Activity Report 2012

Team GALEN

Organ Modeling through Extraction, Representation and Understanding of Medical Image Content

RESEARCH CENTER
Saclay - Île-de-France

THEME
Computational Medicine and Neuro-sciences

Table of contents

Team GALEN

Keywords: Computer Vision, Image Processing, Medical Images, Discrete Optimization, Machine Learning

Creation of the Team: February 15, 2008 .

1. Members

Faculty Members

Andreas Argyriou [R-BUCE Fellow, Ecole Centrale de Paris]
Matthew Blaschko [Associate Professor, Ecole Centrale de Paris]
Iasonas Kokkinos [Associate Professor, Ecole Centrale de Paris]
Pawan Kumar [Associate Professor, Ecole Centrale de Paris]
Nikos Paragios [Team Leader, Professor, Ecole Centrale de Paris, HdR]

External Collaborators

Jean-Francois Deux [Associate Professor of Diagnostic Radiology - Paris 12 School of Medicine]
Nikos Komodakis [Associate Professor, Ecole des Ponts-ParisTech]
Alain Rahmouni [Professor of Diagnostic Radiology - Paris 12 School of Medicine]
Dimitris Samaras [Associate Professor, Stony Brook University - USA / Digiteo Subsample Chair]

PhD Students

Stavros Alchatzidis [DIOCLES ERC-Grant]
Pierre-Yves Baudin [SIEMENS CIFRE]
Haithem Boussaid [Centrale-Paris Grant]
Xiang Bo [MENESR Grant]
Enzo Ferrante [MENESR Grant]
Vivien Fecamp [Inria-DGA]
Katerina Gkirtzou [Centrale-Paris Grant]
Nicolas Honnorat [GE-HealthCare CIFRE]
Puneet Kumar [MENESR Grant]
Helene Langet [GE-HealthCare CIFRE]
Fabrice Michel [MEDICEN Stereos+ Grant]
Sarah Parisot [Intrasense CIFRE]
Stavros Tsogkas [ANR HiCore Grant]

Administrative Assistants

Christine Biard
Carine Morroti-Delorme [Ecole Centrale de Paris]

2. Overall Objectives

2.1. GALEN@Centrale-Paris

Computational vision is one of the most challenging research domains in engineering sciences. The aim is to reproduce human visual perception through intelligent processing of visual data. The application domains span from computer aided diagnosis to industrial automation & robotics. The most common mathematical formulation to address such a challenge is through mathematical modeling. In such a context, first the solution of the desired vision task is expressed in the form of a parameterized mathematical model. Given such a model, the next task consists of associating the model parameters with the available observations, which is often called the model-to-data association. The aim of this task is to determine the impact of a parameter choice to the observations and eventually maximize/minimize the adequacy of these parameters with the visual observations. In simple words, the better the solution is, the better it will be able to express and fit the data. This is often achieved through the definition of an objective function on the parametric space of the model. Last, but not least given the definition of the objective function, visual perception is addressed through its optimization with respect to the model parameters. To summarize, computation visual perception involves three aspects, a task-specific definition of a parametric model, a data-specific association of this model with the available observations and last the optimization of the model parameters given the objective and the observations.

Such a chain processing inherits important shortcomings. The curse of dimensionality is often used to express the importance of the model complexity. In simple words, the higher the complexity of the model is, the better its expressive power will be with counter effect the increase of the difficulty of the inference process. Non-linearity is another issue to be addressed which simply states that the association between the model and the data is a (highly) non-linear function and therefore direct inference is almost infeasible. The impact of this aspect is enforced from the curse of non-convexity that characterizes the objective function. Often it lives in high-dimensional spaces and is ill posed making exact inference problematic (in many cases not possible) and computationally expensive. Last, but not least modularity and scalability is another important concern to be addressed in the context of computational vision. The use of task-specific modeling and algorithmic solutions make their portability infeasible and therefore transfer of knowledge from one task to another is not straightforward while the methods do not always scale well with respect either to the dimensionality of the representation or the data.

GALEN aims at proposing innovative techniques towards automatic structuring, interpretation and longitudinal modeling of visual data. In order to address these fundamental problems of computational perception, GALEN investigates the use of discrete models of varying complexity. These methods exhibit an important number of strengths such as their ability to be modular with respect to the input measurements (clinical data), the nature of the model (certain constraints are imposed from computational perspective in terms of the level of interactions), and the model-to-data association while being computational efficient.

2.2. Highlights of the Year

- **BIOMED Summer School**: Galen has organized the Biomedical Image Analysis Summer School : Modalities, Methodologies & Clinical Research at Paris between July 9th and July 14th, 2012 involving international leaders/contributors in the field of biomedical image analysis as instructors where approx 100 participants were selected from an outstanding number of applications.
- **China Research Council Award**: Chaohui Wang was the recipient of the Chinese Government Award for Outstanding (self-financed) PhD. In 2012, a total of 495 awards were given worldwide in all disciplines, with 17 Chinese students in France receiving awards.
- **CVPR Participation**: GALEN has participated in the 2012 annual IEEE Conference in Computer Vision and Pattern Recognition (CVPR'12) conference, the leading event in the field of computer vision with five papers (double blind full submissions, acceptance rate %25).
- **EU FP7 Success**: GALEN has secured cutting edge research funding from the European Union through the highly competitive 2012 "Cognitive Vision and Robotics" FP7-ICT-9 call (5% acceptance) through two accepted grants (out of 12 for the entire call): MOBOT (Intelligent Active MObility Assistance RoBOT integrating Multimodal Sensory Processing, Proactive Autonomy and Adaptive Interaction) and RECONFIG (Cognitive, Decentralized Coordination of Heterogeneous Multi-Robot Systems).
- **MICCAI Participation**: GALEN has participated in the 2012 annual Medical Image Computing and Computer Assisted Intervention (MICCAI'12) conference one of the leading events in the field of medical image analysis with four papers (double blind full submissions, acceptance rate %30) and two invited talks in the associated workshops.

3. Scientific Foundations

3.1. Structured coupled low- and high-level visual perception

A general framework for the fundamental problems of image segmentation, object recognition and scene analysis is the interpretation of an image in terms of a set of symbols and relations among them. Abstractly stated, image interpretation amounts to mapping an observed image, X to a set of symbols Y. Of particular interest are the symbols Y^* that *optimally explain the underlying image*, as measured by a scoring function s that aims at distinguishing correct (consistent with human labellings) from incorrect interpretations:

$$Y^* = \mathrm{argmax}_Y s(X, Y) \tag{1}$$

Applying this framework requires (a) identifying which symbols and relations to use (b) learning a scoring function s from training data and (c) optimizing over Y in Eq. 1.

A driving force behind research in GALEN has been the understanding that these three aspects are tightly coupled. In particular, efficient optimization can be achieved by resorting to sparse image representations that 'shortlist' putative solutions and/or by working with scoring functions that can be efficiently optimized. However, the accuracy of a scoring function is largely affected by the breadth of relationships that it accommodates, as well as the completeness of the employed image representation. Determining the tradeoff between these two requirements is far from obvious and often requires approaches customized to the particular problem setting addressed. Summarizing, even though the three problems outlined above can be addressed in isolation, an integrated end-to-end approach is clearly preferable, both for computational efficiency and performance considerations.

Research in GALEN has therefore dealt with the following problem aspects: first, developing a generic and reliable low-level image representation that can be used transversally across multiple tasks. The use of learning-based techniques has been pursued for boundary detection and symmetry detection in [32], yielding state-of-the-art results, while in [27] trajectory grouping was used to come up with a mid-level representation of spatio-temporal data. Complementary to the detection of geometric structures, we have also explored methods for their description both for image and surface data [17]. We are currently pursuing the formulation of the task in structured prediction terms, which will hopefully allow us to exploit the geometrical interdependencies among symmetry and boundary responses.

Second, we have worked on learning scoring functions for detection with deformable models that can leverage upon the developed low-level representations, while also being amenable to efficient optimization. Building on our earlier work on using boundary and symmetry detector responses to perform groupwise registration within categories we used discriminative learning to train hierarchical object models that rely on shape-based representations; these were successfully applied to the detection of shape-based categories, while we are currently pursuing their integration with appearance-based models.

Third, efficient optimization for deformable models was pursued in [18], where we have developed novel techniques for object detection that employ combinatorial optimization tools (A^* and Branch-and-Bound) to tame the combinatorial complexity; in particular our work has a best-case performance that is logarithmic in the number of pixels, while our work in [18] allows us to further accelerate object detection by integrating low-level processing (convolutions) with a bounding-based object detection algorithm. Working on a different approach, in [10] we have pursued the exploitation of reinforcement-learning to optimize over the set of shapes derivable from shape grammars. We are currently pursuing a full-fledged bounding-based inference algorithm, which will integrate the tasks of boundary detection and grouping in a single, integrated object detection algorithm.

3.2. Machine Learning & Structure Prediction

The foundation of statistical inference is to learn a function that minimizes the expected loss of a prediction with respect to some unknown distribution

$$\mathcal{R}(f) = \int \ell(f, x, y) dP(x, y), \tag{2}$$

where $\ell(f, x, y)$ is a problem specific loss function that encodes a penalty for predicting $f(x)$ when the correct prediction is y. In our case, we consider x to be a medical image, and y to be some prediction, e.g. the segmentation of a tumor, or a kinematic model of the skeleton. The loss function, ℓ, is informed by the costs associated with making a specific misprediction. As a concrete example, if the true spatial extent of a tumor is encoded in y, $f(x)$ may make mistakes in classifying healthy tissue as a tumor, and mistakes in classifying diseased tissue as healthy. The loss function should encode the potential physiological damage resulting from erroneously targeting healthy tissue for irradiation, as well as the risk from missing a portion of the tumor.

A key problem is that the distribution P is unknown, and any algorithm that is to estimate f from labeled training examples must additionally make an implicit estimate of P. A central technology of empirical inference is to approximate $\mathcal{R}(f)$ with the empirical risk,

$$\mathcal{R}(f) \approx \widehat{\mathcal{R}}(f) = \frac{1}{n} \sum_{i=1}^{n} \ell(f, x_i, y_i), \tag{3}$$

which makes an implicit assumption that the training samples (x_i, y_i) are drawn i.i.d. from P. Direct minimization of $\widehat{\mathcal{R}}(f)$ leads to overfitting when the function class $f \in \mathcal{F}$ is too rich, and regularization is required:

$$\min_{f \in \mathcal{F}} \lambda \Omega(\|f\|) + \widehat{\mathcal{R}}(f), \tag{4}$$

where Ω is a monotonically increasing function that penalizes complex functions.

Equation (4) is very well studied in classical statistics for the case that the output, $y \in \mathcal{Y}$, is a binary or scalar prediction, but this is not the case in most medical imaging prediction tasks of interest. Instead, complex interdependencies in the output space leads to difficulties in modeling inference as a binary prediction problem. One may attempt to model e.g. tumor segmentation as a series of binary predictions at each voxel in a medical image, but this violates the i.i.d. sampling assumption implicit in Equation (3). Furthermore, we typically gain performance by appropriately modeling the inter-relationships between voxel predictions, e.g. by incorporating pairwise and higher order potentials that encode prior knowledge about the problem domain. It is in this context that we develop statistical methods appropriate to structured prediction in the medical imaging setting.

3.3. Self-Paced Learning with Missing Information

Many tasks in artificial intelligence are solved by building a model whose parameters encode the prior domain knowledge and the likelihood of the observed data. In order to use such models in practice, we need to estimate its parameters automatically using training data. The most prevalent paradigm of parameter estimation is supervised learning, which requires the collection of the inputs x_i and the desired outputs y_i. However, such an approach has two main disadvantages. First, obtaining the ground-truth annotation of high-level applications, such as a tight bounding box around all the objects present in an image, is often expensive. This prohibits the use of a large training dataset, which is essential for learning the existing complex models. Second, in many applications, particularly in the field of medical image analysis, obtaining the ground-truth annotation may not be feasible. For example, even the experts may disagree on the correct segmentation of a microscopical image due to the similarities between the appearance of the foreground and background.

In order to address the deficiencies of supervised learning, researchers have started to focus on the problem of parameter estimation with data that contains hidden variables. The hidden variables model the missing information in the annotations. Obtaining such data is practically more feasible: image-level labels ('contains car','does not contain person') instead of tight bounding boxes; partial segmentation of medical images. Formally, the parameters **w** of the model are learned by minimizing the following objective:

$$\min_{\mathbf{w}\in\mathcal{W}} R(\mathbf{w}) + \sum_{i=1}^{n} \Delta(y_i, y_i(\mathbf{w}), h_i(\mathbf{w})). \tag{5}$$

Here, $\mathcal{W}$ represents the space of all parameters, n is the number of training samples, $R(\cdot)$ is a regularization function, and $\Delta(\cdot)$ is a measure of the difference between the ground-truth output y_i and the predicted output and hidden variable pair $(y_i(\mathbf{w}), h_i(\mathbf{w}))$.

Previous attempts at minimizing the above objective function treat all the training samples equally. This is in stark contrast to how a child learns: first focus on easy samples ('learn to add two natural numbers') before moving on to more complex samples ('learn to add two complex numbers'). In our work, we capture this intuition using a novel, iterative algorithm called self-paced learning (SPL). At an iteration t, SPL minimizes the following objective function:

$$\min_{\mathbf{w}\in\mathcal{W},\mathbf{v}\in\{0,1\}^n} R(\mathbf{w}) + \sum_{i=1}^{n} v_i \Delta(y_i, y_i(\mathbf{w}), h_i(\mathbf{w})) - \mu_t \sum_{i=1}^{n} v_i. \tag{6}$$

Here, samples with $v_i = 0$ are discarded during the iteration t, since the corresponding loss is multiplied by 0. The term μ_t is a threshold that governs how many samples are discarded. It is annealed at each iteration, allowing the learner to estimate the parameters using more and more samples, until all samples are used. Our results already demonstrate that SPL estimates accurate parameters for various applications such as image classification, discriminative motif finding, handwritten digit recognition and semantic segmentation. We will investigate the use of SPL to estimate the parameters of the models of medical imaging applications, such as segmentation and registration, that are being developed in the GALEN team. The ability to handle missing information is extremely important in this domain due to the similarities between foreground and background appearances (which results in ambiguities in annotations). We will also develop methods that are capable of minimizing more general loss functions that depend on the (unknown) value of the hidden variables, that is,

$$\min_{\mathbf{w}\in\mathcal{W},\theta\in\Theta} R(\mathbf{w}) + \sum_{i=1}^{n} \sum_{h_i\in\mathcal{H}} \Pr(h_i|x_i, y_i; \theta) \Delta(y_i, h_i, y_i(\mathbf{w}), h_i(\mathbf{w})). \tag{7}$$

Here, θ is the parameter vector of the distribution of the hidden variables h_i given the input x_i and output y_i, and needs to be estimated together with the model parameters **w**. The use of a more general loss function will allow us to better exploit the freely available data with missing information. For example, consider the case where y_i is a binary indicator for the presence of a type of cell in a microscopical image, and h_i is a tight bounding box around the cell. While the loss function $\Delta(y_i, y_i(\mathbf{w}), h_i(\mathbf{w}))$ can be used to learn to classify an image as containing a particular cell or not, the more general loss function $\Delta(y_i, h_i, y_i(\mathbf{w}), h_i(\mathbf{w}))$ can be used to learn to detect the cell as well (since h_i models its location).

3.4. Discrete Biomedical Image Perception

A wide variety of tasks in medical image analysis can be formulated as discrete labeling problems. In very simple terms, a discrete optimization problem can be stated as follows: we are given a discrete set of variables $\mathcal{V}$, all of which are vertices in a graph $\mathcal{G}$. The edges of this graph (denoted by $\mathcal{E}$) encode the variables' relationships. We are also given as input a discrete set of labels $\mathcal{L}$. We must then assign one label from $\mathcal{L}$ to each variable in $\mathcal{V}$. However, each time we choose to assign a label, say, x_{p_1} to a variable p_1, we are forced to pay a price according to the so-called *singleton* potential function $g_p(x_p)$, while each time we choose to assign a pair of labels, say, x_{p_1} and x_{p_2} to two interrelated variables p_1 and p_2 (two nodes that are connected by an edge in the graph $\mathcal{G}$), we are also forced to pay another price, which is now determined by the so called *pairwise* potential function $f_{p_1 p_2}(x_{p_1}, x_{p_2})$. Both the singleton and pairwise potential functions are problem specific and are thus assumed to be provided as input.

Our goal is then to choose a labeling which will allow us to pay the smallest total price. In other words, based on what we have mentioned above, we want to choose a labeling that minimizes the sum of all the MRF potentials, or equivalently the MRF energy. This amounts to solving the following optimization problem:

$$\arg\min_{\{x_p\}} \mathcal{P}(g, f) = \sum_{p \in \mathcal{V}} g_p(x_p) + \sum_{(p_1, p_2) \in \mathcal{E}} f_{p_1 p_2}(x_{p_1}, x_{p_2}). \tag{8}$$

The use of such a model can describe a number of challenging problems in medical image analysis. However these simplistic models can only account for simple interactions between variables, a rather constrained scenario for high-level medical imaging perception tasks. One can augment the expression power of this model through higher order interactions between variables, or a number of cliques $\{C_i, i \in [1, n] = \{\{p_{i^1}, \cdots, p_{i^{|C_i|}}\}\}$ of order $|C_i|$ that will augment the definition of $\mathcal{V}$ and will introduce hyper-vertices:

$$\arg\min_{\{x_p\}} \mathcal{P}(g, f) = \sum_{p \in \mathcal{V}} g_p(x_p) + \sum_{(p_1, p_2) \in \mathcal{E}} f_{p_1 p_2}(x_{p_1}, x_{p_2}) + \sum_{C_i \in \mathcal{E}} f_{p_1 \cdots p_n}(x_{p_{i^1}}, \cdots, p_{x_{i^{|C_i|}}}). \tag{9}$$

where $f_{p_1 \cdots p_n}$ is the price to pay for associating the labels $(x_{p_{i^1}}, \cdots, p_{x_{i^{|C_i|}}})$ to the nodes $(p_1 \cdots p_{i^{|C_i|}})$. Parameter inference, addressed by minimizing the problem above, is the most critical aspect in computational medicine and efficient optimization algorithms are to be evaluated both in terms of computational complexity as well as of inference performance. State of the art methods include deterministic and non-deterministic annealing, genetic algorithms, max-flow/min-cut techniques and relaxation. These methods offer certain strengths while exhibiting certain limitations, mostly related to the amount of interactions which can be tolerated among neighborhood nodes. In the area of medical imaging where domain knowledge is quite strong, one would expect that such interactions should be enforced at the largest scale possible.

4. Application Domains

4.1. Clinical Projects

- **MR & Muscular Diseases**: The use of MR and Diffusion Tensor Imaging are investigated in collaboration with the Henri Mondor University Hospital and Institut of Myology towards automatic quantification of muscular mass loss and non-invassive biopsy. The aim is to provide tools that could be used to automatically analyze MR imaging and extract useful clinical measurements (Insitut of Myology), and assess the potential impact of diffusion tensor imaging towards automatic quantification either of muscular diseases progression.

- **Image-driven Radiotherapy Treatment & Surgery Guidance** : The use of CT and MR imaging for cancer guidance treatment in collaboration with the Oscar Lambert Center. The aim is to provide tools for automatic dose estimation as well as off-line and on-line positioning guidance through deformable fusion between imaging data corresponding to perioding patient treatment. The same concept will be explored in collaboration with the Saint-Antoine University Hospital towards image-driven surgery guidance through 2D to 3D registration between interventional and pre-operative annotated data.
- **MR Brain Imaging towards Low-Gliomas Tumor Brain Understanding**: The use of contrast enhanced imaging is investigated in collaboration with the Montpellier University Hospital towards better understanding of low-gliomas positioning, automatic tumor segmentation/identification and longitudinal (tumor) growth modeling.
- **CT/MR Perfusion Imaging**: The use of perfusion imaging is investigated in collaboration with the Georges Pompidou European Hospital towards compartmental analysis and measuring tissue perfusion and capillary permeability in liver tumors.

5. Software

5.1. Deformable Registration Software

Participant: Nikos Paragios [Correspondant].

deformable image and volume registration, is a deformable registration platform in C++ for the medical imaging community (publicly available at http://www.mrf-registration.net) developed mainly at Ecole Centrale, Technical University of Munich and University of Crete. This is the first publicly available platform which contains most of the existing metrics to perform registration under the same concept. The platform is used for clinical research from approximately 3,000 users worldwide.

5.2. Dense image and surface descriptors

Participant: Iasonas Kokkinos [Correspondant].

Scale-Invariant Descriptor, Scale-Invariant Heat Kernel Signatures DISD (publicly available at http://vision.mas.ecp.fr/Personnel/iasonas/descriptors.html) implements the SID, SI-HKS and ISC descriptors. SID (Scale-Invariant Descriptor) is a densely computable, scale- and rotation- invariant descriptor. We use a log-polar grid around every point to turn rotation/scalings into translation, and then use the Fourier Transform Modulus (FTM) to achieve invariance. SI-HKS (Scale-Invariant Heat Kernel Signatures) extract scale-invariant shape signatures by exploiting the fact that surface scaling amounts to multiplication and scaling of a properly sampled HKS descriptor. We apply the FTM trick on HKS to achieve invariance to scale changes. ISC (Intrinsic Shape Context) constructs a net-like grid around every surface point by shooting outwards and tracking geodesics. This allows us to build a meta-descriptor on top of HKS/SI-HKS that takes neighborhood into account, while being invariant to surface isometries.

5.3. Dissimilarity Coefficient learning

Participant: Pawan Kumar [Correspondant].

weakly supervised learning, dissimilarity coefficient, structured prediction DISC (publicly available at http://www.centrale-ponts.fr/personnel/pawan/code/DISCAPI.zip) software provides a convenient API for dissimilarity coefficient (DISC) based learning. DISC allows the use of weakly supervised datasets (with missing information) by jointly learning a structured prediction classifier and a conditional probability distribution of the missing information. The parameters of the classifier and the distribution are learned by minimizing a user-specified dissimilarity coefficient between them.

5.4. Efficient bounding-based object detection

Participant: Iasonas Kokkinos [Correspondant].

branch-and-bound, parts detection, segmentation, DPMS implements branch-and-bound object detection, cutting down the complexity of detection from linear in the number of pixels to logarithmic (publicly available at http://vision.mas.ecp.fr/Personnel/iasonas/dpms.html). The results delivered are identical to those of the standard deformable part model detector, but are available in 5 to 20 times less time. This website has been visited 1500 times in 10 months.

5.5. Fast Primal Dual Strategies for Optimization of Markov Random Fields

Participant: Nikos Komodakis [Correspondant].

discrete optimization, Markov random field, duality, graph cuts, FASTPD is an optimization platform in C++ for the computer vision and medical imaging community (publicly available at http://www.csd.uoc.gr/~komod/FastPD/) developed mainly at Ecole Centrale and University of Crete. This is the most efficient publicly available platform in terms of a compromise of computational efficiency and ability to converge to a good minimum for the optimization of generic MRFs. The platform is used from approximately 1,500 users worldwide.

5.6. imaGe-based Procedural Modeling Using Shape Grammars

Participant: Iasonas Kokkinos [Correspondant].

procedural modeling, image-based building reconstruction, shape grammars GRAPES is a generic image parsing library based on re-inforcement learning (publicly available at http://cvc.centrale-ponts.fr/). It can handle grammars (binary-split, four-color, Hausmannian) and image-based rewards (Gaussian mixtures, Randomized Forests) of varying complexity while being modular and computationally efficient both in terms of grammar and image rewards. The platform is used from approximately 500 users worldwide.

5.7. Learning-based symmetry detection

Participant: Stavros Tsogkas [Correspondant].

Scale-Invariant Descriptor, Scale-Invariant Heat Kernel Signatures LBSD (publicly available at http://www.centrale-ponts.fr/personnel/tsogkas/code.html implements the learning-based approach to symmetry detection published in [32]. It includes the code for running a detector, alongside with the ground-truth symmetry annotations that we have introduced for the Berkeley Segmentation Dataset (BSD) benchmark.

5.8. Texture Analysis Using Modulation Features and Generative Models

Participant: Iasonas Kokkinos [Correspondant].

Texture, modulation, generative models, segmentation, TEXMEG is a front-end for texture analysis and edge detection platform in Matlab that relies on Gabor filtering and image demodulation (publicly available at http://cvsp.cs.ntua.gr/software/texture/). Includes frequency- and time- based definition of Gabor- and other Quadrature-pair filterbanks, demodulation with the Regularized Energy Separation Algorithm and Texture/Edge/Smooth classification based on MDL criterion. The platform is used from approximately 250 users worldwide.

5.9. Sparse Prediction

Participant: Andreas Argyriou [Correspondant].

Sparse prediction, K-support norm, SPARSE_K is a sparse prediction code (publicly available at http://www.centrale-ponts.fr/personnel/andreas/code/sparse_k/sparse_k.tar) using regularization with the k-support norm, which we have introduced [36]. The algorithm uses an accelerated first-order method similar to Nesterov's method.

6. New Results

6.1. Machine Learning & Optimization

Participants: Andreas Argyriou, Matthew Blaschko, Pawan Kumar.

- **Sparse Prediction & Convex Optimization Decomposition** [*Andreas Argyriou*]
  In [36], we have introduced a new regularization penalty for sparse prediction, the k-support norm. This norm corresponds to the tightest convex relaxation of sparsity combined with an ℓ_2 penalty. We have shown that this new norm provides a tighter relaxation than the elastic net, and is thus a good replacement for the Lasso or the elastic net in sparse prediction problems. In [41], motivated by learning problems we proposed a novel optimization algorithm for minimizing a convex objective which decomposes into three parts: a smooth part, a simple non-smooth Lipschitz part, and a simple non-smooth non-Lipschitz part.
- **Learning Optimization for NP-complete Inference** [*Matthew Blaschko*]
  In [14] an optimization strategy for learning to optimize boolean satisfiability (SAT) solvers is given. Applications to real-world SAT problems show improved computational performance as a result of the learning algorithm.
- **Max-Margin Min-Entropy Models & Dissimilarity Coefficient based Learning** [*Pawan Kumar*]
  In [22] we proposed the family of max-margin min-entropy (M3E) models, which predicts a structured output for a given input by minimizing the Renyi entropy. The parameters of M3E are learned by minimizing an upper bound on a user-defined loss. We demonstrated the efficacy of M3E on two problems using publicly available datasets: motif finding and image classification. In [19] we proposed a novel structured prediction framework for weakly supervised datasets. The framework minimizes a dissimilarity coefficient between the predictor and a conditional distribution over the missing information. We demonstrated the efficacy of our approach on two problems using publicly available datasets: object detection and action detection.

6.2. Computational Vision & Perception

Participants: Matthew Blaschko, Iasonas Kokkinos, Pawan Kumar, Nikos Paragios.

- **Structured Output Ranking & Detailed Understanding of Objects in Computer Vision** [*Matthew Blaschko*]
  In [23] we proposed a novel method for efficiently optimizing an objective that ranks structured outputs by their loss. Based on the observation that structured output spaces [9] in computer vision problems can be well-modeled by a small number of loss values, our algorithm is able to optimize a quadratic number of pairwise constraints in linear time. In [38] we detail the research activities of a summer workshop hosted by Johns Hopkins University on learning a detailed understanding of objects and scenes in natural images. We worked on automatic verification of annotations provided through Amazon Mechanical Turk [35], texture categorization, and dependence modeling for bottom up proposals.
- **Efficient inference and learning for structured probabilistic models of deformable objects** [*Iasonas Kokkinos*, Haithem Boussaid & Stavros Tsogkas]

We have developed novel features to describe surface points intrinsically through the Intrinsic Shape Context (ISC) descriptor published in [17]. This method has delivered state-of-the-art results in surface point matching and we will explore its use for surface correspondence. The implementation of these descriptors is publicly available. In [32] we proposed a learning-based approach to symmetry detection by fusing multiple cues related to image intensity, color and texture, which delivered state-of-the-art results. We intend to extend this approach to 3D image analysis, and in particular for medical images. The implementation of these detectors is publicly available. In [27] we introduce a grouping-based method to learn and detect action classes in spatio-temporal data. Our method can both classify actions and indicate the spatio-temporal structures which provide support for the decision. The implementation of our front-end is publicly available. In [40] we have extended our work on efficient algorithms for object detection to accommodate fast methods for computing the part scores in a principled optimization framework, while he have thoroughly presented it in [40] and made the implementation publicly available.

- **Multi-view Image Segmentation & Parsing** [*Nikos Paragios*]
In [28] a method for image matching was proposed that exploits hierarchical image representations through higher order graphs. The matching was achieved through a graph-based theoretical framework where the similarity and spatial consistency of the image semantic objects is encoded in a graph of commute times that is also endowed with singleton terms through shape descriptors. Many-to-many matching of regions are specially challenging due to the instability of the segmentation under slight image changes, and we explicitly handle it through high order potentials. These ideas were further explored in the context of co-segmentation [29] where a method to determine a consistent partition of multiple images was introduced through a multi-scale multiple-image generative model based on region matching that exploits inter-image information and establishes correspondences between the common objects that appear in the scene. Last, but not least in [24] a method that combines bottom up (visual information, visual descriptors, elements detection) information and top-town models (hierarchical shape grammars) was considered towards automatic facade parsing though reinforcement learning while in [30] a method for 3D image parsing was proposed based on a hierarchical grammar that was performing explicit 3D modeling of the scene through a combination of multi-image segmentation and a depth reconstruction process. The problem optimal combination of these two concurrent terms was addressed trough a pareto-driven criterion while the optimization was addressed through an evolutionary computation algorithm.

6.3. Biomedical Image Analysis

Participant: Nikos Paragios.

- **Image Reconstruction** [*Nikos Paragios* & Hellene Langet]
In [21] a novel iterative reconstruction algorithm based on compressed sensing was proposed for Digital Subtraction Rotational Angiography (DSRA) that exploits both spatial and temporal sparsity through a proximal implementation that accommodates multiple$L-1$-penalties. These ideas was further explored in [20] where we introduced a three-dimensional reconstruction of tomographic acquisitions in C-arm-based rotational angiography was proposed that was able to deal with the temporal variations due to intra-arterial injections through a compressed-sensing approach leading to significant motion artifacts reduction in spite of the cone-beam geometry, the short-scan acquisition, and the truncated and subsampled data.
- **Image Segmentation** [*Nikos Paragios*, Pierre-Yves Baudin, Xiang Bo & Sarah Parisot]
In [11] the problem of human skeletal muscle segmentation was considered through a graph-based approach (random walker). An automatic seed placement framework was introduced through a graph-theoretic formulation. Towards accounting for anatomical constraints, the Random Walker algorithm was endowed with a liner sub-space statistical prior towards improving segmentation robustness on missing and incomplete data [12]. The same formulation was extended to cope with non-linear priors through a Gaussian-like local prior model penalizing the deviations of the

coefficients of the random walker diffusion matrix from the ones learned from the training data [13]. In [25] a novel graph-based prior was considered towards modeling the distribution of low-glioma brain tumors and spatially characterizing them through a sparse hierarchical graph. Such a prior model was integrated to an image-driven voxel-like segmentation framework where image separation was achieved through a machine learning method towards automatic detection, characterization and segmentation of brain tumors. Furthermore, towards encoding pose invariance in the context of knowledge-based segmentation in [33] where a higher order graph-based implicit pose invariant formulation was introduced for cardiac segmentation. The formulation was endowed with higher order cliques allowing (i) the estimation of boundary and regional image support and (ii) the implicit modeling of local deformations with respect to a prior statistical model while being invariant to linear transformations.

- **Image Registration** [*Nikos Paragios*, Nicolas Honnorat & Sarah Parisot]
  In [15] the problem of organ-driven registration was addressed through simultaneous combined fusion of multi-modal images in the context of guide-wire segmentation through fluoroscopic and contrast enhanced images. To this end, a graphical model was considered that was segmenting and registering the guide-wire in the two modalities while establishing correspondences between the associated curves as well. Similar philosophy was used in the [26] where a method for one shot deformable brain registration and tumor segmentation was proposed between a healthy anatomical atlas and a diseased patient. Both tasks were addressed through a discrete formulation (pair-wise MRF using grid-like deformation models and machine learning discriminative frameworks for the separation of healthy versus diseased tissues) while interconnections between the two graphs were used to alleviate the registration requirement on tumor areas. The problem of symmetric registration was studied in [31] through a common grid deforming in both directions according to a symmetric manner towards minimizing the image similarity criterion between the source and the target image while guaranteeing the expected diffeomorphic nature of the deformation field.
- **Computational Anatomy** [*Nikos Paragios*]
  In [16] we introduced a novel approach for detecting the presence of white matter lesions in periventricular areas of the brain using manifold-constrained embeddings. The proposed method uses locally linear embedding (LLE) to create ”normality” distributions of the brain where deviations from the manifolds are estimated by calculating geodesic distances along locally linear planes in the embedding. Experiments highlight the need of nonlinear techniques to learn the studied data leading to outstanding detection rates when comparing individuals to a specific pathological pattern.

7. Bilateral Contracts and Grants with Industry

7.1. General Electric HealthCare

- Compressed Sensing Digital Subtraction Rotational Angiography [PhD thesis H. Langet: 2009-2012]:
- Guide-wire Segmentation and Tracking of in interventional Imaging [PhD thesis N. Honnorat: 2008-2012]

7.2. Intrasene

Modeling, segmentation and registration of low gliomas brain tumors [PhD thesis S. Parisot: 2010-2013]

7.3. Siemens

Graph-based Knowledge-based Segmentation of the Human Skeletal Muscle in MR Imaging [PhD thesis P-Y. Baudin: 2009-2012]

8. Partnerships and Cooperations

8.1. Regional Initiatives

- Program: DIGITEO CHAIR

  Project acronym: SuBSAmPLE

  Project title: identification and prediction of Salient Brain StAtes through ProbabiListic structure learning towards fusion of Imaging and Genomic data

  Duration: 1/2012-12/2015

  Coordinator: ECP-FR

- Program: DIGITEO OMTE

  Project acronym: CURATEUR

  Project title: Real-time Multi-sensor (2D/3D) Elastic Image Fusion towards Computer-assisted Tumor Removal Surgery

  Duration: 1/2012-6/2014

  Coordinator: ECP-FR

8.2. National Initiatives

8.2.1. ANR

- Program: ANR Blanc International

  Project acronym: ADAMANTIUS

  Project title: Automatic Detection And characterization of residual Masses in pAtients with lymphomas through fusioN of whole-body diffusion-weighTed mrI on 3T and 18F-flUorodeoxyglucoSe pet/ct

  Duration: 9/2012-8/2015

  Coordinator: CHU Henri Mondor - FR

- Program: ANR JCJC

  Project acronym: HICORE

  Project title: HIerarchical COmpositional REpresentations for Computer Vision

  Duration: 10/2010-9/2013

  Coordinator: ECP - FR

8.2.2. Competitivity Clusters

- Program: MEDICEN

  Project acronym: ADOC

  Project title: ADOC – Diagnostic peropératoire numérique en chirurgie du cancer

  Duration: 11/2011-10/2014

  Coordinator: LLTECH - FR

8.3. European Initiatives

8.3.1. FP7 Projects

- Project acronym: MOBOT

  Project title: Intelligent Active MObility Assistance RoBOT integrating Multimodal Sensory Processing, Proactive Autonomy and Adaptive Interaction

Duration: 36 months

Coordinator: TUM - DE

- Project acronym: RECONFIG

  Project title: Cognitive, Decentralized Coordination of Heterogeneous Multi-Robot Systems

  Duration: 36 months

  Coordinator: KTH - SE

8.3.2. Collaborations in European Programs, except FP7

- Program: European Research Council

  Project acronym: DIOCLES

  Project title: Discrete bIOimaging perCeption for Longitudinal Organ modEling and computEr-aided diagnosiS

  Duration: 9/2011-8/2016

  Coordinator: ECP - FR

8.3.3. Collaborations with Major European Organizations

Technical University of Munich, Chair for Computer Aided Medical Procedures & Augmented Reality - Computer Science Department (DE): Mono and Multi-modal image fusion using discrete optimization and efficient linear programming.

Università della Svizzera Italiana, Institute of Computational Science (CH), Construction of deformation-invariant surface descriptors [39] and meta-descriptors for surfaces [17].

8.4. International Initiatives

8.4.1. Inria Associate Teams

8.4.1.1. SPLENDID

Title: Self-Paced Learning for Exploiting Noisy, Diverse or Incomplete Data

Inria principal investigator: Pawan Kumar

International Partner (Institution - Laboratory - Researcher):

Stanford University (United States) - Artificial Intelligence Lab

Duration: 2012 - 2014

The goal of the project is to develop methods for learning accurate probabilistic models using diverse (consisting of fully and weakly supervised samples), incomplete (consisting of partially labeled samples) and noisy (consisting of mislabeled samples) data. To this end, we will build on the intuitions gained from self-paced human learning, where a child is first taught simple concepts using simple examples, and gradually increasing the complexity of the concepts and the examples. In the context of machine learning, we aim to impart the learner with the ability to iteratively adapt the model complexity and process the training data in a meaningful order. The efficacy of the developed methods will be tested on several real world computer vision and medical imaging applications using large, inexpensively assembled datasets.

8.4.2. Inria International Partners

- **Department of Diagnostic Radiology, University of Pennsylvania:** The GALEN and the Section of Biomedical Image Analysis - SBIA group (Pr. Christos Davatzikos) have an established collaboration during the past three years in the area of deformable image fusion. In this context, PhD candidates of the GALEN group spend time visiting the SBIA group, while Pr. Paragios participates at a Nantional Institute Health grant led by SBIA. Such a collaboration led to a number of outstanding rank journal and conference publications.

- **Department of Computer Science, StonyBrook, State University of New York:** The GALEN and the Image Analysis Lab - CBL (Pr. Dimitris Samaras) have an established collaboration during the past years in the area of graph-based methods in medical imaging and computer vision. Pr. Samaras holds a research professor position (DIGITEO chair) at Ecole Centrale de Paris. Such a collaboration led to a number of outstanding rank conference publications during the last years.
- **Chang Gung Memorial Hospital – Linkou, Taiwan:** In the context of France-Taiwan program sponsored from the French Science Foundation, GALEN (in collaboration with the department of radiology of Henri Mondor University Hospital), a project (ADAMANTIUS) was initiated with the Chang Gung Memorial Hospital – Linkou that is the largest private hospital in Taiwan. The aim of the project is to study the Automatic Detection And characterization of residual Masses in pAtients with lymphomas through fusioN of whole-body diffusion-weighTed mrI on 3T and 18F-flUorodeoxyglucoSe pet/ct.

8.5. International Research Visitors

8.5.1. Internships

- Aseem BEHL (from Nov 2012 to Dec 2012)

  Subject: Optimizing Average Precision using Weakly Supervised Data. The average-precision support vector machine (AP-SVM) optimizes an upper bound on the average-precision (AP) loss, which is often used as a measure of accuracy for binary classification. However, it does not handle partially annotated datasets. To address this shortcoming of AP-SVM, we proposed a novel latent AP-SVM formulation, which allows us to learn an accurate set of classifier parameters by minimizing a carefully designed difference-of-convex upper bound on the AP loss.

  Institution: International Institute of Information Technology (IIIT), Hyderabad (India).

  Enzo FERRANTE (from June 2012 until October 2012)

  Subject: Plane+Deformation 2D-3D multimodal data fusion. The goal of the internship was to study the use of discrete optimization methods in the context of 2D to 3D registration in biomedical image analysis. In particular the aim was to define a metric free graphical model formulation that is able to determine for a given 2D image the corresponding 3D volume plane along with the in plane deformation. The case of computer assisted surgery was considered as a test case between 2D interventional images and 3D pre-operative high resolution annotated data.

  Institution: Universidad Nacional del Centro de la Provincia de Buenos Aires (Argentina)

- Danny GOODMAN (Aug 2012)

  Subject: Discriminative Parameter Estimation for Random Walks Segmentation. While random walks (RW) provide an efficient formulation for segmentation, there use is restricted by the lack of an accurate learning framework that estimates its parameters. The main difficulty is that a user can only provide a hard segmentation of a training sample, instead of the optimal probabilistic segmentation. We overcome this deficiency by treating the optimal probabilistic segmentation as latent variables, which allows us to employ the latent SVM formulation for parameter estimation.

  Institution: Stanford University (USA).

  Ishan MISRA (from May 2012 until Aug 2012)

  Subject: Shape-from Shading analysis for Object Categories. The goal of the internship was to see whether shape-from-shading techniques can be used to recover the 3D geometry within an object category. Mr. Misra experimented with techniques for shape-from-shading under unknown illumination as well as surface recovery from a single image. Mr. Misra has delivered the source code for his software to our team, and we intend to use it in our on-going research.

Institution: IIIT HYDERABAD (India)

Bharat SINGH (from May 2012 until September 2012)

Subject: Sub-space real-time Deformable Registration. The aim of this internship was to investigate the use of sub-space image representations towards defining an appropriate metric in the context of mono-modal and multi-modal fusion. Furthermore, it was studied their integration in a graph-theoretic framework for deformable fusion that can benefit from its implementation on modern parallel architectures like graphics processing units.

Institution: IIT MADRAS (India)

Eduard TRULLS (from January 2012 until April 2012)

Subject: Segmentation-Aware Image Descriptors. The goal of the internship was to construct appearance descriptors that can exploit segmentation information in order to achieve invariance to background changes. Mr. Trulls implemented a dense descriptor that uses soft segmentation masks, and demonstrated that this results in substantially more invariant descriptors; he evaluated these descriptors on image registration (optical flow) and wide-baseline matching (stereo) where state-of-the-art results were obtained. This work has been submitted for publication and is under evaluation.

Institution: Universidad Polytecnica de Catalunia (UPC) (Spain)

8.5.2. Visits to International Teams

Matthew BLASCHKO & Iasonas KOKKINOS (from June 2012 until August 2012)

Subject: Center for Language and Speech Processing: Towards a Detailed Understanding of Objects and Scenes in Natural Images Workshop. The objective of this workshop was to develop novel methods to reliably extract from images a diverse set of attributes, and to use them to improve the accuracy, informativeness, and interpretability of object models. The goal is to combine advances in discrete-continuous optimisation, machine learning, and computer vision, to significantly advance our understanding of visual attributes and produce new state-of-the-art methods for their extraction.

Institution:John Hopkins University (USA)

Pawan KUMAR (from April 2012 until May 2012)

Subject: SPLENDID Associate Team

Institution: Stanford University (United States)

9. Dissemination

9.1. Scientific Animation

- **Andreas Argyriou**
  - **Conference Committee:** International Conference on Machine Learning (ICML), Neural Information Processing Systems (NIPS), Workshop on Optimization for Machine Learning (in NIPS).
  - **Invited Seminars/Presentations:** Télécom ParisTech - FR, École des Mines de Paris - FR, Imperial College - UK, Queen Mary University London - UK.
- **Matthew Blaschko**
  - **Conference Committee:** International Conference on Robotics and Automation (ICRA), British Machine Vision Conference (BMVC - area chair), Asian Conference on Computer Vision (ACCV), Neural Information Processing Systems (NIPS), Medical Image Computing and Computer Assisted Intervention (MICCAI), Medical Computer Vision Workshop (at MICCAI), European Conference on Computer Vision (ECCV)

  - **Journal Reviewing Services:** Journal of Machine Learning Research, International Journal of Computer Vision, IEEE Transactions on Pattern Analysis and Machine Intelligence, Computer Vision and Image Understanding
  - **Invited Seminars/Presentations:** Katholieke Universiteit Leuven - BE, Stanford University - USA, Google - Mountain View, Institute of Science and Technology - Austria, Johns Hopkins University - USA, Schlumberger - Paris - FR, Ecole Normale Superieure - FR.
- **Iasonas Kokkinos**
  - **Editorial Activities:** Image and Vision Computing Journal.
  - **Conference Committee:** IEEE International Conference on Computer Vision (CVPR - area chair), Perceptual Organization in Computer Vision (POCV - organizer) , Asian Conference on Computer Vision (ACCV), ACCV workshop on Detection and Tracking in Challenging Environments.
  - **Journal Reviewing Services:** IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Image Processing, Computer Vision and Image Understanding, SIAM Journal on Imaging Sciences, EURASIP Journal on Image and Video Processing.
  - **Invited Seminars/Presentations:** Ecole Normale Superieure - FR, Carnegie Mellon University - USA, Johns Hopkins University - USA, Schlumberger - Paris - FR, National Technical University of Athens - GR.
- **Pawan Kumar**
  - **Conference Committee**: IEEE Conference on Computer Vision and Pattern Recognition (ICCV), European Conference on Computer Vision (ECCV), Advances in Neural Information Processing Systems (NIPS).
  - **Journal Reviewing Services**: IEEE Transactions on Pattern Analysis and Machine Intelligence, Journal of Machine Learning Research.
  - **Workshop and Tutorials Organization**: European Signal Processing Conference tutorial on *Learning with Inference for Discrete Graphical Models*, Biomedical Image Analysis Summer School.
  - **Invited Seminars/Presentations**: University of Cambridge - UK, University of Oxford - UK, Institute of Science and Technology - AT, University of Heidelberg - DE, Stanford University - USA.
- **Nikos Paragios**
  - **Editorial Activities:** IEEE Transactions on Pattern Analysis and Machine Intelligence, International Journal of Computer Vision, Medical Image Analysis, Computer Vision and Image Understanding, Image and Vision Computing Journal, Machine Vision and Applications, SIAM Journal in Imaging Sciences.
  - **Conference Committee:** IEEE International Conference in Computer Vision (CVPR), IEEE Computer Vision and Pattern Recognition (ICPR - area chair), Medical Image Computing and Computer Assisted Intervention (MICCAI - area chair), IEEE International Symposium on Biomedical Imaging (ISBI), International Symposium on Visual Computing (ISVC), International Conference on Functional Imaging and Modeling of the Heart (FIMH), Medical Computer Vision Workshop (at MICCAI).
  - **Workshop and Tutorials Organization**: European Signal Processing Conference tutorial on *Learning with Inference for Discrete Graphical Models*, Biomedical Image Analysis Summer School.
  - **Journal Reviewing Services:** IEEE Transactions on Medical Imaging.

- **Invited Seminars/Presentations**: University of Pennsylvania- USA, Stony-Brook University - USA, Clermont-Ferrand University - FR, IEEE International Symposium on Biomedical Imaging (ISBI) - ES, Medical Computer Vision Workshop (MICCAI) - FR, Centre Oscar Lambret - FR.
- **Distinctions:** Member of the SAFRAN Conglomerate Scientific Council.

9.2. Teaching - Supervision - Committees

Participants: Matthew Blaschko, Iasonas Kokkinos, Pawan Kumar, Nikos Paragios.

9.2.1. Teaching

Master : Structure Prediction, 24, M1, Ecole Centrale de Paris [M. Blaschko]

Master : Discrete Optimization, 12, M1, Ecole Centrale de Paris [P. Kumar]

Master : Signal Processing, 36, M1, Ecole Centrale de Paris, France [I. Kokkinos]

Master : Computer Vision, 36, M1, Ecole Centrale de Paris, France [I. Kokkinos]

Master : Pattern Recognition, 24, M2, Ecole Centrale de Paris/Ecole Normale Superieure-Cachan, France [I. Kokkinos]

Master : Advanced Mathematical Models in Computer Vision, 24, M2, Ecole Centrale de Paris/Ecole Normale Superieure-Cachan, France [N. Paragios]

N. Paragios is in charge of the option Medical Imaging, Machine Learning and Computer Vision at the Department of Applied Mathematics of Ecole Centrale de Paris. This option consists of 7 classes in the above mentioned fields, 180 hours of teaching and is also directing the associated M.Sc. (M2) program of the ENS-Cachan in Applied Mathematics, Machine Learning and Computer Vision at Ecole Centrale de Paris.

9.2.2. Supervision

PhD in progress : Stavros Alchatzidis, "Message Passing Methods, Parallel Architectures & Visual Processing", 2011-2014, Nikos Paragios (supervisor)

PhD in progress : Pierre-Yves Baudin, "Knowledge-based Segmentation of the Human Skeletal Muscle through Learning & Inference of Randon Walks", 2009-2013, Nikos Paragios & Pierre Carlier (supervisors)

PhD in progress : Xiang Bo, "Pose-Invariant Knowledge-based Segmentation with Higher Order Graphs", 2009-2013, Nikos Paragios (supervisor)

PhD in progress : Haithem Boussaid, "Learning-based mid-level processing for computer vision and medical imaging", 2010-2014, Iasonas Kokkinos (supervisor)

PhD in progress : Enzo Ferrante, "2D-to-3D Multi-Modal Deformable Image Fusion", 2012-2015, Nikos Paragios (supervisor)

PhD in progress : Vivien Fecamp, "Linear-Deformable Multi-Modal Deformable Image Fusion", 2012-2015, Nikos Paragios (supervisor)

PhD in progress : Katerina Gkirtzou, "Kernels, Machine Learning & Biomedical Imaging-driven Computational Anatomy", 2009-2013, Nikos Paragios (supervisor)

PhD in progress : Nicolas Honnorat, "Detection, Segmentation & Tracking of Guide-Wires in Interventional Imaging", 2009-2013, Nikos Paragios (supervisor)

PhD in progress : Puneet Kumar, "Weakly Supervised Learning for Object Detection and Semantic Segmentation", 2010-2013, Pawan Kumar (supervisor)

PhD in progress : Helene Langet, "Sampling and Motion Reconstruction in Three-dimensional X-Ray Interventional Imaging", 2010-2013, Gilles Fleury & Nikos Paragios (supervisors)

PhD in progress : Fabrice Michel, "Metric Learning & Mono/Multi-modal Data Fusion", 2009-2013, Nikos Paragios (supervisor)

PhD in progress : Sarah Parisot, "Graph-based Detection, Characterization & Segmentation of Brain Tumors", 2010-2013, Nikos Paragios (supervisor)

PhD in progress : Stavros Tsogkas, "Learning-based mid-level processing for computer vision and medical imaging", 2011-2014, Iasonas Kokkinos (supervisor)

9.2.3. Committees

Nikos Paragios: Emannuel Caruyer (Reviewer - University of Nice-Sophia Antipolis), Yangming Ou (external member - University of Pennsylvania), Karima Ouji (Chair - Ecole Centrale de Lyon), Olivier Whyte (Chair - ENS-Cachan), Xiang Zheng (external member - StonyBrook University)

10. Bibliography

Major publications by the team in recent years

[1] M. BLASCHKO, C. LAMPERT. *Learning to Localize Objects with Structured Output Regression*, in "Proceedings of the 10th European Conference on Computer Vision: Part I", ECCV '08, 2008, p. 2–15.

[2] B. GLOCKER, A. SOTIRAS, N. KOMODAKIS, N. PARAGIOS. *Deformable Registration: Setting the State of the Art with Discrete Methods*, in "Annual Reviews on Biomedical Engineering", 2011, p. 219-244.

[3] I. KOKKINOS. *Rapid Deformable Object Detection using Dual-Tree Branch-and-Bound*, in "Neural Information Processing Systems (NIPS)", Granada, Spain, December 2011.

[4] I. KOKKINOS, A. YUILLE. *Inference and Learning with Hierarchical Shape Models*, in "International Journal of Computer Vision", 2011, vol. 93, n^{o} 2, p. 201-225.

[5] N. KOMODAKIS, N. PARAGIOS, G. TZIRITAS. *MRF Energy Minimization and Beyond via Dual Decomposition*, in "IEEE Trans. Pattern Anal. Mach. Intell.", 2011, vol. 33, n^{o} 3, p. 531-552.

[6] M. P. KUMAR, V. KOLMOGOROV, P. TORR. *An analysis of convex relaxations for MAP estimation of discrete MRFs*, in "JMLR - Journal of Machine Learning Research", 2009, http://hal.inria.fr/hal-00773608.

[7] M. P. KUMAR, P. TORR, A. ZISSERMAN. *OBJCUT: Efficient segmentation using top-down and bottom-up cues*, in "PAMI - IEEE Transactions on Pattern Analysis and Machine Intelligence", 2010, http://hal.inria.fr/hal-00773609.

[8] C. LAMPERT, M. BLASCHKO, T. HOFMANN. *Efficient Subwindow Search: A Branch and Bound Framework for Object Localization*, in "IEEE Trans. Pattern Anal. Mach. Intell.", December 2009, vol. 31, n^{o} 12, p. 2129–2142.

Publications of the year

Articles in International Peer-Reviewed Journals

[9] M. BLASCHKO, C. LAMPERT. *Guest Editorial: Special Issue on Structured Prediction and Inference*, in "International Journal of Computer Vision", 2012, vol. 99, n^{o} 3, p. 257–258.

[10] O. TEBOUL, I. KOKKINOS, L. SIMON, P. KOUTSOURAKIS, N. PARAGIOS. *Parsing Facades with Shape Grammars and Reinforcement Learning*, in "IEEE Transactions on Pattern Analysis and Machine Intelligence", 2013, Accepted for publication.

International Conferences with Proceedings

[11] P.-Y. BAUDIN, N. AZZABOU, P. CARLIER, N. PARAGIOS. *Automatic skeletal muscle segmentation through random walks and graph-based seed placement*, in "ISBI'12", Barcelone, Spain, 2012, p. 1036-1039.

[12] P.-Y. BAUDIN, N. AZZABOU, P. CARLIER, N. PARAGIOS. *Manifold-enhanced Segmentation through Random Walks on Linear Subspace Priors*, in "Proceedings of the British Machine Vision Conference", 2012, p. 51.1-51.10, http://hal.archives-ouvertes.fr/hal-00773635.

[13] P.-Y. BAUDIN, N. AZZABOU, P. CARLIER, N. PARAGIOS. *Prior Knowledge, Random Walks and Human Skeletal Muscle Segmentation*, in "Medical Image Computing and Computer-Assisted Intervention - MICCAI 2012", France, 2012, p. 569-576, http://hal.archives-ouvertes.fr/hal-00773665.

[14] A. FLINT, M. BLASCHKO. *Perceptron Learning of SAT*, in "Neural Information Processing Systems", Lake Tahoe, United States, 2012, http://hal.inria.fr/hal-00738219.

[15] N. HONNORAT, R. VAILLANT, N. PARAGIOS. *Graph-based guide-wire segmentation through fusion of contrast-enhanced and fluoroscopic images*, in "ISBI", 2012, p. 948-951.

[16] S. KADOURY, G. ERUS, E. ZACHARAKI, N. PARAGIOS, C. DAVATZIKOS. *Manifold-constrained embeddings for the detection of white matter lesions in brain MRI*, in "ISBI", 2012, p. 562-565.

[17] I. KOKKINOS, M. BRONSTEIN, R. LITTMAN, A. BRONSTEIN. *Intrinsic Shape Context Descriptors for Deformable Shapes*, in "IEEE Conf. on Computer Vision and Pattern Recognition (CVPR 12)", Providence, USA, June 2012.

[18] I. KOKKINOS. *Bounding Part Scores for Rapid Detection with Deformable Part Models*, in "2nd Parts and Attributes Workshop, in conjunction with ECCV 12 ", Florence, Italy, October 2012.

[19] M. P. KUMAR, B. PACKER, D. KOLLER. *Modeling latent variable uncertainty for loss-based learning*, in "ICML - International Conference on Machine Learning", Edinburgh, UK, 2012, http://hal.inria.fr/hal-00773605.

[20] H. LANGET, C. RIDDELL, Y. TROUSSET, A. TENENHAUS, E. LAHALLE, G. FLEURY, N. PARAGIOS. *Compressed Sensing Dynamic Reconstruction in Rotational Angiography*, in "MICCAI'12", Nice, France, 2012, vol. 7510, p. 223-230 [*DOI :* 10.1007/978-3-642-33415-3_28], http://hal.inria.fr/hal-00756119.

[21] H. LANGET, C. RIDDELL, Y. TROUSSET, A. TENENHAUS, E. LAHALLE, G. FLEURY, N. PARAGIOS. *Compressed sensing subtracted rotational angiography with multiple sparse penalty*, in "ISBI'12", Barcelone, Spain, 2012, p. 884-887, http://hal.inria.fr/hal-00725950.

[22] K. MILLER, M. P. KUMAR, B. PACKER, D. GOODMAN, D. KOLLER. *Max-margin min-entropy models*, in "AISTATS - International Conference on Artificial Intelligence and Statistics", La Palma, Canary Islands, 2012, http://hal.inria.fr/hal-00773602.

[23] A. MITTAL, M. BLASCHKO, A. ZISSERMAN, P. TORR. *Taxonomic Multi-class Prediction and Person Layout using Efficient Structured Ranking*, in "European Conference on Computer Vision", 2012.

[24] D. OK, M. KOZINSKI, R. MARLET, N. PARAGIOS. *High-Level Bottom-Up Cues for Top-Down Parsing of Facade Images*, in "3DIMPVT", 2012, p. 128-135.

[25] S. PARISOT, H. DUFFAU, S. CHEMOUNY, N. PARAGIOS. *Graph-based Detection, Segmentation & Characterization of Brain Tumors*, in "CVPR - 25th IEEE Conference on Computer Vision and Pattern Recognition 2012", Providence, United States, 2012, http://hal.inria.fr/hal-00712714.

[26] S. PARISOT, H. DUFFAU, S. CHEMOUNY, N. PARAGIOS. *Joint Tumor Segmentation and Dense Deformable Registration of Brain MR Images*, in "MICCAI (2)", 2012, p. 651-658.

[27] M. RAPTIS, I. KOKKINOS, S. SOATTO. *Discovering Discriminative Action Parts from Mid-Level Video Representations*, in "IEEE Conf. on Computer Vision and Pattern Recognition (CVPR 12)", Providence, USA, June 2012.

[28] J. RUBIO, J. SERRAT, A. LÓPEZ, N. PARAGIOS. *Image Contextual Representation and Matching through Hierarchies and Higher Order Graphs*, in "ICPR", 2012.

[29] J. RUBIO, J. SERRAT, A. LÓPEZ, N. PARAGIOS. *Unsupervised co-segmentation through region matching*, in "CVPR", 2012, p. 749-756.

[30] L. SIMON, O. TEBOUL, P. KOUTSOURAKIS, L. VAN GOOL, N. PARAGIOS. *Parameter-free/Pareto-driven procedural 3D reconstruction of buildings from ground-level sequences*, in "CVPR", 2012, p. 518-525.

[31] A. SOTIRAS, N. PARAGIOS. *Discrete symmetric image registration*, in "ISBI", 2012, p. 342-345.

[32] S. TSOGKAS, I. KOKKINOS. *Learning-Based Symmetry Detection in Natural Images*, in "European Conference on Computer Vision (ECCV 12) ", Florence, Italy, October 2012.

[33] B. XIANG, C. WANG, J.-F. DEUX, A. RAHMOUNI, N. PARAGIOS. *3D cardiac segmentation with pose-invariant higher-order MRFs*, in "ISBI", 2012, p. 1425-1428.

Conferences without Proceedings

[34] H. LANGET, C. RIDDELL, A. TENENHAUS, Y. TROUSSET, E. LAHALLE, G. FLEURY, N. PARAGIOS. *Tomographic Reconstruction with L1-Minimization Strategy for Rotational Angiography*, in "20th International Conference on computational statistics (COMPSTAT'12)", Limassol, Cyprus, August 2012, http://hal.inria.fr/hal-00756446.

[35] N. SAPHRA, J. KANNALA, E. RAHTU, M. BLASCHKO. *Automatic Verification of Boundary Annotations from Mechanical Turk*, in "Women in Machine Learning Workshop", 2012.

Scientific Books (or Scientific Book chapters)

[36] A. ARGYRIOU, R. FOYGEL, N. SREBRO. *Sparse Prediction with the k-Support Norm*, in "Advances in Neural Information Processing Systems 25", Lake Tahoe, United States, P. BARTLETT, F. PEREIRA, C. BURGES, L. BOTTOU, K. WEINBERGER (editors), 2012, p. 1466–1474, http://arxiv.org/abs/1204.5043.

Books or Proceedings Editing

[37] N. PARAGIOS, J. DUNCAN, N. AYACHE. , N. PARAGIOS, J. DUNCAN, N. AYACHE (editors)*Biomedical Image Analysis: Methodologies And Applications*, Springer, 2013, -, 590 pages, http://hal.inria.fr/inria-00616178.

Research Reports

[38] M. BLASCHKO, R. GIRSHICK, J. KANNALA, I. KOKKINOS, S. MAHENDRAN, S. MAJ, S. MOHAMED, E. RAHTU, N. SAPHRA, K. SIMONYAN, B. TASKAR, A. VEDALDI, D. WEISS. *Towards a Detailed Understanding of Objects and Scenes in Natural Images*, Johns Hopkins University, Center for Language and Speech Processing Summer Workshop, 2012.

[39] I. KOKKINOS, M. BRONSTEIN, A. YUILLE. *Dense Scale Invariant Descriptors for Images and Surfaces*, Inria, March 2012, n^o RR-7914, http://hal.inria.fr/hal-00682775.

[40] I. KOKKINOS. *Rapid Deformable Object Detection using Bounding-based Techniques*, Inria, April 2012, n^o RR-7940, http://hal.inria.fr/hal-00696120.

[41] F. ORABONA, A. ARGYRIOU, N. SREBRO. *PRISMA: PRoximal Iterative SMoothing Algorithm*, Inria, November 2012, http://arxiv.org/abs/1206.2372.

[42] A. SOTIRAS, C. DAVATZIKOS, N. PARAGIOS. *Deformable Medical Image Registration: A Survey*, Inria, September 2012, n^o RR-7919, http://hal.inria.fr/hal-00684715.

[43] C. WANG, N. PARAGIOS. *Markov Random Fields in Vision Perception: A Survey*, Inria, September 2012, n^o RR-7945, http://hal.inria.fr/hal-00734983.