



**Proceedings fra den 13. Danske Konference i  
Mønstergenkendelse og Billedanalyse**

**Eds. Søren I. Olsen**

**Technical Report no. 2004/10  
ISSN: 0107-8283**

**DIKU**

University of Copenhagen • Universitetsparken 1  
DK-2100 Copenhagen • Denmark

## TABLE OF CONTENTS:

<i>Color Signal Processing</i> Professor Reiner Lenz, Linköping University	3
<i>Image Relighting: Getting the Sun to Set in an Image Taken at Noon</i> Claus B. Madsen, Rune Laursen	13
<i>Using Mixtures of Gaussians to Compare Approaches to Signal Separation</i> Kaare B. Petersen	21
<i>Stochastic differential equations in image warping</i> Bo Markussen	31
<i>Probabilistic Model-based Background Subtraction</i> J. Anderson, T. Prehn, V. Krüger, A. Elgammal	38
<i>Some Transitions of Extrama-Based Multi-Scale Singularity Trees</i> K. Somchaipeng, J. Sporning, S. Kreiborg, P. Johansen	48
<i>A Test Statistic in the Complex Wishart Distribution and its Application to Change Detection in Polarimetric SAR Data</i> Knut Conradsen, Allan A. Nielsen, Jesper Schou, Henning Skriver	56
<i>Using the LZW to Detect Primitives</i> Lars Reng	72
<i>Medical image sequence coding</i> Mo Wu, Søren Forchhammer	77
<i>Integrity Improvement in Classically Deformable Solids</i> Micky K. Christensen, Anders Fleron	87
<i>The Thin Shell Tetrahedral Mesh</i> Kenny Erleben, Henrik Dohlmann	94
<i>Talking Faces - a State Space Approach</i> Tue Lehn-Schiøler	103
<i>Live Interpretation of Conductors' Beat Patterns</i> Declan Murphy	111
<i>Testing for difference between two groups of functional neuroimaging experiments</i> Finn Å. Nielsen, Andrew C.N. Chen, Lars K. Hansen	121
<i>An image based system to automatically and objectively score the degree of redness and scaling in psoriasis lesions</i> David Delgado, Bjarne Ersbøll, Jens M. Carstensen	130
<i>Survey and Assessment of Advanced Feature Extraction Techniques and Tools for OE Applications</i> Mikael K. Sørensen, Michael S. Rasmussen, Henning Skriver, Peter Johansen, Arthur Pece, Jesper H. Thygesen	138
<i>Automatic change detection for validation of digital map databases</i> Brian P. Olsen, Thomas Knudsen	148
<i>Spatial-temporal Disambiguation of Multi-modal Descriptors</i> Norbert Krüger, Florentin Wörgötter	154

# Color Signal Processing

Reiner Lenz  
Department Science and Technology  
Linköping University  
SE-60174 Norrköping  
Sweden  
email:reile@itn.liu.se

## Abstract

*In this paper we describe some of our recent investigations in the field of color signal processing. We will describe the application of group theoretical methods and illustrate them with some results from low-level filter design, statistical properties of color signal spaces and the construction of invariants.*

## 1 Introduction

Everybody knows what color is. But as a scientific concept color is not at all well-defined. This is perhaps natural since "color" has been of interest in such different scientific disciplines as art, physics, chemistry, biology, mathematics, psychology, physiology and technology to name only a few. We are very well aware of this but here we choose to restrict us to a very narrow definition of what color is. In the first part we identify color simply with the measurement vectors that are produced by a (RGB or multi-channel) camera. In the rest of the paper we define color (or rather the color signal) as a function of wavelength that has non-negative function values. This signal processing approach completely ignores all perceptual aspects of the problem (color scientists would thus probably argue that this paper is not at all about color, since in their view color is in the brain of the observer).

Group theory on the other hand originates partly in the study of symmetries. Roughly one can say that symmetry is a property that is preserved under the application of a set of rules. A simple example: A circle is a geometric figure that is preserved under all 2-D rotations around the center of the circle. Or one could also say that the rotations are the transformations that preserve the circles. Group theoretical methods have had an enormous impact on theoretical physics and mathematics and the purpose of this paper is to show how they can be used to solve problems involving color. We will not go into technical details but try to give an overview of the wide range of color related problems that can be investigated with group theoretical tools.

We start in the first section with a description how the (geometrical) symmetries of the sampling grid and the (permutational) symmetries of the color channels can be used to design low-level filters (corresponding to receptive fields in human color vision). Then we show why spaces of color signals have a conical structure and we illustrate this with an investigation of databases of daylight spectra. The final topic is the application of the Lie-theory of differential equations to color problems. We will show that transformation groups define systems of differential equations in a natural way. Solving these differential equations can, for example, be used to construct invariants. Using the Lie-theory of differential equations gives a number of advantages. One of them is the availability of symbolic packages like Maple that can solve these problems completely automatically.

## 2 PCA without data - Harmonic Analysis of RGB images

Group theoretically defined transforms are important tools in many low-level digital signal processing applications. The most widely used examples are the DFT and the FFT but also wavelet transforms can be understood in terms of group

theory. The example we describe here is the transform based on the dihedral groups [14]. The case that is of greatest use in applications concerns RGB images on a square grid. Other important cases are RGB vectors on hexagonal grids and four-channel sensors on square or hexagonal grids. Potential applications are digital signal processors in digital cameras with varying sampling and spectral resolution properties.

Consider a square sampling grid of the plane. The transformations that map this grid into itself form a group, the dihedral group  $D(4)$ . It consists of all rotations with angle  $90k$  degrees ( $k=0,1,2,3$ ) and all reflections around the symmetry axis of a square. It has eight elements. The dihedral group  $D(6)$  consists of all transformations that preserve the hexagonal sampling grid. The group of all permutations of three elements is  $S(3)$  and consists of six elements. The direct product  $G = D(4) \times S(3)$  consists of all pairs  $g = (d, \pi)$  where  $d$  is an element in  $D(4)$  and  $\pi \in S(3)$ . This product forms a group that has 48 elements. The group elements  $g$  operate on RGB vectors defined on square grids by application of the dihedral part  $d$  on the grid and the permutation part  $\pi$  on the spectral bands. A typical operation would be a 90 degrees rotation combined with a flipping of the R and B channels. Since these transformations are linear we can describe them (given a basis) by a matrix. We will also use  $g$  to denote this matrix. The same symbol will thus be used for the group element, the transformation and the matrix. What is meant should be clear from the context.

The second concept of interest in our context is Principal Component Analysis as Least Mean Squared Error (LMSE) approximation. We start with a stochastic process of signal vectors  $s$  of dimension  $N$ . In our case this could be RGB vectors on an  $n \times n$  window collected in an  $N = 3n^2$ -dimensional signal vector  $s$ . Now change the basis in the signal space and project all signal vectors  $s$  to the first  $M$  basis vectors in the new system (with  $M \ll N$ ) thus creating a new signal  $\tilde{s}$ . Since the information on  $s$  contained in the last  $N - M$  components is lost by the projection this will result in an error for each signal. Selecting the new basis such that the mean approximation error is minimal (in the  $L^2$ -sense) is known as Principal Component Analysis (PCA). It can be shown that the basis to choose consists of the eigenvectors of the correlation matrix  $C = E(ss')$  of the original stochastic process belonging to the  $M$  largest eigenvalues ( $E(\cdot)$  denotes the expectation operator).

PCA requires thus, in general, two operations: (1) compute the correlation matrix  $C$  of the process and (2) compute eigenvectors of the correlation matrix  $C$ . In the following we will show under what conditions we can make general statements about the eigenvectors without actually computing the correlation matrix and the eigenvectors.

We start by selecting a subset of points on the given grid that is closed under the operations of the dihedral transformations. On each location point we have a measurement vector. Given a group element  $g = (d, \pi)$  we apply the geometric transformation  $d$  to the grid points and the permutation  $\pi$  to the vector components. The resulting operation is described by a permutation matrix  $g$  that has exactly one element of value one in each column and row. These permutation matrices have the property that the transpose is equal to the inverse:  $g' = g^{-1}$ . Now assume that each such transformed signal  $gs$  has the same probability as the original signal  $s$ . Then we get for the correlation matrix the property that:

$$C = E(ss') = E((gs)(gs)') = gE(ss')g' = gCg' = gE(ss')g^{-1} = gCg^{-1}$$

where the equality  $E(ss') = E((gs)(gs)')$  holds because the transformation  $g$  merely results in a reordering of the original set of signals.

As a result we find that each group element  $g$  delivers one constraint  $C = E(ss') = gE(ss')g^{-1} = gCg^{-1}$  for the correlation matrix. Such a matrix  $C$  is called a G-symmetric matrix. The properties of G-symmetric matrices are well-known from the theory of group representations and the main results are the following:

- There is **one** matrix  $M$  that block-diagonalizes all G-symmetric matrices  $C$
- The size of the blocks in the diagonal depends only on the group
- The matrix  $M$  can be generated automatically

The detailed description of these facts can be found in [14] and the references mentioned there.

## 2.1 Experimental results

The implementation of the block-diagonalization is similar to the case described in [14] for the dihedral groups. The similarity of the correlation matrix  $C$  and its block-diagonalization  $MCM'$  depends of course on the degree to which the stochastic process really is G-symmetric. To test the usefulness of the method we did the following experiment: From a large image database containing more than 600 000 images we selected 10000 random images and in each image randomly

Representation No.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Dimension	1	1	2	1	1	2	1	1	2	1	1	2	2	2	4
Block size	36	0	72	28	0	56	36	0	72	28	0	56	128	0	256

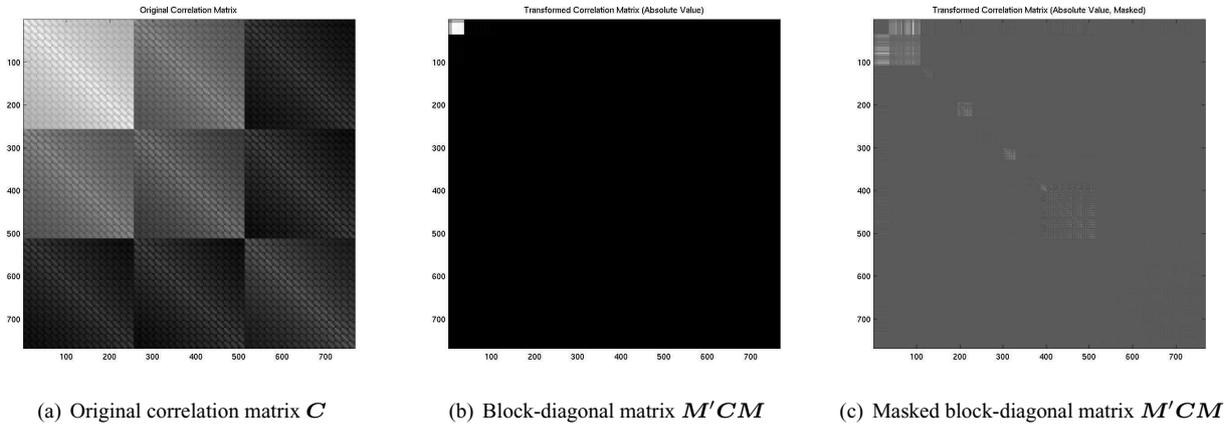
**Table 1. Block structure of transformed correlation matrix**

Eigenvector Number	Original	Block-Diagonal Approximation
1	0.915898	0.914900
2	0.945646	0.945642
3	0.954611	0.954600
4	0.963074	0.963057
5	0.967766	0.967745
6	0.970658	0.970605
7	0.973146	0.973117
8	0.975438	0.975403
9	0.976639	0.976601
10	0.977770	0.977733

**Table 2. Trace( $C$ ) compared with accumulated sum of eigenvalues**

2 patches of size 16 by 16 pixels. The signal vector had thus dimension  $16 \times 16 \times 3 = 768$ . Applying the general theory we find that in the ideal case the transformed correlation matrix  $M'CM'$  had the block structure listed in Table 1. Note that this result holds **for all correlation matrices with this type of symmetry of this size**.

Figure 2.1 shows the original correlation matrix, the absolute value of the block diagonal correlation matrix  $M'CM'$  and the absolute value of the block diagonal correlation matrix  $M'CM'$  where we have masked out the upper left corner to visualize the remaining entries.



One way to evaluate the loss of information introduced by the block-diagonalization is to compare the values of the eigenvalues of the original principal components and the block-diagonal approximations. In Table 2 the ratios of the trace and the accumulated sums of the eigenvalues of the original matrix and the block diagonal are listed: Further inspection shows that each of the blocks has a special transformation property under the elements of the transformation group leading to  $\Gamma$ -functions of a certain type. Some  $\Gamma$ -functions are thus invariant under all dihedral operations (spatially homogeneous) while others are related to edge-type  $\Gamma$ -functions etc..

Finally we want to point out that the block-diagonal lters allow more efficient implementations due to their symmetry properties and that given the domain they can be automatically derived using symbolic mathematical software like Maple.

### 3 PCA and the geometry of the space of color signals

One attempt to understand the properties of biological systems (and to build technological systems) relies on the hypothesis that successful systems are adapted to the structure of the space of signals they are analyzing (see [1, 2, 4, 18] for a few examples and [19] for a collection of articles on natural stimulus statistics.). If this is true, then it is of interest to investigate properties of signal spaces that are often analyzed by natural or artificial systems.

Here we start with the observation that in many interesting cases the signals of interest can only assume non-negative values. A typical example are illuminant spectra. Here the signals  $s$  are functions of the wavelength variable  $\lambda$  and  $s(\lambda)$  is the number of photons of wavelength  $\lambda$ . By definition  $s(\lambda) \geq 0, \forall \lambda$ . Another example from multi-channel color processing are reflectance spectra  $r$  where  $r(\lambda)$  describes the probability that a photon of wavelength  $\lambda$  will be reflected from a point. When an illumination spectrum  $s$  hits an object point with reflection spectrum  $r$  then (in the simplest model) the spectrum that is reflected from that scene point is given by  $c = s \cdot r$  and  $c$  is often called a color signal. The color signal, too, is by definition a non-negative function. The main result we will describe here is the proof that all components of the first eigenfunction of a stochastic process of spectra have the same sign. It can therefore be chosen to have only non-negative entries. This has been observed in many empirical investigations of databases of spectra but to our knowledge it has never been pointed out that this follows from the Perron-Frobenius theory of non-negative matrices and its generalization, the Krein-Rutman theorem.

We will then illustrate a few consequences of this fact for databases of daylight spectra. We will show that it follows that these spaces of daylight spectra have a conical geometry and that the group of Lorentz transformations provides a rich toolbox to investigate problems related to changes of daylight properties.

#### 3.1 Perron-Frobenius and Krein-Rutman Theory

For finite dimensional signal vectors the non-negativity of the first eigenvector follows easily from the Perron-Frobenius theory of non-negative matrices. We will therefore give a brief overview and refer the interested reader to Chapter 13 in [6] for more details

**Definition 1** 1. A matrix  $C$  is non-negative if all elements are non-negative.

2. A matrix  $C$  is positive if all elements are positive.

3. A matrix  $C$  is reducible if there is a permutation matrix  $P$  such that

$$P^{-1}AP = P'AP = \begin{pmatrix} C_1 & 0 \\ C_{21} & C_2 \end{pmatrix} \quad (1)$$

A matrix that is not reducible is called irreducible.

Note the difference between non-negative (positive) and non-negative (positive) **definite** matrices. The first are defined by properties of the elements of the matrix whereas the latter are defined via bilinear products! Note also that for positive (non-negative) matrices we require the transformation matrix  $P$  to be a permutation matrix. For a permutation matrix  $P$  the transposed  $P'$  of the matrix is its inverse  $P^{-1}$ .

One result of the Perron-Frobenius theory of non-negative matrices is the following theorem of Perron (see [6]):

**Theorem 1 [Perron]** A positive matrix has a real, maximal, positive eigenvalue  $r$ . This eigenvalue is a simple root of the characteristic equation and the corresponding eigenvector has only positive elements.

In the context of stochastic processes we will also need the following result about the normal form of non-negative matrices:

**Theorem 2** For a non-negative matrix  $C$  we can find a permutation matrix  $P$  such that:  $\tilde{C} = P'CP =$

$$\begin{pmatrix} C_1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & C_2 & \dots & 0 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & C_g & 0 & \dots & 0 \\ C_{g+1,1} & C_{g+1,2} & \dots & C_{g+1,g} & C_{g+1} & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ C_{s1} & C_{s2} & \dots & C_{s,g} & C_{s,g+1} & \dots & C_s \end{pmatrix} \quad (2)$$

$\tilde{C}$  is the normalform of  $C$ . The matrices  $C_i$  are irreducible and at least one of the matrices  $C_{i,1}, C_{i,2}, \dots, C_{i,i-1}$  is different from zero.

For stochastic processes of color signals we can now argue as follows: color signals have non-negative values, the correlation matrix  $C$  is therefore non-negative. We can find a permutation matrix to bring  $C$  into normalform and since  $C$  is also symmetric we find that  $C$  is block-diagonal with irreducible blocks. A stochastic process is thus the direct sum of uncorrelated subprocesses where each of the subprocesses has a strictly positive first eigenvector. In the following we will therefore only consider processes with strictly positive first eigenvectors. Many of the results of the Perron-Frobenius theory hold also in the more general case where the finite signals are replaced by elements in a Hilbert space and the correlation matrix by an operator. These results are generally known as Krein-Rutman theory (for details see [3], pp. 2129).

### 3.2 Geometry of Spaces of Daylight spectra

Understanding the properties of spectral distributions of daylight and its evolutions at different sites with varying atmospheric conditions is an active research area. Many efforts to measure and model daylight spectra originate in the 60s [10, 11, 17, 23, 24]. The SMARTS model [9] is one example of recent tools to compute clear sky spectral irradiance from a description of atmospheric conditions, time and solar geometries.

In the following we consider spaces of daylight spectra as subsets of vector spaces. We use PCA to approximate the spectra with linear combinations of three principal components. We saw above that the first eigenvector is positive. Therefore the other eigenvectors must have negative contributions somewhere (because of the orthogonality of the eigenvectors). From this we conclude that (given a fixed value for the coefficient of the first eigenvector) the coefficients of the higher order eigenvectors must be bounded (otherwise some negative contribution of this eigenvector will be larger than the corresponding positive contribution from the first eigenvector and the resulting linear combination will be negative, which is impossible since the spectra are non-negative everywhere). The daylight spectra are thus located in a solid cone-like region of the vector space. This allows us to introduce the conical structure of the coordinate vector space of spectra. For all spectra in the set we have (after a suitable re-scaling of the basis vectors):

$$\sigma_0^2 - \sigma_1^2 - \sigma_2^2 > 0; \quad \sigma_0 \geq 0 \quad (3)$$

where  $\sigma_k$  is the  $k$ -th coefficient in the PCA expansion. A conical projection in the space of spectra coordinates is defined as the following perspective projection:

$$x = \sigma_1/\sigma_0; \quad y = \sigma_2/\sigma_0 \quad (4)$$

A spectrum in this space is thus characterized by  $\sigma_0$  (related to the intensity) and the coordinate  $z = x + iy$  (related to the chromaticity). The points  $z = x + iy$  lie on the open unit disk of the complex plane. In the following we will use the terms "intensity" and "chromaticity" in the sense defined here. They are convenient terms but are, of course, different from their meaning in traditional color science. We characterize a spectra sequence by its sequence of projected points on the open unit disk.

It is known that the open unit disk is a two dimensional Poincaré model of hyperbolic geometry, and its isometry transformation group is  $\mathbf{SU(1,1)}$ .  $\mathbf{SU(1,1)}$  is the group consisting of  $2 \times 2$  complex matrices of the following form:

$$\mathbf{M} = \left\{ \begin{bmatrix} a & b \\ \bar{b} & \bar{a} \end{bmatrix} : a, b \in \mathbb{C}, |a|^2 - |b|^2 = 1 \right\} \quad (5)$$

An element  $\mathbf{M} \in \mathbf{SU(1,1)}$  acts as the fractional linear transformation on a point  $z$ :

$$w = \mathbf{M}\langle z \rangle = \frac{az + b}{\bar{b}z + \bar{a}} \quad (6)$$

Special subgroups of  $\mathbf{SU(1,1)}$ , known as one-parameter subgroups  $M(t)$  are given by group elements that are functions of the real values  $t$ , having the following properties:

$$\begin{aligned} M(t_1 + t_2) &= M(t_1)M(t_2); & \forall t_1, t_2 \in \mathbb{R} \\ M(0) &= E = \text{identity matrix} \end{aligned} \quad (7)$$

For a one-parameter subgroup  $M(t)$  we introduce its infinitesimal generator, represented by the matrix  $X$ :

$$X = \left. \frac{dM(t)}{dt} \right|_{t=0} = \lim_{t \rightarrow 0} \frac{M(t) - E}{t} \quad (8)$$

The infinitesimal matrices  $X$  representing one-parameter subgroups  $M(t) \in \mathbf{SU(1,1)}$  form the Lie algebra  $su(1,1)$ , which consists of  $2 \times 2$  complex matrices of the form:

$$su(1,1) = \left\{ \begin{pmatrix} i\gamma & \beta \\ \beta & -i\gamma \end{pmatrix} : \gamma \in \mathbb{R}, \beta \in \mathbb{C} \right\} \quad (9)$$

The Lie algebra  $su(1,1)$  forms a three-dimensional vector space [20], spanned by the basis  $(J_k)$ :

$$J_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}; J_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}; J_3 = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \quad (10)$$

Each infinitesimal matrix  $X \in su(1,1)$  of a one-parameter subgroup  $M(t) \in \mathbf{SU(1,1)}$  has thus a coordinate vector specified by the three real numbers  $\xi_1, \xi_2$  and  $\xi_3$ :  $X = \xi_1 J_1 + \xi_2 J_2 + \xi_3 J_3$ .

Given a start point  $z(0)$  on the unit disk together with a one-parameter subgroup  $M(t)$  we define an  $\mathbf{SU(1,1)}$  curve as the following function of  $t$ :

$$z(t) = M(t)z(0) = e^{tX}z(0); \quad t \in \mathbb{R}, z(t) \text{ on the unit disk} \quad (11)$$

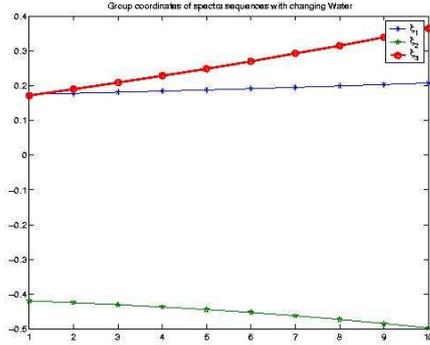
The  $\mathbf{SU(1,1)}$  curves are straight lines in the three dimensional Lie algebra space  $su(1,1)$ , thus the estimation of input chromaticity sequences using  $\mathbf{SU(1,1)}$  curves can be considered as a linearization. The Lie algebra  $\mathbf{SU(1,1)}$  provides a powerful tool to linearize problems involving chromaticity sequences. We developed several methods [15] to compute the group parameters of  $\mathbf{SU(1,1)}$  curves from input data sequences. These methods are used to study the properties of sequences of daylight spectra.

### 3.3 An Illustration

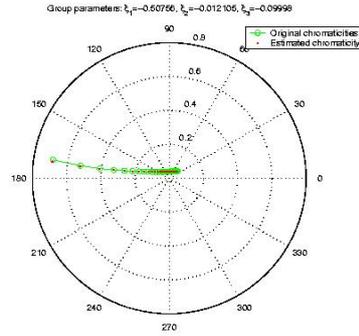
We investigated two different types of daylight spectra:

- A database of 21871 daylight spectra measured at the same location (Norrköping, Sweden) from June 16th, 1992 to July 7th, 1993; between 5:10 and 19:01 (Local time); in 5nm steps from 380nm to 780nm, and
- Daylight spectra sequences generated by the simulation program SMARTS2 [9]. The SMARTS model accepts as input the Sun position and atmospheric parameters including: Angström beta, precipitable water, ozone, and surface pressure. The wavelength range of the generated spectra was 380nm to 780nm in 1nm steps. In a series of experiments, sequences of spectra were generated by changing a single parameter in its allowed range while keeping the others fixed to the default values. In another experiment a large set of simulated daylight spectra with all the feasible combinations of parameters were also created to simulate the investigated space of daylight spectra.

Figure 3.3 shows the results of the  $\mathbf{SU(1,1)}$  estimations of chromaticity sequences generated by SMARTS2. In one experiment we generated spectra with different values of the Angström beta and Precipitable water parameters. They can be seen as the images of a grid in the (Angström beta, precipitable water) plane under the SMARTS2 mapping. The chromaticity sequences corresponding to single parameter (Beta or Water) changes are then created by grouping all chromaticity points having the same value of the other parameter. We choose the value of the precipitable water parameter as constant and estimate the one-parameter group as the Angström beta varies. The group coordinates characterizing these estimated  $\mathbf{SU(1,1)}$  curves are not constant but form a linear function for different settings of the other parameter. In Figure 3.3 we show the computed group coordinates of estimated  $\mathbf{SU(1,1)}$  curves, varying as functions of the other parameter's settings. In another



(d) Group coordinates from varying water values

(e)  $SU(1,1)$  curve estimating changing Sun position

example we generated different daylight spectra by varying the zenith angle with all other parameters set to default values. This corresponds to a sequence of daylight spectra with different positions of the Sun. It can also be considered as a time sequence of daylight spectra. The same  $SU(1,1)$  curve can be used to estimate both sequences of equally spaced zenith angle changes and equally spaced time changes, but the group parameters  $t$  describing the locations of points on curve are different for each sequence.

In all the experiments described above, we found that  $SU(1,1)$  curves provide good estimation for all SMARTS2 spectra sequences generated by changing any single parameter. The method also provides good approximations when the remaining parameters are set to other reasonable values different from the default values.

#### 4 Lie Theory, Differential equations and Invariants

One advantage of the framework described above is the connection to the Lie-theory of differential equations. This connection is provided by the one-parameter groups defined in the last section. Assume we have a one-parameter group  $M(t)$  describing the chromaticity changes of an illumination source. One example could be a sunrise or sunset described by the time-varying sequence of spectra in the experiment mentioned in the last section. Given a chromaticity parameter  $z(0)$  this group generates a sequence of time-varying chromaticity parameters  $z(t) = M(t)\langle z(0) \rangle$  as defined in Eq. (11). Choosing an intensity parameter one can reconstruct the coefficients in the PCA decomposition and by computing the linear combination with the principal components we obtain a sequence of the time-varying illumination spectra  $l(t, \lambda)$  where  $\lambda$  is the wavelength parameter and  $t$  is the group parameter. If we now model the camera as a linear system described by the matrix  $K$  then the time-changing pixel vectors  $p(t)$  for a given object point are given by  $p(t) = Kl(t, \lambda)$ . Since the chromaticity transformations form a one-parameter group it is possible to define the Lie-derivative operator  $\frac{d}{dt}$  with respect to the group parameter  $t$ . This construction can then be used to investigate properties of measurement sequences like the time-varying pixel values recorded by one channel of a color camera.

Another application where these tools can be used is the construction of invariants. With the help of the Lie-theory of differential equations it is possible to give an overview over all invariants and to construct these invariants automatically using programs like Maple. We will illustrate this with the construction of invariants for the dichromatic reflection model.

Name	Matrix Group	Lie-Algebra (in nitesimal element)
Rotation	$\begin{pmatrix} \cos(\tau) & \sin(\tau) \\ -\sin(\tau) & \cos(\tau) \end{pmatrix}$	$X_1 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$
Uniform Scaling	$\begin{pmatrix} e^\sigma & 0 \\ 0 & e^{-\sigma} \end{pmatrix}$	$X_2 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$
Non-Uniform Scaling	$\begin{pmatrix} e^\sigma & 0 \\ 0 & 1 \end{pmatrix}$	$X_3 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$
Shear	$\begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix}$	$X_4 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$

**Table 3. Subgroups of the general linear group of  $2 \times 2$  matrices**

In many applications, for example in color image segmentation, color object recognition etc., the main interest is the physical content of the objects in the scene. Deriving features which are robust to image capturing conditions such as illumination changes, highlights, shadows and geometry changes is a crucial step in such applications. The interaction between lights and objects in the scene is very complicated and sophisticated models such as Transfer Radiative Theory or Monte-Carlo simulation methods are needed to provide realistic results. The complexity of these models (and the fact that many key components are unknown) prevents their application in object recognition. Previous studies of color invariance are, therefore, mostly based on simpler semi-empirical models such as the Dichromatic Reflection Model [21], or the model proposed by Kubelka and Munk [13], together with many additional assumptions [5, 7, 8, 12, 22].

We already mentioned that it is very difficult to describe in detail what happens when light strikes a surface: some of the light will be reflected at the interface producing interface reflection, while another part will transfer through the medium undergoing absorption, scattering, and emission. The Dichromatic Reflection Model (see [21]) describes the relation between the incoming light and the reflected light as a mixture of the light reflected at the material surface and the light reflected from the material body. The model assumes that the light  $L(x, \lambda)$  (of wavelength  $\lambda$ ) reflected from a point of the surface can be decomposed into two additive components: an interface (specular) reflection and a body (diffuse) reflection:

$$L(x, \lambda) = m_S(x)R_S(\lambda)E(\lambda) + m_D(x)R_D(\lambda)E(\lambda) \quad (12)$$

An interpretation is that the geometric properties at  $x$  (such as the angle of incidence light, the angle of remittance light etc.) determine the weights  $m_S(x)$  and  $m_D(x)$ . The relations between the spectral properties of the illumination, the specular and the diffuse reflections depend on the properties of the material and are given by the relations:  $R_S(\lambda)E(\lambda)$  and  $R_D(\lambda)E(\lambda)$  where  $E(\lambda)$  is the spectral power distribution of the incident light. The measured sensor values  $C_n(x)$  at pixel  $x$  in the image using  $N$  filters with spectral sensitivities given by  $f_1(\lambda) \dots f_N(\lambda)$  will be given by the following integral over the visible spectrum:

$$\begin{aligned} C_n(x) &= \int f_n(\lambda) [m_S(x)R_S(\lambda)E(\lambda) \\ &\quad + m_D(x)R_D(\lambda)E(\lambda)] d\lambda \\ &= m_S(x)S_n + m_D(x)D_n \end{aligned} \quad (13)$$

Assume that two object points belong to the same material. They have therefore identical reflection functions and the only difference between the corresponding pixels in the image originates in their geometrical properties. For these two neighboring pixels  $x_1$  and  $x_2$  and one selected channel  $n$  we have then:

$$\begin{bmatrix} C_n(x_1) \\ C_n(x_2) \end{bmatrix} = \begin{bmatrix} m_S(x_1) & m_D(x_1) \\ m_S(x_2) & m_D(x_2) \end{bmatrix} \begin{bmatrix} S_n \\ D_n \end{bmatrix} = M \begin{bmatrix} S_n \\ D_n \end{bmatrix} \quad (14)$$

In the the framework of transformation groups we see that the matrix  $M$  operates on the 2-D vectors  $(S_n \ D_n)'$ . In the group theoretical approach it is now natural to consider groups of matrices  $M$  operating on the vectors  $(S_n \ D_n)'$  and to construct invariants for such groups. Examples of relevant groups of matrices  $M$  are listed in Table 3:

The subgroups in Table 3 form one-parameter groups with infinitesimal elements  $X_k$ . It is now natural to consider all linear combinations of one, two, three or four of the infinitesimal elements  $X_k, k = 1, \dots, 4$ . Combining  $X_2, X_3$  we can thus generate the two-parameter group of scalings and the full group is a four parameter group generated by all linear

combinations of all four in infinitesimal matrices. General Lie-theory requires however not only to consider linear combinations but also the Lie-product of two elements  $X, Y$  defined as  $XY - YX$ . A vector space that is closed under the Lie-product is called a Lie-algebra and its vector-space dimension is the dimension of the Lie-algebra. Computing this Lie product for the shear matrix  $X_4$  and the rotation matrix  $X_1$  we find that the subgroup generated by the shear and the rotation operations must also include the scaling matrices. Closure with regard to the Lie-product implies that the two one-parameter groups of shears and rotations generate the three-parameter Lie-algebra of shears, rotations and uniform scalings. We have now shown how the group of non-singular  $2 \times 2$  matrices (and its subgroups) operate on measurement vectors generated by one sensor type at two different positions. Next we will generalize this to the case where we measure simultaneously  $N$ -channels at each pixel. The most common example of such a generalization corresponds to a transition from intensity- to RGB-images. Since we separated the spectral properties and the non-spectral parameters we see that the transformation matrix  $M$  is the same for all channels. Therefore we obtain in the general case the transformation group:

$$\begin{bmatrix} C(x_1) \\ C(x_2) \end{bmatrix} = M \begin{bmatrix} S \\ D \end{bmatrix} \quad (15)$$

Here  $C(x_1), C(x_2), S, D$  are  $N$ -dimensional vectors. The group now operates on the space  $(\mathbb{R}^2)^N$ . A general result from Lie-theory (see [16]) shows that this implies that there are  $2N - k$  functionally independent invariants where  $k$  is the dimension of the Lie-algebra. For the case of RGB images and the full matrix group we have  $2 \cdot 3 - 4 = 2$  invariants. The general theory does not only give the number of functionally independent invariants but it also provides an algorithm how to construct these invariants. The

A simple Maple program gives the following solutions (with  $C(x_k) = (r_k, g_k, b_k)$ ):

$$f = \_F1 \left( \frac{-g_2 b_1 + b_2 g_1}{r_1 g_2 - r_2 g_1}, \frac{b_2 r_1 - r_2 b_1}{r_1 g_2 - r_2 g_1} \right) \quad (16)$$

If we only start with rotations and shearing then we start with two variables but because of the Lie-structure we have  $k = 3$  and there are three independent invariants:

$$f = \_F1(r_1 g_2 - r_2 g_1, b_2 r_1 - r_2 b_1, \frac{-g_2 b_1 + b_2 g_1}{r_1 g_2 - r_2 g_1})$$

## 5 Conclusions

We used a few examples to illustrate the possible usage of group theoretical methods to solve color related problems. They show that it may be well worth to invest the time and effort to learn the tools of group representations and Lie-theory. The advantages are, at least, twofold: (1) this may lead to new insights into the fundamental properties of spaces of color signals and (2) group theory provides tools and algorithms to solve difficult problems with the help of programs like Maple.

## References

- [1] J. J. Atick and N. Redlich. What does the retina know about natural scenes. *Neural Computation*, 4:449–572, 1992.
- [2] H. B. Barlow. The coding of sensory messages. In W. H. Thorpe and O. L. Zangwill, editors, *Current Problems in Animal Behaviour*, pages 331–360. Cambridge University Press, 1961.
- [3] N. Dunford and J. T. Schwartz. *Part III, Spectral Operators*. Interscience Publishers, New York, 1988.
- [4] D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *josaa*, 4(12):2379–2394, dec 1987.
- [5] G. D. Finlayson and G. Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *International Journal of Computer Vision*, 42(3):127–144, 2001.
- [6] F. R. Gantmacher. *Matrizentheorie*. Springer Verlag, Berlin, Heidelberg, New York, Tokyo, 1986.
- [7] J. M. Geusebroek, R. van den Boomgaard, A. W. M. Smeulders, and H. Geerts. Color invariance. *IEEE Trans. Pattern Anal. Machine Intell.*, 23(12):1338–1350, 2001.
- [8] T. Gevers and A. W. M. Smeulders. Pictoseek: Combining color and shape invariant features for image retrieval. *IEEE Trans. on Image Processing*, 9(1):102–119, 2000.
- [9] C. Gueymard, “Smarts, a simple model of the atmospheric radiative transfer of sunshine: Algorithms and performance assessment,” Tech. Rep. FSEC-PF-270-95, Florida Solar Energy Center, (1995).
- [10] S. T. Henderson and D. Hodgkiss, “The spectral energy distribution of daylight,” *J. Appl. Phys.*, 14, pg. 125–131, (1963).

- [11] D. B. Judd, D. L. MacAdam, and G. Wyszecki, "Spectral distribution of typical daylight as a function of correlated colour temperature," *Journal of the Optical Society of America A*, 54, pg. 1031–1041, (1964).
- [12] G. J. Klinker. *A Physical Approach to Color Image Understanding*. A K Peters Ltd., 1993.
- [13] P. Kubelka and F. Munk. Ein Beitrag zur Optik der Farbanstriche. *Zeitschrift für Technische Physik*, 11a:593–601, 1931.
- [14] R. Lenz. Investigation of receptive fields using representations of dihedral groups. *Journal of Visual Communication and Image Representation*, 6(3):209–227, September 1995.
- [15] R. Lenz, T. H. Bui, and J. Hernández-Andrés. One-parameter subgroups and the chromaticity properties of time-changing illumination spectra. In *Proc. SPIE-2003, Color Imaging VIII, Vol 5008*, pages 237–248, 2003.
- [16] P. J. Olver. *Applications of Lie Groups to Differential Equations*. Springer, New York, 1986.
- [17] Z. Pan, G. Healey, and D. Slater, "Global spectral irradiance variability and material discrimination at boulder, colorado," *Journal of the Optical Society of America A*, 20, 3, pg. 513–521, (2003).
- [18] C. A. Párraga, G. Brelstaff, T. Troscianko, and I. R. Moorehead. Color and luminance information in natural scenes. *J. Opt. Soc. Am. A*, 15(3):563–569, March 1998.
- [19] P. Reinagel and S. Laughlin. Natural stimulus statistics. *Network: Computation in Neural Systems*, 12(3):237–240, 2001.
- [20] D. Sattinger and O. Weaver. *Lie Groups and Algebras with Applications to Physics, Geometry and Mechanics*, volume 61 of *Applied Mathematical Sciences*. Springer, 1986.
- [21] S. A. Shafer. Using color to separate reflection components. *Color Research and Application*, 10(4):210–218, 1985.
- [22] H. M. G. Stokman. *Robust Photometric Invariance in Machine Color Vision*. PhD thesis, Intelligent Sensory Information Systems group, University of Amsterdam, Nov. 2000.
- [23] G. T. Winch, M. C. Boshoff, C. J. Kok, and A. G. du Toit, "Spectroradiometric and colorimetric characteristics of daylight in the southern hemisphere: Pretoria, south africa," *Journal of the Optical Society of America A*, 56, pg. 456–464, (1966).
- [24] G. Wyszecki and W. Stiles. *Color Science*. Wiley & Sons, London, England, 2 edition, 1982.

# Image Relighting: Getting the Sun to Set in an Image Taken at Noon

Claus B. Madsen and Rune Laursen  
 Laboratory of Computer Vision and Media Technology  
 Aalborg University, Aalborg, Denmark  
 cbm/rul@cvmt.auc.dk

## Abstract

Image relighting is a very unique special visual effect which promises to have many important practical applications. Image relighting is essentially the process of, given one or more images of some scene, computing what that scene would look like under some other (arbitrary) lighting conditions, e.g., changing positions and colors of light sources. Image relighting can for example be used for interior light design. This paper describes an approach to image relighting which can be implemented to run in real-time by utilizing graphics hardware, as opposed to other state-of-the-art approaches which at best run at a few frames per second.

## 1 Introduction

This paper addresses the subject of developing relighting techniques, i.e. techniques allowing a user to completely alter the lighting conditions in an image of a real scene. Specifically, the paper focuses on techniques providing *real-time* relighting functionalities, thus enabling the user to interactively change lighting conditions and get "instant" visual feedback. Figure 1 provides an example of the kind of relighting this paper addresses. It should be noted that the work presented here presumes the availability of three things: 1) an original image of the scene, 2) a 3D model of the scene, and 3) a model of the lighting conditions in the scene at the time the original image is acquired. We will return to ways in which the two last pieces of knowledge can be obtained.

Conceptually image relighting in this manner is a two step process. In the first step all effects of the original lighting conditions are *removed*, e.g., highlights, shadows,

and differences in shading across surfaces due to varying light incidence angles. In the second step the scene is subjected to some arbitrary new lighting conditions and the appearance of the scene in these conditions is computed. The second step thus "adds" new highlights, shadows and shading etc. Of these two steps the former is the tricky one, while the latter can be performed using any preferred rendering technique, e.g., ray tracing, radiosity or standard hardware accelerated approaches. Which rendering technique is employed depends on the preferred balance between rendering speed and accuracy in handling various lighting phenomena. In order to achieve true real-time performance we have chosen to use a hardware accelerated approach for step 2, thus sacrificing certain global illumination phenomena.

In our approach step 1 is achieved by a *computational* approach which requires, as stated above, a 3D model of the scene and a model of the original lighting conditions. Alternatively, one could in principle acquire fronto-parallel digital images of the surfaces in the scene under perfectly diffuse, white-balanced lighting conditions and use these images as textures on the 3D model, which is subsequently rendered under novel lighting conditions (step 2). This would be a *mechanical* or *image acquisition* approach to step 1, but in reality acquiring such 'clean' textures devoid of lighting effects is not practical for general scenes.

The contributions in this work lie in the specific manner in which the operations performed in steps 1 and 2 are carried out. With the approach described here the two steps can be combined such that the image relighting becomes a matter of modulating the original image on a pixel by pixel basis with a "relighting map". The relighting map can be computed in real-time using standard techniques,



Figure 1: Left: original image acquired outdoors on a sunny day. Right: same scene as left but this image is a *simulation* of what the scene would look like if the position of the sun were different. The work in this paper enables such changes in lighting conditions to be performed in real-time.

and the modulation can also easily be performed in real-time. Thus, our approach has two advantages: 1) it is directly designed for real-time performance, and 2) the original image is used directly and therefore the final image is not subject to filtering and/or aliasing effects involved in doing reprojections of textures mapped to a 3D model of the scene.

This paper represents the current state of work in progress and will only address the problem for scenes with perfectly diffuse reflectance properties. The paper is organized as follows. In section 2, we present an overview of our approach and show how the relighting effects are achieved. Section 3 then describes related work. Section 4 describes our approach in more detail, followed by section 5 giving some practical details behind the initial experiments we have performed to validate the proposed approach. Section 6 discusses central aspects of the work and points to future research. Finally section 7 offers conclusions.

## 2 Overview of approach

Prior to describing the proposed approach in a more technically rigorous manner this section attempts to provide the reader with an intuitive understanding of the issues

involved and of the process behind our technique.

The approach requires different types of input information. First of all an image of the original scene is required. Secondly, a 3D model of the scene must be available, and the original image must be calibrated to the 3D model such that every pixel corresponds to a known 3D point in the scene model. Third, the original lighting conditions in the scene must be known, i.e., we need to know the sources of light in the scene, and their relative intensities.

The 3D model can be obtained in many different ways [7], e.g., by reconstruction from multiple images using approaches such as [8], by Image-Based Modelling, e.g. [4], or by laser range scanning. Alternatively the scene can be measured and a model constructed manually. The latter is the approach employed for our experimental results, i.e., we have measured the scene, constructed crude polygonal models of the objects, and then calibrated the camera to the 3D model using manually established 2D to 3D point correspondences. Figure 2 shows the scene model used for the relighting illustrated in figure 1.

The required knowledge of the original lighting conditions can most easily be acquired using the popular light probe approach, i.e., by taking high dynamic range images of a reflective sphere placed in the scene, [2, 5]. Alternatively, light source positions, sizes and power can be measured manually as done in [6] or semi-automatically

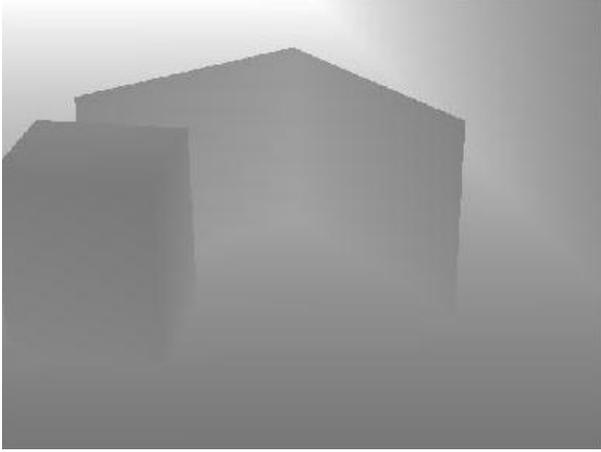


Figure 2: 3D model corresponding to scene shown in figure 1. The image illustrates the model as a depth map, i.e., intensity is a measure of distance from the camera. We have used one plane for the tiled ground plane, a plane for the brick wall on the right, six quadrilaterals for the speaker on the left, and two quads and one triangle for calibration object in the center.

using multiple images as in [12]. For the experimental results in this paper we have done it manually in a manner described in section 5.

In this paper we will limit ourselves to discussing the case of scenes containing surfaces with perfectly diffuse reflectance properties.

Each pixel in the original image is a measurement of the radiance (in the three RGB bands) from a unique 3D point in the scene in the direction of the viewpoint. Thus the original image is a 2D **radiance map**,  $L_o(u, v)$ , where  $u$  and  $v$  are the image coordinates. Because we have the 3D model and knowledge of the original lighting conditions it is trivial<sup>1</sup> to compute the amount of light arriving at the same 3D points in the scene, i.e., it is possible to construct an **irradiance map**,  $I(u, v)$ . When the irradiance map is computed using the known original scene lighting conditions we will call it  $I_o(u, v)$ . Conversely, when the irradiance map is computed using some arbitrary different

<sup>1</sup>Irradiance computation is trivial provided global illumination issues (irradiance contributions from diffuse reflections) are disregarded. If this is not a fair assumption the work by Yu et al., [12], can be used to compute these contributions.

relighting conditions it will be denoted  $I_r(u, v)$ .

For purely diffuse Bidirectional Reflectance Distribution Functions (BRDFs) there is a linear relationship between radiance and irradiance (radiance is proportional to diffuse albedo times irradiance). Therefore diffuse scenes can be very simply relit by dividing the radiance map with the original irradiance map, and then multiplying with the relighting irradiance map. Using the introduced terminology  $L_r(u, v) = L_o(u, v) \cdot I_r(u, v) / I_o(u, v)$ .

Figure 3 shows original and relighting irradiance maps corresponding to the relighting example in figure 1.

Section 1 presented the general concept of relighting as a two step process: 1) removing original light from the image, and 2) adding new light. In the described diffuse case step 1 is represented by the  $L_o(u, v) / I_o(u, v)$  operation, whereas step 2 is performed by multiplying the result from step 1 by the relight irradiance,  $I_r(u, v)$ . Step 1 is a once-only process as it only involves elements that do not change over time (original image and original irradiance map). Step 2 has to be re-performed constantly in response to the user's changes in the desired lighting conditions for the relit scene, so  $I_r(u, v)$  must in general be re-computed for every frame.

Step 1 could be pre-computed and the result stored in an image which is then subsequently modulated by the real-time computed relighting irradiance map. Our approach does not do this; we have designed a solution which embeds the normalization with the original irradiance into the computation of the relighting irradiance map,  $I_r(u, v)$ . Thus, at run-time we actually compute the relight/original light *ratio*  $I_r(u, v) / I_o(u, v)$  directly. This "ratio map" is then used for modulating the original image. By doing this we avoid the non-trivial implementation of a real-time texture division operation.

Computing the ratio map is described in detail in section 4. The idea is based on the observation that if one renders a radiance image of an all-white perfectly diffuse 3D model of the scene under the chosen relighting conditions, this image is then identical to the required irradiance map,  $I_r(u, v)$ . If instead we set the reflectances of the 3D model to be proportional to the inverse of the original irradiances then the rendered image automatically becomes the desired relight/original irradiance ratio map.



Figure 3: Left: irradiance map,  $I_o(u, v)$ , corresponding to original lighting conditions of the scene shown in gure 1. Center: irradiance map,  $I_r(u, v)$ , corresponding to lighting conditions where the dominant light source, the sun, has changed location relative to the scene. These are the lighting conditions valid for the relit image in gure 1. Right: illustrates the "relighting factor", i.e., the ratio of original irradiance to relighting irradiance,  $I_r/I_o$ . Every pixel in the original radiance map is modulated with its corresponding pixel in this map. Dark means "light is removed" from original pixels, and bright means "light is added" to original pixels.

### 3 Related work

There is a small amount of closely related work in the literature. First of all Yu et al., [13], demonstrated how they could acquire reflectance properties of architectural scenes by taking at least two images of each surface of the objects under different lighting conditions. The recovered 3D model combined with the estimated reflectance parameters could then be used to render the scene under changing lighting conditions. The focus of this work is entirely on parameter recovery and relighting is by no means done in real-time.

Similarly, inverse global illumination was proposed by Yu et al., [12] for recovery of reflectance parameters for indoor scenes using multiple images of each surface from different viewpoints. Again this work focuses on reflectance parameter estimation and relighting is done using RADIANCE, [10], which again is far from real-time.

The most closely related work is that of Loscos et al., [6]. This work also enables a user to change the lighting conditions in an image of a scene in an interactive manner, but this work is centered around a radiosity method for irradiance computations. Therefore, the method performs at a few frames per second when the only lighting changes performed are intensity adjustments. If the number of sources or their positions are changed updating takes on the order of 10 seconds.

The work in [6] also employs texture modulation for

efficient relighting, and the modulation texture (irradiance map) is computed using radiosity. We have chosen to focus specifically on true real-time performance and therefore the computation of the relighting irradiance maps does not account for global illumination phenomena such as color bleeding. Nevertheless, with the work currently being done in Pre-computed Radiance Transfer and Photon Mapping, real-time global illumination is coming closer and closer to reality, and our approach can readily be combined with such efficient global illumination techniques.

### 4 The perfectly diffuse case

For perfectly diffuse reflectors the relationship between incident irradiance,  $I$ , and radiance,  $L$ , in any direction is given by the diffuse albedo,  $\rho_d$ :

$$L = \frac{\rho_d}{\pi} I \tag{1}$$

The original image (radiance map),  $L_o(u, v)$ , provides us with measured radiances from a dense set of 3D points in the scene, and these points are known since we assume the camera is calibrated to the scene. In general the diffuse albedo and the irradiance vary for every point in the scene, so the relationship between radiance maps and irradiance

maps becomes:

$$I_o(u, v) = \frac{\rho_d(u, v)}{\pi} I_o(u, v) \quad (2)$$

Here,  $\rho_d(u, v)$  is the "albedo map". When doing re-lighting the albedo stays constant; the only thing that changes is the irradiance at each scene point. Therefore, the radiance map of the relit image/scene can be expressed as:

$$\begin{aligned} L_r(u, v) &= \frac{\rho_d(u, v)}{\pi} I_r(u, v) \\ &= (L_o(u, v)/I_o(u, v)) \cdot I_r(u, v) \\ &= L_o(u, v) \cdot (I_r(u, v)/I_o(u, v)) \end{aligned} \quad (3)$$

Eq. 3 simply shows that the relit image can be computed by modulating the original image with a ratio of two irradiance maps: the relight irradiance map,  $I_r(u, v)$ , corresponding to the user's desired scene (new) lighting conditions, and the original irradiance map,  $I_o(u, v)$ . The key element in our approach is a technique for computing this map in real-time and using it for modulation of the original image.

#### 4.1 Computing the irradiance ratio map

How can we efficiently compute the irradiance ratio map? First, let us describe how simply the relighting irradiance map can be computed using standard local illumination techniques (specially we will use the Phong lighting model of OpenGL, a description of which may be found in books such as [1, 11]). Rendering an image of a scene using the Phong lighting model results in a radiance from a 3D point which can be formulated as (disregarding specular reflection):

$$L = \rho_a I_a + \rho_d \sum_{i=1}^k I_i \cos(\theta_i) \quad (4)$$

Here  $L$  is the radiance from a 3D point in the direction of the viewpoint.  $\rho_a$  and  $\rho_d$  are the ambient and diffuse reflectances, respectively, (eq. 4 is to be evaluated for each of three RGB colors).  $I_a$  is the ambient irradiance at the 3D point.  $I_i \cos(\theta_i)$  is the irradiance at the point caused by the  $i$ th point light source,  $(\theta_i(u, v))$  is the angle between the surface normal and the direction vector to the

$i$ th light source at the 3D point).  $k$  is the number of light sources.

If we set the ambient and diffuse reflectances equal, eq. 4 changes to:

$$L = \rho_d \left( I_a + \sum_{i=1}^k I_i \cos(\theta_i) \right) \quad (5)$$

Eq. 5 states that rendering with OpenGL Phong lighting the radiance from a point equals the reflectance at the point times the total irradiance (ambient plus sum of individual point source contributions) at the point. Thus, by setting unit reflectances, the *radiance equals the irradiance*. This may be self-evident but is important because it shows that we can use the graphics card's efficient lighting computation capabilities to produce irradiance maps needed for relighting.

That is, if we render the 3D scene model from a viewpoint corresponding to the original image, and if all surfaces in the rendered 3D model have unit reflectances, then the resulting image is an irradiance map. This means that relighting irradiance maps,  $I_r(u, v)$ , for any user desired lighting conditions can be rendered simply by rendering a diffuse, all-white 3D scene model under the chosen lighting conditions.

To actually do relighting we not only needed real-time computation of  $I_r(u, v)$ , but we needed the lighting ratio map,  $I_r(u, v)/I_o(u, v)$ . This is accomplished by setting the reflectances of points in the 3D model to the inverse of the original irradiance at that point,  $\rho_d = 1/I_o$ .

To summarize the light ratio maps are generated by doing the following rendering using hardware acceleration:

1. upload 3D scene model to graphics card
2. set ambient and diffuse RGB reflectances of all vertices to the inverse of the original irradiance at that 3D point
3. setup the desired lighting conditions consisting of ambient and point source contributions
4. render the model to a texture using a viewport corresponding to the camera in the original image

## 4.2 Practical issues

In the previous section we described how to use hardware accelerated local lighting rendering to produce irradiance ratio maps for modulating the original image. With this approach there is really no limits to how much the lighting conditions in the scene can be altered.

We are presently implementing the proposed technique but all images in this paper were produced by a non-real-time simulation of the presented approach. Figure 4 shows what the original scene looks like with a (non-existing) light source in the very center of the scene.

For the ongoing implementation of the real-time version the only real issue to contemplate is the resolution of the 3D model of the scene. In order to properly capture gradients in the original irradiances of the scene the resolution of the 3D model has to be high at such gradients. We are working on designing methods for adaptive subdivision based on evaluating irradiance differences between 3D model vertex locations. If differences are too high the surface is subdivided.

Computing original irradiances to be used inversely as reflectances of the (subdivided) 3D scene model is an off-line process which can be done using any preferred rendering technique, for example Monte Carlo ray tracing to enable proper handling of area light sources. This is especially possible if a high dynamic range light probe image of the scene is available, because then an Image-Based Lighting approach, [2, 3], can be used to compute accurate irradiances which properly handle soft shadows in the original image.

In the on-line stage, when rendering the 3D model with the assigned reflectances, cast shadows are important for proper irradiance computation. For this we propose to use a shadow volume approach to detecting shadowed areas.

## 5 Experiments

As mentioned previously the images shown in this paper are produced using a simulation of the presented technique. The original image was acquired with a standard 5 mega pixel digital camera. The scene was measured manually and a simple 3D model of it was constructed (as described in section 2).

The camera was calibrated to the 3D model using man-

ually established 2D to 3D point correspondences, and the estimation of internal and external camera parameters was done using an approach from [9].

The original lighting conditions were modelled as a combination of a point light source (the sun) and an ambient term (the sky). The position of the sun relative to the scene model was determined by orienting the calibration object such that sun rays were parallel to the xz-plane, thus fixing the sun's y-coordinate to zero. The x and z coordinates were then found by measuring the length of a shadow cast by an object of known height. The RGB intensities of the blue ambient sky light was determined from the image colors of the white paper of the calibration object in areas not exposed to sunlight. By comparing RGB values of calibration object cardboard in shadow and in direct light the relative intensities between ambient and sun light were determined, (taking the cosine fall-off for diffuse reflection into account for the sun point source).

It should be made very clear that the original lighting conditions modelled as described above are extremely crude and this was only done to get quick results. In the future we will use light probe images.

For computing the original irradiance map a simple ray tracing approach was implemented which considers local illumination only, by casting primary rays plus shadow feelers. The original radiance map was computed in image resolution. The relighting examples given in the paper basically involve changing the location of the sun source. Given some desired sun position the simple raytracer was used to render a relighting irradiance map, again in image resolution. The relighting irradiance map was divided by the original irradiance map and the result multiplied with the original image to complete the diffuse relighting process.

## 6 Discussion

In this section we will briefly discuss some important points in relation to our proposed method.

Using our approach it is possible to employ arbitrarily complex and accurate computations of the original irradiance map. This is an off-line, once-only computation the results of which are used to set the reflectances of the 3D scene model subsequently used for relighting. We believe handling area light sources to be very important,

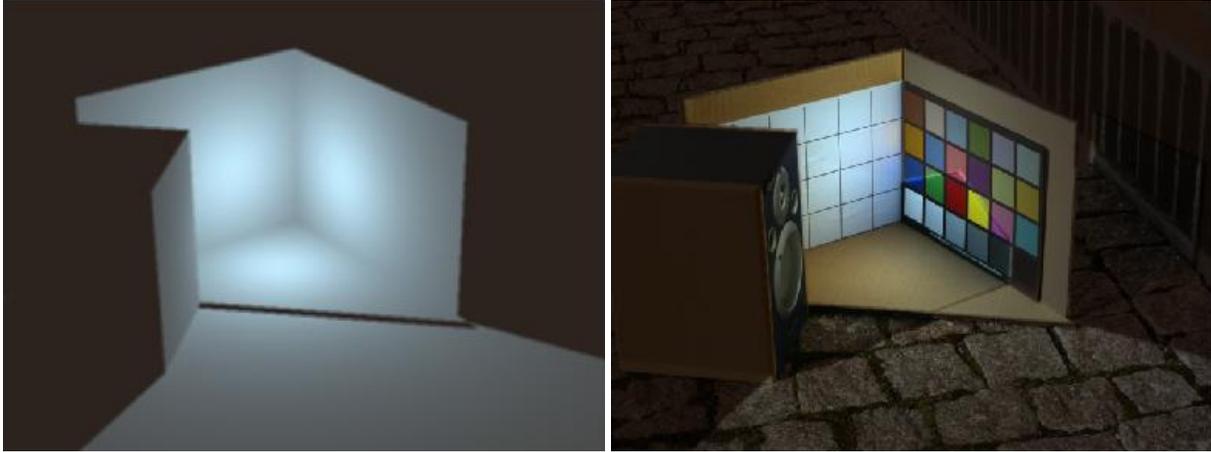


Figure 4: Left: irradiance map corresponding to a user defined lighting environment with a weak ambient term and a point light source in the middle of the scene. Right: resulting relit image.

even for outdoor images, since shadows due to sun light actually do have noticeable penumbra regions. Similarly, we believe taking global illumination phenomena (indirect light) into account is important, especially for indoor scenes, where reflections from other surfaces may be a significant irradiance contribution for a given surface.

Conversely, for the actual on-line, real-time rendering of irradiances during interactive relighting we have here proposed a straight forward local illumination approach. Yet, the basic approach with using the 3D scene model, normalized with original irradiances, can be used in conjunction with any lighting algorithm depending on how accurate one desires the result should be.

Throughout this paper we have assumed scenes to consist entirely of Lambertian materials. Our approach actually does generalize nicely to scenes with glossy BRDFs. It requires an additional rendering pass in the real-time relighting process, in order first to modulate the original image with the diffuse part of the relighting/original irradiance ratio map and subsequently add the specular radiance part. Figure 5 demonstrates the effect of adding a specular component to the surfaces during relighting.

## 7 Conclusions

We have described an approach to image/scene relight which based on a 3D model of the scene and knowledge of the original lighting conditions can compute the appearance of the scene under any arbitrary new lighting conditions, including changing the number of light sources, their positions and radiant powers.

The main contribution of the work is the fact that the approach is directly designed for real-time performance, enabling a user to get instant visual feedback upon having changed the parameters of the lighting environment. A smaller contribution lies in the idea of performing the normalization with the original irradiance by appropriately setting the reflectances of the 3D model used for real-time irradiance computations. This allows the approach to operate directly on the original image, rather than computing the albedo map off-line and modulating it at run-time.

An important aspect of the proposed approach is that relighting is performed as a modulation of the original image. It is believed that doing relighting in this manner is superior to an approach where textures extracted from the image are mapped to the scene geometry and reprojected at run-time, because the multiple re-sampling steps involved will cause the resulting image to be blurred.

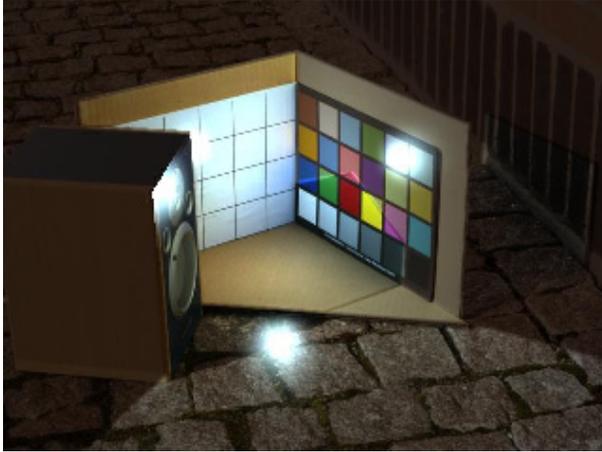


Figure 5: A specular reflection component has been added to each surface during relighting to illustrate the possibility of playing with the reflectance properties.

## Acknowledgments

This research is funded in part by the BENOGO project under the European Commission IST program (IST-2001-39184), and in part by the ARTHUR project (IST-2000-28559). This support is gratefully acknowledged.

## References

- [1] S. R. Buss. *3-D Computer Graphics – A Mathematical Introduction with OpenGL*. Cambridge University Press, 2003.
- [2] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings: SIGGRAPH 1998, Orlando, Florida, USA*, July 1998.
- [3] P. Debevec. Tutorial: Image-based lighting. *IEEE Computer Graphics and Applications*, pages 26 – 34, March/April 2002.
- [4] P. Debevec, C. Taylor, and J. Malik. Modelling and rendering architecture from photographs: a hybrid geometry and image-based approach. In *Proceedings: SIGGRAPH 1996, New Orleans, LA, USA*, pages 11 – 20, August 1996.
- [5] S. Gibson, J. Cook, T. Howard, and R. Hubbard. Rapid shadow generation in real-world lighting environments. In *Proceedings: EuroGraphics Symposium on Rendering, Leuven, Belgium*, June 2003.
- [6] C. Loscos, G. Drettakis, and L. Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, 6(4):289 – 305, October-December 2000.
- [7] M. M. Oliveira. Image-based modelling and rendering: A survey. *RITA - Revista de Informatica Teorica e Aplicada*, 9(2):37 – 66, October 2002. Brazilian journal, but paper is in English.
- [8] M. Pollefeys, R. Koch, and L. VanGool. Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters. *International Journal of Computer Vision*, 32(1):7 – 25, 1999.
- [9] E. Trucco and A. Verri. *Introductory Techniques for 3D Computer Vision*. Prentice Hall, 1998.
- [10] G. J. Ward. The radiance lighting simulation and rendering system. In *Proceedings: SIGGRAPH 1994*, pages 459 – 472, July 1994.
- [11] A. Watt and F. Policarpo. *3D Games: Real-Time Rendering and Software Technology*, volume 1. Addison-Wesley, 2001.
- [12] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proceedings: SIGGRAPH 1999, Los Angeles, California, USA*, pages 215 – 224, August 1999.
- [13] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *Proceedings: SIGGRAPH 1998, Orlando, Florida, USA*, pages 207 – 217, July 1998.

# Using Mixtures of Gaussians to Compare Approaches to Signal Separation

Kaare Brandt Petersen  
Technical University of Denmark

## 1 Introduction

In the signal separation technique called Independent Component Analysis (ICA), one refers to a problem as being "square" if there is as many measurements as sources to separate and "overcomplete" or "under-determined" if there is more sources than measurements. While square ICA is thoroughly investigated through many different approaches with well understood differences and similarities, the overcomplete case is still posing a difficult problem. Below is a short overview of some of the more interesting or illustrative of the approaches.

In 1999, Hagai Attias presented in [1] a Maximum Likelihood approach assuming the generative model  $\mathbf{x} = \mathbf{A}\mathbf{s} + \epsilon$ . In this approach a model distribution is constructed and using Mixture of Gaussians as priors makes it possible to complete the relevant integrals and obtain a closed form expression for the distribution over  $\mathbf{x}$ . The model distribution is approximated to the data distribution through the Kullback Leibler divergence. This approach is extremely flexible while still having appealing analytical properties and the only real drawback is the bad scaling behavior: A sum over  $K^D$  must be computed, which can be rather large for e.g. image data.

In 2000, Lewicki and Sejnowski presented in [5] a Maximum Likelihood approach assuming the generative model  $\mathbf{x} = \mathbf{A}\mathbf{s} + \epsilon$ . In this approach, the log likelihood of  $\mathbf{A}$  is Taylor expanded to second order around the maximum a posteriori estimate of the sources, i.e. one approximate the likelihood with a gaussian, and the sources are in turn are estimated using the updated estimate of  $\mathbf{A}$ . In this approach we see the difficulty that most overcomplete techniques is trying to work around: The likelihood, or some other suitable cost-function, contains some integral involving the source prior and is therefore in general hard to solve. The approach of Lewicki and Sejnowski from 2000 substitutes the integral with a second order expansion and we shall see other possibilities in the following.

In 2001, Girolami presented in [2] a Maximum Likelihood approach assuming the generative model  $\mathbf{x} = \mathbf{A}\mathbf{s} + \epsilon$ . In this approach, Girolami assumes Laplacian priors and the integral associated with the log likelihood is approximated through a variational scheme: The laplacian priors are reformulated in dual space which provides a lower bound of the likelihood to be optimized. The drawback of this approach is that the trick only works for laplacian priors and that the algorithm optimizes a lower bound in stead of the log likelihood it self.

In 2002, Hojen-Soerensen, Winther and Hansen presented in [4], the so-called "Mean Field Approach" assuming the generative model  $\mathbf{x} = \mathbf{A}\mathbf{s} + \epsilon$ . The parameters  $\mathbf{A}$  and possibly the noise covariance  $\mathbf{W}$  are estimated through maximum likelihood assuming knowledge of the mean values of the sources and vice versa. That is, the integral of the log likelihood, translates into mean values of the sources, which are approximated with estimates of the mean values obtained from Mean Field theory. The nice feature of this approach is that we substitute a very complicated integral with an easier non-linear equation and that one can do this for any prior. The problem, of course is that although Mean Field estimates are fairly accurate they are still approximations.

Also in 2002, Shriki, Sampolinski and Lee presented in [7] an interesting variant of Infomax on the filtering model  $\hat{\mathbf{s}} = \mathbf{W}\mathbf{x}$ . In the setup  $\mathbf{y} = g(\mathbf{W}\mathbf{x})$ , Shriki et. al. obtains a relation between  $p(\mathbf{x})$  and  $p(\mathbf{y})$  by assuming a noisy relation  $\mathbf{y} = g(\mathbf{W}\mathbf{x}) + \epsilon$  and letting the noise go to zero. The main problem is that the limit is not taken properly care of and that the certainty of the result therefore is doubtful.

And finally in 2003, Teh, Welling, Osindero and Hinton presented in [8] the Energy Based Model on the filtering model  $\hat{\mathbf{s}} = \mathbf{W}\mathbf{x}$ . Through a setup using inspiration from physics, a model distribution for  $\mathbf{x}$  is constructed and adjusted to be as close to the data distribution as possible through a Kullback Leibler divergence. Again the authors are faced with an intractable integral and this time it is approximated with the so-called "n-step Learning" which is a Hybrid Monte Carlo technique using  $n$  steps in the estimation of the integral. The remarkable claim of the paper is that very few steps such as  $n = 1$ , or  $n = 3$  often will be sufficient to obtain overall convergence of the algorithm.

A common feature for most of the approaches is that they assume a noisy model and obtain a likelihood which involves a difficult integral which is then approximated in some way or another. The exception from this is the approach of Hagai Attias, but one can then argue that assuming priors to be mixtures of gaussians either is a restriction or an approximation.

## 1.1 This Paper

Thus, Independent Component Analysis (ICA) can be performed by a vast range of different methods. These can differ from each other by assumed properties such as noise or time-correlation, but also by the fundamental issue on whether they attempt to estimate the generative mixing matrix, denoted  $\mathbf{A}$ , or a filtering matrix denoted  $\mathbf{W}$ . In case of the same number of observations and sources, the square case, most if not all of these methods can be proven to be equivalent. But in the overcomplete case, where the number of observations are smaller than the number of sources, their differences becomes apparent and it is not easy to compare the results of generative and filtering approaches.

This paper makes an attempt to compare the result of two different methods in the overcomplete case: The Maximum Likelihood (ML), which estimates the generative  $\mathbf{A}$ , and the Energy Based Models (EBM) which seek to estimate a filtering matrix  $\mathbf{W}$ . This is done by assuming the priors to be centered mixtures of gaussians which makes it possible to compare the optimization schemes. This approach, with respect to ML, is closely related to the work of Hagai Attias in

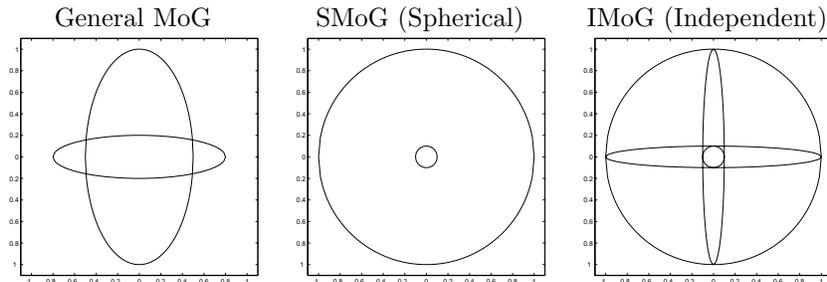


Figure 1: Contour plot of probability densities in 2D source space. and this is the only of the the three which in fact is independent in its variables, i.e. a Independent Mixture of Gaussians, IMoG.

[1], but where Attias is assuming completely general mixtures of gaussians and a noisy mixture model, the mixtures of gaussians in this paper are assumed centered for simplicity and, more importantly, the limit of zero noise is derived to be able to compare with the noiseless EBM.

The structure of the paper is as follows: In Sec 2, we introduce the reader to the mixture of gaussians to be used, and in Sec 3 we apply the mixtures of gaussians to the EBM and the ML. In Sec 4 we compare the results found and finally in Sec 5, we make a short summary.

## 2 Mixture of Gaussians

In this paper we make extensive use of the family of distributions known as mixtures of gaussians (MoG). To clear up some common misconceptions about MoG we introduce the general MoG and then discuss two important subsets of distributions. The density of a  $D$ -dimensional centered MoG is

$$p(\mathbf{s}) = \sum_{\kappa} \frac{\rho_{\kappa}}{\sqrt{|2\pi\mathbf{D}_{\kappa}|}} \exp \left[ -\frac{1}{2} \mathbf{s}^T \mathbf{D}_{\kappa}^{-1} \mathbf{s} \right] \quad (1)$$

where the weights  $\rho_{\kappa}$  sum to one. The matrix  $\mathbf{D}_{\kappa}$  can be any positive definite matrix, but in this context often assumed diagonal,  $\mathbf{D}_{\kappa} = \text{diag}(\sigma_{1\kappa}^2, \sigma_{2\kappa}^2, \dots, \sigma_{D\kappa}^2)$ . The marginal distribution for each  $s_i$  is itself a mixture of gaussians,  $s_i \sim \sum_{\kappa} \rho_{\kappa} \mathcal{N}(0, \sigma_{i\kappa}^2)$ , but note about the joint distribution that even if the coordinates  $s_i$  in each component are independent, i.e. if  $\mathbf{D}_{\kappa}$  all are diagonal, the coordinates  $s_i$  are *not* in general independent in the bigger joint distribution  $p(\mathbf{s})$ .

One subset of interest is the MoG constructed as a product of  $D$  1-dimensional mixtures of gaussians, and therefore independent by construction. Since the variables are independent we call this kind of MoG for Independent Mixtures of Gaussians (IMoG). Denoting the parameters of  $i$ 'th marginal by  $\sum_{\kappa} \rho_{i\kappa} \mathcal{N}(0, \sigma_{i\kappa}^2)$ , the density of the joint distribution is given by

$$p(\mathbf{s}) = \sum_{\{\mathbf{k}\}} \frac{\tilde{\rho}_{\mathbf{k}}}{\sqrt{|2\pi\mathbf{D}_{\mathbf{k}}|}} \exp \left[ -\frac{1}{2} \mathbf{s}^T \mathbf{D}_{\mathbf{k}}^{-1} \mathbf{s} \right] \quad (2)$$

The symbol  $\{\mathbf{k}\}$  denotes all combinations of the vector  $\mathbf{k}$  having length  $D$ , consisting of integers from 1 to  $K$ ,  $\mathbf{D}_{\mathbf{k}} = \text{diag}(\sigma_{1k_1}^2, \sigma_{2k_2}^2, \dots, \sigma_{Dk_D}^2)$ , and  $\tilde{\rho}_{\mathbf{k}}$  is defined by  $\tilde{\rho}_{\mathbf{k}} = \prod_{i=1}^D \rho_{ik_i}$ .

Another subset of interest is the MoG, where the matrices are not only diagonal, but all some constant times the identity  $\mathbf{D}_{\kappa} = \sigma_{\kappa}^2 \mathbf{I}$ . In this case the density is spherical symmetric and we call this kind of distributions for Spherical Mixtures of Gaussians (SMoG).

Obvious from Eq. 2, IMoG is itself a MoG, but the reverse is not true in general. As visualized on Figure 1, the density contour of a MoG with diagonal covariances is some weighted sum of axis-aligned ellipsoids. The length of axis of the ellipsoids corresponds to variances of gaussian components, i.e. the diagonal elements of the covariance matrices. This is also true for the product of 1-dimensional MoG, but in this case *all* combinations of the marginal variances are present as sets of axis of some ellipsoid. Thus knowing the density of a MoG with diagonal covariances, its variables are independent if and only if all combinations of axis lengths are present. Therefore, since this by definition is not the case for SMoG, SMoG and IMoG must be disjoint subsets of MoG.

### 3 Models and Derivation

We now use the MoG as source priors for two different approaches to ICA: Energy Based Models (EBM) and Maximum Likelihood (ML).

We consider a situation in which  $D$  sources  $\mathbf{s}_t$  are mixed into a set of  $M$  measurements  $\mathbf{x}_t$ , expressed by the equation  $\mathbf{x}_t = \mathbf{A}\mathbf{s}_t$ . The signals are  $N$  time steps long and can be arranged into the matrices  $\mathbf{S}$  and  $\mathbf{X}$ , such that mixing of all  $N$  vectors can be expressed in one equation  $\mathbf{X} = \mathbf{A}\mathbf{S}$ . The sources are white and since EBM is restricted to square and overcomplete mixing, we assume  $M \leq D$ .

#### 3.1 Energy Based Models

The EBM method, presented in [8], aim at demixing the measurements  $\mathbf{X}$  by a filtering with  $\mathbf{W}$  in the traditional way  $\hat{\mathbf{S}} = \mathbf{W}\mathbf{X}$ . That this, in the overcomplete case, is not producing independent estimated sources, is discussed in more detail in Sec 4. The filtering coefficients is determined through a construction of a model distribution  $p_{\mathbf{W}}(\mathbf{x})$  which is approximated to the data distribution

$$p_0(\mathbf{x}) = \frac{1}{N} \sum_{t=1}^N \delta(\mathbf{x} - \mathbf{x}_t) \quad (3)$$

through the Kullback-Leibler divergence. The model distribution is constructed in the following way: An energy is defined by  $E(\mathbf{x}; \mathbf{W}) = -\ln p_s(\mathbf{W}\mathbf{x})$ , which ensures the property that choosing  $\mathbf{W}$  such that the resulting estimated sources are not too unlikely, is rewarded through low energy levels and unlikely source values penalized with high energy levels. Other definitions of energy could be made, but this is especially appealing due to its calculational properties. The

energy  $E$  is used in a Gibbs distribution, i.e.

$$p_{\mathbf{W}}(\mathbf{x}) = \frac{e^{-E(\mathbf{x}; \mathbf{W})}}{Z(\mathbf{W})} = \frac{p_s(\mathbf{W}\mathbf{x})}{\int p_s(\mathbf{W}\mathbf{x}) d\mathbf{x}} \quad (4)$$

which is chosen as our model distribution. In [8], Teh et. al. make no assumption on the prior, which makes the normalization part more difficult. To deal with this they use so-called  $n$ -step learning, a variant of a hybrid monte carlo approach, do approximately estimate the normalization part of the optimization. In this paper we instead choose the prior to be a MoG, which enables us to calculate the integral of Eq. 4 and obtain an closed form expression for the model distribution. The result is

$$p_{\mathbf{W}}(\mathbf{x}) = \sum_{\kappa} \gamma_{\kappa} \frac{\exp(-\frac{1}{2}\mathbf{x}^T \mathbf{W}^T \mathbf{D}_{\kappa}^{-1} \mathbf{W} \mathbf{x})}{\sqrt{|2\pi(\mathbf{W}^T \mathbf{D}_{\kappa}^{-1} \mathbf{W})^{-1}|}} \quad (5)$$

where  $\gamma_{\kappa}$  in general is dependent on  $\mathbf{W}$  and  $\mathbf{D}_{\kappa'}$  for all  $\kappa'$  through the following relation: Setting  $\xi_{\kappa} = (|2\pi(\mathbf{W}^T \mathbf{D}_{\kappa}^{-1} \mathbf{W})^{-1}|/|2\pi\mathbf{D}_{\kappa}|)^{1/2}$ , we can write  $\gamma_{\kappa} = \rho_{\kappa} \xi_{\kappa} / \sum_{\kappa'} \rho_{\kappa'} \xi_{\kappa'}$ . Note that in the square case  $\xi_{\kappa} = 1$  and therefore  $\gamma_{\kappa} = \rho_{\kappa}$  and in the case of SMOG priors we obtain  $\xi_{\kappa} = (2\pi/\sigma_{\kappa}^2)^{(D-M)/2}$ . The estimated optimal filtering matrix  $\hat{\mathbf{W}}$  is determined by

$$\hat{\mathbf{W}} = \min_{\mathbf{W}} [KL(p_{\mathbf{W}} || p_0)]$$

in which the gradient of the KL-divergence can be calculated analytically when the prior is chosen to be mixtures of gaussians or some other analytically appealing distribution.

### 3.2 Maximum Likelihood

In the ML setup presented here we assume a generative model  $\mathbf{X} = \mathbf{A}\mathbf{S} + \mathbf{\Gamma}$ , where we, in order to be able to deal with the overcomplete case, have added white gaussian noise,  $\mathbf{\Gamma}$ . In the end we let the noise variance go to zero to obtain the noiseless result.

Assuming white gaussian noise and the prior on the sources  $p_s(\mathbf{s})$  to be a MoG with covariance matrices  $\mathbf{D}_{\kappa}$  and weights  $\rho_{\kappa}$ , we can complete the integration and write the distribution over  $\mathbf{x}$  as

$$p(\mathbf{x}) = \int p(\mathbf{x}|\mathbf{s})p(\mathbf{s})d\mathbf{s} = \sum_{\kappa} \rho_{\kappa} \frac{\sqrt{|2\pi\mathbf{\Phi}_{\kappa}^{-1}|} \exp(-\frac{1}{2}\mathbf{x}^T \mathbf{\Psi}_{\kappa} \mathbf{x})}{\sqrt{|2\pi\mathbf{\Sigma}|} \sqrt{|2\pi\mathbf{D}_{\kappa}|}}$$

where  $\mathbf{\Sigma} = \sigma^2 \mathbf{I}$  is the noise covariance matrix,  $\mathbf{\Phi}_{\kappa} = \mathbf{A}^T \mathbf{\Sigma}^{-1} \mathbf{A} + \mathbf{D}_{\kappa}^{-1}$  and  $\mathbf{\Psi}_{\kappa} = \mathbf{\Sigma}^{-1} - \mathbf{\Sigma}^{-1} \mathbf{A} \mathbf{\Phi}_{\kappa}^{-1} \mathbf{A}^T \mathbf{\Sigma}^{-1}$ . We now want to consider the limit of  $\sigma^2 \rightarrow 0$ , but we need to do this with great care, since otherwise crucial details will vanish in the approximation. Using the Woodbury identity, singular value decomposition and some very good approximations (see the appendix for details), we obtain the following limits

$$\begin{aligned} \mathbf{\Psi}_{\kappa} &\rightarrow (\mathbf{A} \mathbf{D}_{\kappa} \mathbf{A}^T)^{-1} \\ \sqrt{|2\pi\mathbf{\Phi}_{\kappa}^{-1}|} / \sqrt{|2\pi\mathbf{\Sigma}|} &\rightarrow w_{\kappa} / \sqrt{|\mathbf{A} \mathbf{A}^T|} \end{aligned}$$

Here  $w_\kappa$  is a constant with respect to  $\sigma^2$ , but depends on  $\mathbf{D}_\kappa$  and on  $\mathbf{A}$  through the unique orthogonal matrix  $\mathbf{V}$ .  $\mathbf{V}$  which fulfills the eigen value equation  $\mathbf{A}^T \mathbf{A} \mathbf{V} = \mathbf{V} \mathbf{\Lambda}$ , such that the eigen values in  $\mathbf{\Lambda}$  are decreasing in size down the diagonal. The expression for  $w_\kappa$  is  $w_\kappa = \prod_{i=M+1}^D (2\pi / (\mathbf{V}^T \mathbf{D}_\kappa^{-1} \mathbf{V})_{ii})^{1/2}$ . With this limit taken care of, we can write the maximum likelihood expression for  $p(\mathbf{x})$  as

$$p_{\mathbf{A}}(\mathbf{x}) = \sum_{\kappa} \rho_{\kappa} \frac{\exp(-\frac{1}{2} \mathbf{x}^T (\mathbf{A} \mathbf{D}_\kappa \mathbf{A}^T)^{-1} \mathbf{x})}{\sqrt{|2\pi \mathbf{A} \mathbf{D}_\kappa \mathbf{A}^T|}} \quad (6)$$

where the weights somewhat surprisingly turns out to be same as those of the prior (see the appendix for details). The estimated generative mixing matrix  $\hat{\mathbf{A}}$  is the matrix maximizing the log likelihood

$$\hat{\mathbf{A}} = \max_{\mathbf{A}} \left[ \ln P(\mathbf{X}|\mathbf{A}) \right]$$

Note that we have not estimated the sources in this process, only the generative mixing matrix.

## 4 Comparing EBM and ML

Now we compare the EBM and ML approaches derived in the previous section and discuss the significance of the differences and similarities. In the square case, we obtain total equivalence of all expressions setting  $\mathbf{W} = \mathbf{A}^{-1}$  and thus not surprisingly we can conclude, as Teh et .al. does in [8], that the two approaches are equivalent when the number of observations equal the number of sources. Therefore the discussion and comparison in this sections is almost entirely concerned with the overcomplete case.

In the overcomplete case it is not obvious how one should compare results on the generative  $\mathbf{A}$  and the filtering matrix  $\mathbf{W}$ . The filtering approach does not retrieve the original sources, since for any matrix  $\mathbf{W}$  we have  $\mathbf{W} \mathbf{A}_g \neq \mathbf{I}$  because of the dimensionality: It is impossible to construct  $D$   $M$ -dimensional orthogonal matrices when  $D > M$ . And we cannot in general compare the filter matrix  $\mathbf{W}$  with the pseudo-inverse of  $\mathbf{A}$ , since this is not the optimal solution in all cases [5]. But using MoG as prior, both the model distribution  $p_{\mathbf{W}}(\mathbf{x})$  of the EBM and the loglikelihood  $p_{\mathbf{A}}(\mathbf{x})$  of the ML becomes MoG's with parameters which must be estimated to fit a common data set  $\mathbf{X}$ . In fact we end up with two optimizations which look rather similar

$$0 = \frac{\partial}{\partial \mathbf{W}} \sum_{t=1}^N \ln p_{\mathbf{W}}(\mathbf{x}_t) \quad 0 = \frac{\partial}{\partial \mathbf{A}} \sum_{t=1}^N \ln p_{\mathbf{A}}(\mathbf{x}_t)$$

The similarity is to some extent both genuine and deceptive: Both distributions are MoG, but the dependency of the weights and covariances on  $\mathbf{W}$  and  $\mathbf{A}$  are different. In this section we compare EBM and ML by comparing the covariances and weights of their MoG's.

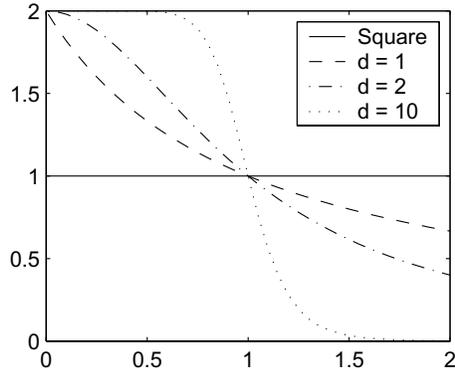


Figure 2: The Weights for EBM and ML. The plot demonstrates for SMOG priors that the differences in weights increase dramatically when the settings gets strongly overcomplete. (See the text for details on the plots)

#### 4.1 Spherical MoG

We now consider the special case when the priors are SMOG, i.e. the variances are given as  $\mathbf{D}_\kappa = \sigma_\kappa^2 \mathbf{I}$ . In this case the sources are not assumed to be independent, which is interesting in its on right, but it also serves as a clear example of properties which holds for the more general cases as well. When we assume the priors to be SMOG, the model distribution simplifies significantly,

$$p_{\mathbf{W}}(\mathbf{x}) = \sum_{\kappa} \gamma_{\kappa} \frac{\exp(-\frac{1}{2} \mathbf{x}^T \mathbf{W}^T \mathbf{D}_{\kappa}^{-1} \mathbf{W} \mathbf{x})}{\sqrt{|2\pi(\mathbf{W}^T \mathbf{D}_{\kappa}^{-1} \mathbf{W})^{-1}|}}, \quad \gamma_{\kappa} = \frac{\rho_{\kappa} (1/\sigma_{\kappa}^2)^d}{\sum_{\kappa'} \rho_{\kappa'} (1/\sigma_{\kappa'}^2)^d} \quad (7)$$

where  $d = (D - M)/2$ .

##### The Covariances

The covariances of EBM and ML respectively can in fact become equal in the case of SMOG. The equation setting the covariances equal

$$(\mathbf{W}^T \mathbf{D}_{\kappa}^{-1} \mathbf{W})^{-1} = \mathbf{A} \mathbf{D}_{\kappa} \mathbf{A}^T \quad \forall \kappa \quad (8)$$

translates into  $\mathbf{W}^T \mathbf{W} \mathbf{A} \mathbf{A}^T = \mathbf{I}$ , which is fulfilled for  $\mathbf{W} = \mathbf{A}^+$ , where  $\mathbf{A}^+$  denotes the pseudo inverse (Moore-Penrose) of the matrix  $\mathbf{A}$ . When  $\mathbf{A}$  has full rank, the pseudo-inverse is given by  $\mathbf{A}^+ = \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^{-1}$ . But the equation is also fulfilled for any matrix  $\mathbf{W} = \mathbf{U} \mathbf{A}^+$ , where  $\mathbf{U}$  is orthogonal and thus, there is an entire family of matrices which would make the covariances of the EBM equal to the covariances of the ML for a given  $\mathbf{A}$ . Conversely for any  $\mathbf{W}$  we can choose  $\mathbf{A} = \mathbf{W}^+$  to obtain the same covariances and in this sense the two approaches have equal flexibility with respect to adjusting the covariances to the data. This is evident in the  $4 \times 2$  example in Fig 3 a) and b).

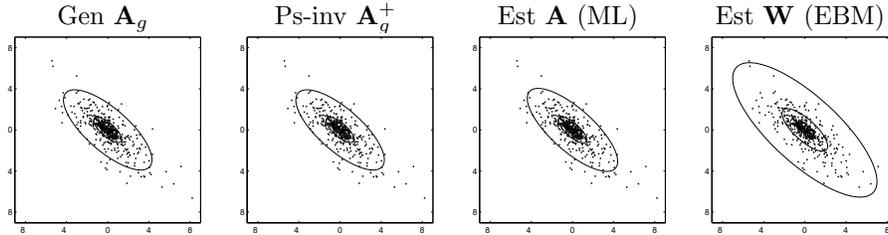


Figure 3: The covariances in case of SMOG priors for different settings. Plot a) Generative  $\mathbf{A}$ . Plot b) Pseudo-inverse of the generative  $\mathbf{A}$ . Plot c) Estimated  $\mathbf{A}$ . Plot d) Estimated  $\mathbf{W}$ .

### The Weights

The weights can be compared by examining the ratio  $\gamma_\kappa/\rho_\kappa$  and as we shall see, this ratio differs strongly from 1 in most cases. From the expression of  $\gamma_\kappa$  in Eq. 7, we see that the constrain making the weights of ML and EBM equal is

$$\frac{(1/\sigma_\kappa^2)^d}{\sum_{\kappa'} \rho_{\kappa'} (1/\sigma_{\kappa'}^2)^d} \stackrel{?}{=} 1 \quad \forall \kappa$$

which is clearly impossible when the weights  $\sigma_\kappa^2$  must be different for different  $\kappa$ . Thus, in the SMOG case, the weights of EBM and ML cannot be equal and furthermore the ratio  $\gamma_\kappa/\rho_\kappa$  becomes relatively large for those  $\kappa$  where  $\sigma_\kappa^2$  is very small and vice versa.

Fig 2 is a general illustration of this. Here we assume a SMOG prior with two components which has  $\sigma_1^2 = 1$  and  $\rho_1 = \rho_2 = 0.5$  and let  $\sigma_2^2$  vary between 0 and 2. The resulting ratio  $\gamma_2/\rho_2$  is shown in Fig 2: The  $x$ -axis is  $\sigma_2^2$ , the  $y$ -axis is  $\gamma_2/\rho_2$  and the 4 curves are plotted for different degrees of overcompleteness  $d=0$  (Square) and  $d = 1, 2, 10$  (Overcomplete). Clearly only in the square case and for  $\sigma_2^2$  values close to 1, is the ratio reasonably close to 1.

Another more specific example is shown in Fig 3 c) and d), which is the estimated covariances for ML and EBM in a  $4 \times 2$  case. In this example the effect of the enhanced weight on smaller covariances is clear: When the weights of the smaller covariance is strong, the points far from origin is considered extreme, and the covariances are expanded accordingly. Thus, for this reason, EBM seem to favor larger covariances compared to ML.

## 5 Summary and Acknowledgements

Conclusively the use of Mixtures of Gaussians made it possible to compare Maximum Likelihood with Energy Based Models. The results show that in the overcomplete case with Spherical Mixtures of Gaussians as priors, the Energy Based Model is biased toward larger covariances compared to the Maximum Likelihood. One can show that this effect is also present for IMoG priors, though not at all as strong.

Finally I need to give due credit: This paper is closely related to the result of earlier work together with Jiucang Hao and Te-Won Lee from University of California San Diego (UCSD).

## A Details of the Calculations

When calculating the limit of  $\Psi_\kappa$ , we first set in the definition  $\Sigma = \sigma^2 \mathbf{I}$  and simplify the expression into

$$\Psi_\kappa = \frac{1}{\sigma^2} \{ \mathbf{I} - \mathbf{A}(\mathbf{A}^T \mathbf{A} + \sigma^2 \mathbf{D}_\kappa^{-1}) \mathbf{A}^T \}$$

Now defining  $\mathbf{Q} = \mathbf{A} \mathbf{D}_\kappa \mathbf{A}^T / \sigma^2$  we can use Woodburys identity for inverse matrices and obtain

$$\mathbf{A}(\mathbf{A}^T \mathbf{A} + \sigma^2 \mathbf{D}_\kappa^{-1}) \mathbf{A}^T = \mathbf{Q} - \mathbf{Q}(\mathbf{I} + \mathbf{Q})^{-1} \mathbf{Q}$$

Since  $\mathbf{Q}$  is symmetric and very large compared to  $\mathbf{I}$ , the right hand side can be approximated by  $\mathbf{I} - \mathbf{Q}^{-1}$ . To see this, write  $\mathbf{Q}$  as  $\mathbf{Q} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$ , for an orthogonal  $\mathbf{V}$  and diagonal  $\mathbf{\Lambda}$  and remember the identity  $x - x^2/(1+x) = 1 - 1/(1+x)$  to obtain  $\mathbf{I} - \mathbf{V} \text{diag}(1/(1 + \Lambda_{ii}/\sigma^2)) \mathbf{V}^T \cong \mathbf{I} - \mathbf{V} \text{diag}(1/(\Lambda_{ii}/\sigma^2)) \mathbf{V}^T = \mathbf{I} - \mathbf{Q}^{-1}$ . Inserting this into the equation containing  $\Psi_\kappa$ , gives the desired result.

When calculating the limit of the fraction containing the determinant  $|2\pi \Phi_\kappa^{-1}|$ , we use the fact that since  $\mathbf{A}^T \mathbf{A}$  is symmetric there exists an orthogonal  $\mathbf{V}$  such that  $\mathbf{A}^T \mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$ . Since  $|\mathbf{V}| = 1$  we get

$$|2\pi \Phi_\kappa^{-1}| / |2\pi \Sigma| = (2\pi \sigma^2)^{(D-M)} / |\mathbf{\Lambda} + \sigma^2 \mathbf{V}^T \mathbf{D}^{-1} \mathbf{V}|$$

Since  $\sigma^2$  are assumed arbitrary small, only the diagonal of the sum of matrices will contribute significantly to the determinant. And since further more  $\mathbf{A}^T \mathbf{A}$  has rank  $M$ , the  $M$  first elements of the diagonal matrix will dominate together with the remaining  $M - D$  factors

$$|\mathbf{\Lambda} + \sigma^2 \mathbf{V}^T \mathbf{D}^{-1} \mathbf{V}| \cong \prod_{i=1}^M \Lambda_{ii} \prod_{j=M+1}^D (\sigma^2 \mathbf{V}^T \mathbf{D}^{-1} \mathbf{V})_{jj}$$

inserting this into the fraction above gives the desired result.

The coefficients  $\alpha_\kappa$  has the structure  $\rho_\kappa w_\kappa \sqrt{|2\pi \mathbf{A} \mathbf{D}_\kappa \mathbf{A}^T|} / \sqrt{|\mathbf{A} \mathbf{A}^T| \cdot |2\pi \mathbf{D}_\kappa|}$ . In the square case much of the difficulties of taking the noiseless limit disappears,  $w_\kappa = 1$  and we easily obtain  $\alpha_\kappa = \rho_\kappa$ . In the case of SMOG, setting  $\mathbf{D}_\kappa = \sigma_\kappa^2 \mathbf{I}$ , we obtain  $w_\kappa = (2\pi \sigma_\kappa^2)^{(D-M)/2}$  and therefore  $\alpha_\kappa = \rho_\kappa$ . Supported by numerical results we conjecture that this is also the case for the general overcomplete case.

## References

- [1] H. Attias, *Independent Factor Analysis*, Neural Computation 11, 803-851, 1999.
- [2] M. Girolami, *A Variational Method for Learning Sparse and Overcomplete Representations*, Neural Computation, 13(11), pp 2517 - 2532, 2001.

- [3] G. E. Hinton, *Training products of Experts by minimizing contrastive divergence*, Neural Computation, 14(8):1771-1800, 2002.
- [4] P. Hoyer-Sorensen, O. Winther, L. K. Hansen, *Mean-Field Approaches to Independent Component Analysis*, Neural Computation Volume 14, Issue 4, April 2002.
- [5] M. S. Lewicki, T. J. Sejnowski, *Learning Overcomplete Representations*, Neural Computation, 12(2):337-65, 2000.
- [6] R. M. Neal, *Probabilistic Inference Using Markov Chain Monte Carlo Methods*, Technical Report CRG-TR-93-I, Department of Computer Science, University of Toronto, 1993.
- [7] O. Shriki, H. Sampolinski, D.D. Lee, *An information maximization approach to overcomplete and recurrent representations* Neural Info. Proc. Sys. 13 (2001).
- [8] Y. W. Teh, M. Welling, S. Osindero, G. Hinton *"Energy-Based Models for Sparse Overcomplete Representations"*, The Journal of Machine Learning Research - Special Issue on Independent Components Analysis, guest edited by Te-Won Lee, Jean-Francois Cardoso, Erkki Oja and Shun-ichi Amari. 2003.

# Stochastic differential equations in image warping

Bo Markussen

Department of Computer Science

University of Copenhagen, DK-2100 Copenhagen

email: boma@diku.dk

## Abstract

This paper is concerned with image warping as applied in *e.g.* non-rigid registration. We use the theory of stochastic differential equations to give a precise mathematical description of the Brownian warps model proposed by Nielsen et al.. Furthermore we introduce a renormalization technique allowing for Bayesian inference of the warp  $\phi$  given *e.g.* the observation of matching landmarks in two images. As a final result we show that the maximum *a posteriori* estimator is equivalent to the minimal energy estimator derived by Joshi & Miller using  $\phi$  ws given by deterministic differential equations.

## 1 Introduction

Image warping is diffeomorphic transformations of the domain on which an image is represented. A typical application is the registration problem. Suppose for instance that several images of the same object have been taken with different devices. Each device records different properties of the object, and we wish to combine the different images a joint description of the object. In practice the images are often not aligned, whence it would be improper simply to stack them on top of each other. However, if we are allowed to deform the images (imagine the images being recorded of pieces of rubber, and that we may deform the pieces of rubber in any possible way without introducing folds or tears), then it might be possible to achieve alignment. In some cases the needed transformation are non-rigid leading to a non-rigid registration problem. This paper is concerned with a probabilistic description of this problem. In particular we consider the statistical problem of finding the maximum *a posteriori* estimator of the warps given a set of matching landmarks. Landmarks are points which are identifiable in all the images. We hence know that these landmarks should be matched by the

warps. Given this information we seek to estimate the remaining registration. A related application would be the estimation of the image  $o \circ w$  given a sequence of matched landmarks in a sequence of images.

This will be a purely theoretical paper. The results are threefold: 1) A stochastic differential equation description of the least committed prior for Brownian warps proposed in [7], 2) the introduction of a mathematical rigorous renormalization technique allowing the computation of a maximum *a posteriori* estimator, 3) equivalence of the maximum *a posteriori* and the minimum energy warp found in [4]. These results have recently been presented in the paper [1] at the second workshop on *Generative Model Based Vision* at the IEEE conference on *Computer Vision and Pattern Recognition*.

Mathematically a warp between two  $d$ -dimensional images (in applications we typically have  $d = 2$  or  $d = 3$ ) is a diffeomorphism  $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ . Points  $x \in \mathbb{R}^d$  are usually understood as column vectors  $x = \{x^i\}_{i=1,\dots,d}$ . Similarly, the coordinates of a matrix  $a \in \mathbb{R}^{d \times d}$  are denoted by  $a = \{a^{ij}\}_{i,j=1,\dots,d}$ . The transpose of a vector/matrix  $A$  is denoted by  $A^*$ , probabilistic expectation is denoted by  $\mathbb{E}$ , and the set of natural numbers is denoted by  $\mathbb{N}$ .

## 2 The first principle model

In accordance with the physical ‘‘piece of rubber’’ analogy it is natural to realize the warping as a dynamical process. At time  $t = 0$  we have the original image, and at time  $t = T$  we have the final image. At a time point  $0 < t < T$  we have an intermediate warp. We thus introduce a double indexed family of warps  $\phi(\cdot, s, t): \mathbb{R}^d \rightarrow \mathbb{R}^d$  with  $0 \leq s \leq t \leq T$ . The warp  $\phi(\cdot, s, t)$  gives the warping between time  $s$  and  $t$ . The physical interpretation implies the following properties

$$\phi(x, s, s) = x, \quad \phi(x, s, t) = \phi(\phi(x, s, r), r, t).$$

The next step is to realize  $\phi(\cdot, s, t)$  as a stochastic process. The axiom stated below was proposed in [7] and parallels the assumptions leading to Gaussian distribution and the Brownian motion. A Brownian motion is a stochastic process  $B$  with independent Gaussian distributed increments  $B(t) - B(s) \sim \mathcal{N}(0, t - s)$ . The differential  $dB(t)$  can be interpreted as Gaussian white noise of order  $\sqrt{dt}$ .

**Axiom.** *The diffeomorphisms  $\phi(\cdot, t_{h-1}, t_h)$  are stochastically independent for every partition  $0 = t_0 < t_1 < \dots < t_H = T$ , and the distribution of  $\phi(\cdot, s, t)$  depends only on the time length  $t - s$ .*

The Axiom together with the following property imply a central limit theorem for the Jacobians  $\{\partial\phi^i(x)/\partial x^j\}_{i,j=1,\dots,d}$ . These Jacobians were employed in [7] for

the landmark matching problem. Unfortunately the limiting distribution is very difficult to analyze, and has to our knowledge only been described in dimension  $d = 2$  in [3]. The underlying probabilistic model can be rigorously defined using stochastic calculus. The physical interpretation of Eq. (3) given below is that the “piece of rubber” is influenced by infinitely many random forces modelled by Brownian motions  $B_n$ ,  $n \in \mathbb{N}$ . The  $n$ 'th force works in the direction  $f_n(x)$  at the spatial position  $x$  and has size  $dB_n(t)$  at time  $t$ . We will use stochastic integrals such as  $\int_0^T X(t) dB(t)$ , which loosely speaking is the limit of the incremental sums

$$\sum_{h=1}^H X(t_{h-1}) (B(t_h) - B(t_{h-1})).$$

However, some care has to be taken. Especially there is a distinction between the so-called Itô and Stratonovich integrals. The essential reference for the results employed in this paper is the monograph [5]. But be warned, this monograph can be quite hard to read. Let  $\circ dB_n(t)$  denote the Stratonovich differential at time  $t$ .

**Theorem 1.** *In addition to the Axiom assume for every  $x, y \in \mathbb{R}^d$  and  $s \in [0, T]$  the existence of the limits*

$$\begin{aligned} & \frac{\mathbb{E}[\phi(x, s, t) - x]}{t - s} \xrightarrow[t \rightarrow s]{t > s} b(x), \\ & \frac{\mathbb{E}[(\phi(x, s, t) - x)(\phi(y, s, t) - y)^*]}{t - s} \xrightarrow[t \rightarrow s]{t > s} a(x, y) \end{aligned} \quad (1)$$

and functions  $f_0$  and  $f_n$ ,  $n \in \mathbb{N}$ , such that

$$\begin{aligned} a(x, y) &= \sum_{n \in \mathbb{N}} f_n(x) f_n(y)^* \in \mathbb{R}^{d \times d}, \\ b(x) &= f_0(x) + \frac{1}{2} \left\{ \sum_{j=1}^d \frac{\partial a^{ij}(x, y)}{\partial x^j} \Big|_{y=x} \right\}_{i=1, \dots, d} \in \mathbb{R}^d. \end{aligned} \quad (2)$$

*Under some additional regularity conditions there exists stochastically independent Brownian motions  $B_n$ ,  $n \in \mathbb{N}$ , such that  $\phi(x, s, t)$  solves the stochastic differential equation given on integral form by*

$$\phi(x, s, t) = x + \int_s^t f_0(\phi(x, s, u)) du + \sum_{n \in \mathbb{N}} \int_s^t f_n(\phi(x, s, u)) \circ dB_n(u). \quad (3)$$

Note that the statistical properties of the flow of diffeomorphisms  $\phi(x, s, t)$  are encoded by the infinitesimal mean and covariance  $b(x)$  and  $a(x, y)$ , and that the actual appearance of the flow is encoded by the Brownian motions  $B_n$ . In remaining of this paper the functions  $b(x)$  and  $a(x, y)$  are assumed to be known a priori.

### 3 The maximum a posteriori estimator

The idea behind the regularized maximum *a posteriori* estimator is to approximate the finite dimensional Brownian motion  $\mathbf{B} = \{B_n(t)\}_{n \in \mathbb{N}, t \in [0, T]}$  by the finite dimensional random variable

$$\mathbf{B}^N = \left\{ B_n \left( \frac{mT}{N} \right) \right\}_{n, m=1, \dots, N} \in \mathbb{R}^{N \times N}.$$

The  $N^2$ -dimensional Gaussian vector  $\mathbf{B}^N$  has probability density

$$p_N(\mathbf{B}^N) = \prod_{n=1}^N \prod_{m=1}^N \sqrt{\frac{N}{2\pi T}} \exp \left( -\frac{N(B_n(t_m^N) - B_n(t_{m-1}^N))^2}{2T} \right).$$

For  $\tilde{\mathbf{B}} \in \mathcal{C}^1([0, T]; \mathbb{R}^{\mathbb{N}})$  we define the renormalized Brownian density as the limit

$$\tilde{p}(\tilde{\mathbf{B}}) = \lim_{N \rightarrow \infty} p_N(\tilde{\mathbf{B}}^N) = \exp \left( -\frac{1}{2} \sum_{n \in \mathbb{N}} \int_0^T \left( \frac{\partial \tilde{B}_n(t)}{\partial t} \right)^2 dt \right). \quad (4)$$

Let  $(x_l, y_l) \in \mathbb{R}^d \times \mathbb{R}^d$ ,  $l \in L$ , be a given set of landmarks. Let the approximative flow  $\phi_N(\cdot, t_m, t_k)$  with  $t_m = \frac{mT}{N}$  and the approximative maximum *a posteriori* estimator  $\hat{\mathbf{B}}^N$  be defined by

$$\phi_N(x, t_m, t_k) = x + \sum_{h=m}^{k-1} \frac{f_0(\phi_N(x, t_m, t_h)) T}{N} + \sum_{n=1}^N \sum_{h=m}^{k-1} f_n(\phi_N(x, t_m, t_h)) (B_n(t_{h+1}) - B_n(t_h)),$$

$$\hat{\mathbf{B}}^N = \arg \max \left\{ p_N(\mathbf{B}^N) \mid \mathbf{B}^N \in \mathbb{R}^{N \times N} : \phi_N(x_l, 0, T) = y_l \text{ for } l \in L \right\},$$

the renormalized flow  $\tilde{\phi}(x, s, t)$  and the renormalized maximum *a posteriori* estimator  $\tilde{\mathbf{B}}$  be defined by

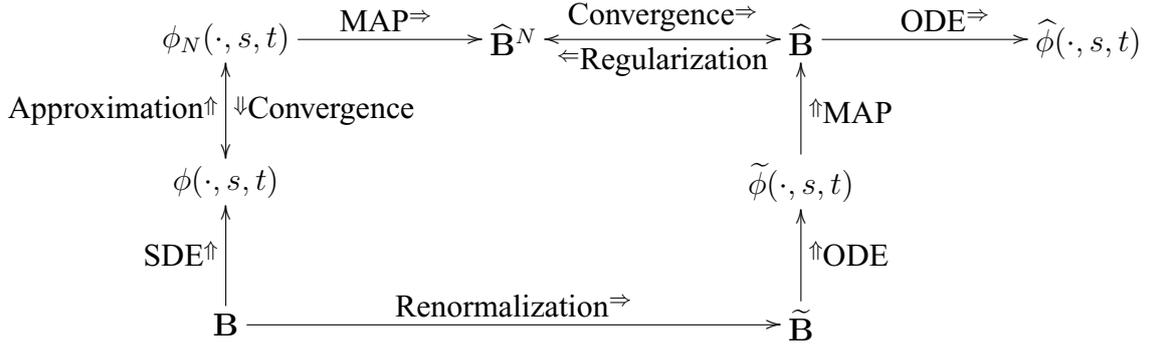
$$\tilde{\phi}(x, s, t) = x + \int_s^t f_0(\tilde{\phi}(x, s, u)) du + \sum_{n \in \mathbb{N}} \int_s^t f_n(\tilde{\phi}(x, s, u)) d\tilde{B}_n(u),$$

$$\tilde{\mathbf{B}} = \arg \max \left\{ \tilde{p}(\tilde{\mathbf{B}}) \mid \tilde{\mathbf{B}} \in \mathcal{C}^1([0, T]; \mathbb{R}^{\mathbb{N}}) : \tilde{\phi}(x_l, 0, T) = y_l \text{ for } l \in L \right\},$$

and the maximum *a posteriori* estimator  $\hat{\phi}(\cdot, s, t)$  be defined by the ordinary differential equation

$$\hat{\phi}(x, s, t) = x + \int_s^t f_0(\hat{\phi}(x, s, u)) du + \sum_{n \in \mathbb{N}} \int_s^t f_n(\hat{\phi}(x, s, u)) d\hat{B}_n(u). \quad (5)$$

Then it can be show that  $\phi_N(x, s, t)$  converges to  $\phi(x, s, t)$  and that  $\widehat{\mathbf{B}}^N$  converges to  $\widehat{\mathbf{B}}$  as  $N \rightarrow \infty$ . These convergence results indeed justify to call  $\widehat{\phi}(\cdot, s, t)$  the maximum *a posteriori* estimator. The main novelty of the paper [1] was the awareness of need to consider the convergence of  $\widehat{\mathbf{B}}^N$  to  $\widehat{\mathbf{B}}$ . The relations between these properties are gathered in the following diagram



Observe that discretizations done in this section only are an intermediate step in the theoretical analysis. However, the differential equation for the  $\widetilde{\phi}(\cdot, s, t)$  will typically be discretized for the numerical implementation. Moreover, working directly with the covariance function  $a(x, y)$  instead of the functions  $f_n(x)$  we can circumvent the infinite sum over  $n \in \mathbb{N}$  in Eq. (5). For these results, a specific algorithm for the landmark matching problem and a discussion of the choice of the functions  $b(x)$  and  $a(x, y)$  see [1].

## 4 Interpretation of energy minimization approaches

The theory of differential geometry gives that there for every diffeomorphism  $\phi$  exists a velocity field  $v: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$  such that  $\phi(x) = \phi(x, 0, T)$  for the solution to the transport equation

$$\phi(x, s, t) = x + \int_s^t v(\phi(x, s, u), u) \, du. \quad (6)$$

An alternative approach for warp estimation is to define some energy functional over velocity fields and choose the minimal energy velocity field fitting the landmarks, *cf.* [4] and [2]. Given some differential operator  $\mathcal{L}$ , *e.g.* the Laplacian, the paper [4] proposes to use the energy functional

$$E(v) = \int_0^T \int_{\mathbb{R}^d} \|\mathcal{L} \tilde{v}(x, t)\|^2 \, dx \, dt.$$

For vanishing  $f_0(x)$  the velocity field corresponding to the landmarks path  $\{\tilde{\phi}(x_l, 0, t)\}_{l \in L}$  is given by

$$\tilde{v}(x, t) = \left\{ a(x, \tilde{\phi}(x_l, 0, t)) \right\}_{l \in L} \left\{ a(\tilde{\phi}(x_k, 0, t), \tilde{\phi}(x_l, 0, t)) \right\}_{k, l \in L}^{-1} \left\{ \frac{\partial \tilde{\phi}(x_k, 0, t)}{\partial t} \right\}_{k \in L}.$$

If the infinitesimal covariance  $a(x, y)$  equals the Greens function for the square of the differential operator  $\mathcal{L}$ , i.e.  $\mathcal{L}^* \mathcal{L} a(x, y) = \delta_{x=y}$ , then

$$\mathcal{L}^* \mathcal{L} \tilde{v}(x, t) = \left\{ \delta_{x=\tilde{\phi}(x_l, 0, t)} \right\}_{l \in L} \left\{ a(\tilde{\phi}(x_k, 0, t), \tilde{\phi}(x_l, 0, t)) \right\}_{k, l \in L}^{-1} \left\{ \frac{\partial \tilde{\phi}(x_k, 0, t)}{\partial t} \right\}_{k \in L}$$

and it is shown in [1] that the energy functional  $E(\tilde{v})$  equals minus the logarithm of the renormalized Brownian density for  $\tilde{\mathbf{B}}$ . In this way the maximum *a posteriori* estimator equals the minimal energy estimator. This result gives 1) a probabilistic interpretation of the energy minimization methods proposed in [4] and [2], 2) the “statistical equivalence” of the energy minimization methods to the Bayesian approach proposed in [7].

**Acknowledgement:** I am grateful to Peter Johansen and Mads Nielsen for introducing me to these intriguing problems.

## References

- [1] B. Markussen, “A Statistical Approach to Large Deformation Diffeomorphisms”, second workshop on Generative Model Based Vision at the 22nd IEEE conference on Computer Vision and Pattern Recognition, Washington D.C., 2004.
- [2] V. Camion and L. Younes, “Geodesic Interpolating Splines,” Third International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition, *EMMCVPR 2001*, LNCS 2134, pp. 513-527, 2001.
- [3] A. D. Jackson, B. Lautrup, P. Johansen, and M. Nielsen, “Products of Random Matrices,” *Phys. Rev. E*, Vol. 66, article 66124, 2002.
- [4] S. C. Joshi and M. I. Miller, “Landmark Matching via Large Deformation Diffeomorphisms,” *IEEE Trans. Image Process*, Vol. 9, pp. 1357-1370, 2000.
- [5] H. Kunita, *Stochastic Flows and Stochastic Differential Equations*, Cambridge University Press, 1990.

- [6] M. I. Miller and L. Younes, "Group Actions, Homeomorphisms, and Matching: A General Framework," *Int. J. Computer Vision*, Vol. 41, pp. 61–84, 2002.
- [7] M. Nielsen, P. Johansen, A. D. Jackson, and B. Laurrup, "Brownian Warps: A Least Committed Prior for Non-rigid Registration," Fifth International Conference on Medical Image Computing and Computer-Assisted Intervention, *MICCAI 2002*, LNCS 2489, pp. 557-564, 2002.
- [8] C. J. Twining and S. Marsland, "Constructing Diffeomorphic Representations of Non-rigid Registrations of Medical Images," 18'th International Conference on Information Processing in Medical Imaging, *IPMI 2003*, LNCS 2732, pp. 413-425, 2003.

# Probabilistic Model-based Background Subtraction

J. Anderson<sup>1</sup>, T. Prehn<sup>1</sup>, V. Krüger<sup>1</sup>, and A. Elgammal<sup>2</sup>

<sup>1</sup> Aalborg University Esbjerg  
Niels Bohrs Vej 8  
6700 Esbjerg  
Denmark

<sup>2</sup> Dept. of Computer Science  
Rutgers, the State University of New Jersey  
110 Frelinghuysen Road  
Piscataway, NJ 08854-8019  
USA

**Abstract.** Usually, background subtraction is approached as a pixel-based process, and the output is (a possibly thresholded) image where each pixel reflects, independent from its neighboring pixels, the likelihood of itself belonging to a foreground object. What is neglected for better output is the correlation between pixels. In this paper we introduce a model-based background subtraction approach which facilitates prior knowledge of pixel correlations for clearer and better results. Model knowledge is being learned from good training video data, the data is stored for fast access in a hierarchical manner. Bayesian propagation over time is used for proper model selection and tracking during model-based background subtraction. Bayes propagation is attractive in our application as it allows to deal with uncertainties during tracking. The system runs in real-time and was extensively tested on suitable outdoor video data.

## 1 Introduction

Companies and scientists work on vision systems that are expected to work in real-world scenarios. Car companies work, e.g., on road sign and pedestrian detection and due to the threat of terrorism, biometric vision systems and surveillance applications are under development.

All these approaches work well in controlled environments, e.g., attempts to recognize humans by their face and gait has proven to be very successful in a lab environment. However, in uncontrolled environments, such as outdoor scenarios, the approaches disgracefully fail, e.g., the gait recognition drops from 99% to merely 30%. This is mainly due to low quality video data, the often small number of pixels on target and visual distractors such as shadows and strong illumination variations.

What is needed are special feature extraction techniques that are robust to outdoor distortions and that can cope with low-quality video data. One of the

most common feature extraction techniques in surveillance applications is background subtraction (BGS) [5; 3; 9; 6]. BGS approaches assume a stable camera. They are able to learn a background as well as possible local image variations of it, thus generating a background model even of non-rigid background objects. During application the model is compared with novel video images and pixels are marked according to the belief that they are fitting the background model. Generally, BGS methods have the following drawbacks:

1. BGS techniques are able to detect “interesting” images areas, i.e., image areas that are sufficiently different from the learned background model. Thus, BGS approaches are able to detect, e.g., a person and the shadow that he/she throws. However, BGS approaches are not able to distinguish between a foreground object and its shadow.
2. Very often, the same objects causes a different output when the scenario changes: E.g. the BGS output for a person walking on green grass or gray concrete may be different.

In this paper we present a Model-based Background Subtracting (MBGS) method that learns not only a background model but also a foreground model. The “classical” background subtraction detects the region of interests while the foreground models are being applied to the classical BGS output to “clean up” possible noise. To reach a maximum of robustness, we omit any thresholding and apply statistical matching techniques to the likelihood measures of the BGS.

Having the recent gait recognition attempts in mind, we have applied the following limitations to our discussion (however, the methodology is general enough that we do not see any limit of generality in these limitations):

1. We consider only humans as objects and ignore objects that look different from humans.
2. We limit our discussion to silhouettes of humans as they deliver a fairly clothing-independent description of an individual.

The Model-based Background Subtraction System (MBGS System) consists of a learning part to learn possible foreground objects and a MBGS part, where the output of a classical BGS is verified using the previously trained foreground object knowledge.

To learn and represent the foreground knowledge (here silhouettes of humans) is not straight forward due to the absence of a suitable vector space. One possibility is to describe the data in a hierarchical manner, using a suitable metric and a suitable representation of dynamics between the silhouettes. We interpret the silhouettes as probabilistic density functions  $S(X, Y)$  where each pixel describes the likelihood of belonging either to foreground or to background. To compare these “probabilistic” silhouettes we use the Kullback-Leibler distance, k-means clustering is used to cluster similar ones. Similar approaches for hierarchical contour representation can be found in [4; 14].

In the second part, we again consider the contours as densities over spatial coordinates and use normalized correlation to compute the similarity between

the silhouette density and the computed one in the input image. Tracking and silhouette selection is being done using Bayesian propagation over time. Bayesian propagation can be applied directly since we are dealing with densities and it has the advantage that it considers the uncertainty in the estimation of the tracking parameters and the silhouette selection. The densities in the Bayesian propagation are approximated using an enhancement of the well-known Condensation method [7]. A similar enhancement of the Condensation method has been applied in video based face recognition [12].

The remainder of this paper is organized as follows: In Sec. 2 we introduce the learning approaches. The actual BGS method is discussed in Sec. 3. We discuss our experimental results in Sec. 4 and conclude with final remarks in Sec. 5.

## 2 Learning and Representation of Foreground Objects

In order to make use of foreground model knowledge, the MBGS system needs to be able to:

- learn a model representation for possibly a number of different and non-rigid objects from video data and
- quickly access the proper model information during application of the MBGS.

Our main idea is the following: Apply the classical BGS to a scenario that is controlled in a manner that facilitates the learning process. In our case, since we want to learn silhouettes of humans, that only humans are visible in the scene during training and that the background variations are kept as small as possible to minimize distortions. Then, use this video data to learn the proper model knowledge.

After the application of a classical BGS, applying mean-shift tracking [1] allows to extract from the BGS output-data a sequence of small image patches containing, centered, the silhouette. This procedure is the same as the one presented in [10], however, with the difference that here we do not threshold the BGS output but use probabilistic silhouettes (instead of binary ones as in [10]) which still contain for each silhouette pixel the belief of being a foreground pixel.

To organize this data we first normalize the exemplars w.r.t scale and position. Then we use, similar to [4], a combination of tree structuring and k-means clustering. We use a top down approach: The first level is the root of the hierarchy which contains all the exemplars. Then the second level is constructed by using a the k-means clustering to cluster the exemplars from the root. The third level is constructed by clustering each cluster from the second level, again, using k-means, see Fig. 1 for an example. The k-means clustering uses the Kullback-Leibler divergence measure which measures the similarity between two density functions  $p$  and  $q$ :

$$\text{KLDist}(p, q) = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x} . \quad (1)$$

$\text{KLDist}(p, q)$  is non-negative and only zero if  $p$  and  $q$  coincide.



**Fig. 1.** An example of our clustering approach: 30 exemplars with  $K=3$  and the algorithm stops after reaching 3 levels.

See Fig. 2 for an example clustering of a training sequence of a single individual (as training data, we chose here non-optimal data for better visualization). The tree structure facilitates a fast search of exemplars along the tree vertices,



**Fig. 2.** The five images show the cluster centers computed from a training sequence of a single individual.

and the cluster centers are either used to apply MBGS on a coarse level or they are used as proto-types, as in [4], to direct the search to a finer level in the hierarchy. Once the tree is constructed, we generate a Markov transition matrix: Assuming that the change over time from one silhouette to a next one can be understood as a first order Markov process, the Markov transition matrix  $M_{ij}^l$  describes the transition probability of silhouette  $s_j$  following after silhouette  $s_i$  at level  $l$  in the hierarchy. During MBGS application particle filtering [8; 11; 2] will be used to find the proper silhouette (see Sec. 3). The propagation of silhouettes over time (see Sec. 3) is non-trivial, as silhouette do not form a vector space. However, what is sufficient, is a (not necessarily symmetric) metric space, i.e., given a silhouette  $s_i$ , only silhouettes that are close according to a given (not necessarily symmetric) metric are needed to be considered. In the tree structure similar contours are clustered which facilitates the propagation process. The Markov transition matrix  $M_{ij}$  on the other hand describes directly the transition likelihoods between clusters.

### 3 Applying Background Subtraction and Recognizing Foreground Objects

The MBGS system is built as an extension to a pixel based BGS approach. It uses foreground models to define likely correlations between neighbored pixels in the output  $P(\mathbf{x})$  of the BGS application.

Each pixel in the image  $P(\mathbf{x})$  contains a value in the range  $[0, 1]$ , where 1 indicates the highest probability of a pixel being a foreground pixel. A model in the hierarchy can be chosen and deformed according to a 4-D vector

$$\theta = [i, s, x, y], \quad (2)$$

where  $x$  and  $y$  denote the position of the silhouette in the image  $P$ ,  $s$  its scale, and  $i$  is a natural number that refers to a silhouette in the hierarchy.

The “matching” is done by normalized correlation between a model silhouette density, parameterized according to a deformation vector  $\theta_t$  and the appropriate region of interest in the BGS image  $P_t(\mathbf{x})$ , appropriately normalized.

In order to find at each time-step  $t$  the most likely  $\theta_t$  in the image  $P_t(\mathbf{x})$ , we use Bayesian propagation over time

$$\begin{aligned} p(\theta_t|P_1, P_2, \dots, P_t) &\equiv p_t(\alpha_t, i_t) \\ &= \sum_{i_{t-1}} \int_{\alpha_{t-1}} p(P_t|\alpha_t, i_t) \\ &\quad p(\alpha_t, i_t|\alpha_{t-1}, i_{t-1})p_{t-1}(\alpha_{t-1}, i_{t-1}) \end{aligned} \quad (3)$$

with  $\alpha_t = [s, x, y]_t$ . The probability images are denoted by “ $P_t$ ” while little “ $p$ ” denotes density functions:  $p_t(P_t|\alpha_t, i_t)$  denotes the likelihood measure of observation  $P_t$ , given the parameters  $\alpha_t$  and  $i_t$  while  $p_t(\alpha_t, i_t)$  denotes the prior at time  $t$ . We approximate the posterior density  $p(\theta_t|P_1, P_2, \dots, P_t)$  with a sequential Monte Carlo method [2; 7; 11; 13].

Using Bayesian propagation allows to take into account the uncertainty in the estimated parameter.

Monte Carlo methods use random samples for the approximation of a density function. Our MBGS system uses separate sample sets for each object in the input image. A new sample set is constructed every time a new object in the video image matches sufficiently well.

As the diffusion density  $p(\alpha_t, i_t|\alpha_{t-1}, i_{t-1})$  in Eq. 3 we use the Brownian motion model due to the absence of a better one. For the propagation of the position and scale parameters,  $x$ ,  $y$ , and  $s$ , this is straight forward. For the propagation of the silhouette we use the following strategy: The likelihood for selecting a silhouette from a certain silhouette cluster in the hierarchy is computed from the Markov transition matrix  $M$  by marginalizing over the silhouettes in that particular cluster. Within a cluster, the new silhouette is then chosen randomly. The reason for this is that our training data is too little so that the Markov transition matrix  $M$  appeared to be specific to the training videos.

## 4 Experiments

In this section we present the experimental results with our MBGS implementation. We have carried out a large number of experiments with both, online and offline video data. The experiments have shown that the drawbacks of the classical BGS approach, that were mentioned in section 1, can be remedied with MBGS:

1. Because shadows are not part of the model information, they are be classified as background by the implemented MBGS approach. In fact, most non-model object types will be classified as background, and therefore MBGS allows for effective object type filtering.
2. The output presented from the MBGS does not vary, even when the scenario changes significantly and objects are always detected completely even when significantly covered up by other objects such as bushes or lamp posts.

The verification is done systematically by comparing the AUE MBGS system with two previously developed pixel-based state-of-art background subtraction approaches: One is the non-parametric approach developed at Maryland (UMD BGS) [3]. The second one (AUE BGS) was previously developed at AUE and utilizes an alternative image noise filtering approach. The comparison is done by eye-sight due to the absense of ground-truth. We have tested live-feed as well as videos from a database of 50 videos with different scenarios such as parkinglot-scenes and park-scenes. The videos contained both side views and frontal views of the person.<sup>3</sup> In Figure 3 one sees a scenario in one of the test movies with a pedestrian walking behind trees, thereby at times being occluded. The output of two pixel-based approaches is shown in the lower part of the figure. One can notice that the shadow casted by the pedestrian is classified as foreground by the pixel-based approaches but as background by the model-based approaches.

Figure 4 shows the same scenario but this time a scene is shown where the pedestrian is heavily occluded. The occlusion causes the pedestrian to more or less disappear when the pixel based approaches are used.

The scenario in figure 5 (see also the provided video) shows two pedestrians walking towards each other, thereby crossing behind lamp posts and a statue. When processing this scenario, a combination of image filtering and the background variation appears where the silhouettes of the pixel-based are too heavily distorted to still be identified as a person.

Also both pixel-based approaches severely distorts the contours of the pedestrians. By only inspecting the pixel-based results, it is hard to tell that the foreground objects are actually pedestrians.

The system has been tested on a 2 GHz Pentium with Linux. We use RGB videos with an images size  $320 \times 240$  pixels. With only a single person to track, the system needs with 350 particles  $\approx 50$  ms/frame. 25ms/frame were used by the classical BGS, 25ms/frame were used for the matching.

<sup>3</sup> We have provided the most representative videos that document the limits of our approach.



Fig. 3. BGS approach comparison of shadow issue.

## 5 Conclusion

The presented model-based background subtraction system combines the classical background subtraction with model knowledge of foreground objects. The application of model knowledge is not applied on a binary BGS image but on the “likelihood image”, i.e. an image where each pixel value represents a confidence of belonging either to the foreground or background. This approach considerably increases robustness as these likelihoods can also be understood as uncertainties which is exploited for the tracking and silhouette selection process. Also, the propagation of densities prevents the need of selecting thresholds (e.g. for binarization of the image  $P$ ) or of maximization. Thresholds are only used for visualization purposes and otherwise for the detection of a new human in the field of view.

In the above application we have chosen silhouettes of humans, but we believe that this choice is without limit of generality since even different object types fit into the tree structure.

The presented experiments were carried out with only a single individual in the database. We have experimented also with different individuals (and thus



**Fig. 4.** BGS approach comparison of heavily occlusion.

varying contours), but the output was instable w.f.t. the choice of the individual. This is under further investigation and the use of our approach for gait recognition is future research.

## References

1. Dorin Comaniciu, Visvanathan Ramesh, and Peter Meer. Real-time tracking of non-rigid objects using mean shift. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, volume 2, pages 142–149, Hilton Head Island, SC, June 13-15, 2000.
2. A. Doucet, S. Godsill, and C. Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and Computing*, 10:197–209, 2000.
3. A. Elgammal and L. Davis. Probabilistic framework for segmenting people under occlusion. In *ICCV*, ICCV01, 2001.
4. D. Gavrilu and V. Philomin. Real-time object detection for "smart" vehicles. In *Proc. Int. Conf. on Computer Vision*, pages 87–93, Korfu, Greece, 1999.
5. I. Haritaoglu, D. Harwood, and L. Davis. W4s: A real-time system for detection and tracking people in 2.5 D. In *Proc. European Conf. on Computer Vision*, Freiburg, Germany, June 1-5, 1998.



Fig. 5. BGS approach comparison of low contrast issue.

6. T. Horprasert, D. Harwood, and L.S. Davis. A statistical approach for real-time robust background subtraction and shadow detection. In *Proceedings of IEEE ICCV'99 FRAME-RATE Workshop*, 1999.
7. M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *Int. J. of Computer Vision*, 1998.
8. M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *Int. J. of Computer Vision*, 29:5–28, 1998.
9. Yuri A. Ivanov, Aaron F. Bobick, and John Liu. Fast lighting independent background subtraction. *Int. J. of Computer Vision*, 37(2):199–207, 2000.
10. A. Kale, A.N. Rajagopalan, N. Cuntoor, V. Krueger, and R. Chellappa. Identification of humans using gait. *IEEE Trans. Image Processing*, 2004. to be published.
11. G. Kitagawa. Monte carlo filter and smoother for non-gaussian nonlinear state space models. *J. Computational and Graphical Statistics*, 5:1–25, 1996.
12. V. Krueger and S. Zhou. Exemplar-based face recognition from video. In *Proc. European Conf. on Computer Vision*, Copenhagen, Denmark, June 27-31, 2002.
13. J.S. Liu and R. Chen. Sequential monte carlo for dynamic systems. *Journal of the American Statistical Association*, 93:1031–1041, 1998.

14. K. Toyama and A. Blake. Probabilistic tracking in a metric space. In *Proc. Int. Conf. on Computer Vision*, volume 2, pages 50–59, Vancouver, Canada, 9–12 July, 2001.

# Some Transitions of Extrema-Based Multi-Scale Singularity Trees

K. Somchaipeng<sup>1,2</sup>, J. Sporning<sup>2</sup>, S. Kreiborg<sup>1</sup>,  
and P. Johansen<sup>2</sup>

<sup>1</sup>3D-Lab, School of Dentistry, Dept. of Pediatric Dentistry, University of Copenhagen Nørre Alle 20 DK-2200 Copenhagen N. Denmark

<sup>2</sup>Dept. of Computer Science, University of Copenhagen, Universitetsparken 1, DK-2100, Copenhagen N, Denmark

## Abstract

Based on the scale-space theory and the singularity theory, Extrema-Based multi-scale singularity trees (MSSTs) representing the deep structure of the images are constructed. Due to the nature of the generic catastrophic events in scale-spaces and the developed tree building scheme, the resulting MSSTs are always rooted, ordered, and binary trees.

An empirical study of the stability of MSSTs is conducted by observing the transitions of the MSSTs extracted from a series of gradually changed generated images having the same number of extrema.

Three categories of transitions are observed in the experiment, i.e. changes of positions, changes of relations, and changes of connections.

## 1 Gaussian Scale-Space

The  $N + 1$  dimensional Gaussian scale-space,  $L : \mathbb{R}^{N+1} \rightarrow \mathbb{R}$ , of an  $N$  dimensional image,  $I : \mathbb{R}^N \rightarrow \mathbb{R}$ , is an ordered stack of images, where each image is a blurred version of the former [3, 12, 4]. The blurring is performed according to the diffusion equation,

$$\partial_t L = \nabla^2 L, \quad (1)$$

where  $\partial_t L$  is the first partial-derivative of the image in the scale direction  $t$ , and  $\nabla^2$  is the Laplacian operator, which in two dimensions reads  $\partial_x^2 + \partial_y^2$ .

An example of the scale-space of a man face image is shown in Figure 1.

The Gaussian kernel is the Green's function of the heat diffusion equation, i.e.

$$L(\cdot; t) = I(\cdot) \otimes g(\cdot; t), \quad (2)$$

$$g(x; t) = \frac{1}{(4\pi t)^{N/2}} e^{-x^T x / (4t)}, \quad (3)$$



Figure 1: A scale-space of a man face image. The latter images are the simplification of the former images.

where  $L(\cdot, t)$  is the image at scale  $t$ ,  $I(\cdot)$  is the original image,  $\otimes$  is the convolution operator,  $g(\cdot; t)$  is the Gaussian kernel at scale  $t$ ,  $N$  is the dimensionality of the problem, and  $t = \sigma^2/2$ , using  $\sigma$  as the standard deviation of the Gaussian kernel. The *Gaussian scale-space* is henceforth called the scale-space in this article.

The information in scale-space is logarithmically degraded, the scale-parameter is therefore often sampled exponentially using  $\sigma = \sigma_0 e^T$ . Since differentiation commutes with convolution and the Gaussian kernel is infinitely differentiable, differentiation of images in scale-spaces is conveniently computed,

$$\partial_{x^n} L(\cdot; t) = \partial_{x^n} (I(\cdot) \otimes g(\cdot; t)) = I(\cdot) \otimes \partial_{x^n} g(\cdot; t). \quad (4)$$

Alternative implementations of the scale-space are multiplication in the Fourier Domain, finite differencing schemes for solving the heat diffusion equation, additive operator splitting [11], and recursive implementation [2, 10].

Each method has its advantages and disadvantages, we prefer the spatial convolution, since it is guaranteed not to introduce spurious extrema in homogeneous regions. Typical border conditions are Dirichlet, Cyclic repetition, and Neumann boundaries. We use Dirichlet boundaries, where the image is extended with zero values in all directions.

Although the dimensionality of the constructed scale-space is one higher than the dimensionality of the original image, *critical points*, in the image at each scale are always points. A critical point is e.g. an extremum,  $\partial_x L = \partial_y L = 0$ . Critical points are classified by the eigenvalues of the Hessian matrix, the matrix of all second derivatives, computed at that point. Critical points with all positive eigenvalues are minima, critical points with all negative eigenvalues are maxima, and critical points with a mixture of both negative and positive eigenvalues are saddles.

As we increase the scale parameter, the critical points move smoothly forming *critical paths*. Along scale, critical points meet and annihilate or are created. Such events are called *catastrophic events*, and the points where they occur are called *catastrophe points* and the collection of events at all scales is called the *deep structure* of the image.

The notion of genericity is used to disregard events that are not likely to occur for typical images, i.e. generic events are stable under slight perturbation



Figure 2: Extrema partitions of a man face image at difference scales. The MSST is constructed based on the nesting of these partitions across scales.

of the image. There are only two types of generic catastrophe events in scale-space namely pairwise *creation events* and *annihilation events* [1], and it has further been shown that generic catastrophe events only involves pairs of critical points where one and only one eigenvalue of the Hessian matrix changes its sign, e.g. the annihilation of a minimum (+, +) and a saddle (+, -). A detailed discussion of a robust method for extracting critical paths and catastrophe points from scale-spaces can be found in [9].

## 2 Extrema-Based Multi-Scale Singularity Trees

There are a few methods for constructing tree structures from the deep structure of two-dimensional images proposed in the literatures so far [6, 5, 7].

To our knowledge, an attempt to construct tree structures from the deep structure of three-dimensional images is first proposed in [9]. The scheme produces rooted ordered binary trees called Extrema-Based multi-scale singularity trees with catastrophe points as nodes. Creation catastrophes are ignored, since we are only interested in the linking of the extrema which present at the lowest scale.

The binary tree structure is obtained by linking catastrophe points, and the algorithm is as follows: Let  $\Omega \subset \mathbb{R}^3$  be a compact connected domain and define  $I : \Omega \rightarrow \mathbb{R}$  to be an image,  $\vec{e} \in \Omega$  as an extremum, and  $\vec{x} \in \Omega$  as a point. Consider the set of continuous functions  $\gamma : [0, P] \rightarrow \Omega$  for which  $\gamma(0) = \vec{e}$ ,  $\gamma(P) = \vec{x}$ , and  $\gamma \in \Gamma_{\vec{e}\vec{x}}$ , where  $\Gamma_{\vec{e}\vec{x}}$  is the set of all possible paths from an extremum  $\vec{e}$  to a point  $\vec{x}$ , and  $\gamma$  is parameterized using Euclidean arc-length. The energy  $E_{\vec{e}}(\vec{x})$  with respect to an extremum  $\vec{e}$  calculated at  $\vec{x}$  is defined as

$$E_{\vec{e}}(x) = \inf_{\gamma \in \Gamma_{\vec{e}\vec{x}}} \int_0^P \sqrt{\left| x(\gamma(p)) \right|^2 + \nu \left( \frac{dI(\gamma(p))}{dp} \right)^2} dp. \quad (5)$$

Let  $\mathcal{E} \subset \Omega$  be the set of all extrema in the image. The *extrema partition*,  $Z_i$ , associated with an extrema  $\vec{e}_i \in \mathcal{E}$  is defined as the set of all points in the images, where the energy  $E_{\vec{e}_i}(\vec{x})$  is minimal,

$$Z_i = \{ \vec{x} \in \Omega \mid E_{\vec{e}_i}(\vec{x}) < E_{\vec{e}_j}(\vec{x}), \forall \vec{e}_j \in \mathcal{E}, i \neq j \}. \quad (6)$$

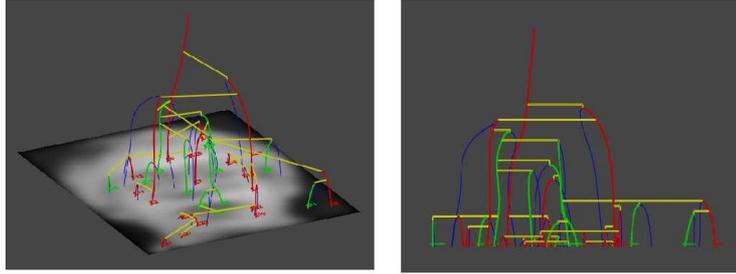


Figure 3: Extracted MSST of a man face. Critical paths of extrema, minima, and saddles are shown in red, green, and blue, respectively. Horizontal yellow lines show the linking of the tree structure.

An approximation of the energy map  $M_i : \Omega \rightarrow \mathbb{R}^+$ , which defines the energy at every point in the image associated with an extremum  $\vec{e}_i$ , can be efficiently calculated using the *Fast Marching Methods* [8]. Figure 2. show the extrema partitions of a man face at different scales.

At an catastrophe point, one extrema annihilates with a saddle and disappears, implying that also one extrema partition vanishes. The algorithm then constructs the tree structure based on the nesting of extrema partitions across scales. At a catastrophe, a node consisting of a catastrophe and two ports is created. The left port will be a reference to an extremum whose extrema partition is present immediately after the catastrophe in scale, and the right node will be the extremum that is lost at the catastrophe. The nodes are linked in a top-down manner, starting from the highest to the lowest catastrophe in scale. There will be as many nodes as there are relevant catastrophes in the constructed tree. A Full discussion of the method for constructing Extrema-Based MSSTs can be found in [9].

An example of the extrema-based multi-scale singularity tree of a man face image is shown in Figure 3.

## 2.1 Multi-Scale Singularity Strings

Due to redundancies in the tree structure and the data stored in the nodes of the tree, MSSTs can be converted to string structures simply by concatenating the nodes of the tree according to the order of the catastrophes. The resulting strings are called multi-scale singularity strings (MSSSs). The information on the structure of the tree is completely kept in the string, a MSST can be converted to a MSSS and vice versa without lose of information. An example of equivalent MSST and MSSS are shown in Figure 4.



Figure 4: A MSST and its MSSS equivalence. The MSSS is simply constructed from a MSST by orderly place nodes of the MSST according to the order of the catastrophe points in scale, from higher scales to lower scales.

### 3 Experiment

To empirically study the stability of MSSTs, a series of generated images of three stationary and one moving Gaussian blobs is used in the experiment. All images in the series have the same number of extrema, i.e. four maxima surrounding a minimum. Examples of the test images are shown in Figure 5.

Smooth changes of the images are forced by smoothly changing the parameters used for generating the images. For each test image, a scale-space is computed, and the MSST is constructed. Changes in the constructed MSSTs between neighbouring images are observed and carefully classified into categories.

### 4 Results

The structure of the constructed MSSTs are generally stable under small perturbation. For those images close to the transitions small perturbation can change the structure of the MSSTs considerably depending on the type of the transitions.

Three categories of transitions are observed in the experiment, i.e. *changes of positions*, *changes of relations*, and *changes of connections*. Changes of positions are the movements of the extrema and catastrophe points, changes of

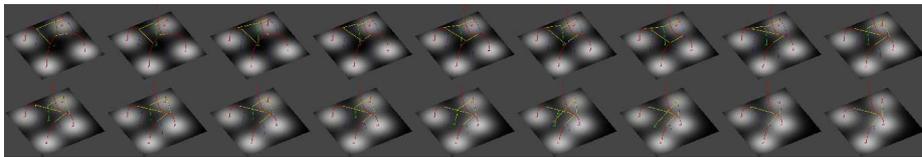


Figure 5: Example images from the series of test images . A Gaussian blob is moving from the bottom-right corner of the image to the top-left corner of the image, while the positions of the other three Gaussian blobs are fixed.

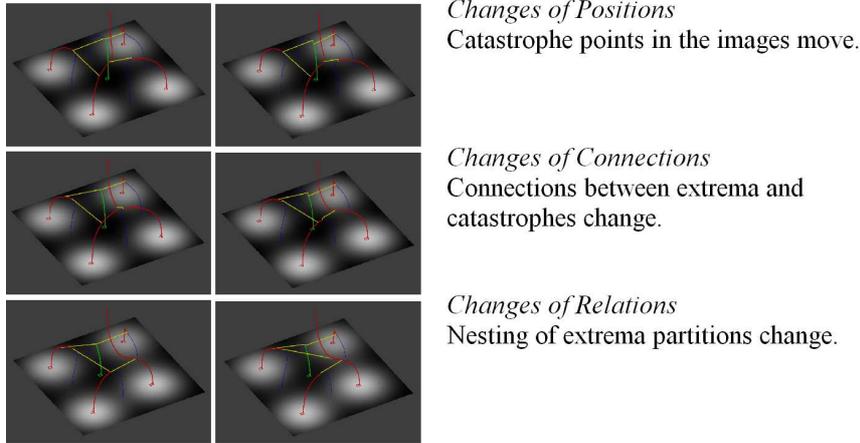


Figure 6: The three observed transitions : changes of positions, changes of relations and changes of connections. The critical paths and the linking are shown on top of the original images.

connections swap the connection between two extremum-catastrophe pairs, and changes of relations are the changes of the nesting of extrema partitions in scales.

Figure 6. shows the before and after images for each transition category, with the critical paths and the linking on top of the original image .

## 5 Discussion

In practice, some of the the observed three transition categories might occur at the same time. In general, changes of positions occur most frequently, and changes of connections and changes of relations usually accompanied by changes of positions.

For both MSSTs and MSSSs, the changes of connections transitions simply modify the data stored in the corresponding nodes of MSSTs and MSSSs, the tree structures are preserved.

The changes of connections transitions imply changes of the ordering of catastrophe points in scale therefore modify heavily the tree structures of MSSTs, since the trees are built top-down according to the order of the catastrophe points in scale. However, in cases of MSSSs, the transitions are simply the swapping of two neighboring nodes in the strings and the modifications of the data stored in the corresponding nodes

The changes of relations transitions move subtrees in MSSTs from one place to another place while only modify data store in nodes of MSSSs.

## 6 Conclusion

MSSTs and MSSSs are equivalent. They represent the same thing, the deep structure of the images. The tree structure in a MSST is completely encoded in its equivalent MSSS. The MSST can then be reconstructed from the MSSS and vice versa.

According to the experiment, MSSSs seem to be more stable than their equivalent MSSTs, in the sense that small changes in image structures induce simple transitions in MSSSs but more complicated transitions in MSSTs. The algorithms we intend to apply to the structure are the edit distance algorithms, finding the cheapest path changing one structure to another structure through transitions. Having simple transitions is an important property for the data structure, since it is easy to find a sensible cost for each transition.

The software for constructing the scale-spaces, extracting the critical paths and the catastrophe points, and constructing the MSSTs is implemented in C/C++. The software is currently capable of handling both two dimensional and three dimensional images.

The future research will be applying the extracted MSSTs or, equivalently, MSSSs, and the knowledge of their transitions, in medical image applications, such as image matching, image database indexing etc. A database of three-dimensional medical images are being collected. We focus primarily on CT and MR images of craniofacial regions.

## References

- [1] James Damon. Local Morse theory for Gaussian blurred functions. In Jon Sparring, Mads Nielsen, Luc Florack, and Peter Johansen, editors, *Gaussian Scale-Space Theory*, chapter 11, pages 147–163. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.
- [2] R. Deriche. Recursively Implementing the Gaussian and its Derivatives. In V. Srinivasan, Ong Sim Heng, and Ang Yew Hock, editors, *Proceedings of the 2nd Singapore International Conference on Image Processing*, pages 263–267. World Scientific, Singapore, 1992.
- [3] T. Iijima. Basic theory on normalization of a pattern (in case of typical one-dimensional pattern). *Bulletin of Electrotechnical Laboratory*, 26:368–388, 1962. (in Japanese).
- [4] J. J. Koenderink. The Structure of Images. *Biological Cybernetics*, 50:363–370, 1984.
- [5] Arjan Kuijper. *The deep structure of Gaussian scale space images*. PhD thesis, Image Sciences Institute, Institute of Information and Computing Sciences, Faculty of Mathematics and Computer Science, Utrecht University, 2002.

- [6] L. M. Lifshitz and S. M. Pizer. A multiresolution hierarchical approach to image segmentation based on intensity extrema. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 12(6):529–541, 1990.
- [7] B. Platel. Multiscale Hierarchical Segmentation. Technical Report BMT-Report no. 2002-04, Department of BioMedical Engineering, Technical Universitet of Eindhoven, 2002.
- [8] J. A. Sethian. Fast Marching Methods. *SIAM Review*, 41(2):199–235, 1999.
- [9] K. Somchaipeng, J Sporning, S. Kreiborg, and P. Johansen. Software for Extracting Multi-Scale Singularity Trees. Technical report, Deliverable No.8, DSSCV, IST-2001-35443, 15. September 2003.
- [10] L.J. van Vliet, I.T. Young, and P.W. Verbeek. Recursive Gaussian Derivative Filters. In *Proceedings of the 14th International Conference on Pateern Recognition ICPR '98*, volume 1, pages 509–514., Brisbane, Australia, 1998. IEEE Computer Society Press.
- [11] J. Weickert, K. J. Zuiderveld, B. M. ter Haar Romeny, and W. J. Niessen. Parallel implementations of AOS schemes: A fast way of nonlinear diffusion filtering. In *Proceedings of the 4th International Conference on Image Processing*, volume 3, pages 396–399, Santa Barbara, CA, USA, October26–29 1997. IEEE Computer Society Press.
- [12] A. P. Witkin. Scale–space filtering. In *Proc. 8th Int. Joint Conf. on Artificial Intelligence (IJCAI '83)*, volume 2, pages 1019–1022, Karlsruhe, Germany, August 1983.

# A Test Statistic in the Complex Wishart Distribution and Its Application to Change Detection in Polarimetric SAR Data

Knut Conradsen, Allan Aasbjerg Nielsen, Jesper Schou, and Henning Skriver

**Abstract**—When working with multilook fully polarimetric synthetic aperture radar (SAR) data, an appropriate way of representing the backscattered signal consists of the so-called covariance matrix. For each pixel, this is a  $3 \times 3$  Hermitian positive definite matrix that follows a complex Wishart distribution. Based on this distribution, a test statistic for equality of two such matrices and an associated asymptotic probability for obtaining a smaller value of the test statistic are derived and applied successfully to change detection in polarimetric SAR data. In a case study, EMISAR L-band data from April 17, 1998 and May 20, 1998 covering agricultural fields near Foulum, Denmark are used. Multilook full covariance matrix data, azimuthal symmetric data, covariance matrix diagonal-only data, and horizontal–horizontal (HH), vertical–vertical (VV), or horizontal–vertical (HV) data alone can be used. If applied to HH, VV, or HV data alone, the derived test statistic reduces to the well-known gamma likelihood-ratio test statistic. The derived test statistic and the associated significance value can be applied as a line or edge detector in fully polarimetric SAR data also.

**Index Terms**—Covariance matrix test statistic, EMISAR, radar applications, radar polarimetry, remote sensing change detection.

## I. INTRODUCTION

**D**UE TO ITS all-weather mapping capability independently of, for instance, cloud cover, synthetic aperture radar (SAR) data hold a strong potential, for example, for change detection studies in remote sensing applications. In this paper, multitemporal SAR images of agricultural fields are used to demonstrate a new change detection method for polarimetric SAR data. It is well known that the development of different crops over time causes changes in the backscatter. The radar backscattering is sensitive to the dielectric properties of the vegetation and the soil, to the plant structure (i.e., the size, shape, and orientation distributions of the scatterers), to the surface roughness, and to the canopy structure (e.g., row direction and spacing and cover fraction) [1], [2].

The polarimetric SAR measures the amplitude and phase of backscattered signals in four combinations of the linear receive and transmit polarizations: horizontal–horizontal (HH),

horizontal–vertical (HV), vertical–horizontal (VH), and vertical–vertical (VV). These signals form the complex scattering matrix that relates the incident and the scattered electric fields [3]. The inherent speckle in the SAR data can be reduced by spatial averaging at the expense of loss of spatial resolution. In this so-called multilook case, a more appropriate representation of the backscattered signal is the covariance matrix in which the average properties of a group of resolution cells can be expressed in a single matrix. The average covariance matrix is defined as [3]

$$\langle \mathcal{C} \rangle = \begin{bmatrix} \langle S_{hh} S_{hh}^* \rangle & \langle S_{hh} S_{lv}^* \rangle & \langle S_{hh} S_{vv}^* \rangle \\ \langle S_{lv} S_{hh}^* \rangle & \langle S_{lv} S_{lv}^* \rangle & \langle S_{lv} S_{vv}^* \rangle \\ \langle S_{vv} S_{hh}^* \rangle & \langle S_{vv} S_{lv}^* \rangle & \langle S_{vv} S_{vv}^* \rangle \end{bmatrix} \quad (1)$$

where  $\langle \cdot \rangle$  denotes ensemble averaging;  $*$  denotes complex conjugation; and  $S_{rt}$  is the complex scattering amplitude for receive polarization  $r$  and transmit polarization  $t$  ( $r$  and  $t$  are either  $h$  for horizontal or  $v$  for vertical). Reciprocity, which normally applies to natural targets, gives  $S_{hv} = S_{vh}$  (in the backscattering direction using the backscattering alignment convention [3]) and results in the covariance matrix (1) with rank 3.  $\langle \mathcal{C} \rangle$  follows a complex Wishart distribution [4]. The components in the covariance matrix containing both co- and cross-polarized scattering matrix elements often contain little information. For randomly distributed targets with azimuthal symmetry, the elements are zero [5].

In this paper, a test statistic for equality of two complex covariance matrices and an associated asymptotic probability measure for obtaining a smaller value of the test statistic are derived and applied to change detection in fully polarimetric SAR data. In [6], a change detection scheme based on canonical correlations analysis is applied to scalar EMISAR data (see also [7]–[10]).

If used with HH, VV, or HV data only, the test statistic reduces to the well-known test statistic for equality of the scale parameters in two gamma distributions.

The derived test statistic and the associated significance measure can be applied as a line or edge detector in fully polarimetric SAR data also [11].

Section II sketches important aspects of the complex Gaussian and Wishart distributions, the likelihood-ratio test statistic in the complex Wishart distribution, and the associated significance measure. Section III gives a case study in which data from the Danish airborne EMISAR [12], [13] are used. Section IV discusses the results from the case study, and Section V concludes.

Manuscript received June 8, 2001; revised October 17, 2002. This work was supported in part by The Danish National Research Councils in part under the Earth Observation Programme and in part under the European Space Agency Follow-On Research Programme. The EMISAR data acquisitions and part of the data processing were supported by the Danish National Research Foundation.

K. Conradsen and A. A. Nielsen are with IMM, Informatics and Mathematical Modelling, Technical University of Denmark, DK-2800 Lyngby, Denmark (e-mail: aa@imm.dtu.dk).

J. Schou and H. Skriver are with EMI, Section of Electromagnetic Systems, Technical University of Denmark (Ørsted DTU), DK-2800 Lyngby, Denmark.

Digital Object Identifier 10.1109/TGRS.2002.808066

## II. THEORY

This section describes the complex normal and Wishart distributions and the likelihood-ratio test for equality of two complex Wishart matrices. For a more thorough description, see the Appendix.

### A. Complex Normal Distribution

We say that a  $p$ -dimensional random complex vector  $\mathbf{Z}$  follows a complex multivariate normal distribution with mean  $\mathbf{0}$  and dispersion matrix  $\Sigma$ , i.e.,

$$\mathbf{Z} = [Z_1 \cdots Z_p]^T \in N_C(\mathbf{0}, \Sigma) \quad (2)$$

if the frequency function is

$$\begin{aligned} f(\mathbf{z}) &= \frac{1}{\pi^p |\Sigma|} \exp \{-\mathbf{z}^{*T} \Sigma^{-1} \mathbf{z}\} \\ &= \frac{1}{\pi^p |\Sigma|} \exp \{-\text{tr} [\Sigma^{-1} \mathbf{z} \mathbf{z}^{*T}]\} \end{aligned} \quad (3)$$

where  $|\cdot|$  denotes the determinant;  $\text{tr}$  denotes the trace of a matrix; and  $^{*T}$  denotes complex conjugation ( $^*$ ) and transpose ( $^T$ ).

### B. Complex Wishart Distribution

We say that a Hermitian positive definite random  $p \times p$  matrix  $\mathbf{X}$  follows a complex Wishart distribution, i.e.,

$$\mathbf{X} \in W_C(p, n, \Sigma) \quad (4)$$

if the frequency function is

$$w(\mathbf{x}) = \frac{1}{\Gamma_p(n)} \frac{1}{|\Sigma|^n} |\mathbf{x}|^{n-p} \exp \{-\text{tr} [\Sigma^{-1} \mathbf{x}]\} \quad (5)$$

where

$$\Gamma_p(n) = \pi^{p(p-1)/2} \prod_{j=1}^p \Gamma(n-j+1). \quad (6)$$

The frequency function is defined for  $\mathbf{w}$  positive definite.

If  $\mathbf{X}$  and  $\mathbf{Y}$  are independent and both follow complex Wishart distributions

$$\mathbf{X} \in W_C(p, n, \Sigma) \text{ and } \mathbf{Y} \in W_C(p, m, \Sigma) \quad (7)$$

then their sum also follows a complex Wishart distribution

$$\mathbf{S} = \mathbf{X} + \mathbf{Y} \in W_C(p, n+m, \Sigma). \quad (8)$$

### C. Test for Equality of Two Complex Wishart Matrices

Let the independent  $p \times p$  Hermitian positive definite matrices  $\mathbf{X}$  and  $\mathbf{Y}$  be complex Wishart distributed, i.e.,  $\mathbf{X} \in W_C(p, n, \Sigma_x)$  with  $\hat{\Sigma}_x = 1/n\mathbf{X}$  and  $\mathbf{Y} \in W_C(p, m, \Sigma_y)$  with  $\hat{\Sigma}_y = 1/m\mathbf{Y}$ . We consider the null hypothesis  $H_0 : \Sigma_x = \Sigma_y$ , which states that the two matrices are equal against the alternative hypothesis  $H_1 : \Sigma_x \neq \Sigma_y$ .

In general, suppose that the observations on which we shall base our test have joint density  $f(x, \theta)$  where  $\theta$  is the set of parameters of the probability function that has generated the

data. Then  $H_0$  states that  $\theta \in \omega_0$  where  $\omega_0$  is a subset of the set  $\Omega$  of all possible  $\theta$ .  $H_1$  states that  $\omega \in \omega_1$  where  $\omega_0$  and  $\omega_1$  are disjoint, and often  $\Omega = \omega_0 \cup \omega_1$ . The likelihood ratio

$$Q = \frac{\max_{\theta \in \omega_0} L(\theta)}{\max_{\theta \in \Omega} L(\theta)} \quad (9)$$

where  $L$  is the likelihood function rejects  $H_0$  for small values.

If  $H_0$  is true (in statistical parlance ‘‘under  $H_0$ ’’), then in our case  $\mathbf{S} = \mathbf{X} + \mathbf{Y} \in W_C(p, n+m, \Sigma)$  with  $\hat{\Sigma} = 1/(n+m)\mathbf{S} = 1/(n+m)(\mathbf{X} + \mathbf{Y})$ . The likelihood-ratio test statistic becomes

$$Q = \frac{L(\hat{\Sigma})}{L_x(\hat{\Sigma}_x) L_y(\hat{\Sigma}_y)}. \quad (10)$$

Here

$$L_x(\Sigma_x) = \frac{1}{\Gamma_p(n)} |\Sigma_x|^{-n} |\mathbf{X}|^{n-p} \exp \{-\text{tr} [\Sigma_x^{-1} \mathbf{X}]\} \quad (11)$$

and

$$L_x(\hat{\Sigma}_x) = \frac{1}{\Gamma_p(n)} \left| \frac{1}{n} \mathbf{X} \right|^{-n} |\mathbf{X}|^{n-p} \exp \{-n \text{tr} \mathbf{I}\} \quad (12)$$

where  $\mathbf{I}$  is the identity matrix ( $\text{tr} \mathbf{I} = p$ ). Similar expressions are valid for  $L_y(\Sigma_y)$  and  $L_y(\hat{\Sigma}_y)$ . For the numerator of  $Q$  we get

$$\begin{aligned} L(\Sigma) &= \frac{1}{\Gamma_p(n) \Gamma_p(m)} |\Sigma|^{-(n+m)} |\mathbf{X}|^{n-p} |\mathbf{Y}|^{m-p} \\ &\quad \cdot \exp \{-\text{tr} [\Sigma^{-1} (\mathbf{X} + \mathbf{Y})]\} \end{aligned} \quad (13)$$

and

$$\begin{aligned} L(\hat{\Sigma}) &= \frac{1}{\Gamma_p(n) \Gamma_p(m)} \left| \frac{1}{n+m} (\mathbf{X} + \mathbf{Y}) \right|^{-(n+m)} \\ &\quad \cdot |\mathbf{X}|^{n-p} |\mathbf{Y}|^{m-p} \exp \{-(n+m) \text{tr} \mathbf{I}\}. \end{aligned} \quad (14)$$

This leads to the desired likelihood-ratio test statistic

$$Q = \frac{(n+m)^{p(n+m)} |\mathbf{X}|^n |\mathbf{Y}|^m}{n^m m^m |\mathbf{X} + \mathbf{Y}|^{n+m}}. \quad (15)$$

If  $n = m$ , which is typically the case for change detection, we get

$$\ln Q = n(2p \ln 2 + \ln |\mathbf{X}| + \ln |\mathbf{Y}| - 2 \ln |\mathbf{X} + \mathbf{Y}|). \quad (16)$$

If

$$\rho = 1 - \frac{2p^2 - 1}{6p} \left( \frac{1}{n} + \frac{1}{m} - \frac{1}{n+m} \right) \quad (17)$$

and

$$\begin{aligned} \omega_2 &= -\frac{p^2}{4} \left( 1 - \frac{1}{\rho} \right)^2 + \frac{p^2(p^2 - 1)}{24} \\ &\quad \cdot \left( \frac{1}{n^2} + \frac{1}{m^2} - \frac{1}{(n+m)^2} \right) \frac{1}{\rho^2} \end{aligned} \quad (18)$$

then the probability of finding a smaller value of  $-2\rho \ln Q$  is

$$\begin{aligned} P \{-2\rho \ln Q \leq z\} &\simeq P \{\chi^2(p^2) \leq z\} \\ &\quad + \omega_2 [P \{\chi^2(p^2 + 4) \leq z\} - P \{\chi^2(p^2) \leq z\}]. \end{aligned} \quad (19)$$

For covariance matrix data,  $p = 3$ . For HH, HV, or VV data,  $p = 1$ . In the latter case,  $\mathbf{X}$  and  $\mathbf{Y}$  are therefore scalars  $X$  and  $Y$ , and  $Q$  reduces to

$$Q = \frac{(n+m)^{n+m}}{n^n m^m} \frac{X^n Y^m}{(X+Y)^{n+m}} \quad (20)$$

which is equivalent to the well-known likelihood-ratio test statistic for the equality of two gamma parameters [14], [15] (see the Appendix).

Fig. 1 shows  $\rho$  and  $\omega_2$  as functions of the number of looks for  $n = m$  and  $p = 3$ .

#### D. Azimuthal Symmetry

By swapping first rows and then columns two and three in  $\langle \mathbf{C} \rangle$  in (1), we obtain in the azimuthal symmetry case

$$\begin{bmatrix} \langle S_{hh} S_{hh}^* \rangle & \langle S_{hh} S_{vv}^* \rangle & 0 \\ \langle S_{vv} S_{hh}^* \rangle & \langle S_{vv} S_{vv}^* \rangle & 0 \\ 0 & 0 & \langle S_{hv} S_{hv}^* \rangle \end{bmatrix} = \begin{bmatrix} \mathbf{X}_1 & 0 \\ 0 & \mathbf{X}_2 \end{bmatrix} = \mathbf{X} \quad (21)$$

where  $\mathbf{X}_1$  is  $p_1 \times p_1$  (here  $2 \times 2$ ), and  $\mathbf{X}_2 = \langle S_{hv} S_{hv}^* \rangle$  is  $p_2 \times p_2$  (here  $1 \times 1$ ). This matrix is not Wishart distributed. We now consider  $\mathbf{X}_1 \in W_C(p_1, n, \Sigma_{x1})$ ,  $\mathbf{X}_2 \in W_C(p_2, n, \Sigma_{x2})$ ,  $\mathbf{Y}_1 \in W_C(p_1, m, \Sigma_{y1})$ , and  $\mathbf{Y}_2 \in W_C(p_2, m, \Sigma_{y2})$ , and we assume that  $\mathbf{X}_1$ ,  $\mathbf{X}_2$ ,  $\mathbf{Y}_1$ , and  $\mathbf{Y}_2$  are mutually independent.

We want to test the hypothesis  $H_0 : \Sigma_{x1} = \Sigma_{y1} (= \Sigma_1)$  and  $\Sigma_{x2} = \Sigma_{y2} (= \Sigma_2)$  against all alternatives. We have the likelihood function

$$L(\Sigma_{x1}, \Sigma_{x2}, \Sigma_{y1}, \Sigma_{y2}) = L^{(1)}(\Sigma_{x1}) L^{(2)}(\Sigma_{x2}) L^{(3)}(\Sigma_{y1}) L^{(4)}(\Sigma_{y2}). \quad (22)$$

The likelihood-ratio test statistic becomes

$$Q = \frac{(n+m)^{p(n+m)}}{n^{pn} m^{pm}} \frac{|\mathbf{X}_1|^n |\mathbf{Y}_1|^m |\mathbf{X}_2|^n |\mathbf{Y}_2|^m}{|\mathbf{X}_1 + \mathbf{Y}_1|^{n+m} |\mathbf{X}_2 + \mathbf{Y}_2|^{n+m}} = \frac{(n+m)^{p(n+m)}}{n^{pn} m^{pm}} \frac{|\mathbf{X}|^n |\mathbf{Y}|^m}{|\mathbf{X} + \mathbf{Y}|^{n+m}} \quad (23)$$

where the latter equality is due to the fact that the determinant of a block diagonal matrix is the product of the determinants of the diagonal elements, i.e., we get the same test statistic as in the full covariance matrix case. In this case  $p^2 = p_1^2 + p_2^2$ . If

$$\rho = \frac{1}{p^2} (p_1^2 \rho_1 + p_2^2 \rho_2) \quad (24)$$

$$\rho_i = 1 - \frac{2p_i^2 - 1}{6p_i} \left( \frac{1}{n} + \frac{1}{m} - \frac{1}{n+m} \right) \quad (25)$$

and

$$\omega_2 = -\frac{p^2}{4} \left( 1 - \frac{1}{\rho} \right)^2 + \frac{p_1^2 (p_1^2 - 1) + p_2^2 (p_2^2 - 1)}{24} \cdot \left( \frac{1}{n^2} + \frac{1}{m^2} - \frac{1}{(n+m)^2} \right) \frac{1}{\rho^2} \quad (26)$$

then the probability of finding a smaller value of  $-2\rho \ln Q$  is

$$P\{-2\rho \ln Q \leq z\} \simeq P\{\chi^2(p^2) \leq z\} + \omega_2 [P\{\chi^2(p^2 + 4) \leq z\} - P\{\chi^2(p^2) \leq z\}]. \quad (27)$$

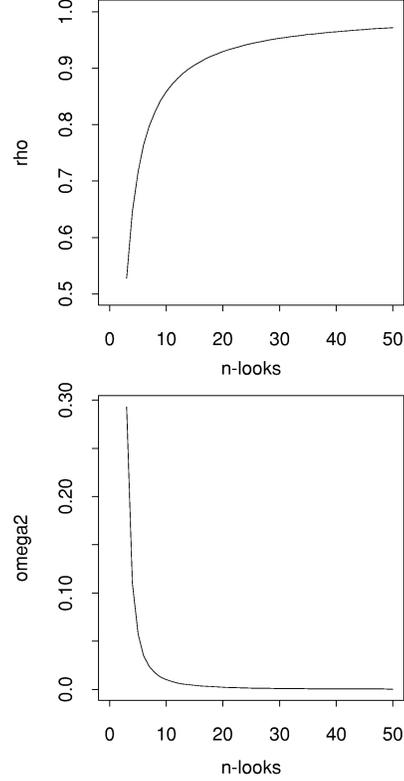


Fig. 1.  $\rho$  and  $\omega_2$  as functions of the number of looks for  $n = m$  and  $p = 3$ .

### III. DATA

To illustrate the change detection capability of the derived test statistic, EMISAR and ground data from an agricultural test site at the Research Center Foulum located in Central Jutland, Denmark are used. Agricultural fields have been selected for the analysis because of the large change in the polarimetric properties for such areas, due to the development of the crops with time. Polarimetric parameters of agricultural crops have previously been analyzed from this area [2].

#### A. SAR Data and Calibration

The EMISAR system is a fully polarimetric airborne SAR system, and it operates at two frequencies: C-band (5.3-GHz/5.7-cm wavelength) and L-band (1.25-GHz/24-cm wavelength) [11], [12]. The SAR system is normally operated from an altitude of approximately 12 500 m; the spatial resolution is 2 m  $\times$  2 m (one-look); the ground range swath is approximately 12 km; and typical incidence angles range from 35° to 60°. The processed data from this system are fully calibrated by means of an advanced internal calibration system. The radiometric calibration is better than  $\pm 0.5$  dB, and the channel imbalance is less than  $\pm 0.5$  dB in amplitude and  $\pm 5^\circ$  in phase [12]. The cross-polarization contamination is generally suppressed by more than 30 dB. The stability of the system is very important in the change detection scheme set up in this paper.

A large number of acquisitions with both the C- and L-band polarimetric SAR from 1994 to 1999 exist for the agricultural test site. To illustrate the change detection capability of the derived test statistic, L-band data from April 17, 1998 and May

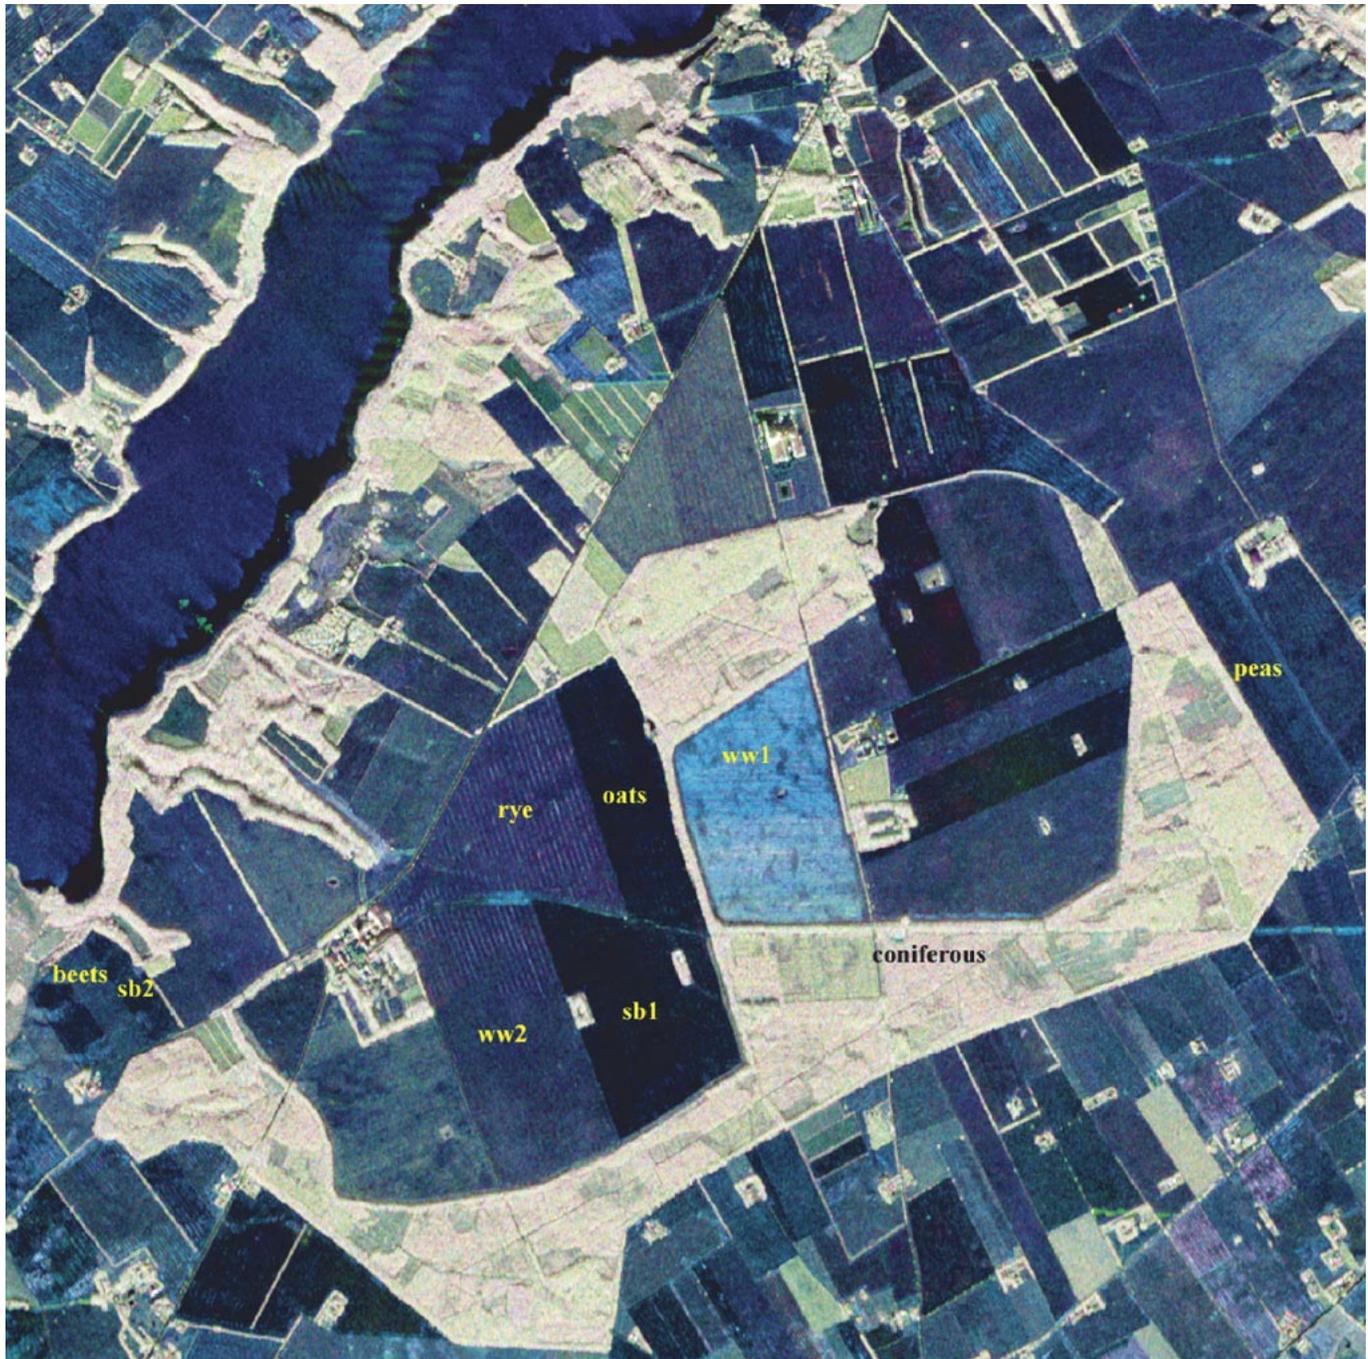


Fig. 2. L-band EMISAR image of the test area acquired on April 17, 1998, 5210 m  $\times$  5120 m.

20, 1998 have been used. The two EMISAR images are shown in Figs. 2 and 3, as color composites of the HH (green), HV (actually the complex addition of HV and VH, red), and VV (blue) channels. The HH and VV channels are stretched linearly between  $-30$  dB and  $0$  dB. The HV channel is stretched linearly between  $-36$  dB and  $-6$  dB. The area is relatively flat, and corrections of the local incidence angle due to terrain slope are not critical in this study, since the acquisition geometry for the two acquisitions are almost identical, and therefore the correction has not been carried out. The geometrical coregistration is, however, very important in a change detection application, where two images are compared on a pixel-by-pixel basis. The polari-

metric images were registered to a digital elevation model generated from interferometric airborne data acquired by EMISAR. The registration was carried out by combining a model of the imaging geometry with few ground control points, and the images were registered to one another with a root mean-square accuracy of better than one pixel [13]. In the study, 13-look covariance matrix data with a  $5 \text{ m} \times 5 \text{ m}$  pixel spacing are used.

#### B. Test Site

The area contains a number of agricultural fields of different sizes with different crops. The lengthy, dark blue feature in the upper left corner of Figs. 2 and 3 is a lake, while the bright



Fig. 3. L-band EMISAR image of the test area acquired on May 20, 1998, 5210 m  $\times$  5120 m.

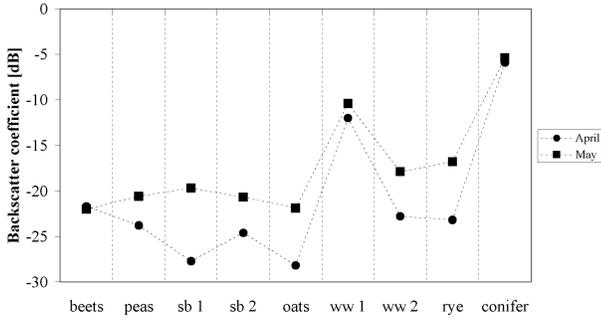
greenish areas seen especially in the lower part of the images are forests (primarily coniferous forests). In the April acquisition, the colors of the agricultural fields are predominantly bluish, due to the larger VV- than HH-backscatter coefficient for the bare fields for the spring crops and the sparsely vegetated fields for the winter crops. For the May acquisition, the spring crops are mainly in the tillering stage, and the winter crops are at the end of the stem elongation stage, in the boot stage, or at the beginning of heading, depending on the crop type.

A number of test areas have been selected for quantitative analysis of the test statistic. These areas are outlined in Figs. 2 and 3, and the development stage and the height of the vegeta-

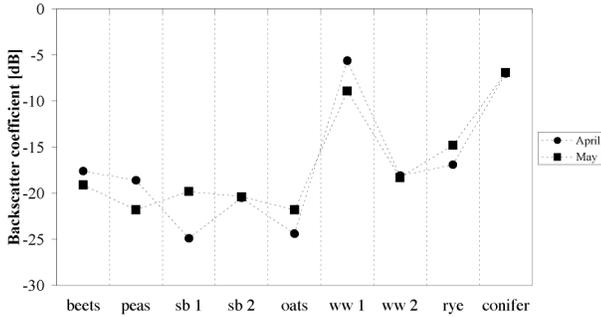
tion are shown in Table I for reference. Five spring crop fields have been used, i.e., one beet field, one pea field, two spring barley fields, and one oats field. All spring crop fields are bare at the April acquisition with the surface being relatively smooth due to sowing or harrowing. At the May acquisition, the beet field is still bare, whereas the other fields have some relatively dense and low vegetation. The three winter crop fields, i.e., two winter wheat fields and one rye field, have low vegetation for the April vegetation and relatively dense and high vegetation for the May acquisition. Finally, a common spruce field, which is virtually unchanged between the two acquisitions, is used in the investigation.

TABLE I  
DEVELOPMENT STAGES AND VEGETATION HEIGHTS (IN PARENTHESES)

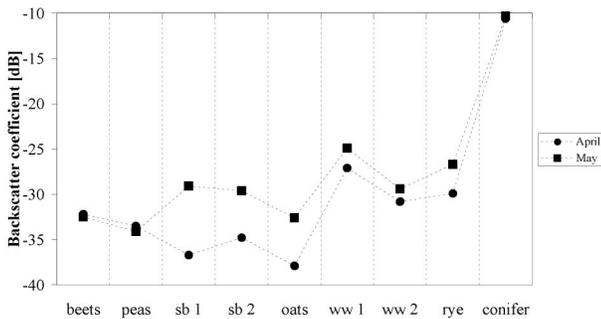
Crop	April	May
Beets	bare (harrowed)	bare
Peas	bare (rolled)	middle of side shoot development (11 cm)
Spring barley 1	bare	end of tillering (14 cm)
Spring barley 2	bare (sowed)	middle of tillering (12 cm)
Oats	bare	end of tillering (22 cm)
Winter wheat 1	beginning of stem elongation	end of stem elongation (44 cm)
Winter wheat 2	end of tillering (14 cm)	end of stem elongation (42 cm)
Rye	beginning of tillering (21 cm)	end of boot stage (51 cm)
Coniferous	common spruce	common spruce



(a)



(b)

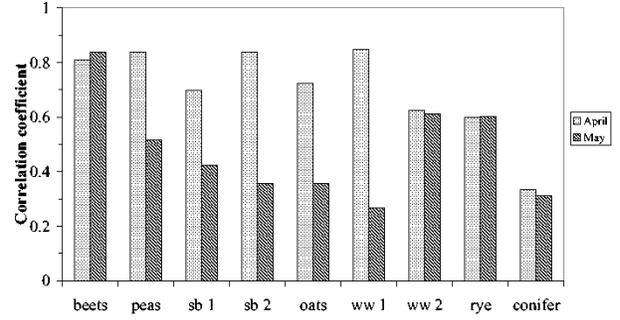


(c)

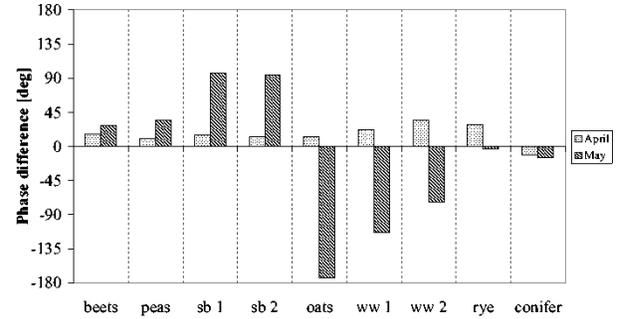
Fig. 4. (a)  $\gamma_{hh}^0$ , (b)  $\gamma_{vv}^0$ , and (c)  $\gamma_{hv}^0$  backscatter coefficients for the test areas shown in Figs. 2 and 3 for L-band in April and in May.

#### IV. RESULTS AND DISCUSSION

In Section IV-A, polarimetric parameters for the fields used in the quantitative evaluation will be presented and discussed to provide the background for interpretation of the test statistic results. The results for the test statistic are presented and discussed in Section IV-B.



(a)



(b)

Fig. 5. (a) Correlation coefficient  $\rho_{hhvv}$  and (b) phase difference  $\phi_{hh-vv}$  between HH and VV for the test areas shown in Figs. 2 and 3 for L-band in April and in May.

#### A. Polarimetric Parameters

The polarimetric parameters used to describe the selected fields are standard parameters derived from the covariance matrix (1) [2]; the  $\gamma^0$  backscatter coefficients  $\gamma_{hh}^0$ ,  $\gamma_{hv}^0$ , and  $\gamma_{vv}^0$  where the  $\gamma^0$  backscatter coefficient is slightly less dependent on the incidence angle than the  $\sigma^0$  backscatter coefficient [2]; the correlation coefficient  $\rho_{hhvv}$  and the phase difference  $\phi_{hh-vv}$  of the HH and VV components, which contain important information about the scattering mechanisms; and the co- and cross-polarized polarization signatures, which are graphical representations of the polarimetric properties [2], [3]. Fig. 4 shows the  $\gamma_{hh}^0$ ,  $\gamma_{hv}^0$ , and  $\gamma_{vv}^0$  backscatter coefficients for the various test fields and for both the April and the May acquisitions. Correspondingly,  $\rho_{hhvv}$  and  $\phi_{hh-vv}$  are shown in Fig. 5, and the polarization responses are shown in Fig. 6.

1) *Spring Crops:* All spring crops (beets, peas, spring barley 1 and 2, oats) show classical behavior for rough surface scattering for the April acquisition, i.e., high  $\rho_{hhvv}$  [Fig. 5(a)], small  $\phi_{hh-vv}$  [Fig. 5(b)], low  $\gamma_{hv}^0$  backscatter [Fig. 4(c)], larger  $\gamma_{vv}^0$  than  $\gamma_{hh}^0$  backscatter [Fig. 4(a) and (b)], and textbook examples of surface scattering polarization responses (which are therefore not shown here). The actual backscatter level from the surface is, of course, controlled by the soil moisture and the surface roughness of the individual fields, and we observe rather weak backscatter from the spring barley 1 and the oats fields (Fig. 4) for the April acquisition due to very smooth surfaces.

The beets field also shows rough surface behavior for the May acquisition [Fig. 6(a)]. The pea field shows some volume scattering behavior for the May acquisition, due to the sparse vegetation, i.e., the  $\rho_{hhvv}$  [Fig. 5(a)] has decreased, and the pedestal of the polarization response has increased [Fig. 6(b)]. This effect is

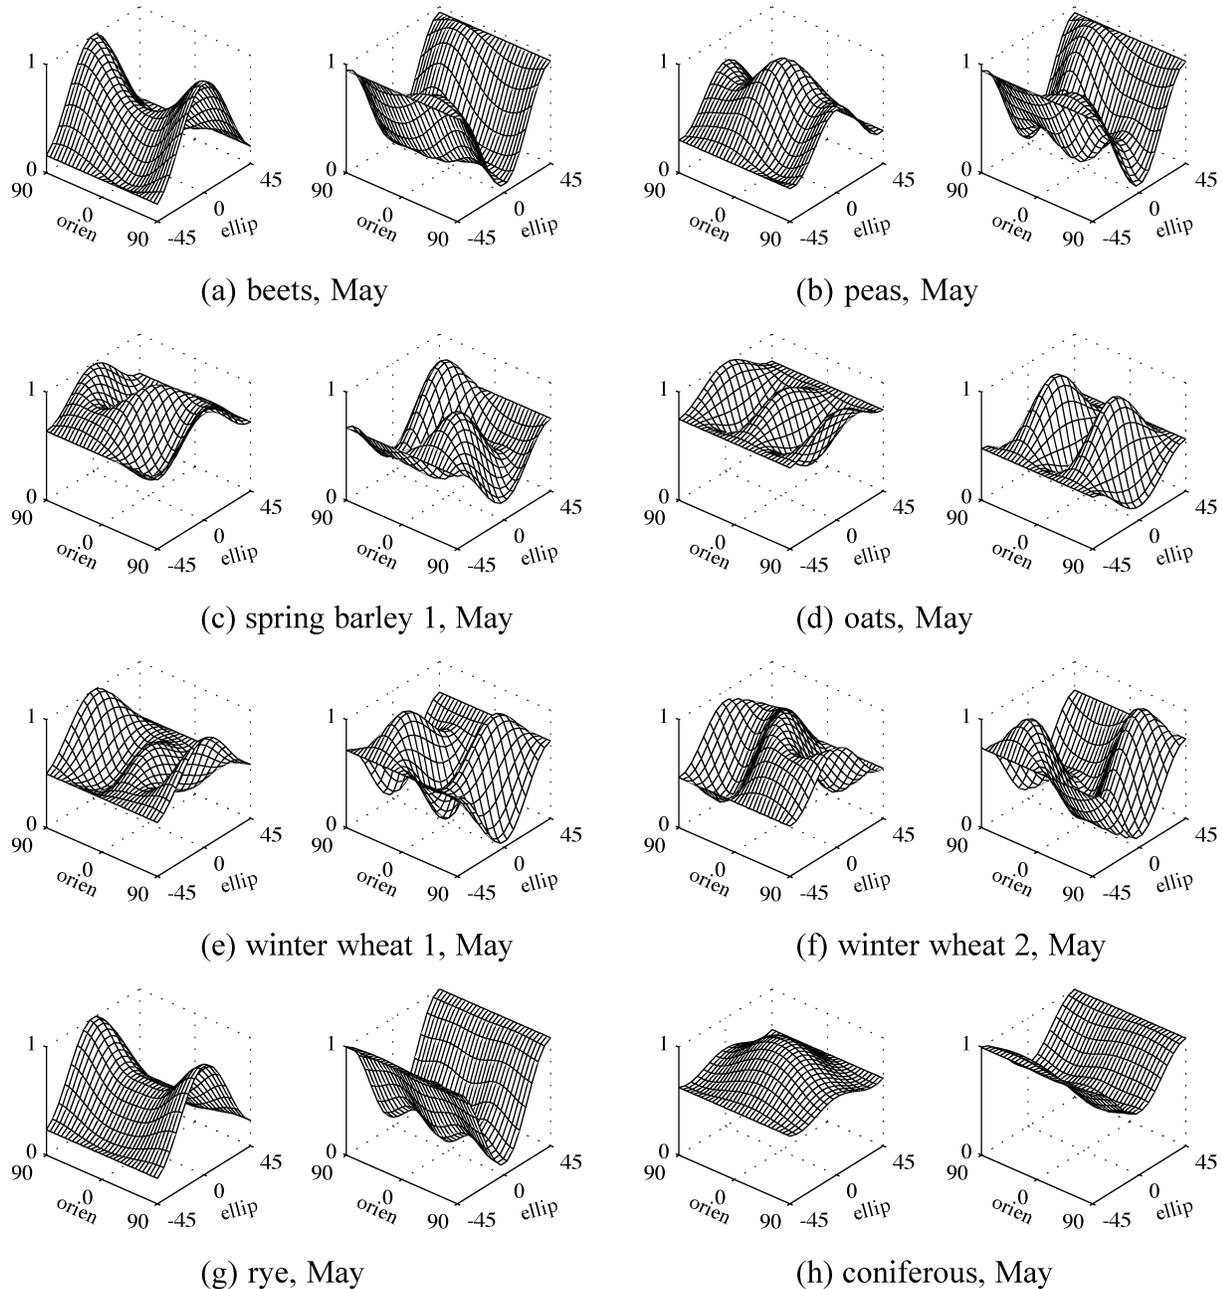


Fig. 6. Polarization signatures for the test areas shown in Figs. 3 for the L-band acquisition in May. Orientation angle of  $0^\circ$  corresponds to HH backscatter and the left-hand signature is the copolarized signature, whereas the right-hand signature is the cross-polarized signature.

even more pronounced for the spring barley and the oats fields, due to a more dense vegetation [Figs. 5(a) and 6(c) and (d)]. For the latter fields, a large  $\phi_{hh-vv}$  is observed too [Fig. 5(b)], and a pronounced double-bounce response is observed, especially for oats [Fig. 6(c) and (d)]. The double-bounce scattering is most likely caused by penetration through the vegetation, scattering from the ground surface, and scattering from the vegetation, or vice versa. This phenomenon has previously been observed early in the growing season for winter crops [2].

2) *Winter Crops*: The backscatter coefficients from the winter crops are, in general, larger than from the spring crops (Fig. 4), due to the contribution from the volume scattering. The behavior of the winter wheat and the rye fields resembles surface scattering for the April acquisition (Fig. 5), indicating

penetration through the vegetation and a large surface scattering component. The cross-polarized backscatter is, however, somewhat larger than for surface scattering, due to the volume scattering contribution [Fig. 4(c)]. The backscattering from the winter wheat 1 field is significantly larger than from the winter wheat 2 field for both acquisitions (Fig. 4). The reason is that the sowing direction for the winter wheat 1 field is exactly perpendicular to the radar look direction. For the May acquisition, the winter wheat fields also show some double-bounce scattering behavior [Figs. 5(b) and 6(e) and (f)]. The rye field shows virtually no change in the polarimetric parameters between the two acquisitions, except for some increase in the  $\gamma_{hh}^0$  backscatter (Fig. 4). The coniferous forest area shows pronounced volume scattering behavior for both acquisitions, i.e.,

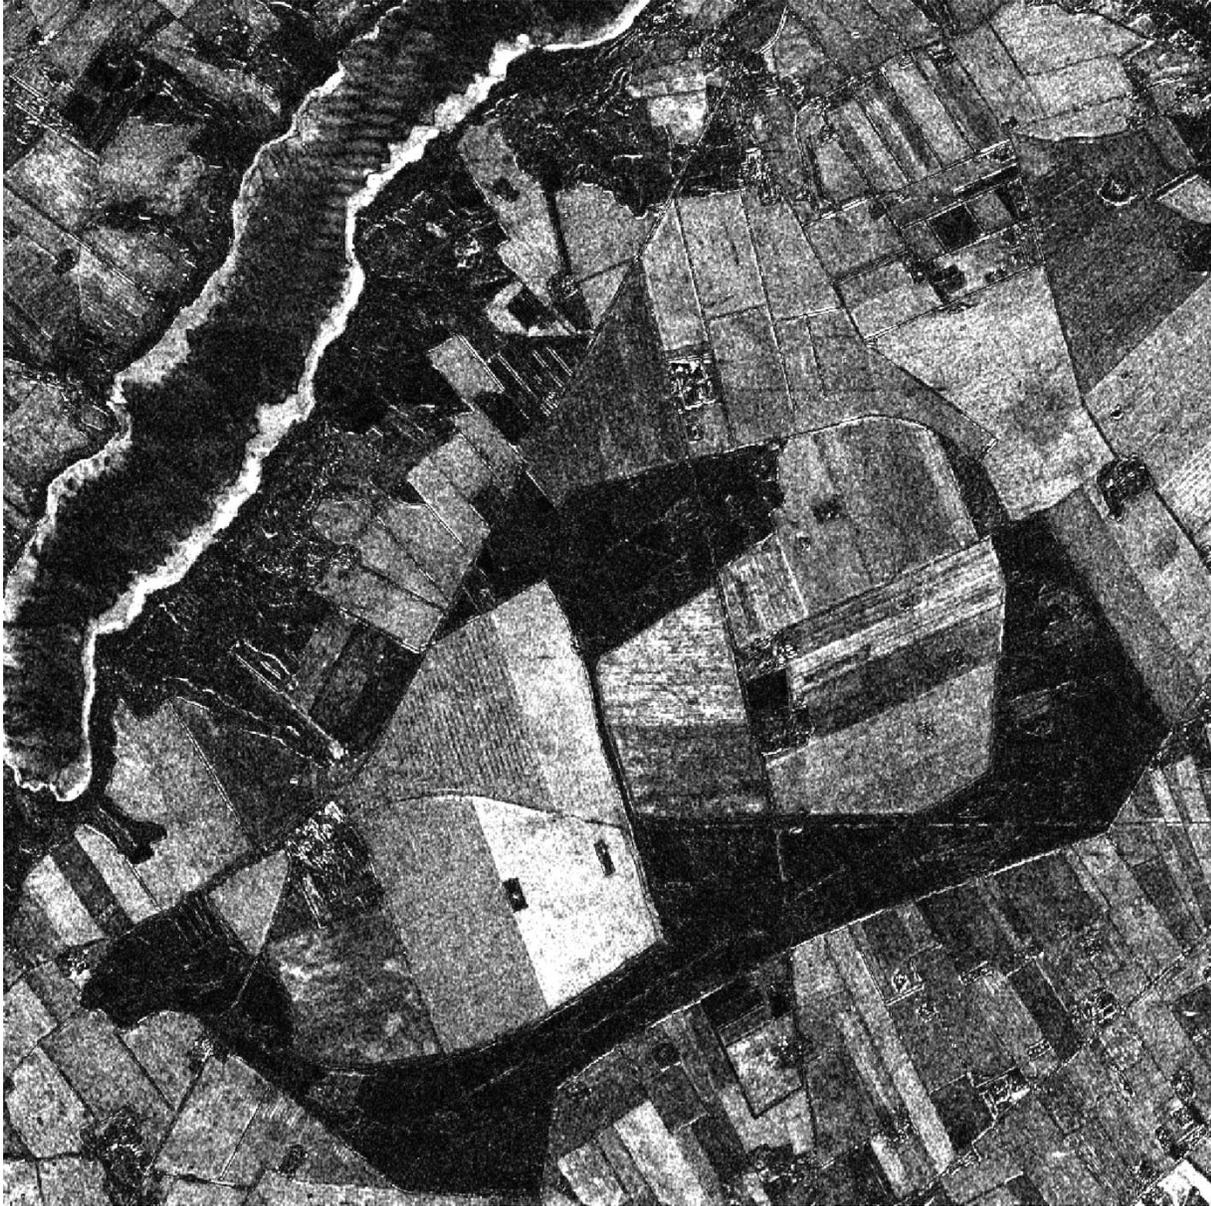


Fig. 7. Logarithm of the test statistic  $\ln Q$  (16) for the images shown in Figs. 2 and 3 in the assumed azimuthally symmetric case.

small  $\rho_{hhvv}$  [Fig. 5(a)], small  $\phi_{hh-vv}$  [Fig. 5(b)], strong  $\gamma_{hv}^0$  backscatter [Fig. 4(c)], and large pedestal for the polarization responses [Fig. 6(h)].

### B. Test Statistic

Figs. 7 and 8 show “ $-\ln Q$ ” (16) for the azimuthally symmetric case and the diagonal-element-only case, respectively, for the two images shown in Figs. 2 and 3. The test statistic is inverted to show areas with large change as bright areas and areas with small change as dark areas. Consequently, when “large values of the test statistic” is mentioned below, it means large values of “ $-\ln Q$ ” and vice versa for small values. We observe that especially the forest areas appear very dark, indicating virtually no change between the two acquisitions. For the agricultural fields, the results range from dark (no change) to bright (large change) areas depending on the crop type. Fig. 9(a) shows

the average “ $-\ln Q$ ” for the test areas outlined in the previous sections in the following four different cases:

- 1) using only the VV channel;
- 2) using only the three diagonal elements of the covariance matrix;
- 3) using the covariance matrix but assuming azimuthal symmetry;
- 4) using the full covariance matrix.

Furthermore, Fig. 9(b) shows the average probability of finding a larger value of “ $-2\rho \ln Q$ ” (derived from (19) and Theorem 6 in the Appendix) for the four cases mentioned above. Fig. 9(b) also indicates the 5% and the 1% significance levels, and the regions with probabilities lower than these levels are the regions where we will typically reject the hypothesis of equal covariance matrices (or VV channels) at the two points in time, i.e., these are regions with major change between the two data acquisitions.



Fig. 8. Logarithm of the test statistic  $\ln Q$  (16) for the test areas shown in Figs. 2 and 3 in the diagonal case.

Figs. 10 and 11 show in white for the images in Figs. 2 and 3, where the hypothesis of equal covariance matrix has been rejected at a 1% significance level for the azimuthally symmetric case and the diagonal case, respectively. Clearly, we observe detection of changes in the azimuthally symmetric case that have not been detected in the diagonal case, as well as improved detection in the azimuthally symmetric case of changes that already to some extent have been detected in the diagonal case.

In general, the test statistic for the full covariance matrix is only slightly larger than that for the assumed azimuthally symmetric case [Fig. 9(a)]. We may conclude that the additional information added by the co- and cross-elements of the covariance matrix is small. Also, the change detection potential of the single VV channel is seen to be much less than for the other three cases. Therefore, the discussion below will concentrate on comparing the results for two cases: the *diagonal case*, where only the three diagonal backscatter coefficient elements of the

covariance matrix are used, and the *polarimetric case*, where azimuthal symmetry is assumed (i.e., all the co- and cross-polarization elements are zero).

1) *Similar Polarimetric Parameters*: The two regions with virtually no change between the acquisitions but with different dominating scattering mechanisms, i.e., beets (surface scattering) and coniferous forest (volume scattering), show both large values of the test statistic and no significant difference between the diagonal and the polarimetric case. It is not possible to reject the hypothesis of equal covariance matrices at a 5% significance level for any of the regions [Fig. 9(b)]. The rye field also has very similar polarimetric parameters for the two acquisitions, except for  $\gamma_{hh}^0$ , as mentioned above, and the test statistic for the diagonal and the polarimetric cases are relatively close [Figs. 7–9(a)]. The hypothesis of equal covariance matrices is rejected at a 5% significance level for both cases. Consequently, in these cases with relatively similar

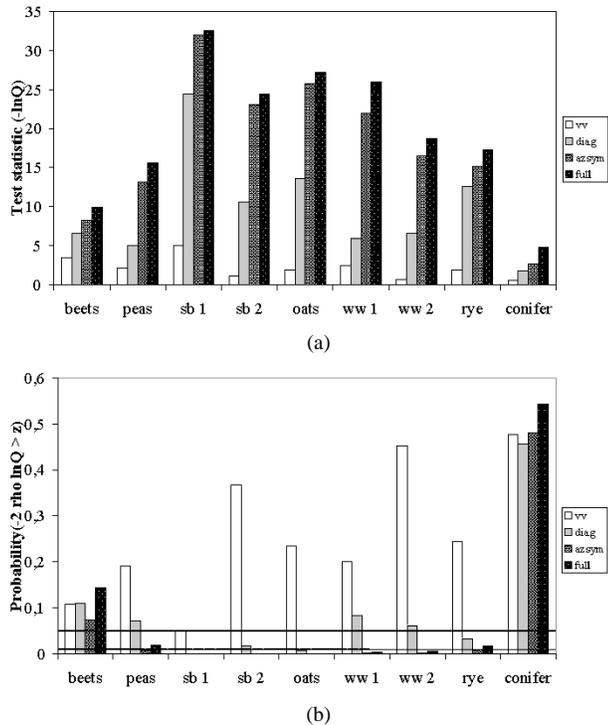


Fig. 9. (a) Average “ $-\ln Q$ ” for the test areas shown in Figs. 2 and 3 in four different cases: 1) using only the VV channel, 2) using only the three diagonal elements of the covariance matrix, 3) using the covariance matrix but assuming azimuthal symmetry, and 4) using the full covariance matrix. (b) Average probability of finding a larger value of “ $-2\rho \ln Q$ ” (derived from (19) and Theorem 6 in the Appendix) for the same four cases.

polarimetric parameters, the new test statistic for polarimetric data performs equally well as the nonpolarimetric test statistic.

2) *Similar Backscatter Coefficients, Large Difference for  $\rho_{hhvv}$  and/or  $\phi_{hh-vv}$* : Three fields have very similar backscatter coefficients for the two acquisitions, whereas a large difference between  $\rho_{hhvv}$  and/or  $\phi_{hh-vv}$  exists between the two acquisitions, i.e., the pea field (where  $\rho_{hhvv}$  decreases between the two acquisitions, due to the sparse vegetation cover at the May acquisition), the winter wheat 2 field (where  $\phi_{hh-vv}$  changes significantly between the two acquisitions), and the winter wheat 1 field (where both  $\rho_{hhvv}$  and  $\phi_{hh-vv}$  change significantly). A significantly larger test statistic is observed in the polarimetric case than in the diagonal case [Figs. 7–9(a)] for all three fields. Also, it is not possible to reject the hypothesis of equal covariance matrices at a 5% significance level for any of the three fields in the diagonal case [Fig. 9(b)]. On the other hand, the hypothesis is rejected in the polarimetric case at a 1% significance level for all three fields [Fig. 9(b)]. Thus, the results clearly show that the new polarimetric test statistic is very sensitive to differences in the full polarimetric information contained in the covariance matrix.

3) *Large Difference for All Polarimetric Parameters*: Finally, three regions show significant changes in all polarimetric parameters, i.e., the backscatter coefficients, the  $\rho_{hhvv}$ , and the  $\phi_{hh-vv}$ , i.e., the two spring barley fields and the oats field, which have a smooth bare surface at the April acquisition and a relatively dense vegetation cover at the May acquisition. For the spring barley 2 and the oats fields, we see a medium test statistic in the diagonal case, whereas a

much larger test statistic is observed for the polarimetric case [Figs. 7–9(a)]. The spring barley 1 field has a very large change in the backscatter coefficients, due to the relatively smooth surface at the April acquisition (Fig. 4), and we observe a very large test statistic for both the diagonal and the polarimetric cases [Figs. 7–9(a)]. The two spring barley fields have almost the same  $\rho_{hhvv}$  and  $\phi_{hh-vv}$  (Fig. 5), whereas the change in the backscatter coefficients is largest for the spring barley 1 field, as mentioned above. This difference is clearly important for both test statistics, where the test statistic for both the diagonal and the polarimetric case is much larger for the spring barley 1 field than for the spring barley 2 field. The hypothesis of equal covariance matrices is rejected for all three fields at the 5% significance level in both the diagonal and the polarimetric case. This is also the case at the 1% significance level, except for the spring barley 2 field in the diagonal case. Consequently, even when large changes in the backscatter coefficients ensure detection with a nonpolarimetric method, the addition of polarimetric information improves the detection of changes with the new polarimetric test statistic.

## V. CONCLUSIONS

In this paper, a test statistic for equality of two complex Wishart distributed covariance matrices and an associated asymptotic probability measure for obtaining a smaller value of the test statistic have been derived. The test statistic provides a unique opportunity to develop optimal algorithms for change detection, edge detection, line detection, segmentation, etc., for polarimetric SAR images. Such algorithms have previously been based on results of applying algorithms to the single channels and subsequently combining these results using some kind of fusion operator.

As a demonstration of the potential of the new test statistic, the derived measures have been applied to change detection in fully polarimetric SAR data for a test area with primarily agricultural fields and forest stands where two images acquired with approximately one-month interval have been used. In the case with areas with only small change in the polarimetric parameters between the two acquisitions, the new test statistic for polarimetric data performs equally well as the nonpolarimetric test statistic. When the backscatter coefficients are virtually unchanged, but either the phase and/or the correlation coefficient between the HH and VV polarizations have changed, the results clearly show that the new polarimetric test statistic is much more sensitive to the differences than test statistics based only on the backscatter coefficients. Also, in the case where all parameters in the covariance matrix have changed between the two polarizations, the new test statistic shows improved change detection capability. Consequently, the results show clearly that the new test statistic offers improved change detection capability for fully polarimetric SAR data.

## APPENDIX ANALYSIS OF COMPLEX WISHART MATRICES

In change detection and edge detection in polarimetric SAR data, it is useful to be able to compare two complex Wishart distributed matrices.



Fig. 10. Rejection of hypothesis of equal covariance matrices at 1% level for the assumed azimuthally symmetric case (white: rejection, black: acceptance).

Most of the standard literature in multivariate analysis only contains references to the real case (e.g., see [16]). This does not mean, however, that results for the complex case do not exist. In [4], the relevant class of complex distributions is introduced, and [17] completed much of the work, either giving results or (indirectly) pointing out how results may be obtained. It is, however, not straightforward to deduce the relevant formulas from their work.

In [18], many of the necessary formulas are deduced in an elegant way using the fact that the problem is invariant under a group of linear transformations. The notation chosen is not straightforward however. Some of their results appear in [19], an unpublished thesis in Danish, including results on comparing covariance matrices. In [20], the theory for linear and graphical models in multivariate complex normal models is covered.

Since the formulas for the distribution of the likelihood ratio do not seem to be available and since no authors seem to have

treated the so-called block diagonal case, we have chosen to give a rather thorough description of the necessary results.

We start with a short introduction to the complex normal and the complex Wishart distributions. We then compare two gamma distributions, which is the one-dimensional test often used in the radar community. Then we give a straightforward (brute force) derivation of the likelihood-ratio criterion for testing the equality in the complex case. Then, we describe the so-called block diagonal case, which among other things covers the case known in the radar community as the azimuthal symmetric case and the total independence case.

After quoting results from [21] on asymptotic distributions, we establish the necessary results on moments by brute force integration. By straightforward but rather cumbersome calculations, the results in [21] yield the desired results. Alternatively, one may use results in [20] and derive expressions involving the product of beta distributed random variables.



Fig. 11. Rejection of hypothesis of equal covariance matrices at 1% level for the diagonal case (white: rejection, black: acceptance).

### A. Complex Normal Distribution

Following [4], we say that a  $p$ -dimensional random complex vector  $\mathbf{Z}$  follows a complex multivariate normal distribution with mean  $\mathbf{0}$  and dispersion matrix  $\Sigma$ , i.e.,

$$\mathbf{Z} = [Z_1 \cdots Z_p]^T \in N_C(\mathbf{0}, \Sigma) \quad (28)$$

if the frequency function is

$$\begin{aligned} f(\mathbf{z}) &= \frac{1}{\pi^p |\Sigma|} \exp \{-\mathbf{z}^{*T} \Sigma^{-1} \mathbf{z}\} \\ &= \frac{1}{\pi^p |\Sigma|} \exp \{-\text{tr} [\Sigma^{-1} \mathbf{z} \mathbf{z}^{*T}]\}. \end{aligned} \quad (29)$$

$\Sigma$  is Hermitian positive definite and of the form

$$\begin{bmatrix} \sigma_1^2 & (\alpha_{12} + i\beta_{12})\sigma_1\sigma_2 & \cdots \\ (\alpha_{12} - i\beta_{12})\sigma_1\sigma_2 & \sigma_2^2 & \cdots \\ (\alpha_{13} - i\beta_{13})\sigma_1\sigma_3 & (\alpha_{23} - i\beta_{23})\sigma_2\sigma_3 & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}. \quad (30)$$

In other words, we have for  $Z_i = X_i + iY_i$

$$E \{Z_j Z_j^*\} = E \{X_j^2 + Y_j^2\} = \sigma_j^2 \quad (31)$$

$$E \{Z_j Z_k^*\} = (\alpha_{jk} + i\beta_{jk})\sigma_j\sigma_k. \quad (32)$$

### B. Complex Wishart Distribution

We say that a Hermitian positive definite random  $p \times p$  matrix  $\mathbf{X}$  follows a complex Wishart distribution, i.e.,

$$\mathbf{X} \in W_C(p, n, \Sigma) \quad (33)$$

if the frequency function is

$$w(\mathbf{x}) = \frac{1}{\Gamma_p(n)} \frac{1}{|\Sigma|^n} |\mathbf{x}|^{n-p} \exp \left\{ -\text{tr} [\Sigma^{-1} \mathbf{x}] \right\} \quad (34)$$

where

$$\Gamma_p(n) = \pi^{p(p-1)/2} \prod_{j=1}^p \Gamma(n-j+1). \quad (35)$$

If confusion concerning  $n$  or  $\Sigma$  may arise, we write  $w(\mathbf{x}; n, \Sigma)$  rather than  $w(\mathbf{x})$ . The frequency function is defined for  $\mathbf{x}$  positive definite, and in evaluating integrals the volume element becomes

$$d\mathbf{x} = (dx_{11}, \dots, dx_{pp}) \prod_{k=1}^p \prod_{j=1}^{k-1} dx_{jkR} dx_{jkI} \quad (36)$$

where  $R$  and  $I$  denote real and imaginary parts. For further description and useful results on Jacobians, etc., see [4] and [17].

It is emphasized that the formulas for the complex normal and Wishart distributions differ from their real counterparts.

Estimation of  $\Sigma$  from a normal sample follows in Theorem 1.

*Theorem 1:* Let  $\mathbf{Z}_1, \dots, \mathbf{Z}_n$  be independent, complex normal random variables, i.e.,

$$\mathbf{Z}_i \in N_C(\mathbf{0}, \Sigma). \quad (37)$$

Then the maximum-likelihood (ML) estimator for  $\Sigma$  is

$$\hat{\Sigma} = \frac{1}{n} \sum_{j=1}^n \mathbf{Z}_j \mathbf{Z}_j^{*T} \quad (38)$$

and  $n\hat{\Sigma}$  is Wishart distributed

$$n\hat{\Sigma} \in W_C(p, n, \Sigma). \quad (39)$$

*Proof:* See [4]. ■

From Theorem 1, we easily obtain Theorem 2.

*Theorem 2:* Let  $\mathbf{X}$  and  $\mathbf{Y}$  be independent Wishart distributed matrices

$$\mathbf{X} \in W_C(p, n, \Sigma) \quad (40)$$

$$\mathbf{Y} \in W_C(p, m, \Sigma) \quad (41)$$

Then the sum will again be Wishart distributed, i.e.,

$$\mathbf{S} = \mathbf{X} + \mathbf{Y} \in W_C(p, n+m, \Sigma). \quad (42)$$

*Proof:* Straightforward. ■

### C. Test on Equality of Two Gamma Parameters

Let the independent random variables  $X$  and  $Y$  (which are real scalars) be gamma distributed  $G(n, \beta_x)$  and  $G(m, \beta_y)$ . The frequency function for  $X$  is

$$f(x; \beta_x) = \frac{1}{\Gamma(n)} \frac{1}{\beta_x^n} x^{n-1} \exp \left\{ -\frac{x}{\beta_x} \right\}, \quad x > 0 \quad (43)$$

and similarly for  $Y$ . The likelihood function for the parameters  $(\beta_x, \beta_y)$  thus becomes

$$L_x(\beta_x) L_y(\beta_y) = \frac{1}{\Gamma(n)} \frac{1}{\Gamma(m)} \frac{1}{\beta_x^n} \frac{1}{\beta_y^m} \cdot X^{n-1} Y^{m-1} \exp \left\{ -\frac{X}{\beta_x} - \frac{Y}{\beta_y} \right\} \quad (44)$$

and under the hypothesis  $\beta_x = \beta_y = \beta$ , we obtain

$$L(\beta) = \frac{1}{\Gamma(n)} \frac{1}{\Gamma(m)} \frac{1}{\beta^{n+m}} \cdot X^{n-1} Y^{m-1} \exp \left\{ -\frac{(X+Y)}{\beta} \right\}. \quad (45)$$

Taking the derivatives of the log likelihoods and setting them equal to zero, we obtain the ML estimates

$$\hat{\beta}_x = \frac{X}{n} \quad (46)$$

$$\hat{\beta}_y = \frac{Y}{m} \quad (47)$$

$$\hat{\beta} = \frac{(X+Y)}{(n+m)}. \quad (48)$$

Therefore the likelihood-ratio test statistic becomes

$$Q = \frac{L(\hat{\beta})}{L_x(\hat{\beta}_x) L_y(\hat{\beta}_y)} = \frac{(n+m)^{n+m}}{n^n m^m} \frac{X^n Y^m}{(X+Y)^{n+m}}. \quad (49)$$

The critical region is given by  $Q \leq q$  or

$$\frac{X^n Y^m}{(X+Y)^{n+m}} \leq c. \quad (50)$$

Straightforward calculations show that this critical region is of the form

$$\frac{Y}{X} \leq c_1 \text{ or } \frac{Y}{X} \geq c_2. \quad (51)$$

Since  $Y/X$  under the null hypothesis is distributed like Fisher's  $F$ , i.e.,

$$\frac{m}{n} \frac{Y}{X} \in F(2m, 2n) \quad (52)$$

$c_1$  and  $c_2$  may be determined by means of quantiles in the  $F$ -distribution.

### D. Test on Equality of Two Complex Wishart Matrices

We consider independent Wishart distributed matrices

$$\mathbf{X} \in W_C(p, n, \Sigma_x) \quad (53)$$

$$\mathbf{Y} \in W_C(p, m, \Sigma_y) \quad (54)$$

and wish to test the hypothesis

$$H_0 : \Sigma_x = \Sigma_y = \Sigma \text{ against } H_1 : \Sigma_x \neq \Sigma_y. \quad (55)$$

We have the likelihood functions

$$L_x(\boldsymbol{\Sigma}_x)L_y(\boldsymbol{\Sigma}_y) = \frac{1}{\Gamma_p(n)\Gamma_p(m)} \cdot |\boldsymbol{\Sigma}_x|^{-n}|\boldsymbol{\Sigma}_y|^{-m}|\mathbf{X}|^{n-p}|\mathbf{Y}|^{m-p} \cdot \exp\{-\text{tr}[\boldsymbol{\Sigma}_x^{-1}\mathbf{X} + \boldsymbol{\Sigma}_y^{-1}\mathbf{Y}]\} \quad (56)$$

$$L(\boldsymbol{\Sigma}) = \frac{1}{\Gamma_p(n)\Gamma_p(m)} \cdot |\boldsymbol{\Sigma}|^{-(n+m)}|\mathbf{X}|^{n-p}|\mathbf{Y}|^{m-p} \cdot \exp\{-\text{tr}[\boldsymbol{\Sigma}^{-1}(\mathbf{X} + \mathbf{Y})]\}. \quad (57)$$

The ML estimates are

$$\hat{\boldsymbol{\Sigma}}_x = \frac{1}{n}\mathbf{X} \quad (58)$$

$$\hat{\boldsymbol{\Sigma}}_y = \frac{1}{m}\mathbf{Y} \quad (59)$$

$$\hat{\boldsymbol{\Sigma}} = \frac{1}{n+m}(\mathbf{X} + \mathbf{Y}). \quad (60)$$

Therefore, the likelihood-ratio test statistic becomes

$$Q = \frac{L(\hat{\boldsymbol{\Sigma}})}{L_x(\hat{\boldsymbol{\Sigma}}_x)L_y(\hat{\boldsymbol{\Sigma}}_y)} = \frac{(n+m)^{p(n+m)}}{n^p m^p} \frac{|\mathbf{X}|^n |\mathbf{Y}|^m}{|\mathbf{X} + \mathbf{Y}|^{n+m}}. \quad (61)$$

Thus, we get the critical region

$$C = \left\{ \mathbf{X}, \mathbf{Y} \mid \frac{(n+m)^{p(n+m)}}{n^p m^p} \frac{|\mathbf{X}|^n |\mathbf{Y}|^m}{|\mathbf{X} + \mathbf{Y}|^{n+m}} \leq q \right\}. \quad (62)$$

### E. Tests in the Block Diagonal Case

In some applications (e.g., remote sensing), there are several independent Wishart matrices in each observation. They will often be arranged in a block diagonal structure like

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_1 & \mathbf{0}_{p_1 p_2} \\ \mathbf{0}_{p_2 p_1} & \mathbf{X}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{X}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_2 \end{bmatrix} \quad (63)$$

where  $\mathbf{X}_i \in W_C(p_i, n, \boldsymbol{\Sigma}_{x_i})$ . This covers the so-called azimuthal symmetric case and the case with independence of co- and cross-polarized signals. If we define

$$\boldsymbol{\Sigma}_x = \begin{bmatrix} \boldsymbol{\Sigma}_{x1} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{x2} \end{bmatrix} \quad (64)$$

it is important to note that  $\mathbf{X}$  is not Wishart distributed  $W_C(p_1 + p_2, n, \boldsymbol{\Sigma}_x)$ .

We now consider similar partitionings of  $\mathbf{X}$  and  $\mathbf{Y}$ , i.e., we have independent random matrices

$$\mathbf{X}_i \in W_C(p_i, n, \boldsymbol{\Sigma}_{x_i}) \quad (65)$$

$$\mathbf{Y}_i \in W_C(p_i, m, \boldsymbol{\Sigma}_{y_i}) \quad (66)$$

for  $i = 1, \dots, I$ . We have  $p = p_1 + \dots + p_I$  and set  $N = n + m$ . We want to test the hypothesis

$$\boldsymbol{\Sigma}_{x_i} = \boldsymbol{\Sigma}_{y_i} = \boldsymbol{\Sigma}_i, \quad i = 1, \dots, I \quad (67)$$

against all alternatives. The likelihood-ratio criterion becomes the product of the criteria, i.e.,

$$Q = \prod_{i=1}^I Q_i = \prod_{i=1}^I \frac{N^{p_i N}}{n^{p_i n} m^{p_i m}} \frac{|\mathbf{X}_i|^n |\mathbf{Y}_i|^m}{|\mathbf{X}_i + \mathbf{Y}_i|^{n+m}}. \quad (68)$$

Since the determinant of a block diagonal matrix is the product of the determinants of the diagonal elements, we obtain

$$Q = \frac{N^{pN}}{n^p m^p} \frac{|\mathbf{X}|^n |\mathbf{Y}|^m}{|\mathbf{X} + \mathbf{Y}|^{n+m}} \quad (69)$$

i.e., the same result as in the general case [see (61)]. Note, however, that the distribution has changed, since  $\mathbf{X}$  and  $\mathbf{Y}$  are no longer Wishart distributed.

### F. Large Sample Distribution Theory

In [21] (as quoted in [16]), a general asymptotic expansion of the distribution of a random variable whose moments are certain functions of gamma functions has been developed. We state the main result as a theorem, and we shall use it in determining the (asymptotic) distribution of the likelihood-ratio criterion.

*Theorem 3:* Let the random variable  $W$  ( $0 \leq W \leq 1$ ) have the  $h$ th moment

$$E\{W^h\} = K \left( \frac{\prod_{j=1}^b y_j^{y_j}}{\prod_{k=1}^a x_k^{x_k}} \right)^h \frac{\prod_{k=1}^a \Gamma(x_k(1+h) + \xi_k)}{\prod_{j=1}^b \Gamma(y_j(1+h) + \eta_j)} \quad (70)$$

where  $K$  is a constant so that  $E\{W^0\} = 1$  and  $\sum_{k=1}^a x_k = \sum_{j=1}^b y_j$ . For an arbitrary  $\rho$  we set

$$\beta_k = (1 - \rho)x_k \quad (71)$$

$$\epsilon_j = (1 - \rho)y_j. \quad (72)$$

The first three Bernoulli polynomials are denoted

$$B_1(h) = h - \frac{1}{2} \quad (73)$$

$$B_2(h) = h^2 - h + \frac{1}{6} \quad (74)$$

$$B_3(h) = h^3 - \frac{3}{2}h^2 + \frac{1}{2}h. \quad (75)$$

We define

$$f = -2 \left[ \sum_{k=1}^a \xi_k - \sum_{j=1}^b \eta_j - \frac{1}{2}(a - b) \right] \quad (76)$$

$$\omega_r = \frac{(-1)^{r+1}}{r(r+1)} \left\{ \sum_{k=1}^a \frac{B_{r+1}(\beta_k + \xi_k)}{(\rho x_k)^r} - \sum_{j=1}^b \frac{B_{r+1}(\epsilon_j + \eta_j)}{(\rho y_j)^r} \right\}. \quad (77)$$

If we select  $\rho$  so that  $\omega_1 = 0$ , we have

$$P\{-2\rho \ln W \leq z\} \simeq P\{\chi^2(f) \leq z\} + \omega_2 [P\{\chi^2(f+4) \leq z\} - P\{\chi^2(f) \leq z\}]. \quad (78)$$

*Proof:* See [16]. ■

*Theorem 4:* Let

$$\mathbf{X} \in W_C(p, n, \Sigma) \quad (79)$$

$$\mathbf{Y} \in W_C(p, m, \Sigma) \quad (80)$$

be independent complex Wishart distributed matrices. Then for  $N = n + m$  and

$$U = \frac{|\mathbf{X}|^n |\mathbf{Y}|^m}{|\mathbf{X} + \mathbf{Y}|^N} \quad (81)$$

we have

$$E\{U^h\} = \frac{\Gamma_p(n(1+h))\Gamma_p(m(1+h))}{\Gamma_p(N(1+h))} \frac{\Gamma_p(N)}{\Gamma_p(n)\Gamma_p(m)}. \quad (82)$$

*Proof:* The joint frequency function of  $(\mathbf{X}, \mathbf{Y})$  is  $w(\mathbf{x}; n, \Sigma)w(\mathbf{y}; m, \Sigma)$ . Therefore

$$E\{U^h\} = \int \int u^h w(\mathbf{x}; n, \Sigma)w(\mathbf{y}; m, \Sigma) d\mathbf{x} d\mathbf{y} \quad (83)$$

and evaluation of this integral gives the desired result. ■

By means of the two previous theorems, we are now able to state the important result on the (asymptotic) distribution of the likelihood-ratio criterion.

*Theorem 5:* We consider the likelihood-ratio criterion

$$Q = \frac{N^{pN}}{n^{pn}m^{pm}} \frac{|\mathbf{X}|^n |\mathbf{Y}|^m}{|\mathbf{X} + \mathbf{Y}|^N} \quad (84)$$

and define

$$f = p^2 \quad (85)$$

$$\rho = 1 - \frac{2f-1}{6\sqrt{f}} \left( \frac{1}{n} + \frac{1}{m} - \frac{1}{N} \right) \quad (86)$$

$$\omega_2 = -\frac{f}{4} \left( 1 - \frac{1}{\rho} \right)^2 + \frac{f(f-1)}{24} \cdot \left( \frac{1}{n^2} + \frac{1}{m^2} - \frac{1}{N^2} \right) \frac{1}{\rho^2}. \quad (87)$$

Then

$$P\{-2\rho \ln Q \leq z\} \simeq P\{\chi^2(f) \leq z\} + \omega_2 [P\{\chi^2(f+4) \leq z\} - P\{\chi^2(f) \leq z\}]. \quad (88)$$

*Proof:* Omitted, straightforward but cumbersome calculations. ■

We now address the block diagonal case and state the main result in Theorem 6.

*Theorem 6:* Let the situation be as described in Section A–E, and define for  $i = 1, \dots, I$

$$f_i = p_i^2 \quad (89)$$

$$\rho_i = 1 - \frac{2f_i-1}{6\sqrt{f_i}} \left( \frac{1}{n} + \frac{1}{m} - \frac{1}{N} \right) \quad (90)$$

$$f = f_1 + \dots + f_I \quad (91)$$

$$\rho = \frac{f_1}{f} \rho_1 + \dots + \frac{f_I}{f} \rho_I \quad (92)$$

$$\omega_2 = -\frac{f}{4} \left( 1 - \frac{1}{\rho} \right)^2 + \frac{1}{24} \cdot [f_1(f_1-1) + \dots + f_I(f_I-1)] \left( \frac{1}{n^2} + \frac{1}{m^2} - \frac{1}{N^2} \right) \frac{1}{\rho^2}. \quad (93)$$

Then the asymptotic distribution of the likelihood-ratio criterion is given by

$$P\{-2\rho \ln Q \leq z\} \simeq P\{\chi^2(f) \leq z\} + \omega_2 [P\{\chi^2(f+4) \leq z\} - P\{\chi^2(f) \leq z\}]. \quad (94)$$

*Proof:* Omitted, straightforward but cumbersome calculations. ■

## REFERENCES

- [1] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing: Active and Passive*. Norwood, MA: Artech House, 1986, vol. 3.
- [2] H. Skriver, M. T. Svendsen, and A. G. Thomsen, "Multitemporal L- and C-band polarimetric signatures of crops," *IEEE Trans. Geosci. Remote Sensing*, vol. 37, pp. 2413–2429, Sept. 1999.
- [3] J. J. van Zyl and F. T. Ulaby, "Scattering matrix representation for simple targets," in *Radar Polarimetry for Geoscience Applications*, F. T. Ulaby and C. Elachi, Eds. Norwood, MA: Artech House, 1990.
- [4] N. T. Goodman, "Statistical analysis based on a certain multivariate complex Gaussian distribution (An introduction)," *Ann. Math. Stat.*, vol. 34, pp. 152–177, 1963.
- [5] S. V. Nghiem, S. H. Yueh, R. Kwok, and F. K. Li, "Symmetry properties in polarimetric remote sensing," *Radio Sci.*, vol. 27, no. 5, pp. 693–711, 1992.
- [6] A. A. Nielsen, R. Larsen, and H. Skriver, "Change detection in bi-temporal EMISAR data from Kalø, Denmark, by means of canonical correlations analysis," in *Proc. 3rd Int. Airborne Remote Sensing Conf. and Exhibition*, Copenhagen, Denmark, July 7–10, 1997.
- [7] A. A. Nielsen, "Change detection in multispectral bi-temporal spatial data using orthogonal transformations," in *Proc. 8th Austral-Asian Sensing Conf.*, Canberra ACT, Australia, Mar. 25–29, 1996.
- [8] A. A. Nielsen and K. Conradsen, "Multivariate alteration detection (MAD) in multispectral, bi-temporal image data: A new approach to change detection studies," Dept. Mathematical Modelling, Tech. Univ. Denmark, Lyngby, Denmark, Tech. Rep. 1997-11, 1997.
- [9] A. A. Nielsen, K. Conradsen, and J. J. Simpson, "Multivariate alteration detection (MAD) and MAF post-processing in multispectral, bi-temporal image data: New approaches to change detection studies," *Remote Sens. Environ.*, vol. 19, pp. 1–19, 1998.
- [10] A. A. Nielsen, "Multi-channel remote sensing data and orthogonal transformations for change detection," in *Machine Vision and Advanced Image Processing in Remote Sensing*, I. Kanellopoulos, G. G. Wilkinson, and T. Moons, Eds. Berlin, Germany: Springer, 1999, pp. 37–48.
- [11] J. Schou, H. Skriver, K. Conradsen, and A. A. Nielsen, "CFAR edge detector for polarimetric SAR images," *IEEE Trans. Geosci. Remote Sensing*, vol. 41, pp. 20–32, Jan. 2003.
- [12] S. N. Madsen, E. L. Christensen, N. Skou, and J. Dall, "The Danish SAR system: Design and initial tests," *IEEE Trans. Geosci. Remote Sensing*, vol. 29, pp. 417–476, May 1991.
- [13] E. L. Christensen, N. Skou, J. Dall, K. Woelders, J. H. Jørgensen, J. Granholm, and S. N. Madsen, "EMISAR: An absolutely calibrated polarimetric L- and C-band SAR," *IEEE Trans. Geosci. Remote Sensing*, vol. 36, pp. 1852–1865, Nov. 1998.

- [14] R. Touzi, A. Lopes, and P. Bousquez, "A statistical and geometrical edge detector for SAR images," *IEEE Trans. Geosci. Remote Sensing*, vol. 26, pp. 764–773, Nov. 1988.
- [15] A. Lopes, E. Nezry, R. Touzi, and H. Laur, "Structure detection and statistical adaptive speckle filtering in SAR images," *Int. J. Remote Sens.*, vol. 13, no. 9, pp. 1735–1758, 1993.
- [16] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*, 2nd ed. New York: Wiley, 1984.
- [17] C. G. Khatri, "Classical statistical analysis based on a certain multivariate complex Gaussian distribution," *Ann. Math. Stat.*, vol. 36, pp. 98–114, 1965.
- [18] S. A. Andersson, H. K. Brøns, and S. T. Jensen, "Distribution of eigenvalues in multivariate statistical analysis," Inst. Mathematical Statistics, Univ. Copenhagen, Copenhagen, Denmark, 1982.
- [19] G. Gabrielsen, "Fordeling af egenverdier i reelle, komplekse og kvaternion Wishart fordelinger," M.S. thesis (in Danish), Inst. Mathematical Statistics, Univ. Copenhagen, Copenhagen, Denmark, 1975.
- [20] H. H. Andersen, M. Højbjerg, D. Sørensen, and P. S. Eriksen, *Linear and Graphical Models for the Multivariate Complex Normal Distribution*. Berlin, Germany: Springer-Verlag, 1995, vol. 101, Lecture Notes in Statistics.
- [21] G. E. P. Box, "A general distribution theory for a class of likelihood criteria," *Biometrika*, vol. 36, pp. 317–346, 1949.



**Knut Conradsen** received the M.S. degree from the Department of Mathematics, University of Copenhagen, Copenhagen, Denmark, in 1970.

He is currently a Professor with the Informatics and Mathematical Modelling, Technical University of Denmark (DTU), Lyngby, Denmark. Since 1995, he has been Deputy Rector (or Vice President) of DTU. His main research interest is the application of statistics and statistical models to real-life problems. He has worked on many national and international projects on the application of statistical methods

in a wide range of applications, recently mainly in remote sensing. He has also conducted extensive studies in the development and application of mathematical and statistical methods concerning spatial, multi/hyperspectral, and multitemporal data, including both optical and radar data.



**Allan Aasbjerg Nielsen** received the M.S. degree from the Department of Electrophysics, Technical University of Denmark, Lyngby, Denmark, in 1978, and the Ph.D. degree from Informatics and Mathematical Modelling (IMM), Technical University of Denmark, in 1994.

He is currently an Associate Professor with IMM. He is currently working with IMM's Section for Geoinformatics. He has been with the Danish Defense Research Establishment from 1977 to 1978. He has worked on energy conservation in

housing with the Thermal Insulation Laboratory, Technical University of Denmark, from 1978 to 1985. Since 1985, he has been with the Section for Image Analysis, IMM. Since then, he has worked on several national and international projects on the development, implementation, and application of statistical methods and remote sensing in mineral exploration, mapping, geology, environment, oceanography, geodesy, and agriculture funded by industry, the European Union, Danida (the Danish International Development Agency), and the Danish National Research Councils.



**Jesper Schou** received the M.S. degree in engineering and the Ph.D. degree in electrical engineering from the Technical University of Denmark, Lyngby, Denmark, in 1997 and 2001 respectively. His primary research interests during his Ph.D. studies was image analysis and the processing of polarimetric SAR data, including filtering, structure detection, and segmentation.

He is currently working with 3-D scanner systems and related 3-D image analysis algorithms.



**Henning Skriver** received the M.S. degree and the Ph.D. degree from the Technical University of Denmark, Lyngby, in 1983 and 1989, respectively, both in electrical engineering.

He has been with the Section of Electromagnetic Systems (EMI), Department Ørsted, Technical University of Denmark (DTU), Lyngby, since 1983, where he is currently an Associate Professor. His work has been concerned primarily with various topics related to the utilization of SAR data for different applications. From 1983 to 1992, his main

area of interest was retrieval of sea ice parameters from SAR data, including SAR data from ERS-1. Since 1992, he has covered different aspects of land applications of SAR data, such as forestry in the MAESTRO-1 project, and agricultural and environmental applications using both satellite SAR data and data from the Danish airborne polarimetric SAR, EMISAR. His interests also include various methods for processing of SAR data, such as SAR image simulation, SAR image filtering, speckle statistics, texture analysis, segmentation, calibration, and polarimetric analysis. He is currently a Project Manager for a project concerning the use of SAR data for cartographic mapping.

# Using the LZW to Detect Primitives

Lars Reng,

Laboratory of Computer Vision and Media Technology,  
Aalborg University

## I. Abstract

*This paper is the result of some early attempts to find methods for detection of primitives. The primitive problem is view in a simple general case that should apply to almost any domain. A hypothesis stating that the LZW algorithm can be used for solving such a problem in a simple domain is made. Some of the problem areas are discussed, and possible solutions suggested. Several test-scenarios are mentioned, but due to lack of final test results no final conclusions can be made. The paper aims not at creating a landmark in this area, but is instead intended as a basis for discussion and creative exchange of ideas.*

## II. Introduction

The use of speech-primitives has in recent years resulted in major improvements in the area of speech-recognition. The primitives are not only used in speech recognition, but have also resulted in better speech synthesizers. Due to this success, there are now being made great efforts in researching if the primitive-concept can be copied to other areas of automatic recognition and synthesizer development.

This paper describes a hypothesis discovered during the earliest stages of a project aiming at developing motion-primitives for describing upper-body body-language. As a pilot research-project the task of deriving primitives will be attempted on a number of simple data sequences. The goal is to find a faster and smarter way to derive the primitives, than using a simple but slower exhaustive search method.

Some types of compression algorithms rely on recurring patterns in the input data sequence, and actually build a dictionary containing the patterns they find. Using parts of such an algorithm might speed up the analysis process of much bigger and more complex sequences. It has therefore been decided to compare the patterns found by both a simple exhaustive search and such an algorithm.

The algorithm chosen for this test is the LZW<sup>1</sup> [3], which is probably the most famous modified version of the landmark LZ<sup>2</sup> algorithms. This paper intends to test the hypothesis that the LZW, which will be described in more detail later in this paper, can find the primitives of a test-sequence given by a normal ASCII string. The paper will also try to cover some of the problems/concerns that the LZW could cause.

## III. The Simple Exhaustive Search

In order to evaluate the results of the LZW algorithms, all training data are also analyzed by a simple exhaustive search. The exhaustive search should be able to find the number of occurrences for all possible sub-sequences combinations, and there by supply a 'ground

---

<sup>1</sup> Lempel, Ziv, and Welch.

<sup>2</sup> The two famous LZ77 and LZ78 made by Lempel and Ziv.

truth' for a given sequence. The runtime for the exhaustive search is also used for comparison, since one of the incentives for trying to use a LZW algorithm was to decrease the runtime.

The exhaustive search will run through the input string (S) for each window-size<sup>3</sup> of (1..n) counting the occurrences of every sub-string. Figure 1 seen below, shows an example of the exhaustive search to illustrate the principle and allow a comparison with the figures illustrating the LZ principles. The input string (S) has a value of (001010210210212021021200), and the input sequence is analyzed for patterns of size = 2.

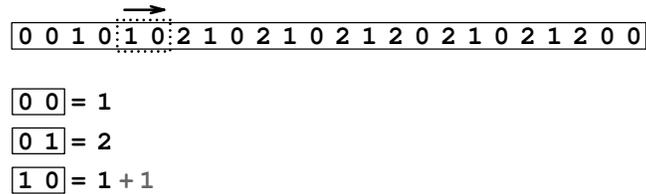


Figure 1 : Exhaustive search.

## IV. The LZ Algorithms

The original LZ77 and LZ78 algorithms was developed and published by Lempel and Ziv back in 1977 and 1978 [1],[2]. They belong to a group of compression algorithms known as 'adaptive dictionary' algorithms. The dictionary coding technique relies upon correlation between parts of data, and is specifically build to detect recurring patterns. It is this characteristic that we hope to benefit form, and not the algorithms ability to compress data.

The first following parts of this paper will describe the LZ77, LZ78 and LZW in more detail, and give a simple example to show the basic principles of these algorithms. To a reader already experienced with these techniques, the parts will bring no new information.

### IV - i. The LZ77 Algorithm

Even though the LZ77 is considered to be an 'adaptive dictionary algorithm' is does in fact not create a dictionary. Instead it uses a 'sliding window' technique, as depicted in Figure 2, to find the recurring patterns and compress the data sequence in a single pass. The sliding window consists of two parts, a lookahead buffer which contains the part of the sequence that has not yet been encoded, and the search buffer which contains the last part of the sequence that has already been encoded.

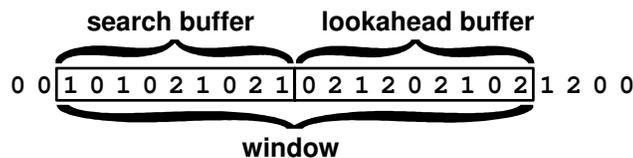


Figure 2 : Sliding window.

The size of each buffer can be set to any given value. A small buffer-size will result in a fast but very pour compression, as a big buffer-size will result in a better compression but also longer run-time. The size of the search buffer is often bigger than the size of the

<sup>3</sup> The value of (n) has in principle only the size of the input-string as limit. But will in practice be set to a value much smaller.

lookahead buffer, since the lookahead buffer is always encoded from the leftmost value, while a match can be found anywhere in the search buffer. The algorithm starts by filling the lookahead buffer with the first values of the input sequence, while the search buffer is filled with a predefined value (i.e. zeros). The left most values of the lookahead buffer is then compared to the values of the search buffer, so the biggest possible match between the leftmost values of the lookahead buffer is matched to a similar sub-string beginning<sup>4</sup> in the search buffer. If a match is found, the position of the search-match counted from the left of the search buffer is appended to the output string with the size of the match, and the value right after the match in the lookahead buffer (the first none-matched value). If no match is found the position and length is set to zero.

Figure 3 will show an example taken from the original LZ77 paper<sup>5</sup>. The input sequence = (001010210210212021021200...), the search and lookahead buffer-size = (9).

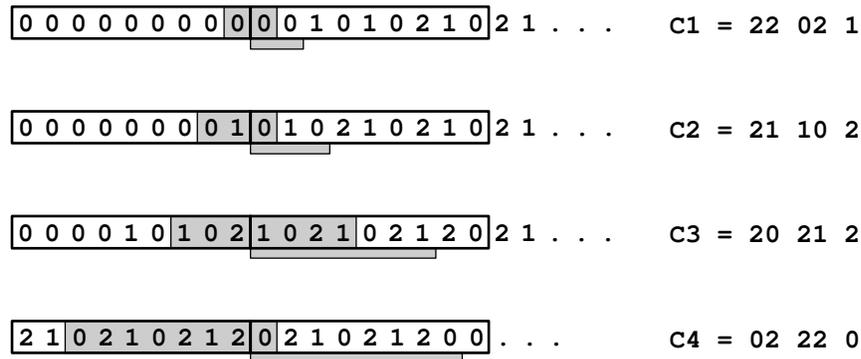


Figure 3 : LZ77 encoding.

There have been made several modified versions of the LZ77. Some of these combine the LZ77 with Huffman encoding, while others alter certain parts to improve the compression when working on a certain type of data. One of the improvements, the LZSS by Storer and Szymanski, avoids adding the first none-matched value to output unless no match is found. Another change made in LZSS is that the first none-matched value is part of the next match.

#### IV - ii. The LZ78 Algorithm

The LZ78 is a real adaptive dictionary-based compression algorithm, and has no sliding windows. It builds a dictionary based on the input sequence, and its only output is a pointer to the longest matching dictionary entry, and the first none-matching element. If no match is found the output consists of an index value of zero and the first none-matching element. The output pair is also added to the dictionary as a new entry, unless the index is zero. In that case only the none-matching element is added.

Figure 4 shows an example of a LZ78 encoding. (Input sequence = 001212121021012101221011).

<sup>4</sup> The match may extend into the lookahead buffer.

<sup>5</sup> Note that all values are written in ternary numbers (radix-3).

Dictionary			Output
#	entry	phrase	
1	0	0	0 0
2	1+1	01	1 1
3	2	2	0 2
4	1	1	0 1
5	3+1	21	3 1
6	5+0	210	5 0
7	6+1	2101	6 1
8	7+2	21012	7 2
9	7+1	21011	7 1

Figure 4 : LZ78 Encoding

The LZ78 has a few weaknesses; one of them is size of the dictionary. If the input sequence is very large and contains a high level of variation, the dictionary might also become extremely large. This problem can be solved by either start over on a new dictionary or simply stop adding new entries. Another problem is that the LZ78 like the LZ77 always adds the first none-matched value to output, and jumps over it when beginning to match the remaining part of the sequence. One of the many modified versions of the LZ78 (the LZW algorithm) takes care of this problem.

#### IV - iii. The LZW Algorithm

The LZW, introduced by Terry Welch in 1984 [3], is a modified version of the LZ78. Where the LZ78 adds both a pointer to a dictionary entry and the first none-matched value to output, the LZW applies the LZSS principle by *only* adding entry-pointers. Without any other changes, it would be impossible to decode the output signal; so to compensate for the lack of sequence values, the dictionary is initialized with all the basic symbols of the input alphabet. The LZW also uses the LZSS principle of not jumping over the first none-matched value when conducting the new match. Besides these few alterations, the LZW works just like the LZ78.

### V. Using the LZW as a Search Algorithm

The LZW algorithm is like the rest of the LZ algorithms designed to compress an input signal as much as possible, and not to keep track of how many times the same pattern occurs in the input sequence. This might be a major problem then trying to use the algorithm for analyzing a sequence to find primitives. Each time a match is made, there are also added a new entry into the dictionary. This means that even though the LZW detects the right recurring pattern, it will at some point also have an entry of that pattern with many different symbols at the end, and therefore most certainly those to match the (pattern+symbol) from the input sequence to a similar (pattern+symbol) in the dictionary. This has the unfortunate outcome, that the same pattern is match with many different *names* (entries) and therefore *not* recognized as being the same recurring pattern.

**The way the standard LZW is designed, it might also have a lot of similar entries, that only differs in the first symbol, see example in**

Figure 5. These similar *sub-strings* will not even be matched once, and can therefore be completely missed as potential primitives. The last and final problem is that a recurring pattern might be divided into two different entries, as a result of the ever expanding

dictionary. This problem might be the hardest to counteract, since this requires a reassembling of the two pairs of successive entries.

3	4	4	6	6	8	8	
5	4	4	6	6	8	8	
1	4	4	6	6	8	8	
7	1	4	4	6	6	8	8

Figure 5 : Unrecognizable entries.

### V - i. Extracting the “Lost” Sub-Strings

The LZW is a modified version of the LZ78, with some improved abilities. The LZW is design for compressing data and not for analyzing a sequence to detect primitives, so modifying the LZW even further might improve it quite a bit.

If a small search for similar sub-strings were conducted on the entries of the dictionary this should solved the problem of different *names* for similar patterns, since the number of occurrences for each entry containing the sub-string can be added to get the total number of occurrences of the sub-string.

Other LZW problems might also be solved by adding extra functionality, taking into consideration that it might have unfortunate effects on the algorithms run-time.

### V - ii. Comparing the LZW and Exhaustive Search Algorithm

The paper should have included an extensive comparison between the exhaustive search approach and the modified versions of the LZW. Test-results were, however, not available before deadline of this paper. These results are instead scheduled to be presented at the DanCom conference in August.

## VI. Discussion and Conclusion

Even though a final conclusion to this experiment is impossible due to the lack of final test results, some conclusions might be attempted never the less. This paper is the result of the every first steps of a project aiming at finding a set of primitives to describe upper-body body-languish. The preliminary testing has enlightened many problems/concerns towards using an algorithm build for a somewhat different task.

The goal of the experiment is not to get complete statistics of how many occurrences of how many recurrent patterns is included in the input sequence, but instead, if the LZW can be used to find primitives. The final test result might surprisingly show that the LZW is superior to the exhaustive search in selecting the ‘right’ entries.

## VII. References

- [1] Ziv, J., and Lempel, A. A Universal Algorithm for Sequential Data Compression. IEEE Transactions on Information Theory 23 (1977), 337-343.
- [2] Ziv, J., and Lempel, A. A Compression of Individual Sequences via Variable-Rate Coding. IEEE Transactions on Information Theory 24 (1978), 530-536.
- [3] Welch, Terry. A Technique for High-Performance Data Compression. Computer, June 1984.

# Medical image sequence coding

Mo Wu and Søren Forchhammer

Research Center COM, Technical University of Denmark

Abstract:

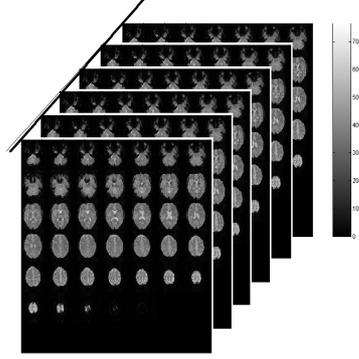
Wavelet-based methods have become popular for two-dimensional (2D) image compression. It has also been used in 3D image sequence compression, especially for medical image sequence compression. Our contribution in this paper is to present a method to compress 4D functional Magnetic Resonance Imaging (fMRI) data (16 bit per pixel) in a 3D setting. In order to acquire lossless compression and improve compression ratio, different reversible integer to integer wavelet transforms have been tried out. In the entropy coding part, we use context based adaptive arithmetic coding and extend the 2D contexts to 3D by including the information from the third dimension slices. The compression ratio can be increased by the new 3D context. A Head Object Oriented Lossless Compression (HOOLC) for 4D fMRI data is also introduced, which can also be applied for other types of objects in lossless compression.

Index Terms - fMRI, wavelet based image coding, medical image compression, lossless compression, adaptive arithmetic coding, HOLC, HOOLC.

## 1. Introduction

Medical ultrasound data, Computer Tomography data (CT) and functional Magnetic Resonance Imaging (fMRI) data may have very large volumes and can only be transmitted efficiently by using compression. They include tremendous volumes of images with a high level of noise as well as a high correlation between them. Using wavelet based image sequence coding can decrease the volume of the data significantly and also provide progressive coding. Based on the lifting [6] scheme integer to integer wavelet transforms can be constructed [5], making efficient lossless compression possible, e.g. for medical image compression.

The fMRI data we consider is 4D data. In the 3D spatial domain, it is composed of slices at different positions. Including time as the fourth dimension, the data become 4D. Functional MR data has its own property compared with e.g. video sequences. It has much more noise than normal video. The noise comes from the capturing and processing of the acquisition device. So even in areas of the image, which is supposed to be (almost) zero, there may be unexpected values. Viewing the images by naked eyes, sometimes we cannot tell the difference between two slices at different time. The popular technique, motion prediction [12] used in video technology, is not suitable for this condition. Lossless compression is very important for medical applications, because the subtle differences can give us information about the patients. But on the other hand many bits are most likely of little importance. So, for the medical data compression, traditional methods need to be revised. Firstly we use 3D wavelets based compression for the 4D dataset ( $64 \times 64 \times 49 \times 32$ ), which has been reformed to 3D setting ( $448 \times 448 \times 32$ ) by aligning slices at different positions and the same time to one slice [see Fig. 1]. Here the data is 2 bytes per pixel. Thereafter optimized context based coding will be used for the 4D data compression.



**Figure 1:** 4D fMRI medical image sequence data in 3D representation

A popular wavelet-based embedded image coding is Set Partitioning In Hierarchical Trees (SPIHT) [10], by Said and Pearlman, which improves upon Embedded Zerotree Wavelet (EZW) [11]. The coding efficiency comes from exploring interband dependencies between the wavelet coefficients. It assumes that the significant status of one region in a subband is highly correlated to its parent node. As a consequence, if the status of descendent subband coefficients is not dependent on its lower bands' status, the cost will be more. 3D SPIHT [19] was developed for 3D data compression.

Another popular coder is the context based layered zero coding [13], which is the core element of JPEG2000 [1] image compression standard. Compared with the zerotree structure coding [10], [11], it concentrates on pursuing the dependency inside the individual subbands. The context of the zero coding will be refined for the 3D structure.

The type of data is important information for us when performing the compression. We can use the extra information to know the interested region of the images and preprocess the data to improve the compression ratio.

The rest of this paper is organized as follows: 3D packet integer to integer wavelet transform is covered in Section 2. Section 3 will talk about 3D context based adaptive arithmetic block coding, concentrating on how to extend the 2D context to 3D. Section 4 will show the object based 3D wavelet coding. It combines the 3D contexts with object information to increase compression of the 4D datasets. Experimental results, in both lossy and lossless cases, are given in Section 5. Conclusions are drawn in Section 6.

## 2. 3D packet integer to integer wavelet transform

Integer to integer lifting is used to construct reversible 1D wavelet transforms. 3D packet [3] integer to integer transform is constructed based on the 1D wavelet transforms.

### 2.1. Reversible integer to integer wavelet transform

Traditional (discrete) wavelet transform decomposes one sequence of data into subbands by using low pass and high pass filters. The data can be further decomposed in subbands to de-correlate the data.

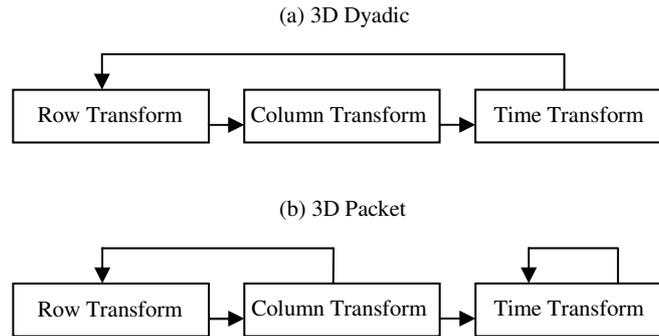
The original medical data is represented by integer numbers (or finite precision). By using the normal wavelet transforms, such as the biorthogonal Daubechies 9/7, output wavelets coefficients will be real numbers, which can not be coded efficiently without loss. But for medical applications, only minor mistakes in the medical data may result in serious problems. Thus, efficient lossless compression requires an integer to integer wavelet transform. The lifting technique is a good way to construct integer to integer wavelet transforms [5], [6].

Every lifting step can be rounded to integer representation, and the process can be reversible. Thus the wavelets transform can be converted to an integer transform, if it is constructed by integer lifting steps. There are many integer transforms [14-16], which can be used. Here we use the filters: (2, 2), (2, 4), (4, 2), (4, 4), (6, 2), (2+2, 2), (2/6), (9/7-F) and (9/7)

These lifting scheme filters do not include the scaling part for the forward transform, except the (9/7). The scaling part [3] can be split into two integer lifting steps and maintain that the transform is reversible. But the scaling is used to make the transform unitary. It is also possible to do the transform without this part in lossless compression, making the transformation simpler and more robust towards lossy conditions, because the more integer lifting steps, the more the (rounding) error will spread from the code to the reconstructed images. For example, the lowest band of a 4 level 2D dyadic wavelets transform [9] combined with a 2 level packet transform [3] in the third dimension will experience 2 times scaling in the third dimension and 4 times scaling in the other two dimensions (see below), so the extra integer lifting steps will be  $2 \times 2 \times 4 = 16$  times. The rounding process at each step will exacerbate the quality of the reconstructed images if truncating for lossy coding..

## 2.2. 3-D packet wavelets transform

We do the spatial 2D dyadic wavelet transform for each slice and then do the wavelet transform in the third dimension. It is different from the 3D dyadic wavelet transform if the number of decomposition levels in the spatial domain is more than one. The order of the steps in the transformation is shown on Fig. 2.



**Figure 2:** The order of the wavelet transform in the three dimensions for: (a) 3D dyadic transformation and (b) 3D packet transform

Different wavelets are good at decorrelating different textures. The relationship between the neighbor pixels in the special domain is different from that in the time domain. Different wavelets can be used for the spatial domain from the one for the time domain.

## 3. 3D Context based adaptive arithmetic block coding

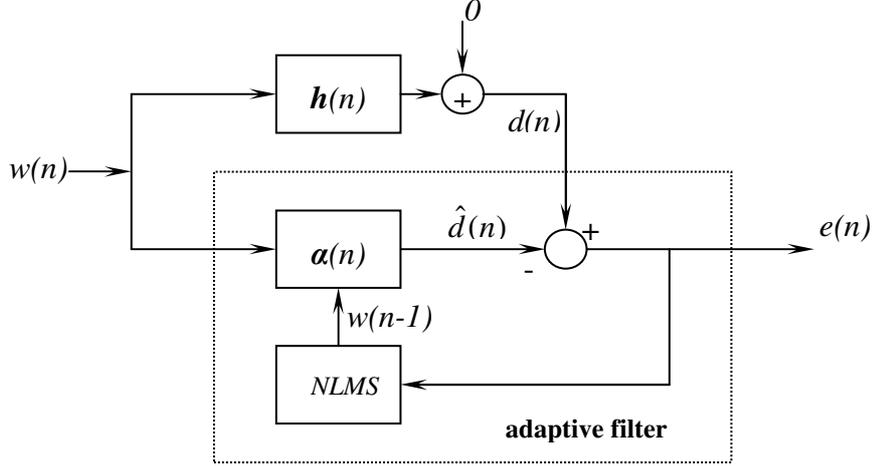
In the 3D context based adaptive arithmetic block coding, we did the entropy coding for the wavelet coefficients based on the Zero Coding (ZC) primitive, Sign Coding primitive (SC) and Magnitude Refinement primitive (MR), which are the core of the Embedded Block Coding by Optimized Truncation (EBCOT) [13] technique for image coding. The basic coding contexts have been extended to 3D by including the status from the front slice and back slice and the most current status from the front point in the third dimension. The main target is finding the best contexts to achieve a higher compression ratio.

### 3.1. Zero coding, Sign coding, Magnitude refinement

We did the entropy coding for the wavelet coefficients by using ZC, SC and MR primitives and thereafter include context information from the front slice and the back slice. There are 17 contexts in total for the normal 2D contexts model

### 3.2. Extension to 3-D contexts

Based on the 2D context model, we extend the contexts to 3D by including the information from the front slice and the back slice. We use a Normalized Least-Mean-Square (NLMS) algorithm [2] [see Fig. 3] to train  $\alpha_i, 1 \leq i \leq 18, i \in Z$  to predict the wavelets coefficient with the 18 neighbors in these two slices. The estimator [3] is defined as,  $\Delta = \sum_{w_i \in S} \alpha_i |\omega_i|$ ,  $\omega_i$  is the wavelets coefficients at the front and the back slice.



**Figure 3:** Adaptive filter structure for linear prediction.  $w(n)$  is the wavelets coefficients;  $d(n)$  is the current wavelet coefficient we need to predict;  $\hat{d}(n)$  is the predicted value;  $e(n)$  is the error in prediction.

$$\alpha_{k+1} = \alpha_k + \mu \frac{e_k \omega_k}{\sum_{m=1}^L \omega_{m,k}^2}$$

We train  $\alpha_i$  for the different subbands in the third dimension and thereafter simply use the mean results of the different subbands as  $\alpha_i$ . If we were to use linear prediction, we would (implicitly) assume the signal is stationary, unfortunately, this is not the case, and  $\alpha_i$  should not converge to one value, but rather fluctuate around some value. So the mean value is used for an approximate division of the original contexts. This prediction is only used for extending the zero coding contexts, because there is no relation with the sign and little relation with magnitude refinement. The estimator is quantized to 2 intervals or 4 intervals. When we further extend these 3D contexts, we only use 2 intervals here. But this prediction is only used to extend the ZC contexts, because this prediction almost has no relationship with the MR and no relationship with SC.

### 3.3. Scan order in coding the blocks

In the coding process, we use a block coding method, but without making the coding independent for every block with respect to the other blocks, not only because we want to extend the contexts to 3D, but also as we can use information from the other blocks. If we have more information available, we can predict the coded value more accurately.

So, in this condition, we want to initialize the counts of the statistics for adaptive arithmetic coding by inheriting the counts from the most similar neighbor block. One way to evaluate the similarity is by the Kullback-Leibler distance [7] [8] between the coded block and its

neighbour block. The blocks that have the smallest distance have the most similarity of their statistics. The average Kullback-Leibler distance  $D$  for this model is defined as:

$$\left\{ \begin{array}{l} D = \frac{\sum_c (n_0^{(c)} + n_1^{(c)}) d_c}{\sum_c (n_0^{(c)} + n_1^{(c)})} \\ d_c = \sum_{k=0}^1 q_k^{(c)} \log_2 \left( \frac{q_k^{(c)}}{p_k^{(c)}} \right) \\ p_k, q_k = \frac{n_k + \delta}{n_0 + n_1 + 2\delta} \end{array} \right.$$

Where  $d_c$  is the Kullback-Leibler distance for one context, which has label  $c$  and a discrete distribution,  $k = 0, 1$ . The probability function for the first block is  $p_k^{(c)}$ , for the second one it is  $q_k^{(c)}$ . The statistics in first block will be passed on to the second block, and the second block will use it to do the arithmetic coding. Thus, the distance  $D$  we calculate for predicting the arithmetic coding performance uses the counts ( $n_0^{(c)}$  and  $n_1^{(c)}$ ) of every context in the second block as a weight for calculating the average Kullback-Leibler distance to get more accurate prediction. From the definition of the distance, we can see it is always nonnegative, and equals zero if all of the probabilities distribution of the contexts are the same.

Experiments show that the distance between the neighbouring blocks in the same subband in the third dimension (time) is the smallest, followed by the neighbour block in the same spatial subband. It is easy to understand if we explain it from another point of view. The medical image sequences have high correlation between neighbour slices at different positions, but even higher between neighbour frames in time. The decomposition in the time dimension can make it into a multi-resolution subband, but the stillness of the frames makes the wavelet coefficients in the same band generally the same.

In the same 3D subband, neighbouring blocks in the third dimension's (time) direction are coded first, thereafter we move to the next neighbouring block in the 2D spatial domain to finish coding all the blocks in the same way. After finishing the coding of one 3D subband, we move to the next subband along the direction shown in the Figure 4. After finishing coding the first packet, we move to the second one and use the same scan order as in the first one. After we finish scanning all the points for one instance, then the next round for the next bit-plane starts. Scanning by this way can also support progressive coding.

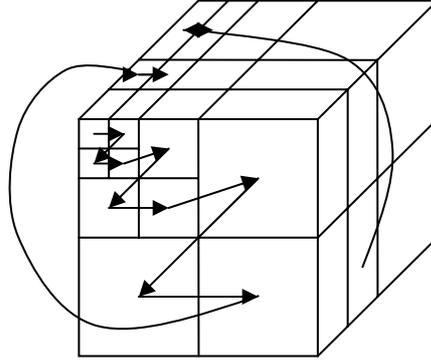


Figure 4, The order in block coding

### 3.4. Initialization of the contexts

We optimize the method of initializing counts of the statistics table by comparing the code lengths in a different way. In order to make it adaptive and localized for small regions in the adaptive arithmetic coding [9], [17], we use the same statistics table in the same subband, which is in the 3D spatial subband and includes at least one slice. The counts of a new subband in the same resolution will be initialized by the scaled old counts from the former subband. It is calculated as follow,

$$new\ count = \left[ \frac{old\ count \times \lambda}{number\ of\ blocks\ in\ last\ subband} \right]$$

$\lambda$  is a factor in controlling the volume of the initialized counts.

In subbands of different resolution or different bit-plane, the counts of new statistics will be force to zeros.

### 3.5. Further extension of the 3-D context

As we have stated before, the average Kullback-Leibler distance between the neighbour slices in the time direction is small. Besides, pixels coded at one bit-plane can also be used as information in coding the same position in the next slice. Status of the pixel at one bit-plane has two types. We can extend the number of zero coding contexts to  $9 \times 2 \times 2 = 36$ . This information can be used to extend magnitude refinement coding contexts and the sign coding contexts. They will have  $3 \times 2 = 6$  and  $5 \times 2 = 10$  contexts separately.

## 4. Object oriented 3D context based coding

The functional fMRI data we used is image slices of a human head. The head object is located at the center of the images and it has a blurred boundary between object and the background. There is a lot of noise in the background. In the medical treatment, the head object is the most important for diagnosis, thus compression may be increased if only the head object information is coded. But we should find a fully safe way to strip the background, but not any part of the head. We call the lossless compression of the head object for Head Object Lossless Compression (HOLC). Another object oriented method is lossless compression of the whole dataset, but using different statistics for different objects. We call it Head Object Oriented Lossless Compression (HOOLC).

### 4.1. HOLC

To demonstrate the concept we use a simple (bounding) rectangle shape to mask the head object. In deciding the size of the rectangle, we use different thresholds to test the acceptable boundary between the head and the background. Because gray level of the bone is much

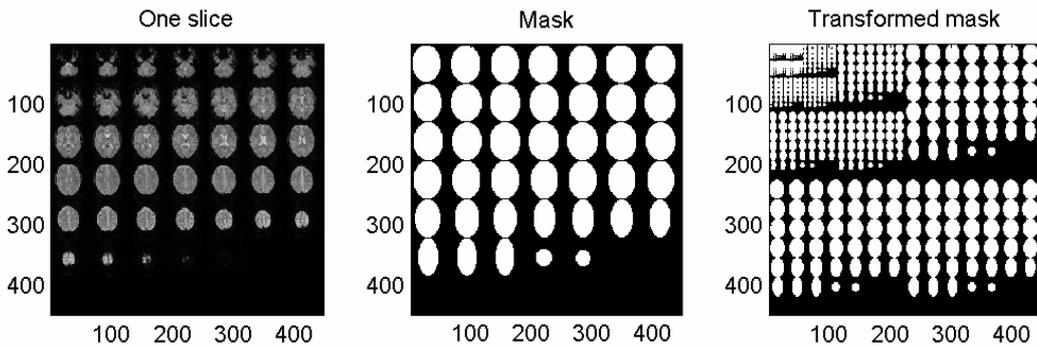
larger than that of the margin, it is easy to find the boundary in slices at the back head position. But it is difficult to find it in slices at the front face position. Thus we test different thresholds until it is safe to force the margin to zero. We code the preprocessed dataset and combine the mask information we have known in the preprocessing. Mask information is processed by transforming the mask from the 2D spatial domain to 2D wavelets domain according to 2D dyadic decomposition level. Because the wavelet decompositions change the resolution to a lower scale, so it is easy to predict the boundary in the wavelets domain without truly using the wavelet transformation. The transformed mask is only used to differentiate the boundary in the wavelets domain and the head object will be coded separately from the almost zero back ground (head edge spread to the background due to long length filters and higher decomposition level).

An ellipse-shaped mask has also been tested in HOLC. It can be very effective for some slices using a high threshold, but it will cut off part of the head object in some slices. In order to make it safe for all the slices, the size would be very large and the area of the object would be larger than that of the simple rectangle shape mask. So, it was not used for the HOLC method.

## 4.2. HOOLC

In HOOLC mode, the ellipse mask is a good choice, because almost all the head slices have shapes similar to an ellipse. It can match the boundary more accurately than a rectangle. Because both the head object and the background object would be compressed, it is tolerable to make mistakes in masking the head. Also the ellipse shape is very easy to be coded with only one point coordinate and two axis' length values and the same position share the same value.

By this method, we separate the statistics of these two objects. Because the textures of these two objects are different, the statistics can be very different. It can increase the compression ratio based on the clear classification. In order to get the highest compression ratio, a good threshold needs to be chosen to separate the different objects or in other words, the different textures. The same coding method is applied as in HOLC except preprocessing. The ellipse mask and the transformed version [see Figure 5] are used for separating the coding content.



**Figure 5:** One slice of the image sequence, the ellipse mask with threshold equal to 50 and mask after four levels decomposition

So HOOLC can fully support lossless compression by compressing the head object and the background separately. HOOLC first codes the head object and then the background, so it can provide extra dimension to progressive coding to the traditional wavelet based coding. Also HOOLC can support lossless reconstruction of the head object and almost lossless reconstruction of the background. It will be flexible to be used for the doctors, such as used in telemedicine [18].

## 5. Experiment results

In the experiments, lossless compression for a dataset of dimensions  $448 \times 448 \times 32$  (16 bits per pixel) was performed. The data is larger than the normal 8 bit data sets, and the maximum value requires 11 bits to be represented. As the data set is mainly used for medical analysis and treatment, we concentrate on the test of lossless compression (2D context, 3D context and HOOLC) and lossy compression HOLC with different reversible wavelets transform. The lossy compression results were compared with 3D\_SPIHT. For HOOLC, an optimal threshold was determined for the (2, 2) reversible wavelet transform.

### 5.1. Lossless compression and HOLC with different filters

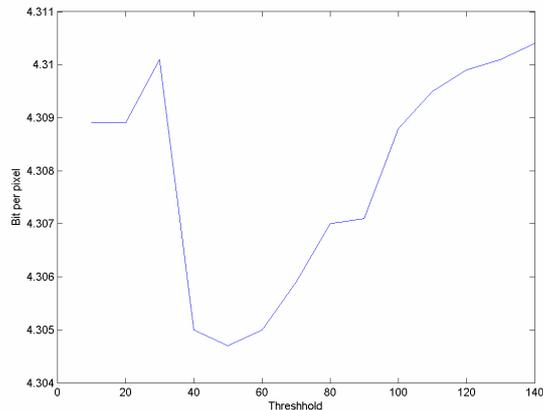
The lossless compression performances of the proposed coders were evaluated with different integer filters. The results [see Table 1] show that extending the 2D context to 3D can increase the lossless compression ratio. The best filter for the head fMRI sequence is 2/6 and then (2, 4), (2, 2). The hybrid filters can also give us good results, and 2D 2/6 combined with (2, 4) has the best performance in the 2D context coding. As mentioned before, HOLC is only lossless for the head object. The result of HOLC shows that many bits have been used to code the noisy background.

**Table 1:** Lossless compression results with different encoders and reversible filters

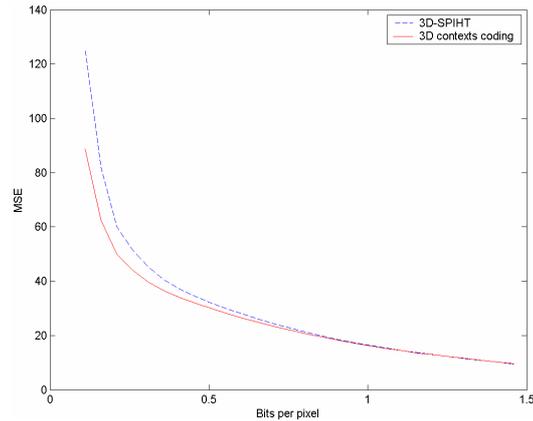
Methods	2D context	3D context	HOOLC (T = 50)	HOLC (T =100)
(2,2)	4.393	4.334	4.305	2.686
(2,4)	4.386	4.329	4.299	2.734
(4,2)	4.445	4.386	4.356	2.837
(2+2,2)	4.439	4.384	4.354	2.816
(4,4)	4.421	4.362	4.332	2.843
(6,2)	4.471	4.413	4.383	2.892
9/7	4.428	4.355	4.324	2.903
9/7-F	4.419	4.346	4.315	2.872
2/6	4.384	<b>4.292</b>	<b>4.262</b>	<b>2.585</b>
2D(2,2) – 3rdD(2,4)	4.387	4.330	4.300	2.685
2D(2,2) – 3rdD 2/6	4.397	4.327	4.298	2.681
2D(2,4) – 3rdD(2,2)	4.391	4.333	4.303	2.735
2D(2,4) – 3rdD 2/6	4.396	4.326	4.297	2.730
2D 2/6 – 3rdD(2,2)	4.378	4.299	4.270	2.590
2D 2/6 – 3rdD(2,4)	<b>4.373</b>	4.294	4.265	2.589

### 5.2. Threshold choosing in lossless 3D object based medical image compression

In HOOLC, different thresholds from 10 to 140 have been tested. The result is shown [see Figure 6]. When the threshold equals to 50, we get the highest compression ratio. And the mask can increase the compression ratio by classifying different textures of the head object and the background.



**Figure 6:** Lossless compression result with different thresholds. Using (2,2,) reversible wavelet transform



**Figure 7:** Rate-distortion comparison between 3D contexts adaptive arithmetic coding and 3D-SPIHT in lossy compression.

But the mask is larger than the clear boundary of the head. One explanation is that near the boundary there are still some areas having a texture similar to that of the head but with a lower gray scale.

### 5.3. Lossy compression

Lossy compression was constructed based on the 3D context coding, and the transformation is a real number to real number wavelet transform, which is also constructed by lifting steps [6] and has scaling to make it unitary. Result was compared with 3D\_SPIHT [see Figure 7]. The performance is better than 3D-SPIHT. In high bit-rate, the performance is also better than 3D-SPIHT, but not as obvious as in low bit-rate. The integer to integer transform based coding can also support progressive reconstruction, but the rate-distortion performance is not better than that of the real number to real number wavelet transform because of rounding errors at every step. If we need to use integer to integer transform based lossy compression, the filters with less lifting steps will be preferable.

## 6. Conclusion

This paper gives an overview of four evolutionary 3D wavelet based coders for 4D medical image sequence, from 2D context to 3D context and then object based coders. The performance is improved and tested after the refinement of the contexts and object classification. Filter 2/6 has the best performance in compressing fMRI data of head, and hybrid filters can also be used. The lossy compression performance is better than 3D-SPIHT.

A good extension from 2D context to 3D has been given. The Kullback-Leibler distance was used to explain the similarity between the different statistics.

HOLC is a coder supporting region-of-interest (ROI). HOOLC is a safe way to compress fMRI image data of a head and it can increase the compression ratio. It also provides an extra dimension for progressive reconstruction, focusing on the object first and then the background. This can increase the compression ratio by mildly throwing away some non-usable bits related to the background.

## References

- [1] *JPEG2000 Image Compression Fundamentals, Standards and Practice*, David S. Taubman, Michael W. Marcellin, Page 356.
- [2] A Family of Normalized LMS Algorithms, Scott C. Douglas, *IEEE Signal Processing*, Vol. 1, NO. 3 March 1994.

- [3] Lossy-to-Lossless Compression of Medical Volumetric Data Using Three-Dimensional Integer Wavelet Transforms, Zixiang Xiong, Xiaolin Wu, *IEEE Transactions on Medical Imaging*, Vol. 22, No. 3, March 2003
- [4] M. Antonini M. Barlaud P. Mathieu I. Daubechies, "Image coding using wavelet transform," *IEEE Trans. on Image Processing*, vol. 1, no. 2, pp. 205--220, 1992.
- [5] M. D. Adams and F. Kossentini, Reversible Integer-to-Integer Wavelet Transforms for Image Compression: Performance Evaluation and Analysis, *IEEE Trans. on Image Processing*, vol. 9, no. 6, Jun. 2000, pp. 1010-1024.
- [6] I. Daubechies and W. Sweldens Factoring Wavelet Transforms into Lifting Steps, *J. Fourier Anal. Appl.*, Vol. 4, Nr. 3, pp. 247-269, 1998.
- [7] Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. New York: Wiley, 1991.
- [8] Qian, H. "Relative Entropy: Free Energy Associated with Equilibrium Fluctuations and Nonequilibrium Deviations." 8 Jul 2000.
- [9] Khalid Sayood, *Introduction to data Compression*, second edition
- [10]: A. Said and W. A. Pearlman. "A New Fast and Efficient Coder Based on Set Partitioning in Hierarchical Trees". *IEEE Transactions on Circuits and Systems for Video Technologies*, pages 243/250, June 1996
- [11] J. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Signal Processing*, vol. 41, pp. 3445--3463, Dec. 1993.
- [12] Ligang Lu<sup>1</sup>, Zhou Wang<sup>2</sup>, and Alan C. Bovik, "Adaptive frame prediction for foveation scalable video coding."
- [13] D. Taubman, "High performance scalable image compression with EBCOT," *IEEE Trans. Image Processing*, vol. 9, pp. 1158--1170, July 2000. *IEEE Trans. Commun.*, vol. 45, pp. 437--444, Apr. 1997.
- [14] R. C. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo, "Wavelet transforms that map integers to integers," *J. Appl. Comput. Harmon. Anal.*, vol. 5, pp. 332--369, 1998.
- [15] A. Bilgin, G. Zweig, and M. W. Marcellin, "Three-dimensional image compression using integer wavelet transforms," in *Appl. Opt.: Inform. Proc.*, vol. 39, Apr. 2000, pp. 1799--1814.
- [16] W. Sweldens, "The lifting scheme: A construction of second generation wavelet," *SIAM J. Math. Anal.*, vol. 29, pp. 511--546, 1997.
- [17] G. Langdon, "An introduction to arithmetic coding," *IBM J. Res. Dev.*, vol. 28, pp. 135--149, 1984.
- [18] "Telemedicine Today Magazine," Westlake Village, CA.
- [19] B.-J. Kim, Z. Xiong, and W. A. Pearlman, "Low bit-rate scalable video coding with 3-D set partitioning in hierarchical trees (3-D SPIHT)," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 10, pp. 1365--1374, Dec. 2000.

# Integrity Improvements in Classically Deformable Solids

Micky K. Christensen\*

Anders Fleront†

Department of Computer Science  
University of Copenhagen  
DIKU

## Abstract

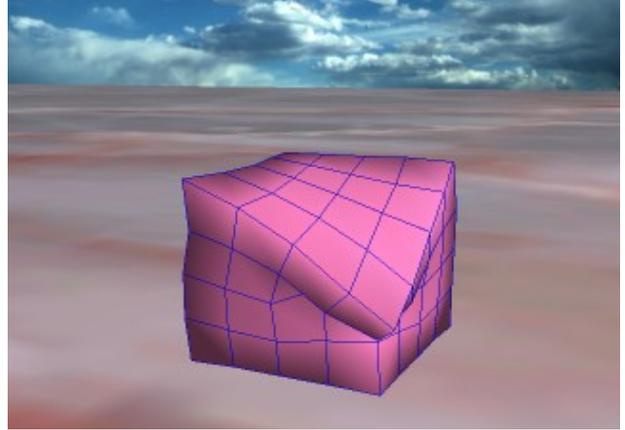
The physically-based model for simulating elastically deformable objects, presented by Terzopoulos et al. in 1987, has significant problems with handling solids. We give a plausible explanation for the reasons of instability that lead to implosion of deformable solids. We propose an extension to the original model, with improvements that result in a model of increased stability. Comparisons are made with the original method to illustrate this point. The improved model is suitable for interactive simulations of deformable solids, with only a small constant decrease of performance.

**Keywords:** Physics-based animation, simulation, dynamics, deformation, real time, deformable solids, continuum elastic model, constraints, variational calculus, finite differences

## 1 Introduction

Deformable models seem to have gained increasing interest during the last years. Part of this success comes from a desire to interact with objects that resembles those in real life, which all seem to be deformable at some level. The fact that CPUs and GPUs today are both advanced and powerful makes it possible to simulate and animate deformable bodies interactively. This paper builds on work previously done by Christensen and Fleron [1], which focused on an implementation of a method for simulating elastically deformable objects, first put forward by Terzopoulos et al. in 1987 [9]. The implementation could simulate various types of deformable surfaces convincingly, but severable instability issues were noticed, with regard to simulating deformable solids. As the authors of the original method claim the model can simulate deformable curves, surfaces, and solids, we will look into why the method was unable to simulate deformable solids satisfactorily. At first sight it seemed that the solids did have a hard time keeping their integrity. Integrity preservation is important to give a realistic impression of a deformable solid. We will try to solve the instabilities using concepts from the original frame work, and propose an improved model that is able to simulate solids, as well as surfaces, without a significant decrease of the overall performance.

In section 2 we revisit the theory of elastically deformable models, with focus on solids, to give an overview of the method. The theory serves as a foundation for understanding the following sections. Section 3 reveals and explains the instabilities of the original model. Sections 4 and 5 focus on solutions and improvements of the local and global problems, respectively. In section 6 we present the



**Figure 1** Large deformation results in convincing material buckling

results of adding the developed improvements to the model. Comparison between the original and improved model will be performed visually.

## 2 Elastically Deformable Solids

At the basis of the theory of deformable models lies elasticity theory. From this physical theory Terzopoulos et al. have extrapolated a model that governs the movements of elastically deformable bodies [9]. In this paper we will concentrate on the theory of deformable solids.

A point in a solid is described by the intrinsic coordinates  $\mathbf{a} = [a_1, a_2, a_3]$ . A deformable solid is thought of as having a natural rest state, where no elastic energy is inherent. The rest state is described by  $\mathbf{r}^0(\mathbf{a}) = [r_1^0(\mathbf{a}), r_2^0(\mathbf{a}), r_3^0(\mathbf{a})]$ , where the positional vector function  $\mathbf{r}$  of the object is defined in Euclidian 3-space. When the solid is deformed, it takes on a different shape than its rest shape, and distances between nearby points are either stretched or compressed with the deformation. This ultimately creates elastic energy that results in internal forces that will seek to minimize the elasticity.

The deformation will evolve over time and as such, it can be described by the time-varying vector function  $\mathbf{r}(\mathbf{a}, t) = [r_1(\mathbf{a}, t), r_2(\mathbf{a}, t), r_3(\mathbf{a}, t)]$ . The evolving deformation is independent of the rigid body motion of the solid. The equations governing the motion of particles in a deformable solid are obtained from Newtonian mechanics, and given by

$$\frac{\partial}{\partial t} \left( \mu \frac{\partial \mathbf{r}}{\partial t} \right) + \gamma \frac{\partial \mathbf{r}}{\partial t} + \frac{\delta \mathcal{E}(\mathbf{r})}{\delta \mathbf{r}} = \mathbf{f}(\mathbf{r}, t), \quad (1)$$

\* micky@mkc.dk

† afleron@knights.dk

where  $\mathbf{r}(\mathbf{a}, t)$  is the position of the particle  $\mathbf{a}$  at time  $t$ ,  $\mu(\mathbf{a})$  is the mass density at particle  $\mathbf{a}$ , and  $\gamma(\mathbf{a})$  is the damping density. The right hand side represents the sum of externally applied forces at time  $t$ . The third term on the left hand side of (1) is called a variational derivative, and this is where the internal elastic energy is dealt with.  $\varepsilon(\mathbf{r})$  is called a functional and it measures the potential energy that builds up when the body is deformed.

## 2.1 Deformation Energy

A method is needed to measure the deformation energies that arise when a solid deforms. For this task, we use the area of differential geometry. It is convenient to look at arc-lengths on curves, running along the intrinsic directions of the solid. A way of measuring the directions is specified by what is known as the first fundamental form. This form takes on a local view, and looks at changes in length between particles in a small neighborhood of a particle. When looking at the first fundamental form, or the metric tensor, we get an idea of the distances between particles. For a deformable solid the metric tensor is

$$G_{ij}(\mathbf{r}(\mathbf{a})) = \frac{\partial \mathbf{r}}{\partial a_i} \cdot \frac{\partial \mathbf{r}}{\partial a_j}, \quad 1 \leq i, j \leq 3, \quad (2)$$

which is a symmetric  $3 \times 3$  tensor. The diagonal of the tensor represents length measurements along the coordinate directions from the particle in question. The off-diagonal elements represent angle measurements between the coordinate directions, which resist shearing within the local particle structure.

When measuring deformation energy in a solid, we are interested in looking at the change of the shape, with respect to the natural rest shape. This rest shape is described by the superscript 0 such that

$$G_{ij}^0(\mathbf{r}(\mathbf{a})) = \frac{\partial \mathbf{r}^0}{\partial a_i} \cdot \frac{\partial \mathbf{r}^0}{\partial a_j}. \quad (3)$$

By using the weighted Hilbert-Schmidt matrix norm of the difference between the metric tensors, a simplified way of describing the energy of deformation becomes

$$\varepsilon(\mathbf{r}) = \int_{\Omega} \sum_{i,j=1}^3 \eta_{ij} (G_{ij} - G_{ij}^0)^2 da_1 da_2 da_3, \quad (4)$$

where  $\Omega$  is the domain of the deformable solid and  $\eta$  is a user defined tensor that weights each of the coefficients of the metric. For a single point the energy function is simply

$$S = \sum_{i,j=1}^3 \eta_{ij} (G_{ij} - G_{ij}^0)^2. \quad (5)$$

The elastic energies that occur when a solid is deformed come from the stretching and compressing of the particles in the body. To discover the natural behavior of a deformed solid seeking towards its rest state, a term that minimizes the deformation energy is desirable. Variational calculus can be applied to find such a minimizing term, which is described by the Euler-Lagrange differential equations [5]. For a

problem defined by a function of three independent variables and their first derivatives, the equations become

$$\frac{\partial F}{\partial \bar{y}} - \frac{d}{dx_1} \frac{\partial F}{\partial \bar{y}_{x_1}} - \frac{d}{dx_2} \frac{\partial F}{\partial \bar{y}_{x_2}} - \frac{d}{dx_3} \frac{\partial F}{\partial \bar{y}_{x_3}} = 0. \quad (6)$$

The symbol  $F$  is the function describing the problem, and  $F$  will be minimized with respect to the symbol  $\bar{y}$ , which in our case is the positional vector field  $\mathbf{r}$ . Substituting (5) into  $F$  in (6) and executing the differentiations yields a neatly compact expression that minimizes the given problem

$$\frac{\delta S}{\delta \mathbf{r}} = - \sum_{i,j=1}^3 \partial_{a_i} \left( \alpha_{ij} \mathbf{r}_{a_j} \right), \quad (7)$$

with

$$\alpha_{ij} = \eta_{ij} \left( \mathbf{r}_{a_i} \cdot \mathbf{r}_{a_j} - G_{ij}^0 \right). \quad (8)$$

The alpha tensors (8) represent the comparison between the deformed state and the rest state of the solid. When an element in an alpha tensor becomes positive, it means that the corresponding constraint has been stretched and that it wants to shrink. Likewise, when an element becomes negative, the constraint has been compressed and it wants to grow.

## 2.2 Discretization

The deformable model is continuous in the intrinsic coordinates. To allow an implementation of deformable solids, the model is discretized into a unit 3D grid structure, representing the particles which will make up a solid. The grid will have three principal directions called  $l$ ,  $m$ , and  $n$ . Particles in the grid are uniformly distributed with spacings in each of the three directions, given by the symbols  $h_1$ ,  $h_2$ , and  $h_3$ . The number of particles in each of the directions are designated  $L$ ,  $M$ , and  $N$ . Each particle property will be discretized using an index  $[l, m, n]$ , which will return the property value of that particular grid entry, e.g. the particle positions will be described by the discrete vector function  $\mathbf{r}[l, m, n]$ .

The model requires that derivatives are calculated in the intrinsic directions of the object. For this purpose we use finite difference operations across the grid, to achieve the desired derivative approximations [2]. An approximation to the first derivative in the  $m$ -direction, can for example be obtained by use of either a forward or backward difference operator. Given an arbitrary grid function  $\mathbf{u}[l, m, n]$  these approximations are

$$D_2^+(\mathbf{u}) = h_2^{-1} (\mathbf{u}[l, m+1, n] - \mathbf{u}[l, m, n]) \quad (9)$$

and

$$D_2^-(\mathbf{u}) = h_2^{-1} (\mathbf{u}[l, m, n] - \mathbf{u}[l, m-1, n]), \quad (10)$$

where the superscript symbols,  $+$  and  $-$ , represent forward and backward operators, respectively.

Using the difference operators, it is possible to discretize (7), by replacing the derivatives with the corresponding

difference operators. The discrete equation for the elastic force  $\mathbf{e}$  becomes

$$\mathbf{e}[l, m, n] = \sum_{i,j=1}^3 -D_i^-(\mathbf{p})[l, m, n], \quad (11)$$

$$\mathbf{p}[l, m, n] = \alpha_{ij}[l, m, n] D_j^+(\mathbf{r})[l, m, n], \quad (12)$$

and the  $\alpha$  tensor field is also discretized using finite differencing,

$$\alpha_{ij}[l, m, n] = \eta_{ij}[l, m, n] (D_i^+(\mathbf{r})[l, m, n] \cdot D_j^+(\mathbf{r})[l, m, n] - G_{ij}^0[l, m, n]). \quad (13)$$

Problems at the boundaries of the grid will occur because no information is readily available about the derivatives in these places. A natural boundary condition can be created by setting to zero all forward difference operators that reach outside the grid [9].

To solve equations for all particles at the same time, the values in the positional grid  $\mathbf{r}$  and in the energy grid  $\mathbf{e}$  can be unwrapped into  $LMN$ -dimensional vectors  $\underline{\mathbf{r}}$  and  $\underline{\mathbf{e}}$ . With these vectors, the entire system of equations can be written as

$$\underline{\mathbf{e}} = \mathbf{K}(\underline{\mathbf{r}}) \underline{\mathbf{r}}, \quad (14)$$

where  $\mathbf{K}(\underline{\mathbf{r}})$  is an  $LMN \times LMN$  sized matrix, called the stiffness matrix, which has desirable computational properties such as sparseness and bandedness. A discussion on how to assemble the stiffness matrix  $\mathbf{K}$  can be found in [4].

We introduce the diagonal  $LMN \times LMN$  mass matrix  $\mathbf{M}$  and damping matrix  $\mathbf{C}$ , assembled from the corresponding discrete values of  $\mu[l, m, n]$  and  $\gamma[l, m, n]$ , respectively. The equations of the elastically deformable model (1) can now be expressed in grid vector form, by the coupled system of second order ordinary differential equations

$$\mathbf{M} \frac{\partial^2 \underline{\mathbf{r}}}{\partial t^2} + \mathbf{C} \frac{\partial \underline{\mathbf{r}}}{\partial t} + \mathbf{K}(\underline{\mathbf{r}}) \underline{\mathbf{r}} = \underline{\mathbf{f}}. \quad (15)$$

### 2.3 Numerical Integration

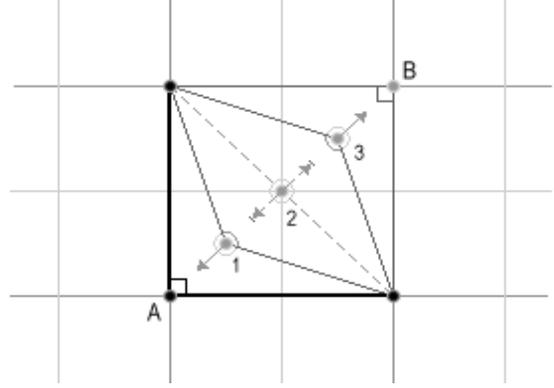
To evolve the deformable model over time, a semi-implicit time integration scheme will be employed. The time interval that the model will evolve in, is subdivided into time steps of equal size  $\Delta t$ . Using central differences, the time derivatives are approximated by

$$\begin{aligned} \frac{\partial^2 \underline{\mathbf{r}}}{\partial t^2} &\approx \frac{\underline{\mathbf{r}}_{t+\Delta t} - 2\underline{\mathbf{r}}_t + \underline{\mathbf{r}}_{t-\Delta t}}{\Delta t^2} \\ \frac{\partial \underline{\mathbf{r}}}{\partial t} &\approx \frac{\underline{\mathbf{r}}_{t+\Delta t} - \underline{\mathbf{r}}_{t-\Delta t}}{2\Delta t}. \end{aligned} \quad (16)$$

By inserting these derivatives into (15) we convert the nonlinear system into the system of linear equations

$$\mathbf{A}_t \underline{\mathbf{r}}_{t+\Delta t} = \underline{\mathbf{g}}_t \Leftrightarrow \underline{\mathbf{r}}_{t+\Delta t} = \mathbf{A}_t^{-1} \underline{\mathbf{g}}_t, \quad (17)$$

where  $\mathbf{A}$  is called the system matrix,



**Figure 2** The surface patch will collapse to a curve, when a particle crosses the opposite diagonal

$$\mathbf{A}_t = \mathbf{K}(\underline{\mathbf{r}}_t) + \left( \frac{1}{\Delta t^2} \mathbf{M} + \frac{1}{2\Delta t} \mathbf{C} \right), \quad (18)$$

and  $\underline{\mathbf{g}}$  is called the effective force vector,

$$\underline{\mathbf{g}}_t = \underline{\mathbf{f}}_t + \left( \frac{1}{\Delta t^2} \mathbf{M} + \frac{1}{2\Delta t} \mathbf{C} \right) \underline{\mathbf{r}}_t + \left( \frac{1}{\Delta t} \mathbf{M} - \frac{1}{2} \mathbf{C} \right) \underline{\mathbf{v}}_t, \quad (19)$$

with the velocity estimated by the normal first order backwards difference

$$\underline{\mathbf{v}}_t = \Delta t^{-1} (\underline{\mathbf{r}}_t - \underline{\mathbf{r}}_{t-\Delta t}). \quad (20)$$

With these equations in hand it is possible to implement real time dynamic simulations of deformable solids. The desirable properties of the system matrix indicate that a specialized solution method, such as the conjugate gradient method [8], can be utilized.

### 3 Instabilities

The instabilities we will address are not numerical in nature, such as the chosen integration method, timesteps, etc. We will turn the interest towards instabilities, regarding the structure of the underlying constraints in the model of elastic solids. The problems can, more or less, be identified as boundary issues, as the integrity instabilities arise from the boundaries of the discrete 3D grid. Real life deformable objects are held together by strong forces at a very small scale. Similarly, we could solve the integrity problems, if we could get away with using extremely high values in the underlying weighting tensors. The reason why we must disregard this option is because it will require us to integrate numerically using “infinitely” small timesteps. As we are interested in interactive simulations, we seek to use large timesteps, thus to simply increase the values of the tensors is not a desired option.

Differential Geometry [7] is used as a tool to measure deformation of an elastic body, in comparison with its resting shape. For solids, the first fundamental forms, or  $3 \times 3$  metric tensors, are sufficient to distinguish between the shapes of two bodies. However, the metric tensor of a solid (2) is not nearly sufficient enough as a tool to compute the complex particle movements of a deformed solid, seeking

towards its resting shape. The discrete components of (2), disregarding the diagonal, describe the cosine to the angle between adjacent directions multiplied by the product of the lengths,

$$\mathbf{v} \cdot \mathbf{w} = |\mathbf{v}| |\mathbf{w}| \cos \theta, \quad 0 \leq \theta \leq \pi. \quad (21)$$

The angle between two vectors is not dependent on their mutual orientation, as verified by the domain of  $\theta$  in (21). This is the main reason why we will get internal structure instabilities. The case of surface patches is depicted in Figure 2. The bold lines and the angle between them form the natural condition. If particle  $A$  is moved towards  $B$ , as in case 1, the angular elasticity constraints will force the particle back to its resting condition, in the direction indicated by the arrow. The behavior of the angular constraint force is defined ambiguously when the angle between vectors is either 0 or 180 degrees. The latter is shown in case 2. If the particle reaches beyond the opposite diagonal, as in case 3, the elasticity will be strengthened and push particle  $A$  into  $B$ . This is clearly a problem, as it will reduce the surface into a curve. The original model [9] suffers from this instability.

The instabilities for solids get even worse. Expanding the square to a cube and focusing on a particle in the grid. Not only does the particle spawn surface instabilities for the 3 directional patches, there is also nothing that prevents the cube from collapsing over the space diagonals. The disability of volume preservation will ruin the integrity of the solid.

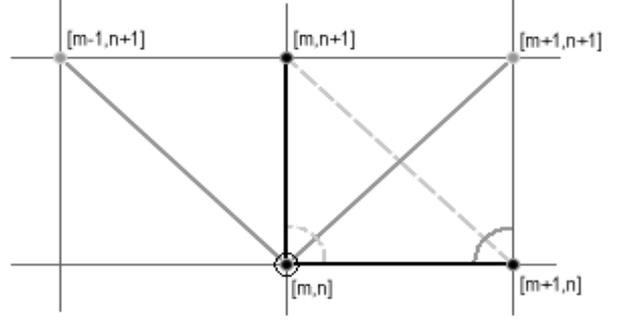
In short, the original elastic model is insufficient to simulate deformable solids. Based on the modeling concept by Terzopoulos et al [9], we will remodel the method, to render it more suitable for simulating elastically deformable solids.

#### 4 Integrity Improvements

To handle integrity instabilities, we are going to design an extension to the original model that can improve its ability to prevent collapsing of the grid cubes. Basically, the extension will be done by adding new constraints to the model, gathered into a new tensor we call the Spatial Diagonal Metric, or SDM. To make things a little easier, we start out by describing the analogy for surfaces.

##### 4.1 Extended Metric

For a surface, the original tension constraints on a given particle is the four constraints given by its  $2 \times 2$  metric tensor  $\mathbf{G}$ . As with solids, when  $i = j$  then  $G_{ij}$  defines the length constraints along the coordinate directions  $m$  and  $n$ , while when  $i \neq j$  then  $G_{ij}$  represents an expression of the same explicit angular constraint between the two directions. The idea is to replace the dual explicit angular constraint with two diagonal length constraints. These constraints will reach from the particle at  $[m, n]$  to the diagonally opposite particles at  $[m+1, n+1]$  and  $[m-1, n+1]$ , as depicted in Figure 3. Besides working as diagonal length constraints, they will also implicitly work as angular constraints that together can account for all 360 degrees. The directions along these new constraints will be



**Figure 3** The dual explicit angular constraint is replaced by two new diagonal constraints that will define the angular constraints implicitly

considered as new directions called  $a_{d1}$  and  $a_{d2}$ . When the model is extended with the new diagonal constraints, variational calculus must be utilized for the expression of the minimizing elastic force. Writing out (5) for the case of surfaces, with focus on the new directions, yields

$$S = \dots + \eta_{12} \left( \mathbf{r}_{a_{d1}} \cdot \mathbf{r}_{a_{d1}} - \mathbf{G}_{12}^0 \right)^2 + \eta_{21} \left( \mathbf{r}_{a_{d2}} \cdot \mathbf{r}_{a_{d2}} - \mathbf{G}_{21}^0 \right)^2 + \dots, \quad (22)$$

where the elements  $G_{12}^0$  and  $G_{21}^0$  of the rest state tensor, now holds the natural state of the new diagonal constraints. With the new directions,  $a_{d1}$  and  $a_{d2}$ , in hand, the Euler-Lagrange equation changes to

$$\frac{\delta S}{\delta \mathbf{r}} = S_{\mathbf{r}} - \partial_{a_1} S_{\mathbf{r}_{a_1}} - \partial_{a_2} S_{\mathbf{r}_{a_2}} - \partial_{a_{d1}} S_{\mathbf{r}_{a_{d1}}} - \partial_{a_{d2}} S_{\mathbf{r}_{a_{d2}}}. \quad (23)$$

Using variational calculus on the terms from (22) results in the following additions to the elastic force  $\mathbf{e}$ ,

$$\mathbf{e}[m, n] = \dots - D_{d1}^-(\mathbf{p}^*)[m, n] - D_{d2}^-(\mathbf{q}^*)[m, n], \quad (24)$$

where

$$\mathbf{p}^*[m, n] = \eta_{12}[m, n] D_{d1}^+(\mathbf{r})[m, n]$$

and

$$\mathbf{q}^*[m, n] = \eta_{21}[m, n] D_{d2}^+(\mathbf{r})[m, n]. \quad (25)$$

Notice that new difference operators arise with the new directions in the discretization of the extended metric. These operators can be considered to be just like the operators in the original directions. For example, the new first order forward difference operators become

$$D_{d1}^+(\mathbf{u}) = \frac{\mathbf{u}[m+1, n+1] - \mathbf{u}[m, n]}{h_{d1}}, \quad (26)$$

$$D_{d2}^+(\mathbf{u}) = \frac{\mathbf{u}[m-1, n+1] - \mathbf{u}[m, n]}{h_{d2}},$$

where  $h_{d1} = h_{d2} = \sqrt{h_1^2 + h_2^2}$  is the grid distance in both diagonal directions. The square root is computationally expensive, but it can easily be avoided in any calculation

involving the new difference operators, or simply be pre-computed as the value does not change at runtime.

We have shown how to extend the metric tensor, and we call the extended metric for *Matrix*. To give a deformable solid a stronger integrity, the *Matrix* can easily be applied to work for solids. The  $3 \times 3$  metric tensor for a solid contains three length constraints along its diagonal. The remaining six elements represent angular constraints between the original coordinate directions, and these can be replaced by the appropriate elements of the *Matrix*.

#### 4.2 Spatial Diagonal Metric

To implement volume preservation, we introduce another new idea called the Spatial Diagonal Metric, which is a  $2 \times 2$  tensor. The four elements will represent four new length constraints that will be spatially diagonal, meaning they will span grid cubes volumetrically. The four new directions can be chosen three-ways, and we have simply chosen them to reach from  $[l, m, n]$  to the four particles at  $[l-1, m+1, n+1]$ ,  $[l-1, m-1, n+1]$ ,  $[l+1, m+1, n+1]$ , and  $[l+1, m-1, n+1]$ . No matter how the four length constraints are chosen to span the cubes, their contributions will end up covering a grid cube symmetrically, as depicted in Figure 4.

The new SDM tensor that represents the volume will be called  $\mathbf{V}$ . It follows the same design ideas as the *Matrix*, and is defined as

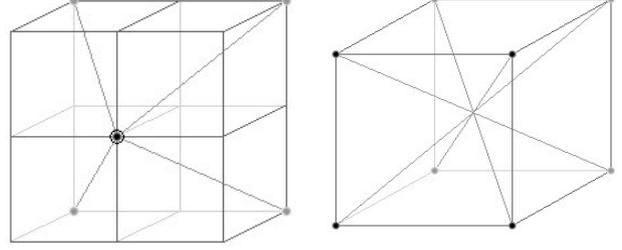
$$\mathbf{V} = \begin{bmatrix} D_{v1}^+ \cdot D_{v1}^+ & D_{v2}^+ \cdot D_{v2}^+ \\ D_{v3}^+ \cdot D_{v3}^+ & D_{v4}^+ \cdot D_{v4}^+ \end{bmatrix}, \quad (27)$$

where  $D_{v1..4}^+$  are the four new first order forward difference operators along the new spatial directions.

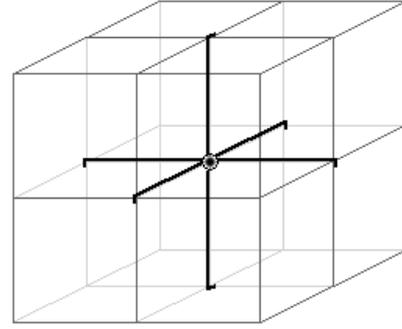
### 5 Global Implosions

With the SDM and *Matrix* contributions added to the elasticity term, the model of deformable solid can prevent the shape of the grid cubes from undergoing local collapsing. This is an important improvement towards keeping the integrity of a deformable solid intact. Another integrity issue still exists that renders a solid unable to prevent implosions. In this matter, we define an implosion as when a grid cube enters one of its adjacent grid cubes. Implosions can happen upon large deformation, which typically are caused by heavy external forces, e.g. reaction to collisions and aggressive user interactions. Global implosions can also be defined as internal self-intersections, thus self-intersection tests can be utilized as a tool to prevent implosions.

Common methods to avoid self-intersection include surrounding each grid cube into an axis aligned bounding box, or AABB, and arrange the AABBs into a bounding volume hierarchy, or BVH tree, that can be updated efficiently as the body deforms [4]. A recent paper introduces image-space techniques that can be implemented on the GPU, to allow performance friendly detection of self-intersections and collision between deformable bodies [6]. Although a method for handling self-intersection can be chosen to perform reasonable, it will still decrease the overall performance.



**Figure 4** Spatial length constraints for solids. On the left it is shown how the four constraints reach from the center. On the right it is shown how the constraint contribution from four particles on one cube side renders symmetric behavior



**Figure 5** Discrete central differences bind adjacent grid cubes together

What we seek is a mechanism that somehow binds adjacent grid cubes together, in such a way that if implosions occur, we can disperse self-intersecting cubes. This is not a method that can prevent self-intersections, but it can restore the integrity of the solid upon implosions. We can reuse what we have been working with so far, and thus reduce the computational cost and memory use significantly, compared to the extra load we would introduce into the system, if we had used a BVH algorithm.

We introduce a new Pillar tensor  $\mathbf{P}$ , which is based upon the discrete metric tensor  $\mathbf{G}$ , but extended to use first order central difference operators. For reasons of clarity we will limit  $\mathbf{P}$  only to use the length constraints

$$\mathbf{P}[l, m, n] = \begin{bmatrix} D_1^2(\mathbf{r}) & 0 & 0 \\ 0 & D_2^2(\mathbf{r}) & 0 \\ 0 & 0 & D_3^2(\mathbf{r}) \end{bmatrix}, \quad (28)$$

where

$$\begin{aligned} D_1(u)[l, m, n] &= \frac{1}{2}h_1^{-1}(u(l+1, m, n) - u(l-1, m, n)), \\ D_2(u)[l, m, n] &= \frac{1}{2}h_2^{-1}(u(l, m+1, n) - u(l, m-1, n)), \\ D_3(u)[l, m, n] &= \frac{1}{2}h_3^{-1}(u(l, m, n+1) - u(l, m, n-1)). \end{aligned} \quad (29)$$

The effect of using central difference operators results in a very convincing way to bind adjacent grid cubes together, see

Figure 5. As every grid particle will be extended with the Pillar tensor, the combined range of  $\mathbf{P}$  will overlap all grid cubes.

To further strengthen the prevention of implosion, the Pillar tensor can easily be extended to use a central difference Metrix tensor.

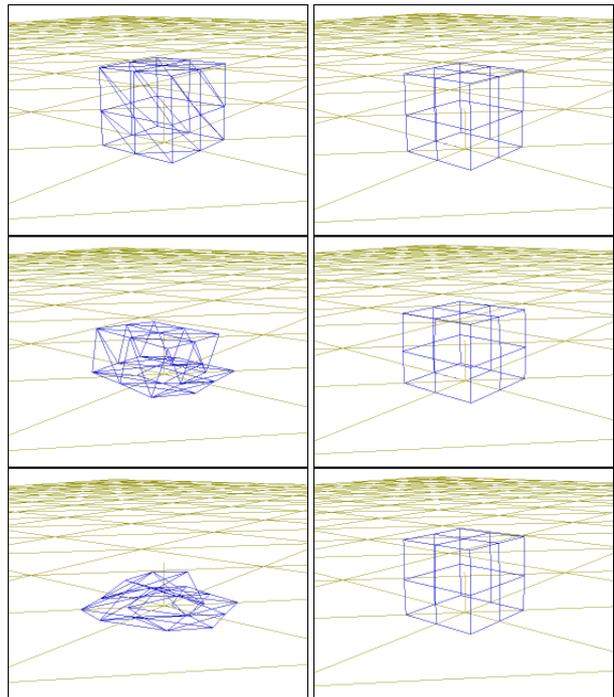
## 6 Results

We have extended the previous implementation of the elastically deformable models [1] to support both the Metrix, SDM, and Pillar contributions, when simulating deformable solids. The implementation can be found in [3]. The deformable solids respond naturally to external forces, e.g. gravity and viscosity. The user can interact with the solids, such as constraining particles to a fixed position and performing pulling operations using spring forces. We have also implemented a scaling mechanism, which allows a user to control the overall uniform strength of the tensors. Adjusting the strength scaling interactively can vary the stiffness of a deformable solid in real time. For example simulating the effect of a solid being inflated.

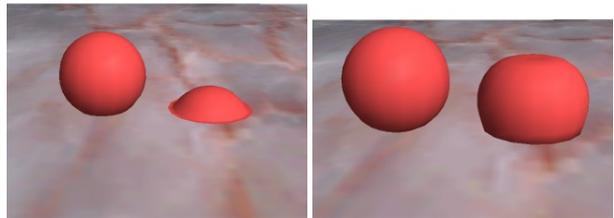
Experiments have revealed that the effects of the Metrix and SDM do not always succeed satisfactorily, in moving particles back to their natural location. In some situations new energy equilibriums arise unnaturally. With the help of the supported visual debugger we have realized that different constraints can work against each other, and the result is that the sum of constraint contributions is zero. To counteract this problem, we have squared the force of the constraints in the SDM (and the Metrix), to make sure they prevail.

To give a reasonable review of how the improved model solves the presented integrity instabilities, we will perform visual comparisons between the original model and the new improved one. In Figure 6, still frames from a small box that is influenced by gravity and collides with a plane, are compared frame to frame between the models. Due to the lack of volume preservation, the constraints of the original model simply cannot keep the shape of the cube. In Figure 7, we have recreated a configuration from [1], comparing two rubber balls with different particle mass. The picture on the left is taken from [1], where the metrics fail to maintain the integrity of the right ball, thus the ball collapses on itself. The picture on the right is simulated using the same parameters but with the improved model, and the integrity of the ball is now strong enough to stay solid. In Figure 8, a test of how well the models can recover from a sudden large deformation is performed.

The improved model enables simulations of situations that are impossible with the original model. In Figure 1, a soft solid is depicted. The solid has been constrained to the ground, and in three of the top corners. Pulling the last corner downwards results in a large deformation and renders convincing material buckling. In Figure 9, the true strength of the Pillar tensor is illustrated, showing an effect of inflation. In Figure 10, some pudding is constrained to the ground, and being twisted by its top face. The sides of the deformable cube skew as is expected of a soft body like pudding. In Figure 11, a large water lily is deformed upon resting on two pearls. The improved model performs a great job in keeping the water lily fluffy.



**Figure 6** A small box is influenced by gravity and collides with a plane. The three stills on the left illustrate the original model, and on the right the frames from the new improved model is shown



**Figure 7** Rubber balls. The left frame illustrates the situation from the old model where the right ball is unable to maintain its integrity. In the right frame the same situation is depicted, but simulated using the improved model



**Figure 8** A wooden box is heavily lifted in one corner. Original vs. improved model, on the left and right frame, respectively

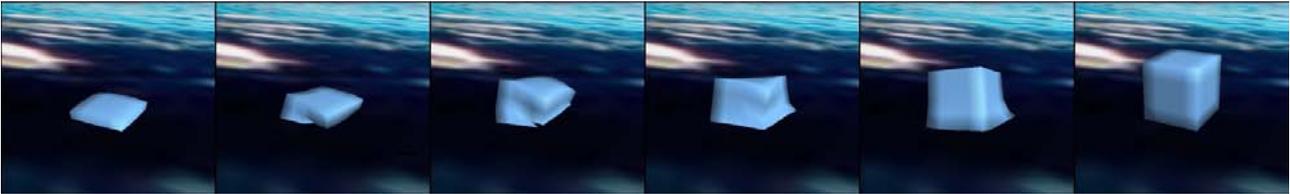


Figure 9 Constraint strength is increased interactively and yields the effect of inflation

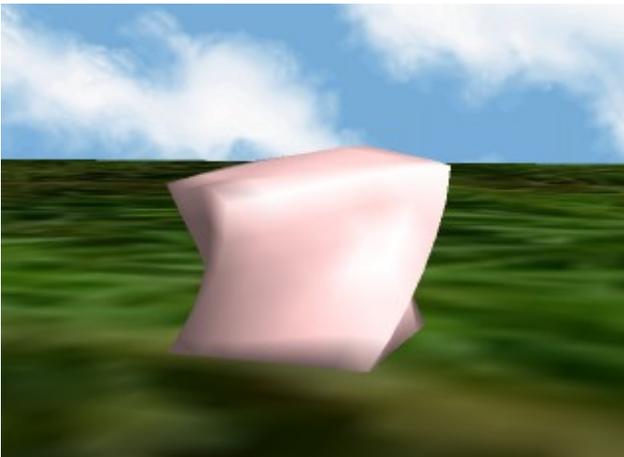


Figure 10 Twisting the pudding renders skewing

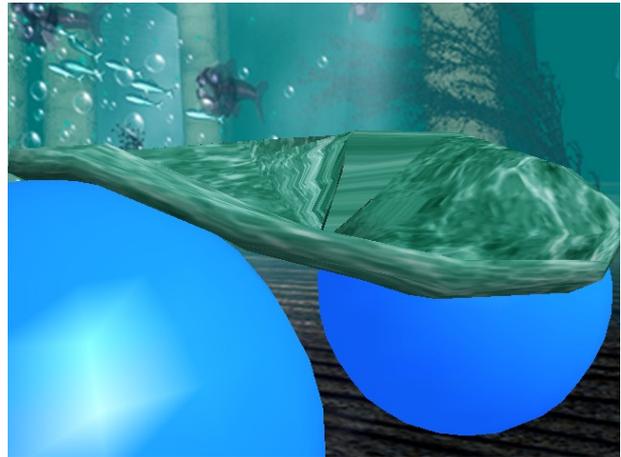


Figure 11 Fluffy water lily modeled using an ellipsoid solid

## 7 Conclusion

The original model for deformable solids, presented in [9], turned out to be insufficient for the sake of realism. Even extremely modest external forces applied to the bodies would ruin their integrity, as can be verified on the left side of Figures 6 to 8. Clearly some extensions to the model were needed. In this paper we have presented three improvements to the model that will give deformable solids a much better ability to keep their original integrity, and thus the ability to handle much larger deformations without collapsing. With the new Metrix, SDM, and Pillar additions to the model, the overall strength of the internal elastic constraints is increased. This reinforces the impression that the elastic bodies are actually solids.

We have shown that the improvements to the original model greatly increase the usability of the method for simulating deformable solids. As the new contributions influence the same particles as the original model, the system matrix  $\mathbf{A}$  is still of size  $LMN \times LMN$ , and it still possess its original pleasant properties that allows an relaxation based solver to invert it, using only a few iterations. The Metrix replaces the metric tensor, thus it is only the SDM and Pillar additions that are new to the calculations, and they can be computed in constant time for each particle.

Newer papers on deformable solids state the importance of preservations of length, surface area, and volume. What we have done with the elastically deformable model is precisely to strengthen the surface preservation using the Metrix, and to implement the missing volume preservation using the SDM.

## References

- [1] M. K. Christensen and A. Fleron (2004), "Implementation of Deformable Objects," Department of Computer Science, University of Copenhagen, DIKU
- [2] D. Eberly (2003), "Derivative Approximation by Finite Differences," Magic Software, Inc., January 21, 2003
- [3] K. Erleben, H. Dohlmann, J. Sporning, and K. Henriksen (2003) "The OpenTissue project," Department of Computer Science, University of Copenhagen, DIKU, November 2003, <http://www.diku.dk/forskning/image/research/opentissue/>
- [4] K. Erleben, J. Sporning, K. Henriksen, and H. Dohlmann (2004), "Physics-based Animation and Simulation," DIKU
- [5] H. Goldstein, C. P. Poole, and J. L. Safko (2002), "Classical Mechanics," Third Edition, Addison-Wesley
- [6] B. Heidelberger, M. Teschner, and M. Gross (2004), "Detection of Collisions and Self-collisions Using Image-space Techniques," Proc. WSCG'04, University of West Bohemia, Czech Republic, pp. 145-152
- [7] J. J. Koenderink (1990), "Solid Shape," MIT Press
- [8] J. R. Shewchuk (1994), "An Introduction to the Conjugate Gradient Method Without the Agonizing Pain," Carnegie Mellon University
- [9] D. Terzopoulos, J. C. Platt, A. H. Barr, and K. Fleischer (1987), "Elastically deformable models," Computer Graphics, volume 21, Number 4, July 1987, pp 205-214

# The Thin Shell Tetrahedral Mesh

Kenny Erleben\* and Henrik Dohlmann†

Department of Computer Science, University of Copenhagen, Denmark

## Abstract

Tetrahedral meshes are often used for simulation of deformable objects. Unlike engineering disciplines graphics is biased towards stable, robust and fast methods, instead of accuracy. In that spirit we present in this paper an approach for building a thin inward shell of the surface of an object. The goal is to devise a simple and fast algorithm yet capable of building a topologically sound tetrahedral mesh. The tetrahedral mesh can be used with several different simulation methods, such as the finite element method (FEM).

The main contribution of this paper is a novel tetrahedral mesh generation method, based on surface extrusion and prism tessellation.

**CR Categories:** I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

**Keywords:** Tetrahedral Mesh, Erosion, Extrusion, Tessellation, Shell, Prism

## 1 Introduction

Given a 3D polygonal model created by a 3D artist, it is often a challenge to create a spatial structure for simulating a deformable object. Besides, polygonal models for visual pleasing pictures tends to be highly tessellated. Thus, even if they do not pose any “errors”, creating a tetrahedral mesh directly from the polygonal model, tends to create an enormous amount of tetrahedra. To achieve real time performance, one seeks a more coarse tetrahedral mesh. These are the kind of problems we aim to solve in this paper.

Given a twofold boundary representation of an object, as a connected triangular mesh (a watertight surface), the tetrahedral mesh is built by extruding each triangle inward, that is in opposite direction of the triangle normal. Thus, for each triangle a prism is generated. The result is a volumetric mesh consisting of connected prisms. These prisms can now be tessellated into tetrahedra, thereby creating the first layer of the thin shell tetrahedral mesh. Succeeding layers can be created by recursively applying this approach. Figure 1 illustrates the basic idea. Although the overall idea is simple, the approach is not without problems. Polygonal models are seldom twofold, but suffers from all kind of degeneracies. The idea we have illustrated is obviously capable of handling an open boundary, but cases where edges share more than two neighboring faces, or self-loop edges, are clearly unwanted, since the generated prisms will overlap each other or degenerate into a zero-volume prism.

The prism generation is reminiscent of an erosion operation with a spherical structural element (mathematical morphology). The radius of the sphere corresponds to the extrusion length. It is well known that working directly on the Brep [Sethian 1999] is fast and simple, but topological problems arises easily, such as swallow tails.

In case the given triangle mesh is not a proper mesh, one can apply a mesh reconstruction algorithm [Nooruddin and Turk 2003].

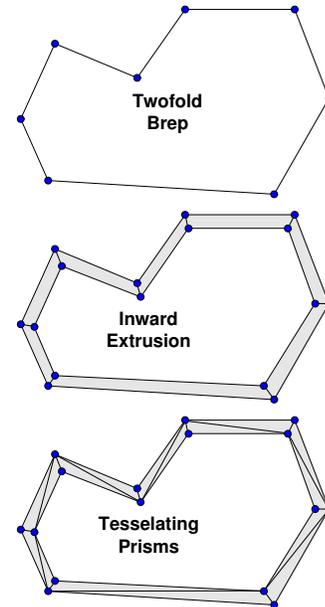


Figure 1: The basic algorithm for generating a thin tetrahedral shell from a boundary representation.

We will require the following properties of the prisms making up a thin shell layer

- No two prisms must be intersecting each other (Neighbors are allowed to touch each other at their common faces).
- No prism must collapse to zero volume or be turned inside out (equivalent to signed volume is always positive).
- All prisms must be convex

Unfortunately, even if we are given a perfect connected twofold triangle mesh, we can get into trouble if we make the inward extrusion too big. This is illustrated in Figure 2. Here, the large extrusion length causes prisms *B* and *C* to become non-convex. Furthermore, *A* and *D*, *B* and *D*, *C* and *A*, and *B* and *C* are overlapping. Fortunately, these degenerate and unwanted prisms can be avoided if the extrusion were made smaller. Thus, we seek an upper bound on how far we can extrude the triangle faces inward, without causing degenerate prisms.

A publicly available implementation of the described algorithm can be found in [OpenTissue 2004].

Existing tetrahedral mesh generation methods create an initial, blocky tetrahedral mesh from a voxelization or signed distance map. Afterwards, nodes are iteratively repositioned, while subsampling tetrahedra to improve mesh quality [Mueller and Teschner 2003; Persson and Strang 2004; Molino et al. 2004]. Our approach differs mainly from these in being surface-based.

\*kenny@diku.dk

†henrikd@diku.dk

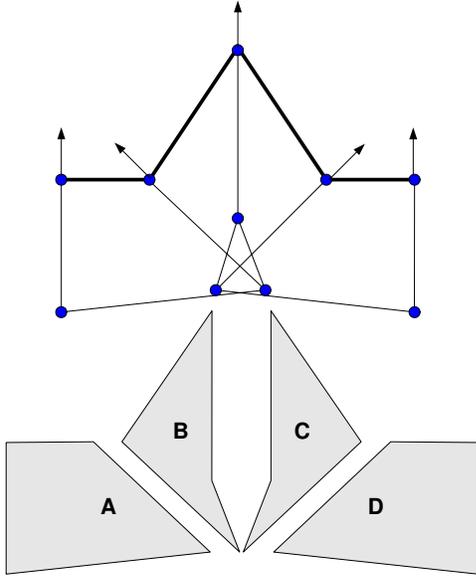


Figure 2: Degenerate prisms results from a too big inward extrusion.

## 2 Inward Extrusion

As a preprocessing step, we compute the pseudo normals for all vertices (the angle-weighted normals [Aanæs and Barentzen 2003]). These will indicate the direction along which a vertex will be extruded.

Given a triangle consisting of three vertices  $\vec{p}_1$ ,  $\vec{p}_2$  and  $\vec{p}_3$ , with angle weighted normals  $\vec{n}_1$ ,  $\vec{n}_2$  and  $\vec{n}_3$ , the inward extruded prism is defined by the six corner points:

$$\begin{aligned} \vec{p}_1 \\ \vec{p}_2 \\ \vec{p}_3 \\ \vec{q}_1(\epsilon) &= \vec{p}_1 - \vec{n}_1 \epsilon \\ \vec{q}_2(\epsilon) &= \vec{p}_2 - \vec{n}_2 \epsilon \\ \vec{q}_3(\epsilon) &= \vec{p}_3 - \vec{n}_3 \epsilon. \end{aligned}$$

The extrusion length is given by  $\epsilon > 0$ . Notation is illustrated in Figure 3. By requiring  $\epsilon$  to be strictly positive, all generated prisms

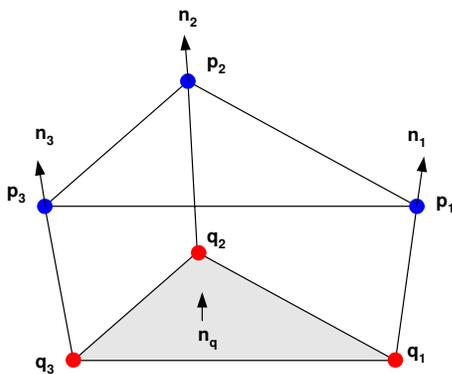


Figure 3: The six corner points defining a prism, and pseudo normals yielding extrusion directions.

must have non-zero volume, as long as prism is not turned inside out. We therefore seek a robust way to determine an upper bound on  $\epsilon$ , such that all generated prisms will be valid.

The direction of the normal of the extruded face,  $\vec{n}_q$ , can be found from  $\vec{q}_1, \vec{q}_2$ , and  $\vec{q}_3$ , using the cross-product:

$$\vec{n}_q(\epsilon) = (\vec{q}_2(\epsilon) - \vec{q}_1(\epsilon)) \times (\vec{q}_3(\epsilon) - \vec{q}_1(\epsilon)).$$

This is a second order polynomial in  $\epsilon$ ,

$$\begin{aligned} \vec{n}_q(\epsilon) &= (\vec{q}_2(\epsilon) - \vec{q}_1(\epsilon)) \times (\vec{q}_3(\epsilon) - \vec{q}_1(\epsilon)) \\ &= ((\vec{p}_2 - \vec{n}_2 \epsilon) - (\vec{p}_1 - \vec{n}_1 \epsilon)) \times ((\vec{p}_3 - \vec{n}_3 \epsilon) - (\vec{p}_1 - \vec{n}_1 \epsilon)) \\ &= ((\vec{p}_2 - \vec{p}_1) + (\vec{n}_1 - \vec{n}_2) \epsilon) \times ((\vec{p}_3 - \vec{p}_1) + (\vec{n}_1 - \vec{n}_3) \epsilon) \\ &= \underbrace{((\vec{p}_2 - \vec{p}_1) \times (\vec{p}_3 - \vec{p}_1))}_{\vec{c}} + \\ &\quad \underbrace{((\vec{p}_2 - \vec{p}_1) \times (\vec{n}_1 - \vec{n}_3) + (\vec{n}_1 - \vec{n}_2) \times (\vec{p}_3 - \vec{p}_1))}_{\vec{b}} \epsilon + \\ &\quad \underbrace{((\vec{n}_1 - \vec{n}_2) \times (\vec{n}_1 - \vec{n}_3))}_{\vec{a}} \epsilon^2 \\ &= \vec{a} \epsilon^2 + \vec{b} \epsilon + \vec{c}. \end{aligned}$$

Observe that  $\vec{c} \neq \vec{0}$ , since its magnitude is equal to twice the area of the triangle being extruded.

To ensure convexity, the dot product of the direction of the normal of the extruded face,  $\vec{n}_q$ , with the pseudo normals,  $\vec{n}_1$ ,  $\vec{n}_2$ , and  $\vec{n}_3$ , must always be positive. That is

$$\begin{aligned} \vec{n}_1 \cdot \vec{n}_q(\epsilon) &> 0 \\ \vec{n}_2 \cdot \vec{n}_q(\epsilon) &> 0 \\ \vec{n}_3 \cdot \vec{n}_q(\epsilon) &> 0. \end{aligned}$$

This yields the following system of constraints,

$$\begin{bmatrix} \vec{n}_1 \cdot \vec{a} & \vec{n}_1 \cdot \vec{b} & \vec{n}_1 \cdot \vec{c} \\ \vec{n}_2 \cdot \vec{a} & \vec{n}_2 \cdot \vec{b} & \vec{n}_2 \cdot \vec{c} \\ \vec{n}_3 \cdot \vec{a} & \vec{n}_3 \cdot \vec{b} & \vec{n}_3 \cdot \vec{c} \end{bmatrix} \begin{bmatrix} \epsilon^2 \\ \epsilon \\ 1 \end{bmatrix} > \mathbf{0}.$$

We solve for the smallest positive  $\epsilon$  fulfilling the system of constraints. That is, each row represents the coefficient of a second order polynomial in  $\epsilon$ , thus for each row we find the two roots of the corresponding polynomial. The three rows yields a total of 6 roots. If no positive root exist, then  $\epsilon = \infty$ , otherwise  $\epsilon$  is set equal to the smallest positive root.

Observe that the third column of the coefficient matrix is always positive (by the property of the angle weighted normals). The first column can be interpreted as an indication of, whether the corresponding extrusion line is trying to “shrink” ( $< 0$ ) or “enlarge” ( $> 0$ ) the extruded face. The middle column is difficult to interpret. As far as we can tell it resembles an indication of the skewness of the resulting prism.

In fact, the tree convexity constraints ensure that no neighboring prism will intersect each other, nor will the prism turn its inside out (ie. flipping the extruded face opposite the original face).

The maximum extrusion length for the entire layer can be found by iterating over each prism. For the  $i$ 'th prism the extrusion length  $\epsilon^i$  is computed. The maximum extrusion length of the layer is found as

$$\epsilon = \min(\epsilon^0, \dots, \epsilon^{n-1}).$$

Afterwards, it is a simple matter to compute the actual extrusion and generating the prisms.

### 3 Prism Generation

The technique in the previous section guarantee that prisms will be convex, and that no intersections occurs between neighboring prisms. However, degenerate prisms can still occur whenever the extruded face collapses to a line or to a point. These are shown in Figure 4. These degenerate cases must be marked, such that the fol-

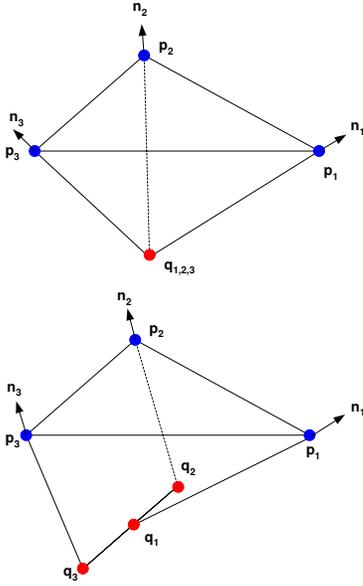


Figure 4: Point-degenerate and line-degenerate prisms.

lowing tessellation can take them into consideration. The problem is how to detect these degenerate cases. If no upper bound were found on the extrusion length for a prism, we can trivially reject the prism. However, if an upper extrusion bound were computed, the prism might degenerate. As a first step in marking the degenerate prisms, we iterate over all those prisms where an upper bound were computed. If the computed upper bound for a prism is equal to the extrusion length, then it might be a degenerate prism. For each of these possible degenerate prisms we first test whether

$$\|\vec{n}_q(\epsilon)\| \leq \gamma,$$

where  $\gamma$  is a user specified threshold to counter numerical precision problems. If this criteria is fulfilled, we clearly know that we are dealing with a degenerate prism.

The degenerate prisms can be classified as either point-degenerate or line-degenerate, as shown in Figure 4. The line-degenerate case is given by the criteria

$$(\vec{q}_2(\epsilon) - \vec{q}_1(\epsilon)) \neq \vec{0} \quad \text{or} \quad (\vec{q}_3(\epsilon) - \vec{q}_1(\epsilon)) \neq \vec{0}.$$

If this criteria is not fulfilled we have a point-degenerate case. A pseudo code version of the algorithm is shown in Figure 5. The degenerate cases do not only influence the prism tessellation (which we will treat in the next section). If another thin shell layer is to be generated, then the original mesh faces can no longer be used. A simple 2D case is shown in Figure 6. Notice that the bold red faces were used when creating layer 0, but they vanish when layer 1 is created. Therefore, if another layer should be generated, a new connected triangular mesh must be formed from the extruded faces of the non-degenerate prisms.

In Figure 7 the surface mesh generation algorithm is shown in pseudo code.

---

```

algorithm markDegeneratePrism( $\epsilon, \gamma$ )
for each prism  $p$  do
  if  $\epsilon^p = \epsilon$  then
    if  $\|\vec{n}_q(\epsilon)\| \leq \gamma$  then
      if  $(\vec{q}_2(\epsilon) - \vec{q}_1(\epsilon)) \neq \vec{0}$ 
        or  $(\vec{q}_3(\epsilon) - \vec{q}_1(\epsilon)) \neq \vec{0}$  then
        mark  $p$  as line-degenerate
      else
        mark  $p$  as point-degenerate
    end if
  end if
end if
next  $p$ 
End algorithm

```

---

Figure 5: Pseudo code for marking degenerate prisms.

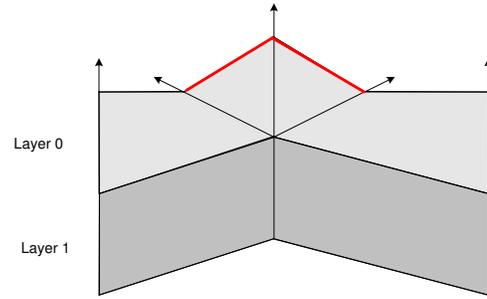


Figure 6: Degenerate cases affect succeeding layers.

### 4 Prism Tessellation

For non-degenerate prisms, having 6 corners, the minimum number of tetrahedra we can tessellate the prism into, is 3. This is shown in Figure 8. The problem with this approach is that the extruded sides of the prism will be triangulated. One therefore has to ensure, that the tessellation of neighboring prisms agree with each other. This is illustrated in Figure 9. As seen in Figure 9 it becomes a global combinatorial problem to match the tessellation of neighbors against each other.

If we use the centroid of the prism as the apex for creating tetrahedra then a prism can be tessellated into eight tetrahedra, as shown in Figure 10. With this approach the tessellation is no longer a global problem, since the tessellation of each prism side can be chosen in-

---

```

algorithm build-surface(Mesh :  $M$ )
for each prism  $p$  do
  for  $i = 1, 2, 3$  do
    if  $\vec{q}_i \notin M$  then
      add  $\vec{q}_i$  to  $M$ 
    end if
  next  $i$ 
  if  $p$  not degenerate then
    add face  $\vec{q}_1, \vec{q}_2, \vec{q}_3$  to  $M$ 
  end if
next  $p$ 
End algorithm

```

---

Figure 7: Pseudo code for generating surface mesh for next shell layer.

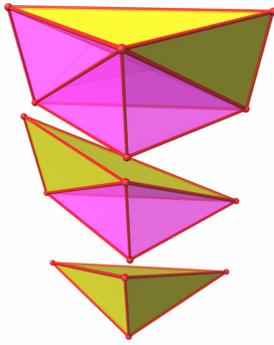


Figure 8: Prism tessellated into 3 tetrahedra.

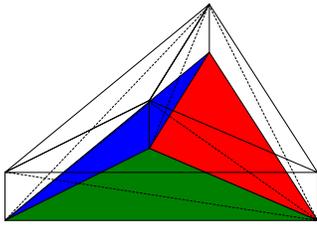


Figure 9: Tessellation of neighboring prisms must be consistent.

dependently of each other.

Both approaches deals nicely with the point-degenerate case. Since a point-degenerate prism is already a tetrahedron, there is no need to tessellate it. However, since the point-degenerate case only have triangular sides, it can never be a direct neighbor with a non-degenerate prism. It can only be neighbor with other point-degenerate cases or line-degenerate cases, which have both rectangular and triangular sides, as shown in Figure 11.

In the following we disregard degeneracies and consider the three-tetrahedra tessellation strategy. This is because it is more attractive due to the lower tetrahedra count.

A prism can be tessellated into three tetrahedra in 6 different ways. In order to classify the 6 types of tessellation, we will mark the rectangular sides of a prism as falling ( $F$ ) or rising ( $R$ ). The edge type depends on whether the tessellation edge is falling or rising as we travel along the extruded prism face in counter clock wise order. See Figure 12. We observe that the three-tetrahedra tessellation strategy will always have two prism sides of the same type, and the last side of opposite type. Thus, we can only have 6 different patterns, as shown in Table 1. The consistency requirement implies that if one side of a prism is marked as  $F$ , then the neighboring prism will have marked the same side as  $R$ . In short, no neighboring prisms will have a side marked with the same type.

A simple tessellation example is shown in Figure 13. Here a tetrahedron mesh is being tessellated. The tetrahedron has been cut up and layed out in 2D. Triangle edges corresponds to rectangular sides of prisms. Let us apply a brute-force strategy to the combi-

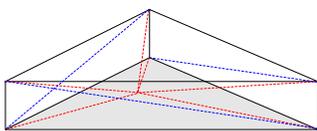


Figure 10: Prism tessellated into eight tetrahedra.

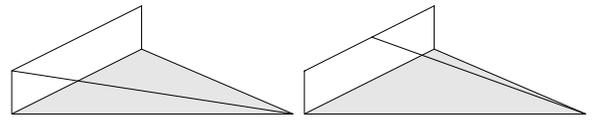


Figure 11: Nice vs. bad line case.

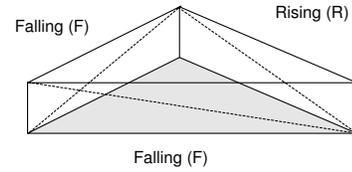


Figure 12: Classification of prism sides as falling ( $F$ ) or rising ( $R$ ).

natorial problem of the tessellation as follows: We start at a single prism and choose one of the 6 tessellation types. Then we visit the neighboring prisms and choose a tessellation type that agree with the immediate neighbor prisms, which have already been tessellated. This is a breadth first traversal over the prisms.

The method is not fail-safe, since inconsistency can arise, as shown in Figure 14. Here, the middle prism is the last prism to be visited by the traversal. Clearly, it is impossible to assign a tessellation type to the prism, since all three sides should have the same type. We can repair the inconsistency by picking one of the neighboring prisms and flipping the type of its shared edge. This action will not change the type of any of the edges marked with arrows in Figure 14. Therefore, the repairing action do not cause a rippling effect through the prisms, and inconsistencies are not introduced at other places.

Fixing inconsistency in this way is attractive, since it offers a local solution to a global problem. However, sometimes we might end up in a dead-lock where no local solution can be found, as is shown in the top of Figure 15.

This time the drawing in the figure resembles a small local view of a larger mesh. Notice that none of the edges shared with the inconsistent prism can be flipped, without creating an inconsistent neighboring prism. The problem is that all the edges marked with arrows are of the same type.

The solution to the problem is shown in the bottom of Figure 15. We let the inconsistency ripple as water waves over to neighboring prisms, in a search for a single prism, where an edge flip does not give rise to a new inconsistency. When such a prism is encountered, we track the trajectory of the ripple wave-front back to the originating inconsistent prism, and flip all shared edges lying on this path. In Figure 16 we have shown the result of the rippling. Notice that two edges are flipped. These are the edges lying on the path to the prism that could be flipped. Also notice that all edges marked with arrows are unaffected by the rippling action. This property ensures that the rippling action will not cause inconsistencies in any prisms elsewhere in the mesh.

A pseudo code of the tessellation-pattern-finding algorithm is shown in Figure 17. Our proposed tessellation pattern algorithm

$F$	$R$	$R$
$R$	$F$	$R$
$R$	$R$	$F$
$R$	$F$	$F$
$F$	$R$	$F$
$F$	$F$	$R$

Table 1: The 6 Four-Tetrahedra Tessellation types.

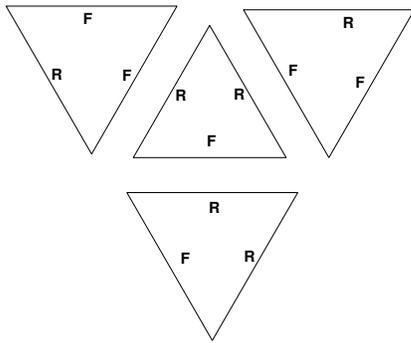


Figure 13: Tesselation Example. A simple 3D mesh (a tetrahedron) have been cut up and layed down in 2D. Triangles correspond to prisms in the thin shell.

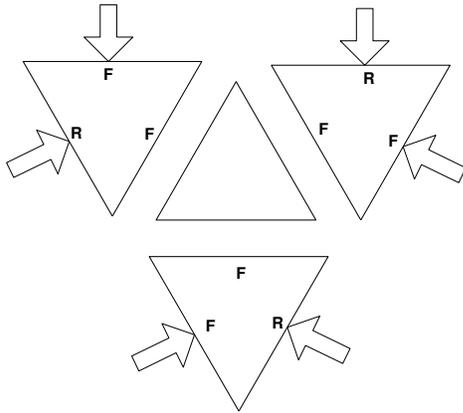


Figure 14: Inconsistent Tesselation Example. The middle prism will have the same type on all sides, which is illegal.

is an ad-hoc solution for the problem at hand. We do not have a formal proof, stating that it is always possible to find a consistent pattern of rising and falling tesselation edges.

## 5 Results

We have implemented the extrusion length computation and tesselation-pattern algorithm. Currently the implementation detects degenerate prisms, but do not tesselate these.

In our test cases we have chosen 14 meshes of increasing size, that were all scaled to be within the unit-cube. A single layer shell were computed, given a user specified maximum extrusion length. In Table 2, performance statistics are listed, together with polygon count, extrusion lengths and rippling action information. The timings for a single layer construction are cheap, and appears to scale lineary with mesh size. The resulting tetrahedral meshes are visualized in Figure 18. As seen in Table 2, the test-cases: diku, teapot, propeller, funnel, cow, and bend have a surprisngly small extrusion limit. The remaining test cases show excellent extrusion limits. Figure 19 shows the prisms corresponding to the minimum extrusion limit. Notice that in all cases, where the limit is unexpected small, the limit is caused by small faces or long slivers on sharp ridges. Figure 20 shows cut-through views of the cylinder, pointy, tube, sphere, teapot, funnel, bowl, and torus meshes. Notice how thin the teapot and funnel are.

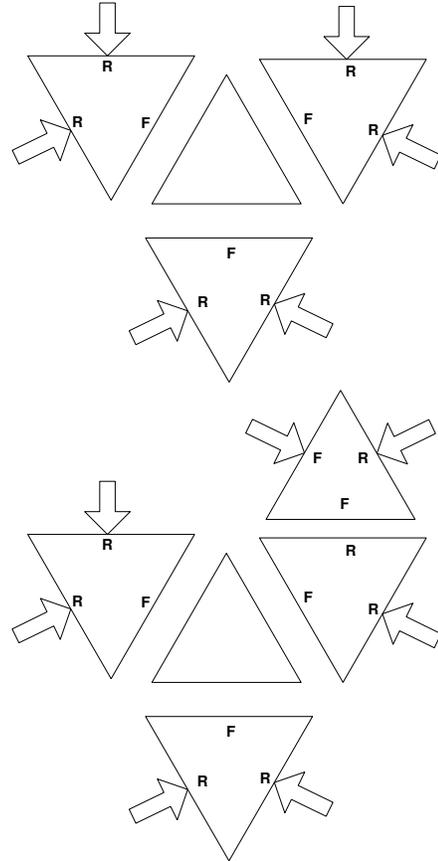


Figure 15: Top picture shows a dead-locked inconsistent tesselation. The bottom picture shows that inconsistency problem have been propagated to neighboring prisms further away.

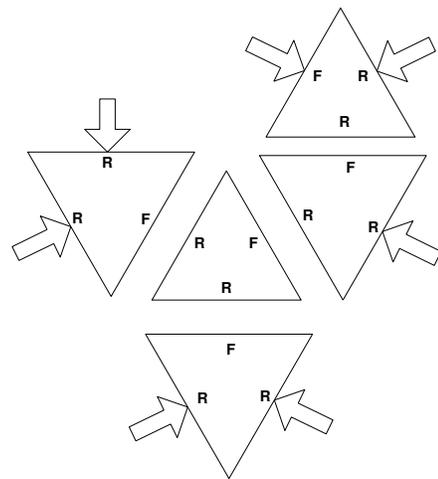


Figure 16: The rippling solution to the dead-locked case shown in Figure 15.

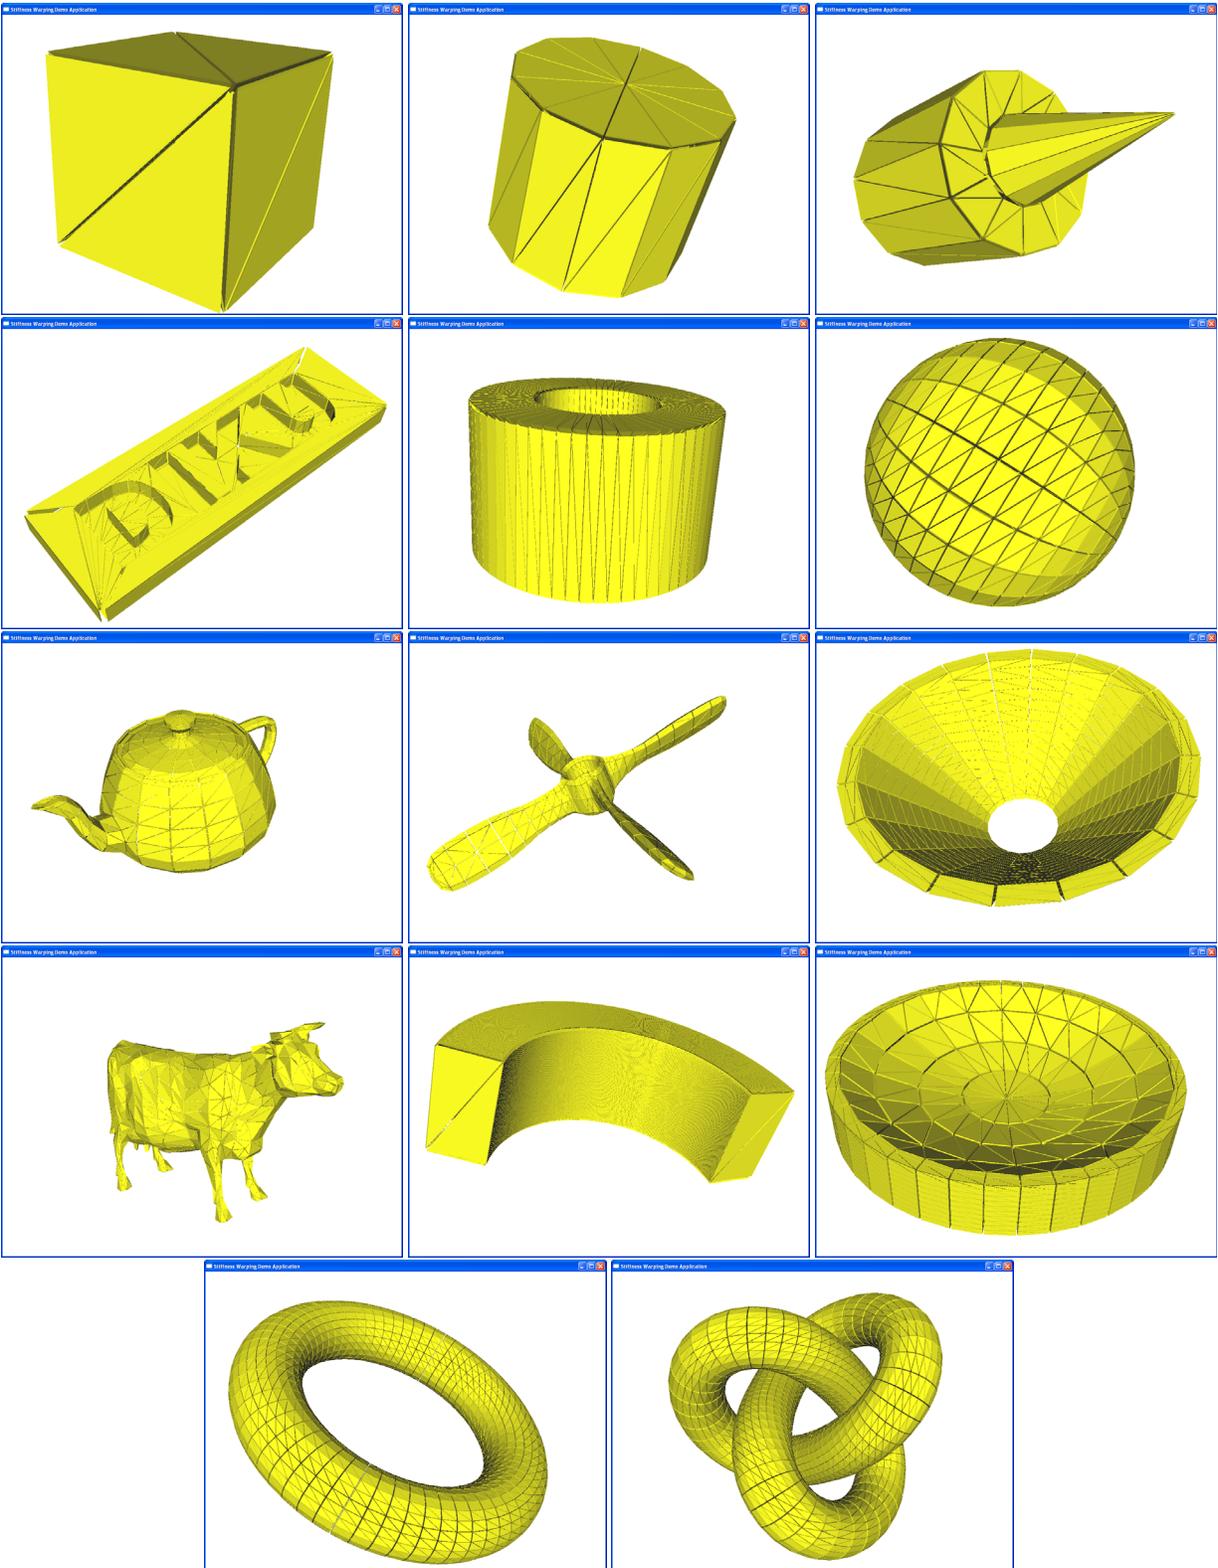


Figure 18: Tessellation results of the 14 meshes.

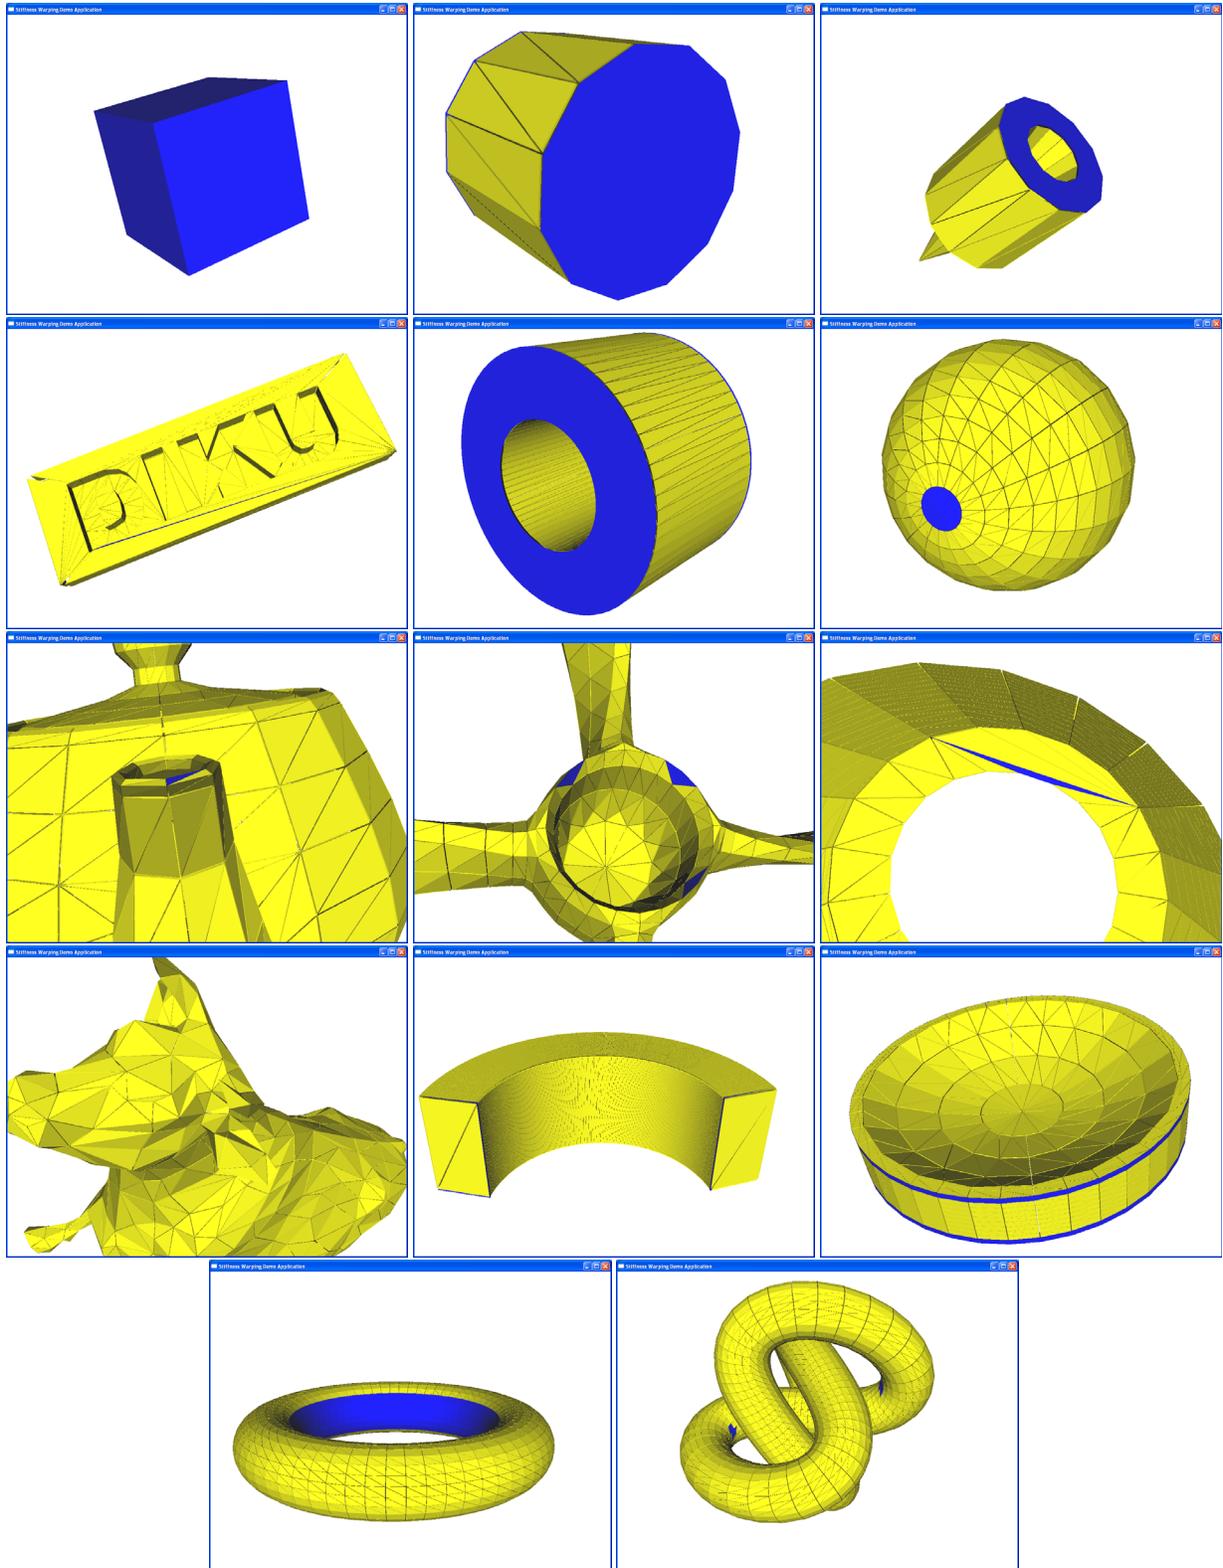


Figure 19: Prisms marked with blue have minimum extrusion limit.

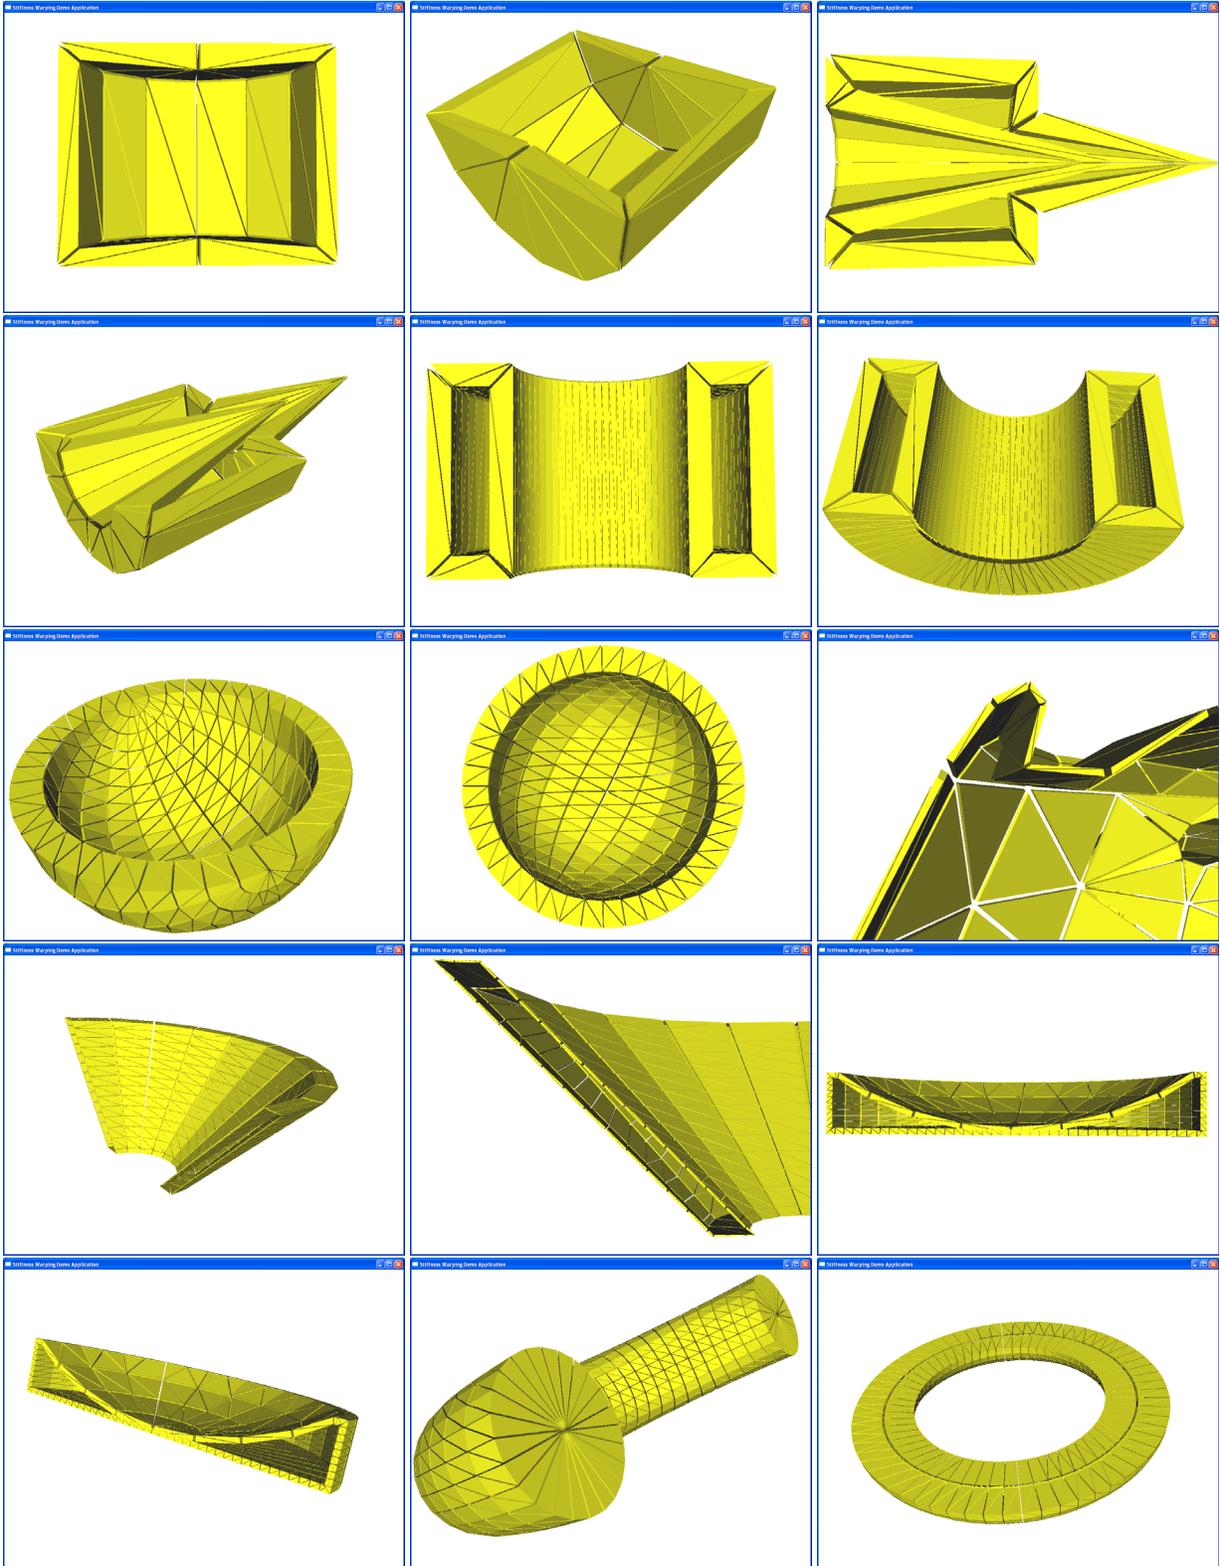


Figure 20: Various cut-through views of a few selected meshes, illustrating the thin shell.

---

```

algorithm tessellation-pattern()
  Queue  $Q$ 
  Push first prism onto  $Q$ 
  While  $Q$  not empty do
    Prism  $p = \text{pop}(Q)$ 
    mark  $p$  as visited
    if neighbors is not tessellated then
      pick random pattern of  $p$ 
    else if exist consistent pattern with neighbors
      assign consistent pattern to  $p$ 
    else
      if exist neighbor that can be flipped
        flip edge type of shared edge with  $p$ 
        assign constant pattern to  $p$ 
      else
        perform-rippling
      end if
    end if
  for all unvisited neighbors  $n$  of  $p$  do
    push( $Q, n$ )
  next  $n$ 
End while
End algorithm

```

---

Figure 17: Pseudo code for determining tessellation pattern.

## 6 Discussion

We have omitted the problem of shell layers overlapping from opposite sides. In Figure 20 the problem is seen in the case of the bowl mesh. Our solution to the problem have been to ignore it. The user can always choose a smaller maximum extrusion length.

Degenerate prisms were ignored, in the sense that we are able to detect if they occur, but our tessellation pattern algorithm is not yet capable of handling them. No degenerate prisms were generated for any of the examples in our result section.

The global computation of the extrusion length work fairly well for some meshes, but for others a surprisingly small extrusion length is found. Long slivers and small faces lying close to sharp ridges are the reason for this phenomena, as can be seen in Figure 19. Thus, we conclude that, not surprisingly, the algorithm is highly dependent on both the shape of the object, but also upon the tessellation of the object surface. It appears that a good uniform tessellation work best. Abrupt tessellation with large aspect ratios results in small extrusion lengths. Mesh reconstruction [Nooruddin and Turk 2003] could be used as a preprocessing step to create a more suitable tessellation.

Another avenue for circumventing the problem of a small global extrusion, might be to investigate the possibility of having non-global extrusion lengths, i.e. a varying extrusion length over the mesh, adapting itself to take the local maximum length without causing degenerate prisms. We believe this is an interesting thought an leave it for future work.

Our results indicate that our tessellation pattern algorithm works: We have not yet encountered an unsolvable problem. We believe this shows, that the combinatorial problem of finding the tessellation pattern, is at least solvable in practice. From a theoretical viewpoint, a proof of existence would be very interesting, and we leave this for future work.

## 7 Conclusion

In this paper we have presented preliminary results, showing that it is possible to generate a thin shell, without any topological errors.

	$ F $	$ R $	$\delta$	$\epsilon$	time(secs.)
box	12	0	0.1	0.866	0
cylinder	48	0	0.1	0.431	0
pointy	96	0	0.083	0.083	0
diku	288	0	0.004	0.004	0
tube	512	0	0.1	0.174	0.01
sphere	760	0	0.1	0.476	0.01
teapot	1056	1	0.001	0.001	0.01
propeller	1200	2	0.004	0.004	0.02
funnel	1280	1	0.003	0.003	0.029
cow	1500	0	0.001	0.001	0.02
bend	1604	0	0.006	0.006	0.029
bowl	2680	0	0.017	0.017	0.04
torus	3072	0	0.099	0.099	0.059
knot	5760	0	0.1	0.102	0.089

Table 2: Performance Statistics on 14 different test cases. In all cases the end user requested a shell thickness of 0.1. The  $\delta$ -column shows the actual shell thickness produced. The  $\epsilon$ -column shows the extrusion limit. The  $|F|$ -column gives the face count of the meshes. The  $|R|$ -column gives the number of times the ripple action were invoked. The zero entries in the time column indicates that the duration were not measureable by the timing method.

As pointed out in the previous section, there are many unsolved issues to be dealt with.

Our motivation for this work were to create a volumetric mesh with low tetrahedra count for animation purpose. Due to the early stage of this work, we have not yet validated whether our approach is useful for animation.

## References

- AANES, H., AND BÆRENTZEN, J. A. 2003. Pseudo-normals for signed distance computation. In *Proceedings of VISION, MODELING, AND VISUALIZATION*.
- MOLINO, N., BRIDSON, R., TERAN, J., AND FEDKIW, R. 2004. Adaptive physics based tetrahedral mesh generation using level sets. (in review).
- MUELLER, M., AND TESCHNER, M. 2003. Volumetric meshes for real-time medical simulations. In *Proc. BVM*, 279–283.
- NOORUDDIN, F., AND TURK, G. 2003. Simplification and repair of polygonal models using volumetric techniques. *IEEE Transactions on Visualization and Computer Graphics* 9, 2, 191–205.
- OPENTISSUE, 2004. <http://www.diku.dk/forskning/image/research/opentissue/>.
- PERSOON, P.-O., AND STRANG, G. 2004. A simple mesh generator in matlab. *SIAM Review* 46, 2 (June), 329–345.
- SETHIAN, J. A. 1999. *Level Set Methods and Fast Marching Methods. Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Cambridge University Press. Cambridge Monograph on Applied and Computational Mathematics.

# Talking Faces – a State Space Approach

Tue Lehn-Schiøler

The Technical University of Denmark  
Informatics and Mathematical Modelling

Email: tls@imm.dtu.dk

July 5, 2004

## Abstract

In this paper a system that transforms speech waveforms to animated faces are proposed. The system relies on continuous state space models to perform the mapping, this makes it possible to ensure video with no sudden jumps and allows continuous control of the parameters in 'face space'.

The performance of the system is critically dependent on the number of hidden variables, with too few variables the model cannot represent data, and with too many overfitting is noticed.

To create a photo realistic image an Active Appearance Model is used. Simulations are performed on recordings of 3-5 sec. video sequences with sentences from the Timit database. From a subjective point of view the model is able to construct an image sequence from an unknown noisy speech sequence even though the number of training examples are limited.

## 1 Introduction

The motivation for transforming a speech signal into lip movements is at least threefold. Firstly, the language synchronization of movies often leaves the actors mouth moving while there is silence or the other way around, this looks rather unnatural. If it was possible to manipulate the face of the actor to match the actual speech it would be much more pleasant to view synchronized movies (and a lot easier to make cartoons). Secondly, even with increasing bandwidth sending images via the cell phone is quite expensive, therefore, a system that allows single images to be sent and models the face in between would be useful. The technique will also make it possible for hearing impaired people to lip read over the phone. If the person in the other end does not have a camera on her phone, a model image can be used to display the facial movements. Thirdly, when producing agents on a computer (like Windows Office Mr. clips) it would make communication more plausible if the agent could interact with lip movements corresponding to the (automatically generated) speech.

The idea of extracting phonemes or similar high-level features from the speech signal before performing the mapping to the mouth position has been widely used in the lip-sync community. Goldenthal [9] suggested a system called "Face Me!". He extracts phonemes using Statistical Trajectory Modeling. Each phoneme is then associated with a mouth position (keyframe). In Mike Talk

[6], phonemes are generated from text and then mapped onto keyframes, however, in this system trajectories linking all possible keyframes are calculated in advance thus making the video more seamless. In "Video rewrite" [2] phonemes are again extracted from the speech, in this case using Hidden Markov Models. Each triphone (three consecutive phonemes) has a mouth sequence associated with it. The sequences are selected from training data, if the triphone does not have a matching mouth sequence in the training data, the closest available sequence is selected. Once the sequence of mouth movements has been determined, the mouth is mapped back to a background face of the speaker. Other authors have proposed methods based on modeling of phonemes by correlational HMM's [21] or neural networks [10].

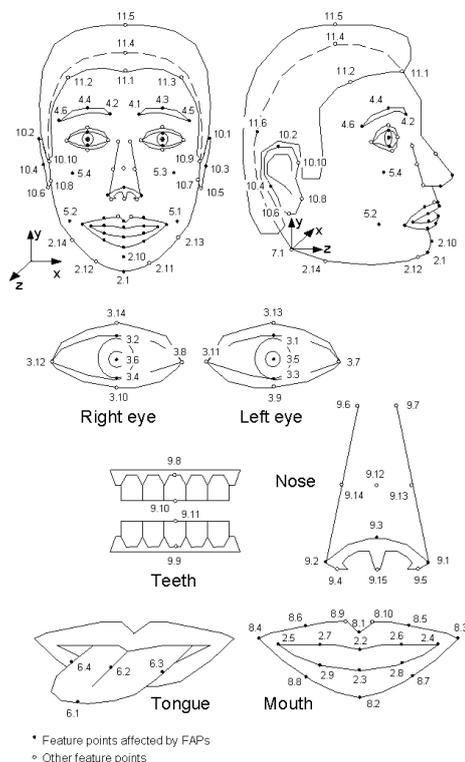
Methods where speech is mapped directly to facial movement are not quite as popular as phoneme based methods. However, in 'Picture my voice' [13], a time dependent neural network, maps directly from  $11 \times 13$  Mel Frequency Cepstral Coefficients (MFCC) as input to 37 facial control parameters. The training output is provided by a phoneme to animation mapping but the trained network does not make use of the phoneme representation. Also Brand [1] has proposed a method based on (entropic) HMM's where speech is mapped directly to images. In [17] Nakamura presents an overview of methods using HMM's, the first MAP-V converts speech into the most likely HMM state sequence and then uses a table lookup to convert into visual parameters. In an extended version MAP-EM the visual parameters are estimated using the EM algorithm. Methods that do not rely on phoneme extraction has the advantage that they can be trained to work on all languages, and that they are able to map non-speech sounds like yawning and laughing.

There are certain inherent difficulties in mapping from speech to mouth positions an analysis of these can be found in [7]. The most profound is the confusion between visual and auditive information. The mouth position of sounds like /b/,/p/ and /m/ or /k/,/n/ and /g/ can not be distinguished even though the sounds can. Similarly the sounds of /m/ and /n/ or /b/ and /v/ are very similar even though the mouth position is completely different. This is perhaps best illustrated by the famous experiment by McGurk [16]. Thus, when mapping from speech to facial movements, one cannot hope to get a perfect result simply because it is very difficult to distinguish whether a "ba" or a "ga" was spoken.

The rest of this paper is organized in three sections, section 2 focuses on feature extraction in sound and images, in section 3 the model are described. Finally experimental results are presented in section 4.

## 2 Feature extraction

Many different approaches has been taken for extraction of sound features. If the sound is generated directly from text phonemes can be extracted directly and there is no need to process the sound track [6]. However, when a direct mapping is performed one can choose from a variety of features. A non-complete list of possibilities include Perceptual Linear Prediction or J-Rasta-PLP as in [1, 5], Harmonics of Discrete Fourier Transform as in [15], Linear Prediction Coefficients as in [12] or Mel Frequency Cepstral Coefficients [9, 10, 13, 17]. In this work the sound is split into 25 blocks per second (the same as the image

Figure 1: Facial feature points <sup>1</sup>.

frame rate) and 13 MFCC features are extracted from each block. To extract features from the images an Active Appearance model (AAM) [3] is used. The use of this model for lipreading has previously been studied by Mathews et al. [14]. AAM's are also useful for low bandwidth transmission of facial expressions [20]. In this work an implementation by Mikkel B. Stegman [19] is used. For the extraction a suitable subset of images in the training set are selected and annotated with points according to the MPEG-4 facial animation standard (Fig. 1). Using these annotations a 14-parameter model of the face is created. Thus, with 14 parameters it is possible to create a photo realistic image of any facial expression seen in the training set. Once the AAM is created the model is used to track the lip movements in the image sequences, at each point the 14 parameters are picked up. In Fig. 2 the result of the tracking is shown for a single representative image.

### 3 Model

In this work the mapping from sound to images is performed by a Kalman filter. The implementation makes use of the toolbox written by Kevin Murphy (<http://www.ai.mit.edu/~murphyk/Software>).

<sup>1</sup>Image from [www.research.att.com/projects/AnimatedHead](http://www.research.att.com/projects/AnimatedHead)

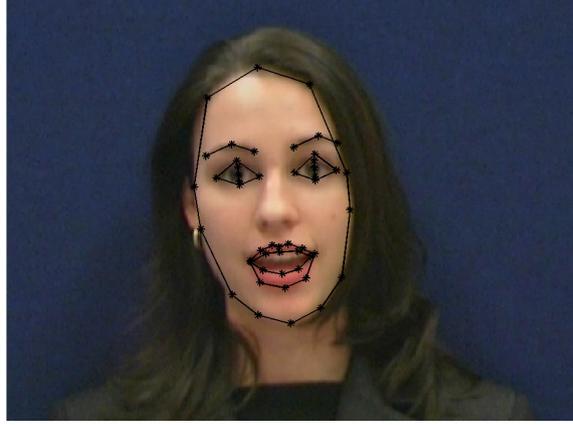


Figure 2: Image with automatically extracted feature points. The facial feature points are selected from the MPEG-4 standard

Normally, HMM's or a Neural network is used for the mapping. In the case of HMM's a series of models are created, each one is trained on a specific subset of data. At each time step model that has the highest likelihood wins and is then responsible for producing the image. In this work the entire sequence is considered at once and only a single state space model is trained. In case of the Kalman filter the model set up is as follows:

$$\mathbf{x}_k = \mathbf{A}\mathbf{x}_{k-1} + \mathbf{n}_k^x \quad (1)$$

$$\mathbf{s}_k = \mathbf{B}\mathbf{x}_k + \mathbf{n}_k^s \quad (2)$$

$$\mathbf{i}_k = \mathbf{C}\mathbf{x}_k + \mathbf{n}_k^i \quad (3)$$

In this setting  $\mathbf{i}_k$  is the image features at time  $k$ ,  $\mathbf{s}_k$  is the sound features and  $\mathbf{x}_k$  is a hidden variable without physical meaning.  $\mathbf{x}$  can be thought of as some kind of brain activity controlling what is said. Each equation has i.i.d. Gaussian noise component  $\mathbf{n}$  added to it.

During training both sound and image features are known, and the two observation equations can be collected in one.

$$\begin{pmatrix} \mathbf{s}_k \\ \mathbf{i}_k \end{pmatrix} = \begin{pmatrix} \mathbf{B} \\ \mathbf{C} \end{pmatrix} \mathbf{x}_k + \begin{pmatrix} \mathbf{n}_k^s \\ \mathbf{n}_k^i \end{pmatrix} \quad (4)$$

By using the EM algorithm [4, 8] on the training data, all parameters  $\{\mathbf{A}, \mathbf{B}, \mathbf{C}, \Sigma^x, \Sigma^s, \Sigma^i\}$  can be found.  $\Sigma$ 's are the diagonal covariance matrices of the noise components.

When a new sound sequence arrives Kalman filtering (or smoothing) can be applied to equations (1,2) to obtain the hidden state  $\mathbf{x}$ . Given  $\mathbf{x}$  the corresponding image features can be obtained by multiplication,  $\mathbf{i}_k = \mathbf{C}\mathbf{x}_k$ . If the intermediate smoothing variables are available the variance on  $\mathbf{i}_k$  can also be calculated.

## 4 results

The data used is taken from the vidtimit database [18]. The database contains recordings of large number of people each uttering ten different sentences while facing the camera. The sound recordings are degraded by fan-noise from the recording pc. In this work a single female speaker is selected, thus 10 different sentences are used, nine for training and one for testing.

To find the dimension of the hidden state ( $\mathbf{x}$ ), the optimal parameters for the KF were found for varying dimensions. For each model the likelihood on training and test sequences were calculated, the result is shown in Fig. 3 and Fig. 4.

The test likelihood provides a statistical measure of the quality of the model and provides a way of comparing models. This allows comparison between different model approaches. Unfortunately the likelihood is not necessarily a good measure of the quality of a model prediction. If the distributions in the model are broad, i.e. the model has high uncertainty, it can describe data well, but, is not a good generative model.

Looking at the results in Fig. 3 and Fig. 4 it is seen that the likelihood of a KF has a peak in the test data around 40 hidden states. In Fig. 5 snapshots from the KF sequence are provided for visual inspection, the entire sequence is available at <http://www.imm.dtu.dk/~tls/code/facedemo.php>, where other demos can also be found.

No precise metric exist for evaluation of synthesized lip sequences. The distance between facial points in the true and the predicted image would be one way, another way would be to measure the distance between the predicted feature vector and the feature vector extracted from the true image. However, the ultimate evaluation of faces can be only provided by human interpretation. Unfortunately it is difficult to get an objective measure this way. One possibility would be to get a hearing impaired person to lipread the generated sequence, another to let people try to guess which sequence was real and which was computer generated. Unfortunately, such test are time and labor demanding. Further more these subjective test does not provide an error function that can be optimized directly.

## 5 Conclusion

A speech to face mapping system relying on state space models is proposed. The system makes it possible to easily train a unique face model that can be used to transform speech into facial movements. The training set must contain all sounds and corresponding face gestures, but there are no language or phonetic requirements to what the model can handle.

In this work a linear model is used but future work will investigate the use of nonlinear models by using particle filtering and Markov Chain Monte Carlo methods. Other future work include extracting emotions from the speech and mapping them to the face.

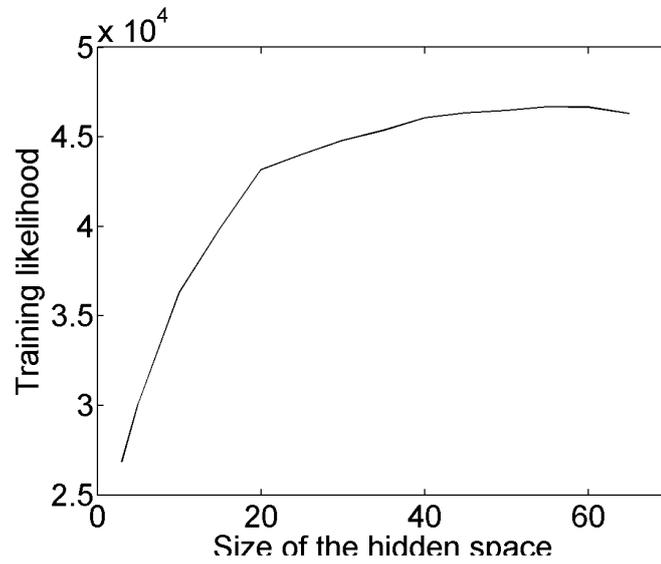


Figure 3: The likelihood evaluated on the training data. The Kalman filter is able to utilize the extra dimension to improve the training result.

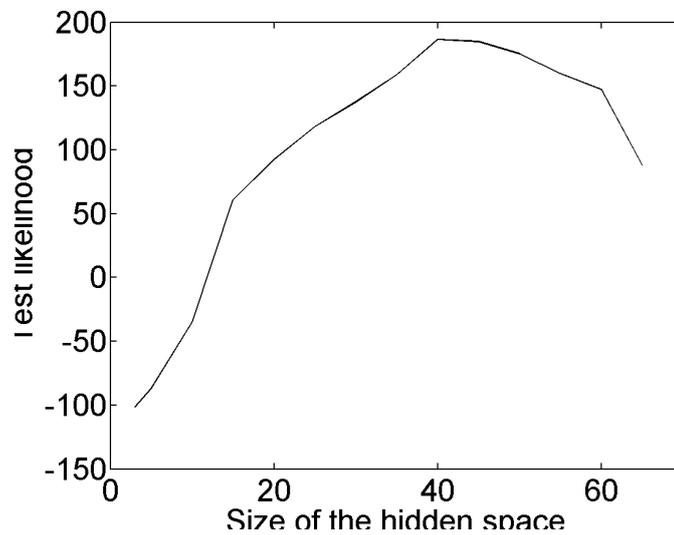


Figure 4: The likelihood evaluated on the test data. The Kalman filter improves performance as more hidden dimensions are added and overfitting is seen for high number of hidden states.

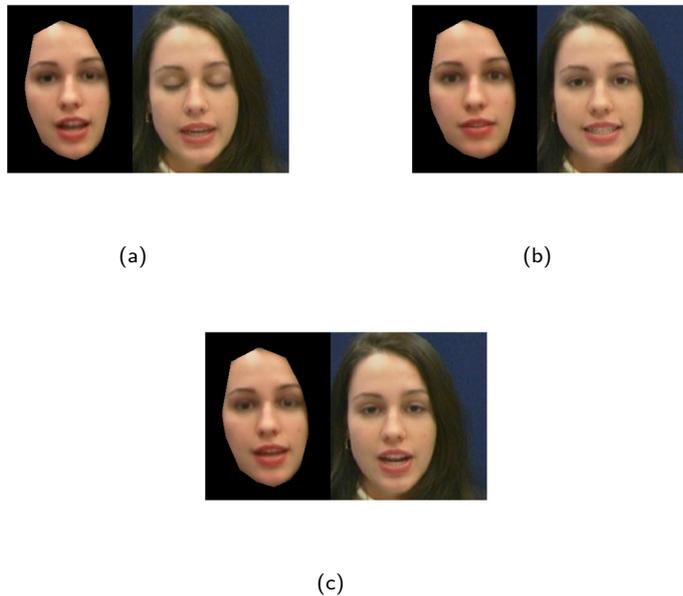


Figure 5: Characteristic images taken from the test sequence when using the Kalman filter. The predicted face is to the left and the true face to the right.

## References

- [1] Matthew Brand. Voice puppetry. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 21–28. ACM Press/Addison-Wesley Publishing Co., 1999.
- [2] Christoph Bregler, Michele Covell, and Malcolm Slaney. Video rewrite: driving visual speech with audio. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 353–360. ACM Press/Addison-Wesley Publishing Co., 1997.
- [3] T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance models. *Proc. European Conference on Computer Vision*, 2:484–498, 1998.
- [4] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *JRSSB*, 39:1–38, 1977.
- [5] S. Dupont and J. Luetttin. Audio-visual speech modelling for continuous speech recognition. *IEEE Transactions on Multimedia*, 2000.
- [6] T. Ezzat and T. Poggio. Mike talk: a talking facial display based on morphing visemes. *Proc. Computer Animation IEEE Computer Society*, pages 96–102, 1998.
- [7] Lavagetto F. Converting speech into lip movements: A multimedia telephone for hard of hearing people. *IEEE Trans. on Rehabilitation Engineering*, 3(1), 1995.

- [8] Z. Ghahramani and G.E. Hinton. Parameter estimation for linear dynamical systems. Technical report, 1996. University of Toronto, CRG-TR-96-2.
- [9] William Goldenthal, Keith Waters, Thong Van Jean-Manuel, and Oren Glickman. Driving synthetic mouth gestures: Phonetic recognition for faceme! In *Proc. Eurospeech '97*, pages 1995–1998, Rhodes, Greece, 1997.
- [10] Pengyu Hong, Zhen Wen, and Thomas S. Huang. Speech driven face animation. In Igor S. Pandzic and Robert Forchheimer, editors, *MPEG-4 Facial Animation: The Standard, Implementation and Applications*. Wiley, Europe, July 2002.
- [11] T. Lehn-Schioler, L. K. Hansen, and J. Larsen. Mapping from speech to images using continuous state space models. In *Joint AMI/PASCAL/IM2/M4 Workshop on Multimodal Interaction and Related Machine Learning Algorithms*, April 2004.
- [12] J. P. Lewis. Automated lip-sync: Background and techniques. *J. Visualization and Computer Animation*, 2, 1991.
- [13] Dominic W. Massaro, Jonas Beskow, Michael M. Cohen, Christopher L. Fry, and Tony Rodriguez. Picture my voice: Audio to visual speech synthesis using artificial neural networks. *Proc. AVSP 99*, 1999.
- [14] I. Matthews, T.F. Cootes, J.A. Bangham, S. Cox, and R. Harvey. Extraction of visual features for lipreading. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(2):198–213, 2002.
- [15] David F. McAllister, Robert D. Rodman, Donald L. Bitzer, and Andrew S. Freeman. Speaker independence in automated lip-sync for audio-video communication. *Comput. Netw. ISDN Syst.*, 30(20-21):1975–1980, 1998.
- [16] H. McGurk and J. W. MacDonald. Hearing lips and seeing voices. *Nature*, 264:746–748, 1976.
- [17] Satoshi Nakamura. Statistical multimodal integration for audio-visual speech processing. *IEEE Transactions on Neural Networks*, 13(4), July 2002.
- [18] C. Sanderson and K. K. Paliwal. Polynomial features for robust face authentication. *Proceedings of International Conference on Image Processing*, 3:997–1000, 2002.
- [19] M. B. Stegmann. Analysis and segmentation of face images using point annotations and linear subspace techniques. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, August 2002. <http://www.imm.dtu.dk/pubdb/p.php?922>.
- [20] B.J. Theobald, S.M. Kruse, J.A. Bangham, and G.C. Cawley. Towards a low bandwidth talking face using appearance models. *Image and Vision Computing*, 21(13-14):1117–1124, 2003.
- [21] Jay J. Williams and Aggelos K. Katsaggelos. An hmm-based speech-to-video synthesizer, 1998.

# Live Interpretation of Conductors' Beat Patterns

Declan Murphy  
 Computer Science Department  
 University of Copenhagen  
 declan@diku.dk

## Abstract

A method is presented for following and interpreting the conductor's gestures known as beat patterns. Taking knowledge of the expected beat pattern and observed baton coordinates as input, the system keeps track of where the baton is in relation to the beat pattern. Tempo, dynamic level and registration updates are made on the fly according to deviations of the observed coordinates from those expected based on the score. A suitable mathematical model of the conducting process is constructed, and an algorithm for following and interpreting the execution of the beat patterns is developed.

## 1 Introduction

The work of this paper builds upon a previous computer vision technique to track an unadorned conductor's baton [6] and takes a step towards a more complete conducting system. The camera-view location of the baton's tip is available in real-time from the tracker, so the next task is to be able to follow and interpret the conductor's beat pattern as executed by the user. Deviations of the user's conducting from that expected for the music are to modulate the playing tempo, dynamics and articulation.

Rudolf [9] elucidated the structure of the conductor's beat patterns and how they have a sort of grammar. The task in this sense is to parse the execution of the beat patterns—as they unfold on the fly—in order to extract the expressive information. A symbolic understanding of the conductor's gestures is attainable by correlating the tracked motion with the standard conductors' beat patterns, some of which are shown in Fig. 1.

While it ought to be possible to recognize all beat patterns (allowing for some delay) without any *a priori* expectation of which particular pattern it

should be [10, 2], this approach would be neither necessary nor realistic. In practice, both the musicians and the conductor are quite familiar with the score before a performance. They all have a fair idea of how the music ought to turn out and of what to expect from each other. It is simply inconceivable that a real conductor would issue a 3-beat pattern where a 4-beat one would have been appropriate: such a degree of inconsistency would just never work in practice. The musicians do not need to be told how many beats are in a bar! Even if they did, it would be too late by the time they got the message. It is rather *how* the particular beat pattern is executed that conveys the important information, and accordingly it is the deviation of the executed beat pattern from what would have been expected that is to be recognized—not the beat pattern *per se*. It is therefore reasonable to annotate a score file with timing, dynamic and articulation indications—just as a real musician's score is—rather than using only raw MIDI values.

The conductor has a conception of an ideal performance interpretation, according to which s/he instructs the players to keep as close as possible to. This is done by issuing instructions to compensate for any drift from this ideal, in order that the performance should tend to return to a presumed state of perfection. For example, if the conductor considers that the musicians are playing too quickly, then a beat pattern will be executed more slowly than would otherwise have been expected. It is important to understand that the role of the conductor is not to *play* the orchestra as such, but instead to modulate their playing towards a coherent ideal interpretation. There are no absolute rules or values to be measured: it is instead a question of establishing expectation and recognizing subsequent deviation from it. (On a cognitive science note as discussed in [5, Ch.6], it may be argued at length that this process—of establishing and maintaining expectation while

deviating from it — is what music is all about. It seems appropriate then that the meta-process of modulating performance should also be of this nature.)

The model then involves a representation of the score with beat pattern annotations, and a system for recognising the deviations from the expected tempo and volume as indicated by the sampled baton locations. These deviation data can then be used to modify live generated output.

## 2 Nature of the Task

Even though the general form of the beat pattern may be known beforehand, the actual performed beat pattern will in general deviate from this curve's exact shape, size, registration and rotation. Some of this deviation is perfectly natural and should not have any bearing on the output, but of course the whole point is that certain other deviations are significant. It becomes important to be able to recognize when a drift from the abstract pattern is just the natural variation of this particular instance of it and when it is a message to alter musical playback. The difficulty of this recognition is compounded by the fact that this recognition must take place in real-time and that the beat-pattern interpreter only receives periodic updates of the tip's position — just isolated sample points — with which to determine a great deal of information on the fly. To make matters even worse, the score values for the nominal expected size and speed of the beat patterns change, e.g., as the music gets softer and slows down. In practice, the beat pattern follower can only rely on receiving sample point updates at the relatively coarse rate of 15 per second. Figure 2 illustrates how this coarseness introduces temporal uncertainty.

The simulation and recognition of cues will not be attempted here. It may safely be expected that a programmed computer will respond without fail at the appropriate time: it will not forget or get lost just because the previous thirty-seven bars have been rests, for example. The expressive timing of a voice's entry is relevant though, and this may be dealt with as a special case of articulation.

### 2.1 Considerations

Fifteen samples per second would not be good enough at all for satisfactory temporal resolution if the baton was to play the music. The resolution of MIDI Time Code (MTC) is 1–2 ms, for example.

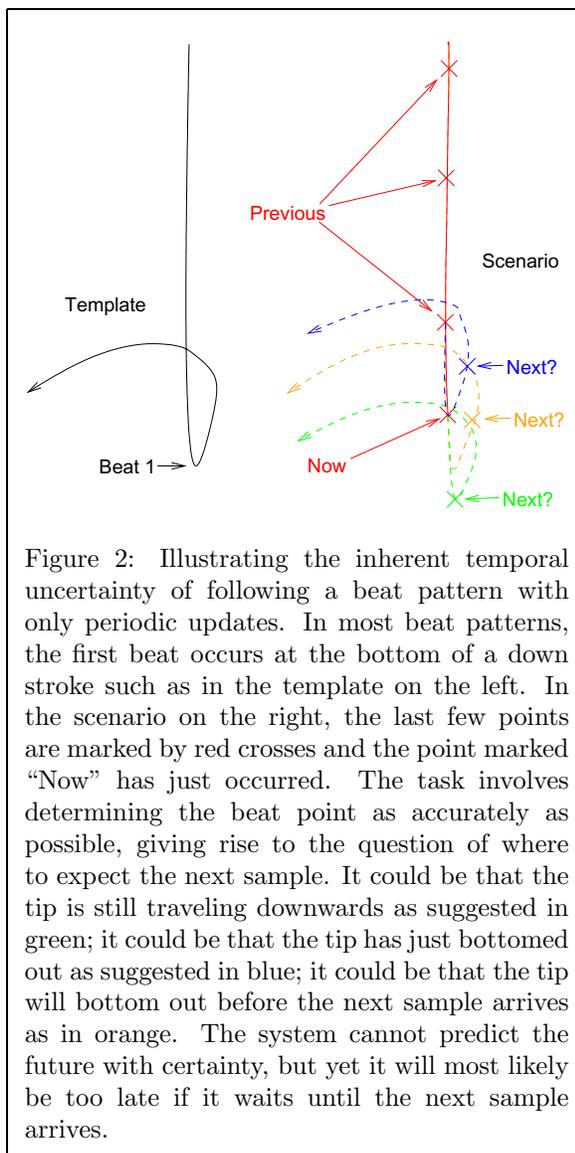
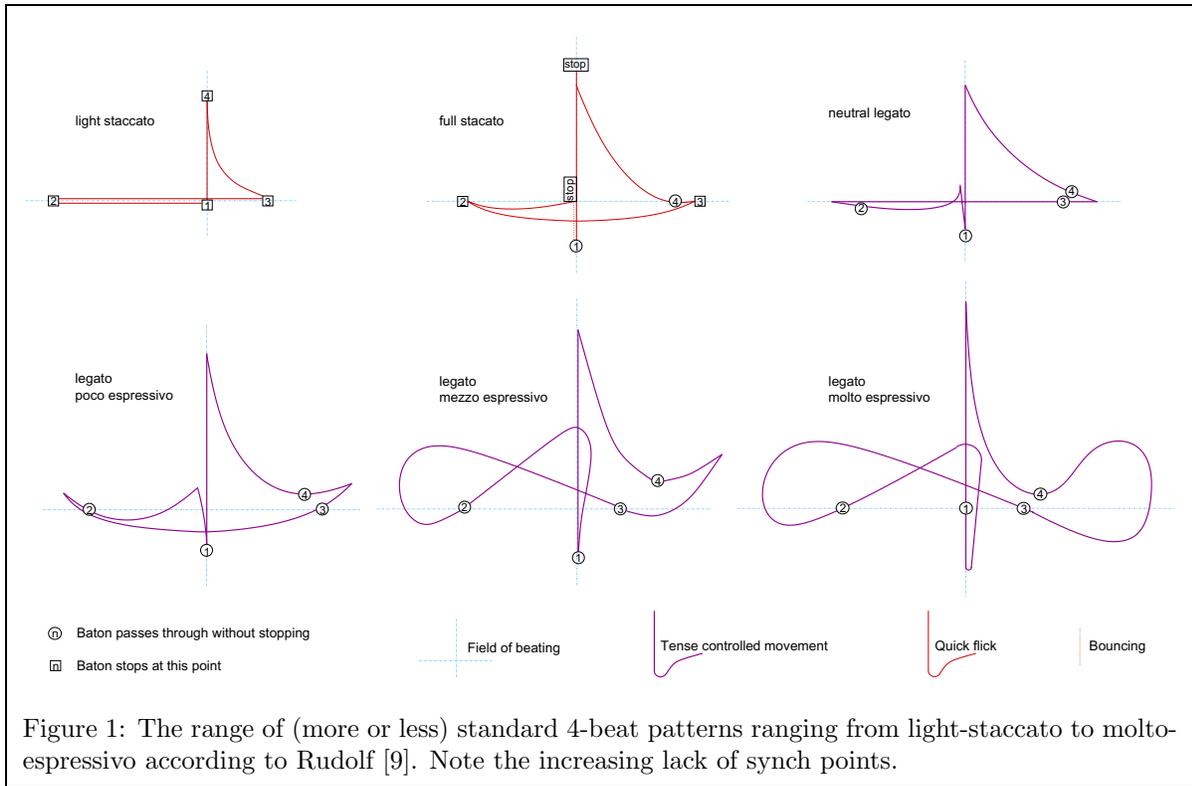


Figure 2: Illustrating the inherent temporal uncertainty of following a beat pattern with only periodic updates. In most beat patterns, the first beat occurs at the bottom of a down stroke such as in the template on the left. In the scenario on the right, the last few points are marked by red crosses and the point marked “Now” has just occurred. The task involves determining the beat point as accurately as possible, giving rise to the question of where to expect the next sample. It could be that the tip is still traveling downwards as suggested in green; it could be that the tip has just bottomed out as suggested in blue; it could be that the tip will bottom out before the next sample arrives as in orange. The system cannot predict the future with certainty, but yet it will most likely be too late if it waits until the next sample arrives.

Since conducting operates by regulating at a supra-note level, this ought to be just about good enough. If the sample rate were much lower then there would be pathological loss of information: the system would break down. There is a lower bound specified by the Nyquist sampling theorem (a corollary of Shannon's far-reaching theorem about entropy in information theory [11]), but greater resolution is required for determining the current location along the beat pattern and for simultaneously recognizing deformations of the template.

Some considerations which were instrumental in formulating the template and designing the following algorithm were:



1. The major phases of all of the beat patterns may be recognized by the tip's rough position  $\mathbf{q}$  and approximations to its first and second time derivatives  $\dot{\mathbf{q}}, \ddot{\mathbf{q}}$ . If there were on-board accelerometers or if numerical computation of derivatives was not inherently unstable ([1], [5, §A.1.2]), then this would be a preferred way to proceed. It would offer automatic accommodation of individual style and freestyle, morphing between levels of expressivity, ease of encoding and efficient computation.
2. If the spatial match is very good, then the tempo should be adjusted to match the conductor.
3. If  $\mathbf{q}$  lies significantly outside of the expected curve, signal a crescendo; if  $\mathbf{q}$  lies inside, a diminuendo.
4. Extrema may be used to decisively update scale, translation and tempo. In-between extrema, there must be anticipation in order for the system concept to work. This anticipation should be based on the expectation derived from the score file, viz. crescendo, ritardando, etc.

## 2.2 Clarification of the Task

The system is **given**:

- snapshot samples of  $\mathbf{q}(t_i) \in \mathbb{R}^2$ , the tip in flight, for some discrete time values  $t_i$ ,
- a score file annotated with indications of beat patterns, dynamics, tempo, etc.,
- a template of the beat pattern as a parametric continuous closed curve  $\mathbf{p}$  in  $\mathbb{R}^2$ .

The system must **calculate**:

- $\mathbf{q}$ 's parametric position in the template in order to determine the conductor's tempo,
- updates of the scale factor of the template in order to determine the conductor's dynamics,
- updates of the translation of the template in order to continue to match the conductor's execution with the standard template successfully and accurately.

The rotation of the template shall be assumed to be fixed. The camera can simply be rotated if necessary, but even this simple once-off calibration should not normally be necessary. People do not

normally lose their sense of up and down, so rotation matching is considered to be not worth implementing. Without rotation, any drift or translation of the conductor’s beating may be completely described by a vector  $\mathbf{b} \in \mathbb{R}^2$ .

### 2.3 Representation of the Template

The first approach considered was to encode each of the standard beat patterns (from Rudolf [9]) as normalized continuous closed parametric curves  $\mathbf{p} : [0, 1) \rightarrow \mathbb{R}^2$  constructed out of simple polynomials as splines. To this would be added a time curve  $\kappa : [0, 1) \rightarrow [0, 1)$  such that the composition  $\mathbf{p} \circ \kappa$  maps at the same rate as the gesture.

This idea of standardizing the beat patterns with simple polynomials was replaced by tailored splines constructed from recordings of the user’s execution of the beat pattern. Tailored splines were preferred for a couple of reasons. Firstly, there is a natural variation of style of execution between conductors, between text books on the subject, and between pieces with the same conductor. Secondly, although the first approach was more elegant from an analytic point of view, the recorded splines are more practical from a pragmatic engineering perspective. This approach also allows the user to pre-program special articulation, such as off-beat entry cues, with relative ease.

The convenience of considering the template’s spatial  $\mathbf{p}$  and temporal  $\kappa$  dimensions separately was retained for several reasons:

- much more smooth and accurate splines are attained in regions of low baton movement (otherwise spurious points linger in a small neighbourhood which should only have one point),
- invertability is not feasible otherwise (since some patterns require the baton to hold steady) and the guaranteed invertability of  $\kappa$  allows a more efficient implementation of locally inverting  $\mathbf{p} \circ \kappa$ ,
- direct independent scaling of size (amplitude) and speed (tempo),
- easier analytic simulation,
- easier analysis and adjustment of recorded templates.

$\mathbf{p} : [0, 1) \rightarrow \mathbb{R}^2$  is continuous, and  $\mathbf{p}$  is piecewise  $C^1$ . For each beat pattern template, there exist points  $\rho_0, \dots, \rho_{n_r-1} \in [0, 1)$  such that  $\mathbf{p}$  is not

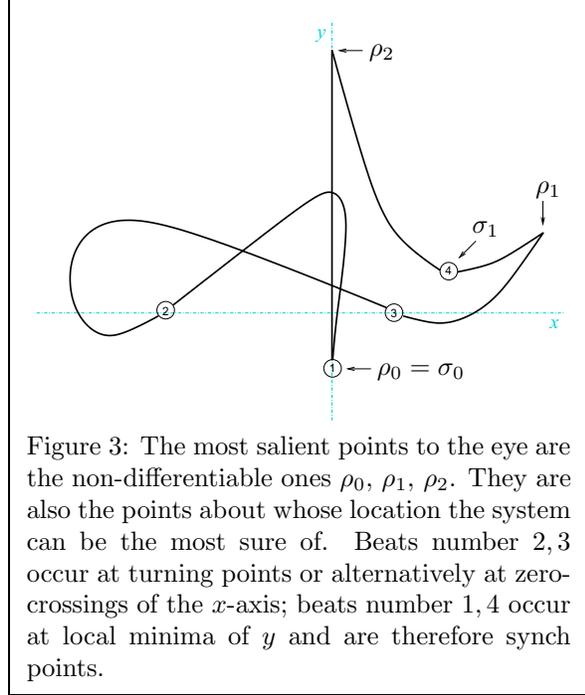


Figure 3: The most salient points to the eye are the non-differentiable ones  $\rho_0, \rho_1, \rho_2$ . They are also the points about whose location the system can be the most sure of. Beats number 2,3 occur at turning points or alternatively at zero-crossings of the  $x$ -axis; beats number 1,4 occur at local minima of  $y$  and are therefore synch points.

differentiable at  $\kappa(\rho_j)$ ,  $0 \leq j < n_r$ . These  $r = \{\rho_j\}_{j=0}^{n_r-1}$  are the most visually salient points of the beat pattern to the human eye, called “rough” points in this paper. There also exist  $\sigma_0, \dots, \sigma_{n_s-1} \in [0, 1)$  such that, for all  $\sigma_j \in s = \{\sigma_j\}_{j=0}^{n_s-1}$ , a beat point occurs at  $\mathbf{p} \circ \kappa(\sigma_j)$  and either  $\sigma_j \in r$ ,  $\dot{p}_x(\kappa(\sigma_j)) = 0$  or  $\dot{p}_y(\kappa(\sigma_j)) = 0$  where  $\mathbf{p}(t) = (p_x(t), p_y(t))$  and the dot denotes the time derivative as usual. In other words, a beat occurs at  $\mathbf{p}(\kappa(\sigma_j))$  that may be located accurately in both space and time. The  $\sigma_j \in s$  will be called “synch” points. Figure 3 shows a legato-espessivo 4-beat pattern with its rough points and synch points marked.

## 3 Model and Algorithm

### 3.1 Model

Each beat pattern has its own template  $\{\mathbf{p}, \kappa, r, s\}$  where  $r = \{\rho_0, \dots, \rho_{n_r-1}\}$  is the set of rough points and  $s = \{\sigma_0, \dots, \sigma_{n_s-1}\}$  are the synch points. The template curve  $\mathbf{p}$  is centered about the origin  $\mathbf{0} \in \mathbb{R}^2$  which corresponds to the intersection of the two axes which define the conductor’s field of beating. It is also normalized with respect to the templates for the other beat patterns to  $\mathbf{mf}$ . The template time function  $\kappa$  is both continuous and bijective on  $[0, 1)$ . It is proved in [5, §A.1.1] that  $\kappa$  must

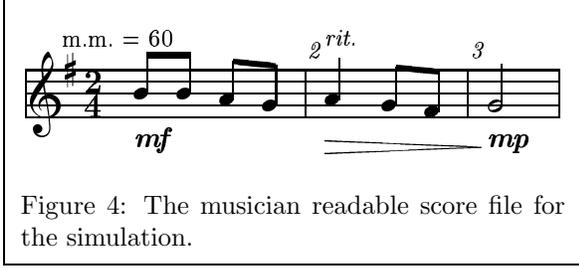


Figure 4: The musician readable score file for the simulation.

therefore be monotonic. Without loss of generality,  $\kappa(0) = 0$  and  $\lim_{t \rightarrow 1} \kappa(t) = 1$ .  $r = \{t \in [0, 1) : \mathbf{p}(\kappa(t)) \notin C^1\}$  is a finite set (by virtue of the shape of the beat patterns) so without loss of generality they can be ordered such that  $j < k \implies \rho_j < \rho_k$ ,  $0 \leq j, k < n_r$ . The set  $s$  shall be ordered similarly.

The normalized template is to be scaled and translated to fit to the incoming data, so a point  $\mathbf{p}(\kappa(t))$  for some  $\kappa(t)$  will have the general form  $a\mathbf{p} + \mathbf{b}$ .  $a \in (0, \infty)$  scales the template closed curve  $\mathbf{p}[0, 1)$ , centered about the origin, as a multiplicative scalar.  $\mathbf{b} \in \mathbb{R}^2$  translates the scaled template as an additive vector. The remaining variable in the model is the current tempo  $m \in [0, \infty)$  where one bar lasts  $m$  seconds so that, for the example in Fig. 4, if there are nominally 60 crochet beats per minute and two such beats per bar then  $m = 2$ . The expected location of  $\mathbf{q}$  at time  $t$  can now be written as

$$\mathbb{E}_t(\mathbf{q}) = a\mathbf{p}\left(\kappa\left(\frac{t \bmod m(t)}{m(t)}\right)\right) + \mathbf{b}.$$

The algorithm consists of comparing the observed value  $\mathbf{q}(t_i)$  with  $\mathbb{E}_{t_i}(\mathbf{q})|_{t_{i-1}}$  at each time step, and updating  $m, a, \mathbf{b}$  such that  $\mathbf{q}(t_i) \approx \mathbb{E}_{t_i}(\mathbf{q})|_{t_i}$ . This procedure can now be stated formally.

### 3.2 Have

The system is equipped with the following information:

- $\mathbf{p} : [0, 1) \rightarrow \mathbb{R}^2 : t \mapsto (p_x(t), p_y(t))$ ,  $\lim_{t \rightarrow 1} \mathbf{p}(t) = \mathbf{p}(0)$  piecewise differentiable, continuous, closed, parametric curve in  $\mathbb{R}^2$ .
- $\kappa : [0, 1) \rightarrow [0, 1)$ ,  $\kappa(0) = 0$  continuous, piecewise differentiable, bijective time spline.
- $r = \{\rho_0, \dots, \rho_{n_r-1}\} \subset [0, 1)$  with  $\rho_j < \rho_k$  for  $0 \leq j < k < n_r$ .  $\mathbf{p}$  is differentiable on the interval  $\kappa(\rho_j, \rho_{j+1})$  but  $\mathbf{p}$  is non-differentiable at the point  $\kappa(\rho_j)$ .  $\mathbf{p}$  is also differentiable on  $\kappa(0, \rho_0)$  and  $\kappa(\rho_{n_r-1}, 1)$ .

- $s = \{\sigma_0, \dots, \sigma_{n_s-1}\} \subset [0, 1)$  with  $\sigma_j < \sigma_k$  for  $0 \leq j < k < n_s$ .  $\forall \sigma_j \in s$ ,  $\mathbf{p}(\kappa(\sigma_j))$  is a beat point and either  $\sigma_j \in r$ ,  $\dot{p}_x(\kappa(\sigma_j)) = 0$  or  $\dot{p}_y(\kappa(\sigma_j)) = 0$ .
- $m_s(t_i), a_s(t_i) \in \mathbb{R}^+$  the nominal score values for tempo and dynamics at time  $t_i$ .
- $m_c(t_{i-1}), a_c(t_{i-1}) \in \mathbb{R}^+$ ,  $\mathbf{b}_c(t_{i-1}) \in \mathbb{R}^2$  the previous conductor's values for tempo, dynamics and offset respectively.

### 3.3 Know

The trace of hypothetical neutral conducting, following only the score, is given in the model by:

$$\mathbf{s}(t) = a_s(t)(\mathbf{p} \circ \kappa)\left(\frac{t \bmod m_s(t)}{m_s(t)}\right).$$

The expected location of the conductor's trace at time  $t_i$  is based on the most recent values of the tempo, dynamics and offset according to the conductor,  $m_c(t_{i-1}), a_c(t_{i-1}), \mathbf{b}_c(t_{i-1})$ . This information is combined with the values from the score for the current time, to effect any changes. Let

$$\bar{a}(t_i) = \frac{a_c(t_{i-1})a_s(t_i)}{a_s(t_{i-1})}, \quad \bar{m}(t_i) = \frac{m_c(t_{i-1})m_s(t_i)}{m_s(t_{i-1})}. \quad (1)$$

Then the locus of  $\mathbf{q}(\tau)$  as calculated at  $t_i$  is modelled as

$$\bar{\mathbf{c}}_{t_i}(\tau) = \bar{a}(t_i)(\mathbf{p} \circ \kappa)\left(\frac{\tau \bmod \bar{m}(t_i)}{\bar{m}(t_i)}\right) + \mathbf{b}_c(t_{i-1}), \quad (2)$$

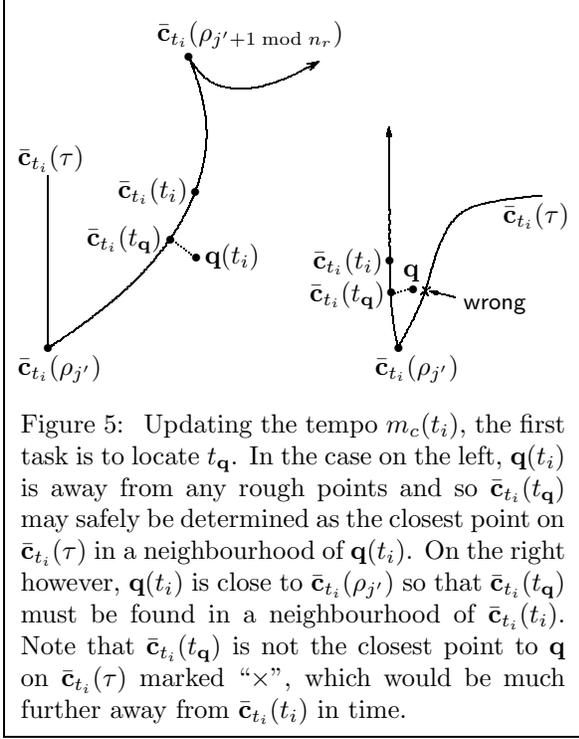
and the best estimation of  $\mathbf{q}$  at  $t_i$  is given by  $\bar{\mathbf{c}}_{t_i}(t_i)$ .

### 3.4 Want

For every time step  $t_i$ , the algorithm has to calculate the most plausible values for  $m_c(t_i), a_c(t_i), \mathbf{b}_c(t_i)$  such that, with these new values,  $\mathbf{q}(t_i)$  lies on (or very nearly on) the modelled conductor's trace. This may be stated by  $\mathbf{c}_{t_i}(t_i) \approx \mathbf{q}(t_i)$  where

$$\mathbf{c}_{t_i}(\tau) = a_c(t_i)(\mathbf{p} \circ \kappa)\left(\frac{\tau \bmod m_c(t_i)}{m_c(t_i)}\right) + \mathbf{b}_c(t_i). \quad (3)$$

In other words,  $\bar{\mathbf{c}}_{t_i}(t_i)$  are the expected points using the latest information from the last frame and the current nominal score values, and  $\mathbf{c}_{t_i}(t_i)$  are the adjusted points.



### 3.5 First Step: Try to Align Tempo

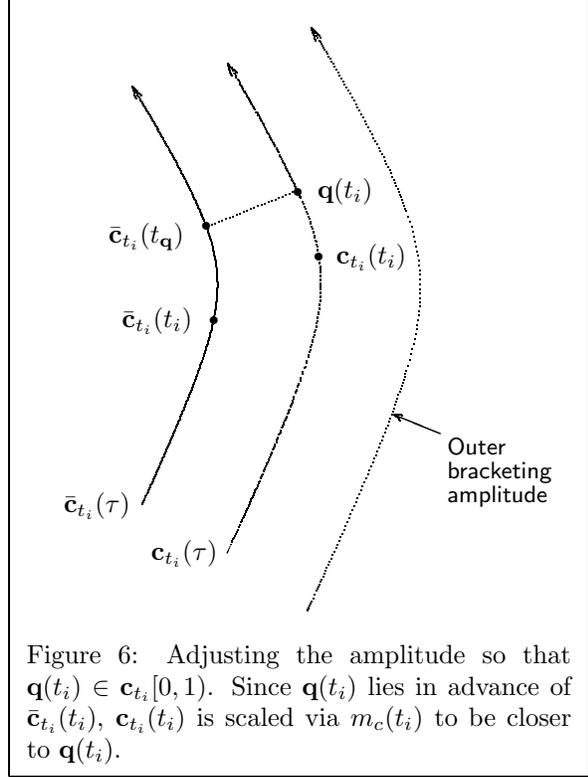
In the simple case, only the tempo needs to be updated slightly. The first task is to find  $t_q$  such that, with a certain interval  $(a, b)$ ,

$$\|\mathbf{q}(t_i) - \bar{\mathbf{c}}_{t_i}(t_q)\|_2 = \min_{a \leq \tau \leq b} \|\mathbf{q}(t_i) - \bar{\mathbf{c}}_{t_i}(\tau)\|_2. \quad (4)$$

If  $\mathbf{q}$  is not in the vicinity of a rough point, i.e.,  $\forall \rho \in r$ ,  $\|\mathbf{q}(t_i) - \bar{\mathbf{c}}_{t_i}(\rho)\|_2 > ad_r$  for an implementation dependent constant  $d_r \in \mathbb{R}^+$ , then  $(a, b)$  might just as well be  $[0, 1)$  except for numerical efficiency. If say  $\|\mathbf{q}(t_i) - \bar{\mathbf{c}}_{t_i}(\rho_j)\|_2 \leq ad_r$  for some  $0 \leq j < n_r$  then let  $\hat{t}_i = (t_i \bmod \bar{m}(t_i)) / \bar{m}(t_i)$  and choose  $j' \in \{j, j-1 \bmod n_r\}$  such that  $\rho_{j'} < \hat{t}_i < \rho_{j'+1 \bmod n_r}$ , where the choice of  $j'$  depends on whether or not  $\bar{\mathbf{c}}_{t_i}(t_i)$  has passed  $\bar{\mathbf{c}}_{t_i}(\rho_j)$  yet in this cycle. Now  $a, b$  are chosen such that  $(a, b) \subset (\rho_{j'}, \rho_{j'+1 \bmod n_r})$  as illustrated in Fig. 5.

### 3.6 Basic Case: Tempo Only

Now that  $t_q$  has been calculated, the remaining task is to adjust  $m_c(t_i)$  so as to tend to align  $t_q$  and  $t_i$ . The distance in time which has to be aligned is factored over the length of a beat at tempo  $\bar{m}(t_i)$ . Let  $l$  be the denominator of the time signature for



the current bar, and as before let

$$\hat{t}_i = \frac{t_i \bmod \bar{m}(t_i)}{\bar{m}(t_i)}. \quad (5a)$$

Then the updated tempo is

$$m_c(t_i) = \frac{\bar{m}(t_i)}{\bar{m}(t_i) + (t_q - \hat{t}_i)l}. \quad (5b)$$

### 3.7 Next Case: Amplitude and Tempo

The criterion for needing to adjust the amplitude  $a_c(t_i)$  is that  $\mathbf{q}$  is too far away from the expected locus, formally that  $\|\mathbf{q}(t_i) - \bar{\mathbf{c}}_{t_i}(t_q)\|_2 \geq d_a \log(\bar{a}(t_i) + 1)$  for an implementation dependent constant  $d_a \in \mathbb{R}^+$ . If this is not the case then the amplitude is simply updated as expected by putting  $a_c(t_i) = \bar{a}_c(t_i)$ .

The task is to adjust  $a_c(t_i)$  by the least amount such that  $\mathbf{q}(t_i) \in \mathbf{c}_{t_i}[0, 1)$ . This can be achieved without undue inefficiency by the technique known as *bracketing* and *bisecting*. Figure 6 sketches the scenario. The prior adjustment of  $m_c(t_i)$  carries through to the new scale and does not need to be recalculated.

### 3.8 Full Case: Offset Adjustment

When the conductor's time passes through a rough point  $\rho \in r$  then any necessary adjustment of the registration  $\mathbf{b}$  is made. Formally the condition is  $\mathbf{c}_{t_i}(\rho) \in \mathbf{c}_{t_i}(t_i, t_{i+1})$  for some  $\rho \in r$ , but the implementation should perform the test in  $[0, 1]$  rather than along a curve in  $\mathbb{R}^2$  for reasons of both accuracy and efficiency. Let

$$\widehat{t}_i = \frac{t_i \bmod m_c(t_i)}{m_c(t_i)} \quad \text{and} \quad \widehat{t}_{i+1} = \frac{t_{i+1} \bmod m_c(t_i)}{m_c(t_i)}. \quad (6a)$$

Then the condition is that

$$\exists \rho \in r \text{ such that } \rho \in \begin{cases} (\widehat{t}_i, \widehat{t}_{i+1}) & \widehat{t}_i < \widehat{t}_{i+1}, \\ (\widehat{t}_i, 1) \cup [0, \widehat{t}_{i+1}) & \widehat{t}_i > \widehat{t}_{i+1}. \end{cases} \quad (6b)$$

The hypothesis here is that rough points will always be correctly placed spatially. Say that  $\rho_j \in r$  satisfies the above condition (Eqn 6b), and let  $\rho_{j'} = \rho_{j-1 \bmod n_r}$ . First the amplitude  $a_c(t_{i+1} \simeq t_{\rho_j})$  is set such as to satisfy

$$\begin{aligned} & \left\| a_c(t_{\rho_j}) \mathbf{p}(\kappa(\rho_j)) - a_c(t_{\rho_{j'}}) \mathbf{p}(\kappa(\rho_{j'})) \right\|_2 \\ & = \left\| \mathbf{q}(t_{\rho_j}) - \mathbf{q}(t_{\rho_{j'}}) \right\|_2, \end{aligned} \quad (7)$$

which reduces to a somewhat messy but readily solvable quadratic in  $a_c(t_{\rho_j})$ . Any necessary updating of the score amplitude value  $a_s$  is propagated through the time steps from  $t_{\rho_{j'}}$  to  $t_{\rho_j}$  by  $\bar{a}$  in Eqns 1, 2, 3 and is implicit in Eqn 7 as such. It only remains to set  $\mathbf{b}$  which is given simply by rearranging  $a\mathbf{p} + \mathbf{b} = \mathbf{q}$ :

$$\mathbf{b}_c(t_{i+1} \simeq t_{\rho_j}) = \mathbf{q}(t_{\rho_j}) - a_c(t_{\rho_j}) \mathbf{p}(\kappa(\rho_j)). \quad (8)$$

In particular, this procedure ensures that the  $x$ -component of  $\mathbf{c}_i(\tau)$  is set on the upbeat-downstroke of each beat pattern.

The case for the conductor's time passing through a synch point is treated analogously except that the hypothesis only assumes reliable spatialization along the axis that has the extremum. Say that  $\sigma_j \in s$  satisfies the above condition in Eqn 6b and that  $\dot{p}_y(\kappa(\sigma_j)) = 0$ . Let

$$s_y = \left\{ \sigma \in s : \frac{dp_y}{dt}(\kappa(\sigma)) = 0 \right\}$$

and let

$$\rho_{j'} = \begin{cases} \max\{\rho \in r \cup s_y : \rho < \sigma_j\} & \text{if it exists,} \\ \max\{r \cup s_y\} & \text{otherwise.} \end{cases}$$

Then the equation for calculating the amplitude  $a_c(t_{\sigma_j})$  becomes

$$\begin{aligned} & \left| a_c(t_{\sigma_j}) p_y(\kappa(\sigma_j)) - a_c(t_{\rho_{j'}}) p_y(\kappa(\rho_{j'})) \right| \\ & = \left| q_y(t_{\sigma_j}) - q_y(t_{\rho_{j'}}) \right| \end{aligned} \quad (9)$$

where  $\mathbf{q}(t) = (q_x(t), q_y(t))$ , instead of Eqn 7.

## 4 Implementation

The composition function  $\mathbf{p} \circ \kappa$  needs to be inverted at selected points in a given neighbourhood in  $\mathbb{R}^2$  in order to determine  $t_{\mathbf{q}}$  from Eqn 4. It is actually a form of bracketing and bisection which was also used for determining  $a_c(t_i)$  in §3.7. Because of the low image resolution used, it takes only a small few iterations to establish the result in the implementation. For most beat patterns,  $\mathbf{p}$  is not a simple curve, i.e., it crosses itself and is not injective, but it is injective on the interval  $(a, b)$  of Eqn 4. The time spline  $\kappa$  does not need to be inverted, only the spatial spline  $\mathbf{p}$ . By setting constants  $\kappa_a, \kappa_b \in [0, 1]$  so that

$$\kappa_a = \kappa \left( \frac{a \bmod \bar{m}(t_i)}{\bar{m}(t_i)} \right), \quad \kappa_b = \kappa \left( \frac{b \bmod \bar{m}(t_i)}{\bar{m}(t_i)} \right),$$

then only half the amount of computation is required in order to calculate  $\mathbf{c}_{t_i}$  between  $\kappa_a$  and  $\kappa_b$  in Eqn 2.

The model and algorithm were implemented in C in order to maximize real-time performance by minimizing the overhead of processing time. C was also chosen for portability reasons: this code was developed in particular to be directly incorporated into the *PatternPlay* system [5, 4] or to be encapsulated as an EyesWeb module [12], both systems being platforms for investigating the use of gestural expression for making music. The code is open source and could provide the basis for many other conducting or gesture recognition systems. It is available from [7] and on the CD-ROM accompanying [5].

In order to make up a complete basic conducting system, the remaining system components are: a baton tracker [6], a facility to record the user's own beat patterns, a sequencer for live output and a module for imposing performance patterns onto the score file. The module for recording the user's beat patterns has to automatically locate the rough points, mark the synch points, and integrate them into a score file database. The module for imposing

performance patterns is implemented by means of a general framework for structural navigation and manipulation based on topos theory [3]. The full system is described in [5, §16.3.2].

## 4.1 Simulation

The algorithm presented in §3 was, at least in its finer details, worked out by building a simulation. The level of detail is somewhat less than a full implementation, but there should be no new problems in principle scaling up to a full system. A hypothetical beat pattern  $\mathbf{p}(t)$  was constructed simply by

$$\mathbf{p}(t) = \begin{cases} (2t - 4t^2)\mathbf{i} + 2t\mathbf{j} & 0 \leq t \leq \frac{1}{2}, \\ 2(1-t)\mathbf{j} & \frac{1}{2} \leq t < 1, \end{cases} \quad (10)$$

where  $\mathbf{i}, \mathbf{j}$  are the usual basis vectors in  $\mathbb{R}^2$ . This pattern along with some simulation data is plotted in Fig. 7. The associated time spline  $\kappa$  is plotted in Fig. 8. Figure 4 shows the musician-readable score for which this simulation was run.

Articulation could be handled by matching beat pattern attributes to performance parameters. It is not incorporated into the work of this paper as it is envisaged that the user would manually make these associations, whether they be global for an entire piece or once off for a single bar. Cues for expressive timing of voice entry could also be implemented as special cases of articulation.

## 4.2 Discussion

The recurring pattern of

$$f : (m, t) \mapsto \frac{t \bmod m(t)}{m(t)}$$

in several of the equations behaves well in that  $f$  is onto  $[0, 1)$ .  $f(m, t)$ 's points of discontinuity as  $t \rightarrow nm$  for any  $n \in \mathbb{N}$  match  $\kappa(\tau)$ 's point of discontinuity when  $\tau \rightarrow 1$ , which together result in a seamless continuous closed curve  $\mathbf{p} \circ \kappa \circ f$  in  $\mathbb{R}^2$ . As long as  $t \neq nm$  for all  $n \in \mathbb{N}$ , then

$$\frac{\partial f}{\partial m}(m, t) = -\frac{t \bmod m}{m^2}$$

which is suitable for the sizes and ratio of  $t$  and  $m$ . The fact that  $\partial f / \partial m$  depends on  $t$  is not as bad as it may sound at first: it is  $\mathbf{q}$  which is used to adjust  $m$ ,  $m$  does not adjust  $t$  as such.

It may of course happen that  $m_s, a_s$  change discontinuously depending on the score, which

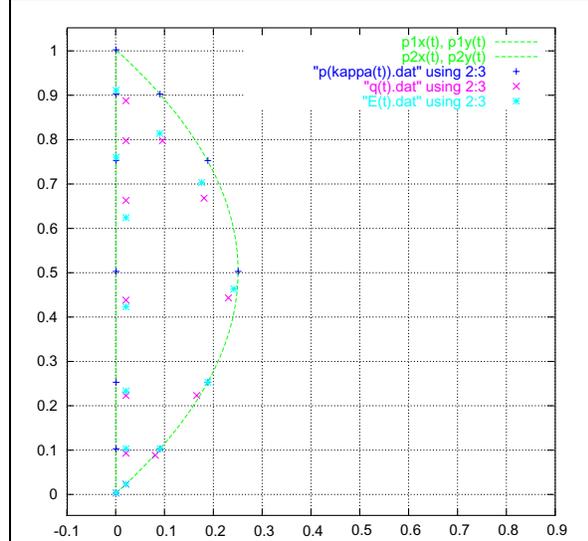


Figure 7: A plot of the simulation pseudo beat pattern as given by Eqn 10 and some early developmental fitting results.

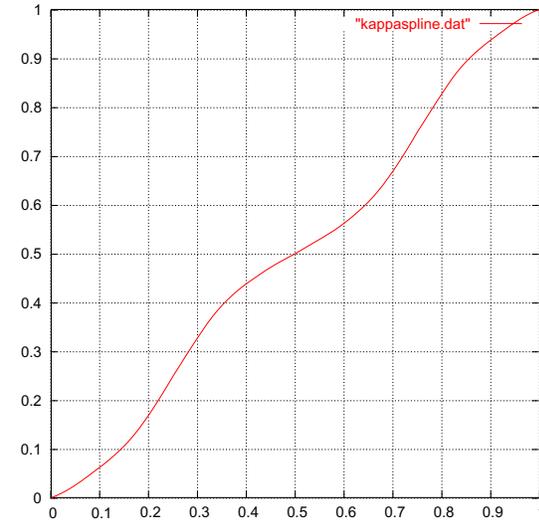


Figure 8: The time spline  $\kappa$  used for the simulation. This is a plot of actual spline interpolation values calculated by the implementation.

will suddenly change the systems behaviour for constant user input. However, as well as the conductor modulating the music output, there is an additional feedback loop consisting of the generated music telling the conductor what is going on with the performance: the conductor knows the instant there should be any such change and behaves accordingly. In this way, even discontinuities arising from the score should not pose any problem

to an alert conductor.

The method presented is one approach to the problem, but other strategies are quite possible. As mentioned at the outset, there is a certain unavoidable uncertainty inherent in the nature of the task. The approach presented here is chosen as the one considered to be the most promising of success and musically meaningful, at least until further user feedback becomes available.

In particular, if reliable acceleration information was available (discussed in [5, §A.1.2]) then a different—or indeed complimentary—method might well prove fruitful. In-between a beat pattern’s rough points, the motion is predominantly either up, down, left, right, or changing direction. The motion through the beat pattern may be parsed by knowing its approximate  $\mathbf{i}, \mathbf{j}$  velocity components. These  $\dot{p}_x, \dot{p}_y$  also give slope which may also be used as a fitting criterion. The motion can also be parsed in terms of its slope and direction.

The hypotheses used as registration criteria in §3.8 are not guaranteed to hold. It may be that they are hallmarks of good style; it may be that they vary according to a conductor’s style. Such data is not to be found in the literature. Upon reflection and without further data, they seem to be at least as strong as any rival hypothesis. Such quantities as  $a_c, |\dot{\mathbf{b}}|, |\mathbf{b}|$  and  $\dot{m}$  could be monitored and used to regulate any adjustments, but it is impossible in general to discern the user’s intention from drift or noise with certainty on the fly.

One aspect of conducting which is not modelled is how different ensembles generally require different conducting technique from the same conductor, and how the same ensemble may require (or at least receive) a different style from different conductors, by virtue of their different musical background, preferred style, social character and mutual rapport. A modelling of a certain sense of personality might give the system a more human feel, but that is left for future research.

The lack of any absolute reference frame or absolute measurable values for free gestures poses a problem if the task should become *playing* of music. The degree of articulateness necessary for skillful playing of a musical instrument seems to require both haptic and aural feedback. In terms of using gestures captured by computer vision to control music, this problem has been circumvented by focusing on expectation instead. Indeed the very lack of preciseness of the beat points of the more expressive and rubato beat patterns is presumably quite deliberate: to allow rubato. *Rubato* means

“stolen” in Italian, the semantics being that there should be local expressive timing deviations over a steady underlying pulse. One can steal a moment from the start of a phrase boundary in order to stall just before it, for example. Figure 1 shows a range of 4-beat patterns increasing in both expressivity and rubato.

### 4.3 Filtering

One of the problems encountered with previous related work on conducting audio files [8] was that sudden changes in audio playback rate are distracting to the listener whereas slow changes seem unresponsive to the user. Ideally one would like a system which can overlook the human lack of precision from one beat to the next but yet respond instantly if, e.g., the user suddenly stops. A routine for calculating the magnitude of a local deviation was envisaged. This magnitude could be used to reduce or stop the low pass filtering of the beat timing.

These issues turn out to be less problematic at the note level. The adjustment of the timing takes place in relation to the length of the last beat (Eqn 5b). This means that the music playback can comfortably come to an unexpected complete halt in one beat, a response time which rivals an alert musician. The distracting attributes arising from audio manipulation such as frequency shift, disproportionate note attack, time scaling irrespective of note length, etc. are not a problem when the playback is *generated*. When it is only the sustain part of a note which is scaled then the time scaling is perceived as being more natural. Anyway, generating the music in response to the conductor is more realistic than post-processing a recording which already has one particular interpretation stamped on it: the audio conductor must undo a pre-existing interpretation before asserting another, instead of working directly from the score.

## 5 Summary

A method for live following and interpreting of conductors’ beat patterns is presented. Before play starts, the system has been provided with a score which is annotated with the conductor’s own expected neutral execution of the beat patterns. During play, the system receives a live stream of the baton’s coordinates. At each sample time-frame, an updated value is output of how advanced along the

current beat pattern the conductor is, which can be used to regulate sequenced playback accordingly. Values are also output for the current dynamic level and any specially marked cues.

The baton's progress through the given beat pattern, and—more importantly—any meaningful deviation from the expected beat pattern, is calculated by trying to fit each sample received to the latest beat-pattern curve. At every time step, this curve is updated in terms of tempo, dynamics and registration. This update is based upon the best estimation of these parameters as indicated by the conductor in relation to their nominal score values.

A suitable computational representation of the beat patterns is devised under consideration of the nature of the conducting task and of the variability of the movement from one part of a beat pattern to another and of the various beat patterns from each other. Independent piecewise differentiable splines are used for the curves through space and time. The alignment of the expected and indicated tempo is performed by locally inverting the composition of the two splines. Significant drift in or out of the curve is aligned by updating the amplitude accordingly. Registration adjustment is accomplished at the visually salient “rough” and “synch” points.

The method is implemented into a complete computer music conducting system under GPL source code [7].

## References

- [1] Francis Begnaud Hildebrand. *Introduction to Numerical Analysis*. McGraw-Hill, New York, Second edition, 1974.
- [2] M. Lee, G. Garnett, and D. Wessel. An Adaptive Conductor Follower. In A. Strange, editor, *Proceedings of the International Computer Music Conference*, Pages 454–455, San Francisco, USA, 1992. ICMA.
- [3] Guerino Mazzola et al. *The Topos of Music: Geometric Logic of Concepts, Theory, and Performance*. Birkhäuser Verlag, 2002. In collaboration with Stefan Göller and Stefan Müller.
- [4] Declan Murphy. Pattern Play. In Alan Smaill, editor, *Additional Proceedings of the Second International Conference on Music and Artificial Intelligence*, On-line technical report series of the Division of Informatics, University of Edinburgh, Edinburgh, Scotland, UK, September 2002. <http://dream.dai.ed.ac.uk/group/smaill/icmai/b06.pdf>.
- [5] Declan Murphy. *Expressive Manipulation of Musical Structure: From Gesture Tracking to Analysis of Music*. Ph.D thesis, Copenhagen University, 2003. In Submission.
- [6] Declan Murphy. Tracking a Conductor's Baton. In Søren Olsen, editor, *Proceedings of the 12th Danish Conference on Pattern Recognition and Image Analysis*, volume 2003/05 of *DIKU report series*, Pages 59–66, Copenhagen, Denmark, August 2003. DSAGM, HCØ Tryk.
- [7] Declan Murphy. Beat Pattern Recognition. Anonymous FTP, May 2004. [ftp://ftp.diku.dk/diku/users/declan/beat\\_ptn/](ftp://ftp.diku.dk/diku/users/declan/beat_ptn/).
- [8] Declan Murphy, Tue Haste Andersen, and Kristoffer Jensen. Conducting Audio Files via Computer Vision. In *Proceedings of the Fifth International Gesture Workshop*, LNAI, Genoa, Italy, April 2003. Springer. In press.
- [9] Max Rudolf. *The Grammar of Conducting: A Comprehensive Guide to Baton Technique and Interpretation*. Macmillan, Third edition, 1993. prepared with Michael Stern.
- [10] D. Rumelhart and J. McClelland. *Parallel Distributed Processing*. The MIT Press, 1986. Two volumes.
- [11] Claude Elwood Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 1948.
- [12] University of Genoa The InfoMus lab, DIST. The EyesWeb Project. WWW, October 2003. [http://musart.dist.unige.it/sito\\_inglese/research/r\\_current/eyesweb.html](http://musart.dist.unige.it/sito_inglese/research/r_current/eyesweb.html).

# Testing for difference between two groups of functional neuroimaging experiments

Finn Årup Nielsen<sup>\*†</sup>, Andrew C. N. Chen<sup>‡</sup>, Lars Kai Hansen<sup>†</sup>

July 5, 2004

## Abstract

We describe a meta-analytic method that tests for the difference between two groups of functional neuroimaging experiments. We use kernel density estimation in three-dimensional brain space to convert points representing focal brain activations into a voxel-based representation. We find the maximum in the subtraction between two probability densities and compare its value against a resampling distribution obtained by permuting the labels of the two groups. As such it appears as a general method for comparing the local intensity of two non-stationary spatial point processes. The method is applied on data from thermal pain studies where “hot pain” and “cold pain” form the two groups.

## 1 Introduction

Human functional neuroimaging examines the relationship between cognitive functions and brain areas with positron emission tomography (PET) or magnetic resonance imaging (MRI) brain scanners. Experiments typically investigate a specific brain function and determine its “activation” in the brain volume. This is usually done by scanning multiple subjects while they are under two different conditions (e.g., “activation” and “rest”). Statistical analysis of the scanning often employing the general linear model results in a statistical parametric image volume, that is summarized by the significant local maxima (Friston et al., 1995). These local maxima are presented in scientific articles by their 3-dimensional coordinates and, e.g., their  $z$ -score or  $p$ -value.

Before the statistical analysis the brain scans are spatially normalized to a standard brain atlas, — the so-called “Talairach atlas” (Talairach and Tournoux, 1988). This allows the 3-dimensional coordinates of the local maxima — the “Talairach coordinates” — to be compared across studies.

For meta-analysis we should use the statistical parametric image volume for optimal results. Although there are beginning to appear neuroimaging databases that contain such data, e.g., the fMRI Data Center (Van Horn et al., 2001) and Neuro-Generator (Roland et al., 2001), the image volumes are typically not available and we have to resort to the Talairach coordinates.

The access to the Talairach coordinates is made easier when they are represented in a database. Two such databases exist: The BrainMap database (Fox and Lancaster, 1994; Fox and Lancaster, 2002) and the Brede database (Nielsen, 2003). A number of

---

<sup>\*</sup>Neurobiology Research Unit, Rigshospitalet, Copenhagen

<sup>†</sup>Informatics and Mathematical Modeling, Technical University of Denmark, Lyngby

<sup>‡</sup>Center of Sensory-Motor Interaction, Aalborg University, Aalborg

studies have modeled the distribution of the Talairach coordinates in these databases, e.g., (Fox et al., 1997; Nielsen and Hansen, 2002).

If the Talairach coordinates are restricted to a specific area their distribution may be approximated with a Gaussian distribution and inference can be made with “parametric models” (Fox et al., 1997) and, e.g., the Hotelling’s  $T^2$  can be employed to test for difference between two groups of coordinates (Christoff and Gabrieli, 2000; Nielsen et al., 2004). However, in many cases the distribution of the Talairach coordinates will have several spatial modes, and therefore it has been suggested to use Gaussian mixture models (Nielsen, 2001) or kernel density estimators (Nielsen and Hansen, 2002; Turkeltaub et al., 2002; Chein et al., 2002; Wager et al., 2003).

The statistical analysis is usually performed in a mass-univariate setting where the number of statistical tests corresponds to the number of voxels in the volume. This results in a massive multiple comparison problem that is most often countered by employing random field theory for the statistical inference (Cao and Worsley, 2001). But permutation tests can also be used by constructing the null distribution for the maximum statistic, where the maximum is taken across all the voxel in the statistical parametric image (Holmes et al., 1996; Nichols and Holmes, 2001).

Below we will describe a meta-analytic method that uses permutation tests together with maximum statistic and kernel density estimation to give a statistical value for the difference between two groups of Talairach coordinates and thereby testing if two groups of functional neuroimaging experiments are different. We will use Talairach coordinates from hot and cold pain experiments. The pain modality is of special interest for our particular method since it typically causes a spatially multimodal activation pattern where several distinct brain regions are involved: Thalamus, the somatosensory cortex, insula and anterior cingulate cortex (Ingvar, 1999).

## 2 Method

A probability density volume is constructed by convolving a local maximum  $l$  (in the following called “location”) positioned in Talairach space at  $\mathbf{x}_l$  with a 3-dimensional Gaussian kernel with isotropic variance (Nielsen and Hansen, 2002; Turkeltaub et al., 2002; Chein et al., 2002).

$$p(\mathbf{x}|l) = (2\pi\sigma^2)^{-3/2} \exp \left[ -\frac{(\mathbf{x} - \mathbf{x}_l)^\top (\mathbf{x} - \mathbf{x}_l)}{2\sigma^2} \right], \quad (1)$$

where we select the kernel width to  $\sigma = 1$  cm. When we construct the probability density corresponding to a group of experiments  $p(\mathbf{x}|g)$  we combine the contributions from all the individual locations associated with the group (i.e., all the locations in all the experiments associated with the group)

$$p(\mathbf{x}|g) = \sum_{l \in g} p(\mathbf{x}|l) p(l|g), \quad (2)$$

where the prior is simply set to  $p(l|g) = 1/|L_g|$ , i.e., inversely proportional to the number of locations in the  $g$  group. The continuous probability density is converted to a vector by sampling it in a regular grid

$$\mathbf{v}_g \equiv p(\mathbf{x}|g). \quad (3)$$

In the present application we use a coarse grid of  $(8 \text{ mm})^3$  resulting in 7752 voxels. As a statistic for the difference between two volumes ( $\mathbf{v}_1$  and  $\mathbf{v}_2$ ) we simply use the subtraction performed separately for each voxel

$$\mathbf{t} = \mathbf{v}_1 - \mathbf{v}_2. \quad (4)$$

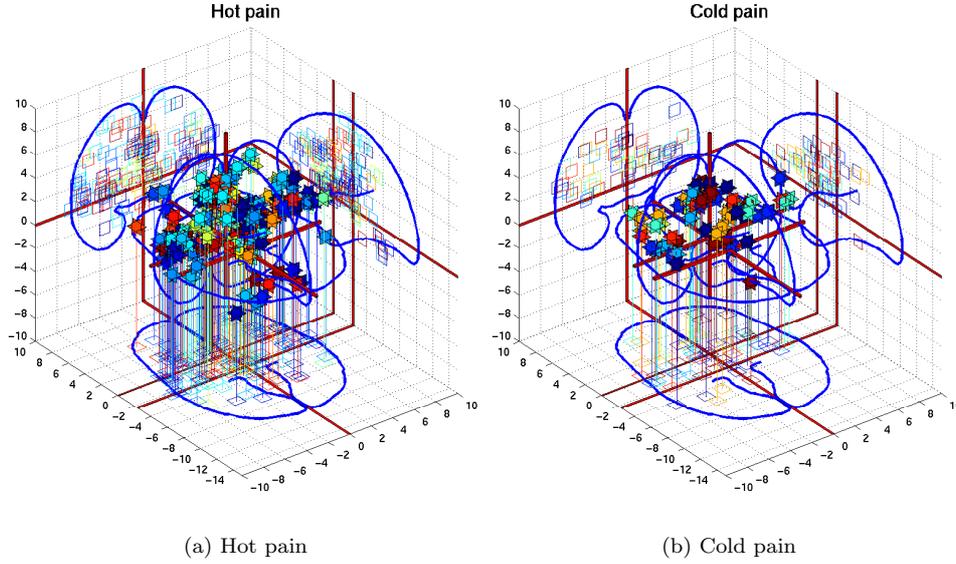


Figure 1: Visualization of the Talairach coordinates for hot pain and cold pain studies in a corner cube environment (Rehm et al., 1998). The glyphs are colored according to the experiment they belong to and are projected onto the walls. The blue curves are the outline of the brain. The thick red lines are the axes of the Talairach atlas and the view point makes the upper left part of the back of the brain the closest point. The units on the axes are centimeters.

To counter the multiple comparison problem a null distribution uses the maximum value across voxels

$$t = \max_i(t_i) \quad (5)$$

The null distribution of this maximum statistic is established by permutation: A distribution is built up by permuting the assignment of experiments to the two groups resulting in two new groups  $\mathbf{v}_1^*$  and  $\mathbf{v}_2^*$ , that each comprises of the same number of experiments as the original two groups. Thus the null distribution of the maximum statistic appears as

$$t^* = \max_i(v_{1,i}^* - v_{2,i}^*). \quad (6)$$

The permutation is randomized sufficiently many times ( $n = 1, \dots, N$ ) to generate a “smooth” and stable distribution. We may assign a conservative  $p$ -value to the  $i$ th voxel by counting the proportion its statistic  $t_i$  is exceeded by the  $N$  values of the permutation maximum statistic

$$P_i = 1/N \sum_n |t_i < t_n^*|. \quad (7)$$

The  $p$ -values enable us to choose a statistically based threshold in the subtraction image.

To demonstrate the method we invoke data from thermal pain studies, where the two groups are hot and cold pain. Such studies will typically employ a  $45 - 50^\circ\text{C}$  or  $0 - 5^\circ\text{C}$  stimulus to the subjects. The studies were added to the Brede database

1	WOEXT: 183 - Hot pain	(Tracey et al., 2000)
2	WOEXT: 186 - Attended heat pain on right hand	(Brooks et al., 2002)
3	WOEXT: 187 - Distracted heat pain on right hand	(Brooks et al., 2002)
4	WOEXT: 188 - Attended heat pain on left hand	(Brooks et al., 2002)
5	WOEXT: 189 - Distracted heat pain on left hand	(Brooks et al., 2002)
6	WOEXT: 217 - Hot pain in right hand	(Craig et al., 1996)
7	WOEXT: 225 - Hot pain on left hand (group 1)	(Becerra et al., 1999)
8	WOEXT: 227 - Hot pain on left hand (group 2)	(Becerra et al., 1999)
9	WOEXT: 230 - Painful heat on right fingers	(Gelnar et al., 1999)
10	WOEXT: 233 - Hot pain on right hand in rest, mental imagery and hypnosis	(Faymonville et al., 2000)
11	WOEXT: 234 - Hot pain on right hand in rest and mental imagery	(Faymonville et al., 2000)
12	WOEXT: 235 - Hot pain on right hand during hypnosis	(Faymonville et al., 2000)
13	WOEXT: 237 - Interaction between hypnosis and hot pain on right hand	(Faymonville et al., 2000)
14	WOEXT: 238 - Correlated with pain ratings in hot pain on right hand in rest, mental imagery and hypnosis	(Faymonville et al., 2000)
15	WOEXT: 240 - Interaction between hypnosis and pain ratings in hot pain on right hand.	(Faymonville et al., 2000)
16	WOEXT: 245 - Heat pain on right arm	(Tölle et al., 1999)
17	WOEXT: 246 - Positive correlation with pain threshold	(Tölle et al., 1999)
18	WOEXT: 248 - Correlation with pain intensity	(Tölle et al., 1999)
19	WOEXT: 249 - Correlation with pain unpleasantness	(Tölle et al., 1999)
20	WOEXT: 298 - Early phase heat pain	(Casey et al., 2001)
21	WOEXT: 299 - Late phase heat pain	(Casey et al., 2001)
22	WOEXT: 312 - Heat pain on left hand	(Vogt et al., 1996)
23	WOEXT: 314 - Heat pain on left volar forearm	(Adler et al., 1996)
24	WOEXT: 319 - Heat pain on left arm	(Casey et al., 1996)
1	WOEXT: 182 - Cold pain	(Tracey et al., 2000)
2	WOEXT: 184 - Cold pain in left hand	(Petrovic et al., 2000)
3	WOEXT: 213 - Cold pain in right hand	(Craig et al., 1996)
4	WOEXT: 263 - Cold pain on right foot	(Frankenstein et al., 2001)
5	WOEXT: 264 - Cold pain on right foot masked by silent word reading	(Frankenstein et al., 2001)
6	WOEXT: 265 - Silent word reading while cold pain on right foot	(Frankenstein et al., 2001)
7	WOEXT: 266 - Cold pain versus cold pain with silent word reading	(Frankenstein et al., 2001)
8	WOEXT: 320 - Cold pain on left arm	(Casey et al., 1996)

Table 1: List of included hot and cold pain experiments.

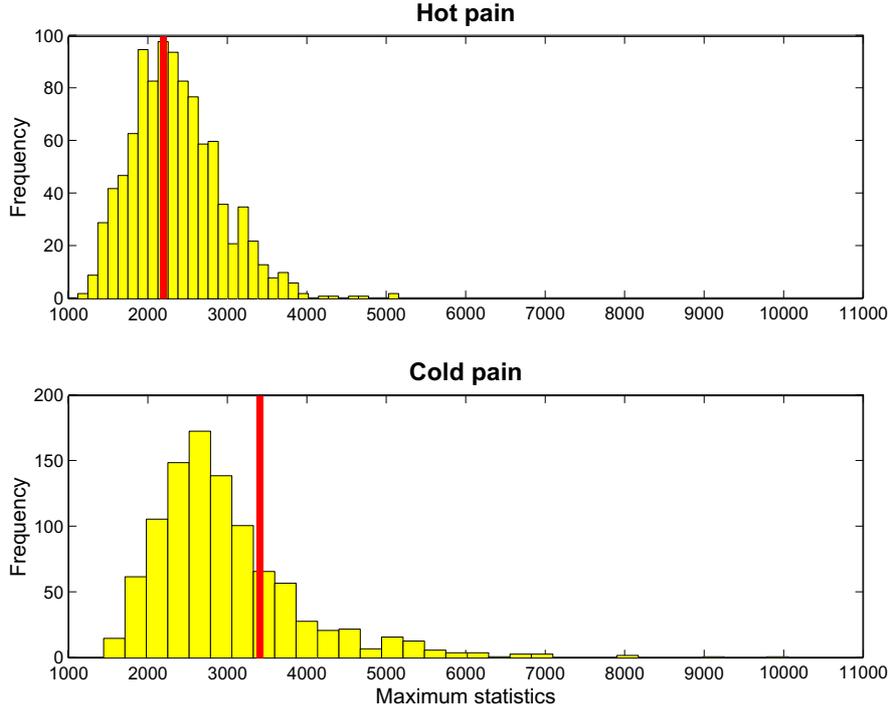


Figure 2: Empirical histograms of the maximum statistics  $t^*$  after 1000 permutations. The thick red lines indicate the maxima for the hot and cold pain statistics  $t_{\text{hot}}$  and  $t_{\text{cold}}$ .

(Nielsen, 2003). Slight variations among the studies appear in the application of the Talairach atlases, and locations that conform to the so-called MNI space are adjusted before entry (Brett, 1999). All the locations from the pain experiments are shown in two panels in figure 1, where the color indicates the experiments the locations originate from. Table 1 lists all the included 24 hot and 8 cold pain experiments. The list is automatically generated from the information in the Brede database. Note that the pain stimulus is induced under varying contexts, at different sites on the body and some studies contribute with several experiments, e.g., a study by Faymonville *et al.* contributes with 6 experiments.

Both the statistics for hot and cold pain are considered, — simply by reversing the subtraction (or actually, in the practical implementation, by finding the minimum)

$$t_{\text{hot}} = \max_i (v_{\text{hot},i} - v_{\text{cold},i}) \quad (8)$$

$$t_{\text{cold}} = \max_i (v_{\text{cold},i} - v_{\text{hot},i}). \quad (9)$$

Many of the operations performed in our method described above are implemented in the Brede Neuroinformatics Toolbox (Nielsen and Hansen, 2000).

### 3 Results and discussion

Figure 2 displays the empirical histograms of the null distributions of the maximum statistics  $t^*$ . The thick red lines indicate the maximum statistics of the comparisons of

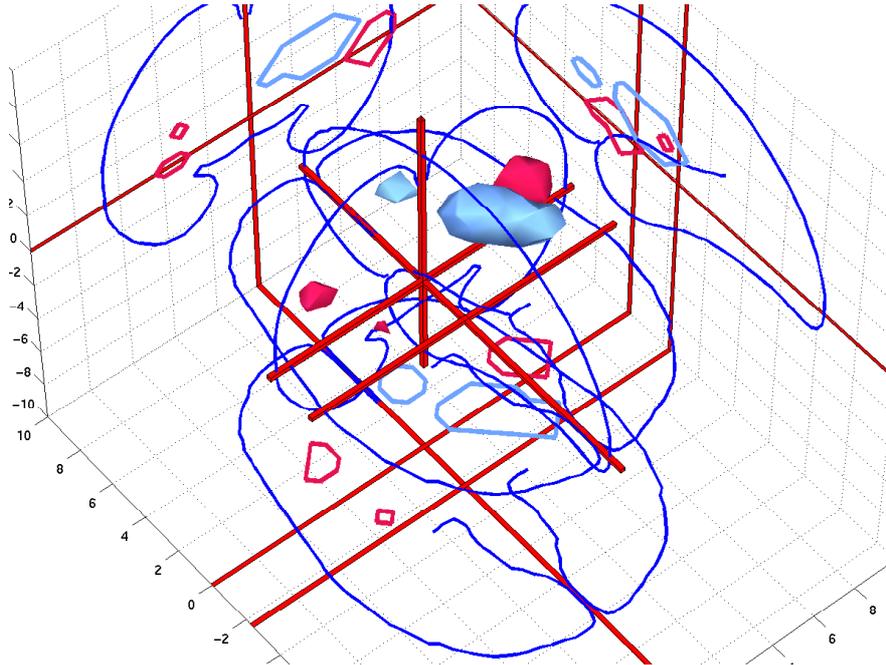


Figure 3: Results from the permutation test in a corner cube environment. The red isosurfaces are for hot pain  $t_{\text{hot}}$  and the light blue isosurfaces are for cold pain  $t_{\text{cold}}$  based on very low thresholds at  $P = 0.95$ .

interest  $t_{\text{hot}}$  and  $t_{\text{cold}}$ . They show that our method does not find any larger differences between hot and cold pain: The statistics of both hot and cold pain are found in the middle of the null distributions. The distribution for the cold pain has a heavier tail than the hot pain distribution. This is caused by the fewer experiments that form the cold pain group.

Figure 3 shows isosurfaces with a very liberal threshold in the subtraction images of hot  $t_{\text{hot}}$  and cold pain  $t_{\text{cold}}$ . The right hemisphere appears with the first difference between the pain modalities, but only with very low support: For cold pain the most “significant” voxel has a  $p$ -value of  $P \approx 0.25$ .

We have previously proposed a method that uses a database of experiments to generate a null distribution of the correlation coefficient between two volumes (Nielsen and Hansen, 2004a; Nielsen and Hansen, 2004b). That method requires a database of dissimilar experiments to build up a null hypothesis, e.g., a pain experiment is compared with a memory or language experiment. The method we present in this contribution does not need this extra data, but relies only on data from the two groups that are being compared. Furthermore, our previous method is a global method performing an omnibus test for the entire volume and the entire set of locations, while our new method allows us to make inference on the voxel-level. This is also an advantage over the Hotelling’s  $T^2$  test. However, our presented method does require — to gain sufficient statistical power — that there are many experiments in each group.

## 4 Acknowledgment

Jørgen Kold is acknowledged for collection of the data. Finn Årup Nielsen is supported by the Villum Kann Rasmussen Foundation.

## References

- Adler, L. J., Gyulai, F. E., Diehl, D. J., Mintun, M. A., Winter, P. M., and Firestone, L. L. (1996). Regional brain activity changes associated with fentanyl analgesia elucidated by positron emission tomography. *Anesthesia & Analgesia*, 84(1):120–126.
- Becerra, L. R., Breiter, H. C., Stojanovic, M., Fishman, S., Edwards, A., Comite, A. R., Gonzalez, R. G., and Borsook, D. (1999). Human brain activation under controlled thermal stimulation and habituation to noxious heat: An fMRI study. *Magnetic Resonance in Medicine*, 41(5):1044–1057.
- Brett, M. (1999). The MNI brain and the Talairach atlas. <http://www.mrc-cbu.cam.ac.uk/Imaging/mnispace.html>. Accessed 2003 March 17.
- Brooks, J. C. W., Nurmikko, T. J., Bimson, W. E., Singh, K. D., and Roberts, N. (2002). fMRI of thermal pain: effects of stimulus laterality and attention. *NeuroImage*, 15(2):293–301.
- Cao, J. and Worsley, K. J. (2001). Applications of random fields in human brain mapping. In Moore, M., editor, *Spatial Statistics: Methodological Aspects and Applications*, volume 159 of *Lecture notes in Statistics*, chapter 8, pages 170–182. Springer, New York.
- Casey, K. L., Minoshima, S., Morrow, T. J., and Koeppe, R. A. (1996). Comparison of human cerebral activation patterns during cutaneous warmth, heat pain, and deep cold pain. *Journal of Neurophysiology*, 76(1):571–581.
- Casey, K. L., Morrow, T. J., Lorenz, J., and Minoshima, S. (2001). Temporal and spatial dynamics of human forebrain activity during heat pain: analysis by positron emission tomography. *Journal of Neurophysiology*, 85(2):951–959.
- Chein, J. M., Fissell, K., Jacobs, S., and Fiez, J. A. (2002). Functional heterogeneity within Broca’s area during verbal working memory. *Physiology & Behavior*, 77(4-5):635–639.
- Christoff, K. and Gabrieli, J. D. E. (2000). The frontopolar cortex and human cognition: Evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology*, 28(2):168–186.
- Craig, A. D., Reiman, E. M., Evans, A., and Bushnell, M. C. (1996). Functional imaging of an illusion of pain. *Nature*, 384(6606):258–260.
- Faymonville, M. E., Laureys, S., Degueldre, C., Del Fiore, G., Luxen, A., Franck, G., Lamy, M., and Maquet, P. (2000). Neural mechanisms of antinociceptive effects of hypnosis. *Anesthesiology*, 92(5):1257–1267.
- Fox, P. T. and Lancaster, J. L. (1994). Neuroscience on the net. *Science*, 266(5187):994–996.

- Fox, P. T. and Lancaster, J. L. (2002). Mapping context and content: the BrainMap model. *Nature Reviews Neuroscience*, 3(4):319–321.
- Fox, P. T., Lancaster, J. L., Parsons, L. M., Xiong, J.-H., and Zamarripa, F. (1997). Functional volumes modeling: Theory and preliminary assessment. *Human Brain Mapping*, 5(4):306–311.
- Frankenstein, U. N., Richter, W., McIntyre, M. C., and Remy, F. (2001). Distraction modulates anterior cingulate gyrus activations during the cold pressor test. *NeuroImage*, 14(4):827–36.
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-B., Frith, C. D., and Frackowiak, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2:189–210.
- Gelnar, P. A., Krauss, B. R., Sheehe, P. R., Szeverenyi, N. M., and Apkarian, A. V. (1999). A comparative fMRI study of cortical representations for thermal painful, vibrotactile, and motor performance tasks. *NeuroImage*, 10(4):460–482.
- Holmes, A. P., Blair, R. C., Watson, J. D. G., and Ford, I. (1996). Non-parametric analysis of statistic images from functional mapping experiments. *Journal of Cerebral Blood Flow and Metabolism*, 16(1):7–22.
- Ingvar, M. (1999). Pain and functional imaging. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 354(1387):1347–1358.
- Nichols, T. E. and Holmes, A. P. (2001). Nonparametric permutation tests for PET functional neuroimaging experiments: A primer with examples. *Human Brain Mapping*, 15(1):1–25.
- Nielsen, F. Å. (2001). *Neuroinformatics in Functional Neuroimaging*. PhD thesis, Informatics and Mathematical Modelling, Technical University of Denmark, Lyngby, Denmark.
- Nielsen, F. Å. (2003). The Brede database: a small database for functional neuroimaging. *NeuroImage*, 19(2). Presented at the 9th International Conference on Functional Mapping of the Human Brain, June 19–22, 2003, New York, NY. Available on CD-Rom.
- Nielsen, F. Å., Balslev, D., and Hansen, L. K. (2004). Mining posterior cingulate. *NeuroImage*, 22. Presented at the 10th Annual Meeting of the Organization for Human Brain Mapping, June 14–17, 2004, Budapest, Hungary. Available on CD-ROM.
- Nielsen, F. Å. and Hansen, L. K. (2000). Experiences with Matlab and VRML in functional neuroimaging visualizations. In Klasky, S. and Thorpe, S., editors, *VDE2000 - Visualization Development Environments, Workshop Proceedings, Princeton, New Jersey, USA, April 27–28, 2000*, pages 76–81, Princeton, New Jersey. Princeton Plasma Physics Laboratory.
- Nielsen, F. Å. and Hansen, L. K. (2002). Modeling of activation data in the BrainMap<sup>TM</sup> database: Detection of outliers. *Human Brain Mapping*, 15(3):146–156.
- Nielsen, F. Å. and Hansen, L. K. (2004a). Assessing the reproducibility in sets of Talairach coordinates. *NeuroImage*, 22. Presented at the 10th Annual Meeting of the Organization for Human Brain Mapping, June 14–17, 2004, Budapest, Hungary. Available on CD-ROM.

- Nielsen, F. Å. and Hansen, L. K. (2004b). Finding related functional neuroimaging volumes. *Artificial Intelligence in Medicine*, 30(2):141–151.
- Petrovic, P., Petersson, K. M., Ghatan, P. H., Stone-Elander, S., and Ingvar, M. (2000). Pain-related cerebral activation is altered by a distracting cognitive task. *Pain*, 85(1-2):19–30.
- Rehm, K., Lakshminarayan, K., Frutiger, S. A., Schaper, K. A., Sumners, D. L., Strother, S. C., Anderson, J. R., and Rottenberg, D. A. (1998). A symbolic environment for visualizing activated foci in functional neuroimaging datasets. *Medical Image Analysis*, 2(3):215–226.
- Roland, P., Svensson, G., Lindeberg, T., Risch, T., Baumann, P., Dehmel, A., Fredriksson, J., Halldorson, H., Forsberg, L., Young, J., and Zilles, K. (2001). A database generator for human brain imaging. *Trends in Neuroscience*, 24(10):562–564.
- Talairach, J. and Tournoux, P. (1988). *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme Medical Publisher Inc, New York.
- Tölle, T. R., Kaufmann, T., Siessmeier, T., Lautenbacher, S., Berthele, A., Munz, F., Zieglgansberger, W., Willoch, F., Schwaiger, M., Conrad, B., and Bartenstein, P. (1999). Region-specific encoding of sensory and affective components of pain in the human brain: a positron emission tomography correlation analysis. *Ann Neurol*, 45(1):40–7.
- Tracey, I., Becerra, L., Chang, I., Breiter, H., Jenkins, L., Borsook, D., and Gonzalez, R. G. (2000). Noxious hot and cold stimulation produce common patterns of brain activation in humans: a functional magnetic resonance imaging study. *Neuroscience Letters*, 288(2):159–162.
- Turkeltaub, P. E., Eden, G. F., Jones, K. M., and Zeffiro, T. A. (2002). Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *NeuroImage*, 16(3 part 1):765–780.
- Van Horn, J. D., Grethe, J. S., Kostelec, P., Woodward, J. B., Aslam, J. A., Rus, D., Rockmore, D., and Gazzaniga, M. S. (2001). The functional magnetic resonance imaging data center (fMRIDC): the challenges and rewards of large-scale databasing of neuroimaging studies. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 356(1412):1323–1339.
- Vogt, B. A., Derbyshire, S., and Jones, A. K. P. (1996). Pain processing in four regions of human cingulate cortex localized with co-registered PET and MR imaging. *European Journal of Neuroscience*, 8(7):1461–1473.
- Wager, T. D., Phan, K. L., Liberzon, I., and Taylor, S. F. (2003). Valence, gender, and lateralization of functional brain anatomy in emotion: a meta-analysis of findings from neuroimaging. *NeuroImage*, 19(3):513–531.

# AN IMAGE BASED SYSTEM TO AUTOMATICALLY AND OBJECTIVELY SCORE THE DEGREE OF REDNESS AND SCALING IN PSORIASIS LESIONS.

David Delgado\*  
Informatics and Mathematical Modelling,  
Denmark  
email: ddg@imm.dtu.dk

Bjarne Ersbøll  
IMM,  
Denmark  
email: be@imm.dtu.dk

Jens Michael Carstensen  
IMM  
Denmark  
email: jmc@imm.dtu.dk

## Abstract

In this work, a combined statistical and image analysis method to automatically evaluate the severity of scaling in psoriasis lesions is proposed. The method separates the different regions of the disease in the image and scores the degree of scaling based on the properties of these areas. The proposed method provides a solution to one of the present problems in dermatology: the lack of suitable methods to assess the lesion and to evaluate the changes during the treatment. An experiment over a collection of psoriasis images is conducted to test the performance of the method. Results show that the obtained scores are highly correlated with scores made by doctors. This and the fact that the obtained measures are continuous indicate the proposed method is a suitable tool to evaluate the lesion and to track the evolution of dermatological diseases.

**Keywords:** *psoriasis, exploratory data analysis, segmentation, decision trees, classification*

## 1 Introduction

One of the main problems in the treatment of dermatological diseases is the difficulty of tracking the evolution of the disease. Physicians are visited by the patients several times to control the evolution of the disease. However, due to the fact that no objective methods to summarize the lesion exist, physicians make scorings and take notes to document the actual condition of the patient. A drawback of this method is the dependency on the individual physician.

The advances in image analysis during the last decade have led to the development of different methods to deal with related problems in the dermatological field. Engström [1] observed the effect of a new enzymatic debrider observing the evolution of the lesion area and the lesion color. These measurements were obtained from digitized photographs analyzed with a computer. Later, Hansen [2], developed an image system that included calibration for increasing the quality of the images. The system diagnoses burns and pressure ulcers in animals but the possibility of being used in humans was mentioned. In a recent paper, Hillebrand [3] used computer analysis in high resolution digital images to compare the skin condition of a group of females.

In this work, a method to objectively score the degree of scaling and redness in psoriasis is proposed. The method realises a hierarchical segmentation to isolate the different structures present in the image: normal skin, red area and scales. Different values are obtained from these areas and they are used to approximate the doctor scorings.

---

\*The dermatologists Lone Skov and Bo Bang of Gentofte Hospital of Denmark and the anonymous patients are gratefully acknowledged for their collaboration during the image acquisition sessions.

## 2 Segmentation of the areas present in the lesion

Psoriasis is a dermatological disease characterized by red, thickened areas with silvery scales [4]. In order to score the degree of scales and redness in psoriasis, the first step is to segment the different areas in the lesion. The wide variety of forms and different levels of severity that psoriasis can exhibit makes this task highly complex.

### 2.1 Segmentation of the lesion

The segmentation of the disease with respect to the healthy area is based on the assumption, that under a suitable projection, both the normal skin and the lesion are distributed approximately as a Gaussian distribution. This assumption was supported by an exploratory data analysis of a small set of psoriasis images where several projections were considered. Furthermore, a principal component analysis [5] and an independent component analysis [6] on a dataset of 115 images indicated that the difference between the green and the blue band exhibits a good contrast to discriminate between the lesion and the normal skin. The distribution of this difference approximately follows a linear mixture of two Gaussians. The estimation of their means and variances makes it possible to identify the lesion by means of discriminant analysis. The parameters of the Gaussians were estimated according to Taxt [7]. Figure 1 shows the segmentation of the lesion.

### 2.2 Extracting the scales

Segmentation of the scales is complicated by the fact that scales may or may not appear in the image. If they appear they may range from a few spots to a large area. Moreover, non-uniformity of the areas with redness (ranging from red to brown) makes the task even harder. This variability implies that the lesion has to be considered in small areas where the change in redness is not significant. This can be accomplished with watersheds [8] to mark the different scales and then locally use a clustering algorithm [9] to segment them. This approach requires specifying the number of watersheds. In this work, the number of watersheds is determined in two steps. First a new image is created based in the watershed regions. Each watershed area is replaced by the minimum value of this area. This new image is then thresholded and the watersheds with values less than the threshold are the areas where the scales are detected. The method was tested on a set of psoriasis images and it demonstrated a good performance. However, the method had difficulties with some images that had problems during acquisition (especially shadows), so the number of watersheds was not found correctly. To solve this problem, the number of watersheds was fixed visually by a tuning parameter. The blue band was used to find the watersheds because a canonical analysis had shown that this band is the best to separate the scales from the red area. Figure 2 displays the segmentation of the scales.

### 2.3 Scoring the disease

Once the different areas have been segmented, a decision tree is created to automatically score the degree of scaling in the different images, approximating the scorings made by the physicians. Three variables are used as input to the model: the area of the scaling, the ratio between the area of scaling and the area of the lesion, and the ratio between the area of scaling and the area of redness. The whole procedure is shown in Figure 3.

In the evaluation of the redness, a canonical discriminant analysis over the difference of the mean values of the spectral values in the red and healthy area points out to approximate the physicians scorings using a clustering method as, e. g., a K-nearest neighbour.

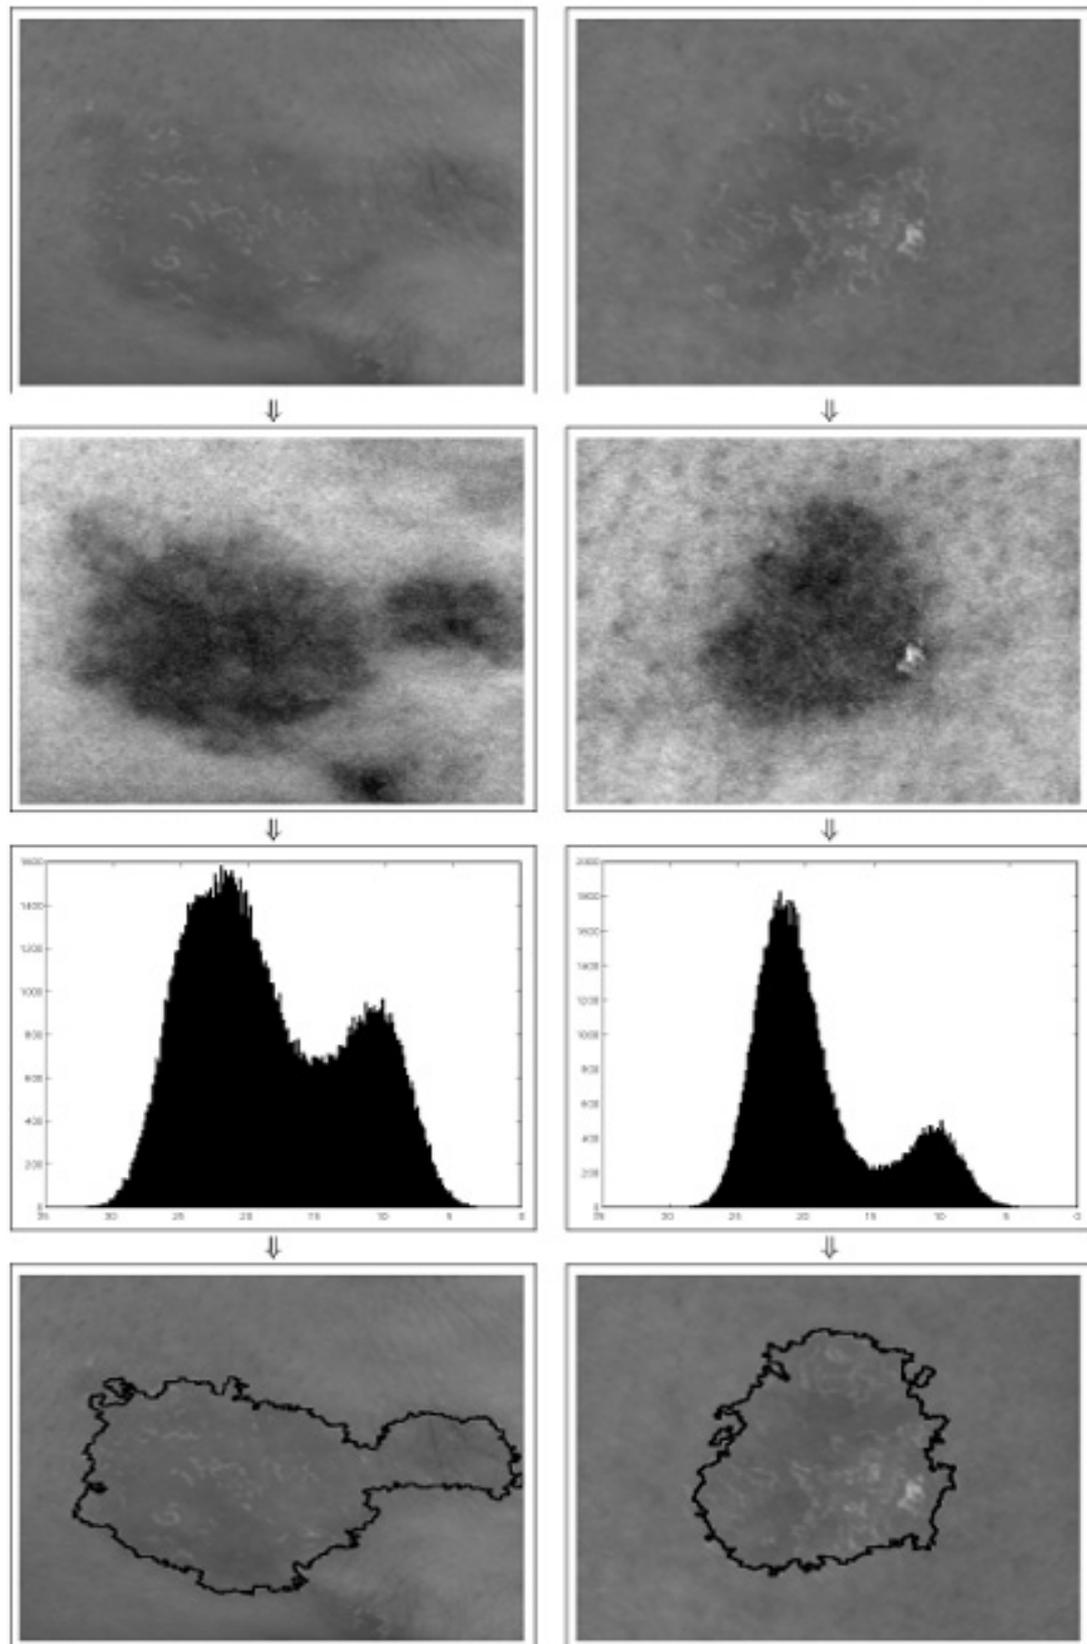


Figure 1: Top row: Two psoriasis images. Second row: Difference between the blue and green bands. Third row: Histogram of the band difference (blue minus green). Bottom row: Lesion segmentation result.

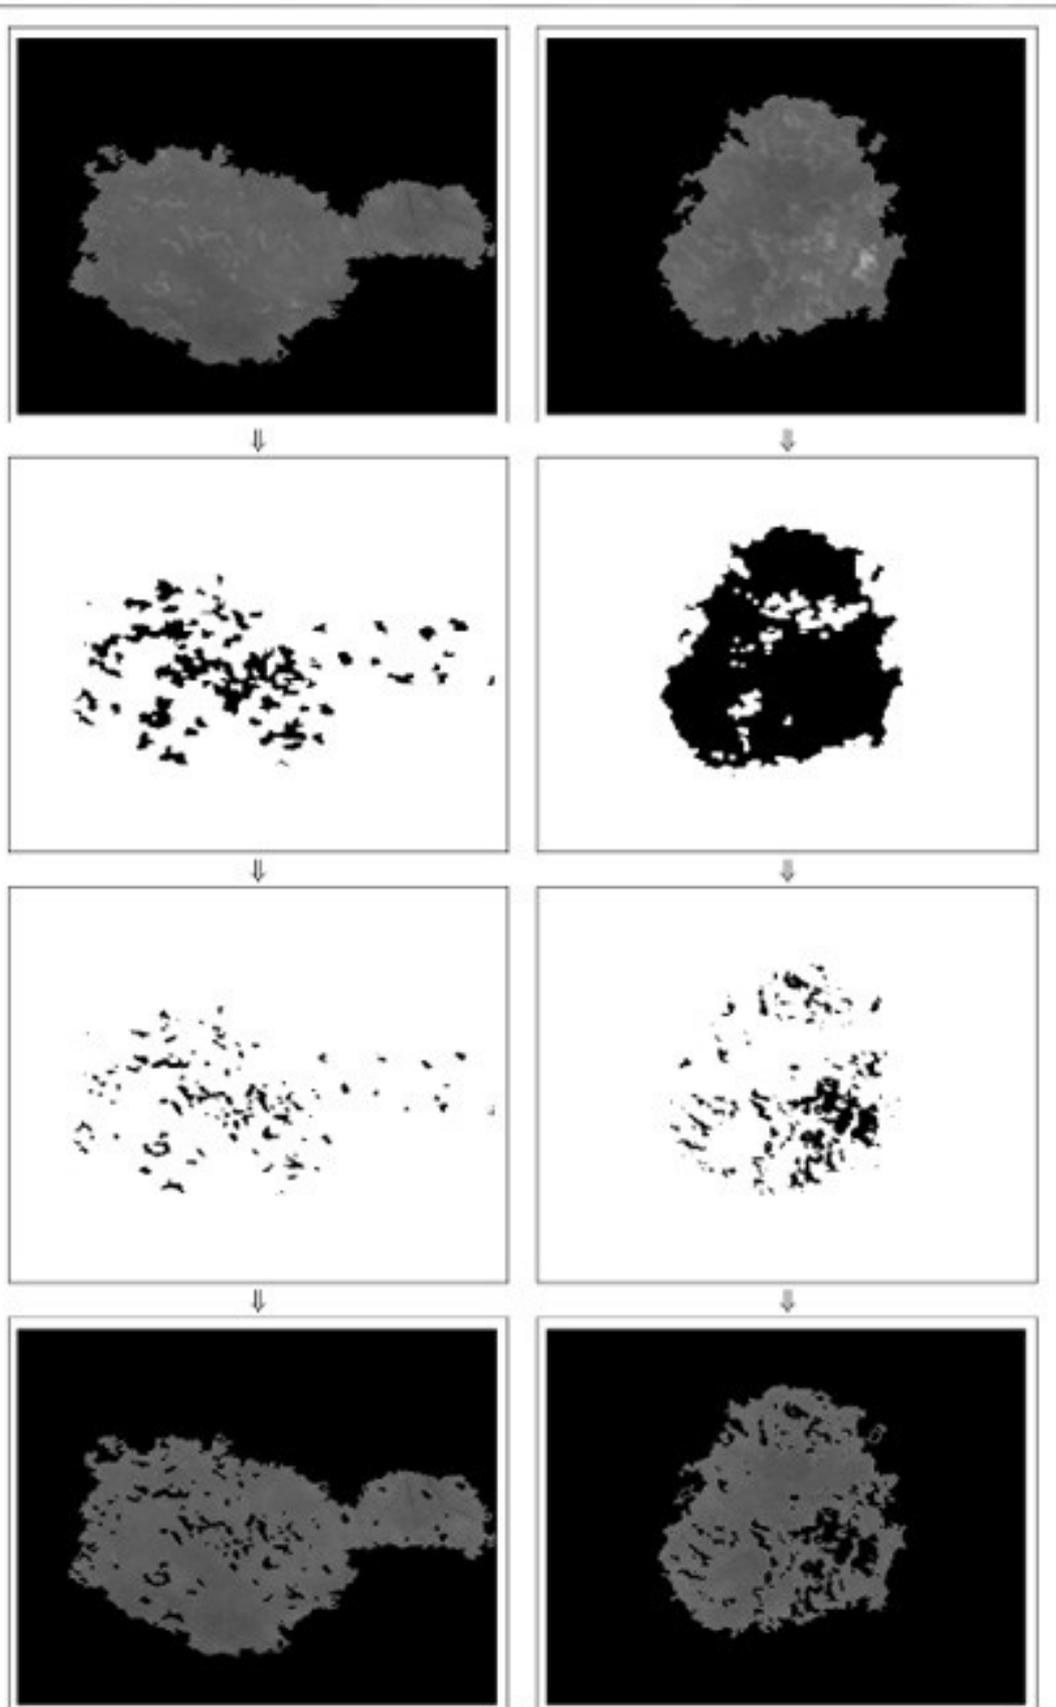


Figure 2: Top row: The original lesion. Second row: Scaling markers. Third row: Scaling segmentation result. Bottom row: A clear display of segmentation on top of the original image.

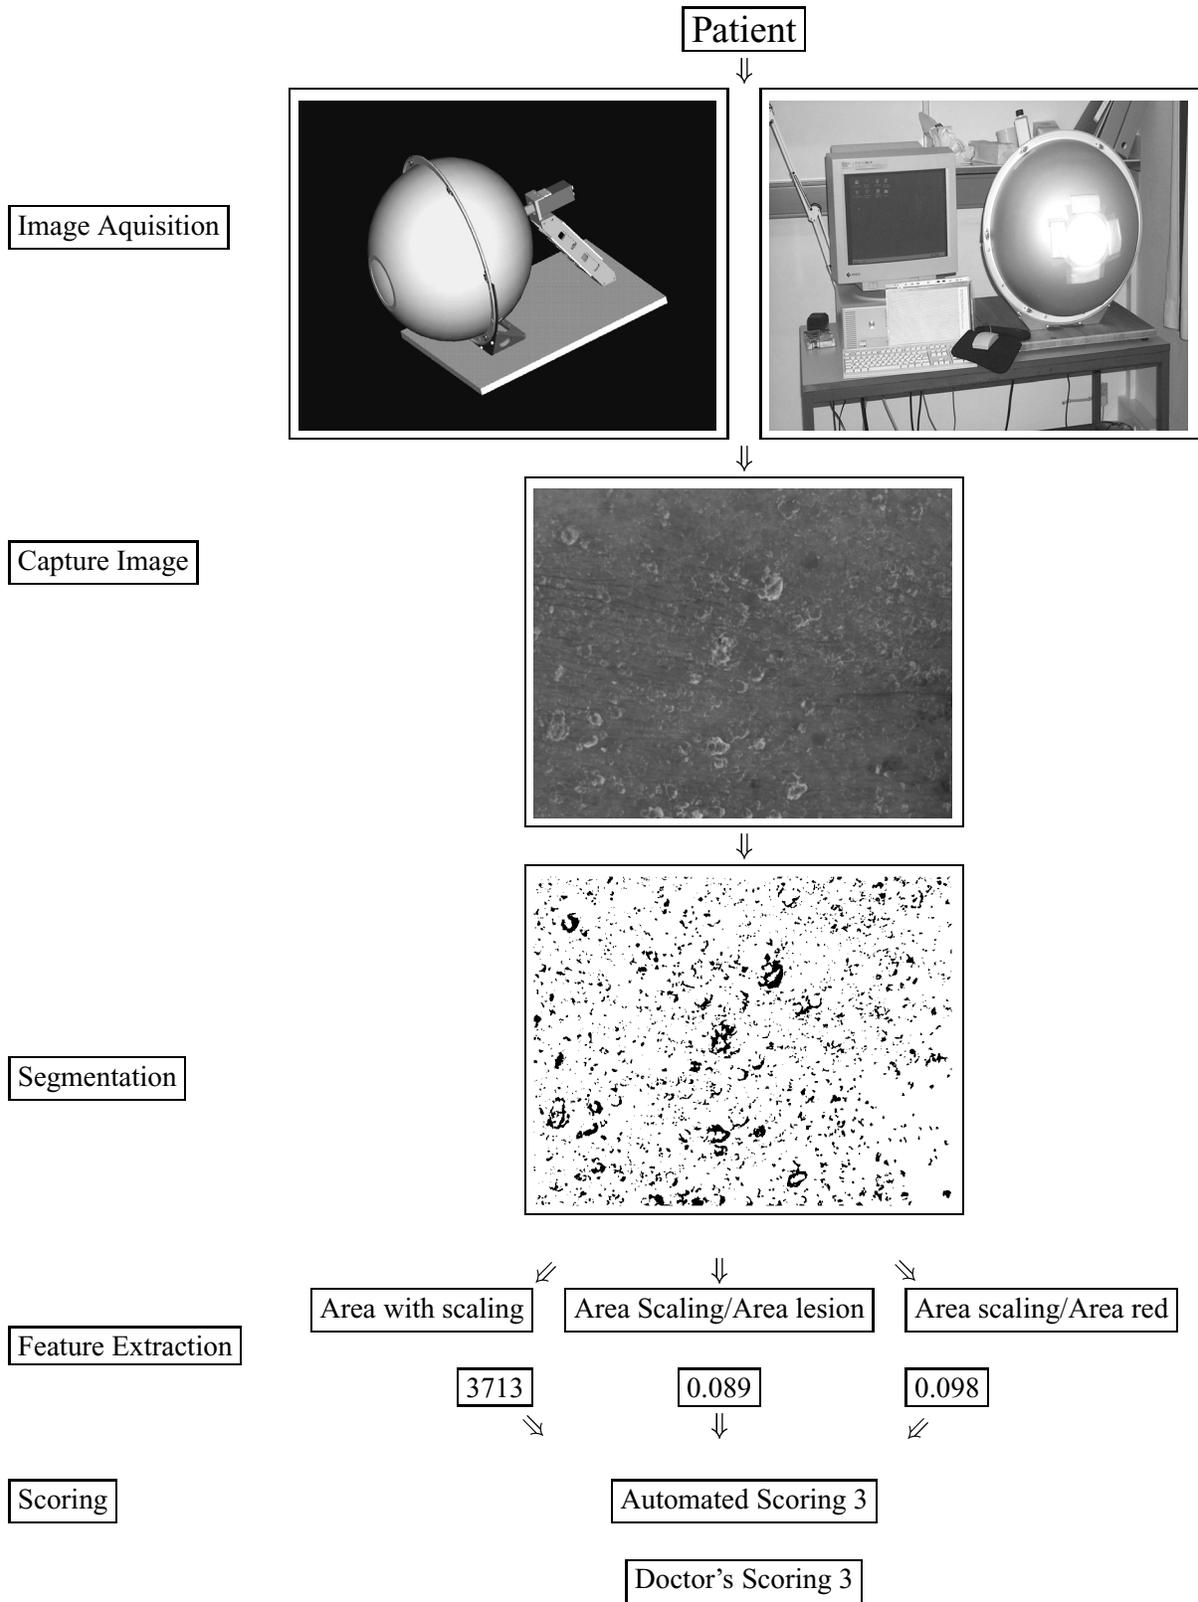


Figure 3: diagram of the method.

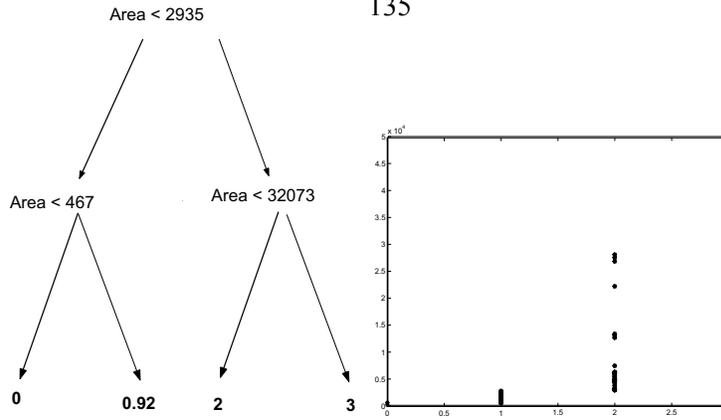


Figure 4: Left: Decision tree for the scoring given the parameters of the segmentation. Right: Dependency of lesion area on physicians' scoring of the lesions

### 3 Experiments

Two experiments are conducted to test the accuracy of the proposed method to score the scaling and redness in psoriasis lesions

#### 3.1 First experiment: Scoring the degree of Scaling

In collaboration with the dermatological department of Gentofte Hospital in Denmark an experiment was conducted. The goal of the experiment is to objectively score the severity of the scaling in psoriasis images. To accomplish this goal, a set of 46 psoriasis images was selected from a database of psoriasis collected from different patients. The physicians scores of these images was also available. The images were selected to cover the maximal possible diversity. The different areas of each image were extracted according to the procedure described in the previous sections and the above mentioned three summary values were obtained. A cross-validation process was used to build 23 decision trees. These decision trees utilized 44 data points to build the tree and two for testing it. Results showed that the *rst* variable, the area of the scaling, is enough to explain the physicians scoring. The automated scoring with our method has proven reliable, and on several occasions even allowed for corrections of physicians' mistakes. In these cases, the physicians were asked to re-score their previous judgements, and in all cases the assessment was changed.

Figure 4 left shows the nal tree generated using all the points. Figure 4 right plots the area of the scaling versus the physicians' scoring.

#### 3.2 Second experiment: Classifying the severity of redness

The second experiment aims to assess the possibility of automatically scoring the degree of the redness of the lesion. To achieve this goal, a set composed of 77 images of psoriasis lesions was selected from a data set with 175 images to perform the experiment. The selected images do not present shadows, scars or other elements that could spoil the result of the experiment. The severity of redness for each image was scored by the physicians in order to have a reference measure. The different areas involved in the chosen images were segmented according to the procedure described previously. The mean of the tri-chromatic bands was calculated in the healthy and in the red skin area for each image. The difference of these two means showed to be a good feature to evaluate the lesion.

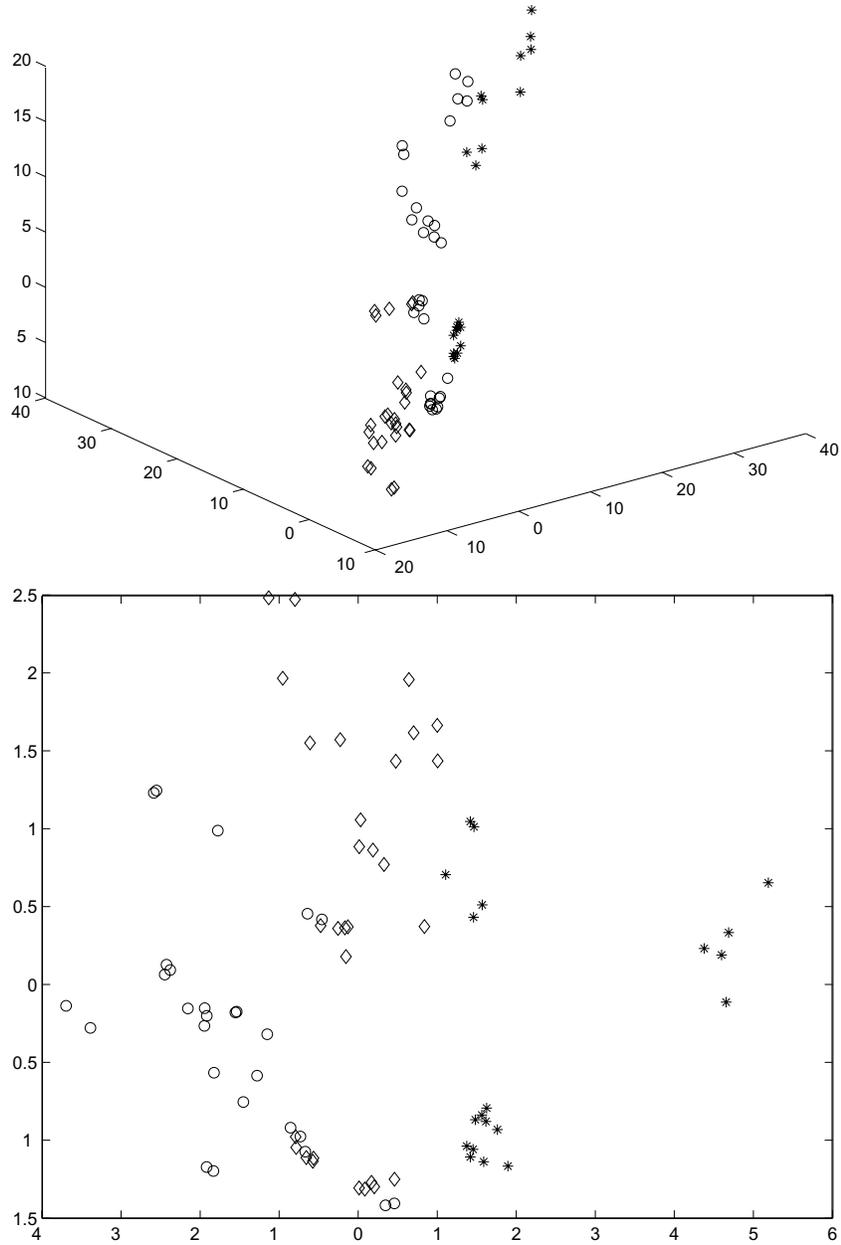


Figure 5: Up: 3D plot of the variables considered to classify the redness. Down: Result of applying a canonical discriminant analysis to the three selected variables to classify the redness

Figure 5 left shows a 3D plot of the differences where the different symbols represent different physicians scorings. The presence of three defined groups indicates the possibility of classify the redness using a clustering algorithm. This suggestion is even more clear if a canonical discriminant analysis is realized as it can be noticed in Figure 5.

## 4 Summary and conclusion

In this work, a procedure to evaluate the severity of the scaling and redness in psoriasis has been developed. The method automatically separates the different parts and extracts different parameters. In certain difficult cases such as uneven illumination it has been noticed that, allowing a manual interaction increases the accuracy notably.

The method provides objective measures that avoid the dependence of the physician in the tracking of dermatological diseases. It has been shown that one of the provided measures is highly correlated with the doctor scoring. Together with the other two measures we expect to be able to provide a better lesion description.

## References

- [1] Engström N., Hansson F., Hellgren L., Tomas J., Nordin B. Vincent J. and Wahlberg A. Computerized Wound Image Analysis In Pathogenesis of Wound and Biomaterial-Associated Infections. Springer-Verlag p 189-193, 1990.
- [2] Hansen G., Sparrow E., Kokate J., Leland K., Iaizzo P. Wound Status Evaluation using Color Image Processing IEEE Transactions on Medical Imaging, vol. 16, no.1, February 1997
- [3] Hillebrand G., Miyamoto K., Schnell B., Ichihashi M., Shinkura R., Akiba S. Quantitative evaluation of skin condition in an epidemiological survey of females living in northern versus southern Japan Journal of Dermatological Science, vol. 27, p 42-52, 2001.
- [4] C. Camisa, *Handbook of psoriasis*, Blackwell Science, 1998.
- [5] Johnson R., Wichern D. Applied Multivariate Statistical Analysis. Chapter 8. Prentice-Hall, 1995.
- [6] Hyvärinen A., Karhunen J., Oja E. Independent Component Analysis Wiley publications, 2001.
- [7] Taxt T., Hjort L., Eikvik L. : Statistical Classification using a Linear Mixture of two Multi-normal Probability Densities. Pattern Recognition Letters.(1991) 12 731-737
- [8] Vincent L., Soille P. Watersheds in digital spaces: an efficient algorithm based on immersion simulations, IEEE Trans. Pattern Anal. Mach. Intell. 13 (6) (1991) 593-598.
- [9] Sonka M., Hlavac V., Boyle R. Image Processing Analysis, and Machine Vision Brooks, p 128-130, 1999.

# Survey and Assessment of Advanced Feature Extraction Techniques and Tools for EO Applications

Mikael Kamp Sørensen<sup>a\*</sup>, Michael Schultz Rasmussen<sup>a</sup>, Henning Skriver<sup>b</sup>, Peter Johansen<sup>c</sup>,  
Arthur Pece<sup>c</sup>, and Jesper Hoyerup Thygesen<sup>d</sup>.

<sup>a</sup>GRAS, c/o Inst. of Geography, Univ. of Copenhagen, Øster Voldg. 10, 1350 Copenhagen, Denmark

<sup>b</sup>Technical Univ. of Denmark, Ørsted-DTU, 2800 Kgs. Lyngby, Denmark

<sup>c</sup>Dept. of Comp. Science, Univ. of Copenhagen, Universitetsparken 1, 2100 Copenhagen, Denmark

<sup>d</sup>ROVSNING A/S, Dyregårdsvej 2, 2740 Skovlunde, Denmark

## ABSTRACT

In addition to the substantial amounts of available Earth Observation (EO) data, there is currently an increasing trend towards the acquisition of larger and larger EO data and image quantities from single satellites or missions, with multiple, higher resolution sensors and with more frequent revisiting. More sophisticated algorithms and techniques than those largely in use today are required to exploit this rapidly growing wealth of data and images to a fuller extent.

The project "Survey and Assessment of Advanced Feature Extraction Techniques and Tools for EO Applications" (SURF) funded by the European Space Agency (ESA) will address these issues. The objective of SURF is to provide an overview of the current state-of-the-art methods within feature extraction and manipulation for EO applications and to identify scenarios and related architectures for exploitation of the most promising EO feature extraction methods. The task is to identify the most promising methods to extract pertinent information from EO data on environment, natural resources and security issues. SURF aims at listing existing methods with the final goal of identifying the three most promising methods to be implemented in prototype solutions. The work includes the development of the concept for the evaluation and rating of methods relative to the users needs for information, the maturity and novelty of the methods, the potential for fusing data and the operational feasibility. Special emphasis will be made regarding the exploitation of state-of-the-art image processing, pattern recognition and classification techniques.

**Key words:** automatic feature extraction, classification, survey, methods, EO, remote sensing, data fusion

## INTRODUCTION

The global archives of remote sensing data are constantly increasing due to a number of factors. First, the amount of operational remote sensing satellites is continuously growing. Second, the spatial and spectral resolutions of the new sensors are constantly improving due to technological advances, resulting in exponentially larger data volumes. Finally, the data transmission capabilities of modern satellites result in vastly increased data transfers.

However, even though the use and exploitation of EO data is increasing, the vast data archives are not being fully exploited. It has been estimated that only a small percentage of the remote sensing data archived annually are actually being utilized by end users. Part of the reason is that most remote sensing applications depend on image processing carried out by experts. Even though the basic image processing such as geometric rectification and radiometric calibration is

---

\* mks@gras.ku.dk; phone +45 35 32 41 75; fax +45 35 32 25 01; www.gras.ku.dk

largely performed automatically for most remote sensing products, the conversion from image data to the thematic information required by end users is a constraining factor for the further proliferation of remote sensing applications.

A number of initiatives have already focused on automatic feature extraction. Through the Information Society Technologies (IST) programme, the European Union seeks to mature technological advances. An example of such a project is Satellite Based Information System on Coastal Areas and Lakes (SISCAL) (SISCAL, 2003), that will provide near-real time data on the marine environment. Examples of data mining projects include the ESA project Knowledge Driven Information Mining in Remote Sensing Image Archives (KIM).

However, there is still a need for developing automatic feature extraction methods that will make remote sensing information readily available to a wide range of end users. The "user" is considered as a broad category from professionals to the general public, from the environmental scientist to the analyst working in a county or a citizen seeking information through the Internet. For this reason EO data products must be exploitable in a user-friendly way through a clear understanding and description of the information and through digital data formats that are easy to handle.

Most end users work with thematic features that change over time and the growing archives of remote sensing data further encourage and facilitate the analysis of time series. Change detection and trend analysis are important techniques that are able to match the user requirements within global change and environmental monitoring.

It is important for the exploitation of remote sensing data to try to identify novel methods. An interesting perspective is the combined use of data from different sensors applying data fusion methods. In the context of the current project, data fusion may both refer to a specific data fusion algorithm (i.e. wavelets, PCA) combining two or more different data sources or it may simply refer to the integrated use of two different image sources. Furthermore, it is the intention to take a broad view of possible candidate methods not only within classic remote sensing sciences, but equally within related fields such as image processing as being explored by medicine and in the studies of materials and minerals using microscopes.

The project "Survey and Assessment of Advanced Feature Extraction Techniques and Tools for EO Applications" (SURF) funded by the European Space Agency (ESA) will deal with these issues. The main objectives of the SURF project are to obtain:

- an up-to-date picture of advanced Methods for the support of Earth Observation applications through Feature Extraction.
- the identification of possible scenarios and related architectures for the exploitation of the most promising EO Feature Extraction Methods.
- the implementation of three Methods using real image data from three cases.

The objectives will be achieved through a survey, evaluation and ranking of Methods applicable to the EO domain. For the three most promising Methods, prototypes will be developed for evaluation using image data.

In the SURF project, a Method (with capital M) is here defined as a way of solving a problem within a given context. A Method may be a specific algorithm or technique or it may be understood in a broader sense, such as a complete software package. Feature extraction can be defined as the conversion from image data to thematic information. In the SURF project, the intention is to support the user in terms of pre-processing through automating the complex part of the image processing by identifying features, patterns, trends, changes etc. and hereby making this information readily available. The approach to the screening of EO research status and trends will start with the identification of thematic areas where EO data have demonstrated their usefulness. Thematic areas where EO data are expected to provide important information in the future will equally be included.

Most scientists maintain the idea of their work becoming useful to the society in one form or the other and the objective of the present paper is to demonstrate how scientific results are evaluated for possible real-world use and hereby propose a batch of important criteria. This will be done by presenting the strategy for the evaluation and the selection of Methods as being done within the SURF project. The work will be placed in a context through the description and evaluation of

three candidate M methods. These are entropy based prediction, NDVI trend map from time series and knowledge-based classification by SAR.

## METHODOLOGY

In order to meet the objectives mentioned above, the survey started with a broad approach, listing the main thematic areas, the applications and M methods. In the consecutive steps, the list of candidate M methods will be progressively narrowed down and in the end, three M methods will be selected for prototype implementation.

At the application level, user considerations are most important. If there is no or little user interest in an application, no further effort will be devoted towards that particular application. At the M method level, the technical considerations are central. If a certain M method is very difficult to implement for any reason, it will get a low rating and will not be considered for prototyping. Table 1 provides an example of the relation between thematic areas, applications and M methods.

Table 1. Example of the relation between thematic areas, applications and M methods.

Thematic area	Thematic sub-area	Application	M method
Environmental monitoring	Environmental degradation	Monitoring deforestation	Regression tree classifier
			Multivariate Anomaly Detection and MAF
			Hyper-columns using Gabor Filters
		Mapping desertification	NDVI trend analysis
			Context-based object-oriented classification

The rating and selection of the most promising M methods will be carried out in close dialog with end users, the scientific community as well as with ESA.

### STEP ONE: SELECTION OF THEMATIC AREAS AND APPLICATIONS

An initial assessment of the user needs within each thematic application is made, considering the following factors:

- a) Users' need for information – The users' need for information within the category is evaluated. Applications that are not considered useful will result in a low score.
- b) Relative provision of information relative to other sources – If the information potentially provided from EO data can easily and cost effectively be obtained through other data sources (non EO data) it should result in a low score. On the other hand, if no other means than EO data is available for obtaining the information, or if complementary data can be achieved using remote sensing data, a high score should be given.
- c) Users' need for automation – This category evaluates the users' need for automation in a given area. The need will be high when working with operational topics or in areas with large amounts of data. Automation within EO data processing may equally be able to provide users with products or semi-products that can be automatically extracted from EO data during the pre-processing and hereby become more easily available and less costly to the users. For ad-hoc applications the need for automation may be less pronounced.

The assessment of user needs is first carried out by the consortium and rated according to the three categories mentioned above. This will be followed by a dialog with users, user communities or representatives to get feedback on the rating and selection and hereby allow adjustment of the ratings.

During the initial phase of the project, five overall thematic areas were identified. Each thematic area is characterized by a number of sub-thematic areas, and within each of these sub-areas, a number of thematic applications are found. The identification of thematic areas and applications is mainly based on the literature and knowledge of existing operational applications.

Table 2 shows the preliminary list of thematic areas. For each sub-area, the relevance of data categories or possible data fusion scenarios is indicated.

Table 2: List of the five thematic areas and related applications				
	Optical data	Radar data	Data fusion	Time series
<b>Environmental monitoring</b>				
Environmental degradation	X	X	X	X
Land cover and land use change	X	X	X	X
Eutrophication and oxygen depletion of the marine environment	X			X
Oil spill monitoring	(X)	X	(X)	X
<b>Global change</b>				
Changes in vegetation productivity (carbon pools)	X	X	X	X
Land cover changes	X	X	X	X
Changes in (patterns of) Sea Surface Temperature (SST)	X			X
Changes in marine Chlorophyll a concentrations	X			X
<b>Security</b>				
Sea-ice monitoring and warning	X	X	X	X
Flooding, mapping and forecast	X	X	X	X
Identifying forest or savannah fires	X			X
Meteorology	X	X	X	X
Volcanic activity	X	X	X	X
<b>Policy and Planning</b>				
Carbon accounting	X	X	X	X
Land cover mapping and changes	X	X	X	X
Agricultural subsistence and control	X	X	X	X
Cartography	X	X	X	X
<b>Resource Assessment</b>				
Geology	X	X	X	
Fisheries	X			X
Forestry	X	X	X	X
Fresh Water	X	X	X	X

## STEP TWO : LISTING OF METHODS

Based on the list of thematic areas and applications, a large number of feature extraction Methods were identified based on the experience of the consortium as well as an initial literature search. During this phase Methods are identified without any further evaluation. An example of Methods listed within a given application is found in Table 3.

Table 3 .Example of Methods for Monitoring Deforestation

RGB -NDVI differencing
Multivariate Autoregression / Maximum Autocorrelation Factors
The VIT auto change Method
Iterative fitting of Principal Component Axis
Regression tree classifier
PCA Differencing
Decision tree fuzzy classification
Hyper-columns using Gabor Filters
Classifications using space-time segmentation
Context-based object oriented classification
Neural networks
Entropy prediction
Knowledge based classification

### STEP THREE : EVALUATION OF METHODS WITHIN THE MOST IMPORTANT APPLICATIONS

For each Method, an evaluation of the technical considerations is carried out. This project phase will be subjected to three iterations in order to secure a thorough evaluation, where new inputs and additional information can be added in the process.

The following categories are used to rate each Method within promising applications:

- a) Maturity of the Method in terms of technical and EO compliance – This parameter considers how robust the Method is. One criteria is to observe the amount and generality of publications that are available on the Method, preferably from different sources and with examples of their implementation. Also, if the Method has been implemented in commercial software packages, operational projects or if it is routinely used by various institutions it may be considered mature. The Method in question does not necessarily need to have a history in EO applications in order to be mature - maturity may also refer to well established applications in other areas of say pattern recognition.
- b) Novelty of the Method – The Method should be first published or reported recently to get a high score. It does not have to be an innovative algorithm, the novelty may equally consist in combining well established Methods or applying well known Methods in an innovative manner.
- c) Operational perspectives (timing, effort, computational requirements, cost etc.) – The perspectives for implementing the Method will be evaluated here. A Method that is computationally demanding (i.e. extensive iterative processes) will be given a low score if it restricts the timely output of results (for instance oil spill mapping where the results have to be delivered in near real time). If the Method is capable of delivering timely outputs within a reasonable financial limit, a high score should be given.
- d) Accuracy of the Method in relation to the specific application – A high absolute accuracy gives a high score. If the relative accuracy is high (providing a correct assessment of the spatial distribution), the Method can be given a high score even if the absolute accuracy is poor. An example is chlorophyll a concentration monitoring, where the absolute accuracy is poor, but the product meets the needs of environmental monitoring institutions to obtain an overview of the spatial distribution of algae.
- e) Potential use of the Method in other areas – If a Method is extremely specific and can only be applied within a single or limited areas a low score should be given. On the other hand, if the Method is of a general nature that allows for its application in other thematic areas, a high score should be given. The number of cross-references will be a useful indicator (cross-references indicate if the same (or nearly the same) Method is listed in other applications).
- f) Perspectives for automation – If a Method can be run completely automatically, a top score should be given. If a substantial amount of manual work is required, a low score should be given. For Methods where some user interaction is required (i.e. specification of constants etc.) a moderate score is appropriate.
- g) SURF feasibility – The feasibility of implementing a given Method within the SURF project will be assessed separately. If a Method is extremely demanding to implement, it will not be possible under the current project.

After the evaluation of more than 50 Methods, approximately 20 Methods will be selected based on the evaluation scores listed above. Selected users and scientists (or groups of) will be consulted to make comments and suggestions on the ratings and decisions being made.

### STEP FOUR : DETAILED ANALYSIS OF SELECTED METHODS

The 20 Methods selected in step three will be narrowed down to three Methods that will be used for prototype development. The process of selecting these three Methods will involve a detailed description of the 20 candidate Methods. The feasibility of the Methods will be assessed from the literature and through contact with the scientific community. Pilot and test cases will be studied with special attention to the potential level of generalisation. User requirements will be

analysed or confirmed through dialogue with users, as well as the technical aspects will be evaluated. ESA will actively take part in the selection process.

#### STEP FIVE: PROTOTYPE DEVELOPMENT

For the three most promising methods, working prototypes will be implemented. As some of the methods may never have been tested in the EO field, the evaluation of prototypes working on EO data is an important step in order to acquire specific working knowledge of the new methods. The prototype implementation will as far as possible be based on existing available software libraries and CO TS.

Criteria for evaluation of the prototypes will be established as a subset of the criteria for the ranking of methods in step three. The operational perspectives, the methods' accuracy, and the level of automation can be tested by running the prototypes on prepared data sets.

Based on the combined knowledge of methods and the user's needs, realistic exploitation scenarios will be described. This includes an assessment of the quality of the information to be provided along with an assessment of the required human resources and the necessary hard- and software.

### METHOD EXAMPLES

In the following section three examples of candidate methods are described and table 4 shows how these were rated.

#### ENTROPY BASED PREDICTION

The purpose is to detect unexpected changes. These could be the possible occurrence of patches of deforestation or changes in agricultural practice (e.g. from fallow to cultivation). The method may even be used to clean up clouded pixels in a time series. Which events are unexpected is a question of prior knowledge of events in the area. The information may in the simplest case be direct prediction of change per pixel, but can equally well be a prediction of a possible new classification scenario.

The principle is to collect statistics on the time sequences at each pixel. Shannon (1948) used this method to quantify information. In a database an account is made of all occurrences of temporal sequences of events. Shannon showed how it would be possible to distinguish between different languages by counting the frequencies of letters in sentences from the language, the frequencies of all consecutive letters, of all triples of letters, etc. The present method is a modern implementation of his method. Shannon showed that his prediction method would be as precise as a program which in advance knew the transition table of a Markov Source which generated the sentences of the language. Lempel and Ziv (1986) showed how to use this principle efficiently for data compression. The principle in data compression is that the program estimates the probability of each symbol before encoding it. If the symbol occurs often, a short code word is used, and if it occurs only infrequently, not much harm is done using a long code word. This coding principle is widely used for text compression in modern personal computers.

When applied to images, statistics per pixel is collected, either as a result of a classification, a spatial filter or the quantization of spectral values. The program then decides if an event is unexpected in the following way. From a query to the database it gets information about time and location where a similar previous sequence of event has occurred before. This way it can make use of the prior information about events in the area. There may be several occasions at which similar phenomena have occurred. The program chooses the longest identical sequence of events. Once it has obtained a time and a place where a similar event has happened, it makes the prediction that the next event to occur is the event which took place at the previous time and location where the optimal match was found.

A surprising aspect is that even if keeping track of all sequences of all lengths seems a titanic job, there is an algorithm which makes it possible to use only one quantum of computation independent of how long the sequence of images is. This means, that given a sufficiently fast computer, the method will work in real time. Furthermore, the method will only need one quantum of storage space for each image it receives. It is not necessary to use increasingly more space for

collecting the statistics. This is equally surprising, since there is no limit on the length of the sequences which the data base keeps track of. The data structure used in the data base is called a suffix tree, and the algorithm was invented by Weiner (1973) quite some years ago. Few read Weiner's paper, a newer version was proposed by Ukkonen (1992).

#### NDVI TREND MAP

Time series of NDVI data derived from EO data are used to identify long term trends in vegetation productivity. NOAA AVHRR data have been used where the integration of NDVI during the growing period is used as estimates of the annual Net Primary Production (NPP). This Method has been applied in dry areas (Rigina and Rasmussen, 2003, Tottrup and Rasmussen, submitted) as well as globally (Nemani et al, 2003). Degrading conditions are identified when integrated NDVI consistently are decreasing. The information is used by ministries of environment in sub-Saharan Africa as a support to the implementation of the UN Rio conventions on desertification and climate change. The information serves as reference data on the state of the environment as well as being used to identify areas subject to specific actions to mitigate consistently degrading environmental conditions.

The Method is based on a few principles. A number of authors have demonstrated the existence of a relationship between NDVI and NPP (Tucker et al, 1986, Prince 1991 and Rasmussen 1992) using the Light Use Efficiency model (LUE) (Kumar & Monteith, 1981). NDVI can be used to assess the fraction of absorbed photosynthetically radiation (fAPAR) and the NPP can be calculated as:

$$NPP = e fAPAR_{NDVI} PAR$$

Where  $e$  is the efficiency of transforming energy into dry matter and  $PAR$  is the photosynthetically active radiation.  $fAPAR_{NDVI}$  is the  $fAPAR$  estimated using NDVI from EO data.  $PAR$  can be obtained from EO data or meteorological data and the  $e$  can be assumed constant for limited areas and during limited time, or values can be determined as a function of ecosystems and different stress parameters. This is the approach applied for the MODIS NPP product (Justice et al., 1998).

From a time series of annual maps of NPP as a function of integrated NDVI, a per-pixel trend map can be computed, where the value of the slope is assigned to each pixel in the image. This will show the overall conditions and areas with consistently decreasing NPP can be identified. The Method was evaluated and further developed in Senegal and results showed that areas experiencing soil exhaustion as well as areas where enabling conditions were prevailing were both captured from a time series of 19 years of Pathfinder AVHRR NDVI data (Tottrup & Rasmussen, submitted).

#### KNOWLEDGE-BASED CLASSIFICATION OF SAR DATA

The Method can be used to map changes within land cover, land use and vegetation. More specifically, detailed information and land cover can be applied in relation to regional and global general circulation models to calculate the fluxes of water-energy-carbon. A knowledge-based, hierarchical classification scheme has been proposed for land-cover classification into relatively broad classes using Synthetic Aperture Radar (SAR) data. The concept has been demonstrated using airborne and spaceborne polarimetric SAR data (Pierce et al., 1994 and 1998) as well as combinations of SAR data from two satellites at different frequencies (C- and L-band) (Dobson et al., 1995 and 1996).

Radar backscatter is sensitive to two types of scene parameters, geometrical and dielectrical properties. The geometrical properties refer to the structure of the vegetation and to the roughness of the bare surface or the surface under the vegetation. For the vegetation, the structure parameters important for radar backscatter are the size distribution, the orientation distribution, and the densities of the vegetation elements (i.e. stem, branches, and leaves). The dielectrical properties are determined by the water content, the phase of the water and the densities of the scattering elements. Based on these sensitivities it is possible to discriminate between different types of vegetation with different structural and dielectrical properties using SAR data. The sensitivity of the radar to these parameters varies with radar frequency and polarisation. Hence, it is possible to resolve potential ambiguities in the discrimination between different vegetation classes using multi-frequency and/or polarimetric SAR data.

In Dobson et al. (1995 and 1996) and Pierce et al. (1994 and 1998) knowledge about the scattering properties from vegetation as well as electromagnetic model results were used to derive a hierarchical classification scheme. First, in the so-called Level I classifier the images are classified into 3 very broad classes: surface, low vegetation and tall vegetation. The classification accuracy of this classifier is found to be very good with accuracies close to 100%. Based on the output of this classification, a Level II classifier is then applied. This classifier may for instance classify the tall vegetation into different forest types, and the low vegetation into different crop types. Classification accuracies for the overall classification scheme of above 90% were normally obtained.

The above-mentioned classification scheme is relatively robust in terms of for instance time of the year, and geographical location, because it is based on knowledge and model results. The classification scheme must, however, to some degree be adapted to the data available and to the requirements for the land-cover classes. After the manual set up of the algorithm, it is very simple and fast.

Table 4: Ratings of three different methods						
Technical Considerations	Entropy based prediction		NDVI Trend Map		Knowledge-based classification of SAR data	
Maturity of the Method in terms of technical and EO compliance	**	The Method has not been used before on images. Widely used for text-processing.	****	Mature and used. Must be assessed for the area considered.	****	The Method is robust because it based on physical modelling and knowledge of backscatter
Novelty of the Method	*****		*****	Demonstration studies have been made and a number of independent science teams works currently with the Method.	**	
Operational perspectives (timing, effort, computational requirements, cost etc.)	**	Somewhat uncertain. In principle, the Method can be implemented to function in time proportional with the length of the sequence, but it will take some care. A prototype program was run some years ago on image (32*32 pixels) with sequences of length 96.	*****	Very good perspectives. Has been demonstrated.	*****	When the classification scheme is set up, the Method is extremely simple and fast, and the computational requirements are small. Good perspectives in handling large data volumes
Accuracy of the Method in relation to the specific application		The accuracy cannot be estimated due to the novelty of this Method.	****	Alternatives are less accurate and will not cover the entire area being monitored.	****	Normally the classification is into broad classes. The actual accuracy depends on the data available (parameters and acquisition time)
Potential use of the Method in other areas	****		***	May be used within different vegetation types such as forest. Long term changes in land surface temperature.	*****	Method already used in other applications
Perspectives for automation	****		*****	Most relevant because the Method is data demanding.	****	The set up of the classification scheme is done manually, and then the classification is fully automatic
Feasibility of implementing the Method within SURF (resources)	**	A prototype could be implemented by the team, but not a real-time version.	*****	The team has experience developing and implementing the Method.	***	Relatively simple Method without any computational complexity. Most of the work is included in the set up of the classification scheme

User Considerations	Entropy based prediction (Application: Mapping deforestation)		NDVI Trend Map (Application: Mapping desertification)		Knowledge-based classification of SAR data (Application: Mapping land cover and land use changes)	
The user's need for information	*****	Strong need especially in countries where conventional statistics are not reliable.	*****	Needed for monitoring the state of the environment and possible degradation of natural resources.	****	Planning and natural resource management.
The potential provision of information relative to the effort of the user to obtain this information through other sources	*****	EO data the only feasible inf. source that provide homogeneous data at the regional level.	****	EO data the only feasible inf. source that provide homogeneous data at the regional level.	****	Homogeneous and repetitive information from EO data.
User's need for automation	****	Large data volumes necessitate a high degree of automation.	****	Large data volumes necessitate a high degree of automation.	****	Assistance to compile the vast database of EO data needed.

## CONCLUSION

A strategy has been presented for a survey and assessment of candidate EO methods to be implemented in automatic or semi-automatic feature extraction. The thematic areas and subsequent applications where EO data can support users with important information have been listed based on the literature and existing knowledge. These are rated according to user's potential need for information. Only thematic areas and applications where the user's interest can be justified have been selected for further consideration. For each application, different methods are listed and a first assessment is made in order to reduce the number of methods to approximately 20. Examples are presented including the corresponding ratings. This rating of the methods are to a large extent subjective, however, it is assessed that it can be used for documentation and a first selection. An alternative would be to use a decision model attributing different weights to the different ratings and using conditional selections. In this case it will be essential to guarantee the independency between the team responsible for the development of the decision model and the team responsible for the rating of the methods. Furthermore, evaluators bias should be assessed through cross checks.

The next step in the process will be the in depth feasibility study of the 20 selected methods. The results of the SURF project will become publicly available at ESRIN's web pages at [earth.esa.int/itd/Projects/mid2004](http://earth.esa.int/itd/Projects/mid2004).

## ACKNOWLEDGEMENTS

The SURF Consortium consists of ROVSING A/S, Geographic Resource Analysis and Science Ltd. (GRAS), Technical University of Denmark, Ørsted-DTU Department and Department of Computer Science, University of Copenhagen. The SURF project is funded by the European Space Agency (ESA), ESRIN contract no.: 17127/03/I-O L with Alessandro Ciarlo as the technical officer for SURF.

## REFERENCES

Dobson, M. C., F.T. Uhlaby, and L.E. Pierce, 1995, Land-Cover Classification and Estimation of Terrain Attributes Using Synthetic Aperture Radar, *Remote Sensing of the Environment*, 51, 199-214.

Dobson, M. C., L.E. Pierce, and F.T. Uhlaby, 1996, Knowledge-Based Land-Cover Classification Using ERS-1/JERS-1 SAR Composites, *IEEE Transactions on Geoscience and Remote Sensing*, 34, 83-99.

Justice, C.O., Vermote, E., Townshend, J.R.G., DeFries, R.S., Roy, D.P., Hall, D.K., Salomonson, V.V., Privette, J.L., Riggs, G., Stahler, A.H., et al., 1998, The Moderate Resolution Imaging Spectroradiometer (MODIS): Land Remote Sensing for Global Change Research. *IEEE Transactions on Geoscience and Remote Sensing*, 36, 1228-49.

Kumar, K., Monteith, J.L. 1981, Remote Sensing of Crop Growth. In *Plants and the Daylight Spectrum*, edited by Monteith, J.L. (London: Academic Press), pp.133-44.

Lempel, A. & Ziv, J., 1986, Compression of 2-dimensional data. *IEEE transactions on information theory*, 32, 2-8.

Nemani, R.R., Keeling, C.D., Hashimoto, H., Jolly, W.M., Piper, S.C., Tucker, C.J., Myneni, R.B., Running, S.W., 2003, Climate-driven increases in global terrestrial net primary production from 1982 to 1999. *Science*, 300, 1560-3.

Pierce, L.E., F.T. Uhlaby, K. Sarabandi, and M.C. Dobson, 1994, Knowledge-Based Classification of Polarimetric SAR Images, *IEEE Transactions on Geoscience and Remote Sensing*, 32, 1081-1086.

Pierce, L.E., K.M. Bergen, M.C. Dobson, and F.T. Uhlaby, 1998, Multitemporal Land-Cover Classification Using SIR-C/X-SAR Imagery, *Remote Sensing of the Environment*, 64, 20-33.

Prince, S.D., 1991, A model of regional primary production for use with coarse resolution satellite data. *International Journal of Remote Sensing*, 12, 1313-30.

Rasmussen, M.S., 1992, Assessment of millet yields and production in northern Burkina Faso using integrated NDVI from the AVHRR. *International Journal of Remote Sensing*, 13, p 3431-42.

Rigina, O. and Rasmussen, M.S., 2003, Using trend line and principal component analysis to study vegetation changes in Senegal 1986-1999 from AVHRR NDVI 8 km data. *Danish Journal of Geography*, 103, 31-42.

Shannon, C.E., 1948, A Mathematical Theory of Communication. *Bell System's Technical Journal*, 27, 379-423

SISCAL, 2003, URL: [www.SISCAL.NET](http://www.SISCAL.NET), August 2003.

Tottrup, C. and Rasmussen, M.S., Submitted, Mapping long-term changes in crop productivity in Senegal through trend analysis of time series of remote sensing data. Submitted to: *Agriculture, Ecosystems and Environment*,

Tucker, C.J., Justice, C.O., Prince, S.D., 1986, Monitoring the grasslands of the Sahel 1984-1985. *International Journal of Remote Sensing*, 7, 1571-81.

Ukkonen, E., 1992, Constructing suffix trees on-line in linear time. In J. v. Leeuwen, editor, *Algorithms, Software, Architecture*. *Information Processing 92*, Volume 1, pages 484-492. Elsevier, 1992.

Weiner, P., 1973, Linear pattern matching algorithms. In 14th Annual Symposium on Switching and Automata Theory, pages 1-11, The University of Iowa, 15-17 October 1973. IEEE.

## AUTOMATIC CHANGE DETECTION FOR VALIDATION OF DIGITAL MAP DATABASES

Brian Pilemann Olsen

Thomas Knudsen

National Survey and Cadastre, Denmark, 8 Rentemestervej, DK-2400 Copenhagen NV, Denmark, {bpo|thk}@kms.dk

**KEY WORDS:** Photogrammetry, Remote Sensing, Change Detection, Classification, Automation, High Resolution, Infra-red, DEM/DTM

### ABSTRACT

In almost all areas of our society there is an increasing need for up to date digital map databases. Traditionally, different manual, labour intensive and hence costly methods have been used for map updating, with the change detection for the updating being by far the most complex and expensive part. In this paper an automatic change detection method is presented and evaluated. The method only considers changes in the buildings theme, but it can be extended to other object classes. The aim is the development of an efficient change detection procedure for database maintenance in a production environment. The method is based on classification principles and combines an unsupervised and a supervised classification in order to determine the spectral response of the building class and thus locate potential buildings. The result is iterated by a height filter to refine the result. The method is evaluated on building registrations from the Danish TOP10DK map database. The test case presented in the paper is from a residential suburban area. The method detects almost all changes due to demolished buildings whereas only a smaller part of the new buildings are detected. This is primarily due to the use of a very special roofing material. The method leads to a number of false alarms, which to a large degree can be eliminated by refinements of the algorithm or by introduction of additional information e.g. infra-red images or texture measures.

### 1 INTRODUCTION

The digital topographic map database **TOP10DK** is the primary topographic product of the National Survey and Cadastre—Denmark (Kort & Matrikelstyrelsen, KMS). The development and update of TOP10DK is based on aerial photogrammetry: one fifth of Denmark is photographed each spring before foliage, resulting in approx. 1200 photos (about 400 GB of image data).

Map updating can be carried out by a complete remapping of the area for each revision cycle, but much work can be saved by detecting changes from the previous version of the map database and concentrating on these areas of change. Change detection for topographical mapping is on the other hand not a simple task: Although the intention is always to carry out the photographs at approximately the same time of year, the natural, inter annual variations of the vegetation coverage is of a magnitude that hides the (primarily human generated) changes sought for. Furthermore it is almost impossible to take the photos at the same geographical position and with the same attitude as within the previous photo campaign. This means that the change detection must be carried out by comparing a new image directly with the existing map database, rather than by a simpler image-to-image comparison.

In this paper an automatic procedure for change detection concentrating on buildings, which are important mapping objects, is presented. The next step in the update process: the actual 3D object registration, is not considered here. This subject has recently been treated extensively by Niederöst (2003) and Süweg (2003).

As can be seen from figure 1 buildings are often highly diverse both when it comes to size, form and spectral signature. They are therefore hard to describe by spectral information only. Adding height information to the process, e.g. in the form of digital surface models, may improve the distinguishing of buildings from other objects having a similar spectral response.

When introducing automatic change detection procedures, the aim is to detect at least the same percentage of factual changes as a manual operator is capable of. It is, on the other hand, acceptable if the change detection procedure introduces false alarms, as long as they are few, since they can easily be rejected during the actual 3D object registration.



Figure 1: Buildings are typically highly diverse and spectrally ill-defined, when considered as a single group. The last image (Lower Right) is a building as it is represented in a DSM.

#### 1.1 Related work

Other European countries e.g. Germany, Switzerland, and the Netherlands have also established and completed digital map databases with national coverage in the past few years. The National Mapping agencies in these countries therefore face the same problem as the KMS and projects with similar aims considering automatic or semi-automatic map updating have been established.

In Germany the project for updating the ATKIS database focuses on registration of more generic surface types (settlement, grassland, street, water etc.) (Petzold and Walter, 1999, Walter and Fritsch, 2000). The method for change detection uses supervised classification, with training areas automatically generated using the existing registrations in the ATKIS map data base (Walter, 2000). Experiments combining multi-spectral images (RGB, colour infra red – CIR) with height information and reduction of the information to generic surface types, have shown that it is possible to perform automatic change detection with a satisfactory accuracy (Petzold and Walter, 1999, Petzold, 2000) (note, however, that the accuracy requirements for ATKIS are somewhat lower than for TOP10DK (Kort & Matrikelstyrelsen, 2001, AVLBD, 1988)). The change detection leads to a “change map”

where the generic objects are divided in three classes: no change, possible change and change.

In the Swiss ATOMI project, aerial colour photos, a high resolution Digital Elevation Model (DEM) and a Digital Surface Model (DSM) are used aiming at the enhancement of the planimetric accuracy for the 2D VECTOR25 database (Eidenbenz et al., 2000, Niederöst, 2003). The surface model is generated by auto-correlation in aerial photographs in the scale of 1:10.000 and is used as the primary data source. Image information (RGB/CIR) is primarily used to discern man made objects from natural objects (buildings vs. vegetation).

Data from the digital multi spectral camera *High Resolution Stereo Camera—Airborne, HRSC-A* (Neukum, 1999) is evaluated and used within the Dutch project (Asperen, 1996, Hoffmann et al., 2000). The HRSC-A data set includes high-resolution (15 cm) spectral data (RGB and CIR) and an automatically generated high resolution surface model from stereo matching.

A new change detection project within the framework of EuroSDR is about to start up later this year. The emphasis is on development of methods for localising changes in land cover from very high resolution imagery, the integration of change maps in the updating process and finally comparison of different methods for change detection (EuroSDR, 2004).

## 2 DATA

The change detection procedure presented in section 3 below, is evaluated using datasets mainly associated with the development and updating of the Danish TOP10DK topographical map data base.

### 2.1 RGB images

For the establishment and update of TOP10DK traditional RGB aerial photographs have been used. All images are taken from an altitude of approximately 3800 m leading to a scale of 1:25.000. Each image covers an area of 6 km by 6 km and has a forward lap of 60 percent and a side lap of 20 percent. As part of the production work- o w the photos are scanned at a resolution of 21  $\mu\text{m}$ , leading to 350 MB of data, and a spatial pixel resolution of 0.5 m at ground level. The photos were taken as part of a light campaign in April 2000.

### 2.2 Digital Surface Model (DSM)

As was described by Knudsen and Olsen (2003) it is very difficult to locate changes in the building layer using single aerial images and hence only using spectral information in combination with size and form considerations. Therefore a high resolution digital surface model (DSM) with a grid size of 1 meter covering a test area in Lyngby, north of Copenhagen, has been generated to facilitate the building detection. The dataset was collected and made available for these studies by the Danish engineering and mapping company COWI. Data were collected in May-June 2001 using the TOPOSYS1 system (Toposys, 2004, Baltsavias, 1999) which only record first responses of the pulse. The expected height accuracy is approximately 0.15 m.

### 2.3 Digital Map Database

The building theme from TOP10DK has been selected as target for the update procedure. TOP10DK is a fully 3D map database, including 51 object types (building, lake, highway ...) organised

in 8 classes (traffic, water ...). The precision of the database is better than 1 meter both horizontally and vertically. For change detection in the building layer, only new buildings larger than 25  $\text{m}^2$  and changes of building size larger than 10  $\text{m}^2$  are considered.

## 3 METHOD

The method presented is a revision of a method described by Olsen et al. (2002) and Knudsen and Olsen (2003). It is based on classification principles, using existing object registrations in the map database as training areas in order to determine the characteristics of the different classes used to search for and build the object model. As it is very difficult to generate an unambiguous object model for buildings using only spectral information, the revised method also incorporates height information in the form of high resolution DSM data e.g. from LIDAR or photogrammetric auto-correlation to distinguish between objects in terrain from objects above terrain.

### 3.1 The method step by step

The method which consists of three steps, *preparation, classification, and detection* is outlined in figure 2.

Two major assumptions have to be fulfilled for the change detection procedure to be successful:

- (1) The number of changes in a given class (e.g. building) must be much smaller than the number of objects used to describe the class. This is valid for most urban areas.
- (2) New objects must share the same spectral characteristics as the existing objects used to generate the object model. This is often the case as only a small number of roofing materials is in common use.

**3.1.1 Preparation:** The preparation consists of a *data fusion* step to bring the data sets into a common reference frame and a *preprocessing* step where various enhancement methods are applied to the data data to prepare them for the change detection procedure.

**Data fusion:** as objects from the existing digital map database is to be used as training areas for the determination of the class characteristics, image data (raster) and the map database (vector) must be co-registered. Generally co-registration can be done either by registration of the image data to the map database or by registration of the map database to the image data.

The most used method is registration of image data to the map database. However the method has the disadvantage that most image data types (aerial photos) have to be re-sampled as rectified images or orthophotos. For the data sets to fit completely to each other a high precision elevation model, including description of man made objects (buildings, bridges etc.) must be available (i.e. a Digital Surface Model, DSM).

Another way is to project the map database directly to the other data sets, e.g. onto the aerial images using the basic photogrammetric equations (Kraus, 1993). For this to work a database with  $(X, Y, Z)$  coordinates and the orientation parameters for the aerial images have to be available. This method, leads to the most precise co-registration, and eliminates any resampling of image data.

**Preprocessing:** Various algorithms are applied to the data set to prepare them for the change detection process. The three most important processes are: (1) calculation of NDVI (Normalised

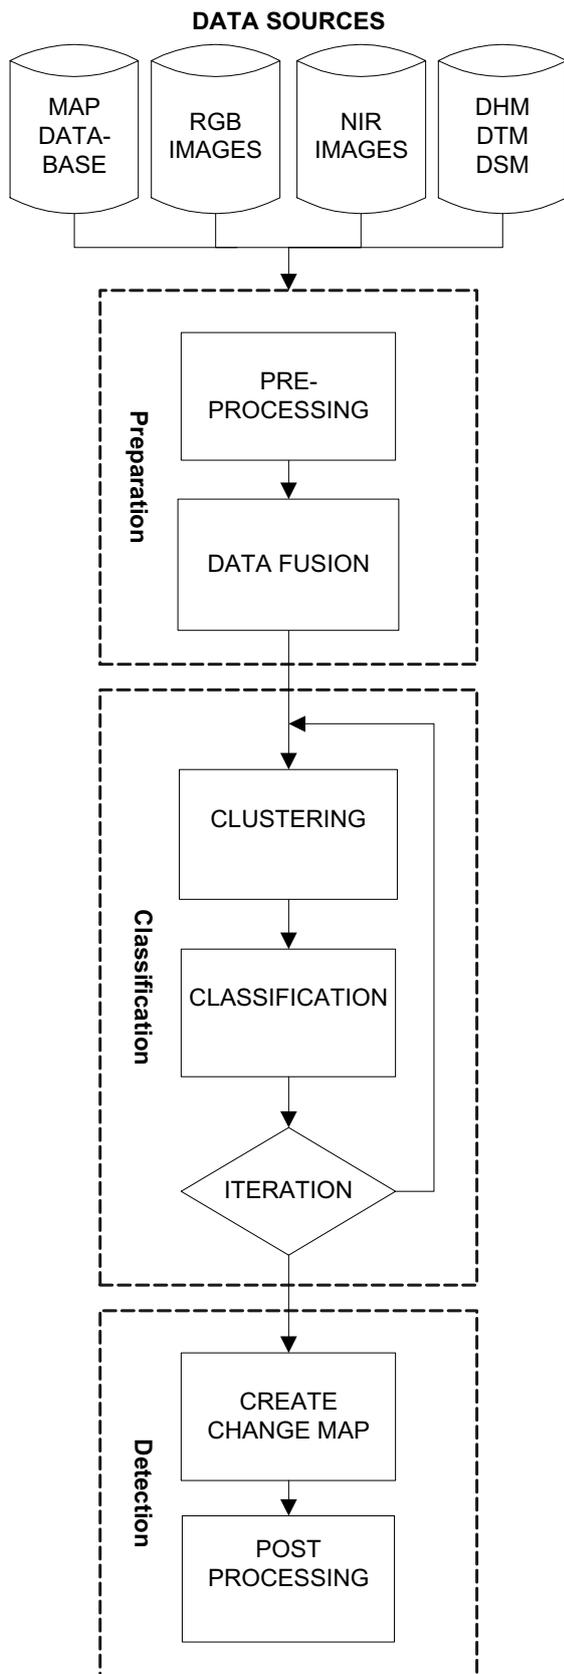


Figure 2: Change detection workflow—cf. section 3 for description

Difference Vegetation Index) images if colour infra-red (CIR) photos are available; (2) generation of a normalized Digital Surface Model (nDSM); and (3) evaluation of training areas.

NDVI is calculated as  $NDVI = \frac{ir-red}{ir+red}$ ; NDVI is well suited for distinguishing vegetated areas from man made objects.

A nDSM only includes objects which stands above terrain and it can be calculated as using a Digital Terrain Model (DTM):  $nDSM = DSM - DTM$ . If a DTM is not available it must be estimated from the DSM. A very simple method for DTM estimation using grey tone morphology is described by Weidner and Förstner (1995), and used in these tests. First a minimum filtering of the DSM is performed using a structuring element  $B$  (with a given size and form). In this way the minimum height in the area determined by the structuring element is assigned to the origin of the structuring element (pixel). This minimum filtering is followed by a maximum filtering, using the same structuring element. Performing the two steps in the described order equals an morphological opening:  $\bar{z} = z \circ B$  and leads to an estimation or approximation of the topographic surface, the DTM. In order to eliminate all elements above terrain (buildings), the size of the structuring element must be chosen in such a way that it is not completely contained in a building. The size depends on the area to be processed. If a priori information concerning existing building sizes in the area is available the size can be fixed using this information. In the test presented in this paper the size of  $B$  is fixed to 25 m. The process is illustrated in figure 3.

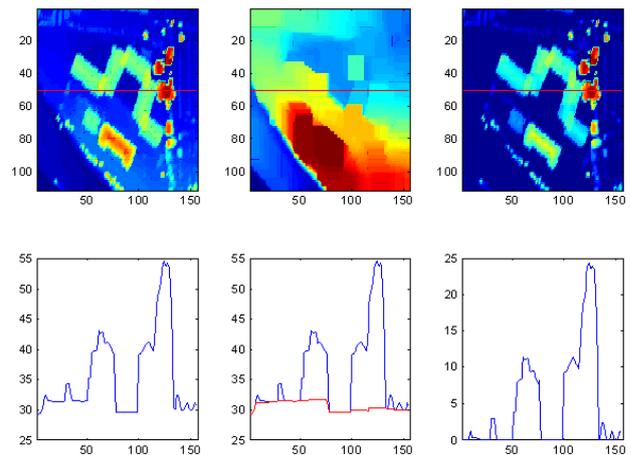


Figure 3: nDSM creation using artificial DTM. UL: DSM. UM: estimated DTM,  $\bar{z} = z \circ B$ . UR: nDSM = DSM - DTM. LL: DSM profile. LM: DTM profile. LR: nDSM profile. All profiles follow the red line in the DSM, DTM and nDSM respectively.

Validation of the training areas is done using the estimated nDSM and/or the NDVI image. An objects above terrain mask can be generated using a height threshold of e.g.  $z \leq 2.5$  meters in the nDSM. With this mask, areas registered as buildings in the existing map database, which no longer stand above terrain are filtered out. Objects covered by vegetation can be eliminated using the NDVI mask (if available), as they can be detected as areas with  $NDVI \geq 0.1$ . The result of the validation is a refined building mask, holding only the buildings which are most likely still buildings.

**3.1.2 Classification:** The first step in the classification part is to perform a clustering process. As stated by e.g. Kressler and Steinbocher (1996) some classes, (e.g. buildings) have to be subdivided into more unique subclasses as they are spectrally highly diverse. This task is handled by splitting up the group of pixels registered as buildings by the building mask into smaller and

more unique sub-classes using a simple migrating means clustering process. The algorithm is based on the ISODATA algorithm (Ball and Hall, 1965), and the number of sub-classes is automatically determined by the algorithm in order to make a best fit to the input dataset.

The clustering process is followed by an actual classification. The sub-classes, which are spectrally more uniform than the base building class, are used (either alone or in combination with other class descriptions (e.g. water, roads, forest, grassland etc.) to perform a Mahalanobis classification of the entire image. This causes all pixels in the image to be assigned to the class having the smallest Mahalanobis distance from the pixel value to the class (Richards and Jia, 1999). Threshold values all being dependent on the class characteristics, are used to assign pixels with a distance too far from the closest class to a garbage class.

The two successive steps are run a number of times as part of an iteration process. This is done mainly due to the fact that the result of the ISODATA algorithm is strongly dependent on the position of the initial cluster centres. After this iteration process it is possible to accept pixels identified as buildings a specific number of times. In the case study presented below all pixels which are classified as a building one or more times are considered to be a "building", leading to an image holding pixels with values of either zero or one: zero indicates no building and one indicates a potential building pixel.

Using the nDSM, the image holding potential buildings are filtered in order to extract only objects (pixels) which stands above terrain.

**3.1.3 Change detection:** First a *change map* is computed by a pixel by pixel comparison of the existing map database (in a raster version) to the classification result. Since the change map includes all potential changes in the building layer it includes noise in the form of single pixels, and some false alarms due to misclassification. The single pixels are removed using morphological opening. The remaining change pixels are segmented, and pixel clusters smaller than the detection requirement (e.g.  $25 \text{ m}^2 = 25 \text{ pixels}$  in the TOP10DK case) and/or not fulfilling the size and shape specifications for buildings are removed from the dataset, leading to a reduction of the false alarms and the final change map.

## 4 CASE STUDY

The procedure is tested on the data described in section 2. The latest update of the TOP10DK database was carried out five years before the photos were taken.

### 4.1 Test area

The test area used for evaluation is situated in Kgs. Lyngby, a suburb 15 km north of Copenhagen. The area contains many different types of buildings and houses since it includes a small industrial area; a cemetery; a church; a small train station; large strip buildings; and a gasoline station. The looks and shapes of the buildings as well as the heights differ a lot. Vegetated areas take up a large part of the area, and since the area also includes a highway, two bridges (one for pedestrians and one for cars) and a rail road, this causes a very special terrain structure. The area is also characterized by the fact that many changes have taken place since the establishment of the TOP10DK database.

*Area:* Approximately  $700\text{m} \times 500\text{m}$ . Lower left corner (E, N) = (718450, 6187050). Upper right corner (E, N) = (719150, 6187550).

(E, N) coordinates are given in UTM zone 32. All images (RGB, TOP10DK and DSM) are 500 rows by 700 columns and approximately 70 buildings are included in the area. 72 registrations are included in the existing map data base. 12 new houses have been build since the last revision and 14 have been demolished.

### 4.2 Test data

All datasets are subsamples from larger datasets and they are brought into the same geographical reference by orthorectification of the aerial photographs using an existing digital terrain model (DTM) with a grid spacing of 20 m.

### 4.3 Results

The results are visualised in figure 4. The first image shows the RGB image with the existing map database superimposed in yellow colour. It can be seen that a lot of development has taken place since the last revision of the map database. This is most pronounced in the right part of the image where 7 buildings have been demolished and 10 new buildings with blue roofing material have been built.

The second image from the top shows the result after the classification step. White pixels indicate potential buildings, and as it can be seen large areas are misclassified as vegetation and roads are classified as potential buildings.

The third image shows the result after the height filtering. It can be seen that all roads are now removed. Some vegetated areas still remain as potential buildings, though.

The last image shows the RGB image with the changes found by the automatic change detection algorithm superimposed in yellow. The results are summarised in table 1.

	Factual	Detected
Demolished Buildings	14	12
New Buildings	12	2
Changes	26	14
False alarms		45

Table 1: Statistical results

Approximately 50 percent of the factual changes in the test area have been detected by the algorithm.

## 5 DISCUSSION

Most success is experienced in the group of demolished buildings where only two demolished buildings have not been detected. This is caused by the method used for change detection where the detected buildings are compared directly to the existing registrations on a pixel wise basis. The two demolished buildings which are not detected are positioned in the upper right area of the test area (marked by a red circle). And as can be seen a new building has been built exactly at the same position as where the two old buildings were positioned. Due to the comparison method neither the new nor the demolished buildings are "highlighted" by the algorithm. Only two new buildings have been detected by the algorithm. The reason for the poor result is that 9 out of the twelve new buildings have completely different spectral responses than the existing buildings in the area, as they are either blue (6 buildings) or still not finished (3). As one of the two hypotheses regarding the change detection procedure are not fulfilled, the algorithm is expected to fail.

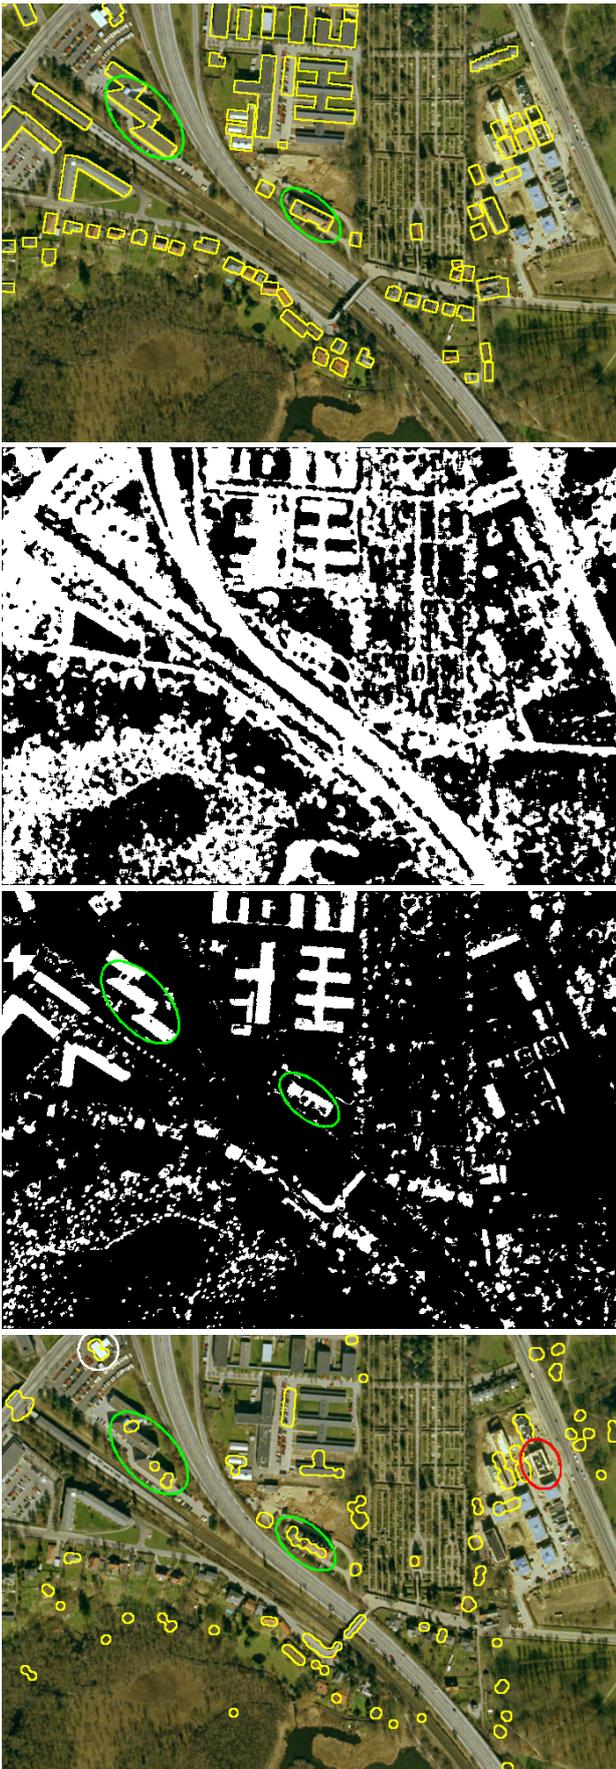


Figure 4: Results from the change detection algorithm.

45 false alarms (3 times the factual changes found) are generated. Of those, 26 are located in vegetated areas, 3 are bridges (above terrain), and 16 caused by existing buildings which apparently have not been re-detected by the algorithm. If infra-red images are available, the majority of the false alarms can be eliminated using the NDVI or by calculation of textural features using the DSM, as it can be expected that the texture for forested areas differs from buildings. The 3 false alarms caused by bridges can only be eliminated by the use of a more “clever” algorithm for nDSM generation. Looking more into the false alarms caused by buildings not re-detected, it can be seen that a large proportion of those buildings are actually detected (figure 4, examples marked by green circles), but these detections are eliminated in a later stage of the change detection algorithm, as part of noise reduction. Refining the noise reduction method may lead to more existing buildings being “re-detected”. One of the false alarms (shown by a white circle in the upper right corner of the area), is a factual difference but *not* a change, since it is a roof covering a gasoline station, and such roofs are not to be registered in the TOP10DK database, according to the map specification. Such false alarms can only be verified by a human operator.

## 6 CONCLUSION

The method presented shows reasonable performance when detecting demolished buildings (12 out of 14 are detected), whereas the number of new buildings detected is poor (only 2 out of 12). A reason for this exists, as the new buildings do not share the spectral response of the existing buildings in the area. One of the hypotheses for the algorithm is not fulfilled. The algorithm introduces a fairly large number of false alarms (3 times the number of factual changes detected). Most of these false alarms can be eliminated by refining some of the processing steps in the algorithm (noise reduction, DSM generation) or by introduction of additional information e.g. infra-red images or textural measures.

**Acknowledgements:** We would like to thank the engineering and mapping company, *COWI*, for letting us use their Digital Surface Model.

## REFERENCES

- Asperen, P. V., 1996. Digital updates of the Dutch topographic service. In: *International Archives of Photogrammetry and Remote Sensing* 31(B4), ISPRS, pp. 891–900.
- AVLBD, 1988. Amtliches topographisch-kartographisches informationssystem (ATKIS). <http://www.adv.online.de/english/products/atkis.htm>.
- Ball, G. H. and Hall, D. J., 1965. A novel method of data analysis and pattern classification. Technical report, Stanford Research Institute, Menlo Park, CA.
- Baltsavias, E. P., 1999. Airborne laser scanner: existing systems and trends and other resources. *ISPRS Journal of Photogrammetry and Remote Sensing* 54(2-3), pp. 164–198.
- Eidenbenz, C., Kaeser, C. and Baltsavias, E., 2000. ATOMI - automated reconstruction of topographic objects from aerial images using vectorized map information. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXXIII (B3), ISPRS, pp. 462–471.
- EuroSDR, 2004. EuroSDR – Change detection project for consideration. <http://www.eurocdr.org/2002/research/ResPoss.asp?ResPosID=27&#pos>.

- Hoffmann, A., van der Vegt, J. W. and Lehmann, F., 2000. Towards automated map updating: is it feasible with new data acquisition and processing techniques? In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXXIII (B2), ISPRS, pp. 295–302.
- Knudsen, T. and Olsen, B. P., 2003. Automated change detection for updates of digital map databases. *Photogrammetric Engineering & Remote Sensing* 69(11), pp. 1289–1296.
- Kort & Matrikelstyrelsen, 2001. TOP10DK. Geometrisk registrering. Technical report, National Survey and Cadastre—Denmark, Rentemestervej 8, 2400 København NV, Denmark.
- Kraus, K., 1993. *Photogrammetry*. Vol. 1, 4 edn, Dümmler, Bonn, Germany.
- Kressler, F. and Steinnocher, K., 1996. Change detection in urban areas using satellite images and spectral mixture analysis. In: *International Archives of Remote Sensing and Photogrammetry*, Vol. 31 number B7, ISPRS, pp. 379–383.
- Neukum, G., 1999. The airborne HRSC—A: Performance results and application potential. *Photogrammerische Woche* pp. 83–88.
- Niederöst, M., 2003. *Detection and Reconstruction of Buildings for Automated Map Updating*. PhD thesis, ETH Zurich, Institut für Geodäsie and Photogrammetrie, Eidgenössische Technische Hochschule, 8093 Zurich.
- Olsen, B. P., Knudsen, T. and Frederiksen, P., 2002. Digital change detection for map database update. In: J. Chen and J. Jiang (eds), *Integrated Systems for Spatial Data Production, Custodian and Decision Support*, Vol. XXXIV (2), ISPRS, pp. 357–363.
- Petzold, B., 2000. Revision of topographic databases by satellite images – experiences and expectations. pp. 15–23.
- Petzold, B. and Walter, V., 1999. Revision of topographic databases by satellite images. In: M. Schroeder, K. Jacobsen, G. Konechy and C. Heipke (eds), *Sensors and mapping from space 1999*, ISPRS, Hanover, Germany, p. 9.
- Richards, J. A. and Jia, X., 1999. *Remote Sensing and Digital Image Analysis: an introduction*. 3 edn, Springer, Berlin, Germany.
- Süweg, I., 2003. *Reconstruction of 3D Building Models from Aerial Images and Maps*. Publications on Geodesy, Netherlands Geodetic Commission, Delft, the Netherlands.
- Toposys, 2004. Toposys1 specification. <http://www.toposys.de/>.
- Walter, V., 2000. Automatic change detection in gis databases based on classification of multispectral data. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXXIII (B4), ISPRS, pp. 1138–1145.
- Walter, V. and Fritsch, D., 2000. Automated revision of gis databases. In: L. K.-J. et al. (ed.), *Proceedings of the Eight ACM Symposium on Advances in Geographic Information Systems*, ACM, pp. 129–134.
- Weidner, U. and Förstner, W., 1995. Towards automatic building extraction from high resolution digital elevation models. *ISPRS Journal of Photogrammetry & Remote Sensing* 50(4), pp. 38–49.

# Spatial-temporal Disambiguation of Multi-modal Descriptors

Norbert Krüger<sup>1</sup>

Florentin Wörgötter<sup>2</sup>

<sup>1</sup> Computer Science  
Aalborg University Esbjerg  
6705 Esbjerg  
nk@cs.aue.auc.dk

<sup>2</sup>Computational Neuroscience  
University of Stirling,  
Stirling FK9 4LA Scotland, UK,  
worgott@cn.stir.ac.uk

## Abstract

In this paper, we describe a new kind of visual representation in terms of local multi-modal Primitives and its role as initiator of disambiguation process based on regularities in visual data. Our Primitives can be characterized by three properties: (1) They describe visual information by different modalities. (2) They are essentially adaptable according to the spatial and temporal context. (3) They give a condensed representation of local image structure that make use of meaningful attributes.

Our Primitives are motivated by human visual processing. They are functional abstractions of hypercolumns in visual cortex [1, 2, 3]. The efficient and generic coding of visual information allows for a wide range of applications. For example, they have been used to investigate the multi-modal character of Gestalt laws in natural scenes [4], to code a multi-modal stereo matching and to investigate the role of different visual modalities for stereo [5] and in a combination of grouping and stereo [6]. In this paper, we use the regularity Rigid Body Motion to acquire reliable 3D information in a spatial-temporal disambiguation process.

## 1 Introduction

The aim of this work is to compute reliable feature maps from natural scenes. We believe that to establish artificial systems that perform reliable actions we need such reliable features. These can only be computed through *integration across the spatial and temporal context and across visual modalities* since local feature extraction is necessarily ambiguous ([7, 8]). The European Project ECOVISION [9] focusses exactly on this issue and the work described here is a central pillar of this ongoing project. In this paper, we describe a new kind of image representation in terms of local multi-modal Primitives (see fig. 1). These Primitives can be characterized by three properties:

**Multi-modality:** Different domains that describe different kind of structures in visual data are well established in human vision and computer vision. For example, a local edge can be analyzed by local feature attributes such as orientation or energy in certain frequency bands. In addition, we can distinguish between line and step-edge like structures (contrast transition). Furthermore, color can be associated to the edge. This image patch also changes in time due to ego-motion or object motion. Therefore

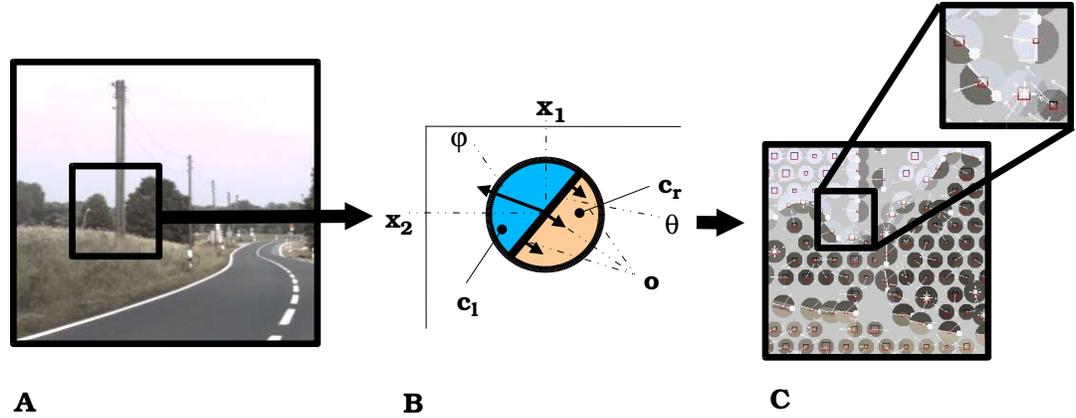


Figure 1: **A:** Image sequence and frame. **B:** Schematic representation of the multi-modal Primitives. **C:** Extracted Primitives.

time specific features such as a 2D velocity vector (optic flow) can be associated to this image patch. Finally by using stereo, information about the 3D origin of an image patch can be computed. In this work we define local multi-modal Primitives that realize these multi-modal relations. The modalities, in addition to the usually applied semantic parameters position and orientation, are contrast transition, color and optic flow (see fig. 1).

**Condensation:** Integration of information requires *communication between Primitives* expressing spatial [4, 5] and temporal dependencies [10]. This communication has necessarily to be paid for with a certain cost. This cost can be reduced by limiting the amount of information transferred from one place to the other, i.e., by reducing the bandwidth. Therefore we are after a *compression* of data. Essentially we only need less than 5% of the amount of the pixel data of a local image patch to code a Primitive that represents such a patch. However, condensation not only means a compression of data since communication and memorization not only require a reduction of information. Moreover, we want to reduce the amount of information within an image patch *while preserving perceptually relevant information*. This leads to *meaningful* descriptors such as our attributes position, orientation, contrast transition, color and optic flow. In [4], we have also shown that these descriptors (in particular when jointly applied) allow for strong mutual prediction that can be related to classical Gestalt laws.

**Adaptability:** Since the interpretation of local image patches in terms of the above mentioned attributes as well as classifications such as ‘edgeness’ or ‘junctionness’ are necessarily ambiguous when based on local processing stable interpretations can only be achieved *through integration* by making use of contextual information [7]. Therefore, all attributes of our Primitives are equipped with a confidence that is essentially *adaptable according to contextual information* expressing the reliability of this at-

tribute. Furthermore, the feature attributes itself adapts according to the context (see section 3). Adaptation occurs by means of recurrent processes in which predictions based on statistical and deterministic regularities disambiguate the locally extracted and therefore necessarily ambiguous data. The instantiation of these processes becomes feasible because of the condensed representation which leads to a manageable amount of meaningful higher order relations of visual events.

In section 2, we describe the Primitive attributes and their extraction. In section 3, we refer to applications of our Primitives for contextual integration where we will especially focus on spatial-temporal disambiguation.

## 2 Multi-modal Primitives

In addition to the position  $\mathbf{x}$ , we compute the following semantic attributes and associate them to our Primitives (see also fig. 1).

**Intrinsic Dimension:** Local patches in natural images can be associated to specific local sub-structures, such as homogeneous patches, edges, corners, or textures. Over the last decades, sub-domains of Computer Vision have extracted and analysed such sub-structures.

The intrinsic dimension (see, e.g., [11]) has proven to be a suitable descriptor that distinguishes such sub-structures. Homogeneous image patches have an intrinsic dimension of zero (i0D); edge-like structures are intrinsically 1-dimensional (i1D) while junctions and most textures have an intrinsic dimension of two (i2D). A continuous definition of intrinsic dimensionality has been given in [12, 13]. There, it has also been shown that *the topological structure of intrinsic dimension essentially has the form of a triangle*. This triangular structure can be used to associate 3 confidences ( $c_{i0D}, c_{i1D}, c_{i2D}$ ) to homogenous-ness, edge-ness, or junction-ness (see also [14]).

This association of confidences to visual attributes is a general design principle in our system. These confidences as well as the attributes themselves are subject to contextual integration via recurrent processes. Aspects with associated low confidences have a minor influence in the recurrent processes or can be disregarded.

**Orientation:** The local orientation associated to the image patch is described by  $\theta$ . The computation of the orientation  $\theta$  is based on a rotation invariant quadrature filter, which is derived from the concept of the *monogenic signal* [15]. Considered in polar coordinates, the monogenic signal performs a *split of identity* [15]: it decomposes an intrinsically one-dimensional signal into intensity information (amplitude), orientation information, and phase information (contrast transition). These features are pointwise mutually orthogonal. The intensity information can be interpreted as an indicator for the likelihood of the presence of a certain structure with a certain orientation and a certain contrast transition (phase, see below). Orientation estimation is further improved by interpolating across the orientation information of the whole image patch to achieve a more reliable estimate.

**Contrast transition:** The contrast transition is coded in the phase  $\varphi$  of the applied filter [15]. The phase codes the local symmetry, for example a bright line on a dark background has phase 0 while a bright/dark edge has phase  $-\pi/2$  (see fig. 21). Therefore,

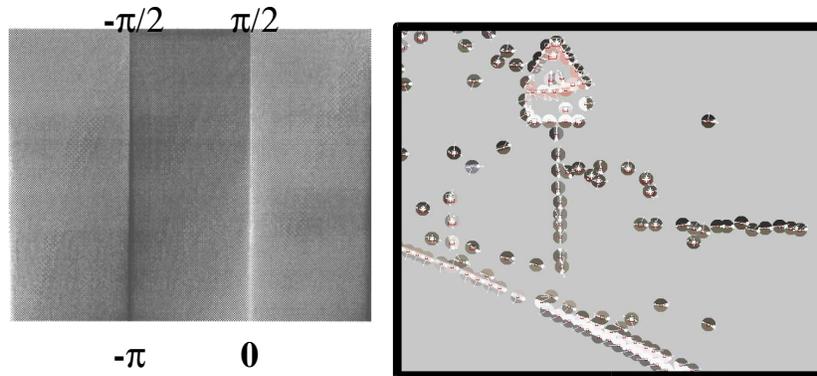


Figure 2: **Left:** Different grey level structures (such as line and step-edge structures) can be associated to different phase values. **Right:** For line-like structures the colour on the line is important which is irrelevant for step edges. Consequently there is a different coding for these different structures: The line indicating the border of the street is first coded as two step-edges and then with greater distance as a line.

the line that marks the border of the street is represented as a line or two edges depending on the distance (see fig. 2II). In case of boundaries of objects phase represents a description of the transition between object and background.

**Color:** Color ( $c^l, c^m, c^r$ ) is processed by integrating over image patches in coincidence with their edge structure (i.e., integrating separately over the left and right side of the edge as well as a middle strip in case of a line structure). In case of a boundary edge of a moving object at least the color at one side of the edge is expected to be stable since (in contrast to the phase) it represents a description of the object. Note that the coding of colour depends on the phase. In case of a line structure (i.e.,  $\varphi \approx 0$  or  $\varphi \approx \pi$ ) besides the colour on the left and right side also a colour for the middle stripe is shown.

**Optic Flow:** Optic flow is associated to a primitive by a vector  $\mathbf{o}$ . There exist a large variety of algorithms that compute the local displacement in image sequences. [16] have them divided into 4 classes: differential techniques, region-based matching, energy based methods and phase-based techniques. After some comparison we decided to use the well-known optic flow technique [17]. This algorithm is a differential technique in which however in addition to the standard gradient constraint equation an anisotropic smoothing term is supposed to lead to better flow estimation at edges (for details see [17, 16, 18]).

**Stereo:** An image patch also describes a certain region of the 3D space and therefore 3D attributes can be associated such as a 3D-position and a 3D-direction (in the following called 3D-Primitives). In [5] we have defined a stereo similarity function that makes use of multiple-modalities to enhance matching performance. We could show

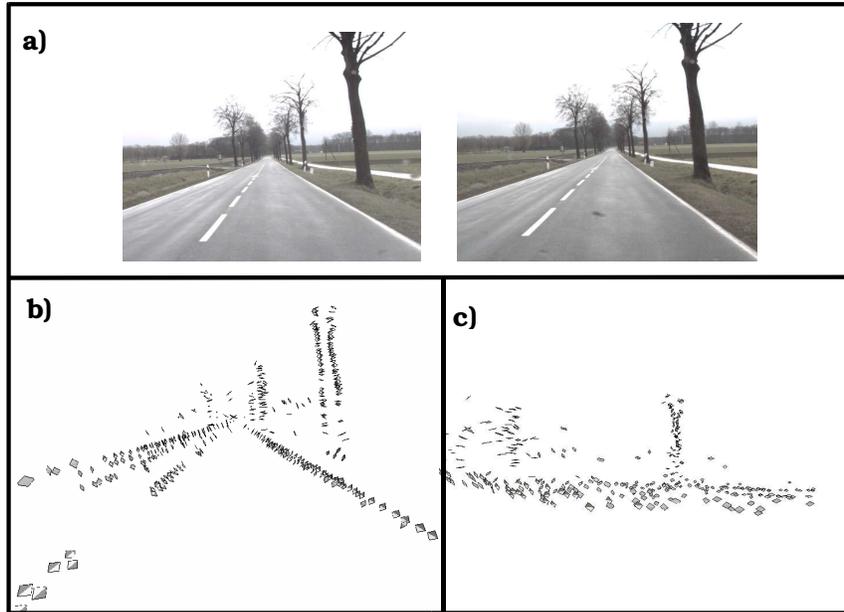


Figure 3: a) Left and right image from stereo frame. b) Front view of extracted spatial primitives (note that this representations has been extracted over multiple frames). The spatial primitives carry beside the actual depth value computable from the disparity also information about 3D orientation, phase and colour. c) Side view of the representation. The structure of the street as well as the two trees are clearly visible.

that it is the joint use of different modalities that gives best results and could compute weights for the importance of the different modalities [5]. The stereo information is coded by the 3D position and the 3D orientation.

Note that the result of the primitive representation is actually not only a disparity map. It is a representation in which the 'symbols' carry beside the depth information also information about other semantic aspects (see figure 3b,c).

We end up with a parametric description of a Primitive as

$$E = (\mathbf{x}, \theta, \varphi, (\mathbf{c}^l, \mathbf{c}^m, \mathbf{c}^r), \mathbf{o}, (c_{i0D}, c_{i1D}, c_{i2D})).$$

In addition, there exist confidences  $c_i, i \in \{\varphi, \mathbf{c}^l, \mathbf{c}^m, \mathbf{c}^r, \mathbf{o}\}$  that code the reliability of the specific sub-aspects that is also subject to contextual adaptation. In addition, information about the underlying 3D-structure is associated to a primitive.

Although a usual image patch that is represented by our Primitives has a dimension of  $3 \times 12 \times 12 = 432$  values (3 color values for each pixel in a  $12 \times 12$  patch), the output of our Primitives has less than 20 parameters. Therefore, the Primitives condense the image information by more than 95%. This condensation is a crucial property of our Primitives that allows to represent meaningful information in a directly accessible and compressed way.

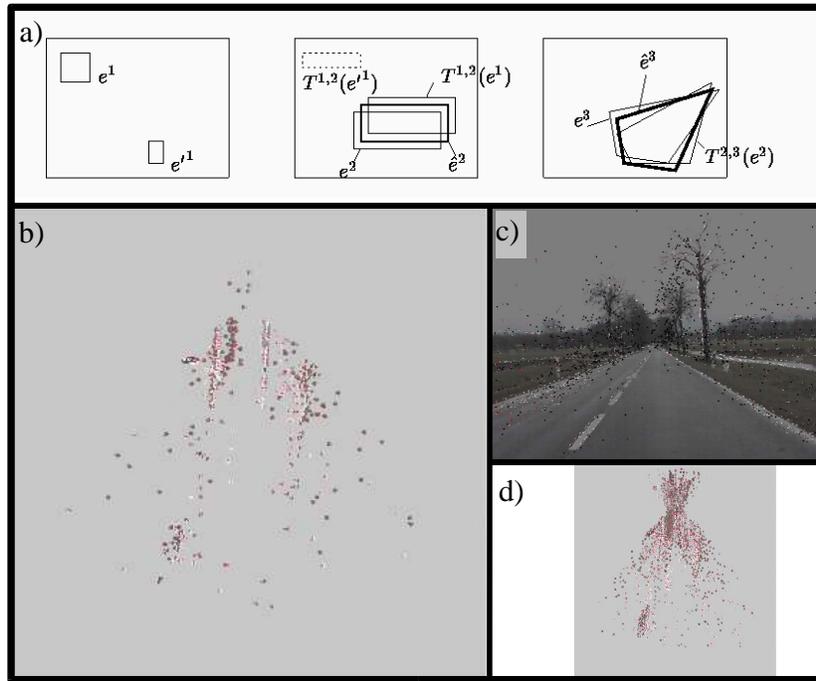


Figure 4: a) Accumulation scheme for spatial-temporal integration. b) Top view of scene in figure 3a after applying the accumulation scheme and elimination of primitives by thresholding over confidences. c) Accumulated representation projected to an image (white line segments represent reliable primitives while dark ones represent unreliable ones). d) Top view as in b) but without thresholding.

The visual modalities coded in our Primitives are processed in early stages of visual processing. Hubel and Wiesel investigated the structure of the first stage of cortical processing that is located in an area called ‘striate cortex’ [1, 2]. The striate cortex is organized in a retinotopic map that has a specific repetitively occurring pattern of substructures called hyper-columns. Hyper-columns themselves contain so called orientation columns and blobs. However, it is not only orientation that is processed in an orientation column but the cells are sensitive to additional attributes such as disparity, color, contrast transition, local motion and disparity (see [19]). Also specific responses to junction-like structures could be measured [20]. Therefore, it is believed that in V1 basic local feature descriptions are processed similar to the feature attributes coded in our Primitives.

### 3 Integration of Contextual Information and Conclusion

It is not only local image processing performed in early visual processing. As mentioned above, there occurs an extensive communication within visual brain areas as well as across these areas. In this communication process the locally necessarily ambiguous data becomes disambiguated in recurrent processes that are based on regularities in visual data. In this sense the Primitives must be understood *as initiators of an disambiguation process* that makes use of contextual knowledge. We will briefly point to the application of the Primitives in such processes now.

We have used the applied image representation in different contexts. Firstly, the Primitives can be subject to a purely spatial contextual modification. We define links between Primitives based on a statistical criterion in [4]. In [6], we apply this linkage structure to improve stereo processing by demanding correspondences that preserve links across groups in the left and right image.

Secondly, we stabilize features according to the temporal context. By making use of the motion of an object to predict feature occurrences across frames we can stabilize stereo processing by modifying the confidences according to the temporal context. For example, in figure 4 the regularity Rigid Body Motion has been used to code predictions of 3D-Primitive occurrences across frames. Schematically, the method depicted in figure 4a) is applied: Assuming a certain occurrence of a primitive in the first frame (4a,i) and having an estimate of the underlying 3D structure we are able, based on the motion between frames, to predict the occurrence of the primitive in the next temporal frame. In case our prediction becomes verified we increase the confidence associated to the primitive (see figure 4b). After a couple of iterations the reliable estimates can be easily detected by their associated confidences. Note that we can also improve the estimate of its parameters by interpolating between our prediction and the actually found feature as schematically shown in figure 4a.

Wrong 3D primitives lead to predictions that can not be verified in consecutive frames which then results in a lowering of the associated confidences (shown in dark colour in figure 4c). After a couple of frames the 3D-Primitives caused by wrong correspondences can be easily detected by their low confidences. While in figure 4d all hypotheses in the process are shown (in a top view of the scene) in figure 4b only Primitives that have been proven to be reliable are displayed. Accordingly, in figure 4, the general structure of the scene is clearly visible.

### 4 Conclusion

We have introduced a novel kind of image representation in terms of visual Primitives. These Primitives are multi-modal and give a dense and meaningful description of a scene. Our Primitives are used as a first stage in the artificial visual system confidences associated to our Primitive attributes adapt according to spatial and temporal context and in this way stabilize the locally unreliable feature extraction.

## References

- [1] D.H. Hubel and T.N. Wiesel, “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex,” *J. Physiology*, vol. 160, pp. 106–154, 1962.
- [2] D.H. Hubel and T.N. Wiesel, “Anatomical demonstration of columns in the monkey striate cortex,” *Nature*, vol. 221, pp. 747–750, 1969.
- [3] N. Krüger, M. Lappe, and F. Wörgötter, “Biologically motivated multi-modal processing of visual primitives,” *The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour*, vol. 1, no. 5, 2004.
- [4] N. Krüger and F. Wörgötter, “Multi modal estimation of collinearity and parallelism in natural image sequences,” *Network: Computation in Neural Systems*, vol. 13, pp. 553–576, 2002.
- [5] N. Krüger and M. Felsberg, “An explicit and compact coding of geometric and structural information applied to stereo matching,” *Pattern Recognition Letters*, vol. 25, no. 8, pp. 849–863, 2004.
- [6] N. Pugeault, F. Wörgötter, and N. Krüger, “A non-local stereo similarity based on collinear groups,” *Fourth International ICSC Symposium on ENGINEERING OF INTELLIGENT SYSTEMS*, 2004.
- [7] Y. Aloimonos and D. Shulman, *Integration of Visual Modules — An extension of the Marr Paradigm*, Academic Press, London, 1989.
- [8] N. Krüger and F. Wörgötter, “Statistical and deterministic regularities: Utilisation of motion and grouping in biological and artificial visual systems,” *Advances in Imaging and Electron Physics*, vol. 131, 2004.
- [9] ECOVISION, “Artificial visual systems based on early-cognitive cortical processing (EU-Project),” <http://www.pspc.dibe.unige.it/ecovision/project.html>, 2003.
- [10] N. Krüger, M. Ackermann, and G. Sommer, “Accumulation of object representations utilizing interaction of robot action and perception,” *Knowledge Based Systems*, vol. 15, pp. 111–118, 2002.
- [11] C. Zetsche and E. Barth, “Fundamental limits of linear filters in the visual processing of two dimensional signals,” *Vision Research*, vol. 30, 1990.
- [12] N. Krüger and M. Felsberg, “A continuous formulation of intrinsic dimension,” *Proceedings of the British Machine Vision Conference*, 2003.
- [13] M. Felsberg and N. Krüger, “A probabilistic definition of intrinsic dimensionality for images,” *Pattern Recognition, 24th DAGM Symposium*, 2003.
- [14] S. Kalkan, D. Calow, M. Felsberg, F. Wörgötter, M.Lappe, and N. Krüger, “Optic flow statistics and intrinsic dimensionality,” *Proceedings of the ‘Early Cognitive Vision Workshop’ on the Isle of Skye (Scotland)*, 2004.

- [15] M. Felsberg and G. Sommer, “The monogenic signal,” *IEEE Transactions on Signal Processing*, vol. 49, no. 12, pp. 3136–3144, December 2001.
- [16] J.L. Barron, D.J. Fleet, and S.S. Beauchemin, “Performance of optical flow techniques,” *International Journal of Computer Vision*, vol. 12, no. 1, pp. 43–77, 1994.
- [17] H.-H. Nagel and W. Enkelmann, “An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, pp. 565–593, 1986.
- [18] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert, “High accuracy optical flow estimates based on a theory for warping,” *Proceedings of the ECCV 2004*, vol. ?, pp. ??, 2004.
- [19] R.H. Wurtz and E.R. Kandel, “Perception of motion, depth and form,” in *Principles of Neural Science (4th edition)*, E.R. Kandel, J.H. Schwartz, and T.M. Messel, Eds., pp. 548–571. 2000.
- [20] I.A. Shevelev, N.A. Lazareva, A.S. Tikhomirov, and G.A. Sharev, “Sensitivity to cross-like figures in the cat striate neurons,” *Neuroscience*, vol. 61, pp. 965–973, 1995.