HAL open science

Discrete Visual Perception

Nikos Paragios, Nikos Komodakis

To cite this version:

Nikos Paragios, Nikos Komodakis. Discrete Visual Perception. International Conference on Pattern Recognition, Aug 2014, Stockholm, Sweden. pp.1-7, 10.1109/icpr.2014.13 . hal-01023406

HAL Id: hal-01023406

https://hal.science/hal-01023406v1

Submitted on 16 Jul 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Discrete Visual Perception

Nikos Paragios
Center for Visual Computing
Ecole Centrale de Paris, Chatenay-Malabry, France 92295
Email: nikos.paragios@ecp.fr

Nikos Komodakis
LIGM (Gaspard Monge Computer Science Laboratory)
Ecole de Ponts-ParisTech, Champs-sur-Marne, France 77455
Email: nikos.komodakis@enpc.fr

***Abstract*—Computational vision and biomedical image have made tremendous progress of the past decade. This is mostly due the development of efficient learning and inference algorithms which allow better, faster and richer modeling of visual perception tasks. Graph-based representations are among the most prominent tools to address such perception through the casting of perception as a graph optimization problem. In this paper, we briefly introduce the interest of such representations, discuss their strength and limitations and present their application to address a variety of problems in computer vision and biomedical image analysis.**

I. Introduction

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.

To address the above challenges, one can resort to discrete optimization formulations, which apply to a wide variety of tasks in medical image analysis. 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 (i.e., 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:

$$\arg\min_{\{x_p\}} \mathcal{P}(\{x_p\}; g, f) = \sum_{p\in\mathcal{V}} g_p(x_p) + \sum_{(p_1,p_2)\in\mathcal{E}} f_{p_1p_2}(x_{p_1}, x_{p_2}).$$

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 $\{C_i\}_{i\in[1,n]}$ with $C_i = \{p_{i^1}, \cdots, p_{i|C_i|}\}$

$$\arg\min_{\{x_p\}} \mathcal{P}(\{x_p\}; g, f) = \sum_{p\in\mathcal{V}} g_p(x_p) + \sum_{(p_1,p_2)\in\mathcal{E}} f_{p_1p_2}(x_{p_1}, x_{p_2}) + \sum_{i=1}^{n} f_{C_i}(\mathbf{x}_{C_i}) \ ,$$

here $\mathbf{x}_{C_i} = \{x_p | p \in C_i\}$ and f_{C_i} is the price to pay for associating the labels $\mathbf{x}_{C_i}$ to the nodes in 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 convex relaxations. 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.

The reminder of this paper reviews briefly our contributions in the field. Section 2 presents the work done in the area of inference algorithms while section 3 is dedicated to computer vision applications of such methods. Biomedical image analysis is discussed in section 4, while the last section concludes the paper and presents perspectives, future directions of our work.

II. Discrete Optimization of Pair-wise & Higher-order MRFs

Discrete optimization has been an research topic in the vision community for the past decade [1].

In [2] we have introduced a novel linear-programming for efficient optimization of arbitrary pair-wise MRFs. The central idea was instead of solving the original problem to cast it as a linear programming one (using an LP relaxation) and then seek the solution that minimizes the primal-dual gap. Such a method exhibits computational efficiency and can handle arbitrary graphs and interactions between their elements (both in terms of connectivity as well as in terms of potentials)[1].

The quality of MRF-MAP inference from LP-optimization methods heavily depends on the associated relaxation. In [3] we have introduced a novel dynamic optimization method that updates the LP relaxation on the fly towards increasing its tightness and potentially achieving a globally optimal solution upon convergence. To this end, pair-wise terms of the graphical model become also inference variables under global constraints driven from the original objective function. The proposed formulation has led to state of the art results.

Despite the enormous speed up over conventional methods, the aforementioned solution is sequential and therefore couldn't benefit from recent progress on parallel architectures. In [4], a novel efficient method was introduced for exchanging messages in the context of belief propagation through the jump flooding algorithm, an approach that was able to handle a significant number of labels while outperforming in terms of speed - assuming a basic GPU implementation - for at least one order of magnitude all existing optimization methods.

In [5] we have re-visited belief propagation through dual decomposition, an approach that provides a common theoretical framework for message passing while at the same time introduces a novel, scalable, modular and flexible framework for the minimization of graphical models. Such a method allows the decomposition of your problem into a number of sub-problems over shared variables that can be solved efficiently, and a principled way of coordinating the sub-problem solutions is employed such that at the end all of them converge to a unique common optimal solution. Interesting results in the context of blind deconvolution were obtained using such a framework [6]. This framework has been further extended to deal with optimization of higher order MRFs [7] where efficient decomposition methods were proposed leading to state of the art results.

[1]FastPD is an optimization platform in C++ for the computer vision and medical imaging community (publicly available at http://cvn.ecp.fr) developed 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.

III. Image Processing & Computer Vision

A. Graph Matching & Image Representations

The problem of graph matching [8] was initially studied in [9]. A Higher Order MRF formulation that is invariant to linear geometric transformations was proposed, which was seeking hierarchically a sparse to dense matching through a many-to-many formulation [9]. Such a formulation was extended in [10] to track evolving/deforming surfaces and was endowed with a powerful local shape similarity metric that was the Mobius transform. This idea was further extended in [11] towards handling non-rigid shapes through a discrete minimum distortion principle that was modular under various classes of deformations. More recently, such a higher order model was seamlessly integrated with a generic deformation model, called Canonical Distortion Coefficients (CDCs), that characterizes the deformation of every point on a surface using the distortions along its two principle directions leading to an efficient inference paradigm that can easily handle different types of geometry/texture similarity metrics. A different approach exploring shape maps [12] was studied in [13] for the problem of multiple graphs being registered together towards understanding the motion dynamics of the different object parts.

Contextual representation of images using either local representations [14] or graphs [15] is a well studied problem in the field. In [16] we introduced a new image representation encoding the general layout of groups of quantized local invariant descriptors as well as their relative frequency. This is achieved by embedding the image into a graph where nodes correspond to interest point clusters and links to commute times between the different clusters, and subsequently mapping this graph to lower dimension by exploiting its spectral properties. The proposed representation can be used for indexing [17] and classification using simple nearest neighbor methods [18] or complex classifiers. In [19] a method for image matching was proposed that exploits hierarchical image representations through higher order graphs. The matching was achieved through a graph-based theoretical framework where the similarity and spatial consistency of the image semantic objects is encoded in a graph of commute times that is also endowed with singleton terms through shape descriptors. These ideas were further explored in the context of co-segmentation [20] where a method to determine a consistent partition of multiple images was introduced through a multi-scale multiple-image generative model based on region matching that exploits inter-image information and establishes correspondences between the common objects that appear in the scene.

B. Image-driven Procedural Modeling of Architectures

Given the deluge of satellite imagery and urban photo collections, an emerging application is computer vision techniques that can produce scalable representations of 3D models of urban environments. Our team has pursued the development of algorithms for inverse procedural modeling, namely the inference of a set of architectural rules that can explain observed building images. Initially, we have investigated the use of graphical models [21] towards inferring the optimal grammar instance and the associated parameters through local self-similarities of the graphical model (structure repetition). Complexity was the main limitation of the method since the optimization problem involves discrete and continuous variables with an unknown number in advances.

C. Pose Estimation, Articulated Models & Illumination/Reflectance

Inferring 3D geometry from 2D images in particular in the context of articulated objects is a highly interesting problem in computer vision. However at the same time it remains a great challenge due to the dimensionality of the model, and the ill-posedeness.

In [22] we have introduced a novel approach to 3D pose estimation of deformable objects from 2D images that simultaneously recovers the optimal model instance and the corresponding image projections. The main innovation of the method was the implicit handling of the camera pose parameters thanks to the use of a higher order graph where the order of cliques (depending on the projection model) allowed to determine a cost that could simultaneously cope with the unknown model deformation, projection matrix and image correspondences. A different approach was used in the context of articulated models (hand-pose estimation) where we first studied part-based representations [23] in which case geometric elements were used to determine the global shape of the object. The inference was then performed by seeking the optimal model projection to the image through a segmentation function that was aiming to separate skin-like properties from the background through an energy that was minimized using the divergence theorem. The simplicity of the representation made the approach sensitive to the initial conditions, as well as to occlusions and self-occlusions.

Understanding the position of light sources is also critical to a number of applications like hand pose estimation. In [24] a graphical model to estimate the illumination environment and detect the shadows of a scene with textured surfaces from a single image and only coarse 3D information was proposed. The central idea was to represent pixel-wise shadow observations as a mixture model of contributions coming from different illumination sources and then alternate between shadows detection and illumination sources parameter estimation. The same concept was further studied in [25], [26] where a higher-order Markov Random Field (MRF) illumination model was employed, which combines low-level shadow evidence with high-level prior knowledge for the joint estimation of cast shadows and the illumination environment. In such a context, a higher order approach is considered where illumination sources are coupled with the observed image and the latent variables corresponding to the shadow detection. Last, but not least approximate geometry estimation of the scene was also incorporated into this framework in [27].

IV. Biomedical Image Analysis

A. Medical Model-free and Model-based Segmentation

In [28] the problem of introducing sparse prior knowledge from small training sets was studied in the context of brain tumor segmentations. The idea was to decompose the population of training examples into small sets with locally consistent behavior, and then build individual statistics for each population locally. This was achieved through a graphical model formulation where the prior model decomposition was expressed as an unsupervised clustering problem involving the number of populations, their most representative elements and the assignments of the training set. Given such a prior model, then inference was expressed through a coupled formulation seeking the optimal class for the prior model and assigning labels to the image [29] according to it (given the class likelihoods [30]). Such an approach was tested on the segmentation of brain tumors (Glioma II). In [31], [32], [33], [34] grouping was used for the segmentation of curvilinear structures in low signal to noise ratio images through an ordering of these segments using the elastica model in a linear programming framework. The problem was also considered in its temporal aspect [35], [36], [37] through a combined iconic-geometric approach [37] seeking to determine the optimal deformation along with correspondences for a subset of interest points along this structure. In the context of multi-label segmentation [38] studied two distinct cases of introducing prior knowledge to the random walker graphical model. The first was aiming to position automatically the seeds with respect to the different classes [39] through the optimization of a graph-based objective function. Later the idea of introducing prior knowledge [38], [40] in this context was studied using statistical models of varying complexity, like mono-modal distributions or linear sub-spaces per class. Multi-class segmentation of striated Muscles in NMR Images was the clinical case being investigated in this context where the problem of optimal parameter setting was also investigated through machine learning methods [41]. Continuous methods have certain strengths but also exhibit a number of limitations in particular during the inference process like for example their strong dependency from the initial conditions. This was an issue that was addressed in [42], [43], [15], [44]. The central idea was to represent prior knowledge through a point distribution representation mapped to a pair-wise probabilistic graphical model. In such a context global shape priors could be built through the concatenation of local constraints resulting on a powerful and flexible model both in terms of learning as well as in terms of inference. The proposed formulation was translation and rotation invariant and was endowed with boundary and regional support (through approximation of the region probabilities based on a Voronoi decomposition of the image domain[42]). The method was able to handle missing correspondences and was tested in the context of lungs segmentation since in that case one has to handle an important number of miss-detections [44].

This idea has been further expanded in [45], [46] where invariance to linear transformations has been introduced through the use of higher order interactions between the nodes of the probabilistic graphical model. This approach was extended to handle the case of 3D objects projected to multiple 2D images in [47], [45], [46]. In order to determine the exact region probabilities, [48], [49] re-expressed the exact integral using third order cliques through the use of Stokes/divergence theorem resulting on a novel local shape model that is invariant to global transformations while being able to determine exactly the associated image support. The model complexity was reduced through sparsity constraints [50]. Last but not least the idea of decoupling image and model-based segmentation was also considered in [51], where the latent variables correspond both to the model deformation towards the image as well as to the individual labels of pixels (that are jointly optimized). Cardiac segmentation both in 2D (MR Tagged images) as well as in 3D (Computed Tomography) of the myocardium was the clinical example considered to demonstrate the interest of such a method. Introducing more complex geometric properties in the context of articulations [52], [53] was a further development in the same direction targeting spine/vertebra segmentation. The idea was to represent the spine through a probabilistic graphical model endowed with higher order interactions encoding the relative geometric positions of successive vertebras. In order to cope with the gap between the dimensionality of the considered representation and the number of examples within the training set, non-linear manifold embedding methods were also introduced [54], [55]. An alternative to point distribution models that has gained recently

important attention refers to atlas-based methods [56]. The idea is to deform an annotated atlas to a new image and then use the deformation to determine the anatomical structures in the new image. Spine segmentation was investigated in [57] while brain segmentation in the presence of tumors [58], [59] was considered as well where the aim was to use a healthy anatomical atlas as a prior and relax the constraint of image-to-atlas registration in the tumorous area. This concept was explored in [60] through a population registration framework [61] seeking to simultaneously deform n-images towards the target image and determine the optimal segmentation map in the new image through the combination of the deformed atlases labels.

B. Deformable Registration & Longitudinal Modeling

Global or linear registration is a preliminary step of almost all medical image analysis tools either for longitudinal modeling or for population studies. In [62] we introduced the first method able to guarantee the globally optimal solution in the mono-modal case. This was achieved through a re-expression of the objective function as a difference of convex functions and the use of cutting planes algorithms to reach the optimal solution. Such a re-expression was based on a dictionary driven representation using a pre-defined base where both images were expressed as a linear combination of the dictionary elements. In order to deal with the multi-modal case, more recently we have investigated the use of graphical models [63] and discrete optimization on the parameters acting directly on the linear transformation [64] that was also considered for 3D-to-2D registration in [65]. Concerning deformable registration [66], which is the pillar of medical image analysis, in [67][2] we have introduced a low graphical model formulation coupled with efficient linear programming to efficiently and precisely address mono and multi-modal iconic registration [68]. The proposed approach was endowed with modularity, scalability and computational efficiency.

Furthermore, inspired from [67] the problem of co-registration of multiple volumes exhibiting partial overlap was addressed using graphical models leading to an efficient deformable stitching algorithm for whole body acquisitions [69]. Improving computational efficiency and bringing the method closer to real clinical settings was investigated in [4] where a novel fully parallelizable approach for the optimization of graphical models in graphics processing units was adopted leading to two orders of magnitude acceleration.

In [70], [71] this framework has been extended to cope with multi-dimensional feature spaces while at the same time the idea of mutual saliency (on invariant Gabor features) [70] was introduced that changes fundamentally the way image similarity costs are taken into account in a qualitative manner. Such a novel qualitative metric is able to handle missing correspondences and drive the registration to the optimal solution even from a very bad initial configuration. In the same direction but through a different concept that relies on objective function residuals with respect to the different labels in [72] we have introduced the notion of uncertainties of the registration solution through the approximation of the solution covariance matrix using the min-marginal energies. This has led to an adaptive discrete sampling of the label space improving substantially the quality of the obtained results (due to the efficient handling of the search space) while providing novel means to appropriately quantify/qualify the physical solution.

Given that deformable registration is an ill-posed problem, it is often handled through the convexification of the objective function using the regularization term. An alternative, clinical context specific approach was introduced [73] where we have proposed a novel learning method that was invariant to linear transformations to capture the relative displacements of organs with a deformable registration framework and consider them as priors subsequently.

Furthermore, integration of anatomical landmarks coupled with iconic registration was another step forward. An inter-connected layer approach was considered in [74], [75] where we are seeking simultaneously to determine the optimal iconic mapping and the optimal landmark association while imposing consistency with respect to the associated transformations. The parameters of the two graphical models were simultaneously inferred providing dense deformation along with landmarks correspondences. Symmetry was also a property that was naturally imposed to the formulation through a deformation from a common space to both modalities with inverse consistency [76]. Furthermore, the same idea was explored in the context of genetic imaging data either for registration [77] or for a model-based segmentation/partition [78], [79].

The problem of image to volume registration, a more challenging component of fusion has been recently investigated [80]. The objective is to determine an optimal plane and within this plane the associated deformations that will map an intra-operative low resolution image (like for example radiography or ultrasound) to a high resolution anatomical data. Examples of use of such an approach refer to image-based surgery guidance and adaptive radiotherapy. The problem was formulated as an over-parameterized pair-wise graphical model in [80] where we seek both the optimal plane and the associated deformation model. Over-parameterization was adopted to impose unique plane consistency. An alternative higher order approach was recently adopted that implicitly imposes unique plane consistency and determines the optimal deformation through third and fourth order potentials estimated on the 3D deformable vectors.

V. Conclusion

Despite their tremendous use, graphical models have been mostly employed to handle low rank interactions that suffer from the ability to cope with rich models as well as encode better priors. Furthermore, the dimensionality of the label set space was a major issue.

In order to deal with a large number of labels, designing hierarchical divide-and-conquer style algorithms is a promising direction through multi-grid methods [81]. While divide-and-conquer will increase the number of problems that will be required to be solved, each problem will be defined on a small set of similar labels, corresponding to a cluster within a hierarchical clustering of labels. In order to allow for general high-order potentials, we will design iterative algorithms that approximate arbitrary potentials with their sparse, non-linear envelope alternatives. The approximation will change dynamically over the iterations to ensure accuracy. Specifically, the parameters of the non-linear envelope will be determined by both the original clique potential as well as the current estimate of the solution. Finally, in order to handle a large number of variables, we will build on the dual decomposition framework which iteratively solves subproblems defined on a subgraph of

[2]DROP is a deformable registration platform in C++ for the medical imaging community (publicly available at http://www.mrf-registration.net) developed at Ecole Centrale de Paris, 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 given model. The subproblem definitions will be changed dynamically in order to encourage faster convergence.

Learning task-specific graphical models is another promising direction. In the early years of graphical models, limited effort was invested on determining the optimal combination of such terms as they were mainly restricted to pair-wise graphical models of very limited complexity involving a small number of global parameters. Learning the optimal parameters of a model given a set of annotated data can be expressed as a problem of modeling dependencies between variables that minimize a loss function between the model prediction and the annotated solutions that in general is a non convex problem. Solving the convex relaxation of the expected loss through algorithms like the cutting plane, was a pioneering step forward to address such a learning problem. In [82] a framework for learning the optimal graph structure for clustering problems was introduced while in [83] a generic framework for learning the potentials of pair-wise as well as high-order graphical models is presented combining the dual decomposition and the max-margin concept resulting on a fully distributed approach to parameter learning composed of subproblems of limited complexity.

Bridging the gap between data-driven & physics-based modeling is a promising step forward in artificial vision. Modeling in biomedical imaging has been either physics-based driven or data-driven. The former case suffers from computational complexity and efficiency limitations, the later from interpretability while being also heavily dependent on the design of the objective function. The use of annotated data along with the progress made in the field of dimensionality reduction and in the optimization of higher order graphical models do make possible the training and use of large scale models. Consequently, reducing the gap between physical and data-driven modeling could be of great interest and could have tremendous impact on the efficient and automatic interpretation of the data in the future.

Acknowledgment

The authors would like to thank all contributors of the papers being cited in the reference section.

References

[1] C. Wang, N. Komodakis, and N. Paragios, "Markov Random Field Modeling, Inference & Learning in Computer Vision & Image Understanding: A Survey," *Computer Vision and Image Understanding*, vol. 117, no. 11, pp. 1610–1627, 2013. [Online]. Available: http://hal.inria.fr/hal-00858390

[2] N. Komodakis, G. Tziritas, and N. Paragios, "Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies," *Computer Vision and Image Understanding*, pp. 14–29, 2008. [Online]. Available: http://hal.inria.fr/hal-00918699

[3] N. Komodakis and N. Paragios, "Beyond Loose LP-Relaxations: Optimizing MRFs by Repairing Cycles," in *ECCV*, France, 2008, pp. 806–820. [Online]. Available: http://hal.inria.fr/hal-00918715

[4] S. Alchatzidis, A. Sotiras, and N. Paragios, "Efficient parallel message computation for MAP inference," in *Computer Vision (ICCV), 2011 IEEE International Conference on*, Spain, 2011, pp. 1379 – 1386. [Online]. Available: http://hal.inria.fr/hal-00858462

[5] N. Komodakis, N. Paragios, and G. Tziritas, "MRF Energy Minimization and Beyond via Dual Decomposition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 3, pp. 531–552, 2011. [Online]. Available: http://hal.inria.fr/hal-00856311

[6] N. Komodakis and N. Paragios, "MRF-Based blind image deconvolution," in *11th Asian Conference on Computer Vision - ACCV 2012*, vol. 7726. Daejeon, Korea, Democratic People'S Republic Of: Springer, 2012, pp. 361–374. [Online]. Available: http://hal.inria.fr/hal-00856285

[7] ——, "Beyond pairwise energies: Efficient optimization for higher-order mrfs," in *IEEE International Conference in Computer Vision (CVPR)*, 2009, pp. 2985–2992.

[8] C. Wang, Y. Zeng, D. Samaras, and N. Paragios, "Modeling Shapes with Higher-Order Graphs: Methodology and Applications," in *Shape Perception in Human and Computer Vision: An Interdisciplinary Perspective*. Sven J. Dickinson and Zygmunt Pizlo, 2013, pp. 459–471. [Online]. Available: http://hal.inria.fr/hal-00858417

[9] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios, "Dense non-rigid surface registration using high-order graph matching," in *23rd IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2010*, San Francisco, United States, 2010, pp. 382–389. [Online]. Available: http://hal.inria.fr/hal-00856064

[10] ——, "Intrinsic dense 3D surface tracking," in *24th IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2011*, Colorado Springs, United States, 2011, pp. 1225–1232. [Online]. Available: http://hal.inria.fr/hal-00856137

[11] Y. Zeng, C. Wang, D. Gu, D. Samaras, and 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, p. /. [Online]. Available: http://hal.inria.fr/hal-00856323

[12] G. Langs and N. Paragios, "Modeling the structure of multivariate manifolds: Shape maps," in *CVPR*, United States, 2008, pp. 1–8. [Online]. Available: http://hal.inria.fr/hal-00918713

[13] G. Langs, N. Paragios, P. Desgranges, A. Rahmouni, and H. Kobeiter, "Learning deformation and structure simultaneously: in situ endograft deformation analysis." *Medical Image Analysis*, vol. 15, no. 1, pp. 12–21, Feb. 2011. [Online]. Available: http://hal.inria.fr/hal-00856306

[14] A. Znaidia, A. Shabou, H. Le Borgne, C. Hudelot, and N. Paragios, "Bag-of-multimedia-words for image classification," in *Proceedings of the 21st International Conference on Pattern Recognition, ICPR 2012*, Tsukuba, Japan, 2012, pp. 1509–1512. [Online]. Available: http://hal.inria.fr/hal-00825187

[15] Y. Ou, A. Besbes, M. Billelo, M. Mansour, C. Davatzikos, and N. Paragios, "Detecting mutually-salient landmark pairs with MRF regularization," in *2010 IEEE 7th International Symposium on Biomedical Imaging - ISBI 2010*, Rotterdam, Netherlands, 2010, pp. 400–403. [Online]. Available: http://hal.inria.fr/hal-00856059

[16] R. Behmo, N. Paragios, and V. Prinet, "Graph commute times for image representation," in *Computer Vision and Pattern Recognition CVPR*, United States, 2008, pp. 1–8. [Online]. Available: http://hal.inria.fr/hal-00918707

[17] ——, "An Application of Graph Commute Times to Image Indexing," in *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International*, United States, 2008, pp. 495–498. [Online]. Available: http://hal.inria.fr/hal-00918717

[18] R. Behmo, P. Marcombes, A. Dalalyan, and V. Prinet, "Towards optimal naive bayes nearest neighbor," in *European Conference in Computer Vision (ECCV) (4)*, 2010, pp. 171–184.

[19] J. Rubio, J. Serrat, A. López, and N. Paragios, "Image contextual representation and matching through hierarchies and higher order graphs," in *21st International Conference on Pattern Recognition - ICPR 2012*, Tsukuba, Japan, 2012, pp. 2664 – 2667. [Online]. Available: http://hal.inria.fr/hal-00856295

[20] ——, "Unsupervised co-segmentation through region matching," in *25th IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2012*, Rhode Island, United States, 2012, pp. 749–756. [Online]. Available: http://hal.inria.fr/hal-00856291

[21] P. Koutsourakis, L. Simon, O. Teboul, G. Tziritas, and N. Paragios, "Single view reconstruction using shape grammars for urban environments," in *International Conference in Computer Vision (ICCV)*, 2009, pp. 1795–1802.

[22] C. Wang, Y. Zeng, L. Simon, I. Kakadiaris, A., D. Samaras, and N. Paragios, "Viewpoint invariant 3D landmark model inference from monocular 2D images using higher-order priors," in *13th International Conference on Computer Vision - ICCV 2011*, Barcelona, Spain, 2011, pp. 319–326. [Online]. Available: http://hal.inria.fr/hal-00856131

[23] M. De La Gorce and N. Paragios, "A variational approach to monocular hand-pose estimation," *Computer Vision and Image Understanding*, pp. 363–372, 2010. [Online]. Available: http://hal.inria.fr/hal-00856083

[24] A. Panagopoulos, D. Samaras, and N. Paragios, "Robust shadow and illumination estimation using a mixture model," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 651–658.

[25] A. Panagopoulos, C. Wang, D. Samaras, and N. Paragios, "Illumination estimation and cast shadow detection through a higher-order graphical model," in *2011 IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2011*, Colorado Springs, United States, 2011, pp. 673 – 680. [Online]. Available: http://hal.inria.fr/hal-00856316

[26] A. Panagoupolos, C. Wang, D. Samaras, and N. Paragios, "Estimating Shadows with the Bright Channel Cue," in *11th European Conference on Computer Vision - ECCV 2010*, vol. 6554, Heraklion, Greece, 2010, pp. 1–12. [Online]. Available: http://hal.inria.fr/hal-00856060

[27] A. Panagopoulos, C. Wang, D. Samaras, and N. Paragios, "Simultaneous Cast Shadows, Illumination and Geometry Inference Using Hypergraphs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 2, pp. 437–449, Feb. 2013. [Online]. Available: http://hal.inria.fr/hal-00855591

[28] S. Parisot, H. Duffau, S. Chemouny, and N. Paragios, "Graph Based Spatial Position Mapping of Low-grade Gliomas," in *MICCAI - 14th International Conference on Medical Image Computing and Computer-Assisted Intervention - 2011*. Toronto, Canada: Springer Berlin Heidelberg, 2011, pp. 508–515. [Online]. Available: http://hal.inria.fr/hal-00775865

[29] ——, "Graph-based Detection, Segmentation & Characterization of Brain Tumors," in *CVPR - 25th IEEE Conference on Computer Vision and Pattern Recognition 2012*, Providence, United States, 2012. [Online]. Available: http://hal.inria.fr/hal-00712714

[30] H. Boussaid, I. Kokkinos, and N. Paragios, "Rapid Mode Estimation for 3D Brain MRI Tumor Segmentation," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, Lund, Sweden, 2013. [Online]. Available: http://hal.inria.fr/hal-00856770

[31] N. Honnorat, R. Vaillant, and N. Paragios, "Guide-Wire Extraction through Perceptual Organization of Local Segments in Fluoroscopic Images," in *13th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2010*, vol. 6363, Beijing, China, Sep. 2010, pp. 440–448. [Online]. Available: http://hal.inria.fr/hal-00856049

[32] ——, "Robust guidewire segmentation through boosting, clustering and linear programming," in *2010 IEEE 7th International Symposium on Biomedical Imaging - ISBI 2010*, Rotterdam, Netherlands, 2010, pp. 924–927. [Online]. Available: http://hal.inria.fr/hal-00856055

[33] N. Honnorat, R. Vaillant, J. Duncan, and N. Paragios, "Curvilinear structures extraction in cluttered bioimaging data with discrete optimization methods," in *2011 IEEE 8th International Symposium on Biomedical Imaging - ISBI 2011*, Chicago, United States, 2011, pp. 1353–1357. [Online]. Available: http://hal.inria.fr/hal-00856120

[34] N. Honnorat, R. Vaillant, and N. Paragios, "Graph-based guide-wire segmentation through fusion of contrast-enhanced and fluoroscopic images," in *2012 IEEE 9th International Symposium on Biomedical Imaging - ISBI 2012*, Barcelona, Spain, 2012, pp. 948 – 951. [Online]. Available: http://hal.inria.fr/hal-00856299

[35] T. H. Heibel, B. Glocker, N. Paragios, and N. Navab, "Needle tracking through higher-order MRF optimization," in *2010 IEEE 7th International Symposium on Biomedical Imaging - ISBI 2010*, Rotterdam, Netherlands, 2010, pp. 624–627. [Online]. Available: http://hal.inria.fr/hal-00856058

[36] T. Hauke Heibel, B. Glocker, M. Groher, N. Paragios, N. Komodakis, and N. Navab, "Discrete tracking of parametrized curves," in *IEEE International Conference in Computer Vision (CVPR)*, 2009, pp. 1754–1761.

[37] N. Honnorat, R. Vaillant, and N. Paragios, "Graph-Based Geometric-Iconic Guide-Wire Tracking," in *14th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2011*, vol. 6891, Toronto, Canada, 2011, pp. 9–16. [Online]. Available: http://hal.inria.fr/hal-00856106

[38] P.-Y. Baudin, N. Azzabou, P. Carlier, G., and N. Paragios, "Manifold-enhanced Segmentation through Random Walks on Linear Subspace Priors," in *British Machine Vision Conference*, United Kingdom, 2012, pp. 51.1–51.10. [Online]. Available: http://hal.inria.fr/hal-00773635

[39] ——, "Automatic skeletal muscle segmentation through random walks and graph-based seed placement," in *International Symposium Biomedical Imaging (ISBI)*, Barcelone, Spain, 2012, pp. 1036–1039. [Online]. Available: http://hal.inria.fr/hal-00773616

[40] ——, "Prior Knowledge, Random Walks and Human Skeletal Muscle Segmentation," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, France, 2012, pp. 569–576. [Online]. Available: http://hal.inria.fr/hal-00773665

[41] P.-Y. Baudin, D. Goodman, P. Kumar, N. Azzabou, P. G. Carlier, N. Paragios, and 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, Sep. 2013. [Online]. Available: http://hal.inria.fr/hal-00856020

[42] A. Besbes, N. Komodakis, G. Langs, and N. Paragios, "Shape priors and discrete mrfs for knowledge-based segmentation," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 1295–1302.

[43] D. Chittajallu, N. Paragios, and I. Kakadiaris, "An Explicit Shape-constrained MRF-based Contour Evolution Method for 2D Medical Image Segmentation," *IEEE Journal of Biomedical and Health Informatics*, pp. 2168–2194, Apr. 2013. [Online]. Available: http://hal.inria.fr/hal-00858880

[44] A. Besbes and N. Paragios, "Landmark-based segmentation of lungs while handling partial correspondences using sparse graph-based priors," in *2011 IEEE 8th International Symposium on Biomedical Imaging - ISBI 2011*, Chicago, United States, 2011, pp. 989–995. [Online]. Available: http://hal.inria.fr/hal-00856125

[45] C. Wang, O. Teboul, F. Michel, S. Essafi, and N. Paragios, "3D Knowledge-Based Segmentation Using Pose-Invariant Higher-Order Graphs," in *13th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2010*, vol. 6363, Beijing, China, Sep. 2010, pp. 189–196. [Online]. Available: http://hal.inria.fr/hal-00856051

[46] C. Wang, H. Boussaid, L. Simon, J.-Y. Lazennec, and N. Paragios, "Pose-Invariant 3D Proximal Femur Estimation through Bi-planar Image Segmentation with Hierarchical Higher-Order Graph-Based Priors," in *14th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2011*, Toronto, Canada, 2011, pp. 346–353. [Online]. Available: http://hal.inria.fr/hal-00856103

[47] H. Boussaid, S. Kadoury, I. Kokkinos, J.-Y. Lazennec, G. Zheng, and N. Paragios, "3D Model-based Reconstruction of the Proximal Femur from Low-dose Biplanar X-Ray Images," in *The 22nd British Machine Vision Conference - BMVC 2011*, J. oey, S. McKenna, and E. Trucco, Eds., Dundee, United Kingdom, 2011, pp. 1–10. [Online]. Available: http://hal.inria.fr/hal-00856315

[48] B. Xiang, C. Wang, J.-F. Deux, A. Rahmouni, and N. Paragios, "3D Cardiac Segmentation with Pose-Invariant Higher-Order MRFS," in *International Symposium on BIOMEDICAL IMAGING: From Nano to Macro (ISBI)*, Barcelona, Spain, 2012. [Online]. Available: http://hal.inria.fr/hal-00776025

[49] ——, "Tagged cardiac MR image segmentation using boundary & regional-support and graph-based deformable priors," in *2011 IEEE 8th International Symposium on Biomedical Imaging - ISBI 2011*, Chicago, United States, 2011, pp. 1706–1711. [Online]. Available: http://hal.inria.fr/hal-00856116

[50] B. Xiang, N. Komodakis, and N. Paragios, "Pose Invariant Deformable Shape Priors Using L1 Higher Order Sparse Graphs," in *9th International Symposium on Visual Computing - ISVC 2013*, Rethymnon, Greece, Jul. 2013. [Online]. Available: http://hal.inria.fr/hal-00856978

[51] B. Xiang, J.-F. Deux, A. Rahmouni, and 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, Sep. 2013. [Online]. Available: http://hal.inria.fr/hal-00856955

[52] S. Kadoury, H. Labelle, and N. Paragios, "Automatic inference of articulated spine models in CT images using high-order Markov Random Fields," *Medical Image Analysis*, vol. 15, no. 4, pp. 426–437, Aug. 2011. [Online]. Available: http://hal.inria.fr/hal-00856308

[53] S. Kadoury and N. Paragios, "Multimodal inference of articulated spine models from higher order energy functions of discrete MRFS," in *2010 IEEE 7th International Symposium on Biomedical Imaging - ISBI 2010*, Rotterdam, Netherlands, 2010, pp. 1393–1396. [Online]. Available: http://hal.inria.fr/hal-00856052

[54] S. Kadoury, H. Labelle, and N. Paragios, "Spine Segmentation in Medical Images Using Manifold Embeddings and Higher-Order MRFs," *IEEE Transactions on Medical Imaging*, vol. 32, no. 7, pp. 1227–1238, Jul. 2013. [Online]. Available: http://hal.inria.fr/hal-00856319

[55] S. Kadoury and N. Paragios, "Nonlinear Embedding towards Articulated Spine Shape Inference Using Higher-Order MRFs," in *13th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2010*, vol. 6363, Beijing, China, Sep. 2010, pp. 579–586. [Online]. Available: http://hal.inria.fr/hal-00856048

[56] A. Besbes, N. Komodakis, and N. Paragios, "Graph-based knowledge-driven discrete segmentation of the left ventricle," in *IEEE International*

Symposium on Biomedical Imaging: From Nano to Macro (ISBI), 2009, pp. 49–52.

[57] S. Kadoury and N. Paragios, "Surface/volume-based articulated 3d spine inference through markov random fields," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI) (1)*, 2009, pp. 92–99.

[58] S. Parisot, H. Duffau, S. Chemouny, and N. Paragios, "Joint Tumor Segmentation and Dense Deformable Registration of Brain MR Images," in *MICCAI - 15th International Conference on Medical Image Computing and Computer-Assisted Intervention - 2012*, vol. 7511. Nice, France: Springer Berlin Heidelberg, 2012, pp. 651–658. [Online]. Available: http://hal.inria.fr/hal-00773618

[59] S. Parisot, W. Wells Iii, S. Chemouny, H. Duffau, and N. Paragios, "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. [Online]. Available: http://hal.inria.fr/hal-00858696

[60] A. Sotiras, N. Komodakis, G. Langs, and N. Paragios, "Atlas-based deformable mutual population segmentation," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI)*, 2009, pp. 5–8.

[61] A. Sotiras, N. Komodakis, B. Glocker, J.-F. Deux, and N. Paragios, "Graphical models and deformable diffeomorphic population registration using global and local metrics," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI) (1)*, 2009, pp. 672–679.

[62] E. Kokiopoulou, D. Kressner, M. Zervos, and N. Paragios, "Optimal similarity registration of volumetric images," in *24th IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2011*, Colorado Springs, United States, 2011, pp. 2449–2456. [Online]. Available: http://hal.inria.fr/hal-00856133

[63] D. Zikic, B. Glocker, O. Kutter, M. Groher, N. Komodakis, A. Kamen, N. Paragios, and N. Navab, "Linear intensity-based image registration by Markov random fields and discrete optimization," *Medical Image Analysis*, vol. 14, no. 4, pp. 550–562, 2010. [Online]. Available: http://hal.inria.fr/hal-00856077

[64] B. Glocker, D. Zikic, N. Komodakis, N. Paragios, and N. Navab, "Linear image registration through mrf optimization," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI)*, 2009, pp. 422–425.

[65] D. Zikic, B. Glocker, O. Kutter, M. Groher, N. Komodakis, A. Khamene, N. Paragios, and N. Navab, "Markov Random Field Optimization for Intensity-based 2D-3D Registration," in *SPIE 2010, Medical Imaging 2010: Image Processing*, United States, 2010, pp. –. [Online]. Available: http://hal.inria.fr/hal-00858402

[66] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable Medical Image Registration: A Survey," *IEEE Transactions on Medical Imaging*, vol. 32, no. 7, pp. 1153–1190, May 2013. [Online]. Available: http://hal.inria.fr/hal-00858737

[67] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios, "Dense image registration through mrfs and efficient linear programming," *Medical Image Analysis*, vol. 12, no. 6, pp. 731–741, 2008.

[68] B. Glocker, N. Komodakis, N. Paragios, and N. Navab, "Non-rigid Registration using Discrete MRFs: Application to Thoracic CT Images," in *Workshop Evaluation of Methods for Pulmonary Image Registration, MICCAI 2010*, China, 2010, pp. –. [Online]. Available: http://hal.inria.fr/hal-00858382

[69] C. Wachinger, B. Glocker, J. Zeltner, N. Paragios, N. Komodakis, M. Sass Hansen, and N. Navab, "Deformable mosaicing for whole-body mri," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI) (2)*, 2008, pp. 113–121.

[70] Y. Ou, A. Sotiras, N. Paragios, and C. Davatzikos, "DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting." *Medical Image Analysis*, vol. 15, no. 4, pp. 622–639, Aug. 2011. [Online]. Available: http://hal.inria.fr/hal-00856309

[71] A. Sotiras, R. Neji, J.-F. Deux, N. Komodakis, G. Fleury, and N. Paragios, "A Kernel-based graphical model for diffusion tensor registration," in *IEEE ISBI'10*, Rotterdam, Netherlands, 2010, pp. 524–527. [Online]. Available: http://hal.inria.fr/hal-00526570

[72] B. Glocker, N. Paragios, N. Komodakis, G. Tziritas, and N. Navab, "Optical flow estimation with uncertainties through dynamic mrfs," in *IEEE Computer Society Conference on Computer Vision (CVPR)*, 2008.

[73] B. Glocker, N. Komodakis, N. Navab, G. Tziritas, and N. Paragios, "Dense registration with deformation priors," in *Information Processing in Medical Imaging (IPMI)*, 2009, pp. 540–551.

[74] B. Glocker, A. Sotiras, N. Komodakis, and N. Paragios, "Deformable Medical Image Registration: Setting the State of the Art with Discrete Methods," *Annual Review of Biomedical Engineering*, vol. 13, pp. 219–244, 2011. [Online]. Available: http://hal.inria.fr/hal-00858380

[75] A. Sotiras, Y. Ou, B. Glocker, N. Paragios, and C. Davatzikos, "Simultaneous Geometric - Iconic Registration," in *13th International Conference on Medical Image Computing and Computer Assisted Intervention - MICCAI 2010*, vol. 6362, Beijing, China, Sep. 2010, pp. 676–683. [Online]. Available: http://hal.inria.fr/hal-00856046

[76] A. Sotiras and N. Paragios, "Discrete symmetric image registration," in *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, Spain, 2012, pp. 342–345. [Online]. Available: http://hal.inria.fr/hal-00858415

[77] U. Kurkure, Y. Le, H., N. Paragios, J. Carson, T. Ju, and I. Kakadiaris, A., "Landmark/image-based deformable registration of gene expression data," in *24th IEEE Conference on Computer Vision and Pattern Recognition - CVPR 2011*, Rhode Island, United States, 2011, pp. 1089–1096. [Online]. Available: http://hal.inria.fr/hal-00856140

[78] Y. H. Le, U. Kurkure, N. Paragios, T. Ju, J. P. Carson, and I. Kakadiaris, A., "Similarity-based appearance-prior for fitting a subdivision mesh in gene expression images." in *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2012*, vol. 7510. Nice, France: Springer, 2012, pp. 577–84. [Online]. Available: http://hal.inria.fr/hal-00856301

[79] U. Kurkure, Y. Le, N. Paragios, T. Ju, J. Carson, and I. Kakadiaris, A., "Markov Random Field-based fitting of a subdivision-based geometric atlas," in *13th International Conference on Computer Vision - ICCV 2011*, Barcelona, Spain, 2011, pp. 2540–2547. [Online]. Available: http://hal.inria.fr/hal-00856128

[80] E. Ferrante and 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*, vol. 8151. Nagoya, Japan: Springer, 2013, pp. 163–170. [Online]. Available: http://hal.inria.fr/hal-00855662

[81] N. Komodakis, "Towards more efficient and effective lp-based algorithms for mrf optimization," in *European Conference in Computer Vision (2)*, 2010, pp. 520–534.

[82] ——, "Learning to cluster using high order graphical models with latent variables," in *IEEE International Conference in Computer Vision*, 2011, pp. 73–80.

[83] ——, "Efficient training for pairwise or higher order crfs via dual decomposition," in *IEEE Computer Vision and Pattern Recognition*, 2011, pp. 1841–1848.