


Frédéric Cao
Jean-Michel Morel
Frédéric Sur

José-Luis Lisani
Pablo Musé

A Theory of Shape Identification

1948



Springer

Frédéric Cao · José-Luis Lisani
Jean-Michel Morel · Pablo Musé
Frédéric Sur

A Theory of Shape Identification

Frédéric Cao
DxO Labs
3 rue Nationale
92100 Boulogne Billancourt, France
fcao@dxo.com

Pablo Musé
Instituto de Ingeniería Eléctrica
Facultad de Ingeniería
Julio Herrera y Reissig 565
11300 Montevideo, Uruguay
pmuse@fing.edu.uy

José-Luis Lisani
Dep. Matemàtiques i Informàtica
University Balearic Islands
ctra. Valldemossa km.7,5
07122 Palma de Mallorca, Balears
Spain
joseluis.lisani@uib.es

Frédéric Sur
Loria Bat. C - projet Magrit
Campus Scientifique - BP 239
54506 Vandoeuvre-lès-Nancy Cédex
France
sur@loria.fr

Jean-Michel Morel
CMLA, Ecole Normale
Supérieure de Cachan
61 av. du Président Wilson
94235 Cachan Cédex, France
morel@cmla.ens-cachan.fr

ISBN 978-3-540-68480-0 e-ISBN 978-3-540-68481-7
DOI 10.1007/978-3-540-68481-7

Lecture Notes in Mathematics ISSN print edition: 0075-8434
ISSN electronic edition: 1617-9692

Library of Congress Control Number: 2008927359

Mathematics Subject Classification (2000): 62C05, 62G10, 62G32, 62H11, 62H15, 62H30, 62H35, 68T10, 68T45, 68U10, 91E30, 94A08, 94A13, 94B70

© 2008 Springer-Verlag Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover design: WMX Design Bender

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

Preface

Recent years have seen dramatic progress in shape recognition algorithms applied to ever-growing image databases. They have been applied to image stitching, stereo vision, image retrieval, image mosaics, solid object recognition and video and web shape retrieval. More fundamentally, the ability of humans and animals to detect and recognize shapes is one of the enigmas of perception. Digital images and computer vision methods open new ways to address this enigma.

Given a dictionary of digitized shapes and a previously unobserved digital image, the aim of shape recognition algorithms is to know whether some of the shapes in the dictionary are present in the image. This book describes a complete method that starts from a query image and an image database and yields a list of the images in the database containing the query shapes.

Technically speaking there are two main issues. The first is extracting invariant shape descriptors from digital images. Indeed, a shape can be seen from various angles and distances and in various lights. A shape can even be partially occluded by other shapes and still be identifiable. Because the extraction step is so crucial, three acknowledged shape descriptors, SIFT (Scale-Invariant Feature Transform), MSER (Maximally Stable Extremal Regions) and LLD (Level Line Descriptor) will be introduced.¹

The second issue is deciding whether two shape descriptors are identifiable as the *same shape* or not. This decision process will derive from a unique paradigm, called the Helmholtz principle. For each decision a background model is introduced. Then one decides whether an event of interest (such as the presence of a shape in the image) has occurred if it has a very low probability of occurring by chance in the background model. Thus from the statistical viewpoint shape identification goes back to *multiple hypothesis testing*.

A shape descriptor is recognized if it is not likely to appear by chance in the background model. At a higher complexity level, a group of shape descriptors is recognized if its spatial arrangement could not occur just by chance. These two decisions

¹ In a recent review paper on affine invariant recognition written by a pool of experts, SIFT and MSER were actually acclaimed as the best shape descriptors [122].

rely on simple stochastic geometry and eventually compute a false alarm number for each shape descriptor. The lower this number, the more secure the identification. In that way most familiar simple shapes or images can be reliably identified. Many realistic experiments show false alarm rates ranging from 10^{-5} to less than 10^{-300} .

All in all these lecture notes prove that many shapes can indeed be identified. For these shapes one needs no *a priori* model and no training, just one sample of the shape and what statisticians call a *background model*, or a *null model*. In the case of shape recognition, the term background is to be taken to the letter. By the Helmholtz principle a shape is conspicuous if and only if it cannot be generated by the image background on which it is perceived. The background model is therefore easily learnt from the image database itself.

The above description should not be taken to suggest that the shape recognition problem is solved. The methods described only apply to solid shapes and not to deformable shapes. They only deal with individual shapes and images such as logos or paintings, and not with wide classes of objects such as all humans, all cats or all cars. This latter problem is known as *categorization* and is still widely open to research.

The authors are indebted to their collaborators for many important comments and corrections, particularly to Andrés Almansa, Yann Gousseau and Guoshen Yu. David Mumford and another anonymous referee made valuable comments which reshaped the book. All experiments were done using the public software MegaWave (authors: Jacques Froment and Lionel Moisan). The SIFT method is also public and downloadable.

The present theory was mainly developed at the Centre de Mathématiques et Leurs Applications, at ENS Cachan, at the Universitat de les Illes Balears and at IRISA, Rennes. It was partially financed for the past eight years by the Centre National d'Etudes Spatiales, the Centre National de la Recherche Scientifique, the Office of Naval research (grant N00014-97-1-0839) and the Ministère de la Recherche (project ISII-RNRT), and the Ministerio de Educación y Cultura (project MTM2005-08567). Special thanks to Bernard Rougé and Wen Masters for their great interest and support. We are indebted to Nick Chriss for numerous stylistic corrections.

Frédéric Cao
José Luis Lisani
Jean-Michel Morel
Pablo Musé
Frédéric Sur

Contents

- 1 Introduction** 1
 - 1.1 A Single Principle 1
 - 1.2 Shape Invariants and Consequences 4
 - 1.2.1 Shape Distortions 4
 - 1.3 General Overview 9
 - 1.3.1 Extraction of Shape Elements 9
 - 1.3.2 Shape Element Encoding 11
 - 1.3.3 Recognition of Shape Elements 11
 - 1.3.4 Grouping 12
 - 1.3.5 Algorithm Synopsis 12

Part I Extracting Image boundaries

- 2 Extracting Meaningful Curves from Images** 15
 - 2.1 The Level Lines Tree, or Topographic Map 15
 - 2.2 Matas *et al.* Maximally Stable Extremal Regions (MSER) 17
 - 2.3 Meaningful Boundaries 18
 - 2.3.1 Contrasted Boundaries 18
 - 2.3.2 Maximal Boundaries 19
 - 2.4 A Mathematical Justification of Meaningful Contrasted Boundaries 21
 - 2.4.1 Interpretation of the Number of False Alarms 21
 - 2.5 Multiscale Meaningful Boundaries 26
 - 2.6 Adapting Boundary Detection to Local Contrast 27
 - 2.6.1 Local Contrast 29
 - 2.6.2 Experiments on Locally Contrasted Boundaries 30
 - 2.7 Bibliographic Notes 32
 - 2.7.1 Edge Detection 32
 - 2.7.2 Meaningful Boundaries vs. Haralick's Detector 33
 - 2.7.3 Level Lines and Shapes 34
 - 2.7.4 Tree of Shapes, FLST, and MSER 35
 - 2.7.5 Extracting Shapes from Images 35

Part II Level Line Invariant Descriptors

3	Robust Shape Directions	41
3.1	Flat Parts of Level Lines	41
3.1.1	Flat Parts Detection Algorithm	42
3.1.2	Reduction to a Parameterless Method	43
3.1.3	The Algorithm	44
3.1.4	Some Properties of the Detected Flat Parts	44
3.2	Experiments	45
3.2.1	Experimental Validation of the Flat Part Algorithm	45
3.2.2	Flat Parts Correspond to Salient Features	47
3.3	Curve Smoothing and the Reduction of the Number of Bitangent Lines	49
3.4	Bibliographic Notes	54
3.4.1	Detecting Flat Parts in Curves	54
3.4.2	Scale-Space and Curve Smoothing	57
4	Invariant Level Line Encoding	61
4.1	Global Normalization and Encoding	61
4.1.1	Global Affine Normalization	61
4.1.2	Application to the MSER Normalization Method	64
4.1.3	Geometric Global Normalization Methods	65
4.2	Semi-Local Normalization and Encoding	66
4.2.1	Similarity Invariant Normalization and Encoding Algorithm	67
4.2.2	Affine Invariant Normalization and Encoding Algorithm	70
4.2.3	Typical Number of LLDs in Images	71
4.3	Bibliographic Notes	73
4.3.1	Geometric Invariance and Shape Recognition	73
4.3.2	Global Features and Global Normalization	74
4.3.3	Local and Semi-Local Features	75

Part III Recognizing Level Lines

5	A <i>Contrario</i> Decision: the LLD Method	81
5.1	A <i>Contrario</i> Models	81
5.1.1	Shape Model or Background Model?	81
5.1.2	Detection Terminology	83
5.2	The Background Model	85
5.2.1	Deriving Statistically Independent Features from Level Lines	87
5.3	Testing the Background Model	89
5.4	Bibliographic Notes	91
5.4.1	Shape Distances	91
5.4.2	A <i>Contrario</i> Methods	92

6	Meaningful Matches: Experiments on LLD and MSER	93
6.1	Semi-Local Meaningful Matches	93
6.1.1	A Toy Example	94
6.1.2	Perspective Distortion	98
6.1.3	A More Difficult Problem	101
6.1.4	Slightly Meaningful Matches between Unrelated Images	104
6.1.5	Camera Blur	105
6.2	Recognition Relative to Context	113
6.3	Testing A <i>Contrario</i> MSER (Global Normalization)	116
6.3.1	Global Affine Invariant Recognition. A Toy Example	116
6.3.2	Comparing Similarity and Affine Invariant Global Recognition Methods	116
6.3.3	Global Matches of Non-Locally Encoded LLDs	118

Part IV Grouping Shape Elements

7	Hierarchical Clustering and Validity Assessment	129
7.1	Clustering Analysis	129
7.2	A <i>Contrario</i> Cluster Validity	131
7.2.1	The Background Model	131
7.2.2	Meaningful Groups	132
7.3	Optimal Merging Criteria	136
7.3.1	Local Merging Criterion	136
7.4	Computational Issues	140
7.4.1	Choosing Test Regions	140
7.4.2	Indivisibility and Maximality	142
7.5	Experimental Validation: Object Grouping Based on Elementary Features	143
7.5.1	Segments	144
7.5.2	DNA Image	146
7.6	Bibliographic Notes	148
8	Grouping Spatially Coherent Meaningful Matches	151
8.1	Why Spatial Coherence Detection?	151
8.2	Describing Transformations	153
8.2.1	The Similarity Case	153
8.2.2	The Affine Transformation Case	154
8.3	Meaningful Transformation Clusters	155
8.3.1	Measuring Transformation Dissimilarity	155
8.3.2	Background Model: the Similarity Case	157
8.4	Experiments	158
8.5	Bibliographic Notes	161

9	Experimental Results	167
9.1	Visualizing the Results	167
9.2	Experiments	168
9.2.1	Multiple Occurrences of a Logo	168
9.2.2	Valbonne Church	173
9.2.3	Tramway	175
9.3	Occlusions	177
9.4	Stroboscopic Effect	179

Part V The SIFT Method

10	The SIFT Method	185
10.1	A Short Guide to SIFT Encoding	185
10.1.1	Scale-Space Extrema	186
10.1.2	Accurate Key Point Detection	187
10.1.3	Orientation Assignment	188
10.1.4	Local Image Descriptor	189
10.1.5	SIFT Descriptor Matching	189
10.2	Shape Element Stability versus SIFT Stability	190
10.2.1	An Experimental Protocol	190
10.2.2	Experiments	191
10.2.3	Some Conclusions Concerning Stability	195
10.3	SIFT Descriptors Matching versus LLD A <i>Contrario</i> Matching	196
10.3.1	Measuring Matching Performance	198
10.3.2	Experiments	201
10.4	Conclusion	207
10.5	Bibliographic Notes	207
10.5.1	Interest Points of an Image	207
10.5.2	Local Descriptors	207
10.5.3	Matching and Grouping	208
11	Securing SIFT with A <i>Contrario</i> Techniques	209
11.1	A <i>Contrario</i> Clustering of SIFT Matches	209
11.2	Using a Background Model for SIFT	210
11.3	Meaningful SIFT Matching	214
11.3.1	Normalization	214
11.3.2	Matching	215
11.3.3	Choosing Sample Points	218
11.4	The Detection Algorithm	219
11.4.1	Experiments: Securing SIFT Detections	220
11.5	Bibliographic Notes	224

A Keynotes 225

 A.1 Cluster Analysis Reader’s Digest 225

 A.1.1 Partitional Clustering Methods 225

 A.1.2 Iterative Methods for Partitional Clustering 227

 A.1.3 Hierarchical Clustering Methods 228

 A.1.4 Cluster Validity Analysis and Stopping Rules 231

 A.2 Three classical methods for object detection based on spatial
 coherence 235

 A.2.1 The Generalized Hough Transform 235

 A.2.2 Geometric Hashing 236

 A.2.3 A RANSAC-based Approach 237

 A.3 On the Negative Association of Multinomial Distributions 239

B Algorithms 243

 B.1 LLD Method Summary 243

 B.2 Improved MSER Method Summary 244

 B.3 Improved SIFT Method Summary 245

References 247

Index 255

Chapter 1

Introduction

1.1 A Single Principle

Digital images became an object of scientific interest in the seventies of the last century. The emerging science dealing with digital images is called *Computer Vision*. Computer Vision aims to give wherever possible a mathematical definition of visual perception. It can be therefore viewed in the realm of perception theory. Images are, however, a much more affordable object than percepts. Indeed, digital images are sampled real or vectorial functions defined on a part of the plane, usually a rectangle. They are accessible to all kinds of numerical, geometric, and statistical measurements. In addition, the results of artificial perception algorithms can be confronted to human perception. This confrontation is both advantageous and dangerous. Experimental results may easily be misinterpreted during visual inspection. The results look disappointing when matched with our perception. Obvious objects are often very hard to find in digital images by an algorithm.

In a recent book by Desolneux et al. [54], a general mathematical principle, the so called Helmholtz principle, was extensively explored as a way to define all visual percepts (gestalts) as large deviations from randomness. According to the main thesis of these authors one can compute detection thresholds deciding whether a given geometric structure is present or not in each digital image. Several applications of this principle have been developed by these authors and others for the detection of alignments [50], boundaries [51, 35], clusters [53, 33], smooth curves and good continuations [30, 31], vanishing points [2] and robust point matches through epipolar constraint [128].

These works make extensive use of a computed function, the so called *number of false alarms* (NFA) of a perception. The NFA of a perception is the expected number of times this perception could have arisen just by chance in the background. An observed configuration in an image can be numerically defined as a perception if and only if its NFA is smaller than 1. Experimental evidence has confirmed that the NFA of many human percepts of geometric figures is actually much smaller than 1,

typically less than 10^{-n} where n ranges from 10 to 100 and more. See [52] and the textbook [54].

Thus theory and experiments give a mathematical and experimental basis to the existence of sure percepts. Their existence had been stated for a long time by phenomenology, in particular the Gestalt theory [95], but without quantitative evidence.

The idea that perceptions are objects unlikely to form just by chance in a background goes back to Helmholtz [83]. This principle could not be tested until images became digital and therefore accessible to computational experiments. Before the above-mentioned works, there had been several attempts to define percepts as exceptional events. Stewart [168] proposed to detect planes in a cloud of points by what he called the MINPRAN method. He computed the probability that MINPRAN will “hallucinate a fit where there is none”. This probability was computed under the *a contrario* assumption that the points were randomly distributed. Lowe [112] proposed a detection framework based on the computation of accidental occurrence.

In other words, we can shift our attention from finding properties with high prior expectations to those that are sufficiently constrained to be detectable among a realistic distribution of accidentals.[...] Even when we do not know the ultimate interpretation for some grouping and therefore its particular a priori expectation, we can judge it to be significant based on the non-accidentalness criteria.

Later Barlow [17] interpreted the perceptual goals of the neocortex as a search for *suspicious coincidences*. In the same spirit, Grimson and Huttenlocher [75, 76] proposed to compute shape recognition thresholds from a null model viewed as *the conspiracy of random*.

There is a common sense objection to applying the Helmholtz principle to shape recognition. Watching the sky, one often sees castles, cats and dogs in the clouds. Humans have a high capacity for hallucinating familiar generic shapes such as faces in rich visual environments. Thus the Helmholtz principle is not suited for all sorts of shapes.

The situation is, however, quite different regarding more specific iconic shapes such as letters, logos and in general solid shapes. One sees faces in the sky, but certainly not this or that particular face. It is to be expected that any complex enough solid shape will be recognizable in the Helmholtz sense: no random arrangement would be able to reproduce it accurately.

The mathematics to prove this are quite simple. Let S and S' be two shapes observed in two different images and which happen to be similar. Denote their (small) Hausdorff distance after registration by $\delta = d(S, S')$. Assume we know enough of the background model to compute the probability $\Pr(S, \delta) = \Pr(d(S, \Sigma) \leq \delta)$ that some shape in the background, Σ be as similar to S as S' is. If this probability is very small one can deduce that S' does not look like S just by chance. Then S and S' will be *identified as the same shape*.

Digital images contain thousands of significant *shape elements* that constitute their shape contents. (Several kinds of *shape elements*, or *descriptors* will be considered in this book.) Controlling the number of wrong matches involves the computation of the probability of a casual match with the background. This probability

should be very small. But also the number of tests, which can be huge, must be under control.

Definition 1. Let \mathcal{I} and \mathcal{I}' be two images and N, N' the number of shape elements in each. Let \mathcal{S} and \mathcal{S}' be two shape elements extracted from \mathcal{I} and \mathcal{I}' respectively, lying at distance δ . We call number of false alarms of the match between \mathcal{S} and \mathcal{S}' the number

$$\text{NFA}(\mathcal{S}, \mathcal{S}') = N \cdot N' \cdot \Pr(d(\mathcal{S}, \Sigma) \leq \delta).$$

If $\text{NFA}(\mathcal{S}, \mathcal{S}')$ is much smaller than 1, one deduces that \mathcal{S}' could not look like \mathcal{S} just by chance and concludes that \mathcal{S} and \mathcal{S}' have *the same shape*.

There is an important phenomenological consequence: Shapes can be defined without learning, that is without empirical knowledge. By definition a shape is any part of an image which has been identified (in the sense of low NFA) at least once in another image.

From an empirical point of view, there are two kinds of shapes. First, any solid physical object can be photographed under many views and illuminations. If, by using the above definition, two snapshots of the same physical object happen to contain recognizable shape elements, one may say that the object itself is identifiable. These shape elements will constitute the object signature.

Second, humans build all kinds of standardized objects. Likewise, two different standard objects can be identified if they stem from the same industrial process. This also applies to the very numerous iconic planar shapes generated by human visual communication, in particular characters and logos. In the experimental parts of this book, we shall study the identifiability of several such iconic shapes: the lower the NFAs they generate at a given Hausdorff distance, the more recognizable they are.

As a consequence of the present study, one can define solid shapes as equivalence classes of recognized pairs without reference to empirical knowledge or *ground truth*. Thus, one should demonstrate the existence of, say, the Coca-Cola logo just by the fact that a certain group of shape elements appears in several images with very low NFA for all pairwise comparisons. Experiments will compare several snapshots of the same painting or poster, various images extracted from the same movie, or various logos of the same firm. The aim in all cases is to single out and group in clusters all shape elements common to both images. Conversely, the same method gives a negative answer when two images have no shape in common. In that case the NFA is above 1, which means that the shape is likely to occur casually in the background.

From the mathematical and numerical point of view, the main challenge in the whole study is the accurate computation of numbers of false alarms (NFA). This requires the computation of very small probabilities. Small probabilities cannot be directly measured from a shape database as frequencies. Thus, a probabilistic model of the set of all possible shapes should be built. Such a realistic experimental *background model* should be made of a large and representative set of all kinds of digital images. Unfortunately, there is no available probabilistic model for a large set of

images. It is as hopeless as building a global model of the world. Even if such a model were available, one would still face the challenge of computing accurately the probability of very rare events in this world model.

Fortunately enough it is possible to overcome or rather to circumvent these two obstacles. The only information needed is the probability for a background shape to be very close to a given query shape. By a geometric independence argument, this probability will be made into a product of much larger probabilities. These probabilities instead become observable as frequencies in a small image database.

1.2 Shape Invariants and Consequences

1.2.1 Shape Distortions

In order to find the shape invariance classes, it suffices to give a rough typology of the transformations that affect images but not our recognition of the shapes they contain. Following Lisani *et al.* [109], the main classes of perturbations which do not affect recognition are:

1. **Changes in the color and luminance scales (contrast changes).** According to Gestaltists Attneave [13] and Wertheimer [179], shape perception is independent of the gray level scale or of measured colors.

The concentration of information in contours is illustrated by the remarkable similar appearance of objects alike in contour and different otherwise. The “same” triangle, for example, may be either white on black or green on white. Even more impressive is the familiar fact that an artist’s sketch, in which lines are substituted for sharp color gradients, may constitute a readily identifiable representation of a person or thing. Attneave, 1954.

I stand at the window and see a house, trees, sky. Theoretically I might say there were 327 brightnesses and nuances of color. Do I have “327”? No. I have sky, house, and trees. It is impossible to achieve “327” as such. And yet even though such droll calculation were possible and implied, say, for the house 120, the trees 90, the sky 117 – I should at least have this arrangement and division of the total, and not, say, 127 and 100 and 100; or 150 and 177. Wertheimer, 1923.

Refer to Fig. 1.1 designed by E. H. Adelson for a striking illustration of illumination invariance.

2. **Occlusions and background modification.** Shape recognition can also be performed in spite of occlusion and varying background, as shown in Fig. 1.2. The phenomenology of occlusion was thoroughly studied by Kanizsa [95] who argues that occlusion is always present in every day vision: most objects are partially hidden by others. Human perception must therefore be able to recognize partial shapes. Conversely, if a shape occludes a background, its recognition is invariant to changes in the background. This independence of shape




Fig. 1.1 Contrast change invariance. In the left hand image, the A and B squares have exactly the same gray level. This incredible fact is easily checked in the right hand image where A and B are linked by two rectangles with the same gray level. This experiment by E.H. Adelson illustrates the unreliability of brightness perception and the invariance of shape recognition with respect to illumination changes. (Courtesy E.H. Adelson, http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html)

recognition from its background is known in perception psychology as the *figure-background problem* (Rubin [153]). The figure-background problem is part of the occlusion problem. A shape is superimposed on a background, which can be made of various objects. How can the shape be singled out from that clutter? This poses a dilemma. Extract the shape and then recognize it or extract it *because* it has been recognized?




Fig. 1.2 Left: According to Kanizsa and his school, shapes can be recognized even when they undergo several occlusions. Our perception is trained to recognize shapes which are only seeable in part. Here the occluded cross can be easily recovered. Right: the *figure-background problem*. Our perception is adapted to recover a figure on the foreground, independently from the background

3. **Classical noise and blur**, inherent to any perception task and to any image generated according to Shannon's theory.
4. **Geometrical distortions or deformations**. Perspective is deeply incorporated in human perception. Humans can recognize objects and shapes under perspective distortion as long as perspective is not too strong. Recognition is also invariant to elastic deformations, always within some limits.

The previous four invariant properties fix requirements a good image representation should comply with. It will be necessary to formulate a mathematical model for each of them and to derive a well adapted image representation.

1. a. **The local contrast invariance requirement**. A digital image is usually defined as a function $u(x)$, where $u(x)$ represents the gray level or luminance at x . The first task is to extract from the image topological information independent from the varying and unknown contrast change function of the optical or biological apparatus. One can model such a contrast change function as any continuous increasing function g from \mathbb{R}^+ to \mathbb{R}^+ . The real datum corresponding to the observed u could be as well any image $g(u)$. This simple argument can lead to select the level sets of the image [161] or its set of level lines as a complete contrast invariant image description [37]. If u is of class say C^1 , then the level lines are the connected components of $u^{-1}(\lambda)$, which are C^1 curves for almost every $\lambda \in \mathbb{R}$. This is the choice adopted by the LLD (level line descriptor) and MSER (maximally stable extremal regions) methods. Another way to handle the contrast invariance requirement is to encode only the direction of the gradient of u , $\frac{Du}{|Du|}$ and not the gradient Du itself. Indeed, the direction of gradient is normal to the level line and is not altered by any increasing contrast change. This is the way adopted by the SIFT method to cope with contrast changes.
- b. **The concentration of information requirement**. Somewhat in contradiction to this contrast invariance principle, the Gestaltist Attneave [13] asserted that "*Information is concentrated along contours (i.e., regions where color changes abruptly)*". Indeed not all the level lines are needed to have a complete shape description. Most of them are due to noise or to tiny illumination changes. Thus, it makes sense to select only the most contrasted level lines. That is to say, those along which the gradient of u is large enough. Such a selection is not invariant to all contrast changes, since it explicitly uses the gradient value. However, it is still invariant to affine global contrast changes. Figure 1.3 shows an example of such level lines selection. The selection of the most contrasted level lines will be the subject of Chap. 2. It will be applied to the LLD and the MSER methods. The SIFT methods actually weights its gradient orientation histograms by the gradient magnitude (see Chap. 10).
2. **The occlusion and figure-background requirements**. Even the best adapted choice of level lines is not totally suited to describing image parts. Indeed, when a shape A partially occludes a shape B , the level lines of the resulting image




Fig. 1.3 Top left: original image, 83,759 level lines. Top right: meaningful boundaries (883 level lines). Bottom: reconstruction from the meaningful boundaries. Only 883 boundaries remain. The structure of the image is preserved and perceptual loss is very weak. The LLD and MSER methods use these boundaries for building up normalized shape descriptors. See Chap. 4

are a concatenation of pieces of the level lines belonging to A and to B . This is shown with a very simple example in Fig. 1.4. Even if a shape is not occluded, but simply occludes its own background, there may be no level line surrounding the whole shape, as shown in Fig. 1.5. These remarks show that whole level lines are too big and too sensitive to occlusion. In order to overcome this obstacle the general idea is to build shape recognition on shape elements as local as possible. The SIFT method takes small image patches. The LLD method splits level lines into small pieces.

3. **The smoothing requirement.** It is an easy experiment to check that shapes are easily recognized in images subject to noise. This means that shape information is not affected by noise. Noise introduces details which are too fine (in relation to the essential shape information) to be perceptually relevant in terms of recognition. Quoting Attneave (ibid., 1954):

It appears, then, that when some portion of the visual field contains a quantity of information grossly in excess of the observer's perceptual capacity, he treats those components of information which do not have redundant representation somewhat as a statistician treats "error variance", averaging out particulars and abstracting certain statistical homogeneities.