IN PARTNERSHIP WITH:
Ecole Centrale Paris

Activity Report 2013

Project-Team GALEN

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

RESEARCH CENTER
Saclay - Île-de-France

THEME
Computational Neuroscience and Medecine

Table of contents

Project-Team GALEN

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

Creation of the Team: 2008 February 15*, updated into Project-Team:* 2014 January 01.

1. Members

Research Scientist
Matthew Blaschko [Inria, Researcher]

Faculty Members
Nikos Paragios [Team leader, Ecole Centrale de Paris, Professor, HdR]
Iasonas Kokkinos [Ecole Centrale de Paris, Associate Professor, HdR]
Andreas Argyriou [Ecole Centrale de Paris]
Pawan Kumar Mudigonda [Ecole Centrale de Paris, HdR]

Engineers
Aline Guegou [Ecole Centrale de Paris, from Feb 2013 until Jul 2013]
Rafael Marini Silva [Ecole Centrale de Paris, from Feb 2013]
Wojciech Zaremba [Ecole Centrale de Paris, until Jun 2013]
Joachim Zentici [Ecole Centrale de Paris, from Feb 2013 until Jul 2013]
Erwan Zerhouni [Inria, granted by Conseil Régional d'Ile-de-France, from Dec 2013]

PhD Students
Bo Xiang [Inria, granted by Conseil Régional d'Ile-de-France, until Nov 2013]
Stavros Alchatzidis [Ecole Centrale de Paris]
Haithem Boussaid [Ecole Centrale de Paris]
Vivien Fecamp [Inria]
Enzo Ferrante [Ecole Centrale de Paris]
Aikaterini Gkirtzou [Ecole Centrale de Paris, until Dec 2013]
Wacha Bounliphone [Ecole Centrale de Paris, from Nov 2013]
Evgenios Kornaropoulos [Ecole Centrale de Paris, from Sep 2013]
Sarah Parisot [Ecole Centrale de Paris, until Nov 2013]
Eduard Trulls [Ecole Centrale de Paris, Mar 2013 until Dec 2013]
Stavros Tsogkas [Ecole Centrale de Paris]
Puneet Kumar [Ecole Centrale de Paris]

Administrative Assistants
Christine Biard [Inria]
Natalia Leclercq [Ecole Centrale de Paris, from Nov 2013]
Carine Morotti-Delorme [Ecole Centrale de Paris, until Oct 2013]

Others
Mateusz Michalkiewicz [Inria, from Nov 2013]
Jose Ignacio Orlando [Inria, from May 2013 until Nov 2013]
Siddhartha Chandra [Inria, M.Sc. from May 2013]
Dimitrios Damopoulos [Ecole Centrale de Paris, from Nov 2013]
Grigorios Chrysos [Ecole Centrale de Paris, from Sep 2013]
Lufang Gao [Inria, from Nov 2013]

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 8^{th} and July 12^{th}, 2013 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.
- **Coursera**: Pawan Kumar Mudigonda & Nikos Paragios introduced a new course on discrete inference and learning in artificial vision on the Coursera platform with approx 15,000 student enrollments.
- **Editor in Chief**: Nikos Paragios was named editor in chief of the Computer Vision and Image Understanding Journal (CVIU). CVIU is published by Elsevier Publishing House and is one of the oldest and leading journals in the field of computer vision and image understanding. In 2009, it was named one of the top 20 journals in computer science by Times Higher Education.

3. Research Program

3.1. Shape, Grouping and Recognition

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. If we phrase image interpretation as mapping an observed image, X to a set of symbols Y, we are interested 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 for image and object representation (b) learning a scoring function s from training data and (c) optimizing over Y in Eq. 1. One of the main themes of our work is the development of methods that jointly address (a,b,c) in a shape-grouping framework in order to reliably extract, describe, model and detect shape information from natural and medical images. A principal motivation for using a shape-based framework is the understanding that shape- and more generally, grouping- based representations can go all the way from image features to objects. Regarding aspect (a), image representation, we cater for the extraction of image features that respect the shape properties of image structures. Such features are typically constructed to be purely geometric (e.g. boundaries, symmetry axes, image segments), or appearance-based, such as image descriptors. The use of machine learning has been shown to facilitate the robust and efficient extraction of such features, while the grouping of local evidence is known to be necessary to disambiguate the potentially noisy local measurements. In our research we have worked on improving feature extraction, proposing novel blends of invariant geometric- and appearance- based features, as well as grouping algorithms that allow for the efficient construction of optimal assemblies of local features.

Regarding aspect (b) we have worked on learning scoring functions for detection with deformable models that can exploit the developed low-level representations, while also being amenable to efficient optimization. Our works in this direction build on the graph-based framework to construct models that reflect the shape properties of the structure being modeled. We have used discriminative learning to exploit boundary- and symmetry-based representations for the construction of hierarchical models for shape detection, while for medical images we have developed methods for the end-to-end discriminative training of deformable contour models that combine low-level descriptors with contour-based organ boundary representations.

Regarding aspect (c) we have developed algorithms which implement top-down/bottom-up computation both in deterministic and stochastic optimization. The main idea is that 'bottom-up', image-based guidance is necessary for efficient detection, while 'top-down', object-based knowledge can disambiguate and help reliably interpret a given image; a combination of both modes of operation is necessary to combine accuracy with efficiency. In particular we have developed novel techniques for object detection that employ combinatorial optimization tools (A^* and Branch-and-Bound) to tame the combinatorial complexity, achieving a best-case performance that is logarithmic in the number of pixels. In our current work [27] we further accelerate object detection by integrating low-level processing (convolutions) with bounding-based object detection, while we have recently started exploring the potential of combinatorial optimization in the medical imaging realm [22]. Working with stochastic optimization tools, in [17] we have pursued the exploitation of reinforcement-learning to optimize over the set of shapes derivable from shape grammars.

In the long run we aim at scaling up shape-based methods to 3D detection and pose estimation and large-scale object detection. One aspect which seems central to this is the development of appropriate mid-level representations. This is a problem that has received increased interest lately in the 2D case and is relatively mature, but in 3D it has been pursued primarily through ad-hoc schemes. We anticipate that questions pertaining to part sharing in 3D will be addressed most successfully by relying on explicit 3D representations. On the one hand depth sensors, such as Microsoft's Kinect, are now cheap enough to bring surface modeling and matching into the mainstream of computer vision - so these advances may be directly exploitable at test time for detection. On the other hand, even if we do not use depth information at test time, having 3D information can simplify the modeling task during training. In on-going work with collaborators we have started exploring combinations of such aspects, namely (i) the use of surface analysis tools to match surfaces from depth sensors (ii) using branch-and-bound for efficient inference in 3D space and (iii) groupwise-registration to build statistical 3D surface models. In the coming years we intend to pursue a tighter integration of these different directions for scalable 3D object recognition.

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 $\mathbf{w}$. 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. Brain Tumors and Neuro-degenerative diseases

: 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. Furthermore, in collaboration with the Neurospin center of CEA and the Brookhaven National Laboratory at StonyBrook University we investigate the use of machine learning methods towards automatic interpretation of functional magnetic resonance imaging between cocaine addicted and normal subjects. Last, but not least in collaboration with the Georges Pompidou European Hospital an effort toward understanding tumor perfusion process through comportemental models is carried out with emphasis given on elastic organs.

4.2. Image-driven Radiotherapy Treatment & Surgery Guidance

The use of CT and MR imaging for cancer guidance treatment in collaboration with the Gustave Roussy Institute of Oncology. 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 prior to each session and the ones used for scheduling/planning and dose estimation. 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.

5. Software and Platforms

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://cvn.ecp.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://vision.mas.ecp.fr/Personnel/teboul/grapesPage/index.php). 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://cvn.ecp.fr/personnel/tsogkas/code.html implements the learning-based approach to symmetry detection. 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://cvn.ecp.fr/personnel/andreas/code/sparse_k/sparse_k.tar) using regularization with the k-support norm, which we have introduced [39]. The algorithm uses an accelerated first-order method similar to Nesterov's method.

6. New Results

6.1. Shape, Grouping and Recognition

6.1.1. Descriptors

Participants: Eduard Trulls, Iasonas Kokkinos.

In [30] we have extended our prior work on dense scale- and rotation- invariant image descriptors to take into account soft segmentation information. This allows us to discard measurements stemming from background structures, and as such renders our descriptors invariant to background changes and occlusions. This has allowed us to obtain state-of-the-art results on tasks such as large-displacement optical flow and wide-baseline stereo. We have made the implementation of these descriptors publicly available.

6.1.2. 3D structure detection

Participants: Haithem Boussaid, Iasonas Kokkinos.

In [22] we have started exploring the potential of combinatorial optimization in the medical imaging realm. We cast the problem of finding a 3D structure (a brain tumor) as that of finding the mode of a nonparametric distribution, constructed through Kernel Density Estimation. Current techniques for doing this (e.g. Mean Shift mode-seeking, Fast Gauss Transforms, etc.) are either iterative, or linear in the number of pixels, with a typically large constant. Instead, we develop a scheme that involves a very low-constant linear-time preprocessing step, and then uses Branch-and-Bound for fast mode estimation. As such it is scalable to large volumes, and serves as a rapid initialization of a region segmentation algorithm.

6.1.3. Facade parsing

Participants: Olivier Teboul, Iasonas Kokkinos, Loic Simon, Panagiotis Katsourakis, Nikos Paragios.

In [17] we pursue a Reinforcement Learning-based approach to couple image observations with a grammar-based method to partitioning a building facade. For this we expressed 2D grammar- based image parsing as a Markov decision process where an agent has to take actions in an environment so as to maximize some notion of cumulative reward (reflecting the segmentation quality). This allowed us to accelerate previous stochastic hill-climbing approaches to image parsing by more than an order of magnitude.

6.1.4. Fast object detection

Participant: Iasonas Kokkinos.

In [27] we extended our previous work on fast object detection by developing a sparse-coding method for the efficient sharing of computation among multiple object models. In particular the first processing step of 'part score' computation was originally performed separate per object category; instead, we propose to do it 'in batch mode', so as to exploit the commonalities that exist among object parts. Building on recent developments in sparse coding we have managed to construct a compact basis for this task, which in the end gave us a two-fold acceleration over our previous fastest algorithms.

6.2. Machine Learning

6.2.1. Discriminative Parameter Estimation for Random Walks Segmentation

Participants: Pierre-Yves Baudin, Puneet Kumar, Noura Azzabou, Pierre Carlier, Nikos Paragios, M. Pawan Kumar Blaschko

In [19], we proposed a a novel discriminative learning framework that estimates the parameters of a random walks segmentation framework using a training dataset. The main challenge we face is that the training samples are not fully supervised. Specifically, they provide a hard segmentation of the medical images, instead of a probabilistic segmentation. We overcome this challenge by treating the optimal probabilistic segmentation that is compatible with the given hard segmentation as a latent variable. This allows us to employ the latent support vector machine (LSVM) formulation for parameter estimation.

6.2.2. Structured Sparsity & Applications

Participants: Katerina Gkirtzou, Wojciech Zaremba, Matthew Blaschko, M. Pawan Kumar, Nikos Paragios

We developed several machine learning applications to fMRI data, including graph representations [25] and structured sparsity regularization [26], [44]. A similar structured sparsity approach was applied in the development of a novel learning algorithm, the k-support regularized SVM, with applications to neuromuscular disease classification from diffusion tensor imaging [24]. Efficient training applications for taxonomic classification were developed in [21], while a fine grained taxonomic image classification task was introduced in [45]. The role of non-maximal suppression in accurate and efficient object detection cascades was elucidated in [20]. A fast, consistent two-sample test based on kernelized statistics was developed in [33].

6.2.3. Learning from M/EEG Data with Variable Brain Activation Delays

Participants: Wojciech Zaremba, Alexander Gramfort, M. Pawan Kumar, Matthew Blaschko

In [34], propose to address the misalignment of M/EEG samples by explicitly modeling time shifts of different brain responses in a classification setup. To this end, we use the LSVM formulation, where the latent shifts are inferred while learning the classifier parameters. The inferred shifts are further used to improve the signal-to-noise ratio of the M/EEG data, and to infer the chronometry and the sequence of activations across the brain regions that are involved in the experimental task.

6.3. Biomedical Image Analysis

6.3.1. Reconstruction

Participants: Helen Langet, Nikos Paragios

In [38] an overview of the methodological foundations of biomedical image analysis as well as their use to provide answers to a variety of clinical problems are presented. The problem of volumes of rotational angiography using non-linear sparsity constraints was studied in [28] where a novel method able to handle highly under-sampled acquisitions was introduced.

6.3.2. Graphical models and Image Segmentation

Participants: Bo Xiang, Nikos Paragios

[18] presents an overview of the use of graphical models in artificial vision where both inference, learning as well as applications are discussed. In [32] a max-margin dual decomposition method was used towards learning the compact, pose invariant shape representation using higher order graphs acting both on the connectivity of the graph as well its potentials. Graphical model was used as prior in [13] under a "curve" propagation principle for generic prior-constrained organ segmentation in 2D images. Similar inspiration driven from a higher order pose invariant graphical model learned according to [32] was considered in [31] where a novel segmentation method was proposed coupling model-based and pixel-based concepts while being pose invariant. The underlying idea was to consider a two-layer interconnected graphical model acting on pixel and on control points where segmentation consistency was imposed throgh penalties on label discrepancies of the different layers. Higher order graphical models were also employed in [14] for spine segmentation using an articulated graphical model where a non-linear approach/embedding towards reducing the complexity of the inference step was considered at training.

6.3.3. Deformable Registration and Fusion

Participants: Enzo Ferrante, Sarah Parisot, Nikos Paragios

In [16] a comprehensive survey of deformable registration was presented. It was organized in three sections: the first was studying the deformation model, the second the similarity criterion while the last section discussed the different optimization strategies. The problem of atlas-based segmentation/registration in the presence of brain tumors was studied in [29] an adaptive uncertainty-driven sampling strategy was proposed coupling segmentation and registration. Both sampling spaces (quantization of the search space, deformation grid) were determined according to the observed optimization min-marginals. The challenging problem of image to slice registration was proposed in [23] where an over-parameterized low rank graphical model acting both on the plan selection as well the in-plane deformations was introduced. The main strength of the method was its ability to simultaneously recover both the plane and the organ deformation.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- **General Electric HealthCare**:
  - Compressed Sensing Digital Subtraction Rotational Angiography [PhD thesis H. Langet]
  - Guide-wire Segmentation and Tracking of in interventional Imaging [PhD thesis N. Honnorat]
- **Intrasene**: Modeling, segmentation and registration of low gliomas brain tumors [PhD thesis S. Parisot]
- **Siemens**: Graph-based Knowledge-based Segmentation of the Human Skeletal Muscle in MR Imaging [PhD thesis P-Y. Baudin]

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. Excellence Clusters

- Program: DIGITEO (OMTE)

  Project acronym: Curator

  Project title: Real-time 2D/3D Deformable Fusion Towards Computer Assisted Surgery

  Duration: 01/2013-01/2014

  Coordinator: ECP - FR
- Program: MEDICEN

  Project acronym: ADOC

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

  Duration: 11/2011-10/2014

  Coordinator: LLTECH - 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: 10/2010-9/2014

  Coordinator: ECP - FR

- Program: ITMOs Cancer & Technologies pour la santé d'Aviesan / INCa

  Project acronym: CURATOR

  Project title: Slice-to-Image Deformable Registration towards Image-based Surgery Navigation & Guidance

  Duration: 12/2013-11/2015

  Coordinator: ECP - 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: 01/2013-12/2015

  Coordinator: TUM - DE

- Project acronym: RECONFIG

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

  Duration: 01/2013-12/2015

  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

8.4. 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: Nikos Paragios

International Partner (Institution - Laboratory - Researcher):

Stanford University (United States) - Artificial Intelligence Lab - Nikos Paragios

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

8.4.2.1. Informal International Partners

Europe

- Technical University of Munich (DE) – Collaborative research with the Chair for Computer Aided Medical Procedures & Augmented Reality at the department of Computer Science. Collaboration Topic: Graph-based methods for linear/deformable registration, segmentation, and tracking.
- University College London (UK) – Collaborative research with the Gatsby Computational Neuroscience Unit. Collaboration Topic: Kernel measures of dependence.
- University of Oxford (UK) – Collaborative research with the Visual Geometry Group of the Department of the Electrical Engineering. Collaboration Topic: Structured prediction, invariance, and parts-based models.
- University of Oulu (Finland) – Collaborative research with the Machine Vision Group at the department of Electrical Engineering. Collaboration Topic: Ranking based learning algorithms for cascaded object detection.

Americas

- University of California at Los Angeles (US) – Collaborative research with the UCLA Vision Lab and the UCLA Center for Cognition, Vision, and Learning Lab at the Departments of Computer Science and Statistics. Collaboration Topic: Action Recognition & Object Detection Parsing.
- University of Pensylvania (USA) – Collaborative research with the section of Biomedical Imaging of the Department of Radiology. Collaboration Topic: Graph-based methods for linear/deformable registration.
- StonyBrook University, Computer Science Department (USA) – Collaborative research with the image analysis lab in the context of the SubSample DIGITEO Chair. Collaboration Topic: Higher Order Graph-based methods in graph-matching, cocaine addiction analysis with sparse graph models, object detection and implicit 3D pose estimation
- Ecole Polytechnique de Montreal (CA) – Collaborative research with the Canada Research Chair in Medical Imaging and Assisted Interventions. Collaboration Topic: Higher Order Graph-based methods in Spine Imaging

Asia

- International Institute of Information Technology, Hyderabad (India) – Collaborative research with Center for Visual Information Technology. Collaboration Topic: Average precision with weak supervision & self-paced learning for deep convolutional neural networks.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

Professor Spyretta Golemati lecturer at the school of medicine at the University of Athens has visited during her sabbatical the team from June 1st, 2013 to July 30th, 2013.

8.5.1.1. Internships

Siddhartha Chandra

Subject: machine learning for 3D reasoning.

Date: from May 2013 until December 2013.

Institution: IIIT Hyderabad (India)

Dimitrios Damopoulos

Subject: Automatic Detection and Characterization of Liver Tumors

Date: from Nov 2013 until Apr 2014

Institution: National Technical University of Athens (Greece)

José Ignacio Orlando

Subject: Machine Learning for Opthalmology

Date: from Apr 2013 until Sep 2013

Institution: National University of the Center of the Buenos Aires Province (Argentina)

Eduard Trulls

Subject: Segmentation-aware descriptors

Date: from March 2013 until July 2013

Institution: Polytechnical University of Catalunia (Spain)

8.5.2. Visits to International Teams

- M. Pawan Kumar (Inria): one week visit to Stanford University (May 2013).
- M. Pawan Kumar (Inria): one week visit to Stanford University (June 2013).
- Matthew Blaschko (Inria): one week visit to Stanford University (December 2013).

9. Dissemination

9.1. Scientific Animation

- **Andreas Argyriou**
  - **Conference Committee:** International Joint Conference on Artificial Intelligence (IJCAI), International Conference on Machine Learning (ICML), Advances in Neural Information Processing Systems (NIPS)
  - **Workshop and Tutorials Organization**: International Workshop on Advances in Regularization, Optimization, Kernel Methods and Support Vector Machines: theory and applications (ROKS).
- **Matthew Blaschko**

  - **Conference Committee:** British Machine Vision Conference (BMVC - area chair), IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Neural Information Processing Systems (NIPS), Medical Image Computing and Computer Assisted Intervention (MICCAI), International Conference in Computer Vision (ICCV)
  - **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:** Stanford University - USA.
- **Iasonas Kokkinos**
  - **Editorial Activities:** Associate Editor, Image and Vision Computing Journal.
  - **Editorial Activities:** Guest Editor, Computer Vision and Image Understanding Journal.
  - **Conference Committee:** International Conference on Computer Vision (ICCV), International Conference on Computer Vision (CVPR), Artificial Intelligence and Statistics (AISTATS),
  - Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR).
  - **Journal Reviewing Services:** International Journal of Computer Vision, IEEE Transactions on Pattern Analysis and Machine Intelligence, Computer Vision and Image Understanding, Machine Vision and Applications.
  - **Invited Seminars/Presentations:** Institute for Pure and Applied Mathematics (IPAM) - USA, University of Oulu - FI, Kondrad-Zuse Center for Computer Science - DE, Stony Brook University - USA.
- **Nikos Paragios**
  - **Editorial Activities:** Editor in Chief, Computer Vision and Image Understanding.
  - **Editorial Activities:** Associate Editor, 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.
  - **Editorial Activities:** Guest Editor, IEEE Transactions on Pattern Analysis and Machine Ingelligence, Medical Image Analysis.
  - **Conference Committee:** IEEE International Conference in Computer Vision (ICCV-area chair), IEEE International Conference in Computer Vision (CVPR), Information Processing in Medical Imaging (IPMI), Medical Image Computing and Computer Assisted Intervention (MICCAI - area chair).
  - **Workshop and Tutorials Organization**: Biomedical Image Analysis Summer School (BIOMED).
  - **Journal Reviewing Services:** International Journal of Computer Vision, IEEE Transactions on Medical Imaging, NeuroImage.
  - **Invited Seminars/Presentations**: Siemens Corporate Research - USA, Colloqium on Data Science in the Big Data Era - FR, SEERSS International Congress in Robotic Surgery - GR.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

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

9.2.2. 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.

9.2.3. Supervision

HdR : Pawan Kumar, Weakly Supervised Learning for Structured Output Prediction, Ecole Normale Supérieure de Cachan - ENS Cachan, 12/2013

HdR : Iasonas Kokkinos, Learning and Optimization for Shape-based Representations, Université Paris-Est, 9/2013

PhD: Pierre-Yves Baudin, Graph-based Segmentation of Skeletal Striated Muscles in NMR Images, Ecole Centrale de Paris, 05/2013, Nikos Paragios

PhD: Katerina Gkirtzou, Sparsity Regularization and Graph-based Representations in Medical Imaging, 12/2013, Ecole Centrale de Paris, Nikos Paragios

PhD: Nicolas Honnorat, Curvilinear Structures Segmentation and Tracking in Interventional Imaging, 01/2013, Ecole Centrale de Paris, Nikos Paragios

PhD: Helene Langet, Sampling and Motion Reconstruction in Three-dimensional X-Ray Interventional Imaging, 03/2013, Ecole Centrale de Paris, Gilles Fleury & Nikos Paragios

PhD: Fabrice Michel, Multi-Modal Similarity Learning for 3D Deformable Registration of Medical Images, 10/2013, Ecole Centralse de Paris, Nikos Paragios

PhD: Sarah Parisot, Graph-based Detection, Characterization & Segmentation of Brain Tumors, 11/2013, Ecole Centrale de Paris, Nikos Paragios

PhD: Bo Xiang, Knowledge-Based Image SegmentationUsing Sparse Shape Priors and High-Order MRFs, Ecole Centrale de Paris, 11/2013, Nikos Paragios

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

PhD in progress : Wacha Bounliphone, Sparse Methods towards data mining in Bio-informatics& Bio-imaging, 2013-2016, Matthew Blaschko

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

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

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

PhD in progress : Evgenios Kornaropoulos, Diffusion Coefficient: a novel computer aided biomarker, 2010-2013, Nikos Paragios

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

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

9.2.4. Juries

- **Andreas Argyriou**
  - **PhD Thesis Participation**: G. Zappella - IT (PhD).
- **Matthew Blaschko**
  - **PhD Thesis Participation**: K. Gkirtzou - FR (PhD).
- **Iasonas Kokkinos**
  - **Grant Reviewing Services:** Swiss National Science Foundation.
- **Nikos Paragios**
  - **PhD Thesis Participation**: P-Y. Baudin - FR (PhD), K. Gkirtzou - FR (PhD), M. Heinrich - UK (PhD), N. Honnorat - FR (PhD), P. Kumar (HDR), H. Langet - FR (PhD), S. Merlet - FR (PhD), F. Michel - FR (PhD), S. Parisot - FR (PhD), B. Xiang - FR (PhD).
  - **Grant Reviewing Services:** Agence National de la Recherche, Austrian Research Council, Danish Research Council, Dutch Research Council, European Research Council, Israel Research Foundation, Swiss National Science Foundation.

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, pp. 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, pp. 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, pp. 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, pp. 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, pp. 2129–2142

Publications of the year

Doctoral Dissertations and Habilitation Theses

[9] P.-Y. BAUDIN. , *De la segmentation au moyen de graphes d'images de muscles striés squelettiques acquises par RMN*, Ecole Centrale Paris, May 2013, http://hal.inria.fr/tel-00858584

[10] N. HONNORAT. , *Segmentation et suivi de structures curvilinéaires en imagerie interventionnelle*, Ecole Centrale Paris, January 2013, http://hal.inria.fr/tel-00801865

[11] I. KOKKINOS. , *Apprentissage et Optimization pour des Representations basées sur la Forme*, Université Paris-Est, September 2013, Habilitation à Diriger des Recherches, http://hal.inria.fr/tel-00857643

[12] M. P. KUMAR. , *Weakly Supervised Learning for Structured Output Prediction*, École normale supérieure de Cachan - ENS Cachan, December 2013, Habilitation à Diriger des Recherches, http://hal.inria.fr/tel-00943602

Articles in International Peer-Reviewed Journals

[13] D. CHITTAJALLU, N. PARAGIOS, I. KAKADIARIS. *An Explicit Shape-constrained MRF-based Contour Evolution Method for 2D Medical Image Segmentation*, in "IEEE Journal of Biomedical and Health Informatics", April 2013, pp. 2168-2194 [*DOI :* 10.1109/JBHI.2013.2257820], http://hal.inria.fr/hal-00858880

[14] S. KADOURY, H. LABELLE, N. PARAGIOS. *Spine Segmentation in Medical Images Using Manifold Embeddings and Higher-Order MRFs*, in "IEEE Transactions on Medical Imaging", July 2013, vol. 32, n^o 7, pp. 1227-1238 [*DOI :* 10.1109/TMI.2013.2244903], http://hal.inria.fr/hal-00856319

[15] A. PANAGOPOULOS, W. CHAOHUI, D. SAMARAS, N. PARAGIOS. *Simultaneous Cast Shadows, Illumination and Geometry Inference Using Hypergraphs*, in "IEEE Transactions on Pattern Analysis and Machine Intelligence", February 2013, vol. 35, n^o 2, pp. 437-449 [*DOI :* 10.1109/TPAMI.2012.110], http://hal.inria.fr/hal-00855591

[16] A. SOTIRAS, C. DAVATZIKOS, N. PARAGIOS. *Deformable Medical Image Registration: A Survey*, in "IEEE Transactions on Medical Imaging", May 2013, vol. 32, n^o 7, pp. 1153-1190 [*DOI :* 10.1109/TMI.2013.2265603], http://hal.inria.fr/hal-00858737

[17] O. TEBOUL, I. KOKKINOS, L. SIMON, K. PANAGIOTIS, N. PARAGIOS. *Parsing Facades with Shape Grammars and Reinforcement Learning*, in "IEEE Transactions on Pattern Analysis and Machine Intelligence", 2013, vol. 35, n^o 7, pp. 1744-1756 [*DOI :* 10.1109/TPAMI.2012.252], http://hal.inria.fr/hal-00855609

[18] C. WANG, N. KOMODAKIS, N. PARAGIOS. *Markov Random Field Modeling, Inference & Learning in Computer Vision & Image Understanding: A Survey*, in "Computer Vision and Image Understanding", 2013, vol. 117, n^o 11, pp. 1610-1627 [*DOI :* 10.1016/J.CVIU.2013.07.004], http://hal.inria.fr/hal-00858390

International Conferences with Proceedings

[19] P.-Y. BAUDIN, D. GOODMAN, P. KUMAR, N. AZZABOU, P. G. CARLIER, N. PARAGIOS, M. PAWAN KUMAR. *Discriminative Parameter Estimation for Random Walks Segmentation*, in "16th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2013", Nagoya, Japan, September 2013, 8 p. , http://hal.inria.fr/hal-00856020

[20] M. BLASCHKO, J. KANNALA, E. RAHTU. *Non Maximal Suppression in Cascaded Ranking Models*, in "Scandanavian Conference on Image Analysis", Espoo, Finland, April 2013, pp. 408-419 [*DOI :* 10.1007/978-3-642-38886-6_39], http://hal.inria.fr/hal-00815374

[21] M. BLASCHKO, W. ZAREMBA, A. GRETTON. *Taxonomic Prediction with Tree-Structured Covariances*, in "European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases", Prague, Czech Republic, June 2013, pp. 304-319 [*DOI :* 10.1007/978-3-642-40991-2_20], http://hal.inria.fr/hal-00839775

[22] H. BOUSSAID, I. KOKKINOS, N. PARAGIOS. *Rapid Mode Estimation for 3D Brain MRI Tumor Segmentation*, in "Energy Minimization Methods in Computer Vision and Pattern Recognition", Lund, Sweden, 2013, http://hal.inria.fr/hal-00856770

[23] E. FERRANTE, N. PARAGIOS. *Non-rigid 2D-3D Medical Image Registration using Markov Random Fields*, in "16th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2013", Nagoya, Japan, Springer, 2013, vol. 8151, pp. 163-170, http://hal.inria.fr/hal-00855662

[24] K. GKIRTZOU, J.-F. DEUX, G. BASSEZ, A. SOTIRAS, A. RAHMOUNI, T. VARACCA, N. PARAGIOS, M. BLASCHKO. *Sparse classification with MRI based markers for neuromuscular disease categorization*, in "4th International Workhop on Machine Learning in Medical Imaging", Nagoya, Japan, Springer, September 2013, http://hal.inria.fr/hal-00845126

[25] K. GKIRTZOU, J. HONORIO, D. SAMARAS, R. GOLDSTEIN, M. BLASCHKO. *fMRI Analysis with Sparse Weisfeiler-Lehman Graph Statistics*, in "4th International Workhop on Machine Learning in Medical Imaging", Nagoya, Japan, Springer, September 2013, http://hal.inria.fr/hal-00845068

[26] K. GKIRTZOU, J. HONORIO, D. SAMARAS, R. GOLDSTEIN, M. BLASCHKO. *FMRI Analysis of Cocaine Addiction Using K-Support Sparsity*, in "International Symposium on Biomedical Imaging", San Francisco, United States, IEEE, January 2013, http://hal.inria.fr/hal-00784386

[27] I. KOKKINOS. *Shufflets: Shared Mid-level Parts for Fast Multi-Category Detection*, in "ICCV - International Conference on Computer Vision", Sydney, Australia, 2013, http://hal.inria.fr/hal-00857578

[28] H. LANGET, A. RESHEF, C. RIDDELL, Y. TROUSSET, A. TENENHAUS, E. LAHALLE, G. FLEURY, N. PARAGIOS. *Nonlinear diffusion constraints for reconstructing subsampled rotational angiography data*, in "Fully3D 2013", Lake Tahoe, California, United States, June 2013, pp. 38-41, http://hal.inria.fr/hal-00933983

[29] S. PARISOT, W. WELLS III, S. CHEMOUNY, H. DUFFAU, P. NIKOS. *Uncertainty-driven Efficiently-Sampled Sparse Graphical Models for Concurrent Tumor Segmentation and Atlas Registration*, in "ICCV - 14th International Conference on Computer Vision", Sydney, Australia, 2013, http://hal.inria.fr/hal-00858696

[30] E. TRULLS, I. KOKKINOS, A. SANFELIU, F. MORENO-NOGUER. *Dense Segmentation-aware Descriptors*, in "Computer Vision and Pattern Recognition", Portland, Oregon, United States, 2013, http://hal.inria.fr/hal-00856023

[31] B. XIANG, J.-F. DEUX, A. RAHMOUNI, N. PARAGIOS. *Joint Model-Pixel Segmentation with Pose-invariant Deformable Graph-Priors*, in "16th International Conference on Medical Image Computing and Computer Assisted Intervention", Nagoya, Japan, September 2013, http://hal.inria.fr/hal-00856955

[32] B. XIANG, N. KOMODAKIS, N. PARAGIOS. *Pose Invariant Deformable Shape Priors Using L1 Higher Order Sparse Graphs*, in "9th International Symposium on Visual Computing - ISVC 2013", Rethymnon, Greece, July 2013, http://hal.inria.fr/hal-00856978

[33] W. ZAREMBA, A. GRETTON, M. BLASCHKO. *B-tests: Low Variance Kernel Two-Sample Tests*, in "Neural Information Processing Systems", Lake Tahoe, United States, December 2013, http://hal.inria.fr/hal-00842098

[34] W. ZAREMBA, M. P. KUMAR, A. GRAMFORT, M. BLASCHKO. *Learning from M/EEG data with variable brain activation delays*, in "International Conference on Information Processing in Medical Imaging", Asilomar, United States, March 2013, http://hal.inria.fr/hal-00803981

[35] Y. ZENG, W. CHAOHUI, D. GU, D. SAMARAS, N. PARAGIOS. *A Generic Deformation Model for Dense Non-Rigid Surface Registration: a Higher-Order MRF-based Approach*, in "IEEE International Conference on Computer Vision 2013 - ICCV 2013", Sydney, Australia, 2013, 8 p. , http://hal.inria.fr/hal-00856323

Scientific Books (or Scientific Book chapters)

[36] K. DANIILIDIS, P. MARAGOS, N. PARAGIOS. , *International Journal of Computer Vision Special Issue on Novel Representations, Methods, and Algorithms in Computer Vision*, Springer, July 2013, 90 p. , http://hal.inria.fr/hal-00858678

[37] C. WANG, Y. ZENG, D. SAMARAS, N. PARAGIOS. *Modeling Shapes with Higher-Order Graphs: Methodology and Applications*, in "Shape Perception in Human and Computer Vision: An Interdisciplinary Perspective", S. J. DICKINSON, Z. PIZLO (editors), 2013, pp. 459-471, http://hal.inria.fr/hal-00858417

Books or Proceedings Editing

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

Research Reports

[39] A. ARGYRIOU, S. CLÉMENÇON, R. ZHANG. , *Learning the Graph of Relations Among Multiple Tasks*, October 2013, http://hal.inria.fr/hal-00940321

[40] P.-Y. BAUDIN, D. GOODMAN, P. KUMAR, N. AZZABOU, P. G. CARLIER, N. PARAGIOS, M. PAWAN KUMAR. , *Discriminative Parameter Estimation for Random Walks Segmentation: Technical Report*, September 2013, http://hal.inria.fr/hal-00830564

[41] M. PAWAN KUMAR, H. TURKI, D. PRESTON, D. KOLLER. , *Parameter Estimation and Energy Minimization for Region-based Semantic Segmentation*, August 2013, http://hal.inria.fr/hal-00857918

Scientific Popularization

[42] M. BLASCHKO. *Machine Learning for Neurological Disorders*, in "Centraliens", March 2014, n° 632, pp. 30-31, http://hal.inria.fr/hal-00940262

Other Publications

[43] A. ARGYRIOU, L. BALDASSARRE, C. A. MICCHELLI, M. PONTIL. , *On Sparsity Inducing Regularization Methods for Machine Learning*, 2013, 12 p. , arXiv admin note: text overlap with arXiv:1104.1436, http://hal.inria.fr/hal-00855984

[44] M. BLASCHKO. , *A Note on k-support Norm Regularized Risk Minimization*, March 2013, http://hal.inria.fr/hal-00804592

[45] S. MAJI, E. RAHTU, J. KANNALA, M. BLASCHKO, A. VEDALDI. , *Fine-Grained Visual Classification of Aircraft*, 2013, http://hal.inria.fr/hal-00842101