NVPHBV Spring Meeting 2010

Tuesday, May 25, 2010
Berlage-zaal 1
Faculty of Architecture
Delft University of Technology
Julianalaan 134, 2628 BL Delft

hosted by
Pattern Recognition Laboratory
Delft University of Technology
The Netherlands

Program

0930–0950	Reception + coffee
0950–1000	Opening
1000–1025	Fully Automatic Registration and Segmentation for First-Pass Myocardial Perfusion MR Images
1025–1050	Optimizing Hierarchical Classifiers using ROC Analysis
1050–1110	Coffee + tea break
1110–1135	Closed Form of the Steered Elongated Hermite-Gauss Wavelets
1135–1200	The Quest for Natural Intelligence, looking back on 40 years of PHBV
1200–1300	Lunch
1300–1330	Member meeting [ledenvergadering]
1330–1355	Digital Image Forensics
1355–1420	Fast Online Training of Object Detection for Maritime Surveillance
1420–1435	Coffee + tea break
1435–1535	Invited Lecture : Visualisation by Multidimensional Scaling: an Introduction and an Application to Visually Compare Dutch Political Parties
1535–1545	Closing remarks

Abstracts

Fully Automatic Registration and Segmentation for First-Pass Myocardial Perfusion MR Images

Rationale and Objectives: Derivation of diagnostically relevant parameters from first-pass myocardial perfusion MR images involves the tedious and time consuming manual segmentation of the myocardium in a large number of images. To reduce the manual interaction and expedite the perfusion analysis, we propose an automatic registration and segmentation method for the derivation of perfusion linked parameters.

Materials and Methods: A complete automation was accomplished by first registering misaligned images using a method based on Independent Component Analysis (ICA), and then using the registered data to automatically segment the myocardium with Active Appearance Models (AAM). We used 18 perfusion studies (100 images per study) for validation wherein the automatically obtained (AO) contours were compared with expert drawn contours on the basis of point-to-curve error, Dice index, and relative perfusion upslope in the myocardium.

Results: Visual inspection revealed successful segmentation in 15 out of 18 studies. Comparison of the AO contours with expert drawn contours yielded 2.23 ± 0.53 mm and 0.91 ± 0.02 as point-to-curve error and Dice index, respectively. The average difference between manually and automatically obtained relative up-slope parameters was found to be statistically insignificant ($p = 0.37$). Moreover, the analysis time per slice was reduced from 20 minutes (manual) to 1.5 minutes (automatic).

Conclusion: We proposed an automatic method that significantly reduced the time required for analysis of first-pass cardiac MR perfusion images. The robustness and accuracy of the proposed method were demonstrated by the high spatial correspondence and statistically insignificant difference in perfusion parameters, when AO contours were compared with expert drawn contours.

Vikas Gupta[1,2] v.gupta@lumc.nl
Emile A. Hendriks[2]
Julien Milles[1]
Rob J. van der Geest[1]
Michael Jerosch-Herold[3]
Johan H.C. Reiber[1]
Boudewijn P.F. Lelieveldt[1,2]

[1]Division of Image Processing
Department of Radiology
Leiden University Medical Center
The Netherlands

[2]Vision Lab
Delft University of Technology
The Netherlands

[3]3Department of Radiology
Brigham and Women's Hospital
USA

Optimizing Hierarchical Classifiers using ROC Analysis

Pattern recognition problems often exhibit high complexity. A powerful approach to tackle complexity employs problem decomposition. Instead of building one complicated classifier in a large feature space, we may design a hierarchical classifier focusing on separate sub-problems. Each sub-problem classifier may leverage specific features and hence a simpler model.

Although hierarchical classifiers provide much needed simplification, they are difficult to optimize. Receiver Operating Characteristic (ROC) analysis became the standard tool for tuning of trained two- and multi-class classifiers according to performance requirements. The benefits of joint ROC optimization of a two-stage detector/classifier system have been also demonstrated. However, in case of general classifier hierarchies, the designer currently faces a dilemma: Either to decompose the system in parts and optimize each step independently or to build a monolithic classifier.

In this research we are focusing on optimization of decoupled classifier hierarchies where training sets for each node are defined a priori. We propose an algorithm for cost-sensitive optimization of such classifier hierarchies. It allows the designer to view the entire hierarchy of individually-tuned classifiers as a system, and optimize it using the system-wide tools.

Pavel Paclik pavel@prsysdesign.net
Carmen Lai

PR Sys Design
Delft
The Netherlands

Closed Form of the Steered Elongated Hermite-Gauss Wavelets

The wavelet theory plays a central role in image processing. An important complete family of orthogonal wavelets is given by the Hermite-Gauss filters, which are very close to the optimal for feature detection and have interesting causality properties for regularization in the scale-space. Moreover, these filters are steerable, i.e., their rotated versions with every angle can be expressed as linear combinations of fixed bases. This makes them suitable to detect oriented features. These wavelets find many image analysis applications, such as local features detection, texture modeling, and astronomical image compression, just to cite a few.

An intrinsic limit of the Hermite-Gauss filters, pointed out by Canny, is that high noise rejection implies low localization accuracy and vice versa. Such a limit is outrun by considering elongated filters which, however, are no longer steerable. A powerful framework to design steerable approximations of a given profile, which is widely applied in image processing, consists in expanding filters into series of orthogonal functions (single value decomposition). However, closed forms of the resulting bases are hardly available.

We provide a closed form, both in the spatial and in the frequency domain, of a family of wavelets which arise from steering elongated Hermite-Gauss filters. These wavelets have interesting mathematical properties, as they form new dyadic families of eigenfunctions of the 2D Fourier transform, and generalize the well known Laguerre-Gauss harmonics. A special notation introduced here greatly simplifies our proof and unifies the cases of even and odd orders. Applying these wavelets to edge detection increases the performance of about 12.5% with respect to standardmethods, in terms of the Pratt's figure of merit, both for noisy and noise-free input images.

Giuseppe Papari[1] g.papari@rug.nl
P. Campisi[2]
Nicolai Petkov[1]

[1]Johan Bernoulli institute of Mathematics and Computing Science
University of Groningen
The Netherlands

[2]Dipartimento di Elettronica Applicata

Universita' degli Studi di Roma Tre
Italy

The Quest for Natural Intelligence, looking back on 40 years of PHBV

For a motivation on the choice of subject see: Nils J. Nilsson , 2010, 'the quest for Artificial Intelligence', Cambridge. A link with Delft is found in his chapter on Pattern Recognition, "the first Pattern Recognition workshop [of IEEE PAMI] was held in October 1966, A second one was held in Delft, in August 1968." see: http://tab.computer.org/pamitc/. In 1970 I joined the Pattern Recognition Group of the Department of 'Technische Natuurkunde' of the 'Technische Hogeschool', now called TUDelft. The group was headed by prof. Verhagen.

My quest for 'Natural Intelligence' as contrasted with 'Artificial Intelligence' was based mainly on the conviction that the engineering of available knowledge would be much more effective and efficient than running simulated evolution models.

Looking back on 40 years of Pattern Recognition and Image Processing, [PH & BV of NVPHBV], has as a purpose, finding insights worth keeping and passing on. It also helps to keep the size of the personal library to a manageable size and compress lectures on PH and BV.

The following subjects and themes stand out:
- Artificial neural networks, inefficient and not always effective 'learning' algorithms. Keep the understanding of optimisation, borrow algorithms of operations research.
- Genetic programming, annealing , keep the understanding of finding local optima using Monte Carlo type of algorithms.
- Fuzzy logic, keep the understanding that logic is an extreme (artificial) case of probability.
- False analogies and wrong applications of models: optical sensors are basically photon counters. Understanding this avoids blunders such as application of Gaussian models instead of Poisson type models, use quantum models instead of 'energy' or 'entropy'.
- Principal Components, KL transforms, Hadamard Walsh transforms, are handy for data compression but eigenvalues of covariance should not be confused with information content.
- Careless use of the name of Bayes [Bayesian Artificial Intelligence]. Keep the understanding that in mage analysis the 'priors' are proportional to the occurrence of 'pixels' in an arbitrary ROI.
- The Bayes relation between Class Ci and Measurement Xj is a generalisation of normalising the combined frequency Freq(Ci, Xj) over rows or columns.
- Natural Intelligence is usually represented in physical models. Pattern Recognition is used to find the most likely models with the most likely parameters. In areas with weak models it is useful to systematically collect local / conditional models or cases using techniques of Knowledge engineering.

Nanno J. Mulder nanno.j.mulder@misc.utwente.nl

M3X
Enschede
The Netherlands

Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente
The Netherlands

Digital Image Forensics

Along with the rapid replacement of analog cameras with their digital counterparts in the beginning of this century came the rapid distribution of digital images (and videos). Also, images can be easily manipulated without leaving visual traces. Hence, the authenticity of the image can often be questioned. In digital image forensics, we broadly want to answer two questions. The first question pertains to whether the image (photograph) has been manipulated or not. The answer to this question is strictly necessary to answer the second question: which specific camera was used to make the photograph? Both questions are not trivial. In general, it can be very hard to detect whether an image has been manipulated. Certain special cases (e.g. minimally compressed, minimally resized, only certain types of manipulation) can be well detected, but in general with the presence of strong compression the results become unreliable. Methods implemented from the literature often handle compression as an afterthought, and large-scale performance evaluations are scarce or non-existent. This also severely limits the ability to report an objective measure of certainty. The second question can be answered with the help of the Photo Response Non-Uniformity (PRNU). This noise-like sensor fingerprint is present (to a certain extent) in all images that are produced by each camera. By comparing the pattern from a questioned image with the pattern from a set of reference images made with a suspect camera, we can answer whether the questioned image was produced with the suspect camera or not. This works when the image is authentic, but fails when the image underwent spatial transformations (e.g. rotation, shearing, resizing) because the 'fingerprint' is desynchronized. Other approaches often involve timeconsuming feature extraction and SVM training, and only works on the classification level.

Wiger van Houten w.van.houten@nfi.minjus.nl

Netherlands Forensic Institute
The Netherlands

Fast Online Training of Object Detection for Maritime Surveillance

Typically, radar is used for the detection and tracking of ships for traffic management. This technique is quite expensive and cannot detect small ships reliably. Within the Dutch WATERVisie project, we use commercially available cameras to support the detection task. We apply a generic object detection framework to the task of ship detection and present first results.

Training of complex classifiers for object detection is typically very computationally expensive because many training samples are considered and the optimization requires many iterations over the complete training set. We propose the use of the Stochastic Gradient Descent (SGD) learning algorithm that considers each training sample only once and can update the final solution in a stream-like fashion. We have incorporated SGD in an object detection framework and show detection results comparable to state-ofthe- art Support Vector Machine (SVM) classifiers, while reducing the computational complexity by two to three orders of magnitude.

We apply the object detection framework for the application of waterway surveillance. The system is trained to detect ships in video from a PTZ camera scanning a canal in the port of Rotterdam. Because of the enormous variations in the overall shape of the ships, we have annotated the cabin of each ship. We show that the system can detect ships with high recall with a limited number of false detections. Learning with only a few images of ships is sufficient to already obtain a reasonable detection performance.

Rob Wijnhoven[1,2] rob.wijnhoven@vinotion.nl / r.g.j.wijnhoven@tue.nl
Peter H.N. de With[2]

[1]ViNotion
The Netherlands

[2]Technische Universiteit Eindhoven
The Netherlands

Invited Lecture : Visualisation by Multidimensional Scaling: an Introduction and an Application to Visually Compare Dutch Political Parties

Multidimensional scaling (MDS) is a graphical method that represents objects as points in a low dimensional space such that each distance between a pair of points matches the (dis)similarity between the objects as closely as possible. Te strength of the technique lies in its simplicity: nearby points are similar, those far apart are different. In this presentation, we introduce MDS and discuss several properties and possibilities through applications on empirical data. Amongst these applications is one that compares the election programs of Dutch political parties and another one that allows MDS for large scale data sets.

Patrick Groenen groenen@ese.eur.nl

Erasmus School of Economics
Erasmus University Rotterdam
The Netherlands