We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

4,800
Open access books available

122,000
International authors and editors

135M
Downloads

Our authors are among the

154
Countries delivered to

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Computer Vision for Microscopy Applications

Nikita Orlov, Josiah Johnston, Tomasz Macura, Lior Shamir, Ilya Goldberg
Laboratory of Genetics, National Institute on Aging/NIH
USA

1. Introduction

The tremendous growth in digital imagery has introduced the need for accurate image analysis and classification. The applications include content based image retrieval in the World Wide Web and digital libraries (Dong & Yang, 2002; Heidmann, 2005; Smeulders et al., 2000; Veltkamp et al., 2001) scene classification (Huang et al., 2005; Jiebo et al., 2005), face recognition (Jing & Zhang, 2006; Pentland & Choudhury, 2000; Shen & Bai, 2006) and biological and medical image classification (Awate et al., 2006; Boland & Murphy, 2001; Cocosco et al., 2004; Ranzato et al., 2007). Although attracting considerable attention in the past few years, image classification is still considered a challenging problem in machine learning due to the complexity of real-life images. This chapter discusses an approach to computer vision using automated image classification and similarity measurement based on a large set of general image descriptors. Classification results as well as image similarity measurements are presented for several diverse applications.

1.1. Image classification and computer vision

Image analysis can be partitioned into two major approaches. In one, it is assumed that the image is of something that can be modeled a-priori and recognized within the image. This approach uses a model of the subject to drive segmentation followed by extraction of features from the segmented data that correspond to model parameters (size, shape, intensity, distribution, etc). This approach lends itself very well to using imaging for quantitatively measuring defined aspects of a pre-conceived model (Dong & Yang, 2002; Huang et al., 2005; Smeulders et al., 2000). However, a model is not always available, or when available is not always easily reconciled with image data or is not readily useable to extract the relevant subject out of the image (i.e. segmentation). In many applications of image processing, the observed parameters of a model are used to answer questions about the degree to which a particular observation differs from previous observations (i.e. image similarity), or the degree to which an observation agrees with several alternative models (i.e. image classification). Thus an alternative to model-based image analysis for the purposes of computing image similarity or classification is to use pattern recognition and supervised machine learning to answer these questions directly. The focus of this chapter is an approach where a model of what is imaged is built up out of examples consisting of training images rather than an image-independent pre-conceived notion of what is being imaged.

Although the image plane is the carrier of various patterns, the pixels themselves are not normally used directly as inputs to machine learning algorithms. Instead, image content is derived through computation of numerical values that represent quantitative measures of various pixel patterns (Gurevich & Koryabkina, 2006; Heidmann, 2005). These numerical features of the image are based on different algorithms that extract a wide variety of patterns present in the image, such as edges, color (Funt & Finlayson, 1995; Stricker & Orengo, 1995; Tieu & Viola, 2004), textures (Ferro & Warner, 2002; Livens et al., 1996; Smith & Chang, 1994; Smith & Chang, 1996), shapes (Mohanty et al., 2005), histograms (Chapelle et al., 1999; Flickner et al., 1995; Qiu et al., 2004), etc.

Biological microscopy is an emerging application for pattern recognition that presents many diverse problems and image modalitites (Awate et al., 2006; Boland et al., 1998; Boland & Murphy, 2001; Duller et al., 1999; Murphy, 2004; Orlov et al., 2006; Ranzato et al., 2007; Rodenacker & Bengtsson, 2003; Swedlow et al., 2003). When pattern recognition has been used, the tendency is to tailor the image descriptors as well as the classification algorithm to a specific type of imaging problem. Biological microscopy can produce images of many kinds ranging from structural studies of sub-cellular compartments, to the morphology of cells, to tissues and entire organisms. Methods for generating contrast (i.e. the imaging techniques) vary as much as the scale – from fluorescently-tagged protein-specific probes, to various colorimetric stains, to the differential scattering properties of molecular structures. For these reasons, there is no typical imaging problem in biological microscopy and therefore no typical set of image content descriptors. The very nature of the application field requires using a broad variety of algorithms for describing relevant image content (Awate et al., 2006; Boland et al., 1998; Boland & Murphy, 2001; Cocosco et al., 2004; Livens et al., 1996; Ranzato et al., 2007).

A growing demand in pharmaceutical as well as basic research is the use of high-throughput image analysis to score High Content Screens (HCS). In these experiments, a large bank of manipulations (tens of thousands of genes or chemical compounds) is applied one by one to cells grown under defined conditions. Generally the screen is a hunt for genes or compounds that mimic a particular cellular response that can be pre-arranged using positive controls. These screens are typically highly automated using robots for plate and liquid handling, as well as image acquisition. The variety of possible visual assays, combined with the very high demands on the robustness of the processing algorithms makes image analysis in these types of screens the primary bottleneck.

Rodenacker and Bengtsson (Rodenacker & Bengtsson, 2003) have surveyed a large collection of content descriptors for the analysis of grayscale microscopy images. They differentiated feature types into four major categories: intensity, size and shape, texture, and structure. Their suggested scheme for computing signatures includes two pre-processing steps, segmentation (selection of ROIs) and transforms. The use of image transforms is seen as an essential part of feature extraction, where the next-order extraction algorithms (histograms and others) would operate on transforms to produce feature vectors. Many of the feature algorithms given by Rodenacker and Bengtsson could also be used without prior segmentation, and are applicable outside of biological microscopy. The number of descriptors discussed in the paper is quite large, so the authors provide suggestions about which features they found most useful and recommend avoiding textural and structural features for data with strong variation in size and intensity. For feature selection, they recommend hand picking features instead of using independent statistical methods.

Manual feature selection relies on considerable expertise, because of its dependence on the specifics of the experiment as well as the preprocessing steps used.

Lehmann *et al.* (Lehmann et al., 2005) have developed an automated system for categorization and retrieval of images in a medical context. The system includes feature computation and selection as well as classification based on supervised learning using a k-NN algorithm. The set of descriptors they used was limited to Tamura textures as well as several other texture-based descriptors applicable in this domain. The sensitivity of the system was quite good, being able to distinguish 81 distinct categories.

Gurevich and Koryabkina (Gurevich & Koryabkina, 2006) undertook probably the most ambitious survey of existing image descriptors. They developed and adopted from the literature a broad range of features and classified them by scope, method, purpose, etc. While the authors made suggestions of applicability of descriptor types to specific domains, no automated mechanism of feature selection was proposed.

1.2. Digital images: properties and meaning

A given image is not merely an undifferentiated 'bag of features'. The meaning of the image, or the relevant information it contains, can be derived from these features only once their relative importance is determined in a specific context. In supervised machine learning, context is determined by associating a given image with others in a class. The set of classes and the example images they are comprised of defines the context of a specific imaging problem. A given image may be viewed in different contexts by associating it with different groups, which results in a different relative importance of the features, and consequently different interpretations of the image's meaning.

Typical approaches to machine learning emphasize optimizing classification in just one particular problem. Because of this, typical implementations of pattern recognition algorithms only allow for a limited set of descriptors (Awate et al., 2006; Boland & Murphy, 2001; Cocosco et al., 2004; Dong & Yang, 2002; Jing & Zhang, 2006; Ranzato et al., 2007; Rosenfeld, 2001; Shen & Bai, 2006; Smeulders et al., 2000). A limited number of features is desirable because it lowers the computational cost, and reduces the dimensionality of the feature space used in classification. The features selected and their relative weights are problem-specific. The feature set can become inapplicable when new images deviate significantly from those the classifier was trained on, or if they are from a different imaging modality.

A general computer-vision approach requires an alternative to task-specific or manual feature selection (Rosenfeld, 2001). It should use a large feature set in an application-specific context to automatically pick patterns crucial for the given recognition problem (Fig. 1)(Felsenstein, 1989).
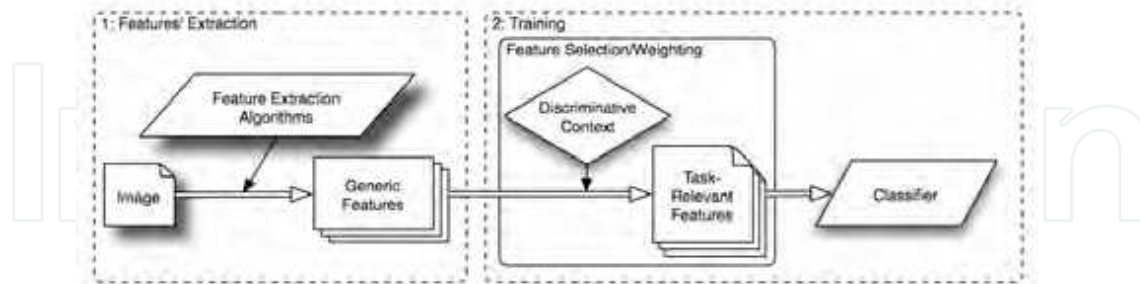
Fig. 1. Image classification scheme. Panel 1: feature extraction; panel 2: feature selection and pattern recognition.

Two antagonistic principles play important roles: context-independent and context-dependent. On the one hand, the system must not ignore details capable of discriminating patterns, which requires having a comprehensive set of context-independent descriptors. On the other hand, irrelevant information (weak features) should be discarded depending on the image context. A general approach to computer vision must balance these two principles. Automating the selection and weighing of features prevents the introduction of anthropogenic bias, which combined with a comprehensive set of descriptors, leads to both generality and objectivity.

Three alternative pathways are available for introducing discriminative context into the initial feature set. The first is to evaluate the discriminative power for all features of the initial set independently of a classifier, keeping features with the highest classification power and discarding the rest. Examples of classifier-independent feature reduction include principal component analysis (PCA) and linear discriminant analysis (LDA, e.g. Fisher discriminant). The second approach is to combine feature selection with classifier building and training, resulting in a feature subset concurrently with the classifier itself. Lastly, there are classifiers capable of performing in high dimensionality feature spaces, essentially by being tolerant of many features with low weights. These include weighted nearest neighbor methods (Parades & Vidal, 2006; Ricci & Acesani, 1999), though these have not been evaluated in feature spaces higher than a few hundred dimensions.

The variety of images available from biological microscopy sets it apart from typical pattern classification problems. This motivated taking a broader look at the principles of image descriptors to measure a wider variety of image content. Because of the generality afforded by addressing the field of biological microscopy as a whole, this approach also proved effective in completely unrelated fields.

1.3. Chapter outline

This chapter describes an approach to compile and recombine traditional and new image analysis algorithms into a general-purpose hyper-dimensional feature set, and the use of an automated feature selection and training method to reduce this feature bank to a context-dependent subset. This effort resulted in a multi-purpose image classifier that can be applied to a variety of image classification problems.

Section 2 introduces the algorithms used in the feature extraction scheme, and Section 3 describes feature reduction and classifier training. Section 4 presents classification results on a set of diverse image types, and Section 5 discusses techniques for computing context-

specific image similarity. Section 6 presents a computing framework used to calculate the features described in Section 2.

2. Extraction of Image Features

Features fall into four categories: polynomial decompositions, textures, object descriptors and pixel statistics. In the first category, a polynomial approximation of the image is computed, and the polynomial coefficients serve as image features. Three kinds of polynomials are computed: Zernike, Chebyshev and Chebyshev-Fourier. Texture features report inter-pixel variation in intensity for several directions and resolutions. These include Haralick textures, as well as Tamura features and Gabor filters. Object features are calculated from one or more object identification algorithms and comprise statistics about object number, spatial distribution, size, shape, etc. Pixel statistics consist of multi-scale intensity histograms, edge statistics, radon histograms and comb-4 moments. In addition to calculating these features on the original image pixels, they are also calculated on several image transforms (Fourier, wavelet, Chebyshev), as well as transform combinations (see Fig. 2).
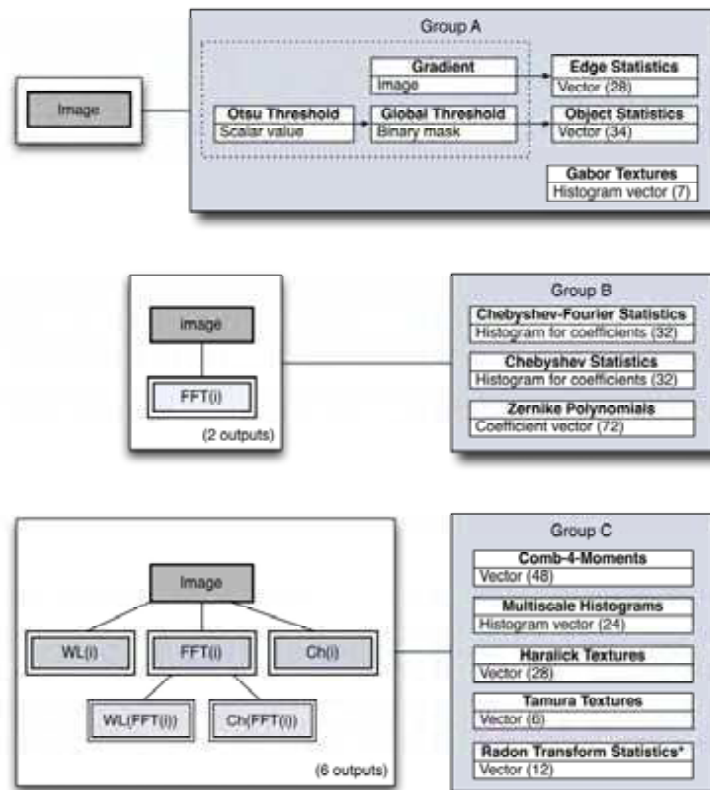

Fig. 2. Proposed scheme for feature extraction. Features are extracted from the original image and several transforms. The Radon Transform Statistics are not computed from the wavelet transform.

Figure 2 illustrates a feature bank (image descriptors) composed of 1025 variables from 11 algorithms, three transforms, and two transform combinations. Each feature measures a different aspect of image content though they cannot be considered strictly orthogonal or independent of each other. All features are based on grayscale images, so that color information is not currently used. Also, though these features cover a broad spectrum of image content, this set cannot be considered to be complete.

2.1. Transforms
In general, transforms allow feature extraction algorithms to be reused to measure very different image content than using them on the original pixels. This algorithm reuse leads to a large expansion of the feature space with a corresponding increase in the variety of image content that can be measured. Each of the three transforms results in a 2D matrix of the same size as the original image pixels; features are computed on image transforms in the same manner as they are on the original images. The Fourier transform is a standard implementation (FFTW), where only the absolute value of the complex transform is used. For the wavelet transform, the standard MATLAB Wavelet toolbox functions 'wavedec2.m' and 'detcoef2.m' were used to compute coefficients for a two-dimensional wavelet decomposition of the image. The Chebyshev transform was implemented by our group and is described in section 2.3 below.

2.2. Pre-processing and color images.
Image pre-processing is a common way to limit noise and improve classification (Hoggar, 2006). Pre-processing is quite common as a prelude for model-based segmentation, but is often unnecessary for the type of scene-based pattern recognition presented in this chapter. All but one of the examples presented in Section 5 were classified without preprocessing. In the example of age-related degeneration of the body-wall muscle, the muscle fibers of the worm's body wall contain significant contributions from the worm's internal structures which were irrelevant for this study. Because the fibers make a regular repeating pattern, a Hamming filter (Hamming, 1989) was utilized to dampen contributions from these less regular structures (see Figure 3).
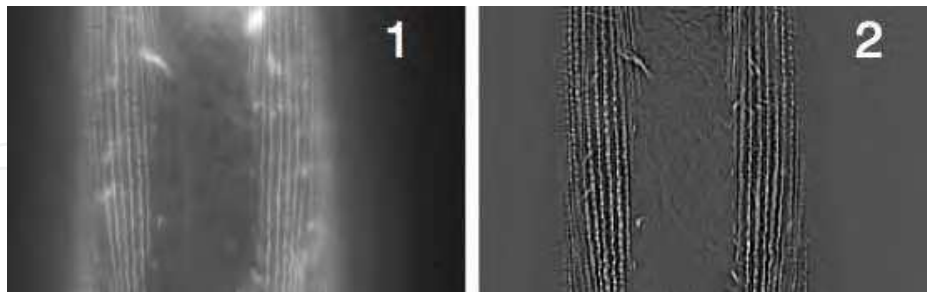

Fig. 3. Image pre-processing. 1: original image (contains irrelevant internal structures); 2: filtered image allows concentrating on morphology of body wall muscle.

Color images also require preprocessing because the feature bank operates on grayscale images. Color information is commonly expressed as several separate color planes in RGB or some other color space. In some cases, the color information is superfluous for classification purposes, and the planes can be combined into a single gray-scale image using

the color's luminance from the NTSC video conversion formula (rgb2gray in MATLAB). In pathology, tissue biopsies are often stained with a pair of dyes called Hematoxylin and Eosin (H&E stain), producing purple cell nuclei with other structures in varying shades of pink. These are normally imaged with an RGB camera, which convolves the H&E channels into an RGB space which is dependent on the camera's spectral response. There are three ways to overcome these difficulties. First, one can compute features on R, G and B channels separately using the existing scheme, with the drawback of treating the 3 channels as independent entities even though each channel is a convolution of H&E. Second, one can use a color deconvolution algorithm (Ruifrok & Johnston, 2001) or similar techniques to approximate the original 2D color space and then use the feature bank on the resulting H&E channels. Lastly, one can introduce feature extraction algorithms that specifically measure color information (e.g. color histograms).

2.3. Chebyshev transform and related features

Chebyshev polynomials (Gradshtein & Ryzhik, 1994)

$$T_n(x) = \cos(n \arccos(x)) \tag{1}$$

are widely used for approximation purposes. For any given smooth function one can generate its Chebyshev approximant, like

$$f(x) \cong \sum_{n=0}^{N} \alpha_n T_n(x). \tag{2}$$

Chebyshev polynomials are orthogonal (with a weight (Gradshtein & Ryzhik, 1994)); therefore, the expansion coefficients α_n can be expressed via the scalar product:

$$\alpha_n = \langle f(x), T_n(x) \rangle. \tag{3}$$

By analogy with Fourier space (formed by the transform coefficients), one can consider the collection of coefficients $\{\alpha_n\}$ as members of some spectral space – the Chebyshev space. Similarly to the 1D case (2), for a given image I_{ij} its two-dimensional approximation through the Chebyshev polynomials is

$$I_{ij} = I(x_i, y_j) \cong \sum_{n,m=0}^{N} \alpha_{nm} T_n(x_i) T_m(y_j). \tag{4}$$

The fast algorithm was used in the transform computation; it takes two sequential 1D transforms, first for rows of the image, then for columns of the resulting matrix (similarly to the implementation of 2D FFT).

The Chebyshev transform is designed to characterize all ranges of the image spectral domain – from low to high frequency features. The idea is to retain a finite number of expansion terms, with the expansion coefficients being used as image descriptors. Chebyshev is used both as a transform (with orders matching image dimensions) and as a set of statistics. The maximum transform order does not exceed N = 20, so that the resulting coefficient vector has dimensions (1x400). The feature vector is produced from the coefficients by applying a 32-bin histogram.

2.4. Features based on Chebyshev-Fourier transform

This 2D transform is defined in polar coordinates, and it uses two different kinds of orthogonal transforms for its two variables: the radial coordinate of the image is approximated with Chebyshev polynomials, and Fourier harmonics are used for the azimuth variable:

$$\Omega_{nm}(r,\phi) = T_n(2r/R - 1) \times \exp(im\phi), \quad 0 \le r \le R. \tag{5}$$

In this sense it shares similarities with Zernike transform where the power polynomials are used in radial direction, and harmonic functions for the angle. For the given image (I_{ij}) the transform generates an image approximant in the form

$$I_{ij} \Rightarrow I(r_k, \phi_l) \cong \sum_{m=-N/2}^{N/2} \sum_{n=0}^{N} \beta_{nm} \, \Omega_{nm}(r_k, \phi_l) \, . \tag{6}$$

In the presented descriptor system (Fig. 2), features based on coefficients β_{nm} of the Chebyshev-Fourier transform capture low-frequency components of the image content (large-scale areas with smooth intensity transitions). The highest order of polynomial used is N = 23, and the coefficient vector is then reduced by binning to 1x32 length.

2.5. Features based on Gabor wavelets

Gabor wavelets (Gabor, 1946) are used to construct spectral filters for segmentation or detection of certain image texture and periodicity characteristics. Gabor filters are based on Gabor wavelets, and the Gabor transform of an image $I(x,y)$ is defined as

$$GT(x,y;f) = \iint_{W_X, W_Y} I(x - w_X, y - w_Y) \, G(w_X, w_Y; f) \, dw_X dw_Y \tag{7}$$

where the kernel $G(w_X, w_Y; f)$ often takes the form (Gregorescu et al., 2002) of a convolution of the Gaussian with the harmonic function:

$$G(w_X, w_Y; f_0) = \exp\{-(X^2 + \gamma Y^2)/2\sigma^2\} \times \exp\{j(f_0 X + \phi)\},$$
$$\begin{cases} X = w_X \cos\theta + w_Y \sin\theta \\ Y = -w_X \sin\theta + w_Y \cos\theta \end{cases} \tag{8}$$

The parameters of the transform are rotation (θ), ellipticity (γ), frequency (f_X, being related to the wavelength) and σ (related to the bandwidth). The parameters γ and σ were chosen to be $\gamma = 0.5, \sigma = 0.56 \times 2\pi/f_X$ (Gregorescu et al., 2002). In the feature bank (see Fig. 2) the Gabor Features (GF) are defined as an area occupied by the Gabor-transformed (GT) image:

$$GF(f_X) = \frac{1}{G_L} \iint_{x,y: GT>0} GT(x,y;f_X) \, dx \, dy \approx \frac{1}{G_L} \sum_{i,j: GT>0} GT_{ij}(f_X), \quad GT_{ij}(f_X) = GT(x_j, y_i; f_X). \tag{9}$$

To minimize frequency bias, these features are computed in a frequency range ($f_X \in \{f_k\}_{k=1}^{K}$), and normalized with the low frequency component $G_L = GF(f_L)$. The frequency values used were $f_L = 0.1$ and $f_X = [1\ 2\ ...\ 7]$.

In the feature bank in Figure 2, the Gabor features belong to a group of textural descriptors and measure image content corresponding to the high and highest spectral frequencies, especially grid-like image textures.

2.6. Radon transform based features

The Radon transform computes a projection of pixel intensity onto a radial line from the image center at a specified angle (radon.m is a built in MATLAB function). The transform is typically used for extracting spatial information where pixels are correlated to a specific direction or angle. The Radon feature computes a series of Radon transforms for angles 0, 45, 90, and 135, and then convolves each transform into a 3-bin histogram; the resulting vector therefore totals 12 entries.

2.7. Multi-scale histograms

This set of features computes histograms with varying numbers of bins (3, 5, 7, and 9) (Hadjidementriou et al., 2001). Each frequency range best corresponds to a different histogram, and thus variable binning allows measuring content in a large frequency range. The maximum number of counts is used to normalize the resulting feature vector, which is 1x24 elements.

2.8. Four-way oriented filters for first four moments

For this set of features, the image is subdivided into a set of stripes in four different orientations (0°, 90°, +45° and -45°). The first four moments (mean, variance, skewness, and kurtosis) are computed for each "stripe", and each set of stripes is sampled as a 3-bin histogram. Four moments in four directions with 3-bin histograms results in a 48-element feature vector.

2.9. Tamura features

Three basic textural properties of an image – contrast, coarseness, and directionality– were proposed by Tamura in 1978 (Tamura et al.). We used these definitions as they were given in Tamura's paper and coded them without modifications. Coarseness gave 4 values (1 for total coarseness and 3 for the histogram), directionality and contrast each contributed 1 entry, totaling 6 features for this group.

2.10. Edge, Zernike and Haralick features

These features were computed as described in (Murphy et al., 2001). Briefly, Edge features measure several statistics on the image's Prewitt gradient. Zernike features are the coefficients of the Zernike polynomial approximation of the image. Haralick features are statistics computed on the image's co-occurrence matrix.

2.11. Object Statistics

Object statistics are calculated from a binary mask of the image resulting from applying a global threshold using Otsu's method. Thirty-four basic statistics about the segmented objects are extracted with MATLAB's 'regionprops.m' function. The statistics include: number of objects, "Euler Number" (the number of objects in the region minus the number of holes in those objects), and image centroid (x and y). Additionally, minimum, maximum, mean, median, variance, and a 10-bin histogram are calculated on both the objects' areas and distances from objects' centroids to the image centroid.

3. Feature Evaluation and Training

Feature extraction is followed by evaluation of each feature's classification power in a given training context. Many classification algorithms such as neural networks, Bayesian belief networks, Markov chain networks, support vector machines, etc, operate in low-dimensional space and neither need nor can function with extraneous or irrelevant features (Bishop, 1996). The problem of mismatch between number of features used for image description and the number useable by these classifiers is often referred to as the 'curse of dimensionality' (Bishop, 1996). Common techniques for reducing dimensionality include Fisher Discriminant and Linear Discriminant Analysis, as well as Principal Component Analysis, Independent Component Analysis, etc. (Bishop, 1996; Fukunada, 1990; Jain & Zongker, 1997; Kudo & Sklansky, 2000; Yang & Wu, 2004; Yu & Yang, 2001). Dimensionality reduction is still considered an active research topic.

Having a large collection of features implies that for every particular classification problem a majority of the features will be sensitive to irrelevant image content and therefore represent noise. Such features unnecessary add to computational complexity and degrade the performance of the classifier when a finite number of training samples are used (Kudo & Sklansky, 2000). Including non-representative features in a classification problem can also lead to over-training with a resulting loss of predictive power.

Dimensionality reduction is a form of hard thresholding, where features below a certain classification power are completely eliminated from subsequent training and classification. An alternative approach (soft thresholding) can be realized with a family of weighting algorithms (Parades & Vidal, 2006; Ricci & Acesani, 1999). The examples presented in Section 4 use a hard thresholding approach which is discussed below.

One way to evaluate the expected usefulness (i.e. discriminative power) of each feature is by computing its ability to separate data between classes while minimizing its within-class variation. This scoring is based on the Fisher linear discriminant which can be formulated as follows (Fukunada, 1990)

$$F = S_B / S_W, \quad S_B = \underset{c=1..C}{\mathrm{mean}} \left(\mu - \mu_c \right)^2, \quad S_W = \underset{c=1..C}{\mathrm{mean}} \ \sigma_c^2 \tag{10}$$

where S_B is the variance of class means from the pooled mean, S_W is the mean of within-class variances, μ_c is the mean of class c, μ is the pooled mean (i.e. mean of all samples), and σ_c^2 is the variance of class c.

The first round of dimensionality reduction calculates the Fisher score separately for each feature, and eliminates 80% of the features with the lowest scores. The second round of

dimensionality reduction tests each feature's discriminative power in a classifier context and iteratively builds the classifier using a greedy-hill climbing algorithm. Naïve Bayesian networks were chosen because they are quickly trained and do not require optimizing the network topology when adding new features to a growing network. In the initial pass, a single-node Bayesian network is constructed from each of the features, and its classification performance and predictive power is scored. To score each network, the initial set of training images is split into a training and test set multiple times, and the classification performance is averaged for all splits. Typically, 35% of the images are reserved for testing in each split, and 35 splits are averaged together for the network score. These parameters have not been systematically calibrated, but proved effective for the datasets evaluated. The feature producing the highest-scoring network is retained for the second pass. In subsequent passes, each of the remaining features is added to the network one by one and the network score is determined. These passes continue until the network score no longer improves by the addition of new features. Generally the number of features selected depends on the number of classes in a classification problem, yielding classifiers with about as many nodes/features as there are classes. This also varies to some extent depending on the separability of the image classes.

4. Experimental Results

In order to evaluate the efficiency of the proposed approach four distinct image data sets were chosen based on application diversity. This small diverse set includes images of fluorescently-labeled cells grown in culture, optical sections of aerogel used to characterize particle traces in the Stardust space probe experiment, fluorescently-labeled muscle tissue, and an organ imaged using differential interference contrast.

The first pair of examples illustrates the effectiveness of the classifier using two very different image types that are particularly challenging for segmentation-based image analysis. The first example (Figure 4) is from an HCS experiment looking for genes that disrupt the formation of fringes or "ruffles" surrounding cells. These ruffles are thought to be important for cell migration, and consequently for the invasiveness of cancer cells as well as developmental abnormalities. The images tested are the control images for the screen, where absence of ruffling is induced by a gene known to have this effect, and these images are contrasted with those of normal cells that display the ruffling. The complete experiment would look for additional genes from a 35,000-gene library that mimic the "no ruffling" morphology (i.e. appear similar to the "no ruffling" control). Various aspects of the images are irrelevant for the purposes of the screen – the cell density and distribution, the overall intensity of the images, etc. These irrelevant variations are included in the set of control images for the screen, and are therefore averaged out in the course of training the classifier. It is not immediately obvious what sorts of image descriptors one would manually chose to differentiate these two classes. Additionally, these images do not appear to be amenable to segmentation-based approaches. Instead, a data driven approach of using a training set to identify relevant features from a large diverse feature set effectively addresses this type of imaging problem (see Table 1).

The second application (Figure 5) involves control images from the Stardust comet dust project. In control experiments, iron grains were shot into aerogel using a dust accelerator. Images were generated by collecting optical sections of aerogel on a microscope. The goal of this classification task is to differentiate images that contain tracks left by dust particles (see

Figure 5, box) from images lacking these tracks. This problem is complicated because all images contain artifacts. While this type of problem can probably be easily addressed with a segmentation algorithm and a discriminator based on the form factor of the segmented objects, this would take significant manual tinkering. However, the same algorithms and parameters used in the previous example worked well without requiring any new software or manual parameter adjustment to work equally well on this very different imaging problem.
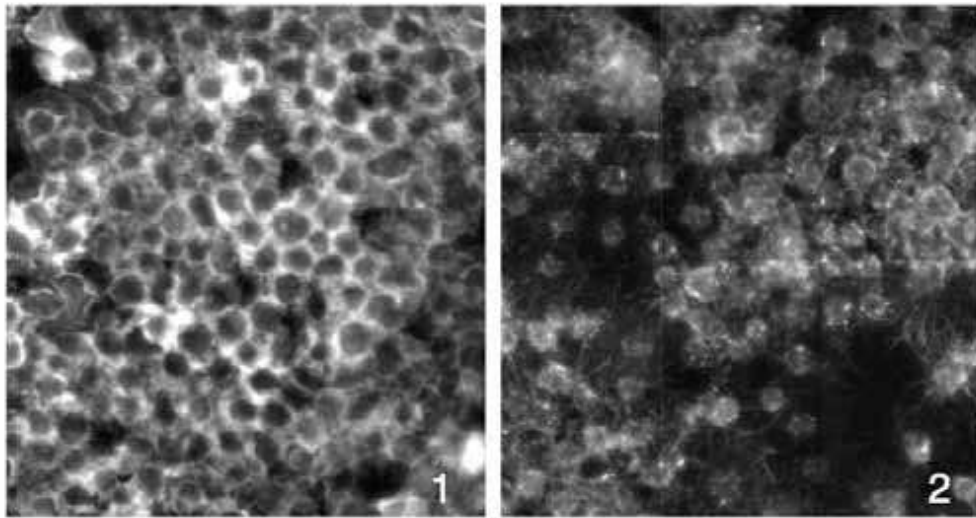
Fig. 4. Images of Ruffling phenotype. Panel 1 shows a field of cells exhibiting the normal ruffling behavior. Panel 2 shows the effect of a knocking down the expression of a gene known to be required for the ruffling phenotype.
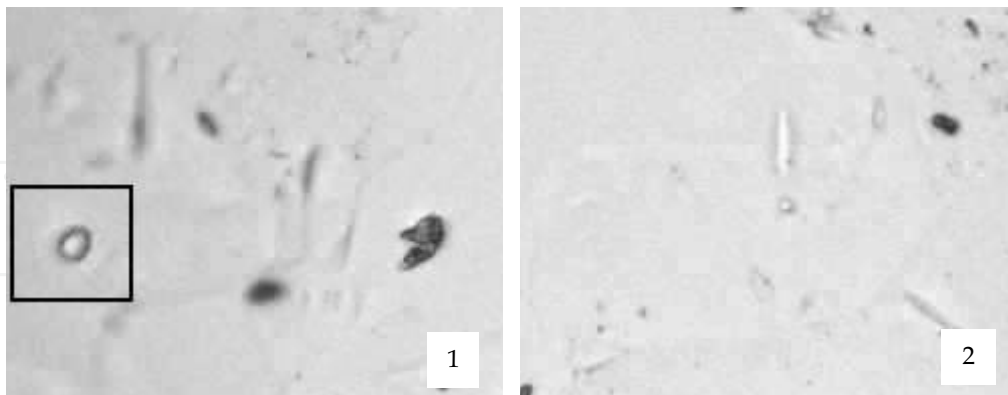
Fig. 5. Images of aerogel from the Stardust project. Panel 1 shows an image containing both a track caused by a dust particle (surrounded by the black frame) and artifacts. Panel 2 shows an image containing only artifacts.

Classification accuracy for these two problems is shown in Table 1. Accuracy is defined as the ratio of correct class assignments to the number of test samples.

Data set	Accuracy
Ruffling cells	0.92
Stardust	0.87

Table 1. Classification results for Ruffling cells and Stardust data sets. These accuracies are averages of 5 separate divisions of training and test images. Ruffling dataset had 616 training and 2648 test images while the Stardust dataset had 2072 training and 1358 test images.

It is of considerable medical interest to identify compounds and genes that affect the aging process. The effects of these compounds or genetic mutations can be quantified if a morphological age can be computed independently of temporal age. *C. elegans* is a small earth-dwelling worm used as a model organism in genetic, development, aging and behavioral studies. Its transparency to visible light, short lifespan, and its well characterized development and genetics make it an ideal organism for aging studies. In the next two studies, the goal is to determine a morphological age based on a microscope image of the worm.
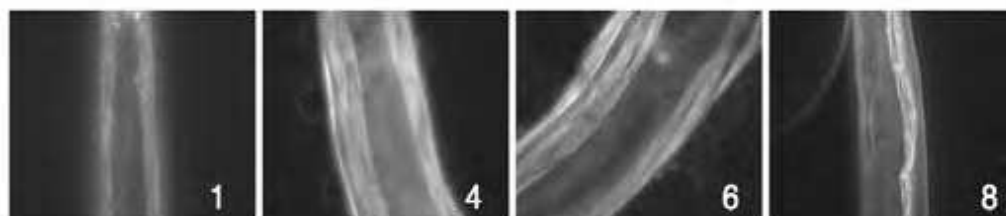

Fig. 6. Aging of body wall muscle in *C. elegans*. Progressive muscle degeneration in the body wall muscle corresponds to 1, 4, 6, and 8 days after molting.
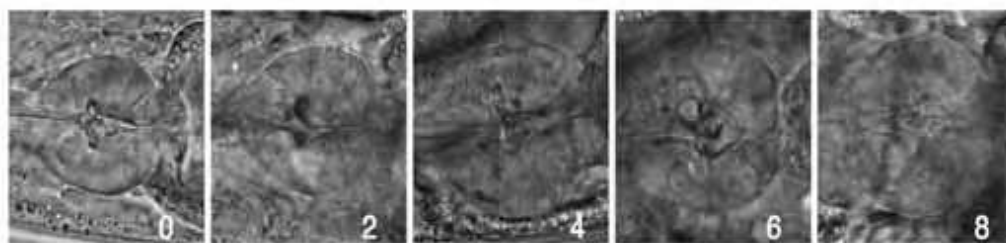

Fig. 7. Aging of pharynx terminal bulb in *C. elegans*. Progressive degeneration of this organ corresponds to 0, 2, 4, 6, and 8 days after molting.

These two studies illustrate a general problem of using images to determine the degree of progression through a continuous morphological process. A distinguishing characteristic of

this type of problem is that there can be considerable uncertainty about the "ground truth" with regard to individual images. On average, the images in a class belong to a known stage along this process, but the exact stage of individual images within the class is far less certain. The individual images in a class may not be precisely synchronized, and the morphological process itself can be subject to substantial stochastic effects. In the specific case presented here, it is known that individual worms can be synchronized at best to within 4 hours. Additionally, individual variability is easily seen in the images collected on a given day, and it is known that a synchronized genetically identical population does not die synchronously, but over a span of 7 to 15 days, which indicates a further loss of synchrony during the process. This supports the prediction that individual images will not be readily classified, and that there will be considerable spill-over between adjacent classes. This would not be a problem in an actual experiment because the treatment would be applied to many individuals, allowing for the use of averaging to evaluate its effect.

These two applications differ in the aging effect being assayed and in the types of images collected. In the first study (see Figure 6), the actin fibers of muscle cells throughout the worm are fluorescently labeled and imaged on a fluorescence microscope. In the second study (see Figure 7), a non-invasive imaging technique that does not require stains or dyes for contrast development (differential interference contrast – DIC) is used to visualize a part of the worm's eating apparatus – the pharynx terminal bulb.

Classification accuracy is not a relevant measure of classifier performance for these two applications because the goal is to measure change rather than class assignment. Instead, the relevant measure is the correlation between the known and computed ages of the classes. The age of each class is computed by averaging the computed ages of its member images. Formula (12) defines an age metric (m) for an individual image in terms of a weighted sum of known class ages (a_c). Weights are the probability (p_c) reported by the Naïve Bayesian classifier of the image belonging to a particular class (c)..

$$m = \sum p_c a_c \qquad (12)$$

The averaged marginal probabilities for each class are reported in Tables 2 and 3 for body wall muscle and pharynx respectively. The corresponding correlation factors were 0.95 and 0.88 for muscle and pharynx images respectively.

Test Data	Day 1	Day 4	Day 6	Day 8
Day 1	0.39	0.27	0.22	0.12
Day 4	0.24	0.47	0.15	0.14
Day 6	0.24	0.15	0.39	0.22
Day 8	0.12	0.15	0.25	0.48

Table2.Confusion matrix of averaged marginal probabilities for body wall muscle.

Test Data	Day 0	Day 2	Day 4	Day 6	Day 8
Day 0	0.95	0.03	0.01	0	0.01
Day 2	0.04	0.48	0.21	0.14	0.13
Day 4	0.01	0.25	0.29	0.22	0.23
Day 6	0.00	0.14	0.24	0.35	0.27
Day 8	0.01	0.16	0.24	0.29	0.30

Table 3. Confusion matrix of averaged marginal probabilities for pharynx.

From inspecting the marginal probabilities in tables 2 and 3, it is clear that there is considerable spill-over to neighbouring classes. However, the probabilities on the diagonal are still strongest, which indicates that accurate classification can still be achieved if a sufficient number of individual images are averaged together.

5. Image Similarity

The last two experiments in the previous section established a baseline of the aging process from which to evaluate effects of experimental manipulations. Aging, being a continuous morphological process, presents a useful means of validating image similarity measures. Simultaneously, image similarity can refine our understanding of aging by answering questions like "Are there distinct stages in this morphological process?" Additionally, reliable image similarity measures are necessary for quantifying class separability, identifying new morphological clusters, or using image queries for content-based image retrieval.

Image similarity is most commonly computed as a distance between vector representations of a pair of images. Each vector defines a point in a high-dimensional space. The feature vectors described in Section 2 form one such space. This raw feature space suffers from the "curse of dimensionality" – a large number of possibly noisy dimensions. This problem can be addressed by constructing a calibrated subspace to emphasize signal and dampen noise.

In such a calibrated space, meaningful distances between images can be determined. Distance matrices can be used to construct dendrograms that visualize population trends. Dendrogram algorithms build a minimal spanning tree from distance information. Branch lengths in a dendrogram indicate similarity between images, and branch angles are inconsequential. Every pair of images is connected by a path; the more similar a pair of images, the shorter the path. In the results given below (Figure 8), a heatmap was applied to represent the known age of each node as a color.
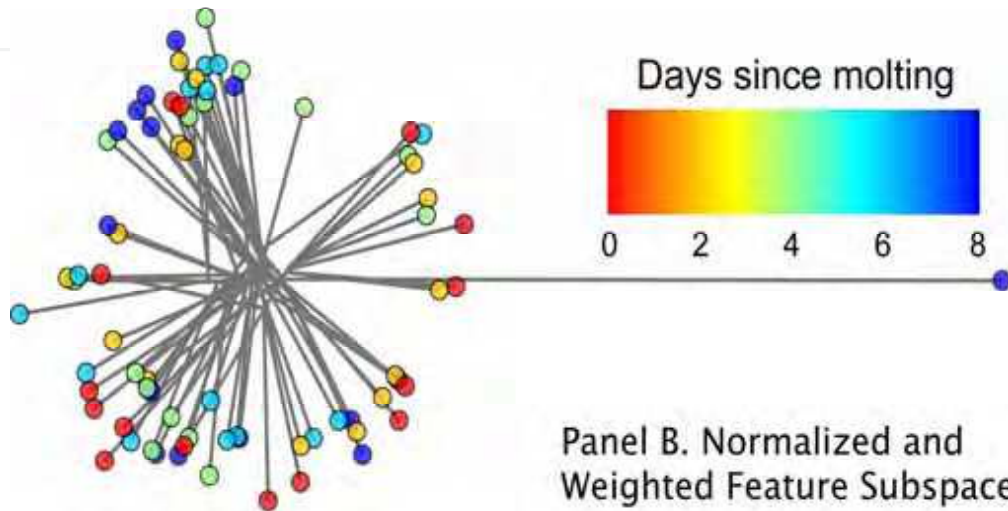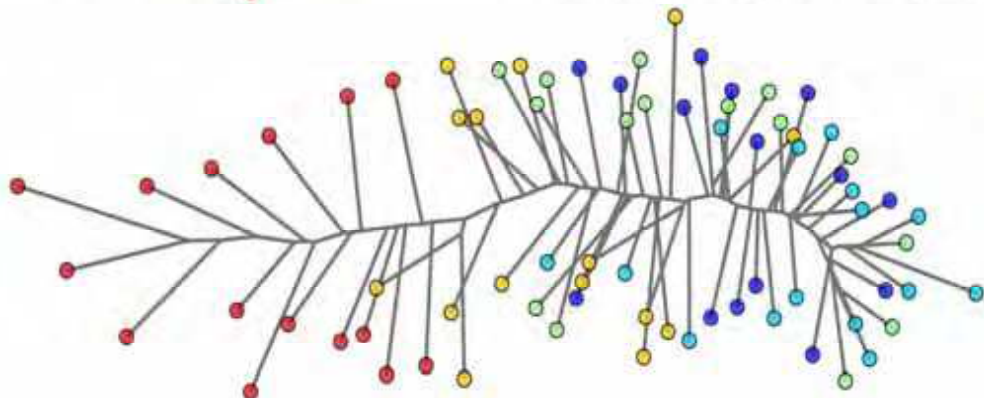
Fig. 8. Individual images of worm pharynx at different ages are represented by colored circles and plotted on dendrograms using distance measures from two types of feature space (see text). The color of each circle represents the known age of each worm in the image. Panel A was constructed from a normalized subspace, and Panel B from a normalized and weighted subspace.

In Figure 8, individual images of worm pharynx (see Figure 7 for image examples) are represented by colored circles in dendrograms computed from two different feature spaces. Normalization of the feature space was accomplished with a linear offset and scaling to a uniform range. Panel B of this figure was constructed from a normalized and weighted

subspace. The subspace was formed by first normalizing then scaling each dimension by its Fisher Discriminant score (see Section 3). Thirty-five percent of the lowest-scoring dimensions were excluded from this space. Dimensional weighting is commonly used in pattern recognition to construct a subspace in which meaningful distances can be determined (Parades & Vidal, 2006; Ricci & Acesani, 1999). Pair-wise Euclidean distance matrices between individual images were computed in these two spaces. The unrooted dendrograms were calculated from the distance matrices with the Fitch-Margoliash method implemented in the Phylip software package (Felsenstein, 1989). The heatmap representing the known age of the worms was applied subsequently.

The normalized full feature space places nearly all images equidistant to each other, which is not meaningful. An exception not shown in the figure is one image whose mean distance to other images was 10,000 times the median value. The FD-weighted feature subspace produces a gradient from young (red, left) to old (blue, right). The FD scores were calculated from unordered class information, yet the subspace yields class ordering. This indicates that the FD-weighted subspace can be used for biologically meaningful measures of similarity.

6. Algorithm Execution Frameworks and the Open Microscopy Environment

Several practical problems emerge when calculating features on large image sets. Full computation of features can take several days for a thousand-image dataset. To avoid recomputation, extracted features need to be stored systematically. To save time when classifying experimental images, it is desirable to calculate only the subset of features identified during dimensionality reduction. Additionally, when a feature bank is extended with more algorithms or permutations, reusing previous results can save considerable amounts of time. A workflow manager and informatics infrastructure addresses these problems better than simply unleashing all of these algorithms on folders full of image files.

The Open Microscopy Environment (OME) provides an image informatics infrastructure which includes facilities for archiving images, meta-data and analysis results. It also provides an Analysis System which executes image-processing algorithms in complex workflows, stores algorithm state and results in a database, and distributes algorithm execution across multiple networked computers. Results from any point in the workflow can be exported out of the OME database for subsequent analysis. The image features presented in Section 2 were computed in OME, and exported into MATLAB. Subsequent dimensionality reduction, signature weighting, dendrograms, and machine learning were performed with other software. Future plans are to integrate this functionality more tightly into OME.

OME is an open source software suite with a thin-client/server architecture (Goldberg et al., 2005; Swedlow et al., 2003). Users access OME using an internet-browser connecting to an extensible, dynamically generated web interface (Johnston et al., 2006). Based around a PostgreSQL database, OME has a middleware layer that provides functionality such as access control, user settings, and image annotation. OME data and computations are performed entirely server-side, optionally using a dedicated cluster. OME is designed to handle gigabytes of high-dimensional microscopy and medical imaging formats, as well as generic TIFF images.

OME algorithm wrappers are called AnalysisModules and are defined using the eXtensible Markup Language (XML) (Achard et al., 2001). AnalysisModule definitions are comprised of two sections: (1) the data modeling section describes the module's name, description,

inputs and outputs and (2) the execution instructions section specifies the interface to the algorithm's implementation (Macura et al., 2005).
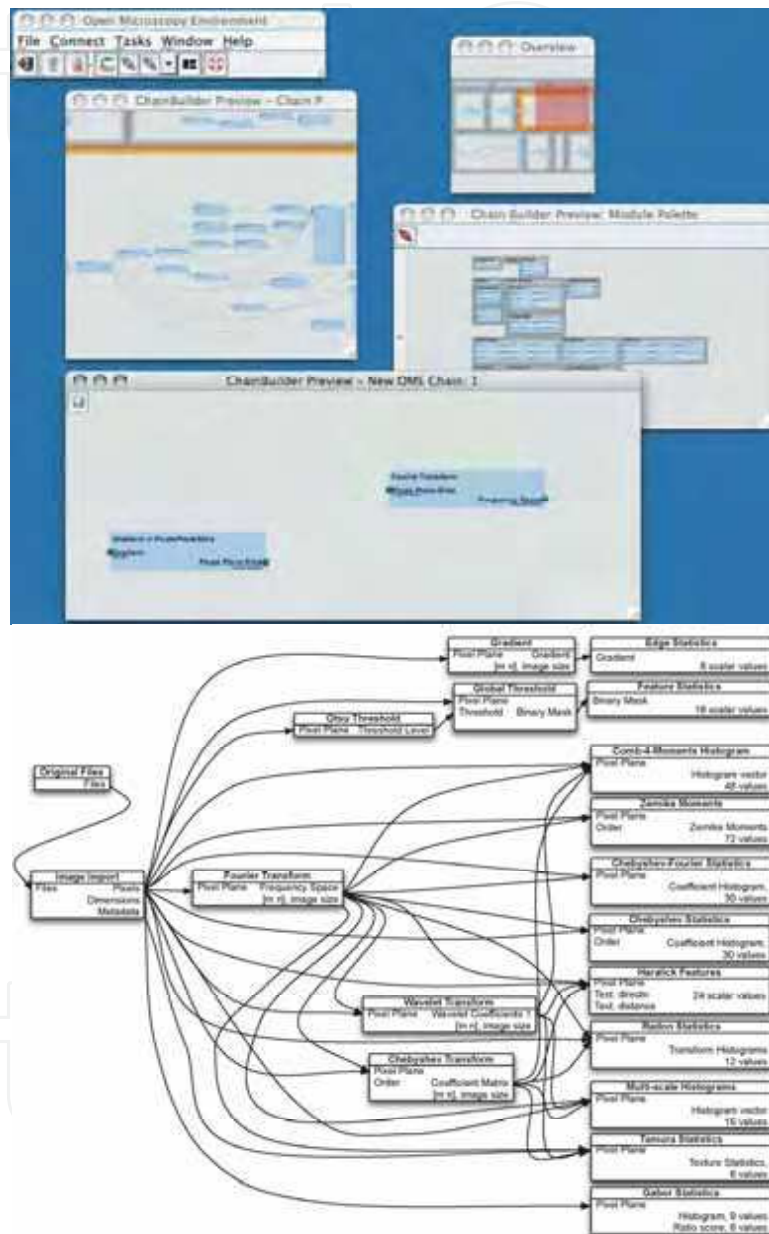

Fig. 9. Top Panel: the ChainBuilder Tool is a user interface for building complex workflows. Bottom Panel: Schematic of the feature extraction workflow (See Fig. 2, Section 2) built using ChainBuilder.

AnalysisModules' inputs and outputs can be connected interactively using the ChainBuilder GUI Tool (See Figure 9), to form workflows called AnalysisChains. After the image processing workflow has been modeled as an analysis chain, the workflow is executed by the OME Analysis System against images managed by OME. The OME Analysis System exploits branching in AnalysisChains to realize that some AnalysisModules can be executed concurrently. This concurrent execution can occur on a local multi-core node as well as remote nodes connected on a network.

OME overhead on a single processor is insignificant (5% of the execution time) and distributed computing is significantly faster (6x using 4 dual-core processors) than executing the algorithms natively in MATLAB. The results from executing algorithms in OME and storing intermediary results in a database and middle-layers agree to native algorithm execution results to the precision expected for 32 bit floating-point representation.

The OME software, implementations of feature extraction algorithms discussed in Section 2, along with AnalysisModule wrappers and AnalysisChains are all available for download from openmicroscopy.org.

7. Summary and Conclusions

This chapter discusses a general computer vision approach to the problem of automatic image classification and similarity measurements. The approach is based on using a large number of image descriptors followed by automated dimensionality reduction and classifier training. Calculation of feature descriptors was implemented as a component of OME, a system for collecting, archiving, annotating and analyzing images and metadata. High dimensional abstract feature sets are needed to allow universal context-independent description of image content. These descriptors in themselves are not sufficient to fully describe image content because their relative importance must first be determined in a given context. Context can be provided by example using training sets and supervised learning to determine which descriptors are important and which are not. This context-dependent weighing of descriptors leads to feature sets that correlate well with independently determined measures of image similarity.

8. Acknowledgements

9. References

Achard, F.; Vaysseix, G. & Barillot, E. (2001). XML, bioinformatics and data integration. *Oxford Bioinformatics,* 17**,** 2, 115-125, 1460-2059.

Awate, S. P.; Tasdizen, T.; Foster, N. & Whitaker, R. T. (2006). Adaptive Markov modeling for mutual-information-based, unsupervised MRI brain-tissue classification. *Medical Image Analysis,* 10, 726-739, 1361-8415.

Bishop, C. (1996) *Neural networks for pattern recognition*. Oxford University Press.

Boland, M.; Markey, M. & Murphy, R. (1998). Automated recognition of patterns characteristic of subsellular structures in florescence microscopy images. *Cytometry,* 33, 366-375, 1097-0320.

Boland, M. V. & Murphy, R. F. (2001). A Neural Network Classifier Capable of Recognizing the Patterns of all Major Subcellular Structures in Fluorescence Microscope Images of HeLa Cells. *Oxford Bioinformatics,* 17**,** 12, 1213-1223, 1460-2059.

Chapelle, O.; Haffner, P. & Vapnik, V. N. (1999). Support vector machines for histogram-based image classification. *IEEE Transactions on Neural Networks,* 10, 1055-1064, 1045-9227.

Cocosco, C. A.; Zijdenbos, A. P. & Evans, A. C. (2004). A fully automatic and robust brain MRI tissue classification method. *Medical Image Analysis,* 7, 5513-527, 1361-8415.

Dong, S. B. & Yang, Y. M. (2002) Hierarchical web image classification by multi-level features. *International Conference on Machine Learning and Cybernetics.* pp.663-668.

Duller, A. W. G.; Duller, G. A. T.; France, I. & Lanb, H. F. (1999). A pollen image database for evaluation of automated identification systems. *Quaternary Newsletter,* 89, 4-9, 0143-2826.

Felsenstein, J. (1989). PHYLIP - Phylogeny Inference Package (Version 3.2). *Cladistics,* 5, 164-166, 0748-3007.

Ferro, C. & Warner, T. A. (2002). Scale and texture in digital image classification. *Photogrammetric engineering and remote sensing,* 68, 51-63, 0099-1112.

Flickner, M.; Sawhney, H.; Niblack, W.; Ashley, J.; Huang, Q.; Dom, B.; Gorkani, M.; Hafner, J.; Lee, D.; Petkovic, D.; Steele, D. & Yanker, P. (1995). Query by image and video content: The QBIC system. *IEEE Computer,* 28, 23-32, 0018-9162.

Fukunada, K. (1990) *Introduction to statistical pattern recognition*. Academic Press, San Diego.

Funt, B. V. & Finlayson, G. D. (1995). Color Constant Color Indexing. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 17, 522-529, 0162-8828.

Gabor, D. (1946). Theory of communication. *Journal Institution of Electrical Engineers,* 93, 429-457,

Goldberg, I. G.; Allan, C.; Burel, J.-M.; Creager, D. A.; Falconi, A.; Hochheiser, H.; Johnston, J.; Mellen, J.; Sorger, P. K. & Swedlow, J. R. (2005). The Open Microscopy Environment (OME) Data Model and XML file: open tools for informatics and quantitative analysis in biological imaging. *Genome Biology,* 6**,** 5, 1465-6914.

Gradshtein, I. S. & Ryzhik, I. M. (1994) *Table of integrals, series and products,* 5 edn. Academic Press, San Diego, NY, Boston, London, Sydney, Tokio, Toronto.

Gregorescu, C.; Petkov, N. & Kruizinga, P. (2002). Comparison of Texture Features Based on Gabor Filters. *IEEE Transactions on Image Processing,* 11, 1160-1167, 1057-7149.

Gurevich, I. B. & Koryabkina, I. V. (2006). Comparative analysis and classification of features for image models. *Pattern Recognition and Image Analysis,* 16, 265-297, 1555-6212.

Hadjidementriou, E.; Grossberg, M. & Nayar, S. (2001) Spatial information in multiresolution histograms. *Computer Vision and Pattern Recognition* pp.702.

Hamming, R. W. (1989) *Digital filters,* 3 edn. Englewood cliffs, NJ: Prentice-Hall.

Heidmann, G. (2005). Unsupervised image categorization. *Image and Vision Computing,* 23, 861-876, 0262-8856.

Hoggar, S. G. (2006) *Mathematics of Image Analysis: Creation, Compression, Restoration, Recognition*. Cambridge University Press, Cambridge.

Huang, J.; Liu, Z. & Wang, Y. (2005). Joint scene classification and segmentation based on hidden Markov model. *IEEE Transactions on Multimedia, 7*, 538-550, 1520-9210.

Jain, A. & Zongker, D. (1997). Feature selection: Evaluation, application, and small sample performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 19*, 153-157, 0162-8828.

Jiebo, L.; Boutell, M.; Gray, R. T. & Brown, C. (2005). Image transform bootstrapping and its applications to semantic scene classification. *IEEE Transactions on Systems, Man and Cybernetics, 35*, 563-570, 0018-9472.

Jing, X. Y. & Zhang, D. (2006). A Face and Palmprint Recognition Approach Based on Discriminant DCT Feature Extraction. *IEEE Transactions on Systems, Man and Cybernetics, 34*, 2405-2414, 0018-9472.

Johnston, J.; Nagaraja, A.; Hochheiser, H. & Goldberg, I. G. (2006). A Flexible Framework for Web Interfaces to Image Databases: Supporting User-Defined Ontologies and Links to External Databases. *Proceedings of 2006 IEEE International Symposium on Biomedical Imaging Proceedings*, Washington D.C., April 2006, IEEE, New York.

Kudo, M. & Sklansky, J. (2000). Comparison of algorithms that select features for pattern classifier. *Pattern Recognition, 33*, 25-41, 0031-3203.

Lehmann, T.; Guld, M.; Deselaers, T.; Keysers, D.; Schubert, H.; Spitzer, K.; Ney, H. & Wein, B. (2005). Automatic categorization of medical images for content-based retrieval and data mining. *Computerized Medical Imaging and Graphics, 29*, 143-155, 0895-6111.

Livens, S.; Scheunders, P.; Wouwer, G.; Dyck, D.; Smets, H.; Winkelmans, J. & Bogaerts, W. (1996). A Texture Analysis Approach to Corrosion Image Classification. *Microscopy Microanalysis Microstructures, 7*, 143-152, 1154-2799.

Macura, T. J.; Johnston, J.; Creager, D. A.; Sorger, P. K. & Goldberg, I. G. (2005) The Open Microscopy Environment Matlab Handler: Combining a BioInformatics Data & Image Repository with a Quantitative Analysis Environment. *http://www.openmicroscopy.org.uk/publications/MatlabHandler.pdf.*

Mohanty, N.; Rath, T. M.; Lea, A.; Manmatha, R.; Lew, M. S.; Chua, T. S.; Ma, W. Y.; Chaisorn, L. & Bakker, E. M. (2005) Lecture Notes in Computer Science. *International Conference on Image and Video Retrieval.* pp.589-598.

Murphy, R. F. (2004). Automated interpretation of protein subcellular location patterns: implications for early detection and assessment. *Annals of the New York Academy of Sciences, 1020*, 124-131, 0077-8923.

Murphy, R. F.; Velliste, M.; Yao, J. & Porreca, G. (2001) Searching Online Journals for Fluorescence Microscopy Images Depicting Protein Subcellular Location Patterns. *2nd IEEE International Symposium on Bio-Informatics and Biomedical Engineering.* pp.119-128. Bethesda, MD.

Orlov, N.; Johnston, J.; Macura, T.; Wolkow, C. & Goldberg, I. (2006) Pattern recognition approaches to compute image similarities: application to age related morphological change. *International Symposium on Biomedical Imaging: From Nano to Macro.* pp.1152-1156. Arlington, VA.

Parades, R. & Vidal, E. (2006). Learning Weighted Metrics to Minimize Nearest-Neighbor Classification Error. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 28*, 7, 1100-1110, 0162-8828.

Pentland, A. & Choudhury, T. (2000). Face recognition for smart environments. *IEEE Computer, 33*, 2, 50-55, 0018-9162.

Qiu, G.; Feng, X. & Fang, J. (2004). Compressing histogram representations for automatic colour photo categorization. *Pattern Recognition,* 37, 2177-2193, 0031-3203.

Ranzato, M.; Taylor, P. E.; House, J. M.; Flagan, R. C.; Lecun, Y. & Perona, P. (2007). Automatic recognition of biological particles in microscopic images. *Pattern Recognition Letters,* 28, 31-39, 0167-8655.

Ricci, F. & Acesani, P. (1999). Data Compression and Local Metrics for Nearest Neighbor Classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 21**,** 4, 380-384, 0162-8828.

Rodenacker, K. & Bengtsson, E. (2003). A feature set for cytometry on digitized microscopic images. *Analytic cellular pathology,* 25, 1-36, 0921-8912.

Rosenfeld, A. (2001). From image analysis to computer vision: an annotated bibliography, 1955-1979. *Computer Vision and Image Understanding,* 84, 298-324, 1077-3142

Ruifrok, A. & Johnston, D. (2001). Quantification of histochemical staining by color deconvolution. *Analytical and Quantitative Cytology and Histology,* 23, 291-299, 0884-6812.

Shen, L. & Bai, L. (2006). MutualBoost learning for selecting Gabor features for face recognition. *Pattern Recognition Letters,* 27, 1758-1767, 0167-8655.

Smeulders, A. W. M.; Worring, M.; Santini, S.; Gupta, A. & Jain, R. (2000). Content-Based Image Retrieval at the End of the Early Years. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 22, 1349-1380, 0162-8828.

Smith, J. R. & Chang, S. F. (1994) Quad-tree segmentation for texture-based image query. *Second Annual ACM Multimedia Conference.* pp.279-286.

Smith, J. R. & Chang, S. F. (1996) Local color and texture extraction and spatial query. *International Conference on Image Processing.* Lausanne, Switzerland.

Stricker, M. A. & Orengo, M. (1995) Similarity of Color Images. *SPIE Storage and Retrieval for Image and Video Databases.* pp.381-392.

Swedlow, J. R.; Goldberg, I.; Brauner, E. & Peter, K. S. (2003). Informatics and Quantitative Analysis in Biological Imaging. *Science,* 3000, 100-102, 0036-8075.

Tamura, H.; Mori, S. & Yamavaki, T. (1978). Textural features corresponding to visual perception. *IEEE Transactions On Systems, Man and Cybernetics,* 8, 460-472, 0018-9472.

Tieu, K. & Viola, P. (2004). Boosting image retrieval. *International Journal of Computer Vision,* 56, 17-36, 1573-1405.

Veltkamp, R.; Burkhardt, H. & Krieger, H. (2001) *State-of-Art in Content-Based Image and Video Retrieval.* Kluwer Academic Publishers, Kluwer.

Yang, Y. & Wu, X. (2004). Parameter tuning for induction algorithm-oriented feature elimination. *IEEE Intelligent Systems,* 19, 40-49, 1541-1672.

Yu, H. & Yang, J. (2001). A direct LDA algorithm for high-dimensional data with application to face recognition. *Pattern Recognition,* 34, 2067-2070, 0031-3203.

Vision Systems: Segmentation and Pattern Recognition

Edited by Goro Obinata and Ashish Dutta

Research in computer vision has exponentially increased in the last two decades due to the availability of cheap cameras and fast processors. This increase has also been accompanied by a blurring of the boundaries between the different applications of vision, making it truly interdisciplinary. In this book we have attempted to put together state-of-the-art research and developments in segmentation and pattern recognition. The first nine chapters on segmentation deal with advanced algorithms and models, and various applications of segmentation in robot path planning, human face tracking, etc. The later chapters are devoted to pattern recognition and covers diverse topics ranging from biological image analysis, remote sensing, text recognition, advanced filter design for data analysis, etc.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Nikita Orlov, Josiah Johnston, Tomasz Macura, Lior Shamir and Ilya Goldberg (2007). Computer Vision for Microscopy Applications, Vision Systems: Segmentation and Pattern Recognition, Goro Obinata and Ashish Dutta (Ed.), ISBN: 978-3-902613-05-9, InTech, Available from:
http://www.intechopen.com/books/vision_systems_segmentation_and_pattern_recognition/computer_vision_for_microscopy_applications

INTECH

open science | open minds