Project acronym	SIMBAD
Project full title	Beyond Features: Similarity-Based Pattern Analysis and Recognition
Deliverable Responsible	Dipartimento di Informatica Università Ca' Foscari di Venezia Via Torino 155 30172 Venezia Mestre Italy http://www.dsi.unive.it/~pelillo
Project web site	http://simbad-fp7.eu
EC project officer	Teresa De Martino
Document title	Second periodic Report
Deliverable	D1.3
Document type	Report
Dissemination level	Public
Contractual date of delivery	M 25
Project reference number	213250
Status & version	Definitive
Work package	WP 1
Deliverable responsible	UNIVE
Contributing Partners	UNIYORK, TUD, IST, UNIVR, ETH Zurich
Author(s)	Marcello Pelillo
Additional contributor(s)	Veronica Giove

Declaration by the scientific representative of the project coordinator

I, as scientific representative of the coordinator of this project and in line with the obligations as stated in Article II.2.3 of the Grant Agreement declare that:

- The attached periodic report represents an accurate description of the work carried out in this project for this reporting period;
- The project (tick as appropriate):

  ☐ has fully achieved its objectives and technical goals for the period;

  ☒ has achieved most of its objectives and technical goals for the period with relatively minor deviations;

  ☐ has failed to achieve critical objectives and/or is not at all on schedule.

- The public website is up to date.
- To my best knowledge, the financial statements which are being submitted as part of this report are in line with the actual work carried out and are consistent with the report on the resources used for the project (section 6).
- All beneficiaries, in particular non-profit public bodies, secondary and higher education establishments, research organisations and SMEs, have declared to have verified their legal status. Any changes have been reported under section 5 (Project Management) in accordance with Article II.3.f of the Grant Agreement.

Name of scientific representative of the Coordinator: Marcello Pelillo [signature]

Date: 20 / 05 / 2010

Signature of scientific representative of the Coordinator: [signature]

PROJECT PERIODIC REPORT

Grant Agreement number: 213250

Project acronym: SIMBAD

Project title: Beyond Features: Similarity-Based Pattern Analysis and Recognition

Funding Scheme: Collaborative Project

Date of latest version of Annex I against which the assessment will be made: 01/02/2008

Periodic report: 1st ☐ 2nd **X** 3rd ☐ 4th ☐

Period covered: from 01/04/2009 to 31/03/2010

Name, title and organisation of the scientific representative of the project's coordinator: Marcello Pelillo, Associate Professor, Computer Science Department, Università Ca' Foscari di Venezia

Tel: +39 041 234 8440

Fax: +39 041 234 8419

E-mail: E-mail: pelillo@dsi.unive.it

Project website address: http://simbad-fp7.eu

1. *Publishable summary*

The field of pattern recognition (or machine learning) is concerned with the automatic discovery of regularities in data through the use of computer algorithms, and with the use of these regularities to take actions such as classifying data into different categories, with a view to endow artificial systems with the ability to improve their own performance in the light of new external stimuli. This ability is instrumental in building next-generation artificial cognitive systems (ACS's), namely systems that can perceive, reason and interact robustly in open-ended environments. The socio-economic implications of this scientific endeavour are enormous, as ACS's will have applications in a wide variety of real-world scenarios ranging from industrial manufacturing to vehicle control and traffic safety, to remote and on-site (environmental) sensing and monitoring, and to medical diagnostics and therapeutics.

This project aims at bringing to full maturation a paradigm shift that is currently just emerging within the pattern recognition and machine learning domains, where researchers are becoming increasingly aware of the importance of similarity information *per se*, as opposed to the classical feature-based (or vectorial) approach. Indeed, the notion of similarity has long been recognized to lie at the very heart of human cognitive processes and can be considered as a connection between perception and higher-level knowledge, a crucial factor in the process of human recognition and categorization.

Traditional pattern recognition techniques are centered around the notion of “feature”. According to this view, each object is described in terms of a vector of numerical attributes and is therefore mapped to a point in a Euclidean (geometric) vector space so that the distances between the points reflect the observed (dis)similarities between the respective objects. This kind of representation is attractive because geometric spaces offer powerful analytical as well as computational tools that are simply not available in other representations. Indeed, classical pattern recognition methods are tightly related to geometrical concepts and numerous powerful tools have been developed during the last few decades, starting from linear discriminant analysis in the 1920’s, to perceptrons in the 1960’s, to kernel machines in the 1990’s.

The geometric approach suffers from a major intrinsic limitation, which concerns the representational power of vectorial, feature-based descriptions. In fact, there are numerous application domains where either it is not possible to find satisfactory features or they are inefficient for learning purposes. This is typically the case when experts cannot define features in a straightforward way, when data are high dimensional, when features consist of both numerical and categorical variables, and in the presence of missing or inhomogeneous data. But, probably, this situation arises most commonly when objects are described in terms of structural properties, such as parts and relations between parts, as is the case in shape recognition. This led in 1960’s to the development of the structural pattern recognition approach, which uses symbolic data structures, such as strings, trees, and graphs for the representation of individual patterns, thereby, reformulating the recognition problem as a pattern-matching problem.

In the last few years, interest around purely similarity-based techniques has grown considerably. For example, within the supervised learning paradigm (where expert-labeled training data is assumed to be available) the now famous “kernel trick” shifts the focus from the choice of an appropriate set of features to the choice of a suitable kernel, which is related to object similarities. However, this shift of focus is only partial, as the classical interpretation of the notion of a kernel is that it provides an implicit transformation of the feature space rather than a purely similarity-based representation. Similarly, in the unsupervised domain, there has been an increasing interest around pairwise algorithms, such as spectral and graph-theoretic clustering methods, which avoid the use of features altogether.

Despite its potential, presently the similarity-based approach is far from seriously challenging the traditional paradigm. This is due mainly to the sporadicity and heterogeneity of the techniques proposed so far and the lack of a unifying perspective. On the other hand, classical approaches are inherently unable to deal satisfactorily with the complexity and richness arising in many real-world situations. This state of affairs hinders the application of machine learning techniques to a whole variety of relevant, real-world problems. Hence, progress in similarity-based approaches will surely be beneficial for machine learning as a whole and, consequently, for the long-term enterprise of building ACS’s. However, by departing from vector-space representations one is confronted with the challenging problem of dealing with (dis)similarities that do not necessarily possess the Euclidean behavior or not even obey the requirements of a metric. The lack of the Euclidean and/or metric properties undermines the very foundations of traditional pattern recognition theories and algorithms, and poses totally new theoretical/computational questions and challenges that we are addressing with this project.

With this project, we are undertaking a thorough study of several aspects of similarity-based pattern analysis and recognition methods, from the theoretical, computational, and applicative perspective, with a view to substantially advance the state of the art in the field, and contribute towards the long-term goal of organizing this emerging field into a more coherent whole. The whole project revolves around two main themes, which basically correspond to the two fundamental questions that arise when abandoning the realm of vectorial, feature-based representations, namely:

- How can one *obtain* suitable similarity information from object representations that are more powerful than, or simply different from, the vectorial?
- How can one *use* similarity information in order to perform learning and classification tasks?

Although the two issues are clearly interrelated, it is advantageous to keep them apart as this allows one to separate the *similarity generation* process (a data modeling issue) from the *learning and classification* processes (a task modeling issue). According to this perspective, the very notion of similarity becomes the pivot of non-vectorial pattern recognition in much the same way as the notion of feature-vector plays the role of the pivot in the classical (geometric) paradigm. This results in a useful modularity, which means that all interactions between the object representation and the learning algorithm are mediated by the similarities, which is where the domain knowledge comes into the scene.

During the second year of the project we continued the work on the above fundamental questions started in the first year. As for the first question, we aimed at devising suitable similarity measures for non-vectorial data, specifically tailored to a given task. We have focussed primarily on structured data (e.g., graphs), because of their expressive power and ubiquity, and on geometric measures as they allow one to employ the whole arsenal of powerful techniques available in the (classical) pattern recognition literature. We have also explored an alternative to this “tailoring” approach, which consists in learning similarities directly from training data.

As concerns the second question, we have both addressed foundational issues related to similarity information and developed practical similarity-based algorithms that do not depend on the actual object representation. In particular, as concerns the latter objective, we have distinguished between the situation where the informational content associated with the violation of the geometric properties is limited, or is simply an artifact of the measurement process, and that where this is not the case. This distinction is important as, depending on the actual situation, two complementary strategies will be pursued: the first attempts to impose geometricity by somehow transforming or re-interpreting the similarity data, the second does not and works directly on the original similarities. Our theoretical analysis is expected to provide means to distinguish between these cases, thereby allowing one to understand which of the two approaches is more appropriate for the problem at hand.

Apart from the potential applications mentioned before, in the second year of the project we devoted substantial effort towards tackling two large-scale biomedical imaging applications. With the direct involvement of leading pathologists and neuroscientists from the University Hospital Zurich and the Verona-Udine Brain Imaging and Neuropsychology Program, we expect to contribute towards the concrete objective of providing effective, advanced techniques to assist in the diagnosis of renal cell carcinoma, one of the ten most frequent malignancies in Western countries, as well as of major psychoses such as schizophrenia and bipolar disorder. These problems are not amenable to be tackled with traditional machine learning techniques due to the difficulty of deriving suitable feature-based descriptions. Indeed, many biomedical applications exhibit precisely the same characteristics. Hence, a successful outcome of our experimentation would provide evidence as to the practical applicability of our approach in biomedicine, thereby fostering further research along the lines set up by SIMBAD, both at the methodological and at the practical level. This would potentially open new opportunities in health and disease management and bring radical improvements to the quality and efficiency of our healthcare systems.

SIMBAD involves the following partners:

- UNIVE: Ca’ Foscari University (Venice, Italy), *coordinator*
- UNIYORK: University of York (England)
- TUD: Delft University of Technology (The Netherlands)
- IST: Insituto Superior Técnico (Lisbon, Portugal)
- UNIVR: University of Verona (Italy)
- ETH Zurich: Eidgenössische Technische Hochschule Zürich (Switzerland)

The consortium has been carefully designed so as to include top-level competences in all relevant areas and problems in pattern recognition. In addition, we strove to set up a highly cohesive network where each research unit not only is well characterized in terms of competences, problems addressed, methodologies used and objectives, but is also tightly coupled with the rest of the network. All members in the consortium have an established international reputation and a long experience in the fields of machine learning and pattern recognition, where in the past few years have also contributed substantially to advance the state-of-the-art of the similarity-based paradigm.

The complementarity of expertises is crucial for our endeavour, as it allows us to attack each problem from different standpoints, thereby fostering cross-fertilization of ideas.

The SIMBAD website is: http://simbad-fp7.eu

2. Project objectives for the period

In this project we aim at advancing the state of the art in similarity-based pattern analysis and recognition from the theoretical, computational, and applicative perspective, and contributing towards the long-term goal of organizing this emerging field into a more coherent whole. As outlined in the "Publishable summary", the project revolves around two main themes, which concerns the issues of how to obtain suitable similarity information from non-vectorial representations, and how to use them, irrespective of the way in which they are obtained. In addition to these two basic themes, a third one arises which concerns the validation of the proposed techniques and their applicability to real-world problems. Pattern recognition is intrinsically an application-based field with well-established validation methodologies. These will be used to quantitatively evaluate the success of the proposed research on large-scale applications with clear social impact.
Accordingly, as for the first year, the objectives for the second year of the project have been structured around the following three strands.

1. Deriving similarities for non-vectorial data. The goal here is to develop suitable similarity measures for non-vectorial data. In the second year of SIMBAD we aimed at integrating the work done in the first year on information-theoretic kernels developed with generative kernels. Besides consolidating the research started in the first year, devoted to Fisher Kernels, we investigated score spaces based on the notion of *free energy* and studied the impact of the normalization step of the score space underlying a typical generative kernel. We also continued working on compression methods (yielding the so-called compression kernels) as a means to approximate information theoretic kernels, and focused on practical applications to a few challenging problems. An alternative to this "kernel tailoring" approach consists in learning good similarities directly from training data. Accordingly, in the second year, we have explored ideas based on ensemble-clustering. Here, we concentrated mainly on four directions, deeply grounded on similarity representations, namely combination methods, learning similarity on temporal data, clustering validity, and scalability. We have also strated studying game-theoretic learning approaches, but the preliminary results are not conclusive yet.

2. Learning and classification with non-(geo)metric similarities. Within this research strand we aim at both addressing foundational issues related to similarity information and developing practical algorithms that do not depend on the actual object representation. In particular, as concerns the latter objective, we have distinguished two cases, which in turn lead to two complementary approaches. On the one hand, we have considered the case where the informational content of non-(geo)metricity is limited or caused by measurement error. In this case it is a plausible strategy to perform some correction on the similarity data (or, alternatively, finding an alternative vectorial representation) in an attempt to impose (geo)metricity, and then use classical geometric techniques. On the other hand, when the information content of non-(geo)metricity is relevant one needs brand new tools, as standard techniques would not work in this case.

More specifically, in the second year of the prject the activity within these themes has been organized around three main areas:

a) Foundations of non (geo)metric similarities. One of the primary objectives here was to explore the origination of non-(geo)metric data. Causes for the non-Euclideaness of dissimilarity data have been identified. Some are non-intrinsic causes and corrections may improve classification results in such cases. For intrinsic causes, the non-Euclidean behavior is likely to be informative and Euclidean corrections or the construction of the Euclidean dissimilarity space may make the standard statistical classifiers applicable to these problems. We identified structural pattern recognition as a significant application domain of non-Euclidean dissimilarity measures, offering a bridge to the tools of statistical pattern recognition.

b) Imposing geometricity on non-geometric similarities (embedding). Despite the growing interest around embedding, the search for robust embeddings procedures on structured data such as weighted graphs has proven elusive, and their geometric and probabilistic characterizations still remains to be explored in depth. In the current period we have made substantial progress on the following topics: (i) Ihara zeta function and graph vectorization, (ii) heat kernel embeddings and spectral geometry, (iii) spherical embeddings, (iv) Ricci flows, (v) generative models of graphs, (vi) classification of attributed graphs, and (vii) structure-preserving embeddings.

c) Learning with non-(geo)metric similarities. When there is significant information content in the non-(geo)metricity of the data, it is hard to define a single, global objective function that satisfactorily accounts for the complexity of the problem at hand and hence alternative approaches are needed. Game theory was developed precisely to overcome the limitations of single-objective optimization as it aims at modeling complex situations where players make decisions in an attempt to maximize their own (mutually conflicting) returns. In the second year of the project we continued our work on game-theoretic models of learning. Specifically, we adapted the matching approach developed in the first year to the problem of inlier selection for robust estimation, and investigated a game-theoretic approach to the selection of distinguishing features from a set of loosely descriptive points, We also extended the general framework to handle continuous k-way interactions and we investigated multi-population settings for both unsupervised and semi-supervised learning problems.

3. Validation. Given the heterogeneity of different approaches that we are pursuing in this project, it is of particular importance to build a real-world testbed that specifically addresses the various difficulties involved with non-metric data. To this end, we are focusing on biomedical datasets that nicely combine high practical relevance of the underlying learning tasks and intrinsically non-metric dissimilarity data. In particular, we are applying the techniques developed in workpackages WP2–WP5 to (i) the analysis of Tissue Micro Array (TMA) images of renal cell carcinoma (RCC) and (ii) the analysis of brain magnetic resonance (MR) images in the context of mental health research (e.g., schizophrenia). As for the former, in the second year of the project we set up the automatic workflow for the RCC-TMA image analysis and worked on nucleus segmentation and the Leave-One-Patient-Out Cross-Validation of the classifiers implemented, so that our results of the nuclei classifiers are comparable with those of other groups. As for the latter, we devised several intermediate representations of the data set and used multi-modal registration to use also the DWI data. We applied various pattern recognition algorithms to the data representations for a better solution to the problem.

Following the outline set forth above, the three major themes form the basis for structuring the project's work plan into coherent work packages (WP's). Specifically, WP2 covers the topics of deriving similarities for non-vectorial data (theme 1), while the second theme concerning learning and classification with non-(geo)metric similarities is addressed by WP3, WP4 and WP5. Finally, the validation phase is undertaken in WP6 and WP7. An additional work package (WP1) deals with management issues, while WP8 deals with dissemination strategies.

More specifically, as far as the scientific and the dissemination part is concerned, the project is articulated in the following work packages:

WP2. Deriving similarities for non-vectorial data (structural kernels)

WP3. Foundations of non-(geo)metric similarities

WP4. Imposing geometricity on non-geometric similarities (embedding)

WP5. Learning with non-(geo)metric similarities

WP6. Analysis of tissue micro-array (TMA) images of renal cell carcinoma

WP7. Analysis of brain magnetic resonance (MR) scans for the diagnosis of mental illness

WP8: Dissemination, communication and exploitation

In the next section we shall describe the activities done within each work package (for more details we refer to related publications and/or deliverables).

3. *Work progress and achievements during the period (moths 1–12)*

Work package WP2:

Deriving similarities for non-vectorial data (structural kernels)

Work package leader: IST

Participants: IST, UNIVR, UNIVE, UNIYORK, ETH Zurich

Start month: 1 *End month: 24* *Overall person-months: 52*

The main goal of WP2 ("Deriving similarities for non-vectorial data") is to develop kernels and more general similarity measures for non-vectorial data. In this second-year report on this work-package, we briefly describe the main achievements in each of the tasks in which WP2 is divided ("Generative Kernels", "Compression Kernels", and "Learning and Combining Similarities"), and point to the resulting publications.

Generative Kernels

Generative kernels and information-theoretic kernels (the topic of WP2.2) are closely related. Both are based on the assumption that the objects of interest where generated by a probabilistic mechanism - a source, in information theoretic terms - and then proceed by defining (dis)similarity measures or kernels between models of these probabilistic sources. In fact, information-theoretic kernels can be considered as generative kernels, although the information-theoretic perspective tends to be more agnostic; i.e., it does not assume that the adopted model reflects the truth. In this section, we will describe the work on the topic traditionally called "generative kernels".

The lead contributor to the WP2.1 task was the UNIVR partner, but in this second year with significant integration with the information-theoretic kernels developed in the first year in the context of WP2.2; this work front will be described in the section devoted to WP2.2.

The research work on generative kernels has mainly focused on three directions: **(i)** the consolidation of the research, started during the first year, devoted to the step of learning the generative model underlying a generative kernel, focusing on Fisher Kernels; **(ii)** the definition and the investigation of score spaces based on the notion of *free energy* of a generative model; **(iii)** the study on the impact of the normalization step of the score space underlying a typical generative kernel.

In the first work direction, a study of the best procedure to build generative models for Fisher kernels was carried out. In particular, we started from the observation that, in the context of generative kernels, most of the research efforts have been devoted to the discriminative step, namely to the definition of a proper score space or of a kernel to be used with a SVM.

Actually, in a typical scenario, a single generative model describing the whole problem is employed. Alternatively, approaches using two models (one for the positive class and another for the rest) or one model per class have shown to increase the performances. We went one step ahead in this direction, allowing a generative framework to freely discover the natural structures or groups in the training set. This is achieved with a preliminary step of clustering, during which a large number of small hidden natural groups is extracted from the data, disregarding class label information. Subsequently, a single and simple generative model is trained for each group (as the groups tend to be small). The underlying intuition is simple: generative models are not used to discriminate between classes (this is left to the discriminative methods), but are used to finely describe the local structure of the data as an ensemble of clusters. Even if the proposed methodology may be general (and applicable to any generative kernel), we explored this direction focusing on the HMM-based Fisher Kernel case, showing promising and comparative results obtained from some experiments. All the details may be found in the publication [Bicego et al. 2009a] and in the D2.2 deliverable.

In the second direction of research, we continued the work carried out during the first year by defining a novel score space exploiting the free energy associated to a generative model, which is a popular score function representing a lower bound on the negative log-likelihood of the visible variables. The free energy permits to embed the uncertainty in the model parameters under the form of entropic terms. Such terms decompose in an entropy set and a cross-entropy set. The former encodes ambiguity within the model, thus considering the cases in which over-fitting or local minima occur during the learning. The latter encodes errors in the model's fit to the data, distributing such discrepancies across several terms, each one focused on a particular factor of the generative joint distribution. The entropy and cross-entropy sets are employed as features of the final score. The resulting score space shows to be highly informative for discriminative learning, allowing to achieve compelling comparative results in heterogeneous classification tasks, overwhelming the best classification accuracy on well-known databases. In particular, our approach was applied to face two bioinformatics problems (exons/introns classification, homology detection), and to deal with a typical computer vision issue (scene/object recognition), and results are compared with the best state-of-the-art outcomes present in the literature in the respective areas. All the details may be found in the publications [Perina et al. 2009a, Perina et al. 2009b] and in the D2.2 deliverable.

The third and final direction of research, mainly carried out in the last part of the second year, was devoted to the investigation of the effect of the normalization of the score spaces. In particular we first investigated the effect of a linear normalization, focusing on the Fisher score and on the so-called trans-space (one of the score spaces introduced for HMM in the first year of the project -- see the deliverable D2.2). Actually, we have shown that a proper normalization is often essential – all details may be found in the D2.2. In the literature, this need for normalization has been also shown for other generative kernels (like the *marginalized kernel*). Nevertheless, all the employed normalizations are based on linear operations, like shifting and linear scaling; on the contrary, we investigated the usefulness of nonlinear transformations. In particular, we focus on a particular class of Score spaces, defined on generative models with latent variables (for example, the states in a Hidden Markov Model). Several nonlinear mappings are indeed possible, and we investigated different ones, based on powering operation, logarithmic and logistic functions. Some experiments on HMM-based problems assessed the validity of the proposed approach, with really promising results. All the details may be found in [Carli et al 2009] and [Carli et al 2010].

Compression kernels

As mentioned in the first year report, this task would be better called "Information theoretic kernels"; its goals are to devise ways to obtain kernels for non-vectorial data, based on information theory.

In this second year of work, we have continued working on compression methods (yielding the so-called compression kernels) as a means to approximate the information theoretic kernels. The work on this front has mainly concentrated on practical applications of this type of kernels to a few challenging problems. In particular, we have we proposed a new compression-based ECG biometric method for personal identification and authentication, based on the Ziv-Merhav cross parsing algorithm for symbol sequences (strings), which works without any feature extraction from the waveforms. This method uses a string similarity measure obtained with a data compression algorithm and yields state-of-the-art performance both in identification and authentication tasks. We believe that this result is a clear proof of concept that compression-based (dis)similarities allow addressing difficult pattern recognition tasks, bypassing the critical feature extraction/selection step. More details and results can be found in [Pereira Coutinho et al 2010a] and [Pereira Coutinho et al 2010b].

A second direction of research studied the combination of the non-extensive information theoretic kernels (namely the Jensen-Shannon and the Jensen-Tsalis kernels, and their weighted counterparts) developed in the first year of this WP, with the generative embeddings developed in the context of WP2.1. This work required a close collaboration between the IST and UNIVR partners, with two visits from UNIVR researchers to IST (one of which a long stay of three months of a PhD student). To explain the direction that we have exploited, we begin by pointing out that using a generative embedding involves two steps: **(i)** defining and learning the generative model and using it to build the embedding; **(ii)** discriminatively learning a (maybe kernel) classifier on the adopted score space. The literature on generative embeddings is essentially focused on step (i), usually using some standard off-the-shelf tool for step (ii). In our work, we have considered a different approach, by focusing also on the discriminative learning step. In particular, we have combined the free energy score space (developed in WP2.1) with the non-extensive information theoretic kernels kernels developed in the first year of WP2.2. Preliminary experimental results in a scene recognition task on the benchmark Corel dataset have shown that this approach yields state-of-the-art performance. This results seem to indicate that the non-extensive information theoretic kernels are very adequate to work on the feature spaces induced by the generative embeddings. Full details of this approach and experimental results can be found in the following publications: [Bicego et al 2009b], [Bicego et al 2009c], and [Martins et al 2010].

Learning and combining similarities

This task is devoted to the problem of deriving similarities from examples and combining them. Two main strategies are used, exploiting: (A) clustering ensembles; (B) game-theoretic approaches.

During this second year, the research work based on the clustering ensemble approach has focused mainly on four directions, deeply grounded on similarity representations: **(i)** combination methods; **(ii)** learning similarity on temporal data; **(iii)** clustering validity; **(iv)** scalability.

Overall, the *evidence accumulation clustering* (EAC) method formed the ground clustering ensemble combination paradigm. This method can be decomposed into three major steps: (a) construction of the *clustering ensemble* (CE); (b) ensemble combination; and (c) extraction of the combined solution. The underlying assumption of EAC is that two objects belonging to the same "natural" cluster will be frequently grouped together. Using a pair-wise frequency count mechanism amongst a clustering committee, the method yields, as an intermediate result in step (b), a co-association matrix that summarizes the evidence taken from the several members in the clustering ensemble, each one taken as independent evidence of pair-wise data organization. This matrix corresponds to the maximum likelihood estimate of the probability of pairs of objects being in the same group, as assessed by the clustering committee, and can be regarded as a pair-wise similarity induced by the CE. A consensus partition is obtained by applying a clustering algorithm over the learned similarity.

The work in the first direction involved a close collaboration between the IST and the UNIVE partners, with one visit from IST researchers to UNIVE, and another from UNIVE researchers to IST. A pair-wise probabilistic clustering method using evidence accumulation was proposed as the result of the collaborative work between these groups. Taking the pair-wise similarity, learned with the EAC method, as estimate of the probability of pairs of objects to belong to the same cluster, this work proposes a probabilistic formulation for the combination process, leading to a consensus soft partition solution, where each object is probabilistically assigned to a cluster. The method reduces the clustering problem to a polynomial optimization in probability domain, which is attacked by means of the Baum-Eagon inequality. This work presents a principled probabilistic solution for consensus clustering, going one step further by extending the EAC paradigm from hard data partitioning to soft clustering solutions. Details and results can be found in [Bulò et al 2010]. In the later work, pair-wise similarities were used; extension of this method to higher order similarities is object of current joint research.

The second work direction consolidates research started during the first year. Taking as test bed the detection of stress states from ECG signals, the methodology previously proposed, combining genetic algorithms-based optimization with the clustering ensembles approach, was further pursued. Besides its temporal nature, these data express the continuous evolution between hidden internal physiological or functional states. As such, without hard separation between states, clustering becomes a very difficult problem. A deeper insight into the application problem revealed that these states may not be simply characterized by a typical heartbeat wave pattern, but rather by distinct combinations of wave patterns. In addition to the previous formulation of a continuous transition between states, meta-clustering was proposed to detect and characterize these more complex state patterns. Details and results can be found in [Medina et al 2010]. In the later work, single heartbeat waves were represented as vectors of features extracted from these. Ongoing work adopts a pure similarity-based framework, direct computation of similarities between pairs of heartbeat waves being issued. Preliminary experimental evaluation using cross-correlation similarity led to very promising results, further confirming the existence of stress-related ECG morphological changes.

The third direction of research addresses the following question: "For *a given data set, which clustering solution should be selected?*". The solution to this problem is based on clustering validation. While there is much work reported in the literature on validating data partitions produced by single clustering algorithms, little has been done in order to validate data partitions produced by clustering combination methods. Most of these works use measures of consistency between consensus solutions and the clustering ensemble, such as the Average Normalized Mutual Information proposed by Strehl and Ghosh. During the second year, we first addressed the validation at the CE level, proposing the Average Cluster Consistency (ACC) index. The main idea consists of measuring how well the clusters in the clustering ensemble fit in the clusters of the consensus partition. The similarity between each partition in the CE and the combined partition is measured based on a weighting of shared samples in matching clusters. The ACC validity index accounts for the average of these similarities over the CE. Details and results on this work can be found in [F. Duarte et al 2010]. Further work in the third research line, undertaken during the second year, proposed the validation of clustering combination results at three levels:

1. *Original data representation* - measure the consistency of clustering solutions with the structure of the data, perceived from the original representation (either feature-based or similarity-based);
2. *Clustering ensemble level* - measure the consistency of consensus partitions with the clustering ensemble;
3. *Learned pairwise similarity* - measure the coherence between clustering solutions and the co-association matrix induced by the clustering ensemble.

Taking pairwise similarities as the underlying representation, traditional clustering validity indices (namely the Silhouette, Dunn's and Davies and Bouldin's validity indices) were adapted to validate consensus solutions, when compared to the original data representation, and the learned similarity. These validity indices roughly account for intra-cluster compactness and inter-cluster separation. We then proposed a statistical validity index based on pair-wise similarity. According to the new index, the quality of the consensus partition is measured in terms of the likelihood of the data constrained to this partitioning. Inspired on the Parzen-window density estimation technique with variable size windows, a k-nearest neighbor density estimate from pair-wise similarities was defined. Taking as reference the learned similarity, the proposed validity index corresponds to a measure of goodness of fit of the consensus partition with the clustering ensemble and the pair-wise information extracted from it. When assessed from the original data representation, this validity index measures the goodness of fit of the combined partition with the statistical properties of the data on the baseline representation. A comparative study of the several validation approaches was undertaken on synthetic and real data. Details and results can be found in [Duarte et al 2010].

The fourth research direction undertaken during the second year, under the clustering ensemble framework, focuses on the scalability of both EAC and Multi-EAC methods to very large data sets. This topic was dealt in collaboration with Prof. Anil K. Jain, from the Michigan State University, USA. The bottleneck of the evidence accumulation paradigm is the quadratic (on the number of samples) space complexity associated with the full representation of the co-association matrix. Taking advantage of the typically sparse nature of this matrix, a split and merge strategy, for building clustering ensembles, combined with a sparse matrix representation was proposed. Experimental results show that linear time and space complexities are thus achieved. Details and results can be found in [Lourenço et al 2010].

Concerning the learning of similarities based on game-theoretic approaches, two main research directions were addressed: (i) learning the pay-off matrices for the clustering game; (ii) imposing Nash conditions. Research efforts involved a close collaboration between the IST and the UNIVE partners.

According to the game-theoretic framework, the pair-wise clustering problem is formalized as a competition between the hypotheses of class membership in a two player game. It assumes a pre-existing matrix of affinities between objects in the data set, corresponding to the players payoff. Within this framework, clusters correspond to evolutionary stable strategies, which is essentially a Nash equilibrium.

In the first work direction, the evidence accumulation clustering framework was further explored for learning the payoff matrix. Intrinsically to this method, the payoff matrix was defined as the learned pair-wise similarity, summarized in the co-association matrix. The clustering game framework assumes that all object within a cluster have high similarity. Given the diversity of clustering contributions for the learned similarity under the EAC approach, in general this assumption does not hold. In order to address this issue, transformations over the co-association matrix were considered, namely its scaling based on the logistic function, and the use of Gaussian kernels. Experimental evaluation of the approach in comparison with the use of payoffs derived from feature-based data representations was encouraging. These, still preliminary, results did not permit, however, the establishment of sound final conclusions about the effectiveness of the approach.

In the second research direction, work undertaken so far on imposing Nash equilibrium based on clusters extracted under the EAC method is still limited and inconclusive, requiring further exploration.

Overall, preliminary experiments on learning the payoff matrices and on using Nash equilibrium performed so far did not provide a clear evidence of the effectiveness of these approaches. Further exploration and validation of the methods is required in order to get definite results, which is the focus of ongoing work.

According to the original plan, outlined in the project proposal, the WP2.3 task envisaged to: extend previous work on the Evidence Accumulation paradigm (EAC), by exploring different data representations, as well as new combination techniques, and inclusion of constraints into the combination process; address the validity of the learned similarities, based on a stability-related index; learn similarities from examples based on game-theoretic approaches by imposing Nash conditions; validation of proposed techniques on local real world data sets.

From the reported above, concerning the work undertaken during the second year of the project, and the work developed during the project's first year, the main goals have been achieved, and new directions, not envisaged at the project proposal stage, were explored.

As stated before, the work on learning similarities from examples based on game-theoretic approaches by imposing Nash conditions, although partially addressed during the second year, did not lead to definitive conclusions. This was the main motivation for the request for delay of deliverable D2.3, which was due at the end of the second reporting period, and that was postponed for July 2010, since this delay has no impact on the other tasks. Furthermore, the impact of this deliverable may even be strengthened by the inclusion of results of ongoing work related to other researched directions, described above.

Publications

[Bicego et al. 2009a]: M. Bicego, M. Cristani, V. Murino, E. Pekalska, R. Duin, "Clustering-based construction of Hidden Markov Models for generative kernels", *Proc. of Int. Conf. on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR2009),* pp. 466-479, Bonn, Germany, 2009.

[Bicego et al 2009b] M. Bicego, A. Martins, V. Murino, P. Aguiar, and M. Figueiredo, “2D shape recognition using information theoretic kernels”, International Conference on Pattern Recognition (ICPR), Istanbul, Turkey, 2010 (accepted).

[Bicego et al 2009c] M. Bicego, A. Perina, V. Murino, A. Martins, P. Aguiar, and M. Figueiredo, “Combining free energy score spaces with information theoretic kernels: application to scene classification”, International Conference on Image Processing (ICIP), Hong Kong, China, 2010 (submitted).

[Carli et al 2009] A. Carli, M. Bicego, S. Baldo, V. Murino, "Non-linear Generative embeddings for kernels on latent variable models", *Proceedings of the ICCV’2009 Workshop on Subspace Methods*, Kyoto, Japan, 2009.

[Carli et al 2010] A. Carli, M. Bicego, S. Baldo, V. Murino, "Nonlinear mappings for generative kernels on latent variable models", Proceedings of the International Conference on Pattern Recognition (ICPR), Istanbul, Turkey, 2010 (accepted).

[Martins et al 2010] A. Martins, M. Bicego, V. Murino, P. Aguiar, and M. Figueiredo, "Information Theoretical Kernels for Generative Embeddings Based on Hidden Markov Models", *Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition and Statistical Techniques in Pattern Recognition (S+SSPR)*, Izmir, Turkey, 2010 (accepted).

[Pereira Coutinho et al 2010a] D. Pereira Coutinho, A. Fred, M. Figueiredo, “'Personal identification and authentication based on one-lead ECG by using Ziv-Merhav cross parsing”, 10th International Workshop on Pattern Recognition in Information Systems (PRIS), Funchal, Portugal, 2010 (accepted).

[Pereira Coutinho et al 2010b] D. Pereira Coutinho, A. Fred, M. Figueiredo, “One-lead ECG-based personal identification using Ziv-Merhav cross parsing”, Proceedings of the International Conference on Pattern Recognition (ICPR), Istanbul, Turkey, 2010 (accepted).

[Perina et al 2009a] A. Perina, M. Cristani, U. Castellani, V. Murino, and N.Jojic. “Free energy score space”, *Advances in Neural Information Processing Systems (NIPS)*, Vancouver, Canada, 2009.

[Perina et al 2009b] A. Perina, M. Cristani, U. Castellani, V. Murino, and N.Jojic. A hybrid generative/discriminative classification framework based on free-energy terms”, *Proceedings of the International Conference on Computer Vision (ICCV),* Kyoto, Japan, 2009.

[Bulò et al 2010] S. Bulò, A. Lourenço, A. Fred, M. Pelillo, “Pairwise probabilistic clustering using evidence accumulation”, *Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition and Statistical Techniques in Pattern Recognition (S+SSPR)*, Izmir, Turkey, 2010 (accepted).

[Medina et al 2010] L. Medina, A. Fred. “Clustering temporal data: application to electrophysiological signals”, Accepted for publication In: *Communications in Computer and Information Science*. Springer 2010.

[F. Duarte et al 2010] F. J. Duarte, J. Duarte, A. Fred. “Average cluster consistency for cluster ensemble selection”, Accepted for publication In: *Communications in Computer and Information Science*. Springer 2010.

[Duarte et al 2010] J. Duarte, A. Fred, A. Lourenço, F. J. Duarte, “On consensus clustering validation” , *Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition and Statistical Techniques in Pattern Recognition (S+SSPR)*, Izmir, Turkey, 2010 (accepted).

[Lourenço et al 2010] A. Lourenço, A. Fred, A. K. Jain, “On the scalability of evidence accumulation clustering”, Proceedings of the International Conference on Pattern Recognition (ICPR), Istanbul, Turkey, 2010 (acc

Work package WP3:

Foundations of non-(geo)metric similarities

Work package leader: TUD

Participants: TUD, ETH Zurich

Start month: 1 *End month: 18* *Overall person-months: 29*

In this year the workpackage was finished. We summarize the objectives, repeat and rephrase the conclusions from the first year studies and summarize the deliverable of the second year.

The objective of the workpackage is to study both the causes of the lack of (geo)metricity in dissimilarity data and its effect on traditional machine learning algorithms. For non-metric dissimilarity data, the triangle inequality is violated for some subsets of three data points. For metric data, a basic structure is available in the dataset. However, it is still possible that sets of more than three points would not fit into a Euclidean space. Consequently, they cannot be isometrically embedded in such a space. As a result, the assumptions are violated for many of the traditional procedures for analyzing and classifying data. For that reason the target of this work package is to study non-Euclidean data. We distinguish the following questions:

- Why is there non-Euclidean data?
- How can it be analyzed? What representation should be used?
- What are examples of datasets?

These are studied and presented in the deliverables. A summary of the work package results follows

The cause of non-Euclidean data

We identified two main causes for non-Euclidean behavior : **non-intrinsic** ones and **intrinsic** ones.

Non-intrinsic causes are related to computational and observational problems. Examples are measurement noise, computational inaccuracies, algorithmic shortcuts in optimization procedures and missing data.

Intrinsic causes are deliberately chosen for non-Euclidean distance measures like the $l1$ norm used for comparing spectra, the edit distance and the Hausdorff distance used for comparing shapes and the single-linkage procedure used in cluster analysis. An important group of the set of intrinsic causes are the pairwise object comparisons that define their own subspace or normalization.

A basic reason for experts to design non-Euclidean dissimilarity measures is the need to incorporate object structure into a measure. Thereby the gap between structural and statistical pattern recognition is bridged.

The use of training objects as well as test objects (i.e. the objects to be classified) for constructing the representation may be necessary in case the non-Euclidean dissimilarities are essential for a good classification. Transductive learning becomes then a consequence of using informative non-Euclidean object proximities.

The representation of non-Euclidean data

We can distinguish three approaches to construct a vector space given a full set of dissimilarities.

Embedding into a pseudo-Euclidean space. This is possible without any error, so there is no loss of information. A problem is however that the inner product definition and the distance definition demand different classifiers in this space than the traditional ones. The SVM suffers from non-Mercer or indefinite kernels, density estimation is not well defined and objects can have negative square distances to other objects. Some classifiers can be defined on the basis of distances or by using kernelized versions.

Euclidean corrections, either based on the pseudo-Euclidean embedding, or directly on the given dissimilarity matrix are available as well. Examples are neglecting the specific distance measure defined for the pseudo-Euclidean space or using a Euclidean subspace of this space. An interesting alternative is the transformation studied in report D3.2: enlarging the off-diagonal dissimilarities by some constant before embedding.

The dissimilarity space. This approach postulates a Euclidean space using the vector of dissimilarities from a given object to the so-called representation set (e.g. the training set) as ’features’. This can always be done without any computational overhead, except when dimension reduction is desired.

Examples of datasets

The datasets collected in the project are described in the SIMBAD report 2009_9. We collected 64 dissimilarity matrices, almost all having a non-Euclidean behavior. We developed software for a systematic analysis of the datasets and compared the performance of some classifiers in various spaces as discussed above.

In general, dissimilarity spaces do somewhat better than pseudo-Euclidean embedded spaces and spaces derived by Euclidean corrections. The main finding however is that for some datasets a direct removal of the non-Euclidean characteristics of the data (neglecting eigenvectors with negative eigenvalues in the pseudo-Euclidean embedding) is counterproductive for some classifiers. So, the non-Euclidean behavior is informative for these datasets and these classifiers. This does not necessarily imply that classifiers used in the pseudo-Euclidean space or based on indefinite kernels are only beneficial. Euclidean corrections may bring the relevant information into the domain of Euclidean classifiers, as well. Some extreme, artificially examples could be given in which all information for the separation between the classes is located in the negative part of the pseudo-Euclidean space

One of the datasets that has become available in the SIMBAD project (WP7) is the Verona set of 182 dissimilarity matrices derived from the MRI brain scans of 124 subjects. They are constructed using 13 different dissimilarity measures applied to ROIs related to 14 different regions in the brain. Majority of the measures is non-Euclidean. The data is analyzed with the same procedures as mentioned above. In addition, they are combined into a single dissimilarity matrix which showed a significantly better performance than any of the individual constituting matrices.

Conclusions of WP3

Causes for the non-Euclideaness of dissimilarity data have been identified. Some are non-intrinsic causes and corrections may improve classification results in such cases. For intrinsic causes, the non-Euclidean behavior is likely to be informative and Euclidean corrections or the construction of the Euclidean dissimilarity space may make the standard statistical classifiers applicable to these problems. We identified structural pattern recognition as a significant application domain of non-Euclidean dissimilarity measures, offering a bridge to the tools of statistical pattern recognition. Examples are the many applications using dissimilarity measures between spectra or histograms like a shape distance or the earth mover distance.

Publications

W.J. Lee and R.P.W. Duin, A Labelled Graph Based Multiple Classifier System, in: J.A. Benediktsson, J. Kittler, F. Roli (eds.), *Multiple Classifier Systems* (Proc. 8th Int. Workshop, MCS 2009, Reykjavik, Iceland, June 10- 12, 2009), Lecture Notes in Computer Science, vol. 5519, Springer, Berlin, 2009, 201-210.

S.W. Kim and R.P.W. Duin, A Combine-Correct-Combine Scheme for Optimizing Dissimilarity-Based Classifiers, in: E. Bayro-Corrochano, J.O. Eklundh (eds.), *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Proc. CIARP 2009, Guadalajara, Jalisco, Mexico, Nov. 15-18, 2009), Lecture Notes in Computer Science, vol. 5856, Springer, Berlin, 2009, 425-432.

Y. Plasencia, E. Garcıa-Reyes, R. P. W. Duin, H. Mendez-Vazquez, C. San- Martin, and C. Soto, A Study on Representations for Face Recognition from Thermal Images, in: E. Bayro-Corrochano, J.O. Eklundh (eds.), *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Proc. CIARP 2009, Guadalajara, Jalisco, Mexico, Nov. 15-18, 2009), Lecture Notes in Computer Science, vol. 5856, Springer, Berlin, 2009, 185-192.

D. Porro Munoz, R.P.W. Duin, I. Talavera, and N. Hernandez, The Representation of Chemical Spectral Data for Classification, in: E. Bayro- Corrochano, J.O. Eklundh (eds.), *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Proc. CIARP 2009, Guadalajara, Jalisco, Mexico, Nov. 15-18, 2009), Lecture Notes in Computer Science, vol. 5856, Springer, Berlin, 2009, 513-520.

S. Klein, M. Loog, F. van der Lijn, T. den Heijer, A. Hammers, M. de Bruijne, A. van der Lugt, R.P.W. Duin, M.M.B. Breteler, and W.J. Niessen, Early diagnosis of dementia based on intersubject whole-brain dissimilarities, *Proc. of IEEE International Symposium on Biomedical Imaging: Macro to Nano* (Rotterdam, 14-17 April 2010), 2010.

W.J. Lee, R.P.W. Duin, and H. Bunke, Selecting Structural Base Classifiers for Graph-based Multiple Classifier Systems, in: N. El Gayar, J. Kittler, F. M. Roli (eds.), *Multiple Classifier Systems* (Proc. 9th Int. Workshop, MCS 2010, Cairo, Egypt), Lecture Notes in Computer Science, vol. 5997, Springer, Berlin, 2010, 155-164.

W.J. Lee, R.P.W. Duin, M. Loog, and A. Ibba, An Experimental Study On Combining Euclidean Distances, *Proc. The 2nd International Workshop on Cognitive Information Processing* (14-15-16 June, 2010 Elba Island, Tuscany - Italy), 2010, accepted

Millan-Giraldo, R.P.W. Duin, and J.S. Sanchez, Dissimilarity-based Classification of Data with Missing Attributes, *Proc. The 2nd International Workshop on Cognitive Information Processing* (14-15-16 June, 2010 Elba Island, Tuscany - Italy), 2010. accepted

L. Sørensen, R.P.W. Duin, M. de Bruijne, W.J. Lee, D.M.J. Tax, and M. Loog, Dissimilarity-based Multiple Instance Learning, *Proc. SSSPR, 2010*. accepted

D. Porro Munoz, R.P.W. Duin, M. Orozco-Alzate, I. Talavera-Bustamante, and J.M. Londo˜no-Bonilla, The Dissimilarity Representation as a Tool for Three-way Data Classification: a 2D Measure, *Proc. SSSPR, 2010*. accepted

S.W. Kim and R.P.W. Duin, An emperical comparison of kernel-based and dissimilarity-based feature spaces, *Proc. SSSPR, 2010*, accepted.

D. Porro-Munoz, R.P.W. Duin, M. Orozco-Alzate, I. Talavera, and J.M. Londo˜no- Bonilla, Classifying Three-way Seismic Volcanic Data by Dissimilarity Representation, *ICPR 2010*, 2010. accepted

A. Ibba, R.P.W. Duin, and W.J. Lee, A study on combining sets of differently measured dissimilarities, *ICPR 2010*, accepted

Y. Plasencia-Calana, E. Garcıa-Reyeso, M. Orozco-Alzate, and R.P.W. Duin, Prototype selection for dissimilarity representation by a genetic algorithm, *ICPR 2010*, 2010. accepted

D. Porro-Munoz, I. Talavera, R.P.W. Duin, N. Hernandez, and M. Orozco- Alzate, Dissimilarity Representation on Functional Spectral Data for Classification, *Journal of Chemometrics*, 2010, submitted.

A. Ulas, R.P.W. Duin, U. Castellani, M. Loog, M. Bicego, V. Murino, M. Bellani, S. Cerruti, M. Tansella, and P. Brambilla, Dissimilarity-based Detection of Schizophrenia, *Proc. ICPR 2010 workshop on Pattern Recognition Challenges in FMRI Neuroimaging*, 2010, accepted.

L. Sørensen, M. de Bruijne, R.P.W. Duin, M. Loog, P. Lo, and A. Dirksen, Image Dissimilarity-Based Quantification of Pathology, *MICCAI, 2010*. accepted

Work package WP4:

Imposing geometricity on non-geometric similarities (embedding)

Work package leader: UNIYORK

Participants: UNIYORK, ETH Zurich, UNIVE, TUD

Start month: 7 *End month: 30* *Overall person-months: 62*

The overall aim in this work package is that, given similarity data, possibly in the form of a weighted graph, we aim at developing algorithms for transforming them into instance-specific vectorial representations (embedding) that are suitable for traditional geometric learning algorithms. Over the period covered Deliverables 4.1 and 4.2 have been completed. Workpackage 4.1 continued to produce interesting scientific results, and Workpackage 4.2 got underway and produced its first findings. There has been a strong collaboration between York and Delft on spherical embeddings which resulted in a CVPR 2010 paper. There are also ongoing collaborations between York and Verona on assessing the non-metricity of the MRI data-sets, and between York-Lisbon-Verona on issues relating to kernelisation. Joint work with Venice is planned on the topic of establishing links between generative super-graph models and compression kernels. Results from the workpackage have been diseminated via the usual publication channels (papers in journals and conferences). Talks and seminars describing aspects of the work have been presented at the NICTA/ANU Summer School on Computer Vision, at UIUC (Beckman Institute) and QMW (London). Material from the workpackage will also be presented at Invited Keynote talks at ICIAR 2010, MCPR 2010 and the TextGraphs workshop at ACL 2010 in Upsala.

Ihara zeta functions and graph vectorisation

In the previous reporting period, we have explored a new structural characterisation of graphs based on the Ihara zeta function. Our motivation for embarking on this study was prior work at York, where we demonstrated that the zeta function is the moment generating function of the heat kernel trace and can be used to cluster graphs. The Ihara zeta function which is an extension of Riemann zeta function from prime numbers to prime cycles in a graph. It is constructed by first transforming a graph into its directed line-graph. The Ihara zet function is reciprocal of the characteristic polynomial of the transition matrix for the directed linegraph. Our characterisation of graphs is based on coefficients of the polynomial. The coefficients of the polynomial are related to the cycle frequencies in the graph, and are easily computed using the eigenvalues of the directed line-graph transition matrix (the so-called Froebenius operator). In the last reporting period, we made a number of original developments to the study of the Ihara zeta function including generalising it to weighted graphs.

In the current period we have explored in some depth how the Ihara zeta function can be generalised to hyper-graphs, and have obtained extensive experimental results on data-sets derived from image data. We have also undertaken an analysis of the log-determinant definining the Ihara zeta function, and re-expressed it in exponential-trace form. This opens up the possibity of simplified spectral analysis and the use of the

Ihara zeta function to define a new class of cycle-length kernels. We aim to undertake this work in collaboration with Lisbon and Verona.

Papers describing this work appear in ICCV 2009 and have been submitted to IEEE TNN and Pattern Recognition.

Heat kernel embeddings and spectral geometry

In the previous period, we concentrated on the use of the heat-kernel for graph embedding, and explored its links with spectral geometry. During the current period, this work has been extended to investigate the use of the wave-kernel as an alternative. Graph embeddings have found widespread use in machine learning and pattern recognition for the purposes of clustering, analyzing and visualizating relational data. However, they have also proved to be useful as a means of graph characterization. Examples in the literature including ISOmap, the Laplacian eigenmap, and the heat-kernel embedding. Once embedded, a graph can be characterised using a feature-vector that characterises the point-set distribution resulting from the embedding. This kind of representation is convenient since a Euclidean vector space makes available powerful geometric analysis tools for data analysis, not available for discrete or structural representations. However, such an embedding assumes that the original relational data is metric. Sometimes, however, this is not the case. For instance when the matrix characterisation of the relational graph contains negative eigenvalues, i.e. it is not positive semi-definite. Under these circumstances the graph embeds not into a Euclidean space, but into pseudo-Euclidean or Krein space. This problem has attracted relatively little attention in the literature. Our aim in this paper is to embed the nodes of a graph as points on the surface of a pseudo-Riemannian manifold in a pseudo-Euclidean space, and to use the resulting point-set as the basis from which to compute graph characteristics. To provide a framework for our study, we turn to the wave kernel. This is the solution of a wave equation, which is an important second-order linear partial differential equation that describes the propagation of a variety of waves. Crucially, the solutions are complex and therefore reside in a pseudo-Euclidean space. Although the wave equation has been extensively studied in the continuous domain, there has been relatively little effort devoted to understanding its behavior on a graph. In common with the heat kernel, the wave kernel can be defined in terms of a combinatorial Laplacian. However, in the case of the wave kernel this is the edge-based Laplacian, introduced by Friedman.

In this period we have explored how to solve the edge-based wave equation, in terms of the eigensystem of the edge-based Laplacian. Since the solution is a sinusoid, it contains both real and imaginary parts. Hence, we embed the nodes of the graph as points residing on a pseudo-Riemannian manifold, determined by the eigenvalues and eigenvectors of the edge-based Laplacian. In an experimental investigation on graphs extracted from 2D image data, we use this matrix for the purpose of graph visualization.

A paper describing the early formation of these ideas appears in S+SSPR. In the future, we plan to apply the method to the Verona and Zurich datasets.

Spherical embbeddings

There are two challenges when analysing patterns defined in terms of the disimilarity between objects rather than ordinal values. First,the objects can not clustered or classified using standard pattern recognition techniques, since they are not characterised by pattern-vectors. Instead, pairwise rather than central clustering techniques must be used. Alternatively, the objects can be embedded into a vector-space using techniques such as multidimensional scaling or IsoMap. Once embedded in such a space then the objects can be characterised by their embedding co-ordinate vectors, and analysed in a conventional manner. Most embedding methods produce an embedding that is Euclidean. However, dissimilarity data cannot always be embedded exactly into a Euclidean space. This is the case when the similarity matrix (the equivalent of a kernel matrix) contains negative eigenvalues, and where the embedding (which depends on the square-root of the eigenvalue matrix) is non-real. Examples of such dissimilarity data occur in a number of data sources furnished by applications in computer vision. For instance, shape-similarity measures and graph-similarity measures are rarely Euclidean. Previous work has shown that there is potentially useful information in the non-Euclidean part of the dissimilarities. Such data can be embedded in a pseudo-Euclidean space, i.e. one where certain dimensions are characterised by negative eigenvalues and the squared-distance between objects has positive and negative components which sum together together to give the total distance. A pseudo-Euclidean space is however non-metric which makes it difficult to correctly compute geometric properties. Another alternative, which we explore in this workpackage, is to embed the data on a Riemannian manifold, which is metric but non-Euclidean. Specifically, we provide a means of embedding such data onto a hypersphere whose radius of curvature is determined by the dissimilarity data. The hypersphere can be either of positive curvature (i.e. an elliptic surface) or of negative curvature (i.e. a hyperbolic surface). We show how to approximate a distribution of dissimilarity data by a suitable hypersphere. Our analysis commences by defining the embedding in terms of a co-ordinate matrix that minimises the Frobenius norm with a similarity matrix. We show how the curvature of the embedding hypersphere is related to the eigenvalues of this matrix. In the case of an elliptic embedding, the radius of curvature is given by an optimisation problem on the smallest eigenvalues of a similarity matrix, while in the case of a hyperbolic embedding it is dependent on the second-smallest eigenvalue. Under the embedding, the geodesic distances between points are metric but non-Euclidean. Once embedded, we can characterise the objects using a revised dissimilarity matrix based on geodesic distances on the hypersphere. We apply our method to a variety of data including shape-similarities, graph comparison and gesture interpretation data. In each case the embedding maintains the local structure of the data while placing the points in a metric space. similarity data so that it conforms either to elliptic or hyperbolic geometry. In practice the former corresponds to a scaling of the distance using a sine function, and the latter scaling the data using a hyperbolic sine function.

With the embedding to hand we develop classifiers tailored to the hypersphere. To this end we develop an optimisation-based procedure for embedding objects on hyperspherical manifolds. The purpose of this embedding is to faithfully represent the dissimilarities between objects in a metric space. A metric space is important because is allows us to compute statistics and defne geometric constructs such as boundaries, in contrast to a non-metric space where nonlocality is a problem. We also define the nearest mean classifier on the manifold. The optimisation approach we use employs the Lie group representation of the hypersphere and its associated Lie algebra to define the exponential map between the manifold and its local tangent space. We can then solve the optimisation problem locally in Euclidean space. This process is efficient enough to allow us to embed datasets of several thousand objects.

This work has been conducted in collaboration with Delft (Duin and Pekalska), and a joint paper appears in CVPR 2010.

Ricci flow

Embedding provides a way to apply feature-based classifiers to dissimilarity data and produces a vectorial representation by projecting dissimilarity data into a fixed dimensional vector space. Multidimensional scaling (MDS) is one of the earliest embedding techniques to find a vectorial representation in a Euclidean space. More recent approaches such as ISOMAP attempt to simultaneously reduce the dimensionality of the embedded space whilst minimizing distortion. This is achieved by inferring a low dimensional manifold on which the data resides. The common aim of the above embedding methods is to locate a low-dimensional representation. In order to apply non-Euclidean dissimilarity data with traditional geometry based learning techniques, we must attempt to rectify the data so as to minimize the non-Euclidean artifacts. One way is to consider the positive definite subspace of the distances. An alternative route adopted by Pekalska et al. is to add a suitable constant amount to the off-diagonal elements of the dissimilarity matrix. This is equivalent to adding a certain constant all eigenvalues of the related Gram matrix, and thus compensating for the effect of the negative eignevalues, while maintaining the same eigenvector structure. An important question that arises is how to correct the dissimilarity such that the new Euclidean distances are not less discriminating than the non-Euclidean dissimilarity. To address this problem we have explored the idea of using Ricci flow on a constant curvature Riemannian manifold to evolve the distance measures through updating Gaussian curvatures on the edges of a weighted graph, so that non-Euclidean distance measures can be rectified (i.e. made Euclidean).

We are applying this methd to the Verona data to assess the degree of non-metricy present and to determine whether we can rectifiy this. We hope to explore this in depth at the Castelbrando workshop.

Generative models of graphs

Relational graphs provide a convenient means of representing dissimilarity data.When abstracted in this way then complex data can be compared or matched using graph-matching techniques. Although matching problems, such as subgraph isomorphism or inexact graph matching provide a computational bottleneck, there are a number of effective algorithms based on probabilistic, optimization or graph spectral techniques that can give reliable results in polynomial time. However, despite considerable progress in the problems of representing and matching data using graph structures, the issue of how to capture variability in such representations has received relatively little attention. For vectorial patterns on the other hand, there is a wealth of literature on how to construct statistical generative models that can deal with quite complex data including that arising from the analysis of variability in shape. The main reason for the lack of progress is the difficulty in developing representations that capture variations in graph-structure. This can manifest itself either as variations in a) edge-connectivity or as variation in b) node-composition and c) node or edge attributes. This trichotomy provides a natural framework for analyzing the state-of-the-art in the literature. Most of the literature can be viewed as modeling variations in node or edge attributes. In fact, most of the work on Bayes nets in the graphical models literature falls into this category. There are also some well documented studies in the structural pattern recognition literature that also fall into this category including the work on Christmas et al. and Bagdanov et al. who both use Gaussian models to capture variations in edge attributes. The problems of modeling variations in node and edge composition are more challenging since they focus on modeling the structure of the graph rather than its attributes. For the restricted class of trees, Torsello and Hancock use a description length criterion to recover the node composition of trees from samples with unknown correspondences. Torsello and Dowe have recently made some progress in extending this method to graphs using importance sampling techniques to overcome some of the computational bottlenecks. The problem of learning edge structure is probably the most challenging of those listed above. Although this is finessed in the case of trees since the operations of node and edge insertion or deletion are equivalent. Broadly speaking there are two approaches to characterizing variations in edge structure for graphs. The first of these is graph spectral, while the second is probabilistic. In the case of graph spectra, many of the ideas developed in the generative modeling of shape using principal components analysis can be translated relatively directly to graphs using simple vectorization procedures based on the correspondences conveyed by the ordering of Laplacian eigenvectors. Although these methods are simple and effective, they are limited by the stability of the Laplacian spectrum under perturbations in graph-structure. The probabilistic approach is potentially more robust, but requires accurate correspondence information to be inferred from the available graph structure. If this is to hand, then a representation of edge structure can be learned. To date the most effective algorithm falling into this category exploits a part-based representation.

We have conduced a more ambitious study where we aim to capture variations in edge structure by learning a probabilistic model of the adjacency matrix when correspondences are not available and must be inferred from the data. We follow Torsello and Hancock and work with a supergraph representation from which each sample graph can be obtained by edit operations.The supergraph is represented by matrix in which each element reflects the probability of connection between a pair of nodes. To furnish the required learning framework, we extend the work of Luo and Hancock and develop an EM algorithm in which the node correspondences and the supergraph edge probability matrix are treated as missing data. This novel technique has been applied to a large database of object views,and used to learn class prototypes that can be used for the purposes of recognition.

This work will form the basis of a collaborative effort with Lisbon and Verona on the topic of generative kernels. Once learned our supergraph provides a generative model that acts in a compressive way, succinctly capturing the variations in a sample of graphs so as to minimise a description length criterion.

Classification of attributed graphs

In the first year of the project various graph matching procedures are worked out, resulting in various non-Euclidean dissimilarity measures. Some are compared on a set of non-attributed graphs generated from the

Coil image dataset. Classifiers in various spaces derived from pseudo-Euclidean embeddings are used. They are compared in the dataset report D3.3 delivered by WP3 earlier this year.

We focussed further on the classification of attributed graphs and started with symbollically labeled nodes. Separate classifiers are trained for sub-graphs defined by the same node labels using the above mentioned procedures for non-attributed graphs. As many classifiers are found as different labels are used for the node attributes. These classifiers are always defined on similar dissimilarity matrices between all training objects. They can thereby be combined in various ways. A study has been presented at the Multiple Classifier Systems workshop MCS 2010. This work has been perfomed in coooperation with Bunke, University of Bern, Switzerland.

In the next year we will focus on graphs with multi-dimensional feature vectors used as node attributes. They will be transformed into non-attributed graphs with weighted edges. Dissimilarity measures will be based on either an edit distance or on spectral graph embedding procedures.

Structure preserving embeddings

In this task we analyze embeddings into Euclidean spaces which leave the group- or clustering structure invariant. The clustering approaches discussed here aim at identifying subsets or clusters of objects represented as ``blocks'' in a permuted dissimilarity matrix. The underlying idea is that objects grouped together in such a cluster can be reasonably well described as a homogeneous sub-population. Our focus on dissimilarity matrices implies that we do not have access to a vectorial representation of the objects, and in general, no such representation will exist, since we do not assume that the dissimilarity matrix fulfills the axioms of a valid metric.

For the class of pairwise clustering methods that are related to minimizing a shift-invariant cost function, the constant shift embedding procedure leads to exact preservation of group structure: the original non-metric pairwise clustering problem can be restated as a grouping problem over points in a vector space, yielding identical assignments of objects to clusters. As a consequence of this constant-shift embedding principle, we could show the equivalence between the pairwise clustering cost function and the classical k-means grouping criterion in the embedding space. The conclusion is that the the k-means cost function (or its dissimilarity-based counterpart) is essentially ``blind'' against metric violations.

In the current period we have considered more general settings where the hard-clustering scenario with fixed number of clusters is replaced by a probabilistic approach which is capable of selecting the number of clusters in a data-adaptive way. We have shown that the Wishart-Dirichlet cluster process for partitioning dissimilarity matrices is shift invariant in an approximate sense, and in particular we have shown that exact shift invariance and data-adaptive selection of the number of clusters define two conflicting goals: shifting increases the tendency to introduce new clusters, since under the shift the mutual similarities between all objects decrease. It seems that strict shift invariance can only be achieved if the number of clusters is fixed, which somehow contradicts our efforts to generalize the k-means setting.

Considering the relevance of structure preserving embedding for the overall goal of the SIMBAD project, namely the development of a new theory of similarity-based pattern recognition, our current view is ambivalent: strict structure preservation could be proved only for a small set of clustering methods, like pairwise k-means and certain graph-based cut/association algorithms. All these algorithms require the user to fix the number of clusters in advance. A ``relaxed'' version of shift invariance holds for a probabilistic version of the pairwise k-means method, but we have to admit that shift invariance and estimation of the number of clusters might be two conflicting goals.

When it comes to building a theory on similarity-based pattern recognition, all these algorithms may be seen as ``negative results'', since they are essentially blind against Euclidean- or even metric violations. In other words: if one wants to learn something about clustering similarity data, one should look at different clustering procedures.

A paper describing the Wishart-Dirichlet cluster process appears in ICML 2010: (Julia E. Vogt, Sandhya Prabhakaran, Thomas J. Fuchs, Volker Roth. The Translation-invariant Wishart-Dirichlet Process for Clustering Distance Data. ICML 2010: Proceedings of the 27th international conference on Machine Learning. In press.) A paper on information theoretic model selection appears in ISIT 2010: (Joachim M. Buhmann. Information theoretic model validation for clustering. In International Symposium on Information Theory, Austin Texas. IEEE, 2010. In press).

Publications

I. Patras and E.R. Hancock. Coupled prediction-classification for robust visual tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence, to appear, 2010.

Bai Xiao, Hang Yu and E.R. Hancock. Manifold Embedding for Shape Analysis, Neurocomputing, to appear, 2010.

DongJun Yu, W.A.P. Smith and E.R. Hancock. Learning a Self-Organising Map Model on a Riemannian Manifold. Mathematics of Surface, LNCS, 5654, pp. 375--390, 2009.

Shenping Xia and E.R. Hancock. Learning Class Specific Hypergraphs. ICIAP 2009, LNCS, 5716, pp. 269--277, 2009.

DongJun Yu, W.A.P. Smith and E.R. Hancock. A Riemammian Self-Organizing Map. ICIAP 2009, LNCS, 5716, pp. 229--238, 2009.

Chengming Zou, Guanghui Zhao, and E.R. Hancock. Reconstructing 3D Facial Shape using Spherical Harmonics. ICIAP 2009, LNCS, 5716, pp. 949--957, 2009.

Peng Ren, R.C. Wilson and E.R. Hancock. Hypergraphs, Characteristic Polynomials and the Ihara Zeta Function. CAIP 2009, LNCS, 5702, pp. 369--376, 2009.

Shenping Xia and E.R. Hancock. Graph based Object Class Discovery. CAIP 2009, LNCS, 5702, pp. 385--393, 2009.

Peng Ren, R.C. Wilson and E.R. Hancock. Weighted Graph Characteristics from Oriented Line Graph Polynomials. International Conference on Computer Vision, pp. 2280 - 2287, 2009.

S. Xia and E.R. Hancock. Incrementally Discovering Object Classes using Similarity Propagation and Graph Clustering. ACCV, LNCS 5996, pp. 373--383, 2009.

R. C. Wilson, E.R. Hancock, E. Pekalska and R. Duin. Spherical Embeddings for non-Euclidean Dissimilarities. CVPR, to appear, 2010.

S. Ceolin and E.R. Hancock. Using the Fisher-Rao Metric to Compute Facial Similarity. ICIAR 2010, LNCS to appear 2010.

W. Xu, R.C. Wilson and E.R. Hancock. Rectifying Non-Euclidean Similarity Data using Ricci Flow Embedding. ICPR, to appear, 2010.

L. Han, R.C. Wilson and E.R. Hancock. A Supergraph-based Generative Model. ICPR, to appear, 2010.

S. Ceolin and E.R. Hancock. Characterising facial gender difference using Fisher-Rao metric. ICPR, to appear, 2010.

F. Escolano, M.A. Lozano and E.R. Hancock. Heat Flow-Thermodynamic Depth Complexity in Networks, ICPR, to appear, 2010.

S. Xia and E.R. Hancock. Mining Exemplars for Object Modelling using Affinity Propagation. ICPR, to appear, 2010.

P. Ren, T. Aleksic, R.C. Wilson and E.R. Hancock. Ihara Coefficients: A Flexible Tool for Higher Order Learning. S+SSPR, to appear, 2010.

H. ElGhawalby and E.R. Hancock. Graph Embedding using an Edge-based Wave Kernel. S+SSPR, to appear, 2010.

W. Xu, R.C. Wilson and E.R. Hancock. Regularisng the Ricci Flow Embedding. S+SSPR, to appear, 2010.

R.C. Wilson and E.R. Hancock. Spherical Embedding and Classification. S+SSPR, to appear, 2010.

S. Xia and E.R. Hancock. Mining Exemplars for Object Recognition using Affinity Propagation. S+SSPR, to appear, 2010.

N. Rahman and E.R, Hancock. Commute Time Convolution Kernels for Graph Clustering. S+SSPR, to appear, 2010.

F. Escolano, E.R. Hancock, D. Giorgi and M.A. Lozano. What is the Complexity of a Network? The Heat Flow-Thermodynamic Depth Approach. S+SSPR, to appear, 2010.

Julia E. Vogt, Sandhya Prabhakaran, Thomas J. Fuchs, Volker Roth. The Translation-invariant Wishart-Dirichlet Process for Clustering Distance Data. ICML 2010: Proceedings of the 27th international conference on Machine Learning. In press.

Joachim M. Buhmann. Information theoretic model validation for clustering. In International Symposium on Information Theory, Austin Texas. IEEE, 2010. In press.

Work package WP5:

Learning with non-(geo)metric similarities

Work package leader: UNIVE

Participants: UNIVE, IST

Start month: 7 *End month: 30* *Overall person-months: 45*

All approaches developed in WP2 and WP4 are based on the assumption that the non-geometricity of similarity information can be eliminated or somehow approximated away. When this is not the case, i.e., when there is significant information content in the non-geometricity of the data, alternative approaches are needed. The objective of this work package is to develop novel, general learning models which do not require the (geo)metric assumption, thereby working directly on the original data. Game theory offers an attractive and unexplored perspective that serves well our purpose. In this report, we briefly describe the main achievements and point to the relevant publications where these achievements are described in greater detail.

We aim at developing a game-theoretic framework based on a formalization of the competition between the hypotheses of class membership. According to this perspective, we shifted the focus from optima of objective functions to equilibria of (non-cooperative) games. A preliminary attempt towards this general goal was developed earlier by the Venice group, where (a variation of) the concept of Nash equilibrium is proposed as a formalization of the notion of a cluster in a non-symmetric pairwise clustering context. In the pervious year we concentrated on extending the proposed approach to grouping and matching, developing new efficient algorithms for extracting Nash equilibria, generalizing the approach to high-order and contextual similarities, and exploring new equilibrium concepts applicable to machine learning problems. In the second year of the project we concentrated on adapting the matching approach to the problem of inlier selection for robust estimation, we investigated a game-theoretic approach to the selection of distinguishing features from a set of loosely descriptive points, We extended the general framework to handle continuous k-way interactions and we investigate multi-population settings both for unsupervised and semi-supervised learning problems.

Inlier selection

A first development of the unsupervised learning approach was to develop a generic matching framework formulated as a competition between correspondence hypotheses [2]. e framework offers a natural equilibrium between representativeness and sparsity, leading to the ability to create a matcher geared to very low false positive rate, a characteristic that makes it particularly well suited for an inlier selection process for robust parameter estimation. Contrary to the usual approach which only use unary relation offered by the similarity between the features to be matched and enforce global consistency only as a afterthought, we use binary relations to drive the selection toward a set of globally consistent correspondences. In [6] the approach is used to create a coarse surface registration approach with precision comparable to fine registration. This work is generalized in [13] where the selection process is shown to be able to extract robust estimates even in conjunction with extremely poor unary relations. The inlier selection mechanism was applied to the extraction of 3d point correspondences for bundle adjustment in [5,8,11], using different global consistency models. Paper [6] resulted in a best student paper award at the Fifth International Symposium on 3D Data Processing, Visualization and Transmission – 3DPVT 2010. The selectivity properties of game-theoretic models was used to extract uncommon or highly distinctive features in [7,13].

Context-dependent similarities

We addressed the problem of extracting context-dependent similarities with an application to shape recognition [10]. The problem was cast as an EM-like iterative refinement process on similarity space, where both the cluster structure and the participation of each shape to the clusters was extracted using the game-theoretic framework and a generative shape-similarity model was used to update the contextual similarities.

High order k-way interactions

In [4] we explored the problem of extracting maximally coherent groups from a set of objects in the case when high-order (rather than pairwise) similarities are given. We showed that the hypergraph clustering problems can be naturally cast into a non-cooperative multi-player “clustering game”, whereby the notion of a cluster is equivalent to a classical game-theoretic equilibrium concept.

In [9] we studied the problem of clustering data objects in a similarity-based context, with the aim of grouping them into a given number of classes without imposing a hard partition, but allowing for a soft assignment of objects to clusters. The proposed approach uses the assumption that similarities reflect the likelihood of the objects to be in a same class in order to derive a probabilistic model for estimating the unknown cluster assignments. This led to a polynomial optimization problem in the probability domain, which was tackled using an evolutionary process with interactions among multiple-populations. This idea was further explored in [12] with an aim to devise a consensus clustering approach built upon the evidence accumulation framework. In a first learning stage, multiple clusterings of the data are combined into a single similarity matrix, called co-occurrence matrix. The probabilistic clustering approach is then used on this matrix in order to softly assign objects to clusters.

Semi-supervised learning

The multi-population approach was adopted to address semi-supervised learning problems, where graph transduction is formulated in terms of a non-cooperative multi-player polymatrix game where any equilibrium of the proposed transduction game corresponds to a consistent labeling of the data.

An attractive feature of our formulation is that it is inherently a multi-class approach and imposes no constraint whatsoever on the structure of the pairwise similarity matrix, being able to naturally deal with asymmetric and negative similarities alike. Further, some well-known approaches in the literature are shown to be special cases of the proposed approach [15].

References

1. A. Torsello and M. Pelillo, “Hierarchical Pairwise Segmentation Using Dominant Sets and Anisotropic Diffusion Kernels.” In EMMCVPR 2009.
2. A. Albarelli, S. Rota Bulò , A. Torsello, and M. Pelillo, “Matching as a Non-Cooperative Game.” In ICCV 2009.
3. M. Pelillo ”What is a Cluster? Perspectives from Game Theory.” In NIPS Workshop on Clustering 2009.
4. S. Rota Bulò and M. Pelillo “A Game-Theoretic Approach to Hypergraph Clustering.” In NIPS 2009.
5. A. Albarelli, E. Rodolà, and A. Torsello, “Robust Game-Theoretic Inlier Selection for Bundle Adjustment.” In 3DPVT, 2010.
6. A. Albarelli, E. Rodolà, and A. Torsello, “A Game-Theoretic Approach to Fine Surface Registration without Initial Motion Estimation.” In CVPR, 2010.
7. A. Albarelli, E. Rodolà, A. Cavallarin, and A. Torsello, “Robust Figure Extraction on Textured Background: a Game-Theoretic Approach.” In ICPR, 2010.
8. E. Rodolà, A. Albarelli, and A. Torsello, “A Game-Theoretic Approach to Robust Selection of Multi-View Point Correspondence.” In ICPR, 2010.
9. S. Rota Bulo and M Pelillo, “Probabilistic Clustering using the Baum-Eagon Inequality.” In ICPR, 2010.
10. A. Erdem and A. Torsello, “A Game Theoretic Approach To Jointly Learn Shape Categories and Contextual Similarities.” In S+SSPR2010.
11. E. Rodolà, A. Albarelli, and A. Torsello, “A Game-Theoretic Approach to Robust Selection of Multi-View Point Correspondence.” In S+SSPR, 2010.
12. S. Rota Bulò, A. Louren, A. Fred, M. Pelillo “Pairwise Probabilistic Clustering Using Evidence Accumulation” In S+SSPR2010.
13. A. Albarelli, E. Rodolà, and A. Torsello, “Loosely Distinctive Features for Robust Surface Alignment.” In ECCV2010 (accepted).
14. S. Rota Bulo', M. Pelillo and I. M. Bomze, “Graph-Based Quadratic Optimization: A Fast Evolutionary Approach.” Computer Vision and Image Understanding (conditionally accepted upon minor revision).
15. A. Erdem and M. Pelillo. Graph Transduction as a Non-cooperative Game (in preparation).

Work package WP6:

Analysis of tissue micro-array (TMA) images of renal cell carcinoma

Work package leader: ETH Zurich

Participants: UNIYORK, UNIVE, UNIVR, IST, TUD

Start month: 13 End month: 36 Overall person-months: 48

The objective of this workpackage is to apply the techniques developed in workpackages WP2 - WP5 to the analysis of Tissue Micro Array (TMA) images of renal cell carcinoma (RCC), which is one of the ten most frequent malignancies in Western societies. Current diagnostic rules rely on exact counts of cancerous cell nuclei that are manually counted by pathologists.

We consider an automated processing pipeline for tissue micro array analysis (TMA) of renal cell carcinoma. It consists of several consecutive tasks, which can be mapped to machine learning challenges. We investigate three of these tasks, namely nuclei segmentation, nuclei classification and staining estimation. We argue for a holistic view of the processing pipeline, as it is not obvious whether performance improvements at individual steps improve overall accuracy.

Empirical validation of the advances in the SIMBAD project would provide convincing evidence that similarity based methods can have advantages over traditional Euclidean feature based methods. In this work package, we investigate a medically important task of diagnosing renal clear cell carcinoma. We aim to supplement the currently labour intensive task of estimating staining proportions for various biomarkers. To make the task more convenient for the non medically inclined audience, we break this down to two sub-tasks that are more traditional machine learning tasks; namely, classification of cell nuclei and estimation of staining proportion. The relevant data and details of the individual tasks are provided at http://ml2.inf.ethz.ch/simbad/. Our efforts to date have been focused on creating easy to use datasets and estimating the performance of traditional feature based techniques.

In collaboration with the other project partners, in particular UNIVR, we organised a workshop in Zurich to discuss the two main validation platforms of the SIMBAD project, WP6 and WP7. The workshop also included presentations and discussions with the medical partners. This workshop resulted in a more streamlined approach to data processing and a clearer understanding of the various application tasks than previously. We also came to a consensus on the baseline kernels to be used with SVMs to provide an indication of the current state of the art performance on current data.

Accurate classification of cell nuclei and staining estimation are two of the most challenging tasks in computational pathology. Even for trained human experts the tasks remain difficult and yield large inter and intra pathologist variability. Both tasks are crucial for the overall performance of a TMA analysis pipeline.

Our progress to date are as follows:

(i) We demonstrated that graph cuts can be employed in conjunction with a circular shape prior to segment cell nuclei in a robust fashion.

(ii) The classification of nuclei into malignant and benign is not only feasible but additional shape features boost the classification performance.

(iii) We investigated and validated a large number of kernels, distance function and combination thereof. The lesson learned is that all extracted features are necessary for optimal performance and the best kernels perform significantly better than chance and comparable to human domain experts.

(iv) Finally, we demonstrated the influence of the classification task on the subsequent staining estimation problem. To this end improving classification performance on nuclear level is a challenging but worthwhile goal.

This status provides an opportunity in the final 12 months of the SIMBAD project to empirically investigate whether any of the advances improves classification accuracy. In particular, the advances in structural kernels from WP2 may provide better ways to compare the nuclei shapes, leading to more accurate similarities. These kernels are still positive semidefinite, and so provides an embedding into a Hilbert space. An alternative route is to use the advances of WP4 to investigate non-metric similarities by embedding them into Hilbert spaces. It would be very interesting to contrast the two approaches to obtaining kernels since they can be directly compared using SVMs. Furthermore, algorithms that directly work on non-metric similarities as developed in WP5 can be applied to the same data, and a deeper understanding of the importance of metricity may be obtained. In summary, we have provided a test bed for the various advances in the SIMBAD project by creating a curated set of data for a medically important task along with the performance of traditional machine learning approaches. We hope that in the near future, this test bed can be used to demonstrate the usefulness of non-metric similarity methods for machine learning.

Publications

T. J. Fuchs and J. M. Buhmann. Inter-Active Learning of Randomized Tree Ensembles for Object Detection. Workshop on On-line Learning for Computer Vision ICCV 2009.

I. Cima, R. Schiess, P. Wild, M. Kaelin, P. Schüffler, V. Lange, P. Picotti, R. Ossola, A. Templeton, O. Schubert, Th. Fuchs, Th. Leippold, S. Wyler, W. Jochum, J. Buhmann, Th. Cerny, H. Moch, S. Gillessen, R. Aebersold, W. Krek. Cancer genetics-guided discovery of serum biomarker signatures for diagnosis and prognosis of human prostate cancer [in submission]

P. Schüffler, Th.J. Fuchs, C.S. Ong, V. Roth, J.M. Buhmann. Computational TMA Analysis and Cell Nucleus Classification of Renal Cell Carcinoma [in submission]

Work package WP7:

Analysis of brain magnetic resonance (MR) scans for the diagnosis of mental illness

Work package leader: UNIVR

Participants: UNIYORK, UNIVE, IST, ETH Zurich, TUD

Start month: 13 *End month: 36* *Overall person-months: 51*

The objective of work package 7 (WP7) is to apply similarity-based techniques and algorithms developed in the other work packages to the analysis of brain magnetic resonance (MR) images in the context of mental health research (e.g., schizophrenia). Data consists of morphological MR images (3DA), instrumental to exploring the content of grey and white matter tissues, the volume of specific structures and the 3D shape morphology of particular brain regions, and diffusion weighted imaging (DWI) images, providing information on the microstructural integrity of the brain.

In the first year of the SIMBAD project, for WP7 we have released a preprocessed version of the data set for our partners to apply their algorithms and to act as baseline experiments for the problem. We devised several intermediate representations of the data set in this year; also we used multi-modal registration to use also the DWI data. We applied various pattern recognition algorithms to the data representations for a better solution to the problem. We established several collaborations with our partners in the project.

Some of the methodologies we apply are still under investigation and some have already been accepted (or submitted) to major conferences and journals in the area. We have 3 accepted conference papers and we have submitted one journal and one conference paper during this period. The details of all the methods are below; all % accuracies reported are Leave-One-Out accuracies.

Data processing

The dataset used in this work is composed of MRI brain scans of 64 patients affected by schizophrenia and 60 healthy control subjects. Images were acquired and transferred to PC workstations in order to be processed for Region of Interest (ROI) *tracing*. This latter procedure is the manual annotation of the images, performed by drawing contours enclosing the intended subpart of the brain. It is carried out by a trained expert following a specific protocol for each ROI. The data set consists of 7 ROIs manually traced by medical experts following a guideline and totally 14 ROIs have been extracted from the two hemispheres of the brain. These brain subparts are known to effect cognitive processes in the brain and we apply our methods on this subdivided data. From the various ROIs, gray value histograms are determined and in order to reduce the effect of inter-subjects intensity variations, the extracted histograms are properly normalized [Cheng *et* al. 2009b]. These histograms have been used in our preliminary baseline experiments; also this data (and a number of new representations created by applying several preprocessing techniques) has been delivered to our partners to apply their algorithms. Subsequently, several dissimilarity measures are defined between pairs of histograms which are described in the next section. We also delivered this data on the SIMBAD web site for our partners to apply their algorithms.

Dissimilarity matrices

With the histograms of intensities, we have applied several histogram and pdf dissimilarity functions to come up with distance matrices to be used with our partners in the SIMBAD project. The data set has been released on the SIMBAD web site along with an additional number of dissimilarities based on non-rigid registration of brains.

We used the following measures for calculating the dissimilarity between two histograms/pdfs: Bhattacharyya distance (pdfs), $\chi 2$ distance (histograms), Diffusion distance (histograms), Earth mover's distance (histograms + pdfs), Euclid distance (histograms + pdfs), Histogram intersection (histograms), Original Kullback-Leibler divergence which creates a nonsymmetric dissimilarity matrix (pdfs), Symmetrized KL divergence (pdfs), Jensen-Shannon divergence (pdfs), L1 distance (histograms + pdfs).

We also calculated the sum of deformations at each voxel and sum of length of deformations at each voxel as a dissimilarity measure for the non-rigid brain registration.

Using 13 (different measures) x 14 (different ROIs), we provided 182 dissimilarity matrices based on histograms and released this data on SIMBAD web site along with dissimilarities based on non-rigid brain registration for our partners to apply their algorithms.

Single-ROI based methods

These consist of our baseline experiments using this data. This part of the work has been done in the last five months of the first year, before WP7 started officially. Releasing the intensity histograms, we also did some preliminary classification experiments using histograms as features for classification algorithms [Cheng et al. 2009b]. In these preliminary experiments, we treat the ROIs independently and do not use the complementary nature of the ROIs. We used 6 different classification algorithms on histograms and preprocessed representations and reached 73.4 per cent accuracy with a single classifier. We then continued to make efforts on more advanced pattern recognition techniques. We have also calculated baseline 1nn accuracies for the dissimilarity matrices to act as a baseline for further dissimilarity based analysis. We reached 66.13% with histograms and 58.87% with registrations.

Multiple-ROI based methods

In these experiments, we applied multi-classifier pattern recognition techniques to combine the complementary nature of the ROIs. We designed experiments using various base classifiers and to combine the classifiers trained with these classifiers we have used. We applied fixed rules (Sum, Min/Max and Median) and an incremental algorithm for combining classifiers. We applied 13 classifiers using 5 different classification algorithms and different parameter sets.

As results of these experiments, we have seen that it's better to combine ROIs, not classifiers using the same ROI. In this setup, we have exploited the complementary nature of the ROIs and we have reached 75.81 % accuracy (single best classifier is 69.35). We have also tried subset selection on ROIs and classifiers to discard ROI-classifier pairs which do not contribute to the solution of the problem but make it less accurate. In this way, we have *again* seen that it does not change too much when we use a subset of classifiers with the same ROI but when we select a small subset of ROIs, we reach better accuracies (77.42) [Ulas *et* al. 2009].

Ongoing work

We have applied various methods to attack this problem. In this section, we will introduce some of them.

Generative kernels

These experiments are an application of methods developed in the context of WP2.1 (Generative kernels) to our data. In these experiments, we combine the power of generative and discriminative methods to come up with a hybrid generative/discriminative method and apply this method to our data set. We exploit state of the art generative models via Fisher kernel and support vector machines (SVM).

Generative and discriminative approaches are the two broad categories within which learning and classification methods fall: a generative approach will estimate the joint probability density function (pdf) of the data and class labels and will classify using the posterior probabilities obtained by Bayes' rule, while a discriminative approach will estimate a classification function directly. The basic idea is to employ a generative model to define feature vectors and project objects to the resulting feature space. Therefore, a meaningful similarity/distance measure is defined, leading to a kernel.

For the generative part, the general idea is to choose a generative model capable of considering all the ROI at the same time, together with the relations between them. To this end, we based our framework on the same concepts behind the constellation probabilistic model, which foresees the encoding of one object in terms of a fixed number *N* of object subparts *Mj*, and relative spatial relationships. We used the ROIs for this purpose.

Fisher kernels allow an effective general way of mixing generative and discriminative models for classification. In particular, the Fisher kernel approach measures the similarity between the objects by comparing them in the tangent space induced by the trained generative model, which is considered as a point in the Riemannian manifold defined by the chosen family of generative models.

We have seen that our hybrid approach is better than the pure discriminative methods. The best single ROI accuracy is 70.97 and when we combine all ROIs, we have 77.42 % accuracy. Our hybrid method achieves 80.65 % accuracy [Cheng *et* al. 2009c].

Feature based morphometry

In this part, we used the DLPFC region of the brain for our experiments. Few and significant landmarks are detected and characterized by local region descriptors. We focus on the Scale Invariant Feature Transform (SIFT) operator, which has been proposed in computer vision for object recognition and it has been successfully used also on medical applications for 3D deformable image registration. The underlying hypothesis consists of relaxing the common constraint that morphological anomalies appear at the same voxel location for all the population. Therefore, a new kernel of a SVM is designed in order to allow the comparison between a pair of brains represented by an unordered set of features. The proposed method is inspired by the Bag-of-Words paradigm which implicitly implements the feature matching within the SVM framework. Such kernels are known as local kernels in order to emphasize the fact that local information is used to characterize the involved objects. Finally, a weighting function is introduced to define the relevance of the detected groups of features, namely the visual words, in discriminating among the two populations (i.e., patients and controls). Moreover, the proposed approach is able to take into account of different morphological abnormalities at the same time, which are possible spread onto the analyzed ROI. We used a smaller subset of our data set because the preliminary parts of these experiments were done before the data was released. We used 54+54 healthy and control subjects and achieved 75 % accuracy. One other aspect of these experiments is that we divided the data set according to demographic information and we have seen that we can reach 84.09 % accuracy when we use only women and the left hemisphere of the brain. The results are also promising for seniors; the accuracy is 81.25 % [Castellani *et* al. 2009].

Multi-modal registration and analysis

The second modality of our data set (DWI) does not come with segmented ROIs. In order to combine the two modalities and exploit the complementary nature of the structural and functional parts of the brain, we first used registration between the morphological and diffusive parts of the brain. We extracted histograms for the ADC (Apparent Diffusion Coefficient) and used these in our preliminary experiments. We registered 114 of the 124 subjects and we trained 1nn, SVM (with linear and RBF kernels) for preliminary experiments. Although single DWI accuracies are not as good as the accuracies of the morphological representations, we believe that combining both modalities, we will achieve better results. After our preliminary experiments with this data, we are ready to deliver it to our partners.

Collaborations

With the release of our data and further intermediate representations and the SIMBAD workshop in Zurich, we established/improved collaborations with our partners. In this section we will summarize these collaborations.

IT Kernels

In this part of our experiments, we applied IT kernels developed by IST group in the context of WP2 and have preliminary results using the kernels as similarities for a knn classifier. We compared this method with linear an RBF kernels. We applied four kernels which are: Jensen Shannon (JS), Jensen Tsallis (JT), and Weighted JT-v1 (Weighted version of JT kernel) and Weighted JT-v2 kernels.

By applying this method and using knn as base classifier, we have carried out our experiments on intensity histograms and geometric shape descriptor histograms (shape index, min/max curvature) we achieved results which were better than using the data directly. Although the accuracies are not as large as our previous experiments (best accuracy is 66.32 %), knn is not the best way to use the similarities and kernels in this context and we plan to apply generative embeddings and support vectors for better accuracy.

Dissimilarity space

This work is a joint work with TUD group and it has been submitted to ICPR workshop on Brain Decoding. We applied the dissimilarity combination methods developed under the context of WP3. In this direction, we used the 182 distance matrices based on histogram and pdf similarities. In a second step we transformed the dissimilarity representations into a vector space in which traditional statistical classifiers can be employed. Unlike the related kernel approach, whose application is often restrained by technicalities like fulfilling Mercer's condition, basically any dissimilarity measure can be used in this new vector space. Since in this space, number of features is equal to the number of subjects, we used support vectors as classifiers. We also used knn algorithms using the original dissimilarities as comparison purposes. We applied two different sets of experiments using this setup. In the first set of experiments, we used a Single-ROI approach. We have seen that the dissimilarity space approach clearly has better results than directly using the classifiers. Best 1nn result is 66.90 %, whereas best SVM result is 76.60 %. In the second set of experiments we used multiple ROIs and we combined the dissimilarities by simple averaging. Before averaging, every dissimilarity matrix was scaled so that the average dissimilarity was 1. This time, we combine every ROI using one similarity measure. We again see that the SVM results are better than 1nn when we combine the ROIs using the same dissimilarity measure. The best result is achieved when we combine all measures and all ROIs [Ulas *et* al. 2010]. This encourages us to further investigate this direction.

Other collaborations

In the context of WP7, we have established collaborations and we are trying to improve these efforts. The workshop for WP6 and WP7 in Zurich was a success for data format standardization, sharing knowledge and implementations, setting up new collaborations and improving existing ones, discussing possible new ideas for the solution of the problems and trying to integrate the medical knowledge into machine learning world. With the York group, we started to work on the dissimilarity matrices of the pair-wise non-rigid registrations. With the Venice group, we started to work on graph based representations and dissimilarities, as well as Feature-based morphometry classification by proposing new descriptors.

Future work

As a future work, we plan to investigate in detail the multiple modalities approach to gain from both structural and functional subparts of the brain and exploit the complementary information. We plan to release those data.

We are currently working on deformation based registration of brain images and extracting distances directly from these representations. This method has the advantage of releasing dissimilarity information without doing further processing. Currently we are doing pair-wise registration but we also plan to divide the data set according to demographic information (age/sex) and create one baseline image for every set and register all the subjects to these base images.

We will try 3D shape descriptors; in particular Surface based morphometry, to define new features from the MRI images. We also plan to apply generative kernels based on topic models and coordinate and consolidate the efforts of our partners towards a solution to this problem.

Publications

[Cheng *et* al. 2009a] Cheng, D. S., Bicego, M., Castellani, U., Cerruti, S., Bellani, M., Rambaldelli, G., Atzori, M., Brambilla, P., Murino, V., 2009. Schizophrenia classification using regions of interest in brain MRI. Tech. rep., Dipartimento di Informatica, University of Verona, Italy.

[Cheng *et* al. 2009b] Cheng, D. S., Bicego, M., Castellani, U., Cerruti, S., Bellani, M., Rambaldelli, G., Atzori, M., Brambilla, P., Murino, V., 2009. Schizophrenia classification using regions of interest in brain mri. In: Proceedings of Intelligent Data Analysis in Biomedicine and Pharmacology, IDAMAP '09. pp. 47–52.

[Cheng *et* al. 2009c] Cheng, D. S., Bicego, M., Castellani, U., Cristani, M., Cerruti, S., Bellani, M., Rambaldelli, G., Atzori, M., Brambilla, P., Murino, V., 2009. A hybrid generative/discriminative method for classification of regions of interest in schizophrenia brain mri. In: Proceedings of workshop on Probabilistic Models for Medical Image Analysis, MICCAI '09. pp. 174–184.

[Ulas *et* al. 2009] Ulas, A., Castellani, U., Bicego, M., Cheng, D. S., Rambaldelli, G., Bellani, M., Brambilla, P., Murino, V., 2009. A multiple region of interest analysis for identifying schizophrenia using mri, submitted to Neuroimage, special issue on “Multivariate Decoding and Brain reading”.

[Ulas *et* al. 2010] Ulas, A., Duin, R. P., Castellani, U., Loog, M., Bicego, M., Bellani, M., Rambaldelli, G., Murino, V., 2010. Dissimilarity-based detection of schizophrenia, submitted to ICPR workshop on “Brain Decoding: Pattern Recognition Challenges in FMRI Neuroimaging”.

[Castellani *et* al. 2009] Castellani, U., Rossato, E., Murino, V., Bellani, M., Rambaldelli, G., Tansella, M., Brambilla, P., 2009. Local kernel for brains classification in schizophrenia. In: AI*IA '09:: Proceedings of the XIth International Conference of the Italian Association for Artificial Intelligence Reggio Emilia on Emergent Perspectives in Artificial Intelligence. Springer-Verlag, Berlin, Heidelberg, pp. 112–121.

Work package WP8:

Dissemination, communication and exploitation

Work package leader: UNIVE

Participants: UNIVE, IST

Start month: 7 *End month: 36* *Overall person-months: 7*

The consortium is ensuring the highest diffusion of the research results, both inside and outside SIMABD.

In order to strengthen the communication among the partners, we have set up a blog on the project site (https://www.dsi.unive.it/~simbadweb/blog/wp-login.php). This is a tool for exchanging information and a way to keep track of interactions among the groups and having them accessible from a centralized repository. We are also making extensive use of the *SIMBAD Technical Report Series* as a tool to provide a timely access of information within the consortium, and increase interactions among the SIMBAD partners (in the first half of 2010 we produced 28 TR's, and in 2009 we produced 12).

To strengthen internal collaborations, during the second year we had many exchange visits among partners which resulted a in a few joint publications. We run two project meetings (Zurich, 26-27 June 2009; Verona, 19-20 November 2009) and and a one-week preparatory meeting focused on our biomedical applications (Zurich, 1-5 February 2010). In July 2010 we'll also be running a "hands-on" internal workshop in Castelbrando, Treviso, Italy.

The external dissemination of the project's results during this period took place mainly through publications in the top-level conferences and journals in the fields of machine learning, pattern recognition and computer vision. We also plan to disseminate the results related to the project's main applications to the medical and the chemometrics communities. The entire list of publications is available in the Annex – List of publications, publications referred to each work package are also listed at the end of the work packages described above).

The following dissemination activities have also been accomplished:

- In September 2009 we presented SIMABD at the Project Exhibition of ECML PKDD 2009, The European Conference on Machine Learning and Principle and Practice of Knowledge Discovery in Databases, which was held in Bled, Slovenia (http://www.ecmlpkdd2009.net/).
- In January 2010 we run a special session devoted to "Learning and Intelligent Optimization in Structured Domains (LIONS)" at LION 2010: The 4th International Conference on Learning and Intelligent OptimizatioN, which was held in Venice (http://www.intelligent-optimization.org/LION4).
- In June 2010 we'll be running a workshop at ICML 2010 (the 27th International Conference on Machine Learning -- one of the leading conferences in machine learning) devoted to "learning in non-(geo)metric spaces (see http://www.dsi.unive.it/~icml2010lngs/). The meeting is sponsored by PASCAL 2 and will have videolectues coverage.

- In Augustt 2010 we'll be running a Special session at SSPR 2010, the IAPR Workshop on Statistical, Structural and Syntactic Pattern Recognition (see: http://www.rvg.ua.es/ssspr2010/special.php), which will be held in Cesme, Izmir, Turkey. The meeting will have videolectures coverage.
- A tutorial will be presented at ICPR 2010 on WP5's Topic: *Game Theory in Pattern Recognition and Machine Learning* (speakers: M. Pelillo and A. Torsello) (see http://www.icpr2010.org/tutorials.php).

We aim at establishing a series of workshops specifically devoted to the project's theme. In the hope that the topics covered within the SIMBAD project will have a lasting and substantial impact within scientific community, the workshops will possibly continue even after the end of the project activities. We are planning to run the first edition of the workshop "*SIMBAD 2011"* in Italy, in the spring of 2011, at the end of the project (in conjuction with the final SIMBAD meeting).

We're also planning to run a journal special issue devoted to the project's theme and at the end of project we plan to publish a book containing the major results achieved by the consortium, to be published by a major international scientific publisher.

Finally, we mention that partners UNIVE and TUD are now member of the PASCAL 2 Network of Excellence.

Project planning and status

The work carried out during the second year of the project is essentially in line with Annex I.

As concerns, WP2, the work in task WP2.3, although partially addressed during the second year, did not lead to definitive conclusions. Further, a fundamentally new information-theoretic model validity perspective emerged at ETH. Hence, we request a **one-year extension** for this WP in order to study the implication of these ideas. There will be no additional cost for the project, and there will be no impact on the other tasks.

As concerns WP3, as anticipated at the 1st review meeting, the production of D3.4 was postponed of a few months. Indeed, being a summary of what we achieved in WP3, it was scheduled for the same date as D3.3, the creation of a collection of datasets and tools (a milestone), which was produced on time (in october 2009). But we wondered whether after this milestone some other datasets in SIMBAD would pop up that could be considered as well and integrated in the final report. This in fact happened: the Verona BrainMRI dataset. It was however so huge that we wrote a separate report on it. The final report, D3.4, just mentions it and further summarizes the previous reports. We also took the opportunity to list a number of other papers (some converted into reports as well, i.e. TR 2009 n. 9 and 10, 2010 n. 19, 20, 21and 22) as output of WP3.

Concerning WP4, we request a zero-cost **6-month extension** in order for the two PhD students to take their work to conclusion on a) Ricci flow regularisation and b) learning generative and compressive graph kernels (with IST and UNIVR). This will have added benefit that they can complete their theses under stipend.

Concerning WP5, some delay was caused by: (1) post-doc (A. Erdem) who gave up the project prematurely; diificulty in hiring post-doc for WP6/7 (J. Hou) we request a zero-cost **6-month extension** in order to satisfactorily complete the job.

Cocnernng WP7, the preprocessing phase was under-estimated and took more time than expected to prepare data to be used by the partners. We also spent considerable person-power in this period. One reason is that the hiring of junior scientists made more difficult to reach the envisaged objectives in the planned period.

Personnel

UNIVE: Assigned to the SIMBAD project are Marcello Pelillo (associate professor), heading the UNIVE contribution, Andrea Torsello (assistant professor) and, Aykut Erdem (post-doc since November 2008) and Jian Hou (post-doc since October 2009). As a result of a previous cooperation on the topics of SIMBAD, Samuel Rota Bulò (postdoc at the University of Venice) and Andrea Albarelli are also actively participating in the project. Veronica Giove, project administrator, assisting the project coordinator and Sonia Barizza, the Administrative Responsible of the Computer Science Department, deal with management activities .

UNIYORK: Assigned to Simbad project are Professor Edwin Hancock, full professor; Lin Han, (Since 1/10/08 PhD student), Eliza Xu (Since 1/10/08 PhD student).
The following people were involved in the SIMBAD project research, but have not been paid with SIMBAD funds: Dr Richard Wilson, Reader; Peng Ren, PhD student; Howaida El Ghawalby, PhD student.

TUD: Assigned to the SIMBAD project are Wan-Jui Lee (100%, post-doc) from the start of the project and Robert P.W. Duin (partially, associated professor), heading the Delft contribution. As a result of a historical cooperation also Elzbieta Pekalska (research associate at the University of Manchester, UK) is actively participating in the project. Starting in September 2008 Alessandro Ibba is working as a PhD student in the field of dissimilarity based pattern recognition as well. He is not payed from the project. Finally, Marco Loog (assistant professor) and Marcel J.T. Reinders (full professor) have been following the project.

IST: Assigned to the SIMBAD project are Mario Figueiredo, Ana Fred, Pedro Aguìar (faculty members) and André Martins, David Pereira-Coutinho, João Duarte, André Lourenço and Liliana Medina (PhD students)

UNIVR: Assigned to the SIMBAD project are Vittorio Murino, Umberto Castellani, Manuele Bicego, Marco Cristani (faculty members), Dong Seon Cheng (post-doc), Anna Carli, Alessandro Perina, and Cerruti Stefania, Ulas Aydin (PhD students).

ETH Zurich: During the second year the two PhD students Sharon Wulff (since Nov. 15, 2008) and Peter Schüffler (since Dec. 1, 2008) worked on the SIMBAD project in the ETH Zurich group. This work has been jointly supervised by Volker Roth, who is a professor at the University of Basel, Dr. Cheng Soon Ong, research associate at ETH Zurich and Prof. Joachim M. Buhmann. Peter Schüffler has worked fulltime on the SIMBAD project and Sharon Wulff was employed part-time on the SIMBAD project.

4. Deliverables and milestones tables

TABLE 1. DELIVERABLES									
Del. no.	**Deliverable name**	**WP no.**	**Lead beneficiary**	**Nature**	**Dissemination level**	**Delivery date from Annex I (proj month)**	**Delivered Yes/No**	**Actual / Forecast delivery date**	**Comments**
D1.2	First Periodic Report	1	1	R	PU	13	Yes	13	
D2.2	Generative Kernels	2	5	R	PU	18	Yes	18	
D8.4	First open scientific event	8	1	O	PU	18	No	27	NIPS2009 SIMBAD workshop was not accepted
D4.1	Spectral and geometric embeddings	4	2	R	PU	18	Yes	18	
D3.3	Database of non-(geo)metric data	3	3	O	PU	18	Yes	18	
D3.4	Final Work package report	3	3	R	PU	18	Yes	25	See page 38 of this section
D8.5	Second dissemination plan	8	1	R	PU	22	Yes	22	
D2.3	Learning similarities from examples	2	4	R	PU	24		36	See page 10 of this

									section
D2.4	Final work package report	2	4	R	PU	24		36	See page 10 of this section
D4.2	Structure preserving embeddings	4	6	R	PU	24	Yes	24	
D7.1	Mid-term report	7	5	R	PU	24	Yes	24	
D6.1	Mid-term report	6	6	R	PU	24	Yes	24	

Milestones

TABLE 2. MILESTONES							
Milestone no.	**Milestone name**	**Work package no**	**Lead Beneficiary**	**Delivery date from Annex I**	**Achieved Yes/No**	**Actual / Forecast achievement date**	**Comments**
MS 3	Structural Kernels	2	4	18	Yes	18	
MS 4	Database of non-geometric similarities	3	3	18	Yes	18	
MS 5	Learning similarities	2	4	24	No	24/36	We have only partially achieved this milestone. There will be no substantial impact on the other parts of the project, though. We plan to achieve this milestone in the third year of the project

MS 6	Embeddings	4	2	24	Yes	24	
MS 7	Assessment of approached to the TMA domain	6	6	24	Yes	24	
MS 8	Assessment of approached to the MA domain	7	5	24	Yes	24	

5. Project management

Management tasks and achievements

Financial distribution: University Ca' Foscari of Venice (UNIVE) acting as the Coordinator, distributed the first Community instalment among the beneficiaries, on the basis of each cost claim declared by each partner, without any delay.

Bank account: Even though the Computer Science Department keeps one single bank account, it is possible, at any time, to determine the actual balance and the flow of money and the quarterly return of interests gained on the prefinancing coming from European Commission funds, which is still generating interests.

Monitoring the project implementation: UNIVE is monitoring the implementation of the project, looking at the compliance of the partners obligations: activities progress, sharing of the results, settlement of the working groups, deliverables (D1.2 First Periodic Report, D2.2 Generative Kernels, D8.4 First open scientific event, D4.1 Spectral and geometric embeddings, D3.3 Database of non-(geo)metric data, D3.4 Final Work package report, D8.5 Second dissemination plan, D2.3 Learning similarities from examples, D2.4 Final work package report, D4.2 Structure preserving embeddings, D7.1 Mid-term report, D6.1 Mid-term report).

Communication of data for evaluation: UNIVE together with all the SIMBAD partners collaborated with the DG Information Society & Media to assess the progress in 2009 towards the achievement of the IST-RTD implementation objectives, gathering the results of the SIMBAD project.

Problems which have occurred

No problem arose from the management of the consortium. All discussions have been carried out during the project meeting, and/or via e-mail.

Changes in the consortium

No amendment of the Consortium has been necessary

List of project meetings, dates and venues

Project meetings:

1. The first Review meeting took place in Venice on the 24th June, 2009

2. The third project meeting took place in Zurich (CH), at the ETH Zurich, on the 26 and 27th June 2009.
3. The fourth project meeting took place in Verona (IT), at the UNIVR, on the 19 and 20th November 2009.
4. A one-week preparatory meeting took place in Zurich, during the 1st and -5th of February 2010, it was focused on WP6 & WP7 research exchanges

Project planning and status

The work carried out during the second year of the project is essentially in line with Annex I.

As concerns, WP2, the work in task WP2.3, although partially addressed during the second year, did not lead to definitive conclusions. Further, a fundamentally new information-theoretic model validity perspective emerged at ETH. Hence, we request a **one-year extension** for this WP in order to study the implication of these ideas. There will be no additional cost for the project, and there will be no impact on the other tasks.

As concerns WP3, as anticipated at the 1st review meeting, the production of D3.4 was postponed of a few months. Indeed, being a summary of what we achieved in WP3, it was scheduled for the same date as D3.3, the creation of a collection of datasets and tools (a milestone), which was produced on time (in october 2009). But we wondered whether after this milestone some other datasets in SIMBAD would pop up that could be considered as well and integrated in the final report. This in fact happened: the Verona BrainMRI dataset. It was however so huge that we wrote a separate report on it. The final report, D3.4, just mentions it and further summarizes the previous reports. We also took the opportunity to list a number of other papers (some converted into reports as well, i.e. TR 2009 n. 9 and 10, 2010 n. 19, 20, 21and 22) as output of WP3.

Concerning WP4, we request a zero-cost **6-month extension** in order for the two PhD students to take their work to conclusion on a) Ricci flow regularisation and b) learning generative and compressive graph kernels (with IST and UNIVR). This will have added benefit that they can complete their theses under stipend.

Concerning WP5, some delay was caused by: (1) post-doc (A. Erdem) who gave up the project prematurely; difficulty in hiring post-doc for WP6/7 (J. Hou) we request a zero-cost **6-month extension** in order to satisfactorily complete the job.

Concerning WP7, the preprocessing phase was under-estimated and took more time than expected to prepare data to be used by the partners. We also spent considerable person-power in this period. One reason is that the hiring of junior scientists made more difficult to reach the envisaged objectives in the planned period.

Personnel figures

TUD

TUD has spent already more than the originally estimated man-months to the project. This is caused by the fact that the SIMBAD research strongly overlaps with the free university research and that they cannot be distinguished in the university administration. It is of course the intention of TUD to contribute further to SIMBAD and focus on the applications in WP6 and WP7 and to the new insights this may produce to the topics studied in WP3 and WP4.

In order to formally let us give our contribute to the project we would like to move EUR 20.000 from "Other costs" to "Personnel costs", without any additional busget.

The following table summarizes the request:

	Balance (III Year)	Updated figures for the III year
Personnel costs	€ 4.741	€ 24.741
Subcontracting	€ -	
Other direct costs	€ 30.574	€ 10.574
Indirect costs	-€ 290	-€ 290
Total costs	**€ 35.025**	**€ 35.025**
Requested EC	€ 26.269	€ 26.269

UNIVR

In the second year of the project, we concluded the work in WP2 with a slight increase of the person-power originally planned (31 vs 28).This is mainly due to the efforts devoted to finalise the collaborative work with the other partners of the WP2 (IST, mainly).

Concerning the WP7, we have spent considerable person-power in this period. One reason is that the hiring of junior scientists made more difficult to reach the envisaged objectives in the planned period. Moreover, we underestimated the work needed to pre-process the data and make it available to the other project partners. To conclude the work in WP7, we estimate to spend approximately a further extra 3 person months.

In order to partially cover these changes, we plan to move 20.000 EURO from the item "Other costs" to the item "Personnel", so that we can finalise our own work in collaboration with the other partners.

However, in any case, we foresee that the total cost of the activities carried out by UNIVR will remain the same of the original budget.

The following table shows the updated budget figures

	Balance (III Year)	Updated figures for the III year
Personnel costs	€ 25.279	€ 45.279
Subcontracting	€ -	
Other direct costs	€ 40.477	€ 20.477
Indirect costs	€ 39.454	€ 39.454
Total costs	**€ 105.211**	**€ 105.211**
Requested EC	€ 78.908	€ 78.908

ETH Zurich

Concerning ETH Zurich, during the second year of the SIMBAD-project we decided to allocate more than the planned resources to complete WP6 in the projected time period. We are expecting more data from our medical collaborators and it takes more time than expected to evaluate the whole pipeline process for the Tissue-Micro-Array Analysis. In addition, our hired junior scientists needed some training time before being able to work on WP6. The publication process for two SIMBAD publications was also time intensive. Therefore, we need to spend approximately four additional person months to finish the work package. One further man month is required to fulfil the unexpectedly time intensive reporting and SIMBAD-project administration.

To cope with the changes we would like to move EUR 20500 (claim) from ”Other costs” to “Personnel costs”. Nevertheless we are not intending to exceed the ETH budget for the project.

The following table summarizes the request:

	III Year budget	Updated figures for the III year
Personnel costs	€ 32.250	€ 52.750
Subcontracting	€ -	€ -
Other direct costs	€ 41.970	€ 21.470
Indirect costs	€ 44.532	€ 44.532
Total costs	**€ 118.752**	**€ 118.752**
Requested EC	€ 89.064	€ 89.064

Comparison between actual person months spent in the first 24 months and the total planned resources

Actual **Planned**

		WP 1	WP 1	WP 2.1	WP 2.1	WP 2.2	WP 2.2	WP 2.3	WP 2.3	WP 3.1	WP 3.1	WP 3.2	WP 3.2	WP 4.1	WP 4.1	WP 4.2	WP 4.2	WP 5.1	WP 5.1	WP 5.2	WP 5.2	WP 6	WP 6	WP 7	WP 7	WP 8	WP 8	TOTAL PM
TOT spent vs Balance (Tot. Allocated - tot. Spent)	UNIVE	9,7	12,0					0,5	5,0					0,3	5,0			3,7	15,0	16,5	25,0	3,2	6,0	3,2	6,0	0,5	2,0	**37,5**
	UNIYORK						2,0							37,5	40,0		2,0						11,0		11,0		1,0	**37,5**
	TUD			1,0						13,5	12,0	4,4	6,0	2,0	2,0	7,0	2,0					3,5	4,0	3,0	4,0	0,5	1,0	**34,9**
	IST			5,0	5,0	5,0	5,0	5,0	5,0											2,4			5,0	3,4	5,0		1,0	**20,8**
	UNIVR			18,6	18,6	12,4	12,4															3,0	7,0	13,0	23,0		1,0	**47,0**
	ETH Zurich					2,0	2,0			3,0	3,0	8,0	8,0	1,8	2,0	3,4						8,5	15,0	1,0	5,0	0,5	1,0	**28,2**
	tot	**9,7**	**12,0**	**24,6**	**23,6**	**19,4**	**21,4**	**5,5**	**10,0**	**16,5**	**15,0**	**12,4**	**14,0**	**41,6**	**49,0**	**10,4**	**4,0**	**3,7**	**15,0**	**18,8**	**25,0**	**18,2**	**48,0**	**23,6**	**54,0**	**1,5**	**7,0**	**205,9**

Comparison between actual expenses and budget of the action

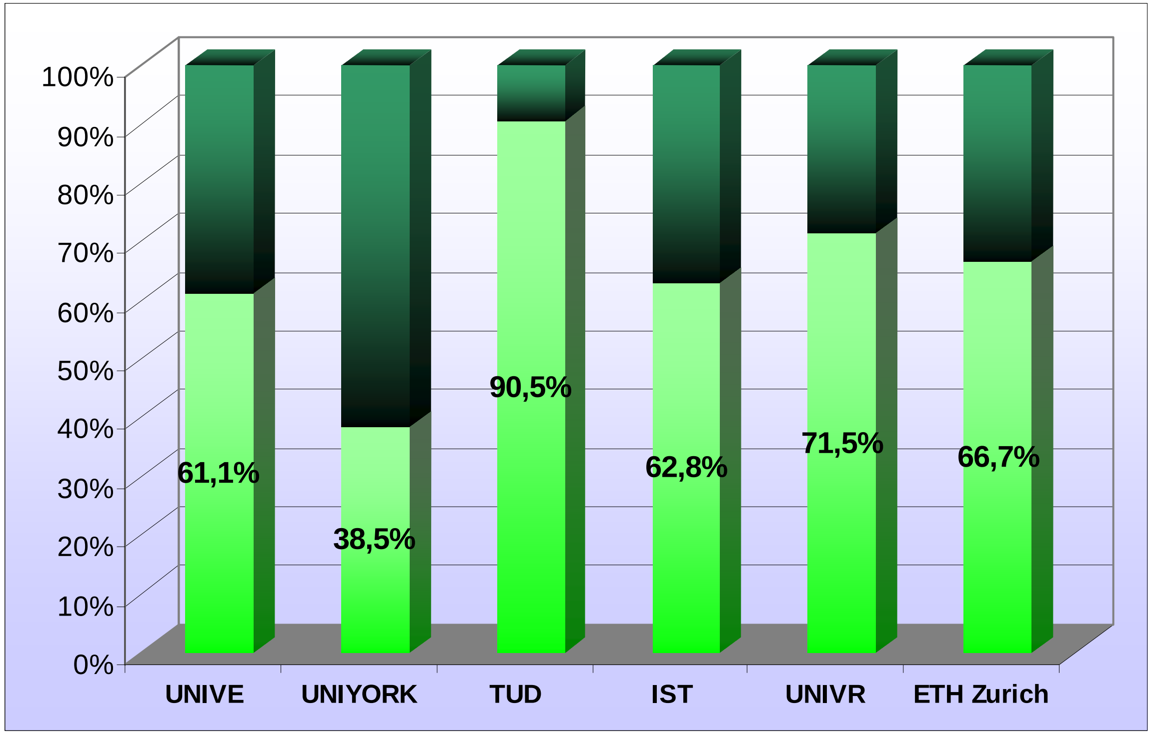

Activities foreseen in the project, in the near future are the following:

1. Fourth meeting, scheduled for the 4th July 2010, in Castelbrando (Treviso, Italy).
2. Hands on meeting from 5th to 10th July 2010

Impact of possible deviations from the planned milestones and deliverables, if any

The deviation from the planned milestone and deliverables will have no significant impact on the rest of the project.

All other deliverables and milestones foreseen in the project have been achieved.

Changes to the legal status of any of the beneficiaries

No change in the legal status of the beneficiaries is to be reported

Development of the Project website

The dedicated web site has been developed during this second year. In the website project technical reports, software, data, and results, both theoretical and technological, have been published and illustrated. An important part of the site is the collection of similarity data creating a common reference for similarity-based computational models and algorithms. simbad-fp7.eu serves as a knowledge-base for similarity-based approaches to pattern recognition and machine learning.

Information about project development are updated by each consortium member thanks to his own restricted web site area.

Two sections have been added after the first year,

1. Technical reports: it is a repository where all the papers presented at different international conferences would be issued as Technical Report.
2. Simbad Blog: this is a tool for exchanging information and a way to keep track of interactions among the groups and having them accessible from a centralized repository.

The other sections have been duly updated during the second year.

Use of foreground and dissemination activities during this period

Participation in International Conferences: the positive results of the project activities have been disseminated within the scientific community through the participation of some of the members of the consortium in three International conferences:

May 2009, GameNets 2009 – Istanbul (Turkey)
June 2009, MCS 2009 – Reyjkjavik (Iceland)
June 2009, ICML – Montreal (Canada)
July 2009, IDAMAP 2009 – Verona (Italy)
August 2009, EMMCVPR 2009 – Bonn (Germany)
September 2009, ECML PKDD 2009 – Bled (Slovenia)
October 2009, ICCV 2009 – Kyoto (Japan)
October 2009, ICCV 2009 SUBSPACE workshop – Kyoto (Japan)
September 2009, PMMIA 2009 workshop (MICCAI 2009) – London (United Kingdom)
September 2009, CAIP 2009 – Münster (Germany)
December 2009, participation to the Conference NIPS 2009 and associated Workshops– Vancouver (Canada)
December 2009, participation in the Conference AIIA 2009 – Reggio Emilia (Italia)

The coordinator, Marcello Pelillo, requested (and obtained) to the Project Officer the agreement to attend to conference held in non European Countries.

Details on the participation of Simbad staff in the before listed conferences:

- From **UNIVE**:

Marcello Pelillo presented one paper at NIPS 2009
Andrea Torsello presented one paper at ICCV 2009 and one at EMMCVPR 2009
Samuel Rota Bulò presented one paper at NIPS 2009 (he also participated in the Game Theory Summer School – 31/05-06/06/2009 – in Bertinoro, Italy).

- From **UNIYORK**:

S. Xia presented one paper at CAIP 2009
Peng Ren presented his paper at ICCV 2009 and at CAIP 2009

- From **TUD**:

Wan-Jui Lee presented one paper at MCS 2009,

- From **UNIVR**:

Dong Seon Cheng presented his paper at IDAMAP 2009
Manuele Bicego presented one paper at EMMCVPR 2009 and one paper at ICCV 2009
Umberto Castellani presented one paper at PMMIA 2009
Alessandro Perina presented his paper at NIPS 2009 and another paper at ICCV 2009
Umberto Castellani presented one paper at AIIA 2009

- From **ETH Zurich**:

Joachim Buhmann presented his paper at NIPS 2009
Cheng Soon Ong presented his paper at ICML 2010

List of publications:

In the second year the Simbad partners produced 39 publications related to the project topics. In the "Annex A – SIMBAD list of publications", we listed only the publications carried out thanks to the SIMBAD funds (which include proper acknowledgements), some of them have been issued as Technical report, included in the Simbad TR series.

Press Coverage: The activities related to the SIMBAD project have been advertised via several interviews which are available at the project website (http://simbad-fp7.eu/news.php).

6. *Explanation of the use of the resources*

TABLE 3.1 PERSONNEL, SUBCONTRACTING AND OTHER MAJOR DIRECT COST ITEMS FOR BENEFICIARY 1 UNIVERSITÀ CA' FOSCARI VENEZIA FOR THE PERIOD FROM 01/04/2009 TO 31/03/2010

Work Package	Item description	Amount	Explanations
1	Personnel costs	*€ 16.566*	Salaries of: Assistant Professor for 0,47 PM (Marcello Pelillo) administrative responsible for 0,7 months (Sonia Barizza) recruited collaborator for 3,7 months (Veronica Giove)
2.3, 4.1, 5.1, 5.2, 6, 7, 8	Personnel costs	*€ 55.336*	Salaries of: Assistant Professor (Marcello Pelillo) for 1,37 PM Assistant Professor (Andrea Torsello) for 1,41 PM one recruited post-doc (Aykut Erdem) for 12 PM one recruited post-doc (Jian Hou) for 6 PM
2.3, 4.1, 5.1, 5.2, 6, 7, 8	Travel costs	*€ 17.841*	Participation in project meetings, national and international meetings and conferences
2.3, 4.1, 5.1, 5.2, 6, 7, 8	Equipment Purchase	*€ 4.727*	Purchase of one server, laptop, desktop computer to implement the specif research
1	Remaining direct costs	*€ 1.758*	Travel for the project meeting
2.3, 4.1, 5.1, 5.2, 6, 7, 8	Remaining direct costs	*€ 2.663*	Minor items
	TOTAL DIRECT COSTS	***€ 98.890***	

Personnel

Marcello Pelillo, associate professor, heading the UNIVE contribution was involved for 0,47 PM in the Management activities (WP1) and for 1,37 PM in the RTD activities (WP2.3: 0,5 pm; WP5.1: 0,19PM; WP5.2: 0,19PM; WP8: 0,5PM)

.

Andrea Torsello, assistant professor, worked for 1,41 PM on the project (WP4.1: 0,3PM; WP5.1: 0,5PM; WP5.2: 0,2PM; WP6: 0,2PM; WP7: 0,2PM)

Aykut Ibrahim Erdem, post-doc is 100% charged on the project (WP5.2)

Jian Hou,, post-doc since October 2009 is 100% charged on the project (WP6: 3PM; WP7: 3PM)

Veronica Giove, project administrator, worked for 3,7 PM on the project.

Sonia Barizza is the Administrative Responsible charged for 0,7 PM on the project

Travels

Marcello Pelillo

Travel to Zurich Swiss (25-27/06/2009): SIMBAD project meeting;

Travel to Parma (Italy) (21-22/05/2009): meeting with prof. Ramin Zabih (Editor in chief IEEE-TPAMI)

Travel to Trento (Italy) (9-10/06/2009): meeting connected with SIMBAD research field

Travel to Milan - Italy (6-8/05/2009): participation in PRIS 2009

Travel to Verona - Italy (8/06/2009): participation in research project meeting

Participation in project review meeting Venice - Italy (24/06/2009):

Travel to Istambul Turkey (9-16/05/2009): participation in GamesNet 2009 Conference

Tavel to Bled (Slovenia) (6-9/09/2009): participation in ECML PKDD 2009 Conference

Travel to Verona - Italy (19-20/11/2009): participation in project meeting

Travel to Lisbon - Portugal (14-20/02/2010): participation in project research meeting

Andrea Torsello:

Travel to Bonn Germany (24-25/08/2009): Participation in EMMCVPR 2009 Conference

Travel to Zurich Swiss (25-27/06/2009): RTD meeting for joint activities within SIMBAD project

Travel to Kyoto -Japan (26/09/2009-7/10/2009):participation in ICCV 2009 conference

Travel to Verona - Italy (19-20/11/2009): participation in project meeting

Samuel Rota Bulò: travel to Berinoro - Italy for participation in Game Theory Summer School (31/05-06/06/2009),

Travel to Verona - Italy (19-20/11/2009): participation in project meeting

Travel to Vancouver and Whistler - Canada (6-14/12/2008): participation in NIPS 2009 conference and Workshop

Travel to Zurich - Switzerland (1-8/02/2010): participation in project research meeting

Travel to Lisbon - Portugal (14-20/02/2010): participation in project research meeting

Erdem Ibrahim Aykut

travel to Zurich - Switzerland (25-28/06/2009): participation in project meeting

Travel to Verona - Italy (19-20/11/2009): participation in project meeting

Jian Hou

Travel to Verona, ITALY (19-20/11/2009): Participation in SIMBAD project meeting

Travel to Zurich - Switzerland (1-6/02/2010): participation in project research meeting

Veronica Giove

Travel to Zurich Swiss (25-27/06/2009): SIMBAD project meeting;

Report on additional resources

Samuel Rota Bulò (4 PM) and Andrea Albarelli (2 PM) is post-doc not paid with SIMBAD funds, but involved in the research activities.

TABLE 3.1 PERSONNEL, SUBCONTRACTING AND OTHER MAJOR DIRECT COST ITEMS FOR BENEFICIARY 2, UNIVERSITY OF YORK, FOR THE PERIOD FROM 01/04/2008 TO 31/03/2009			
Work Package	Item description	Amount	Explanations
4.1 – 4.2	Personnel costs	€ 36.457	Salaries of one full professor, two PhD students
4.1 – 4.2	Travel	€ 6.741	
	Other costs	€85	
	TOTAL DIRECT COSTS	*43.283*	

Personnel

Professor Edwin Hancock, he is full professor leading York strand and WP4, he is charged for 0,75 PM paid with Simbad funds, and for 11,25 PM not paid with Simbad funds

Lin Han, Since 1/10/08 PhD student, 8,2 PM

Eliza Xu, Since 1/10/08 PhD student, 8,2 PM

Travels

Edwin Hancock and Richard Wilson participated in the Zurich project meeting (18-21.11.09).

Edwin Hancock and Richard Wilson participated in the Verona project meeting (25-27/06/2009).

Report on additional resources

The following people were involved in the SIMBAD project research, but have not been paid with SIMBAD funds: Dr Richard Wilson, Reader; Peng Ren, PhD student; Howaida El Ghawalby, PhD student.

TABLE 3.1 PERSONNEL, SUBCONTRACTING AND OTHER MAJOR DIRECT COST ITEMS FOR BENEFICIARY 3, TECHNISCHE UNIVERSITEIT DELFT, FOR THE PERIOD FROM 01/04/2009 TO 31/03/2010			
Work Package	Item description	Amount	Explanations
3,4,6,7,8	Personnel costs	*94710 €*	*Salaries for one researcher (11.8 PM) and one Associate Professor (4.2 PM) and one Assistant Professor (0.5 PM)*
3,4,6,7,8	Travel Costs	*9752 €*	*Iceland, Switzerland, Italy*
3,4,6,7,8	Computer parts	*643 €*	
3,4,6,7,8	Remaining direct costs	*392 €*	
	TOTAL DIRECT COSTS	***105497 €***	

Personnel:

Robert P.W. Duin, associate professor, heading the TUD contribution.
Marco Loog, assistant professor
Wan-Jui Lee, researcher totally paid with Simbad funds
Elzbieta Pekalska, research associate of Manchester University, not paid with Simbad funds
Alessandro Ibba, PhD student not paid with Simbad funds
Marcel Reinders, full professor

Travels:

Wan-Jui Lee, Robert Duin, Marco Loog attended the MCS2009
(Multiple Classifier Systems) in Reijkjavik, June 2009 where Wan-Jui Lee presented results from the project.
Robert Duin: SIMBAD review meeting, Venice, June 2009.
Wan-Jui Lee, Robert Duin: SIMBAD project meeting, Zürich, June 2009.
Wan-Jui Lee, research visit to Horst Bunke c.s., Bern, July 2009.
Wan-Jui Lee, Robert Duin, Marco Loog: SIMBAD project meeting, Verona, November 2009.
Wan-Jui Lee, Marco Loog, Alessandro Ibba: SIMBAD workshop, Zürich, February 2009.

TABLE 3.1 PERSONNEL, SUBCONTRACTING AND OTHER MAJOR DIRECT COST ITEMS FOR BENEFICIARY 4, INSTITUTO SUPERIOR TECNICO, FOR THE PERIOD FROM 01/04/2009 TO 31/03/2010			
Work Package	Item description	Amount	Explanations
2	Personnel costs	€ 57,570	*Salaries of 3 associate professors*
2	Travel	*€ 7,566*	*Travel Costs: Venice project review meeting, and project meetings in Zurich and Verona; visit to Venice(September 2009) for collaborative research activities; participation in ICAART 2010*
	TOTAL DIRECT COSTS	**€ 65 136**	

Personnel:

Mário Figueiredo: associate professor, heading the IST contribution (4,21 PM).

Ana Fred: assistant professor. (4,21 PM)

Pedro Aguiar: assistant professor (3,21 PM)

André Martins: not paid with SIMBAD funds, but involved in the research activities in WP2.2

David Pereira-Coutinho: not paid with SIMBAD funds, but involved in the research activities in WP2.2

Liliana Medina: collaborator, not paid with SIMBAD funds, but involved in the research activities in WP2.3

João Duarte: PDh student, not paid with SIMBAD funds, but involved in the research activities in WP2.3

André Lourenço: PDh student, not paid with SIMBAD funds, but involved in the research activities in WP2.3

Helena Aidos: PDh student, not paid with SIMBAD funds, but involved in the research activities in WP2.3 since February 2010.

Travel:

Ana Fred and Mário Figueiredo participated in the Venice project review meeting, and in the Zurich project meeting that followed, in June 2009.

Ana Fred and PhD student André Martins participated in the Verona project meeting, in November 2009

Ana Fred visited the Venice, for development of collaborative work on learning similarities from clusterings, group in September 2009.

TABLE 3.1 PERSONNEL, SUBCONTRACTING AND OTHER MAJOR DIRECT COST ITEMS FOR BENEFICIARY 5, UNIVERSITÀ DEGLI STUDI DI VERONA FOR THE PERIOD FROM 01/04/2009 TO 31/03/2010			
Work Package	Item description	Amount	Explanations
2, 7	Personnel costs	*69.158 €*	*Salaries of 4 postdoc for 24 PM, and 3 permanent staff for tot. 6 PM*
2,7	Equipment	*2.970 €*	*n.5 Personal computer depreciation for the 2nd period (total amount of the invoices € 10.357,87) to implement the specific research*
2,7	Travel	*12.956 €*	*Mission expenses*
	TOTAL DIRECT COSTS	***85.084 €***	

People involved in the project are:

Vittorio Murino Full Professor, Person responsible of the work (WP2.1: 0,2 PM; WP2.2: 0,8PM, WP7: 1PM)

Umberto Castellani Researcher (WP2.1: 0,2 PM; WP2.2: 0,8PM, WP6: 1PM)

Manuele Bicego Researcher (WP2.1: 0,6 PM; WP2.2: 2,4PM)

Cerruti Stefania PhD Student, Since 10th November 2008, totally paid with SIMBAD funds (WP7: 5PM)

Ulas Aydin post-doc, since 1st July 2009, totally paid with SIMBAD funds (WP2.1: 0,8 PM; WP2.2: 3,2PM, WP6: 2PM; WP7: 3PM)

Dong Seon Cheng post-doc, since 1st January 2009, totally paid with SIMBAD funds (WP2.1: 0,6 PM; WP2.2: 2,4PM)

Alessandro Perina post-doc, since 1st march 2010, totally paid with SIMBAD funds (WP2.1: 0,2 PM; WP2.2: 0,8PM)

Travels:

- Vittorio Murino, Zurich, 26-27 June 2009, Project Meeting
- Manuele Bicego, Cheng Dong Seon, Carli Anna, Castellani Umberto, Venezia 24 June 2009, Project Review Meeting
- Manuele Bicego, Cheng Dong Seon, Anna, Castellani Umberto, Zurich, 25-27 June 2009, Project Meeting
- Manuele Bicego, Bonn, 24-27 August 2009, participation to the Conference EMMCVPR 2009, presentation of the SIMBAD paper "Clustering-based construction of Hidden Markov Models for generative kernels" (M. Bicego, M. Cristani, V. Murino, E. Pekalska, R. Duin)
- Aydin Ulas, Venezia, 29 October 2009, research activities related to cooperation with UNIVE unit on WP7
- Manuele Bicego, Kyoto, 20 Sept. - 5 Oct. 2009, participation to the Conference ICCV 2009, presentation of the SIMBAD paper “Non-linear Generative embeddings for kernels on latent variable models” (A. Carli, M. Bicego, S. Baldo, V. Murino) to the ICCV2009 SUBSPACE workshop

- Umberto Castellani, London, 19-21 september 2009, participation to the workshop PMMIA 2009 and presentation of the SIMBAD paper “A Hybrid Generative/Discriminative Method for Classification of Regions of Interest in Schizophrenia Brain MRI” (D. S. Cheng, M. Bicego, U. Castellani, M. Cristani, S. Cerruti, M. Bellani, G. Rambaldelli, M. Aztori, P. Brambilla, V. Murino)
- Manuele Bicego, Lisboa, 26 Nov.-9 dec. 2009, research activities related to cooperation with IST unit on WP2
- Alessandro Perina, Vancouver and Whistler, 6-15 December 2009, participation to the Conference NIPS 2009 and presentation of the SIMBAD paper “Free energy score space” (A.Perina, M.Cristani, V.Murino, N.Jojic)
- Umberto Castellani, Reggio Emilia 10 December 2009, participation to the Conference AIIA 2009 and presentation of the SIMBAD paper “Local Kernels for Brains Classification in Schizophrenia” (Umberto Castellani, Elisa Rossato, Vittorio Murino, Marcella Bellani, Gianluca Rambaldelli, Michele Tansella, Paolo Brambilla)
- Aydin Ulas, Stefania Cerrutti. Marcella Bellani, Zurich, 01-03 February 2010, Project Meeting

TABLE 3.1 PERSONNEL, SUBCONTRACTING AND OTHER MAJOR DIRECT COST ITEMS FOR BENEFICIARY 6, EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH, FOR THE PERIOD FROM 01/04/2009 TO 31/03/2010			
Work Package	Item description	Amount	Explanations
2,3,4,6,7,8	Personnel costs	*95.594 €*	*Salaries of 2 doctoral students (5 months 100% each, 7 months Sharon Wulff 85%, Peter Schueffler 75%*
2,3,4,6,7,8	Other Direct Costs	*9.169€*	*Travel expenses (SIMBAD-Meeting Verona, Review Meeting Verona, Basel, Professor NIPS, Research associate ICML) and Organization: SIMBAD Meeting Zurich, SIMBAD Workshop Zurich*
2,3,4,6,7,8	Equipment Purchase	*5.795 €*	
	TOTAL DIRECT COSTS[1]	***110.558***	

People involved in the project are:

Joachim Buhmann	Full Professor, PM
Cheng Soon Ong	Researcher
Volker Roth	Associate Professor, University of Basel
Sharon Wulff	PhD-student since 15 November 2008, 100% paid from 15 Nov 2008 to Aug 2009 85% from Sep 2009 to Mar 2010
Peter Schueffler	PhD-student since 1 December 2008 100% paid from 1 Dec 2008 to Aug 2009 75% paid from Sep 2009 to Mar 2010

Travels:

SIMBAD

Joachim Buhmann, Venice, 23-24 June 2009, SIMBAD Review Meeting

Joachim Buhmann, Verona, 18-21 Nov 2009, SIMBAD Meeting

Joachim Buhmann, Lisboa, 5-8 May 2010, SIMBAD Review Meeting (Hotel)

Cheng Soon Ong, Verona, 18-21 Nov 2009, SIMBAD Meeting

Volker Roth, Verona, 18-21 Nov 2009, SIMBAD Meeting

[1] Total direct costs have to be coherent with the directs costs claimed in Form C

Sharon Wulff, Verona, 18-21 Nov 2009, SIMBAD Meeting

Peter Schueffler, Verona, 18-21 Nov 2009, SIMBAD Meeting

Sharon Wulff and Peter Schueffler, several trips to Basel to meet Volker Roth

CONFERENCES

Cheng Soon Ong, Montreal, participation in ICML/UAI/COLT-conferences, 15-23 June 2009

Joachim Buhmann, Vancouver / Whistler, participation in NIPS conference, 5-14 December 2009

Organization:

SIMBAD Meeting, Zurich, 26-27 June 2009

SIMBAD Workshop, Zurich, 1-5 February 2010

Project acronym	SIMBAD
Project full title	Beyond Features: Similarity-Based Pattern Analysis and Recognition
Name of the document	Annex A – SIMBAD List of publications II Year
Dissemination level	Public
Date of last update	20/06/2010

1. Weiping Xu, Edwin R. Hancock, Richard C. Wilson “Rectifying Non-Euclidean Similarity Data using Ricci Flow Embedding”, (ICPR 2010), Istanbul, Turkey, 2010.
2. Weiping Xu, Edwin R. Hancock, Richard C. Wilson, “Regularising the Ricci Flow Embedding” Structural and Syntactic Pattern Recognition (SSPR 2010).
3. A. Albarelli, E. Rodolà, and A. Torsello, “A Game-Theoretic Approach to Fine Surface Registration without Initial Motion Estimation.” In IEEE International Conference on Computer Vision And Pattern Recognigion (CVPR 2010), IEEE Computer Society, 2010.
4. A. Albarelli, E. Rodolà, and A. Torsello, “Robust Game-Theoretic Inlier Selection for Bundle Adjustment.” In 3D Data Processing, Visualization and Transmission (3DPVT), 2010.
5. M. Bicego, A. Martins, V. Murino, P. Aguiar, and M. Figueiredo, “2D shape recognition using information theoretic kernels”, IAPR International Conference on Pattern Recognition (ICPR 2010), Istanbul, Turkey, 2010.
6. A. Martins, M. Bicego, V. Murino, P. Aguiar, and M. Figueiredo, "Information Theoretical Kernels for Generative Embeddings Based on Hidden Markov Models", Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition and Statistical Techniques in Pattern Recognition (S+SSPR 2010), Izmir, Turkey, 2010.
7. D. Pereira-Coutinho, A. Fred, and M. Figueiredo, “One-lead ECG-based personal identification using Ziv-Merhav cross parsing”, IAPR International Conference on Pattern Recognition (ICPR 2010), Istanbul, Turkey, 2010.
8. H. El-Ghawalby, E. R. Hancock “Geometric Characterizations of Graphs Using Heat Kernel Embeddings” – (IMA Conference on the Mathematics of Surfaces), York, United Kingdom, 2010.

9. P. Ren, R. C. Wilson, and E. R. Hancock "Characteristic Polynomial Analysis on Matrix Representations of Graphs" (S+SSPR 2010), Izmir, Turkey, 2010.
10. S. Xia, E. R. Hancock , "Graph-Based Object Class Discovery" (CAIP 2009) Münster, Germany.
11. P. Ren, T. Aleksić, R. C. Wilson, E. R. Hancock "Hypergraphs, Characteristic Polynomials and the Ihara Zeta Function" (CAIP 2009) Münster, Germany.
12. L. Han, R. C. Wilson , E. R. Hancock, "A Supergraph-based Generative Model "-(ICPR 2010), Istanbul, Turkey, 2010.
13. P. Ren ,R. C. Wilson, E. R. Hancock "Weighted Graph Characteristics from Oriented Line Graph Polynomials" (ICCV 2009) Kyoto, Japan.
14. A. Carli, M. Bicego, S. Baldo, V. Murino "Nonlinear mappings for generative kernels on latent variable models" - (ICPR 2010), Istanbul, Turkey, 2010 .
15. A. Carli, U. Castellani, M. Bicego, V. Murino "Dissimilarity-based Representation for Local parts", Int. Workshop on Cognitive Information Processing (CIP 2010), 14-15-16 June, 2010 Elba Island, Tuscany - Italy, 2010.
16. R. C. Wilson, E.R. Hancock, E. Pekalska and R. Duin, "Spherical Embeddings for non-Euclidean Dissimilarities", (CVPR 2010) San Francisco, California.
17. W.J. Lee, R.P.W. Duin, and H. Bunke, "Selecting Structural Base Classifiers for Graph-based Multiple Classifier Systems", in: N. El Gayar, J. Kittler, F. Roli (eds.), Multiple Classifier Systems (Proc. 9th Int. Workshop, MCS 2010, Cairo, Egypt), Lecture Notes in Computer Science, vol. 5997, Springer, Berlin, 2010, 155-164.
18. W.J. Lee, R.P.W. Duin, M. Loog, and A. Ibba, "An Experimental Study On Combining Euclidean Distances", Proc. The 2nd International Workshop on Cognitive Information Processing (CIP 2010), 14-15-16 June, 2010 Elba Island, Tuscany - Italy, 2010.
19. A. Ibba, R.P.W. Duin, and W.J. Lee, "A study on combining sets of differently measured dissimilarities", (ICPR 2010), Istanbul, Turkey, 2010.
20. R.P.W. Duin and E. Pekalska, "Non-Euclidean Dissimilarities: Causes and Informativeness", (SSSPR, 2010) Izmir, Turkey, 2010.
21. R.P.W. Duin, "Pattern Recognition as a Human Centered non-Euclidean Problem", (invited) (ICEIS 2010), Funchal, Madeira, Portugal, June 2010.
22. U. Castellani, E. Rossato, V. Murino, M. Bellani, G. Rambaldelli, M. Tansella, P. Brambilla, "Local Kernels for Brains Classification in Schizophrenia". In R. Serra and R. Cucchiara (Eds.): AI*IA 2009: Emergent Perspectives in Artificial Intelligence, LNAI 5883, pp. 112-121 Springer Berlin/ Heidelberg 2009.
23. M. Bicego, P. Lovato, A. Ferrarini, M. Delledonne "Biclustering of expression microarray data with topic models", ICPR 2010 - Istanbul, Turkey, 2010.
24. Joachim M. Buhmann, "Information theoretic model validation for clustering", ISIT 2010.
25. Samuel Rota Bulò, André Louren, Ana Fred, Marcello Pelillo "Pairwise Probabilistic Clustering Using Evidence Accumulation" Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition and Statistical Techniques in Pattern Recognition (S+SSPR 2010), Izmir, Turkey, 2010.

26. João Duarte, Ana Fred, André Lourenço, and Fernando Duarte "On consensus clustering validation". Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition and Statistical Techniques in Pattern Recognition (S+SSPR 2010), Izmir, Turkey, 2010.
27. André Lourenço, Ana Fred and Anil K. Jain. "On the scalability of evidence accumulation clustering". IAPR International Conference on Pattern Recognition (ICPR 2010), Istanbul, Turkey, 2010 .
28. A. Albarelli, E. Rodolà, A. Cavallarin, and A. Torsello, "Robust Figure Extraction on Textured Background: a Game-Theoretic Approach." In 20th International Conference on Pattern Recognition. (ICPR 2010), Istanbul, Turkey, 2010.
29. E. Rodolà, A. Albarelli, and A. Torsello, "Robust Figure Extraction on Textured Background:." In 20th International Conference on Pattern Recognition. (ICPR2010) , Istanbul, Turkey, 2010.
30. E. Rodolà, A. Albarelli, and A. Torsello, "A Game-Theoretic Approach to the Enforcement of Global Consistency in Multi-View Feature Matching." In Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition (SSPR 2010) and Statistical Techniques in Pattern Recognition (SPR 2010), Izmir, Turkey, 2010.
31. S. Ceolin and E.R. Hancock, "Using the Fisher-Rao Metric to Compute Facial Similarity", ICIAR 2010, LNCS 2010.
32. S. Ceolin and E.R. Hancock, "Characterising facial gender difference using Fisher-Rao metric", (ICPR 2010), Istanbul, Turkey, 2010.
33. F. Escolano, M.A. Lozano and E.R. Hancock, "Heat Flow-Thermodynamic Depth Complexity in Networks ", (ICPR 2010), Istanbul, Turkey, 2010.
34. P. Ren, T. Aleksic, R.C. Wilson and E.R. Hancock "Ihara Coefficients: A Flexible Tool for Higher Order Learning", (S+SSPR 2010), Izmir, Turkey, 2010.
35. H.El Ghawalby and E.R. Hancock, "Graph Embedding using an Edge-based Wave Kernel, (S+SSPR 2010), Izmir, Turkey, 2010.
36. R.C. Wilson and E.R. Hancock, "Spherical Embedding and Classification", (S+SSPR 2010) Izmir, Turkey, 2010.
37. S. Xia and E.R. Hancock, "Mining Exemplars for Object Recognition using Affinity Propagation", (S+SSPR 2010) Izmir, Turkey, 2010.
38. N. Rahman and E.R, Hancock, "Commute Time Convolution Kernels for Graph Clustering", (S+SSPR 2010) Izmir, Turkey, 2010.
39. F. Escolano, E.R. Hancock, D. Giorgi and M.A. Lozano, "What is the Complexity of a Network? The Heat Flow-Thermodynamic Depth Approach", (S+SSPR 2010) Izmir, Turkey, 2010.
40. A. Ulas, R.P.W. Duin, U. Castellani, M. Loog, M. Bicego, V. Murino, M. Bellani, S. Cerruti, M. Tansella, P. Brambilla, "Dissimilarity-based Detection of Schizophrenia" (ICPR2010)
41. A. Carli, M. Bicego, S. Baldo, V. Murino, "Non-linear Generative embeddings for kernels on latent variable models" ICCV2009 SUBSPACE workshop
42. D. S. Cheng, M. Bicego, U. Castellani, M. Cristani, S. Cerruti, M. Bellani, G. Rambaldelli, M. Aztori, P. Brambilla, V. Murino, "A Hybrid Generative/Discriminative Method for Classification of Regions of Interest in Schizophrenia Brain MRI" MICCAI2009 - PMMIA 2009 Workshop
43. A.Perina, M.Cristani, V.Murino, N.Jojic "Free energy score space" NIPS 2009

44. A. Erdem and A. Torsello, “A Game Theoretic Approach To Jointly Learn Shape Categories and Contextual Similarities” In Joint IAPR International Workshops on Structural and Syntactic Pattern Recognition (SSPR 2010).
45. A. Torsello and M. Pelillo “Hierarchical Pairwise Segmentation using Dominant Sets and Anisotropic Diffusion” EMMCVPR 2009.
46. S. Rota Bulò, A. Albarelli, A. Torsello and M. Pelillo “A Hypergraph-based Approach to Affine Parameters Estimation” ICPR 2008.
47. S. Rota Bulo` and M. Pelillo “A Game-Theoretic Approach to Hypergraph Clustering” NIPS 2009.
48. M. Pelillo ”What is a Cluster? Perspectives from Game Theory” NIPS 2009 workshop on Cllustering
49. A. Perina, M. Cristani, U. Castellani, V. Murino, N. Jojic “Free energy score-space” NIPS 2009.
50. M. Bicego, M. Cristani, V. Murino, E. P,ekalska, R.P.W. Duin, “Clustering-based construction of Hidden Markov Models for Generative Kernels” EMMCVPR09
51. A.Perina1 M.Cristani, U.Castellani, V.Murino, N.Jojic, A hybrid generative/discriminative classification framework based on free-energy terms”, ICCV09.
52.
53. M. Bicego, M. Cristani, V. Murino, E. Pekalska, R. Duin, "Clustering-based construction of Hidden Markov Models for generative kernels", *Proc. of Int. Conf. on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR2009),* pp. 466-479, Bonn, Germany, 2009Clustering-based construction of Hidden Markov Models for generative kernels
54. D. Pereira Coutinho, A. Fred, M. Figueiredo, “'Personal identification and authentication based on one-lead ECG by using Ziv-Merhav cross parsing”, 10th International Workshop on Pattern Recognition in Information Systems (PRIS), Funchal, Portugal, 2010 (accepted)
55. D. Pereira Coutinho, A. Fred, M. Figueiredo, “One-lead ECG-based personal identification using Ziv-Merhav cross parsing”, Proceedings of the International Conference on Pattern Recognition (ICPR), Istanbul, Turkey, 2010 (accepted).
56. L. Medina, A. Fred. “Clustering temporal data: application to electrophysiological signals”, Accepted for publication In: *Communications in Computer and Information Science*. Springer 2010.
57. F. J. Duarte, J. Duarte, A. Fred. “Average cluster consistency for cluster ensemble selection”, Accepted for publication In: *Communications in Computer and Information Science*. Springer 2010.
58. W.J. Lee and R.P.W. Duin, “A Labelled Graph Based Multiple Classifier System”, in: J.A. Benediktsson, J. Kittler, F. Roli (eds.), *Multiple Classifier Systems* (Proc. 8th Int. Workshop, MCS 2009, Reykjavik, Iceland, June 10- 12, 2009), Lecture Notes in Computer Science, vol. 5519, Springer, Berlin, 2009, 201-210
59. S.W. Kim and R.P.W. Duin, “A Combine-Correct-Combine Scheme for Optimizing Dissimilarity-Based Classifiers”, in: E. Bayro-Corrochano, J.O. Eklundh (eds.), *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Proc. CIARP 2009, Guadalajara, Jalisco, Mexico, Nov. 15-18, 2009), Lecture Notes in Computer Science, vol. 5856, Springer, Berlin, 2009, 425-432.
60. Y. Plasencia, E. Garcıa-Reyes, R. P. W. Duin, H. Mendez-Vazquez, C. San- Martin, and C. Soto, “A Study on Representations for Face Recognition from Thermal Images”, in: E. Bayro-Corrochano,

J.O. Eklundh (eds.), *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Proc. CIARP 2009, Guadalajara, Jalisco, Mexico, Nov. 15-18, 2009), Lecture Notes in Computer Science, vol. 5856, Springer, Berlin, 2009, 185-192.

61. D. Porro Munoz, R.P.W. Duin, I. Talavera, and N. Hernandez, The Representation of Chemical Spectral Data for Classification, in: E. Bayro- Corrochano, J.O. Eklundh (eds.), *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications* (Proc. CIARP 2009, Guadalajara, Jalisco, Mexico, Nov. 15-18, 2009), Lecture Notes in Computer Science, vol. 5856, Springer, Berlin, 2009, 513-520.
62. S. Klein, M. Loog, F. van der Lijn, T. den Heijer, A. Hammers, M. de Bruijne, A. van der Lugt, R.P.W. Duin, M.M.B. Breteler, and W.J. Niessen, Early diagnosis of dementia based on intersubject whole-brain dissimilarities, *Proc. of IEEE International Symposium on Biomedical Imaging: Macro to Nano* (Rotterdam, 14-17 April 2010), 2010.
63. Millan-Giraldo, R.P.W. Duin, and J.S. Sanchez, Dissimilarity-based Classification of Data with Missing Attributes, *Proc. The 2nd International Workshop on Cognitive Information Processing* (14-15-16 June, 2010 Elba Island, Tuscany - Italy), 2010. accepted
64. L. Sørensen, R.P.W. Duin, M. de Bruijne, W.J. Lee, D.M.J. Tax, and M. Loog, Dissimilarity-based Multiple Instance Learning, *Proc. SSSPR, 2010*. accepted
65. D. Porro Munoz, R.P.W. Duin, M. Orozco-Alzate, I. Talavera-Bustamante, and J.M. Londo˜no-Bonilla, The Dissimilarity Representation as a Tool for Three-way Data Classification: a 2D Measure, *Proc. SSSPR, 2010*. accepted
66. S.W. Kim and R.P.W. Duin, An emperical comparison of kernel-based and dissimilarity-based feature spaces, *Proc. SSSPR, 2010*, accepted.
67. D. Porro-Munoz, R.P.W. Duin, M. Orozco-Alzate, I. Talavera, and J.M. Londo˜no- Bonilla, Classifying Three-way Seismic Volcanic Data by Dissimilarity Representation, *ICPR 2010*, 2010. accepted
68. Y. Plasencia-Calana, E. Garcıa-Reyeso, M. Orozco-Alzate, and R.P.W. Duin, Prototype selection for dissimilarity representation by a genetic algorithm, *ICPR 2010*, 2010. accepted
69. L. Sørensen, M. de Bruijne, R.P.W. Duin, M. Loog, P. Lo, and A. Dirksen, Image Dissimilarity-Based Quantification of Pathology, *MICCAI, 2010*. accepted
70. I. Patras and E.R. Hancock. Coupled prediction-classification for robust visual tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence, to appear, 2010.
71. DongJun Yu, W.A.P. Smith and E.R. Hancock. Learning a Self-Organising Map Model on a Riemannian Manifold. Mathematics of Surface, LNCS, 5654, pp. 375--390, 2009.
72. Shenping Xia and E.R. Hancock. Learning Class Specific Hypergraphs. ICIAP 2009, LNCS, 5716, pp. 269--277, 2009.
73. DongJun Yu, W.A.P. Smith and E.R. Hancock. A Riemammian Self-Organizing Map. ICIAP 2009, LNCS, 5716, pp. 229--238, 2009.
74. A. Albarelli, E. Rodolà, and A. Torsello, “Loosely Distinctive Features for Robust Surface Alignment.” In ECCV2010 (accepted).
75. S. Rota Bulo and M Pelillo, “Probabilistic Clustering using the Baum-Eagon Inequality.” In ICPR, 2010.

76. A. Albarelli, E. Rodolà, and A. Torsello, “A Game-Theoretic Approach to Fine Surface Registration without Initial Motion Estimation.” In CVPR, 2010.

77. A. Torsello and M. Pelillo, “Hierarchical Pairwise Segmentation Using Dominant Sets and Anisotropic Diffusion Kernels.” In EMMCVPR 2009.

78. T. J. Fuchs and J. M. Buhmann. Inter-Active Learning of Randomized Tree Ensembles for Object Detection. Workshop on On-line Learning for Computer Vision ICCV 2009.

79. Cheng, D. S., Bicego, M., Castellani, U., Cerruti, S., Bellani, M., Rambaldelli, G., Atzori, M., Brambilla, P., Murino, V., 2009. Schizophrenia classification using regions of interest in brain mri. In: Proceedings of Intelligent Data Analysis in Biomedicine and Pharmacology, IDAMAP ’09. pp. 47–52.

80. Castellani, U., Rossato, E., Murino, V., Bellani, M., Rambaldelli, G., Tansella, M., Brambilla, P., 2009. Local kernel for brains classification in schizophrenia. In: AI*IA ’09:: Proceedings of the XIth International Conference of the Italian Association for Artificial Intelligence Reggio Emilia on Emergent Perspectives in Artificial Intelligence. Springer-Verlag, Berlin, Heidelberg, pp. 112–121.