IN PARTNERSHIP WITH:
Ecole Centrale Paris

Activity Report 2014

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 Scientists

Andreas Argyriou [Ecole Centrale de Paris, until Sep 2014]
Matthew Blaschko [Inria, Researcher/Ecole Centrale de Paris, Associate Professor, HdR]

Faculty Members

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

Engineers

Rafael Marini Silva [Ecole Centrale de Paris]
Erwan Zerhouni [Inria, until Oct 2014]

PhD Students

Stavros Alchatzidis [Ecole Centrale de Paris]
Eugène Belilovsky [Ecole Centrale de Paris, from Mar 2014]
Maxim Berman [Ecole Centrale de Paris, from Sep 2014]
Diane Bouchacourt [Ecole Centrale de Paris, from Oct 2014]
Haithem Boussaid [Ecole Centrale de Paris]
Wacha Bounliphone [Ecole Centrale de Paris]
Siddhartha Chandra [Inria]
Puneet Dokania [Ecole Centrale de Paris]
Vivien Fecamp [Inria]
Enzo Ferrante [Ecole Centrale de Paris]
Hariprasad Kannan [Ecole Centrale de Paris, from Oct 2014]
Evgenios Kornaropoulos [Ecole Centrale de Paris]
Huu Dien Khue Le [Ecole Centrale de Paris, from Sep 2014]
Stavros Tsogkas [Ecole Centrale de Paris]
Jiaqian Yu [Ecole Centrale de Paris, from Mar 2014]

Administrative Assistants

Alexandra Merlin [Inria]
Natalia Leclercq [Ecole Centrale de Paris]

Others

Aimilia Gastounioti [Ecole Centrale de Paris, Visitor, PhD Student, from Feb until Jun. 2014]
Eduard Trulls [Ecole Centrale de Paris, Visitor, PhD Student, from Jun until Oct. 2014]
Shanmukha Ramakrishna Vedantam [Inria, Visitor, PhD Student, from Jun 2014 until Aug 2014]
Grigorios Chrysos [Ecole Centrale de Paris, Intern, until Jun 2014]
Dimitrios Damopoulos [Ecole Centrale de Paris, Intern, until Jul 2014]
Lufang Gao [Inria, Intern, until Feb 2014]
Mateusz Michalkiewicz [Inria, Intern, until May 2014]
Julia Rackerseder [Ecole Centrale de Paris, Intern, from Jun 2014 until Dec 2014]

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.

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. 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.

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 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 & Structured 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 $f_{p_1p_2}(x_{p_1}, x_{p_2})$. 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 online 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.

4.3. Fundus Image Analysis

Retinal images–also known as fundus images or retinographies–are projective color im- ages of the inner surface of the human eye. In collaboration with Pladema Institute, UNCPBA, Argentina, we are developing a suite of software tools for automatic analysis of retinal images driven by statistical learning approaches.

5. New 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. Ranking with High-Order Information

Participant: Puneet Dokania [Correspondant].

Average precision optimization, high-order information, ranking The software (publicly available at http://cvn.ecp.fr/projects/ranking-highorder/) provides a convenient API for learning to rank with high-order information. The samples are ranked according to a scorethat is proportional to the difference of max-marginals of the positive and the negative class. The parameters of the score function are computed by minimizing an upper bound on the average precision loss. The software also provides an instantiation of the API for ranking samples according to their relevance to an action, using the poselet features.

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.

6. New Results

6.1. Highlights of the Year

- Handbook of Biomedical Imaging: Methodologies and Clinical Research [38] - co-edited from Nikos Paragios, James Duncan and Nicholas Ayache - has been published from Springer Publishing house.
- Nikos Paragios was admitted as a senior fellow at the Insitut Universitaire de France and has been awarded an IBM Faculty award. He has also been one of the plenary invited lecturers at the IARP International Conference in Pattern Recognition (ICPR'2015, Stockholm).

BEST PAPER AWARD :

[26] **Predicting cross-task behavioral variables from fMRI data using the k-support norm in Sparsity Techniques in Medical Imaging (STMI)**. M. MISYRLIS, A. KONOVA, M. BLASCHKO, J. HONORIO, N. ALIA-KLEIN, R. GOLDSTEIN, D. SAMARAS.

6.2. Rounding-based Moves for Metric Labeling

Paticipants: M. Pawan Kumar

Metric labeling is an important special case of energy minimizaton in Markov random fields. While the best known polynomial-time algorithm for the problem is the linear programming (LP) relaxation, in practice it is slow to solve it. In [25], we introduced a new family of efficient move-making algorithms for metric labeling. These algorithms mimic the rounding procedues used for converting a fractional LP solution to a feasible integral solution. Our algorithms provide a matching theoretical guarantee to the LP relaxation, while requiring significantly less computational time.

6.3. Optimizing Average Precision

Paticipants: Puneet Kumar Dokania, Aseem Behl, Pritish Mohapatra, C.V. Jawahar, M. Pawan Kumar

Average precision (AP) is one of the most commonly used measures for ranking. However, due to the inefficiency of optimizing it during learning, a common approach is to use surrogate loss functions such as 0-1 loss. In [27], we proposed a new optimization algorithm for AP-SVM that allows training in a similar time to binary SVM. In [23], we extended the AP-SVM framework to score the samples according to high-order information, as opposed to simple first-order information used in prior work. Finally, in [19], we proposed a novel latent AP-SVM formulation that allows learning from weakly supervised datasets. The advantage of learning with high order and missing information is demonstrated on challenging computer vision problems such as action classification and object detection using standard benchmark datasets.

6.4. Discriminative Training of Deformable Contour Models

Paticipants: Haithem Boussaid, Iasonas Kokkinos and Nikos Paragios

Deformable Contour Models (DCMs) are a main workhorse for medical image analysis - but are not commonly studied from a machine learning perspective. In [21], [20] we haved proposed an integrated machine learning and optimization framework to deploy DCMs in medical image analysis.

Our technical contributions are two-fold: firstly, we use an efficient decomposition-coordination algorithm to solve the optimization problems resulting from Loopy DCMs, by means of the Alternating Direction Method of Multipliers (ADMM); this yields substantially faster convergence than plain Dual Decomposition-based methods.

Secondly, we use structured prediction to exploit loss functions that better reflect the performance criteria used in medical image segmentation. By using the mean contour distance (MCD) as a structured loss during training, we obtain clear test-time performance gains.

We demonstrate the merits of exact and efficient inference with rich, structured models in a large X-Ray image segmentation benchmark, where we obtain systematic improvements over the current state-of-the-art.

6.5. Improved Deformable Part Models for Object Detection

Paticipants: Iasonas Kokkinos, Stavros Tsogkas, Eduard Trulls, Pierre-Andre Savalle, George Papandreou.

In [30] and [36] we have worked on improving the classification accuracy of Deformable Part Models (DPMs) for object detection in two distinct manners. Firstly, in [30] we propose a technique to combine bottom-up segmentation, coming in the form of SLIC superpixels, with sliding window DPM detectors. The merit of our approach lies in 'cleaning up' the low- level features by exploiting the spatial support indicated by segmentation. - tion, for both the root and part filters of DPMs. We use these masks to construct enhanced, background- invariant features to train DPMs. We test our approach on the PASCAL VOC 2007, outperforming the standard DPM in 17 out of 20 classes, yielding an average increase of 1.7AP. Additionally, we demonstrate the robustness of this ap- proach, extending it to dense SIFT descriptors for large dis- placement optical flow.

Secondly, in [36] we have explored the potential of convolutional neural networks as feature extractors for detection with DPMs. In particular, we substitute the Histogram-of-Gradient features of DPMs with Convolutional Neural Network (CNN) features, and demonstrate that we thereby obtain a substantial boost in performance (+14.5 mAP) when compared to the baseline HOG-based models. Some more recent extensions to this work are included in [41] where we explore the potential of explicit scale and aspect ratio search in the context of sliding window detection with CNNs.

6.6. Fine-Grained models of objects and texture

Paticipants: Iasonas Kokkinos, Matthew Blaschko, Stavros Tsogkas, Andrea Vedaldi, Mircea Cimpol, Subhransu Maji, Ross Girshick, Juho Kannala, Esa Rahtu, David Weiss, Ben Taskar, Karen Simonyan.

In [31] and [22] we explore methods for the fine-grained understanding of objects and textures, respectively.

In [22] we introduce a texture dataset that is accompanied by descriptions that capture the essence of the textures in terms of attributes. We explore a broad range of classification techniques for these texture attributes and demonstrate that the learned classifiers help improve generic texture recognition methods.

In [31] we introduce a large-scale dataset of airplanes that is accompanied by thorough human annotations at different levels: airplane types, segment lineouts, attributes, and part descriptions are provided for more than 7000 airplane images. We explore the potential of these rich annotations for the task of constructed fine-grained image descriptions using discriminative training techniques on top of standard image representations (Histogram-of-gradient features).

6.7. Large Scale Video Segmentation

Paticipants: Matthew Blaschko

Spatio-temporal cues are powerful sources of information for segmentation in videos. In [24] we present an efficient and simple technique for spatio-temporal segmentation that is based on a low-rank spectral clustering algorithm. The complexity of graph-based spatio-temporal segmentation is dominated by the size of the graph, which is proportional to the number of pixels in a video sequence. In contrast to other works, we avoid oversegmenting the images into super-pixels and instead generalize a simple graph based image segmentation. Our graph construction encodes appearance and motion information with temporal links based on optical flow. For large scale data sets naïve graph construction is computationally and memory intensive, and has only been achieved previously using a high power compute cluster. We make feasible for the first time large scale graph-based spatio-temporal segmentation on a single core by exploiting the sparsity structure of the problem and a low rank factorization that has strong approximation guarantees.

6.8. Higher Order Graph Matching

Paticipants: Chaohui Wang, Dimitris Samaras, Nikos Paragios

In [42] a generic framework for sparse and dense graph/3D surface matching has been introduced. The framework is endowed with a novel mathematical formulation regarding the matching process along with a novel deformation model. It exploits the power of invariance of higher order clique potentials and through a low to high resolution approach determines optimal correspondences between two sets of 3D points while taking advantage of Mobius tranformation to measure local similarity of shapes/graphs/surfaces. Graph matching of objects undergoing non-rigid deformations along with temporal 3D surface tracking demonstrated the potentials of our method. Inference is solved through an efficient dual decomposition scema.

6.9. Inference of Procedural Grammars from Images

Paticipants: Nikos Paragios

Grammar-like representations are powerful modeling and inference tools in computational vision. In [39] a novel approach towards automatic inference of typology specific building grammars has been introduced. The central idea was to consider that such grammars could be derived through a bottom up approach of common sub-tree reasoning of derivation trees determined through parsing using elementary shape (binary split) grammars. Such an approach performs common subtree reduction within the entire training set and identifies meta-rules (corresponding to the same subtrees) which are then clustered together towards producing a compact, typology specific grammar. Promising results both in terms of grammar compactness as well as in terms of inference demonstrated the potentials of the method that could be used beyond the considered scoped.

6.10. Fully connected CRFs for blood vessel segmentation in retinal images

Paticipants: Matthew Blaschko, José Ignacio Orlando

In [28], we present a novel method for blood vessel segmentation in fundus images based on a discriminatively trained, fully connected conditional random field model. Retinal image analysis is greatly aided by blood vessel segmentation as the vessel structure may be considered both a key source of signal, e.g. in the diagnosis of diabetic retinopathy, or a nuisance, e.g. in the analysis of pigment epithelium or choroid related abnormalities. Blood vessel segmentation in fundus images has been considered extensively in the literature, but remains a challenge largely due to the desired structures being thin and elongated, a setting that performs particularly poorly using standard segmentation priors such as a Potts model or total variation.

6.11. Graph-based Segmentation

Paticipants: Sarah Parisot, Deepak Chittajallu, Ioannis Kakadiaris, Nikos Paragios

In [17] we revisited explicit contour-evolution segmentation methods driven from a graph-based shape prior. Prior knowledge through geometric constraints has been encoded to the model within pair-wise interactions between control points. The segmentation process was driven through an objective function seeking to move the control points towards image locations optimizing the expected visual properties of the organ while satisfying the prior geometric constraints being learned at training. In [18] we have proposed a mathematical formalism for automatic tumor segmentation which was taking advantage of conventional segmentation likelihoods and atlas-based segmentation methods. The central idea was to jointly deform and segment an atlas such that the tumor likelihoods are maximized once projected to the targeted image while relaxing the registration constraints in this area. Furthermore we have endowed to this framework explicit estimation of uncertainties allowing the dynamic sampling of the graph structure resulting on significant speed up of the process while producing quantitative means for the interpretation of the final result.

6.12. Multi-atlas Segmentation

Paticipants: Stavros Alchatzidis, Aristeidis Sotiras, Nikos Paragios

In [33] a novel approach that couples pair-wise deformable registration with multi-atlas segmentation using graphical models was proposed. The method exploits both spaces and seeks to determine the optimal solution which will create the best possible visual agreement between atlases and target image along with their label consistency. The approach optimizes the deformation models and the segmentation labels jointly through an interconnected graph allowing either to relax registration constraints when segmentation labels do indicate or the opposite. The joint optimization of both spaces allowing the “implicit” automatic selection of atlases and therefore improves significantly segmentation performance.

6.13. Higher Order Graph Training throuh Dual Decomposition and Max Margin Principles

Paticipants: Nikos Komodakis, Bo Xiang, Nikos Paragios

In [40] a novel framework based on the structure margin principle was introduced for training higher order graphical models. The idea was to reduce the training of a complex high-order MRF in the parallel training of a series of simple slave MRFs through a principled dual decomposition approach. The theoretical properties of the framework have been studied while the method has been experimentally tested using 2d/3d segmentation problems involving higher order geometric priors that are linear-invariant. The proposed formulation benefits from theoretical guarantees as it concerns performance, computational simplicity while being modular and scalable.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- **Microsoft Research, Cambridge, UK**: Large Scale Diverse Learning for Structured Output Prediction [Ph.D. thesis D. Bouchacourt]
- **General Electrric HealthCare, Buc, FR**: Patient-Specific Optimization of Computed Tomography Acquisition Protocols [Ph.D. thesis H. Pasquier]

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. Excellence Clusters

- 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 date

  Duration: 01/2012-12/2015

  Coordinator: ECP - FR
- Program: DIGITEO (OMTE)

  Project acronym: Curator

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

  Duration: 01/2013-01/2015

  Coordinator: ECP - FR
- Program: DIGITEO
  - Project acronym: SOPRANO
  - Project title: Structured Output Prediction on Large Scale Neuroscience Data
  - Duration: 3/2013-3/2016
  - Coordinator: Ecole Centrale Paris - FR
- Program: MEDICEN

  Project acronym: ADOC

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

  Duration: 11/2011-09/2015

  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: ANR JCJC

  Project acronym: LearnCost

  Project title: Learning Model Constraints for Structured Prediction

  Duration: 2014-2018

  Coordinator: Inria Saclay - 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 & H2020 Projects

8.3.1.1. DIOCLES

Type: FP7

Instrument: European Research Council

Duration: September 2011 - August 2016

Coordinator: Nikos Paragios

Partner: Ecole Centrale de Paris (FR)

Inria contact: Nikos Paragios

8.3.1.2. MOBOT

Type: FP7

Defi: Cognitive Systems and Robotics

Instrument: Specific Targeted Research Project

Objectif: Cognitive Systems and Robotics

Duration: February 2013 - January 2016

Coordinator: Angelika Peer

Partner: University of Bristol (UK)

Inria contact: Iasonas Kokkinos

8.3.1.3. I-SUPPORT

Type: H2020

Defi: Cognitive Systems and Robotics

Instrument: Specific Targeted Research Project

Objectif: Cognitive Systems and Robotics

Duration: March 2015 - February 2018

Coordinator: Rafa Lopez

Partner: Robotnik Automation (Spain)

Inria contact: Iasonas Kokkinos

8.3.1.4. RECONFIG

Type: FP7
Defi: Cognitive Systems and Robotics
Instrument: Specific Targeted Research Project
Objectif: Cognitive Systems and Robotics
Duration: February 2013 - January 2016
Coordinator: Dimos Dimarogonas
Partner: KTH (SE)
Inria contact: Iasonas Kokkinos

8.3.1.5. Strategie

Type: FP7
Instrument: Career Integration Grant
Duration: January 2014 - December 2017
Coordinator: Inria
Inria contact: Matthew Blaschko

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
International Partner (Institution - Laboratory - Researcher):
Stanford University (ÉTATS-UNIS)
Duration: 2012 - 2014
See also: http://cvn.ecp.fr/personnel/pawan/research/splendid.html

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 Engineering Science. Collaboration Topic: Structured prediction 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
- University of Colorado, Department of Computer Science (USA) - Research with the Autonomous Robotics & Perception Group. Collaboration topic: Large scale video segmentation using efficient approximations to a graph Laplacian.

Asia

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

8.5. International Research Visitors

8.5.1. Visits of International Scientists

- Professor Maragos, Petros: Technical University of Athens, GR (October 2014)

8.5.1.1. Internships

- Gastounioti, Aimilia: Technical University of Athens, GR (from February until June 2014)
- Trulls, Eduard: Universitat Politècnica de Catalunya, ES (from June until October 2014)
- Vedantam, Shanmukha Ramakrishna: Virginia Tech, USA (from June 2014 until August 2014)]

8.5.2. Visits to International Teams

- Ferrante, Enzo: Stanford University, USA (from June to September 2014)

8.5.2.1. Research stays abroad

- Boussaid, Haithem: University of Pennsylvania, USA (from June to September 2014)
- Togkas, Stavros: Oxford University, UK (from August to November 2014)

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific events organisation

9.1.1.1. Member of the organizing committee

- Blaschko, Matthew: Co-Organizer of Learning and inference in discrete graphical models tutorial, in conjunction with IEEE Computer Vision and Pattern Recognition (CVPR).

- Kokkinos, Iasonas: Co-Organizer of BASes for Images and Surfaces (BASIS) tutorial, in conjunction with IEEE Computer Vision and Pattern Recognition (CVPR).
- Paragios, Nikos: (i) Co-Organizer of Bayesian and grAphical Models for Biomedical Imaging (BAMBI) workshop, in conjunction with the Medical Image Computing and Computer Assisted Intervention (MICCAI), (ii) Co-Organizer of the Learning and inference in discrete graphical models tutorial, in conjunction with IEEE Computer Vision and Pattern Recognition (CVPR).

9.1.2. Scientific events selection

9.1.2.1. Member of the conference program committee

- Blaschko, Matthew: Neural Information Processing Systems (NIPS), British Machine Vision Conference (BMVC), Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP).
- Kumar, Pawan: Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP).
- Paragios, Nikos: IEEE Computer Vision and Pattern Recognition (CVPR), Medical Image Computing and Computer Assisted Intervention (MICCAI).

9.1.2.2. Reviewer

- Argyriou, Andreas: Neural Information Processing Systems (NIPS).
- Blaschko, Matthew: Artificial Intelligence and Statistics (AISTATS), Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR), IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- Kokkinos, Iasonas: European Conference on Computer Vision (ECCV), IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Neural Information Processing Systems (NIPS), Artificial Intelligence and Statistics (AISTATS), Asian Conference on Computer Vision (ACCV).
- Kumar, Pawan: European Conference on Computer Vision (ECCV), Advances in Neural Information Processing Systems (NIPS).
- Paragios, Nikos: European Conference on Computer Vision (ECCV).

9.1.3. Journal

9.1.3.1. Editor-in-Chief

- Paragios, Nikos: Computer Vision and Image Understanding Journal (CVIU).

9.1.3.2. Member of the editorial board

- Kumar, Pawan: Computer Vision and Image Understanding (CVIU).
- Kokkinos, Iasonas: Image and Vision Computing Journal (IVC), Guest Editor Special Issue on Generative Models in Computer Vision - Computer Vision and Image Understanding Journal (CVIU).
- Paragios, Nikos: Medical Image Analysis Journal (MedIA), SIAM Journal on Imaging Sciences, Guest Editor Special Issue on Discrete Graphical Models in Biomedical Image Analysis - Medical Image Analysis Journal (MedIA).

9.1.3.3. Reviewer

- Kokkinos, Iasonas: International Journal of Computer Vision (IJCV), IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI), IEEE Transactions on Image Processing (T-IP), Image and Vision Computing (IVC), Computer Vision and Image Understanding (CVIU).
- Kumar, Pawan: IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI), Computer Vision and Image Understanding (CVIU).
- Paragios, Nikos: International Journal of Computer Vision (IJCV), IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI), IEEE Transactions on Medical Imaging (T-MI).

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Masters

Blaschko, Matthew

- Master: Foundations of Machine Learning, 36, M1, Ecole Centrale Paris, FR
- Master: Structured Prediction, 24, M2, Ecole Centrale Paris, FR

Kokkinos, Iasonas

- Master: Machine Learning for Computer Vision, 24, M2, Ecole Normale Superieure-Cachan, FR
- Master: Introduction to Deep Learning, 24, M2, Ecole Centrale de Paris, FR
- Master: Introduction to Signal Processing, 36, M1, Ecole Centrale de Paris, FR
- Master: Introduction to Computer Vision, 36, M1, Ecole Centrale de Paris, FR

Kumar, Pawan

- Master: Introduction to Discrete Optimization, 12, M2, Ecole Centrale de Paris, FR
- Master: Discrete Optimization and Learning, 12, M2, Ecole Normale Superieure-Cachan, FR

Paragios, Nikos

- Master: Advanced Mathematical Models in Computer Vision, 24, M2, Ecole Normale Superieure-Cachan, FR

E-learning

MOOC: Coursera

Pedagogical resources : Kumar, Pawan & Paragios, Nikos, Discrete Inference and Lerning in Artificial Vision, M2, https://www.coursera.org/course/artificialvision

9.2.2. Supervision

- HdR : Matthew Blaschko, Advances in Empirical Risk Minimization for Image Analysis and Pattern Recognition, École Normale Supérieure de Cachan, 7 novembre 2014
- PhD in progress : Puneet Kumar Dokania, Learning to Rank with Missing and High-Order Information, 2012-2015, M. Pawan Kumar
- PhD in progress : Diane Bouchacourt, Large Scale Diverse Learning for Structured Output Prediction, 2014-2017, M. Pawan Kumar
- PhD in progress: Haithem Boussaid, Efficient Inference and Learning in Graphical Models for Multi-organ Shape Segmentation, 2011-2015, I. Kokkinos
- PhD in progress: Stavros Tsogkas, Learning structured mid-level representations for object recognition, 2011-2015, I. Kokkinos
- PhD in progress: Siddhartha Chandra, Efficient Learning and Optimization for 3D Visual Data, 2013-2016, Iasonas Kokkinos, Pawan Kumar
- PhD in progress: Stefan Kinauer, Surface-based representations for high-level vision tasks, 2013-2016, Iasonas Kokkinos.
- PhD in progress : Wacha Bounliphone, Statistical tools for Imaging-Genetics data integration, 2013-2016, Matthew Blaschko & Arthur Tenenhaus
- PhD in progress : Jiaqian Yu, Structured Prediction Methods for Computer Vision and Medical Imaging, 2014-2017, Matthew Blaschko

- PhD in progress : Eugene Belilovsky, Structured Output Prediction on Large Scale Neuroscience Data, 2014-2017, Matthew Blaschko
- PhD in progress : Stavros Alchatzidis, Message Passing Methods, Parallel Architectures & Visual Processing, 2011-2014, Nikos Paragios
- 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 : Maxim Berman, Learning Higher Order Graphical Models, 2014-2017, Nikos Paragios
- PhD in progress : Hariprasad Kannan, Efficient Inference on Higher Order Graphs, 2014-2017, Nikos Paragios
- PhD in progress : Huu Dien Khue Le, Graph-based Visual Perception : Theories and Applications, 2014-2017, Nikos Paragios

9.2.3. Juries

- **Matthew Blaschko**
  - **PhD Thesis Participation**: K. Gkirtzou - FR (PhD).
  - **Grant Reviewing Services**: European Research Council (ERC).
- **Iasonas Kokkinos**
  - **PhD Thesis Participation**: N. Dimitriou - GR (PhD).
  - **Grant Reviewing Services:** Swiss National Science Foundation.
- **Kumar, Pawan**
  - **PhD Thesis Participation**: K. Park - Australia (PhD), G. Lin - Australia (PhD).
- **Paragios, Nikos**
  - **PhD Thesis Participation**: M. Blascho - FR (PhD), D. Fortun - FR (PhD), A. Gastounioti - GR (PhD), B. Romain - FR (PhD), J. Tang - CA (PhD), J. Weissenberg - CH (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.

9.3. Popularization

- **Blaschko, Matthew**
  - **Presentations**: Third School on Machine Learning and Knowledge Discovery in Databases (BR), Computer Vision and Pattern Recognition Tutorial (US), KU Leuven (BE), Machine Learning Challenge MICCAI Workshop (US), Agence Nationale de la Recherche (FR)
- **Kokkinos, Iasonas**
  - **Presentations**: Imagenet workshop (in conjunction with ECCV, CH), TTI-Chicago (USA), KTH University (SE), Dagstuhl Seminar on Shape Analysis (DE).
- **Kumar, Pawan M.**
  - **Presentations**: University of Oxford (UK), Ecole des Ponts (FR), Swedish AI Society Workshop (SAIS '14, SE), Xerox Research Center Europe (XRCE) (FR).
- **Paragios, Nikos**

– **Presentations**: Reconnaissance de Formes et l'Intelligence Artificielle (RFIA'14, FR), Medical Imaging Summer School (MISS'14, IT), International Conference on Pattern Recogntion (ICPR'15, SE), Algorithmic issues for Inference in Graphical Models (AIGM'14, FR), University of Patras (GR), Swiss Federal Institute of Technology in Zurich (ETHZ) (CH).

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, M. M. BRONSTEIN, R. LITMAN, A. M. BRONSTEIN. *Intrinsic shape context descriptors for deformable shapes*, in "CVPR - IEEE Conference on Computer Vision and Pattern Recognition", Providence, United States, 2012, pp. 159-166, https://hal.inria.fr/hal-00857572

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

[5] I. KOKKINOS. *Bounding Part Scores for Rapid Detection with Deformable Part Models*, in "ECCV Workshops - Parts and Attributes Workshop, European Conference on Computer Vision", Florence, Italy, 2012, pp. 41-50, https://hal.inria.fr/hal-00857563

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

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

[8] 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

[9] 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

[10] 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

[11] 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

[12] 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

[13] M. RAPTIS, I. KOKKINOS, S. SOATTO. *Discovering Discriminative Action Parts from Mid-Level Video Representations*, in "CVPR - IEEE Conference on Computer Vision and Pattern Recognition", Providence, United States, June 2012, https://hal.inria.fr/hal-00918807

[14] 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], https://hal.archives-ouvertes.fr/hal-00855609

[15] S. TSOGKAS, I. KOKKINOS. *Learning-Based Symmetry Detection in Natural Images*, in "ECCV - 12th European Conference on Computer Vision", Florence, Italy, October 2012, pp. 41-54, https://hal.inria.fr/hal-00856535

Publications of the year

Doctoral Dissertations and Habilitation Theses

[16] M. BLASCHKO. *Advances in Empirical Risk Minimization for Image Analysis and Pattern Recognition*, ENS Cachan, November 2014, Habilitation à diriger des recherches, https://tel.archives-ouvertes.fr/tel-01086088

Articles in International Peer-Reviewed Journals

[17] 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", January 2014, pp. 2168-2194 [*DOI :* 10.1109/JBHI.2013.2257820], https://hal.archives-ouvertes.fr/hal-00858880

[18] S. PARISOT, W. WELLS, S. CHEMOUNY, H. DUFFAU, N. PARAGIOS. *Concurrent tumor segmentation and registration with uncertainty-based sparse non-uniform graphs*, in "Medical Image Analysis", May 2014, vol. 18, n^{o} 4, pp. 647 - 659 [*DOI :* 10.1016/J.MEDIA.2014.02.006], https://hal.inria.fr/hal-01109692

International Conferences with Proceedings

[19] A. BEHL, C. JAWAHAR, M. P. KUMAR. *Optimizing Average Precision using Weakly Supervised Data*, in "CVPR - IEEE Conference on Computer Vision and Pattern Recognition", Columbus, Ohio, United States, 2014, https://hal.inria.fr/hal-00984699

[20] H. BOUSSAID, I. KOKKINOS. *Fast and Exact: ADMM-Based Discriminative Shape Segmentation with Loopy Part Models*, in "IEEE Conference in Computer Vision and Pattern Recognition", Greater Columbus Convention Center in Columbus, Ohio., United States, June 2014, pp. 4058 - 4065 [*DOI :* 10.1109/CVPR.2014.517], https://hal.archives-ouvertes.fr/hal-01109287

[21] H. BOUSSAID, I. KOKKINOS, N. PARAGIOS. *Discriminative learning of deformable contour models*, in "Int.l Symposium on Biomedical Imaging (ISBI)", Bejing, China, April 2014, pp. 624 - 628 [*DOI :* 10.1109/ISBI.2014.6867948], https://hal.inria.fr/hal-01108276

[22] M. CIMPOI, S. MAJI, I. KOKKINOS, S. MOHAMED, A. VEDALDI. *Describing Textures in the Wild*, in "IEEE Conference in Computer Vision and Pattern Recognition", Greater Columbus Convention Center in Columbus, Ohio., United States, June 2014, pp. 3606 - 3613 [*DOI :* 10.1109/CVPR.2014.461], https://hal.archives-ouvertes.fr/hal-01109284

[23] P. K. DOKANIA, A. BEHL, C. V. JAWAHAR, M. PAWAN KUMAR. *Learning to Rank Using High-Order Information*, in "ECCV 2014 - European Conference on Computer Vision (2014)", Zurch, Switzerland, September 2014, pp. 609 - 623 [*DOI :* 10.1007/978-3-319-10593-2_40], https://hal.inria.fr/hal-01076220

[24] M. GHAFARIANZADEH, M. BLASCHKO, G. SIBLEY. *Unsupervised Spatio-Temporal Segmentation with Sparse Spectral Clustering*, in "British Machine Vision Conference (BMVC)", Nottingham, United Kingdom, September 2014, https://hal.inria.fr/hal-01034903

[25] M. P. KUMAR. *Rounding-based Moves for Metric Labeling*, in "NIPS - Advances in Neural Information Processing Systems", Montreal, Canada, 2014, https://hal.inria.fr/hal-01069910

[26] *Best Paper*
M. MISYRLIS, A. KONOVA, M. BLASCHKO, J. HONORIO, N. ALIA-KLEIN, R. GOLDSTEIN, D. SAMARAS. *Predicting cross-task behavioral variables from fMRI data using the k-support norm*, in "Sparsity Techniques in Medical Imaging (STMI)", Boston, United States, September 2014, https://hal.inria.fr/hal-01026303.

[27] P. MOHAPATRA, C. V. JAWAHAR, M. P. KUMAR. *Efficient Optimization for Average Precision SVM*, in "NIPS - Advances in Neural Information Processing Systems", Montreal, Canada, 2014, https://hal.inria.fr/hal-01069917

[28] J. I. ORLANDO, M. BLASCHKO. *Learning fully-connected CRFs for blood vessel segmentation in retinal images*, in "Medical Image Computing and Computer Assisted Intervention (MICCAI)", Boston, United States, Springer, September 2014, https://hal.inria.fr/hal-01024226

[29] N. PARAGIOS, N. KOMODAKIS. *Discrete Visual Perception*, in "International Conference on Pattern Recognition", Stockholm, Sweden, IEEE, August 2014, pp. 1-7, https://hal.archives-ouvertes.fr/hal-01023406

[30] E. TRULLS, S. TSOGKAS, I. KOKKINOS, A. SANFELIU, F. MORENO-NOGUER. *Segmentation-Aware Deformable Part Models*, in "IEEE Conference in Computer Vision and Pattern Recognition", Greater Columbus Convention Center in Columbus, Ohio., United States, June 2014, pp. 168 - 175 [*DOI :* 10.1109/CVPR.2014.29], https://hal.archives-ouvertes.fr/hal-01109286

[31] A. VEDALDI, S. MAHENDRAN, S. TSOGKAS, S. MAJI, R. B. GIRSHICK, J. KANNALA, E. RAHTU, I. KOKKINOS, M. B. BLASCHKO, D. WEISS, B. TASKAR, K. SIMONYAN, N. SAPHRA, S. MOHAMED. *Understanding Objects in Detail with Fine-grained Attributes*, in "IEEE Conference on Computer Vision and Pattern Recognition", Columbus, United States, June 2014, https://hal.inria.fr/hal-00981125

National Conferences with Proceedings

[32] M. BLASCHKO, A. MITTAL, E. RAHTU. *An O(n log n) Cutting Plane Algorithm for Structured Output Ranking*, in "German Conference on Pattern Recognition", Münster, Germany, Springer, September 2014, https://hal.inria.fr/hal-01020943

Conferences without Proceedings

[33] S. ALCHATZIDIS, A. SOTIRAS, N. PARAGIOS. *Discrete Multi Atlas Segmentation using Agreement Constraints*, in "British Machine Vision Conference", Nottingham, United Kingdom, British Machine Vision Association, September 2014, https://hal.inria.fr/hal-01061457

[34] E. BELILOVSKY, A. ARGYRIOU, M. BLASCHKO. *Approximating Combined Discrete Total Variation and Correlated Sparsity With Convex Relaxations*, in "NIPS Workshop on Discrete and Combinatorial Problems in Machine Learning", Montreal, Canada, December 2014, https://hal.inria.fr/hal-01104411

[35] W. BOUNLIPHONE, A. GRETTON, M. BLASCHKO. *Kernel non-parametric tests of relative dependency*, in "NIPS Workshop on Modern Nonparametrics 3: Automating the Learning Pipeline", Montreal, Canada, December 2014, https://hal.inria.fr/hal-01104407

[36] P.-A. SAVALLE, S. TSOGKAS, G. PAPANDREOU, I. KOKKINOS. *Deformable Part Models with CNN Features*, in "European Conference on Computer Vision, Parts and Attributes Workshop", Zurich, Switzerland, September 2014, https://hal.archives-ouvertes.fr/hal-01109290

[37] J. YU, M. BLASCHKO. *Lovasz Hinge for Learning Submodular Losses*, in "NIPS Workshop on Representation and Learning Methods for Complex Outputs", Montreal, Canada, December 2014, https://hal.inria.fr/hal-01104413

Books or Proceedings Editing

[38] N. PARAGIOS, J. DUNCAN, N. AYACHE (editors). *Handbook of Biomedical Imaging: Methodologies and Clinical Research*, Springer, 2014, 590 p. , https://hal.inria.fr/inria-00616178

Research Reports

[39] R. GADDE, R. MARLET, P. NIKOS. *Learning grammars for architecture-specific facade parsing*, September 2014, n^{o} RR-8600, https://hal.inria.fr/hal-01069379

[40] N. KOMODAKIS, B. XIANG, N. PARAGIOS. *A Framework for Efficient Structured Max-Margin Learning of High-Order MRF Models*, Ecole de Ponts-ParisTech ; Ecole Centrale de Paris ; Inria Saclay Ile de France, 2014, n^{o} RR-8645, pp. 1 - 34 [*DOI :* 10.1109/TPAMI.2014.2368990], https://hal.inria.fr/hal-01090971

[41] G. PAPANDREOU, I. KOKKINOS, P.-A. SAVALLE. *Untangling Local and Global Deformations in Deep Convolutional Networks for Image Classification and Sliding Window Detection*, Toyota Technological Institute at Chicago ; Ecole Centrale Paris ; Inria Saclay Ile de France, November 2014, https://hal.archives-ouvertes.fr/hal-01109289

[42] Y. ZENG, C. WANG, D. SAMARAS, X. GU, N. PARAGIOS. *Higher-order Graph Principles towards Non-rigid Surface Registration*, Inria Saclay Ile de France, September 2014, n^{o} RR-8607, 31 p. , https://hal.inria.fr/hal-01086052

Scientific Popularization

[43] M. BLASCHKO. *Machine Learning for Neurological Disorders*, in "Centraliens", March 2014, n^{o} 632, pp. 40-42, https://hal.inria.fr/hal-00940262

Patents and standards

[44] P. NIKOS, E. FERRANTE, R. MARINI SILVA. *Method and device for elastic registration between a two-dimensional digital image and a slice of a three-dimensional volume with overlapping content*, July 2014, n^o^ US 2014/0192046 A1, 12 p. , https://hal.inria.fr/hal-01044647

Other Publications

[45] W. BOUNLIPHONE, A. GRETTON, M. BLASCHKO. *A low variance consistent test of relative dependency*, July 2014, https://hal.inria.fr/hal-01005828