the scene, this part of the image contents –which, in wide-lens cameras can take around 90% of the area– has been deliberately left apart for only using it in order to approximate the original images by re-projection of the reconstructed surfaces.

When the background occludes any of the foreground objects for any of the views, this cannot be ignored and its effect has to be considered in the reconstruction of the dynamic part of the scene. In order to be able to obtain a 3D representation of the background, expensive laser-scan techniques could be used at a pre-processing stage. However, more flexible methods for reconstructing wellstructured large regions without rich texture could also be applied, related to the approach presented in [Furukawa et al., 2009].

11.3 Perspectives

Observing the impact of multi-view systems in the research community, it seems reasonable to imagine a future arrival of systems exploiting the enriched available information of multi-view recording to the mainstream audio-visual industry.

Following the evolution of TV systems –first from black-and-white to color systems, then from color to high-definition systems, now from high-definition to stereoscopic (3D) systems–, a likely step ahead in the future is free-viewpoint video. In order to create realistic applications of this type, it will be necessary to augment the audio-visual transmitted information. Whereas there are on-going works about the efficient transmission of multi-view streams, with this thesis we clearly recall the alternative solution, consisting in representing the dynamic part of the multi-view scene (foreground) by using a unique 3D support, treated separately from the rest of the elements of the scene included in the background.

In the last years, an increasing part of the bi-directional communication has jumped from telephonic calls to video-conferencing, providing an enriched telecommunication experience. Related research on haptic technology attempts to provide a new dimension of realism to this experience, as well as in the field of humancomputer interaction (HCI). There is also on-going research in systems for bidirectional video-communication with multi-view systems [3D Presence, 2010, Vision, 2010]. These can also contribute to the realistic perception for the interlocutors by allowing free mobility in a virtual conferencing space.

In order to achieve this, an important challenge will have to be faced by the audiovisual industry: the introduction of multi-view settings as part of mainstream audiovisual setups not only in recording studios, but also in homes. In this aspect, we have already seen how monocular cameras have been added to all types of devices -e.g. laptops, cell phones or video-gaming systems–. Also, one of the most quickly market-penetrating new technology devices in history –Microsoft's Kinect– has open the door to visual communication increased with depth and advanced HCI to the mainstream.

The next step is to increase the number of sensors in the user space. Although this is a critical step, it has already proven to be accessible in audio applications: starting from the last decade of the 20th century, a large amount of users have switched from stereo reproduction systems to multi-channel ones, despite of the higher space, cabling and cost requirements of the latter. This shows that, given a sufficient amount of new applications, quality and features of interest for the target user, may be motivating enough to adopt multi-view bi-directional communication.

In this second application field, the restriction on the smallest achievable baseline shows up even more clearly than in the scenarios considered in this thesis. It is in these new scenarios where the segregation of foreground and background parts of the scene is not only convenient, but also necessary in order to successfully process the rather sparse multi-view data that is expected to be available, at least in a near future.

To sum up, we expect to see a growing number of systems exploiting multi-view data, either for communication, analysis for HCI or free-viewpoint television, where some of the concepts in this thesis find an application towards economical, accessible systems for the mainstream audiovisual industry. Part IV

Appendices

Appendix A

Camera Model

n order to relate the contents of the real world with that of a number of cameras capturing views of it, a parametric model of the camera has to be considered. The parameters of such a model can be obtained by different camera calibration methods, the working principle of which lies beyond the scope of this appendix. In the following, the availability of calibrated parameters is assumed, with a re-projection error smaller than one pixel.

A camera model is generally described by means of two types of parameters, which either model the position and orientation of the camera *-extrinsic parameters-* or geometric aspects regarding the image creation process of the device *-intrinsic parameters-*. Then, few concepts about back-projection and epipolar geometry are discussed. A more complete introduction to camera geometry can be found in [Hartley and Zisserman, 2000].

A.1 Extrinsic Parameters

The extrinsic parameters describe the orientation and position of a camera with respect to a reference coordinate system. Therefore, this part of the camera model can be properly described by means of a rigid transformation, consisting in a rotation followed by a translation. Such transformation, once applied onto a 3D point described by means of the reference coordinate system, provides the new coordinates of the point referred to a new coordinate system aligned with the camera with its origin at its center of projections –the "camera position"–.

Let **X** be a 3D point in the reference coordinate system of the real world, R the axis rotation, **T** the translation and **X'** the same 3D point referred to the coordinate

system aligned with the camera. Then,

$$\mathbf{X}' = R\mathbf{X} + \mathbf{T} = \begin{pmatrix} R_{11} & R_{12} & R_{13} \\ R_{21} & R_{22} & R_{23} \\ R_{31} & R_{32} & R_{33} \end{pmatrix} \mathbf{X} + \begin{pmatrix} T_1 \\ T_2 \\ T_3 \end{pmatrix}$$
(A.1)

Although it could also be equivalently described in reverse order –a translation followed by a rotation–, it is more convenient the former option. The reason is that, when using homogeneous coordinates, the previous expression can be compacted to a single matrix operation. Let $\tilde{\mathbf{X}}$ be the point \mathbf{X} in homogeneous coordinates, obtained as

$$\tilde{\mathbf{X}} = \begin{pmatrix} \mathbf{X} \\ 1 \end{pmatrix}. \tag{A.2}$$

Using homogeneous coordinates, Equation A.1 can be written as

$$\mathbf{X}' = \begin{pmatrix} R_{11} & R_{12} & R_{13} & T_1 \\ R_{21} & R_{22} & R_{23} & T_2 \\ R_{31} & R_{32} & R_{33} & T_3 \end{pmatrix} \tilde{\mathbf{X}}.$$
 (A.3)

A.1.1 Camera position

The position of the camera can be described by the position in world coordinates at which its center of projections lies. Let \mathbf{X}_{CoP} be such point and \mathbf{X}'_{CoP} the same point described in the coordinate system aligned with the camera. By construction of the camera-aligned coordinate system,

$$\mathbf{X}_{CoP}' = \begin{pmatrix} 0\\0\\0 \end{pmatrix}. \tag{A.4}$$

Then, \mathbf{X}_{CoP} can be obtained as

$$\mathbf{X}_{CoP} = R^{-1} \left(\mathbf{X}_{CoP}' - \mathbf{T} \right) = -R^{-1} \mathbf{T}.$$
 (A.5)

Using the property that a rotation matrix is orthogonal $(RR^{\top} = R^{\top}R = I)$, the previous expression can be written as

$$\mathbf{X}_{CoP} = -R^{\top}\mathbf{T}.\tag{A.6}$$

A.1.2 Camera orientation

The camera-aligned coordinate system has the z'-axis orthogonal to the image plane, oriented towards the scene. The x'-axis has the usual horizontal direction on the image plane, oriented from left to right, whereas the y'-axis corresponds to the

vertical axis on the image plane, oriented from top to bottom. Thus, the camera orientation \mathbf{z}' equals the third row of the rotation matrix R:

$$\mathbf{z}' = \begin{pmatrix} R_{31} \\ R_{32} \\ R_{33} \end{pmatrix}.$$
 (A.7)

A.2 Intrinsic Parameters

The intrinsic parameters describe the geometry of the imaging process that takes place in the camera. Some models assume the absence of image distortion that appear when using optical lenses, whereas others take this effect into account and model it conveniently.

A.2.1 Pinhole camera model

The pinhole camera consists in a small aperture without a lens through which light rays traverse. Therefore, it does not contain optical distortion that appears when using a lens. This model can be conveniently expressed in terms of the *calibration* matrix K:

$$K = \begin{pmatrix} f \cdot s_x & \lambda & c_x \\ 0 & f \cdot s_y & c_y \\ 0 & 0 & 1 \end{pmatrix}.$$
 (A.8)

In this model, $f \cdot s_x$ and $f \cdot s_y$ are the horizontal and vertical focal lengths, respectively, expressed in pixels units. Since most cameras nowadays do not present centering imperfections, the *skewness* parameter λ can be assumed to be zero. Finally, (c_x, c_y) are the coordinates of the camera's principal point (ideally at the image center) also in pixel units. With this, the pixel coordinates $\mathbf{x} = (u, v)$ corresponding to the projection of a 3D point expressed in the camera coordinate system \mathbf{X}' can be obtained as

$$z' \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = K\mathbf{X}', \tag{A.9}$$

where the natural perspective correction (scaling the pixel coordinates by the point's depth with respect to the camera z') is also applied.

A.2.2 Lens distortion

However, real lenses usually have some distortion, mostly radial and slight tangential. Let $z'\tilde{\mathbf{x}}' = z'(\tilde{x}', \tilde{y}', 1)^{\top}$ be the perspective-corrected camera coordinates before scaling with the calibration matrix K. The distorted equivalent, $z'\tilde{\mathbf{x}}'_d = z'(\tilde{x}'_d, \tilde{y}'_d, 1)^{\top}$, can be obtained as:

$$\tilde{x}'_{d} = \tilde{x}'(1 + k_{1}r^{2} + k_{2}r^{4} + \dots) + 2p_{1}\tilde{x}'\tilde{y}' + p_{2}(r^{2} + 2\tilde{x}'^{2})$$

$$\tilde{y}'_{d} = \tilde{y}'(1 + k_{1}r^{2} + k_{2}r^{4} + \dots) + p_{1}(r^{2} + 2\tilde{y}'^{2}) + 2p_{2}\tilde{x}'\tilde{y}'$$
(A.10)

where $r^2 = \tilde{x}'^2 + \tilde{y}'^2$, k_1 and k_2 are the normalized radial distortion coefficients, and p_1 and p_2 are the normalized tangential distortion coefficients. The distorted pixel coordinates $\mathbf{x}_d = (u_d, v_d)$ can be obtained as

$$z' \begin{pmatrix} u_d \\ v_d \\ 1 \end{pmatrix} = K \begin{pmatrix} z'\tilde{x}'_d \\ z'\tilde{y}'_d \\ z' \end{pmatrix}.$$
 (A.11)

Distortion correction

In order to correct the distortion present in the images captured by a real-world camera –especially those with a wide-angle lens–, the image must be warped by the inverse function applied in order to distort the projected coordinates. This can be easily done for every undistorted pixel $\mathbf{x} = (u, v)$ by computing its distorted coordinate $\mathbf{x}_d = (u_d, v_d)$ using Eqs. A.10 and A.11 and taking the bi-linear interpolation of the values in the four closest pixels to \mathbf{x}_d as the value for the undistorted pixel \mathbf{x} .

With this, undistorted images are slightly blurred but, in exchange, camera calibration can be used without taking distortion into consideration. The complete set of matrices in order to obtain the image coordinates $\mathbf{x} = (u, v)$, corresponding to the projection a 3D point in homogeneous world coordinates $\tilde{\mathbf{X}}$ onto a camera image with calibration matrix K and pose $(R|\mathbf{T})$ is:

$$z'\begin{pmatrix} u\\v\\1 \end{pmatrix} = K(R|\mathbf{T})\tilde{\mathbf{X}} = P\tilde{\mathbf{X}},$$
 (A.12)

where P is called the *projection* matrix.

A.3 Back-Projection

Two different types of back-projection of an image point are considered. If the depth z' of the image point is known (range data, stereo, ...), the exact position of the 3D point originating the image point can be obtained in a straightforward manner, whereas in a more general case where the depth is unknown, the back-projection will result in a ray passing through the camera center and the pixel on the image plane.

A.3.1 Known depth

Given an image point $\mathbf{x} = (u, v)$ with known depth z', the equivalent point before perspective correction –projective coordinates– is $z'\tilde{\mathbf{x}} = z'(u, v, 1)^{\top}$. The corresponding 3D point expressed in world coordinates \mathbf{X} with known camera parameters K and $(R|\mathbf{T})$ can be simply computed as

$$\mathbf{X} = R^{\top} \left(K^{-1} z' \tilde{\mathbf{x}} - \mathbf{T} \right) = z' (KR)^{-1} \tilde{\mathbf{x}} + \mathbf{X}_{CoP},$$
(A.13)

where $\mathbf{X}_{CoP} = -R^{\top}\mathbf{T}$ is the camera center.

A.3.2 Unknown depth

Given a point \mathbf{x} without known depth in an image and known camera parameters K and $(R|\mathbf{T})$, an infinite set of 3D points map to this image point. This set is a ray in 3D space, Γ_x , passing through the camera center $\mathbf{X}_{CoP} = -R^{\top}\mathbf{T}$. The other known point of the ray is the image point.

Since in this case the depth is unknown, we define a parametric equivalent point before perspective correction as $\lambda \tilde{\mathbf{x}} = \lambda(u, v, 1)^{\top}$. Proceeding in the same manner as in the previous case with known depth, we obtain

$$\mathbf{X}(\lambda) = R^{\top} \left(K^{-1} \lambda \tilde{\mathbf{x}} - \mathbf{T} \right) = \lambda (KR)^{-1} \tilde{\mathbf{x}} + \mathbf{X}_{CoP},$$
(A.14)

or, expressed as a set,

$$\Gamma_x = \left\{ \mathbf{X}(\lambda) = \mathbf{X}_{CoP} + \lambda (KR)^{-1} \tilde{\mathbf{x}} | \lambda \in \mathbb{R}^+ \right\}.$$
(A.15)

In this case, the unknown depth z' is used as a parameter, λ , that produces a back-projected ray with its variation.

A.4 Epipolar Geometry

The epipolar geometry between two views is based on three entities:

- The *epipole* is the point of intersection of the line joining the two camera centers (the baseline) with the image plane. Equivalently, the epipole is the image in one view of the camera center of the other view.
- An *epipolar plane* is a plane containing the baseline.
- An *epipolar line* is the intersection of an epipolar plane with the image plane. As all epipolar planes contain the baseline, all epipolar lines must meet at the epipole.

Let $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_2$ be two image points in projective coordinates in cameras 1 and 2, respectively, and let \mathbf{l}_1 and \mathbf{l}_2 be the epipolar lines of $\tilde{\mathbf{x}}_2$ in camera 1 and of $\tilde{\mathbf{x}}_1$ in camera 2, respectively. There exists a rank-2 matrix F_{12} –which is called the fundamental matrix– satisfying

$$\mathbf{l}_2 = F_{12}\tilde{\mathbf{x}}_1. \tag{A.16}$$

By definition, $\tilde{\mathbf{x}} \in \mathbf{l}_2 \Leftrightarrow \tilde{\mathbf{x}}^\top \mathbf{l}_2 = 0$, which results in

$$\tilde{\mathbf{x}}^{\top} F_{12} \tilde{\mathbf{x}}_1 = 0, \quad \forall \tilde{\mathbf{x}} \in \mathbf{l}_2.$$
(A.17)

As a consequence of this, and the symmetry of the problem, it is also true that

$$\mathbf{l}_1 = F_{21}\tilde{\mathbf{x}}_2 = F_{12}^{\top}\tilde{\mathbf{x}}_2. \tag{A.18}$$

Related to the back-projection of an image point, line l_1 corresponds to the image in camera 1 of the back-projected ray of image point \mathbf{x}_2 in camera 2.

A.4.1 Derivation of the fundamental matrix

Let P_1 and P_2 be the projection matrices of cameras 1 and 2, respectively. We denote by $\tilde{\mathbf{X}}_{CoP,1}$ and $\tilde{\mathbf{X}}_{CoP,2}$ the centers of projection of cameras 1 and 2, respectively, in homogeneous world coordinates. Finally, let $\tilde{\mathbf{e}}_1 = P_1 \tilde{\mathbf{X}}_{CoP,2}$ and $\tilde{\mathbf{e}}_2 = P_2 \tilde{\mathbf{X}}_{CoP,1}$ be the epipoles in projective coordinates in cameras 1 and 2, respectively.

Under these assumptions, we can affirm that the following points belong to the back-projected ray of \mathbf{x}_1 :

- The projection center of camera 1: $\tilde{\mathbf{X}}_{CoP,1}$.
- $P_1^+ \tilde{\mathbf{x}}_1$, where $P_1^+ = P_1^\top (P_1 P_1^\top)^{-1}$ is a pseudo-inverse matrix of P_1 , *i.e.* $P_1 P_1^+ = I_3$.

As a consequence, $P_2(P_1^+ \tilde{\mathbf{x}}_1)$ and $\tilde{\mathbf{e}}_2 = P_2 \tilde{\mathbf{X}}_{CoP,1}$ will be two points lying in the epipolar line \mathbf{l}_2 . Thus, the epipolar line \mathbf{l}_2 can be computed as

$$\mathbf{l}_2 = \tilde{\mathbf{e}}_2 \times P_2(P_1^+ \tilde{\mathbf{x}}_1) = F_{12} \tilde{\mathbf{x}}_1.$$
(A.19)

So the final expression of the fundamental matrix is

$$F_{12} = [\tilde{\mathbf{e}}_2]_{\times} P_2 P_1^+, \tag{A.20}$$

where the cross product has been expressed as a matrix product with a skew-symmetric matrix $[\tilde{\mathbf{e}}_2]_{\times}$ constructed as

$$[\tilde{\mathbf{e}}_2]_{\times} = \begin{pmatrix} 0 & -\tilde{\mathbf{e}}_{2,z} & \tilde{\mathbf{e}}_{2,y} \\ \tilde{\mathbf{e}}_{2,z} & 0 & -\tilde{\mathbf{e}}_{2,x} \\ -\tilde{\mathbf{e}}_{2,y} & \tilde{\mathbf{e}}_{2,x} & 0 \end{pmatrix}.$$
 (A.21)

Appendix B

kd-Tree

kd-tree or k-dimensional tree is a space-partitioning data structure for organizing points in a k-dimensional space. It is useful for several applications, such as searches involving a multi-dimensional search key, e.g. range searches or nearest-neighbor searches [Friedman et al., 1977].

The *kd*-tree is a special case of binary partition tree in which every *node* is a *k*dimensional point. Every non-*leaf* node can be thought of as implicitly generating a splitting hyperplane that divides the space into two parts, known as subspaces. Points to the left of this hyperplane represent the left sub-tree of that node and points right of the hyperplane are represented by the right sub-tree.

The hyperplane direction is chosen in the following way: every node in the tree is associated with one of the k-dimensions, with the hyperplane perpendicular to that dimension's axis. So, for example, if for a particular split the x axis is chosen, all points in the subtree with a smaller x value than the node will appear in the left subtree and all points with larger x value will be in the right sub-tree. In such a case, the hyperplane would be set by the x-value of the point, and its normal would be the unit x-axis.

B.1 Construction

Since there are many possible ways to choose axis-aligned splitting planes, there are many different ways to construct kd-trees. The canonical method of kd-tree construction has the following constraints:

• As one moves down the tree, one cycles through the axes used to select the splitting planes. (For example, the root would have an *x*-aligned plane, the



Figure B.1: Steps in the construction of a kd-tree with maximum-variance splitting axis. From left to right and top to bottom, a new point from the set at the left side is added to the tree at the right side of each sub-figure

root's children would both have y-aligned planes, the root's grandchildren would all have z-aligned planes, the next level would have an x-aligned plane, and so on.)

• Points are inserted by selecting the median of the points being put into the subtree, with respect to their coordinates in the axis being used to create the splitting plane. (Note the assumption that we feed the entire set of points into the algorithm up-front.)

This method leads to a *balanced kd*-tree, in which each leaf node is about the same distance from the root. However, balanced trees are not necessarily optimal for all applications.

Given a list of n points, the following algorithm will construct a kd-tree containing those points. The inputs are a list of ordered points and the position of the *first* and *last* elements to process in the list. Then, the recursive algorithm goes as follows:

- 1. If the *first* and *last* element are the same, return immediately.
- 2. Create a node, containing the splitting axis and the location of the median in the specified range.
- 3. Sort the list in the specified range by increasing value in the splitting axis.
- 4. Recursively create the left sub-tree by calling the function with last := medianand keeping the original *first*.
- 5. Recursively create the right sub-tree by calling the function with first := median + 1 and keeping the original *last*.
- 6. Return the newly created node.

In Fig. B.1, a *kd*-tree is created by the criterion of maximum-variance splitting axis. This allows obtaining good search performance in sets including outliers.

B.1.1 Dynamic addition of elements

A new point can also be added to a kd-tree after its construction. First, the tree must be traversed, starting from the root and moving to either the left or the right child depending on whether the point to be inserted is on the *left* or *right* side of the splitting plane. Once the node under which the child should be located is reached, the new point is added as either the left or right child of this leaf node, again depending on which side of the node's splitting plane contains the new node.

However, adding points in this manner can cause the tree to become unbalanced, leading to decreased tree performance. The rate of tree performance degradation is dependent upon the spatial distribution of tree points being added, and the number of points added in relation to the tree size. If a tree becomes too unbalanced, it may need to be re-balanced to restore the performance of queries, such as nearestneighbor searching –introduced below–. A reasonable choice for re-balancing the tree is to construct it from ground, using the current set of points, every time the ratio between the number of dynamic insertions since the last re-balancing and the total number of elements of the tree surpasses a certain threshold.

B.2 Queries

Once the tree is constructed, two types of queries can be processed efficiently, in logarithmic time with respect to the number of elements contained in the tree. The first one, (k)-nearest-neighbors search, is able to find the k closest points to a given one, whereas the second one, range search, is able to find all the points at a distance smaller than a maximum value from a given point.

B.2.1 Nearest-neighbors search

The nearest-neighbors (NN) algorithm aims to find the points in the tree which are nearest to a given input point. This search can be done efficiently, in logarithmic time, by using the tree properties to quickly eliminate large portions of the search space. Searching for the nearest neighbor in a kd-tree proceeds as follows:

- Starting with the root node, the algorithm moves down the tree recursively, in the same way that it would if the search point were being inserted (i.e. it goes right or left depending on whether the point is greater or less than the current node in the split dimension).
- Once the algorithm reaches a leaf node, it saves that node point as the *current* best

The algorithm unwinds the recursion of the tree, performing the following steps at each node:

- If the current node is closer than the current best, then it becomes the current best.
- The algorithm checks whether there could be any points on the other side of the splitting plane that are closer to the search point than the current best. In concept, this is done by intersecting the splitting hyperplane with a hypersphere around the search point that has a radius equal to the current nearest distance. Since the hyperplanes are all axis-aligned this is implemented as a simple comparison to see whether the difference between the splitting coordinate of the search point and current node is less than the distance (overall coordinates) from the search point to the current best.
- If the hypersphere crosses the plane, there could be nearer points on the other side of the plane, so the algorithm must move down the other branch of the tree from the current node looking for closer points, following the same recursive process as for the entire search.
- If the hypersphere does not intersect the splitting plane, then the algorithm continues walking up the tree, and the entire branch on the other side of that node is eliminated from the search.

When the algorithm finishes this process for the root node, then the search is complete.

The algorithm can be extended in several ways by simple modifications. It can provide the k-nearest neighbors to a point by maintaining the list of the k current
bests instead of just one. Branches are only eliminated when they cannot have points closer than any of the k current bests.

B.2.2 Range search

As seen in [Preparata and Shamos, 1985], the complexity of finding all points within a given distance of a target point -i.e., performing a *range search*- is still logarithmic in the size of the tree -i.e., the number of points it contains, n- for a fixed range distance.

Let \mathbf{t} be the target point and r the range distance for the search. Then, this operation can be implemented as a modification of the previous NN algorithm:

- 1. The distances from the visited points to \mathbf{t} are not compared to that of the closest point found, but to the fixed initial value r.
- 2. As in k-nearest-neighbors, all the discovered points within this distance are returned in a list, not just the closest.

B.3 Complexity

Building a static kd-tree from n points takes $O(n \log^2 n)$ time if an $O(n \log n)$ sort is used to compute the median at each level. The insertion of a new point into a balanced kd-tree takes $O(\log n)$ time.

B.3.1 High-dimensional data

Finding the nearest point is an $O(\log n)$ operation in the case of randomly distributed points. However, analysis of binary search trees has found that the worst case search time for a k-dimensional kd-tree containing n nodes is given by the following equation:

$$t_{worst} = O(k \cdot n^{1-1/k}). \tag{B.1}$$

These poor running times only apply when n is in the order of the number of dimensions k. In very high dimensional spaces, the curse of dimensionality causes the algorithm to need to visit many more branches than in lower dimensional spaces. In particular, when the number of points is only slightly higher than the number of dimensions, the algorithm is only slightly better than a linear search of all of the points. As a general rule, if the dimensionality is k, then the number of points in the data n should be such that $n \gg 2k$.

Appendix C

Hausdorff Distance

he *Hausdorff distance*, named after Felix Hausdorff, measures how far two subsets of a metric space are from each other. It turns the set of non-empty compact subsets of a metric space into a metric space in its own right.

Informally, two sets are close in the Hausdorff distance if every point of either set is close to some point of the other set. Although more details about the Hausdorff metric can be found in [Henrikson, 1999], here we will introduce just some basic notions.

C.1 Definition

The Hausdorff distance $d_H(X, Y)$ between two non-empty subsets X and Y of a metric space (M, d) is defined as

$$d_H(X,Y) = \max\{\sup_{x \in X} \inf_{y \in Y} d(x,y), \sup_{y \in Y} \inf_{x \in X} d(x,y)\}.$$
 (C.1)

Equivalently,

$$d_H(X,Y) = \inf\{\epsilon > 0; X \subseteq Y_\epsilon \text{ and } Y \subseteq X_\epsilon\},$$
(C.2)

where

$$X_{\epsilon} := \bigcup_{x \in X} \{ z \in M; d(z, x) \le \epsilon \},$$
(C.3)

that is, the set of all points within ϵ of the set X.

Remark. It is not true in general that, if $d_H(X, Y) = \epsilon$, then

$$X \subseteq Y_{\epsilon}$$
 and $Y \subseteq X_{\epsilon}$.

For instance, consider the metric space of the real numbers \mathbb{R} with the usual metric d induced by the absolute value,

$$d(x,y) := |y - x|, \qquad x, y \in \mathbb{R}.$$

Take

$$X := \{1/n; n \in \mathbb{N}\}$$
 and $Y := \{-1/n; n \in \mathbb{N}\}.$

Then $d_H(X,Y) = 1$. However $X \nsubseteq Y_1$ because $Y_1 = [-2,1)$, but $1 \in X$.

C.1.1 Properties

Some properties can be derived from the definition of the Hausdorff distance:

- In general, $d_H(X, Y)$ may be infinite. If both X and Y are bounded, then $d_H(X, Y)$ is guaranteed to be finite.
- We have $d_H(X, Y) = 0$ if and only if X and Y have the same closure.
- On the set of all non-empty subsets of M, d_H yields an extended pseudometric.
- On the set F(M) of all non-empty compact subsets of M, d_H is a metric. If M is complete, then so is F(M). If M is compact, then so is F(M). The topology of F(M) depends only on the topology of M, not on the metric d.

C.2 Motivation

The definition of the Hausdorff distance can be derived by a series of natural extensions of the distance function d(x, y) in the underlying metric space M, as follows:

• Define a distance function between any point x of M and any non-empty set Y of M by:

$$d(x,Y) = \inf\{d(x,y)|y \in Y\}.$$

For example, d(1, [3, 6]) = 2 and d(7, [3, 6]) = 1.

• Define a distance function between any two non-empty sets X and Y of M by:

 $d(X,Y) = \sup\{d(x,Y) | x \in X\}.$

For example, d([1, 7], [3, 6]) = d(1, [3, 6]) = 2.

• If X and Y are compact, then d(X, Y) will be finite; d(X, X) = 0; and d inherits the triangle inequality property from the distance function in M. As it stands, d(X, Y) is not a metric because d(X, Y) is not always symmetric,

and d(X,Y) = 0 does not imply that X = Y (It does imply that $X \subseteq Y$). For example, d([1,3,6,7],[3,6]) = 2, but d([3,6],[1,3,6,7]) = 0. However, we can create a metric by defining the Hausdorff distance to be:

$$d_H(X,Y) = max\{d(X,Y), d(Y,X)\}$$

C.3 Applications

In feature matching, the Hausdorff distance can be used to find a given template in an arbitrary target image [Alhichri and Kamel, 2002]. The template and target image can be pre-processed via a certain feature extractor giving a set of detected feature points. Next, each detected point in the template image is treated as a point in a set. Similarly, an area of the target image is also treated as a set of points. The algorithm then tries to minimize the Hausdorff distance between the template and some area of the target image. The area in the target image with the minimal Hausdorff distance to the template, can be considered the best candidate for locating the template in the target.

C.3.1 Polygon mesh dissimilarity

In computer graphics, the Hausdorff distance is used to measure the difference between two different representations of the same 3D object. Although by definition the Hausdorff distance is the one-sided maximum of the minimum distance, in practice the mean or RMS measures are quite useful, since these are much less sensitive to outliers.

In [Cignoni et al., 1998], a tool is presented, able to compare the difference between a pair of surfaces (*e.g.* a ground-truth mesh and an estimate obtained from a set of oriented points) by adopting a surface sampling approach. It uses the Hausdorff distance in order to return, among other, maximum and mean error. Recently, this tool has been integrated in the open source mesh editor *Meshlab* [Cignoni et al., 2008].

Appendix D

Related Publications

This appendix contains a list of publications related to the thesis that either have been accepted or are pending acceptance at the moment of finishing the dissertation. These publications are classified by their correspondence to the different part in which this thesis is divided.

D.1 Surface Sampling

Three publications –one corresponding to each proposed sampling strategy– have resulted from the research on efficient methods for extraction of dense sets of samples representing surfaces of foreground objects in scenes with known and static background.

Photo-Consistent Surfaces from a Sparse Set of Viewpoints. In this ICIP 2010 conference article [Salvador and Casas, 2010b], the image-based approach for surface sampling introduced in Chapter 4 is presented as a method capable of obtaining photo-consistent surfaces for foreground objects in arbitrarily wide-baseline scenarios. A comparison with a volumetric approach, also included in the first section of Chapter 7, demonstrates the adequacy of a surface-based approach for 3D reconstruction of the visible features in multi-view settings.

Joint Estimation of Shape and Motion from Silhouettes. In this second ICIP 2010 conference article [Salvador and Casas, 2010a], the strategy for surface sampling based on the deformation of an existing closed surface introduced in Chapter 5 is presented as an efficient technique for obtaining silhouette-consistent surfaces in multi-view video sequences. Besides this result, the possibility of obtaining a dense motion field estimate on the surface samples as a simple position difference between consecutive time instants was also introduced as a possible valuable feature for further analysis applications. The experiments presented in this article are also contained in the second section of Chapter 7.

Statistical Surface Sampling for Multi-View Reconstruction. In the ICCV 2011 submission [Salvador and Casas, 2011], the statistical strategy for surface sampling introduced in Chapter 6 is presented as a valuable tool for efficiently extracting a discrete representation of the silhouettes of foreground objects that can be used for interactive *Free-Viewpoint Video* by using it as input in the meshing algorithm presented in Chapter 8. The experiments, introducing an efficiency improvement in the scouting stage of the algorithm –dynamic or multi-resolution–, are reflected in the third section of Chapter 7. The decision upon its acceptance will be announced to the authors in June 2011.

D.2 Surface Interpolation and Applications

One publication has been made, which introduces the strategy for surface interpolation from a set of oriented surface points.

From Silhouettes to 3D Points to Mesh: Towards Free Viewpoint Video. In this 3DVP 2010 workshop article [Salvador et al., 2010], the meshing algorithm presented in Chapter 8 is presented as part of a complete system for efficiently obtaining surface descriptions suitable for Free-Viewpoint Video applications. Being the statistical sampling stage still under development, the sampling strategy chosen for this approach corresponds to the one presented in the first part of Chapter 5 with the addition of some post-processing stages also used for the statistical sampling approach. These algorithms are also presented in the last sections of Chapter 6. The presented results compare the presented surface reconstruction strategy with a baseline system based on volumetric techniques. Although the differences are not as great as the ones obtained by using the statistical sampling strategy presented in the first section of Chapter 10, the experimental results show the adequacy of the surface-oriented methodology.

Finally, an application for the improvement of foreground silhouettes in multiview settings has also been proposed, partially based on some techniques presented in this thesis.

Joint Multi-view Foreground Segmentation and 3D Reconstruction with Tolerance Loop. In the ICIP 2011 submission [Gallego et al., 2011], an application of the conservative estimate of the visual hull with tolerance to segmentation errors is introduced. The main idea is to iteratively refine noisy foreground silhouettes by introducing *a priori* information from the projections of gradually less conservative visual hulls. In this case, a volumetric representation is chosen in order to introduce spatial regularization using a state-of-the-art energy minimization technique (graph-cuts). Then, the surface extracted using marching cubes is fitted to the noisy silhouettes –with tolerance to errors– by using a technique similar to the one presented in the part corresponding to the dynamic sampling of Chapter 5. This application of the method allows to improve the precision of the projection of the reconstructed low-resolution surface. The results show a clear improvement on the quality of the final estimate of the visual hull. The decision upon its acceptance will be announced to the authors in April 2011.

At the moment of finishing this thesis, a new article intended for its publication in *Transaction on Image Processing* is in preparation. This article will present the complete reconstruction approach (surface sampling and interpolation) with new experiments, comparing the proposed technique to state of the art methods like the concatenation of *Patch-based Multi-view Stereo* [Furukawa and Ponce, 2009a, PMVS, 2010] and *Poisson reconstruction* [Kazhdan et al., 2006].

Bibliography

- [3D Presence, 2010] 3D Presence (2010). Seventh Research Framework Programme. http://www.3dpresence.eu/.
- [4D Repository, 2010] 4D Repository (2010). INRIA. 4drepository.inrialpes.fr.
- [Ahmed et al., 2008] Ahmed, N., Theobalt, C., Rossl, C., Thrun, S., and Seidel, H. (2008). Dense correspondence finding for parametrization-free animation reconstruction from video. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1–8.
- [Alcoverro and Pardàs, 2009] Alcoverro, M. and Pardàs, M. (2009). Voxel occupancy with viewing line inconsistency analysis and spatial regularization. In VISAPP (2), pages 464–469.
- [Alhichri and Kamel, 2002] Alhichri, H. S. and Kamel, M. (2002). Multi-resolution image registration using multi-class hausdorff fraction. *Pattern Recognition Let*ters, 23:279–286.
- [Aspert et al., 2002] Aspert, N., Santa-Cruz, D., and Ebrahimi, T. (2002). MESH: Measuring errors between surfaces using the Hausdorff distance. In *Proc. IEEE Int. Conf. on Multimedia and Expo*, volume 1, pages 705–708.
- [Bernardini et al., 1999] Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., and Taubin, G. (1999). The ball-pivoting algorithm for surface reconstruction. *Visualization and Computer Graphics, IEEE Transactions on*, 5(4):349–359.
- [Botsch et al., 2006] Botsch, M., Pauly, M., Gross, M., and Kobbelt, L. (2006). Primo: coupled prisms for intuitive surface modeling. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, SGP '06, pages 11–20, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.
- [Botsch and Sorkine, 2008] Botsch, M. and Sorkine, O. (2008). On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics*, 14:213–230.

- [Campbell et al., 2008] Campbell, N. D. F., Vogiatzis, G., Hernández, C., and Cipolla, R. (2008). Using multiple hypotheses to improve depth-maps for multiview stereo. In *Proc. European Conf. on Computer Vision*, pages 766–779.
- [Carceroni and Kutulakos, 2002] Carceroni, R. and Kutulakos, K. (2002). Multiview scene capture by surfel sampling: From video streams to non-rigid 3D motion, shape and reflectance. *Int. Journal on Computer Vision*, 49(2-3):175–214.
- [Carceroni and Kutulakos, 2001] Carceroni, R. L. and Kutulakos, K. N. (2001). Multi-view scene capture by surfel sampling: from video streams to non-rigid 3D motion, shape and reflectance. In *Proc. IEEE Int. Conf. on Computer Vi*sion, volume 2, pages 60–67.
- [Casas and Salvador, 2006] Casas, J. R. and Salvador, J. (2006). Image-based multi-view scene analysis using 'conexels'. In VisHCI '06: Proceedings of the HCSNet workshop on Use of vision in human-computer interaction, pages 19–28, Darlinghurst, Australia, Australia. Australian Computer Society, Inc.
- [Chen and Donoho, 1994] Chen, S. and Donoho, D. (1994). Basis pursuit. In Signals, Systems and Computers, 1994. 1994 Conference Record of the Twenty-Eighth Asilomar Conference on, volume 1, pages 41 –44 vol.1.
- [Cheung et al., 2000] Cheung, G., Kanade, T., Bouguet, J.-Y., and Holler, M. (2000). A real time system for robust 3d voxel reconstruction of human motions. *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference* on, 2:714–720 vol.2.
- [Cignoni et al., 2008] Cignoni, P., Corsini, M., and Ranzuglia, G. (2008). Meshlab: an open-source 3D mesh processing system. *ERCIM*, pages 47–48.
- [Cignoni et al., 1998] Cignoni, P., Rocchini, C., and Scopigno, R. (1998). Metro: Measuring error on simplified surfaces. *Computer Graphics Forum*, 17:167–174.
- [Courchay et al., 2009] Courchay, J., Pons, J.-P., Monasse, P., and Keriven, R. (2009). Dense and accurate spatio-temporal multi-view stereovision. In *Proc. Asian Conf. on Computer Vision*.
- [de Aguiar et al., 2008] de Aguiar, E., Stoll, C., Theobalt, C., Ahmed, N., Seidel, H.-P., and Thrun, S. (2008). Performance capture from sparse multi-view video. In *Proc. SIGGRAPH*, pages 1–10, New York, NY, USA. ACM.
- [de Aguiar et al., 2007] de Aguiar, E., Theobalt, C., Stoll, C., and Seidel, H.-P. (2007). Marker-less deformable mesh tracking for human shape and motion capture. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1–8.
- [Donoho, 2006] Donoho, D. (2006). Compressed sensing. Information Theory, IEEE Transactions on, 52(4):1289 –1306.

- [Edelsbrunner and Mücke, 1994] Edelsbrunner, H. and Mücke, E. P. (1994). Threedimensional alpha shapes. ACM Trans. Graph., 13:43–72.
- [Erol et al., 2005] Erol, A., Bebis, G., Boyle, R. D., and Nicolescu, M. (2005). Visual hull construction using adaptive sampling. Application of Computer Vision, 2005. WACV/MOTIONS '05 Volume 1. Seventh IEEE Workshops on, 1:234–241.
- [Esteban and Schmitt, 2003] Esteban, C. and Schmitt, F. (2003). Silhouette and stereo fusion for 3d object modeling. In 3-D Digital Imaging and Modeling, 2003. 3DIM 2003. Proceedings. Fourth International Conference on, pages 46–53.
- [Faugeras and Keriven, 2002] Faugeras, O. and Keriven, R. (2002). Variational principles, surface evolution, pde's, level set methods and the stereo problem. In *Biomedical Imaging*, 2002. 5th IEEE EMBS International Summer School on, pages 83 pp.–.
- [Floater and Hormann, 2001] Floater, M. and Hormann, K. (2001). Parameterization of triangulations and unorganized points. In *Principles of Multiresolution in Geometric Modelling*, pages 127–154. Springer.
- [Franco and Boyer, 2003] Franco, J.-S. and Boyer, E. (2003). Exact polyhedral visual hulls. In *British Machine Vision Conference (BMVC'03)*, volume I, pages 329–338.
- [Franco and Boyer, 2005] Franco, J. S. and Boyer, E. (2005). Fusion of multiview silhouette cues using a space occupancy grid. In *Proc. IEEE Int. Conf. on Computer Vision*, volume 2, pages 1747–1753.
- [Franco and Boyer, 2009] Franco, J. S. and Boyer, E. (2009). Efficient polyhedral modeling from silhouettes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 31:414–427.
- [Franco et al., 2006] Franco, J.-S., Lapierre, M., and Boyer, E. (2006). Visual shapes of silhouette sets. In 3D Processing, Visualization and Transmission 2006 (3DPVT'06), Chapel Hill, USA.
- [Friedman et al., 1977] Friedman, J. H., Bentley, J. L., and Finkel, R. A. (1977). An algorithm for finding best matches in logarithmic expected time. ACM Trans. on Mathematical Software, 3(3):209–226.
- [Furukawa et al., 2009] Furukawa, Y., Curless, B., Seitz, S., and Szeliski, R. (2009). Manhattan-world stereo. In Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, pages 1422–1429.
- [Furukawa and Ponce, 2007] Furukawa, Y. and Ponce, J. (2007). Accurate, dense, and robust multi-view stereopsis. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1–8.

- [Furukawa and Ponce, 2008] Furukawa, Y. and Ponce, J. (2008). Dense 3D motion capture from synchronized video streams. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1–8.
- [Furukawa and Ponce, 2009a] Furukawa, Y. and Ponce, J. (2009a). Accurate, dense, and robust multi-view stereopsis. *IEEE Trans. on Pattern Analysis and Machine Intelligence.*
- [Furukawa and Ponce, 2009b] Furukawa, Y. and Ponce, J. (2009b). Dense 3d motion capture for human faces. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1674–1681.
- [Gallego and Pardàs, 2010] Gallego, J. and Pardàs, M. (2010). Enhanced bayesian foreground segmentation using brightness and color distortion region-based model for shadow removal. In *Proc. IEEE Int. Conf. on Image Processing (to appear)*.
- [Gallego et al., 2011] Gallego, J., Salvador, J., Casas, J. R., and Pardàs, M. (2011). Joint multi-view foreground segmentation and 3d reconstruction with tolerance loop. In *IEEE International Conference on Image Processing (ICIP), pending* acceptance.
- [Geelnard, 2010] Geelnard, M. (2010). OpenCTM, the Open Compressed Triangle Mesh file format. http://openctm.sourceforge.net.
- [Goesele et al., 2006] Goesele, M., Curless, B., and Seitz, S. M. (2006). Multiview stereo revisited. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 2402–2409, Washington, DC, USA. IEEE Computer Society.
- [Goldluecke and Magnor, 2004] Goldluecke, B. and Magnor, M. (2004). Space-time isosurface evolution for temporally coherent 3d reconstruction. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, volume 1, pages 350–355.
- [Gouraud, 1971] Gouraud, H. (1971). Continuous shading of curved surfaces. IEEE Transactions on Computers, 20:623–629.
- [Guan et al., 2008] Guan, L., Franco, J.-S., and Pollefeys, M. (2008). Multi-object shape estimation and tracking from silhouette cues. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1–8.
- [Guennebaud et al., 2004] Guennebaud, G., Barthe, L., and Paulin, M. (2004). Dynamic surfel set refinement for high-quality rendering. *Comput. Graph.*, 28:827– 838.
- [Guggenheimer, 1963] Guggenheimer, H. (1963). Differential Geometry. McGraw-Hill.

- [Haralick et al., 1973] Haralick, R. M., Dinstein, and Shanmugam, K. (1973). Textural features for image classification. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-3:610–621.
- [Hartley and Zisserman, 2000] Hartley, R. I. and Zisserman, A. (2000). Multiple View Geometry in Computer Vision. Cambridge University Press, ISBN: 0521623049.
- [Henrikson, 1999] Henrikson, J. (1999). Completeness and total boundedness of the hausdorff metric. MIT Undergraduate Journal of Mathematics, 1:69–80.
- [Hormann et al., 2007] Hormann, K., Lévy, B., and Sheffer, A. (2007). Mesh parameterization: Theory and practice. In *ACM SIGGRAPH Course Notes*.
- [Horprasert et al., 1999] Horprasert, T., Harwood, D., and Davis, L. S. (1999). A statistical approach for real-time robust background subtraction and shadow detection. In Proc. IEEE Int. Conf. on Computer Vision, pages 1–19.
- [Isidro and Sclaroff, 2003] Isidro, J. and Sclaroff, S. (2003). Stochastic refinement of the visual hull to satisfy photometric and silhouette consistency constraints. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 1335–1342 vol.2.
- [Jia et al., 2006] Jia, Y.-B., Mi, L., and Tian, J. (2006). Surface patch reconstruction via curve sampling. In *Robotics and Automation*, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on, pages 1371–1377.
- [Kauff et al., 2007] Kauff, P., Atzpadin, N., Fehn, C., Müller, M., Schreer, O., Smolic, A., and Tanger, R. (2007). Depth map creation and image-based rendering for advanced 3dtv services providing interoperability and scalability. *Image Communication*, 22:217–234.
- [Kazhdan et al., 2006] Kazhdan, M., Bolitho, M., and Hoppe, H. (2006). Poisson surface reconstruction. In *Eurographics Symposium on Geometry Processing*, pages 61–70.
- [Kettner, 1999] Kettner, L. (1999). Using generic programming for designing a data structure for polyhedral surfaces. *Computational Geometry Theory and Applications*, 13:65–90.
- [Klie et al., 2007] Klie, P., Okon, E., Stefanoski, N., and Ostermann, J. (2007). A framework for scene-flow driven creation of time-consistent dynamic 3d objects using mesh parametrizations. In *3DTV Conference*, pages 1–4.
- [Kolev and Cremers, 2008] Kolev, K. and Cremers, D. (2008). Integration of multiview stereo and silhouettes via convex functionals on convex domains. In Proc. European Conf. on Computer Vision, pages 752–765.

- [Kolluri, 2005] Kolluri, R. (2005). Provably good moving least squares. In ACM SIGGRAPH 2005 Courses, SIGGRAPH '05, New York, NY, USA. ACM.
- [Kolmogorov and Zabih, 2002] Kolmogorov, V. and Zabih, R. (2002). Multi-camera scene reconstruction via graph cuts. In ECCV '02: Proceedings of the 7th European Conference on Computer Vision-Part III, pages 82–96, London, UK. Springer-Verlag.
- [Kolmogorov and Zabih, 2004] Kolmogorov, V. and Zabih, R. (2004). What energy functions can be minimized via graph cuts? *IEEE Trans. on Pattern Analysis* and Machine Intelligence, 26:147–159.
- [Krishnamurthy and Levoy, 1996] Krishnamurthy, V. and Levoy, M. (1996). Fitting smooth surfaces to dense polygon meshes. In *Proceedings of SIGGRAPH 96*, pages 313–324.
- [Kung-Fu Girl, 2005] Kung-Fu Girl (2005). A synthetic test sequence for multiview reconstruction and rendering research. Independent Research Group 3: Graphics, Optics, Vision. Max-Planck-Institut Informatik. http://www.mpiinf.mpg.de/departments/irg3/kungfu/.
- [Kutulakos and Seitz, 1999] Kutulakos, K. N. and Seitz, S. M. (1999). A theory of shape by space carving. In Proc. IEEE Int. Conf. on Computer Vision, volume 1, pages 307–314.
- [Landabaso et al., 2006] Landabaso, J. L., Pardàs, M., and Casas, J. R. (2006). Reconstruction of 3D shapes considering inconsistent 2D silhouettes. In Proc. IEEE Int. Conf. on Image Processing, pages 2209–2212.
- [Laurentini, 1994] Laurentini, A. (1994). The visual hull concept for silhouettebased image understanding. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16:150–162.
- [Laurentini, 1995] Laurentini, A. (1995). How far 3d shapes can be understood from 2d silhouettes. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 17(2):188–195.
- [Lazebnik et al., 2007] Lazebnik, S., Furukawa, Y., and Ponce, J. (2007). Projective visual hulls. Int. J. Comput. Vision, 74(2):137–165.
- [Lipski et al., 2010] Lipski, C., Linz, C., Berger, K., Sellent, A., and Magnor, M. (2010). Virtual video camera: Image-based viewpoint navigation through space and time. *Computer Graphics Forum*, 29(8):2555–2568.
- [Lorensen and Cline, 1987] Lorensen, W. E. and Cline, H. E. (1987). Marching cubes: A high resolution 3D surface construction algorithm. In Proc. ACM SIGGRAPH Computer Graphics, volume 21, pages 163–169.

- [Lowe, 2004] Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vision, 60:91–110.
- [Maneewongvatana and Mount, 1999] Maneewongvatana, S. and Mount, D. M. (1999). It's okay to be skinny, if your friends are fat. In CGC Workshop on Computational Geometry, pages 84–89.
- [Matusik et al., 2001] Matusik, W., Buehler, C., and McMillan, L. (2001). Polyhedral visual hulls for real-time rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 115–126, London, UK. Springer-Verlag.
- [Meagher, 1982] Meagher, D. (1982). Geometric modeling using octree encoding. Computer Graphics Image Process, 19:129–147.
- [Miller et al., 2005] Miller, G., Hilton, A., and Starck, J. (2005). Interactive freeviewpoint video. In Proc. IEEE European Conf. on Visual Media Production, pages 52–61.
- [Mueller et al., 2009] Mueller, K., Smolic, A., Dix, K., Merkle, P., and Wiegand, T. (2009). Coding and intermediate view synthesis of multiview video plus depth. In *Image Processing (ICIP), 2009 16th IEEE International Conference on*, pages 741–744.
- [Mueller et al., 2004] Mueller, K., Smolic, A., Merkle, P., Kaspar, B., Eisert, P., and Wiegand, T. (2004). 3d reconstruction of natural scenes with view-adaptive multi-texturing. In 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. Proceedings. 2nd International Symposium on, pages 116–123.
- [Oztireli et al., 2009] Oztireli, C., Guennebaud, G., and Gross, M. (2009). Feature preserving point set surfaces based on non-linear kernel regression. *Computer Graphics Forum (Proceedings of Eurographics 2009)*, 28(2):493–501. Best Student Paper Award.
- [Pfister et al., 2000] Pfister, H., Zwicker, M., van Baar, J., and Gross, M. (2000). Surfels: surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '00, pages 335–342, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.
- [PMVS, 2010] PMVS (2010). Yasutaka Furukawa and Jean Ponce. Patch-based Multi-View Stereo Software, http://grail.cs.washington.edu/software/pmvs.
- [Preparata and Shamos, 1985] Preparata, F. P. and Shamos, M. I. (1985). Computational geometry: an introduction. Springer-Verlag New York, Inc., New York, NY, USA.

- [Reddy et al., 2008] Reddy, D., Sankaranarayanan, A. C., Cevher, V., and Chellappa, R. (2008). Compressed sensing for multi-view tracking and 3-d voxel reconstruction. In *In Proceedings of ICIP*.
- [Rusinkiewicz and Levoy, 2001] Rusinkiewicz, S. and Levoy, M. (2001). Efficient variants of the icp algorithm. In *International Conference on 3D Digital Imaging* and Modeling, pages 145–152.
- [Salvador and Casas, 2010a] Salvador, J. and Casas, J. (2010a). Joint estimation of shape and motion from silhouettes. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 4069–4072.
- [Salvador and Casas, 2010b] Salvador, J. and Casas, J. (2010b). Photo-consistent surfaces from a sparse set of viewpoints. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 4045–4048.
- [Salvador and Casas, 2011] Salvador, J. and Casas, J. R. (2011). Statistical surface sampling for multi-view reconstruction. In *IEEE International Conference on Computer Vision (ICCV), pending acceptance.*
- [Salvador et al., 2010] Salvador, J., Suau, X., and Casas, J. R. (2010). From silhouettes to 3d points to mesh: Towards free viewpoint video. In *Proceedings of* the 1st International Workshop on 3D Video Processing, 3DVP '10, pages 19–24, New York, NY, USA. ACM.
- [Sardarescu et al., 2008] Sardarescu, M., Paragios, N., Komodakis, N., Raymond, R., Hernigou, P., and Rahmouni, A. (2008). Knee reconstruction through efficient linear programming. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI* 2008. 5th IEEE International Symposium on, pages 664–667.
- [Segal et al., 1992] Segal, M., Korobkin, C., van Widenfelt, R., Foran, J., and Haeberli, P. (1992). Fast shadows and lighting effects using texture mapping. SIG-GRAPH Comput. Graph., 26:249–252.
- [Seitz and Dyer, 1997] Seitz, S. and Dyer, C. (1997). Photorealistic scene reconstruction by voxel coloring. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pages 1067–1073.
- [Seitz et al., 2006] Seitz, S. M., Curless, B., Diebel, J., Scharstein, D., and Szeliski, R. (2006). A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, volume 1, pages 519–528.
- [Shin and Tjahjadi, 2008] Shin, D. and Tjahjadi, T. (2008). Local hull-based surface construction of volumetric data from silhouettes. *Image Processing, IEEE Transactions on*, 17(8):1251–1260.

- [Sinha and Pollefeys, 2005] Sinha, S. and Pollefeys, M. (2005). Multi-view reconstruction using photo-consistency and exact silhouette constraints: a maximumflow formulation. *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on,* 1:349–356 Vol. 1.
- [Starck and Hilton, 2005] Starck, J. and Hilton, A. (2005). Spherical matching for temporal correspondence of non-rigid surfaces. In Proc. IEEE Conf. on Computer Vision, volume 2, pages 1387–1394.
- [Stauffer and Grimson, 1999] Stauffer, C. and Grimson, W. E. L. (1999). Adaptive background mixture models for real-time tracking. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, volume 2, pages 252–259.
- [Suau et al., 2010] Suau, X., Casas, J. R., and Ruiz-Hidalgo, J. (2010). Surface reconstruction by restricted and oriented propagation. In *Proc. IEEE Int. Conf.* on Image Processing (to appear).
- [Szeliski, 1993] Szeliski, R. (1993). Rapid octree construction from image sequences. CVGIP: Image Underst., 58(1):23–32.
- [The Stanford 3D Scanning Repository, 2010] The Stanford 3D Scanning Repository (2010). Stanford Computer Graphics Laboratory. http://graphics.stanford.edu/data/3Dscanrep/.
- [Turk and O'Brien, 2005] Turk, G. and O'Brien, J. F. (2005). Modelling with implicit surfaces that interpolate. In ACM SIGGRAPH 2005 Courses, SIGGRAPH '05, New York, NY, USA. ACM.
- [Vedula et al., 1999] Vedula, S., Baker, S., Collins, R., Kanade, T., and Rander, P. (1999). Three-dimensional scene flow. In *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*, page 722, Washington, DC, USA. IEEE Computer Society.
- [Vedula et al., 2002] Vedula, S., Baker, S., and Kanade, T. (2002). Spatio-temporal view interpolation. In EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering, pages 65–76, Aire-la-Ville, Switzerland. Eurographics Association.
- [Vedula et al., 2005] Vedula, S., Baker, S., and Kanade, T. (2005). Image-based spatio-temporal modeling and view interpolation of dynamic events. ACM Trans. on Graphics, 24(2):240–261.
- [Vision, 2010] Vision (2010). CENIT Project of the Spanish Ministry of Industry, Tourism and Commerce. https://www.cenit-vision.org/.
- [Waizenegger et al., 2011] Waizenegger, W., Feldmann, I., and Schreer, O. (2011). Real-time patch sweeping for high-quality depth estimation in 3d videoconferencing applications. In SPIE 2011 Conference on Real-Time Image and Video Processing.

[Xu et al., 2005] Xu, L.-Q., Landabaso, J. L., and Pardàs, M. (2005). Shadow removal with blob-based morphological reconstruction for error correction. In Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Philadelphia, PA, USA. IEEE Computer Society.