CZECH TECHNICAL UNIVERSITY IN PRAGUE

DOCTORAL THESIS STATEMENT

Czech Technical University in Prague
Faculty of Electrical Engineering
Department of Cybernetics

Mgr. Alexander Shekhovtsov

Exact and Partial Energy Minimization in Computer Vision

Ph.D. Programme: Electrical Engineering and Information Technology, P2612

Branch of study: Mathematical Engineering, 3901V021

Doctoral thesis statement for obtaining the academic title of “Doctor”, abbreviated to “Ph.D.”

Prague, February 2013

The doctoral thesis was produced in combined manner during Ph.D. study at the Center for Machine Perception of the Department of Cybernetics of the Faculty of Electrical Engineering of the CTU in Prague.

Candidate:	Mgr. Alexander Shekhovtsov Department of Cybernetics Faculty of Electrical Engineering of the CTU in Prague
Supervisor:	Prof. Václav Hlaváč Department of Cybernetics Faculty of Electrical Engineering of the CTU in Prague
Supervisor-Specialist:	Tomáš Werner, Ph.D. Department of Cybernetics Faculty of Electrical Engineering of the CTU in Prague
Opponents:	....

The doctoral thesis statement was distributed on ....

The defence of the doctoral thesis will be held on .... at .... a.m./p.m. before the Board for the Defence of the Doctoral Thesis in the branch of study Mathematical Engineering in the meeting room No. .... of the Faculty of Electrical Engineering of the CTU in Prague.

Those interested may get acquainted with the doctoral thesis concerned at the Dean Office of the Faculty of Electrical Engineering of the CTU in Prague, at the Department for Science and Research, Technická 2, 166 27 Prague 6.

....

Chairman of the Board for the Defence of the Doctoral Thesis
in the branch of study Mathematical Engineering
Department of Mathematics, Karlovo náměstí 13, 121 35 Prague 2

Contents

Keywords: combinatorial optimization; energy minimization; maximum a posteriori inference in Markov random fields; partial optimality; persistency; distributed min-cut/maxflow.

Note: Full text of the thesis is available at `http://cmp.felk.cvut.cz/~shekhovt/publications/as-phd-thesis.pdf`

1 Problem Formulation

In this work, we study several techniques for minimization of partially separable functions of discrete variables. This problem is commonly known as *energy minimization*, max-sum labeling problem, or weighted constraint satisfaction problem. Mathematically, it is defined as follows. Given a graph $(\mathcal{V}, \mathcal{E})$ and functions $f_s\colon \mathcal{L}_s \to \mathbb{R}$ for all $s \in \mathcal{V}$ and $f_{st}\colon \mathcal{L}_s \times \mathcal{L}_t \to \mathbb{R}$ for all $st \in \mathcal{E}$, where $\mathcal{L}_s$ are finite sets of *labels*, minimize the *energy*

$$E_f(x) = f_0 + \sum_{s\in\mathcal{V}} f_s(x_s) + \sum_{st\in\mathcal{E}} f_{st}(x_s, x_t)\,, \tag{1}$$

over all assignments $x \in \mathcal{L} = \prod_{s\in\mathcal{V}} \mathcal{L}_s$. The general energy minimization problem is NP-hard and even APX-hard, meaning that there is no polynomial-time approximation scheme.

Goal 1: Partial Optimality We study the linear programing (LP) relaxation approach, in which the problem is formulated as an integer linear programing (ILP) and the integrality constraints are relaxed. For some classes of problems, this relaxation is tight, and the problem can be solved to optimality. In addition, it is known that for some 0-1 ILP problems, their relaxations are *persistent*: whenever a part of the relaxed solution takes binary values, there exists an integer optimal solution taking the same values. This allows to determine optimally an assignment for the *part* of variables which are integer in the relaxation. For energy minimization, one can obtain an ILP reformulation possessing the *persistency* property. However, such an approach depends on the particular reduction applied. There were several methods proposed to find an optimal partial assignment directly for the energy in the form (1). The first goal of this work is to analyze and unify methods for optimal partial assignment.

Goal 2: Distributed MINCUT Many algorithms for energy minimization exploit solvable subproblems in the form of a minimum s-t cut problem (further on called MINCUT), which is formulated as follows. Given a directed graph (V, E), an edge capacity function $c\colon E \to \mathbb{R}_+$, and a source and sink vertices $s, t \in V$, find a partition $V = (S, V\backslash S)$ (a *cut*) that separates the source and the sink and minimizes the total cost of the cut (of the edges with tail in S and head in $V\backslash S$). Mathematically, this problem is written as

$$\min_{\substack{S\subset V\\ s\in S\\ t\notin S}} \sum_{\substack{(u,v)\in E\\ u\in S\\ v\notin S}} c(u,v)\,. \tag{2}$$

There are many algorithms known which solve this problem in polynomial time. Nevertheless, there is still an active ongoing research on algorithms that would be more efficient for computer vision problems. This includes specialized implementations, development of parallel and massively parallel (GPU) implementations. In some settings, *distributed* algorithms are an advantage. Such algorithms divide not only the computation but also the memory between the computational units. One application of such algorithms is to solve the problem by loading and processing only a portion of data at a time. Our second goal was to develop a distributed algorithm for the minimum s-t cut problem that would minimize the number of necessary loads of the problem data. In other words, that would be efficient for solving large problems on a single computer.

2 Contributions

Partial Optimality We develop an unified framework to analyze partial optimality methods. We show that several widely applied but previously unrelated partial optimality methods can be obtained from the same unifying sufficient conditions. These are the roof dual method (Boros et al, 2006), its multi-label extension (Kohli et al, 2008), the method of auxiliary submodular problems (Kovtun, 2004) and the family of local methods known as Dead End Elimination (DEE) (Desmet et al, 1992), originally developed in the context of protein structure design. We show that the different sufficient conditions proposed by these methods can be unified into a more general class. We study the common properties of this class, and show the following. All the above mentioned methods can be derived as local sufficient conditions in a specially constructed reparametrization of the problem. All these methods are connected to the standard LP relaxation. In particular, optimal fixation of the part of variables they provide are automatically satisfied by all solutions of LP relaxation. We also show that all fixed points of the expansion move algorithm (with the move step subproblems solved by roof-dual) are preserved by a subclass of the general method. The new framework suggests a way how to derive new partial optimality methods. We prove a characterization of the new unifying sufficient condition and propose a systematic way to derive partial optimality guarantees. For several subclasses of problems we derive methods to find the maximum partial optimality w.r.t. the unifying sufficient condition. This includes one new non-trivial subclass, for which the maximum partial optimality is found via solving a series of minimum cut problems. The overview of this contribution is given in §4.

The present work generalizes and extends the results of Shekhovtsov and

Hlaváč (2011), where unification of a smaller subset of methods was proposed. Optimality guarantees for the relaxed labellings (restricting the set of optimal relaxed labelings without solving the LP) were first proposed for multilabel QPBO method in the works Shekhovtsov et al (2008) and Kohli et al (2008).

Distributed MINCUT We develop a novel distributed algorithm for MINCUT. Noting that MINCUT is employed as a subroutine in many places in energy minimization (either allowing to solve the full problem or its relaxation or to find an improvement), a general algorithm suitable for solving large-scale sparse problems is required. The sequential version of the proposed algorithm allows to solve large instances of the MINCUT/MAXFLOW problem on a single computer using a disk storage. The parallel version of the algorithm allows to speed-up computations using several processors or to solve the problem in parallel on several computers exchanging messages over the network. We prove superior theoretical properties of both proposed algorithms, develop efficient implementations, and show that they achieve a competitive performance on large-scale computer vision instances while greatly improving on the disk operations in the sequential case and message exchange operations in the parallel case. These results were published in the article (Shekhovtsov and Hlaváč, 2012). The overview of this contribution is given in §5.

3 State of the Art

One of the major advances in computer vision in the past few years is the development of efficient deterministic algorithms for minimization of a partially separable function of discrete variables, commonly known as energy minimization, max-sum labeling problem, or weighted constraint satisfaction. The optimization model of this general form has proved useful in nearly all areas. In computer vision, it arises in particular as the maximum a *posteriori* inference in Markov random fields and conditional random fields, which are used to model a variety of vision problems ranging from the dense stereo and image segmentation to the use of pictorial structures for object recognition.

Partial Optimality Methods for partial optimality recover a "part of the optimal labeling" even in the case when finding the complete optimal labeling is not tractable. Several fundamental results identifying optimal partial assignments are obtained from the properties of linear relaxations of some discrete problems. An optimal solution to the continuous relaxation of a mixed-integer 0-1 programming problem is defined to be *persistent* if the set of $[0,1]$ relaxed variables realizing binary values retains the same binary values in at least one integer optimum (Adams et al, 1998). A mixed-integer program is said to be *persistent* (or possess the *persistency* property) if *every* solution to its continuous relaxation is persistent. Nemhauser and Trotter (1975) proved that the vertex packing problem is persistent. This result was later generalized to optimization of quadratic pseudo-Boolean functions (equivalent to energy minimization with $L_s = \{0, 1\}$) by Hammer et al (1984). *Strong persistency* was also proved, stating that if a variable takes the same binary value in *all* optimal solutions to the relaxation, then it realizes that binary value in *all* optimal integer solutions. It was shown that several approaches to minimization of quadratic pseudo-Boolean functions lead to the same lower bound, called the *roof dual* (Hammer et al, 1984; Boros and Hammer, 2002). It is dual to the LP relaxation and therefore provides partial optimality via the persistence property. This dual problem can be converted to MAXFLOW on a specially constructed graph with a double number of vertices (Boros et al, 1991) and thus can be solved by efficient MAXFLOW algorithms. It was found to be a powerful method for quadratic pseudo-Boolean optimization and was also enhanced by *probing* (Boros et al, 2006; Rother et al, 2007). Kolmogorov and Rother (2007) and Rother et al (2007) proposed a review, an efficient implementations and further improvements. After them, Quadratic Pseudo-Boolean Optimization, abbreviated as QPBO(-P), refers to this particular efficient method (resp. with probing). Kolmogorov (2010) gives an alternative interpretation of this method via a submodular lower bound.

Several works considered generalization of persistency to higher-order pseudo-Boolean functions. Adams et al (1998) considered a hierarchy of continuous relaxations of 0-1 polynomial programming problems. Given an optimal relaxed solution, they derive sufficient conditions on the dual multipliers which ensure that the solution is persistent. This result generalizes the roof duality approach, coinciding with it in the case of quadratic polynomials in binary variables.

Kolmogorov (2010, 2012) showed that bisubmodular relaxations provide a natural generalization of the roof duality approach to higher-order terms and possess the persistency property. He also considered submodular relaxations which form a special case of bisubmodular and showed the following. The roof duality relaxation for quadratic pseudo-Boolean functions is a submodular relaxation, and it dominates all other bisubmodular relaxations. For non-quadratic pseudo-Boolean functions, bisubmodular relaxations can be tighter than submodular ones. Kahl and Strandmark (2011, 2012) proposed a polynomial time algorithm to find the tightest submodular relaxation and evaluated it on problems in computer vision.

The multi-label QPBO (abbreviated as MQPBO) method (Kohli et al, 2008) extends partial optimality properties of QPBO to multi-label problems via the reduction of the problem to binary (Boolean) variables.

The following methods use different sufficient conditions, not based on lower bounds. The family of local methods known as the dead end elimination (DEE), originally proposed by Desmet et al (1992). DEE methods were developed in the context of protein structure design and are not widely known in the machine learning and computer vision communities. They formulate simple sufficient conditions allowing to exclude a label in a given pixel based on its unary and adjacent pairwise terms.

Kovtun (2003, 2004, 2011) proposed to construct a submodular problem such that determining optimal partial assignment for it allows to determine optimal partial assignment for the original problem.

MINCUT MINCUT problems in computer vision can originate from the energy minimization framework in several ways. Submodular energy minimization problems completely reduce to MINCUT (Ishikawa, 2003; Schlesinger and Flach, 2006). When the energy minimization is intractable, MINCUT is employed in relaxation and local search methods. The linear relaxation of pairwise energy minimization with 0-1 variables reduces to MINCUT (Boros et al, 1991; Kolmogorov and Rother, 2007) as well as the relaxation of problems reformulated in 0-1 variables (Kohli et al, 2008). Expansion-move, swap-move (Boykov et al, 1999) and fusion-move (Lempitsky et al, 2010) algorithms formulate a local improvement step as a MINCUT problem.

Despite of the existence of a number of algorithms for MINCUT with polynomial running time bounds and good practical performance, there is an active ongoing research on such algorithm and their specific implementations. Among sequential algorithms, the augmenting path implementation developed by Boykov and Kolmogorov (2004), denoted BK, can be considered as a baseline for simple vision problems. Several authors proposed sequential implementations which achieve better practical performance. Goldberg et al (2011) proposed an algorithm, similar to BK, but having a strongly polynomial running time bound and better practical performance on both simple and hard classes of problems. Jamriška et al (2012) give a cache-efficient implementation of BK, specialized to grid-structured graphs and achieving a significant speed-up. Arora et al (2010) proposed a variant of preflow-push algorithm achieving better performance on simple vision problems. Verma and Batra (2012) proposed an extended experimental comparison of sequential solvers on a wider set of benchmark problems. Several authors proposed novel parallel implementations (Delong and Boykov, 2008; Liu and Sun, 2010; Jamriška et al, 2012) and massive parallel implementations on GPU (Vineet and Narayanan, 2008, 2010).

Distributed MINCUT There were several proposals how to parallelize the algorithm by Boykov and Kolmogorov (2004). Partially distributed implementation (Liu and Sun, 2010) augments paths within disjoint regions first and then merges regions hierarchically. In the end, it still requires finding augmenting paths in the whole problem. Therefore, it cannot be used to solve a large problem by distributing it over several computers or by using a limited memory and a disk storage. For the shared memory model Liu and Sun (2010) reported a near-linear speed-up with up to 4 CPUs for 2D and 3D segmentation problems.

Strandmark and Kahl (2010) obtained a distributed algorithm using a dual decomposition approach. The subproblems are MINCUT instances on the parts of the graph (regions) and the master problem is solved using the subgradient method. This approach requires solving MINCUT subproblems with real valued capacities and does not have a polynomial bound on the number of iterations. The integer algorithm proposed by Strandmark and Kahl (2010) is not guaranteed to terminate.

The push-relabel algorithm (Goldberg and Tarjan, 1988) performs many local atomic operations, which makes it a good choice for a parallel or distributed implementation. A distributed version (Goldberg, 1991) runs in $O(n^2)$ time using $O(n)$ processors and $O(n^2\sqrt{m})$ messages, where n is the number of vertices and m is the number of edges in the problem. However, for a good practical performance it is crucial to implement the gap relabel and the global relabel heuristics (Cherkassky and Goldberg, 1994). The global relabel heuristic can be parallelized (Anderson and Setubal, 1995), but it is difficult to distribute. Delong and Boykov (2008) proposed a coarser granulation of push-relabel operations, associating a subset of vertices (a region) to each processor. Push and relabel operations inside a region are decoupled from the rest of the graph. This allows to process several non-interacting regions in parallel or run in a limited memory, processing few regions at a time. The gap and relabel heuristics, restricted to the regions (Delong and Boykov, 2008) are powerful and distributed at the same time. Our work was largely motivated by Delong and Boykov (2008) and the remark that their approach might be extendible to augmenting path algorithms.

4 Unified Partial Optimality

We show that the different sufficient conditions used in the methods of DEE, (M)QPBO, and auxiliary submodular problems can be unified into a general class. We study the common properties of the unified class and show that all the above methods can be derived as local sufficient conditions in a specially constructed reparametrization of the problem. Furthermore, it is guaranteed that the found strong partial optimalities retain all optimal solutions of the LP relaxation, in other words, solutions of the LP relaxation automatically satisfy the derived partial optimalities. Therefore, LP relaxation cannot be tightened by these methods. A similar result holds for fixed points of the fusion move algorithm: they satisfy strong optimalities of a specific subclass.

LP relaxation Let f denote the vector with components f_0, $f_s(i)$ for $i \in \mathcal{L}_s$ and $s \in \mathcal{V}$; $f_{st}(i,j)$ for $i \in \mathcal{L}_s,\ j \in \mathcal{L}_t$ and $st \in \mathcal{E}$. Clearly, $f \in \mathbb{R}^{\mathcal{I}}$ with the appropriately defined set $\mathcal{I}$.

For a labeling $x \in \mathcal{L}$ let $\delta(x) \in \mathbb{R}^{\mathcal{I}}$ be the vector with components $\delta(x)_0 = 1$, $\delta(x)_s(i) = [\![x_s{=}i]\!]$ and $\delta(x)_{st}(i,j) = [\![x_{st}{=}ij]\!]$, where $[\![]\!]$ is the Iverson bracket. Let $\langle\cdot,\cdot\rangle$ denote the scalar product on $\mathbb{R}^{\mathcal{I}}$. The energy minimization problem can be equivalently written as

$$\min_{x\in\mathcal{L}} \langle f, \delta(x)\rangle \tag{3}$$

or, still equivalently, as

$$\min_{\mu\in\mathcal{M}} \langle f, \mu\rangle \tag{4}$$

where $\mathcal{M} = \mathrm{conv}(\delta(\mathcal{L}))$, the convex hull of $\delta(\mathcal{L})$. The LP relaxation that we consider introduces a polytope Λ (called the *local polytope*) such that $\mathrm{aff}(\Lambda) = \mathrm{aff}(\mathcal{M})$, $\Lambda \cap \{0,1\}^{\mathcal{I}} = \mathcal{M} \cap \{0,1\}^{\mathcal{I}}$ and Λ has a polynomial number of facets. The problem (4) can be written equivalently as the ILP

$$\min_{\mu\in\Lambda\cap\{0,1\}^{\mathcal{I}}} \langle f, \mu\rangle\ . \tag{5}$$

By dropping the integrality constraints in (4), we obtain the LP *relaxation*

$$\min_{\mu\in\Lambda} \langle f, \mu\rangle\ . \tag{6}$$

In the case of two labels, all solutions to this relaxation are persistent: whenever *some* components of an optimal μ are integer in the relaxed problem (6), there exist an optimal solution to the ILP problem (5) taking those same integer values. However, in the case of multi-label problems, no such persistency holds.

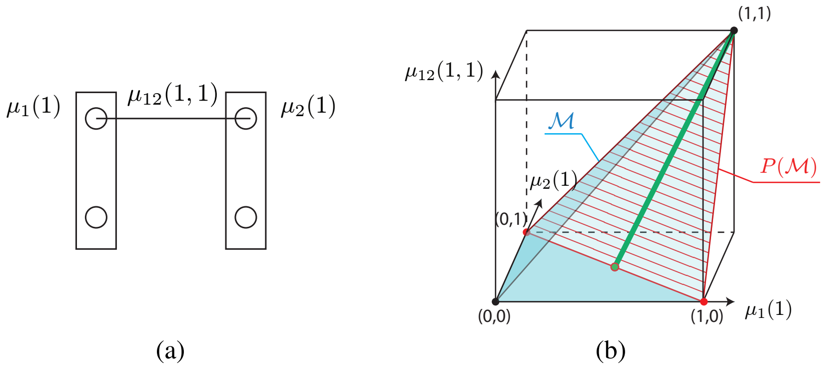

Figure 1: Illustration of the marginal polytope and a projection. (a) An energy minimization problem with two labels. The marginal vector μ is an element of $\mathbb{R}^9$, however, there are only 3 linearly independent coordinates, *e.g.*, the ones shown in the figure. (b) Polytope $\mathcal{M} = \Lambda$ (bluish) in the space of $\mu_1(1)$, $\mu_2(1)$ and $\mu_{12}(1,1)$. Possible labelings are marked as bold dots in this space. Suppose labelings $(1,0)$ and $(0,1)$ are optimal. Using Λ-improving projection P, the search space shrinks to $P(\mathcal{M})$ (the red hashed facet). Using Λ-improving projection P', the search space shrinks to the bold green line, not containing any optimal integer solutions but containing their convex combination.

Projections We observed that the optimality guarantees of the DEE, QPBO, MQPBO, and auxiliary submodular problem methods can be obtained through the following mechanism. An *improving* mapping $p\colon \mathcal{L} \to \mathcal{L}$ of labelings is constructed. Given an arbitrary labeling x, the mapping p provides a labeling that has a better (or at least not worse) energy: $E_f(p(x)) \leq E_f(x)$. In the case of Desmet's DEE, such a mapping changes the label in a single pixel s, replacing label α with label β in s. In the case of QPBO and MQPBO, the mapping is of the form $x \mapsto (x \vee x^{\min}) \wedge x^{\max}$, where $\vee$ and $\wedge$ denote component-wise maximum and minimum, respectively. In the case of auxiliary submodular problems, after a certain reordering of the labels, the mapping is of the form $x \mapsto x \vee x^{\min}$. As soon as mapping p is improving, we know for sure that there exists an optimal labeling in $p(\mathcal{L})$. Indeed, let x be an optimal labeling, then $p(x)$ is an optimal labeling in $p(\mathcal{L})$. The considered mappings are such that $p(\mathcal{L})$ is expressed as pixel-wise domain constraints, eliminating part of the labels in each pixel. For example, in the case of improving mapping $x \mapsto x \vee x^{\min}$, we know there exists a minimizer x such that $x_s \geq x_s^{\min}$ for all

$s \in \mathcal{V}$.

Even if a mapping $p\colon \mathcal{L} \to \mathcal{L}$ is given to us, verification of the improving property is generally NP-hard. So how do the methods we described find such improving mappings? They use certain *sufficient conditions*, which are *different* for each method. We propose the following mechanism to unify these approaches. We extend mappings of labelings to the local polytope Λ as follows. The *linear extension* of mapping $p\colon \mathcal{L} \to \mathcal{L}$ is the linear mapping $P\colon \Lambda \to \Lambda$ satisfying

$$(\forall x \in \mathcal{L}) \quad P\delta(x) = \delta(p(x)) \,. \tag{7}$$

The linear extension merely represents $p(x)$ in the polytope $\mathcal{M}$, where labeling x is represented by vector $\delta(x)$. The extension to Λ is then trivial since Λ resides in the affine space of $\mathcal{M}$. The extension P of an idempotent mapping $p\colon \mathcal{L} \to \mathcal{L}$ is an idempotent linear map: $PP = P$, or a *projection*. We consider such projections that polytope $\mathcal{M}$ is closed under them: $P(\mathcal{M}) \subset \mathcal{M}$.

We define that projection P is Λ-*improving* if

$$(\forall \mu \in \Lambda) \quad \langle f, P\mu \rangle \leq \langle f, P\mu \rangle \,. \tag{8}$$

For a Λ-improving projection P, the problems (6) and (4) can be reduced respectively as

$$\min_{\mu \in \Lambda} \langle f, \mu \rangle = \min_{\mu \in P(\Lambda)} \langle f, \mu \rangle \tag{9}$$

and

$$\min_{\mu \in \mathcal{M}} \langle f, \mu \rangle = \min_{\mu \in P(\mathcal{M})} \langle f, \mu \rangle \,. \tag{10}$$

In other words, the search for the optimal solution, without loss of generality, can be restricted to $P(\Lambda)$ for the LP problem (6) and to $P(\mathcal{M})$ for the energy minimization problem (4). Figure 1 and the following example illustrate improving projections.

Example. *Consider the energy minimization problem with 2 labels and 2 vertices, schematically shown in Figure 1(a). In this case* $\Lambda = \mathcal{M}$*. Let us denote* $\bar{\mu} = (\mu_1(1), \mu_2(1), \mu_{12}(1,1), 1)$*, which is a minimal representation of the vector* $\mu \in \Lambda$*. Let* P *be the following linear mapping of* $\bar{\mu}$*:*

$$P = \begin{pmatrix} 0.5 & -0.5 & 0.5 & 0.5 \\ -0.5 & 0.5 & 0.5 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} . \tag{11}$$

This mapping projects (non-orthogonally) to a facet of the marginal polytope (the red facet in Figure 1(b)). Assume optimal integer solutions are the two labelings $(0, 1)$ *and* $(1, 0)$*. In that case, the projection* P *is* Λ*-improving. Another* *projection*

$$P' = \begin{pmatrix} 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{12}$$

maps $\mathcal{M}$ to a single line (bold green in Figure 1(b)). It is also Λ-improving.

In the thesis we show the following:

- Verification of the Λ-improving property can be solved in polynomial time.
- Sufficient conditions used by DEE, (M)QPBO, and auxiliary submodular problem methods are a special case of Λ-improving projections.
- Λ-improving property has as simple characterization, connected to the dual of (6).

Maximum Projections We study the question when the set of Λ-improving projections with the product operation on them form a semilattice. If it is the case, the maximum element in the semilattice correspond to the projection which has the largest null space. This projection *maximally* reduces the polytopes Λ and $\mathcal{M}$ to $P(\Lambda)$ in (9) and $P(\mathcal{M})$ in (10), respectively. For several such cases we derive polynomial time algorithms for finding maximum projections.

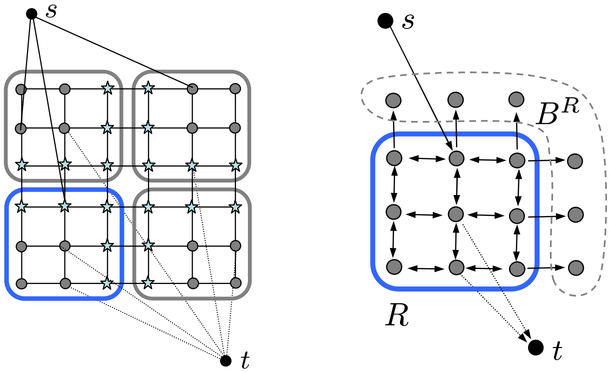

Figure 2: (a) Partition of a network into 4 regions and the boundary set $\mathcal{B}$ depicted by stars. (b) The region network corresponding to the highlighted region in (a).

5 Distributed MINCUT

Problem Partition We take the following approach to split the MINCUT problem into (coupled) subproblems. We consider a fixed collection of regions $(R^k)_{k=1}^K$ forming a partition (disjoint union) of $V \backslash \{s,t\}$. This partition splits the full network on the problem to the subnetworks, as illustrated in Figure 2.

We revisit the algorithm of Delong and Boykov (2008) for the case of a fixed partition into regions. We study a sequential variant and a novel parallel variant of their algorithm, which allows running computations concurrently on neighboring interacting regions using a conflict resolution similar to the asynchronous parallel push-relabel (Goldberg, 1991). Both algorithms perform several sequential iterations (*sweeps*) of discharging all regions.

We prove that both variants have $\Theta(n^2)$ bound on the number of sweeps (meaning there is $O(n^2)$ bound and it cannot be tightened).

Augmenting Path Region Discharge Our new algorithm combines path augmentation and push-relabel approaches. Given a fixed partition into regions, we introduce a distance function, which counts the number of region boundaries crossed by a path to the sink. Intuitively, it corresponds to the amount of costly operations – network communications or loads-unloads of the regions in the streaming mode. Let us denote $\mathcal{B}$ is the set of all *boundary vertices* (incident to inter-region edges), see Figure 2. The *region distance* $d^{*\mathcal{B}}(u)$ in G is the minimal number of inter-region edges contained in a path from u to t, or $|\mathcal{B}|$ if no such path exists, see Figure 3.

The algorithm, called Augmenting Path Region Discharge (ARD), maintains a labeling, which is a lower bound on the distance function. Within a region, we first augment paths to the sink and then paths to the boundary vertices prioritized by the lowest label. Thus the flow is pushed out of the region in the direction given by the distance estimate.

We show that our new algorithm terminates in $O(|\mathcal{B}|^2)$ sweeps. It implies a bound of $O((n')^2 m')$ messages in total, where $n' = |\mathcal{B}|$ and $m' = (\mathcal{B} \times \mathcal{B}) \cap E$. This means that the total number of sent messages depends *only* on the size of the separator set of the partition and not on the total size of the problem. This is a strict improvement over push-relabel algorithms, in particular the algorithm of Delong and Boykov (2008), which as we showed requires $\Omega(n^2)$ sweeps, *i.e.*, it uses $\Omega(n^2 m')$ messages.

Implementation We describe additional heuristics and an efficient implementation of both push-relabel and augmented-path based distributed algorithms. The following heuristics are proposed for ARD:

- Boundary relabel heuristic. Provides an improved estimate of the distance labeling by analyzing only the boundary part of the graph, not looking inside the regions.
- Partial discharge heuristic. Postpones path augmentations to higher boundary vertices to further sweeps. This allows to save a lot of unnecessary work, especially when used in combination with boundary relabeling.
- Boundary search trees. Efficient data structure allowing to for quicker search of the augmenting paths to the boundary vertices at the required height.

Experiments We conducted 3 series of experiments in which we compared our new algorithm against the state-of-the-art sequential and parallel solvers:

- Synthetic experiments, in which we observe general dependencies of the algorithms, with some statistical significance, *i.e.* not being biased to a particular problem instance. It also serves as an empirical validation, as thousands of instances are solved.
- Sequential competition. We study sequential versions of the algorithms, running them on real vision instances (University of Western Ontario web

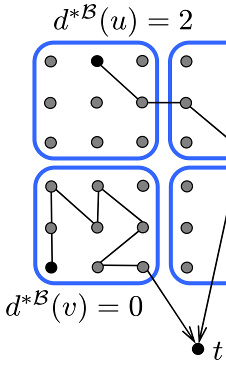

Figure 3: Illustration of the region distance.

pages, 2008). Only a single core of the CPU is utilized. We fix the region partition and study how much disk I/O it would take to solve each problem when only one region can be loaded in the memory at a time.

- Parallel competition. Parallel algorithms are tested on the instances that can fully fit in 2GB of memory. We compare our algorithms with two other state-of-the-art distributed implementations.

Next, we studied the dependency of computation time and number of sweeps on the number of regions in the partition. The results shown in Figure 4 demonstrate that the computation time required to solve test problems is stable over a large range of partitions and the number of sweeps required does not grow rapidly.

We performed additional tests of parallel ARD to determine scalability with number of processors, these results are shown in Figure 5.

We conducted an additional experiment to determine whether a part of optimal solution can be found efficiently by looking into individual regions. It turns out, that for simple problems (*e.g.*, stereo) a large part of the optimal solution can be recovered from the individual parts, while for more complicated problems (3D segmentation) this is no longer the case. The interactions between regions are really necessary.

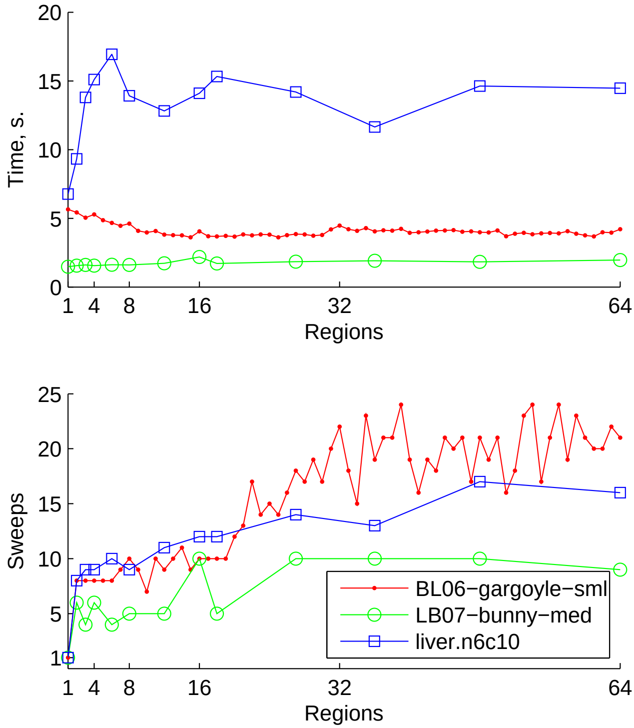

Figure 4: Dependence on the number of regions for the representative instances of multiview, stereo and segmentation. Top: CPU time used. Bottom: number of sweeps.

6 Conclusion

Unified Partial Optimality We presented a new sufficient condition for deriving partial optimal assignment in a multi-label energy minimization. Our framework allows for unified analysis of the methods previously proposed in the literature (Kovtun, 2003, 2004; Boros et al, 2006; Kohli et al, 2008; Desmet et al, 1992). The proposed sufficient condition include the conditions used by the methods in the literature as special cases. At the same time, it is a poly-

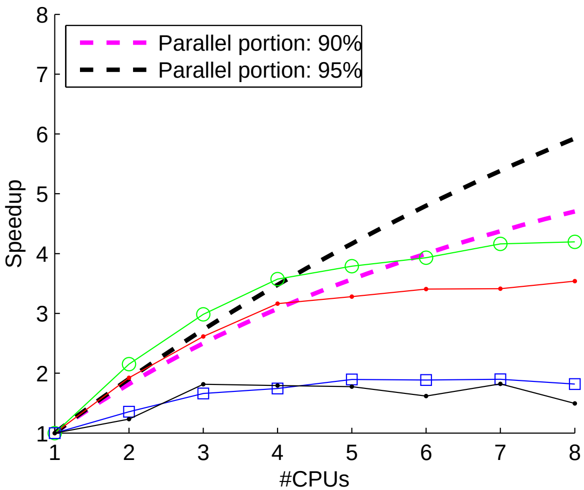

problem	regions	MEM	1CPU	8CPUs
BL06-gargoyle-lrg	16	1.6GB	137.4s	38.82s
BL06-camel-lrg	16	1.8GB	67.14s	16.00s
LB07-bunny-lrg	64	12GB	12.94s	7.11s
liver.n6c100	64	0.9GB	18.41s	12.21s

Figure 5: Speedup of P-ARD with the number of CPUs used. The extended legend shows the time to solve each problem with 1 and 8CPUs (does not include initialization). Dashed lines correspond to the speedup in the ideal case (Amdahl's law) when the parallel portion of the computation is 90% and 95%.

nomially verifiable one. The verification is expressed via a special LP relaxation. A main unifying property for partial optimality methods in this form is that they preserve all solutions of the LP relaxation, *i.e.*, LP relaxation cannot be tightened by these methods. We proved the existence of an equivalent transformation of the problem such that the sufficient condition is expressed in local inequalities. This allows to derive a simplified LP for verification of the condition. We also showed that for a subclass of partial optimality methods (including methods of Kovtun (2003)) that the fixed points of the fusion move algorithm satisfy the derived partial optimality guarantees. We studied subclasses, in which the maximum improving projections can be found efficiently. The new algorithm finds in polynomial time the maximum improving projection of the type "eliminate z by switching to y" for multilabel problems.

Interestingly, DEE and auxiliary submodular problems are not connected with lower bounds. They are nevertheless unified now with (M)QPBO methods, which are derived via lower bounds. In methods of Kovtun (2003), to obtain partial optimalities several sufficient conditions have to be satisfied simultaneously. QPBO method and approaches with submodular lower bounds (Kolmogorov, 2012; Kahl and Strandmark, 2011) maximize a lower bound and obtain optimality guarantees as a by-product. Our approach is more direct: we formalize the maximality of optimal partial solutions in a given class and derive methods achieving it.

We believe that a large part of the approach (including the characterization) is extendible to higher-order models and higher-order relaxations. However, there is still a number of open questions arising from our approach in the pairwise case. We did not fully characterize strictly improving projections and consequently some of the theorems are lacking the strict counterpart. We proposed that the projection may have non-integer weights, but considered only integer projections (except for showing DEE). The set of improving projections for a given function is represented by inequalities which are linear in the projection itself (except for the idempotency, which can be omitted). Hence, in principle, we can optimize some criterion over this convex set. In practice, we are lacking a low-dimensional convex parametric family of projections that would be flexible enough. An ideal criterion would be to maximize the dimensionality of the null space of the projection (equivalent to rank minimization), however even a simpler approximate criterion can provide exact guarantees to the initial energy minimization problem.

Distributed MINCUT We developed a new algorithm for MINCUT problem on sparse graphs, which combines augmenting paths and push-relabel approaches. We proved the worst case complexity guarantee of $O(|\mathcal{B}|^2)$ sweeps for the sequential and parallel variants of the algorithm (S/P-ARD). There are many algorithms in the literature with complexities in terms of elementary arithmetic operations better than we can prove. Nevertheless, we showed that our algorithms are fast and competitive in practice, even in the shared memory model.

We proposed an improved algorithm for the local problem reduction and determined that most of our test instances are difficult enough in the sense that very few vertices can be decided optimally by looking at individual regions. The result that S/P-ARD solves test problems in few tens of sweeps is thus non-trivial. We also gave a novel parallel version of the region push-relabel algorithm of Delong and Boykov (2008). We provided a number of auxiliary results to relate our approach to the state-of-the-art.

Both in theory and practice (randomized test), S-ARD has a better asymptote in the number of sweeps than the push-relabel variant. Experiments on real instances showed that when run on a single CPU and the whole problem fits into the memory S-ARD is comparable in speed with the non-distributed MAXFLOW implementation by Boykov and Kolmogorov (2004), and is even significantly faster in some cases. When only a single region is loaded into memory at a time, S-ARD uses much fewer disk I/O than the algorithm of Delong and Boykov (2008). We also demonstrated that the running time and the number of sweeps are very stable with respect to the partition of the problem into up to 64 regions. In the parallel mode, using 4 CPUs, P-ARD achieves a relative speedup of about $1.5 - 2.5$ times over S-ARD and uses just slightly larger number of sweeps. P-ARD compares favorably to other parallel algorithms, being a robust method suitable for a use in a distributed system.

References

Adams WP, Lassiter JB, Sherali HD (1998) Persistency in 0-1 polynomial programming. Mathematics of Operations Research 23(2):359–389

Anderson R, Setubal JC (1995) A parallel implementation of the push-relabel algorithm for the maximum flow problem. Journal of Parallel and Distributed Computing 29(1):17–26

Arora C, Banerjee S, Kalra P, Maheshwari SN (2010) An efficient graph cut algorithm for computer vision problems. In: European Conference on Computer Vision, Springer-Verlag, Berlin, Heidelberg, pp 552–565

Boros E, Hammer P (2002) Pseudo-Boolean optimization. Discrete Applied Mathematics 1-3(123):155–225

Boros E, Hammer PL, Sun X (1991) Network flows and minimization of quadratic pseudo-Boolean functions. Tech. Rep. RRR 17-1991, RUTCOR

Boros E, Hammer PL, Tavares G (2006) Preprocessing of unconstrained quadratic binary optimization. Tech. Rep. RRR 10-2006, RUTCOR

Boykov Y, Kolmogorov V (2004) An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. In: IEEE Transactions on Pattern Analysis and Machine Intelligence, vol 26, pp 1124–1137

Boykov Y, Veksler O, Zabih R (1999) Fast approximate energy minimization via graph cuts. In: International Conference on Computer Vision, vol 1, pp 377–384

Cherkassky BV, Goldberg AV (1994) On implementing push-relabel method for the maximum flow problem. Tech. rep.

Delong A, Boykov Y (2008) A scalable graph-cut algorithm for N-D grids. In: Computer Vision and Pattern Recognition Conference, pp 1–8

Desmet J, Maeyer MD, Hazes B, Lasters I (1992) The dead-end elimination theorem and its use in protein side-chain positioning. Nature 356:539–542

Goldberg AV (1991) Processor-efficient implementation of a maximum flow algorithm. Information Processing Letters 38(4):179–185

Goldberg AV, Tarjan RE (1988) A new approach to the maximum flow problem. Journal of the ACM 35

Goldberg AV, Hed S, Kaplan H, Tarjan RE, Werneck RF (2011) Maximum flows by incremental breadth-first search. In: European conference on Algorithms, Springer-Verlag, Berlin, Heidelberg, pp 457–468

Hammer P, Hansen P, Simeone B (1984) Roof duality, complementation and persistency in quadratic 0-1 optimization. Mathematical Programming pp 121–155

Ishikawa H (2003) Exact optimization for Markov random fields with convex priors. IEEE Transactions on Pattern Analysis and Machine Intelligence 25(10):1333–1336

Jamriška O, Sýkora D, Hornung A (2012) Cache-efficient graph cuts on structured grids. In: Computer Vision and Pattern Recognition Conference, pp 3673–3680

Kahl F, Strandmark P (2011) Generalized roof duality for pseudo-Boolean optimization. In: International Conference on Computer Vision, pp 255–262

Kahl F, Strandmark P (2012) Generalized roof duality. Discrete Applied Mathematics 160(16-17):2419–2434

Kohli P, Shekhovtsov A, Rother C, Kolmogorov V, Torr P (2008) On partial optimality in multi-label MRFs. In: McCallum A, Roweis S (eds) International Conference on Machine Learning, Omnipress, pp 480–487

Kolmogorov V (2010) Generalized roof duality and bisubmodular functions. In: Advances in Neural Information Processing Systems, pp 1144–1152

Kolmogorov V (2012) Generalized roof duality and bisubmodular functions. Discrete Applied Mathematics 160(4-5):416–426

Kolmogorov V, Rother C (2007) Minimizing non-submodular functions with graph cuts – a review. IEEE Transactions on Pattern Analysis and Machine Intelligence

Kovtun I (2003) Partial optimal labeling search for a NP-hard subclass of (max, +) problems. In: DAGM-Symposium, pp 402–409

Kovtun I (2004) Image segmentation based on sufficient conditions of optimality in NP-complete classes of structural labelling problem. PhD thesis, IRTC ITS National Academy of Sciences, Ukraine, in Ukrainian

Kovtun I (2011) Sufficient condition for partial optimality for (max, +) labeling problems and its usage. Control Systems and Computers (2), special issue

Lempitsky V, Rother C, Roth S, Blake A (2010) Fusion moves for Markov random field optimization. IEEE Transactions on Pattern Analysis and Machine Intelligence 32:1392–1405

Li SZ (2009) Markov Random Field Modeling in Image Analysis, 3rd edn. Springer Publishing Company, Incorporated

Liu J, Sun J (2010) Parallel graph-cuts by adaptive bottom-up merging. In: Computer Vision and Pattern Recognition Conference, pp 2181–2188

Nemhauser G, Trotter L Jr (1975) Vertex packings: Structural properties and algorithms. Mathematical Programming 8:232–248

Rother C, Kolmogorov V, Lempitsky V, Szummer M (2007) Optimizing binary MRFs via extended roof duality. In: Computer Vision and Pattern Recognition Conference

Schlesinger D, Flach B (2006) Transforming an arbitrary minsum problem into a binary one. Tech. Rep. TUD-FI06-01, Dresden University of Technology

Shekhovtsov A, Hlaváč V (2011) On partial opimality by auxiliary submodular problems. Control Systems and Computers (2), special issue

Shekhovtsov A, Hlaváč V (2012) A distributed mincut/maxflow algorithm combining path augmentation and push-relabel. International Journal of Computer Vision p 28

Shekhovtsov A, Kolmogorov V, Kohli P, Hlavac V, Rother C, Torr P (2008) LP-relaxation of binarized energy minimization. Tech. Rep. CTU–CMP–2007–27, Czech Technical University

Strandmark P, Kahl F (2010) Parallel and distributed graph cuts by dual decomposition. In: Computer Vision and Pattern Recognition Conference, pp 2085–2092

University of Western Ontario web pages (2008) Computer vision research group. max-flow problem instances in vision. `http://vision.csd.uwo.ca/maxflow-data/`

Verma T, Batra D (2012) Maxflow revisited: An empirical comparison of maxflow algorithms for dense vision problems. In: BRITISH MACHINE VISION CONFERENCE, BMVA Press, p 12

Vineet V, Narayanan P (2008) CUDA cuts: Fast graph cuts on the GPU. In: Computer Vision GPU workshop at ECCV, pp 1–8

Vineet V, Narayanan PJ (2010) Solving multilabel MRFs using incremental α-expansion on the GPUs. In: Asian conference on Computer Vision, Springer-Verlag, Berlin, Heidelberg, vol 3, pp 633–643

Author's Publications

Publications related to the thesis

Impacted journal articles

Shekhovtsov A, Hlaváč V (2012) A distributed mincut/maxflow algorithm combining path augmentation and push-relabel. International Journal of Computer Vision p 28, authorship: 90-10

Shekhovtsov A, Kovtun I, Hlaváč V (2008) Efficient MRF deformation model for non-rigid image matching. Computer Vision and Image Understanding 112:91–99, authorship: 30-30-30

Publications excerpted by WOS

Kohli P, Shekhovtsov A, Rother C, Kolmogorov V, Torr P (2008) On partial optimality in multi-label MRFs. Omnipress, pp 480–487, authorship: 20-20-20-20-20

Shekhovtsov A, Kovtun I, Hlaváč V (2007) Efficient MRF deformation model for non-rigid image matching, authorship: 33-33-33

Shekhovtsov A, Kohli P, Rother C (2012) Curvature prior for MRF-based segmentation and shape inpainting. In: Axel Pinz HBFL Thomas Pock (ed) DAGM/OAGM, pp 41–51, authorship: 30-30-30

Werner T, Shekhovtsov A (2007) Unified framework for semiring-based arc consistency and relaxation labeling. In: Grabner M, Grabner H (eds) CVWW 2007: Proceedings of the 12th Computer Vision Winter Workshop, Verlag der Technischen Universität Graz, Technikerstr. 4, A-8010, Graz, Austria, pp 27–34, authorship: 60-40-30

Other publications

Shekhovtsov A, Hlaváč V (2008) A lower bound by one-against-all decomposition for Potts model energy minimization. In: Computer Vision Winter Workshop, authorship: 30-30-30

Shekhovtsov A, Kovtun I, Hlaváč V (2006) Efficient MRF deformation model for image matching. Tech. Rep. CTU–CMP–2006–08, Center for Machine Perception, Czech Technical University, authorship: 30-30-30

Shekhovtsov A, Kolmogorov V, Kohli P, Hlavac V, Rother C, Torr P (2008) LP-relaxation of binarized energy minimization. Tech. Rep. CTU–CMP–2007–27, Czech Technical University, authorship: 16-16-16-16-16-16

Shekhovtsov A, Kohli P, Rother C (2011) Curvature prior for MRF-based segmentation and shape inpainting. Tech. Rep. CTU–CMP–2011–11, Center for Machine Perception, K13133 FEE Czech Technical University, Prague, Czech Republic, authorship: 30-30-30

Additional publications

Publications excerpted by WOS

Jancošek M, Shekhovtsov A, Pajdla T (2009) Scalable multi-view stereo. In: Computer Vision Workshops at International Conference on Computer Vision, pp 1526–1533, authorship: 60-10-30

Schlesinger D, Flach B, Shekhovtsov E (2004) A higher order MRF model for stereo reconstruction. In: DAGM, pp 440–446, authorship: 45-45-10

Shekhovtsov A, Hlaváč V (2010) Joint image GMM and shading MAP estimation, authorship: 90-10

Shekhovtsov A, Garcia-Arteaga JD, Werner T (2008) A discrete search method for multi-modal non-rigid image registration. In: Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment at International Conference on, p 6, authorship: 30-30-30

Other publications

Shekhovtsov A, Hlaváč V (2010) Joint image GMM and shading MAP estimation. Tech. Rep. K333–35/10, CTU–CMP–2010–03, Department of Cybernetics, Faculty of Electrical Engineering Czech Technical University, Prague, Czech Republic, authorship: 90-10

Citations of Author's Work

Citations of Author's Work in WoS

[1] K. Alahari, P. Kohli, and P. H. S. Torr. Dynamic Hybrid Algorithms for MAP Inference in Discrete MRFs. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 32(10):1846–1857, OCT 2010.

[2] S. R. Arashloo and J. Kittler. Pose-Invariant Face Matching Using MRF Energy Minimization Framework. In , *ENERGY MINIMIZATION METHODS IN COMPUTER VISION AND PATTERN RECOGNITION, PROCEEDINGS*, volume 5681 of *Lecture Notes in Computer Science*, pages 56–69, 2009.

[3] S. R. Arashloo and J. Kittler. Energy Normalization for Pose-Invariant Face Recognition Based on MRF Model Image Matching. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 33(6):1274–1280, JUN 2011.

[4] S. R. Arashloo, J. Kittler, and W. J. Christmas. Pose-invariant face recognition by matching on multi-resolution MRFs linked by supercoupling transform. *COMPUTER VISION AND IMAGE UNDERSTANDING*, 115(7, SI):1073–1083, JUL 2011.

[5] E. S. Brown, T. F. Chan, and X. Bresson. Completely Convex Formulation of the Chan-Vese Image Segmentation Model. *INTERNATIONAL JOURNAL OF COMPUTER VISION*, 98(1): 103–121, MAY 2012.

[6] T. Brox and J. Malik. Large Displacement Optical Flow: Descriptor Matching in Variational Motion Estimation. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 33(3):500–513, MAR 2011.

[7] D. Chrysostomou, A. Gasteratos, L. Nalpantidis, and G. C. Sirakoulis. Multi-view 3D scene reconstruction using ant colony optimization techniques. *MEASUREMENT SCIENCE & TECHNOLOGY*, 23(11), NOV 2012.

[8] L. Cordero-Grande, G. Vegas-Sanchez-Ferrero, and C. Casaseca-de-la Higuera, Pabloan Alberola-Lopez. Topology-Preserving Registration: A Solution via Graph Cuts. In , *COMBINATORIAL IMAGE ANALYSIS*, volume 6636 of *Lecture Notes in Computer Science*, pages 420–431, 2011.

[9] C. Domokos, J. Nemeth, and Z. Kato. Nonlinear Shape Registration without Correspondences. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 34(5): 943–958, MAY 2012.

[10] R. Donner, G. Langs, B. Micusik, and H. Bischof. Generalized sparse MRF appearance models. *IMAGE AND VISION COMPUTING*, 28(6):1031–1038, JUN 2010.

[11] O. Duchenne, A. Joulin, and J. Ponce. A Graph-Matching Kernel for Object Categorization. In *2011 IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION (ICCV)*, pages 1792–1799, 2011.

[12] A. El-Baz and G. Gimel'farb. Global image registration based on learning the prior appearance model. In *2008 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, VOLS 1-12*, PROCEEDINGS - IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, pages 3132–3138, 2008.

[13] H. Fathi and I. Brilakis. Automated sparse 3D point cloud generation of infrastructure using its distinctive visual features. *ADVANCED ENGINEERING INFORMATICS*, 25(4):760–770, OCT 2011.

[14] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski. Towards Internet-scale Multi-view Stereo. In *2010 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 1434–1441, 2010.

[15] B. Glocker, N. Komodakis, N. Paragios, and N. Navab. Approximated Curvature Penalty in Non-rigid Registration Using Pairwise MRFs. In , *ADVANCES IN VISUAL COMPUTING, PT 1, PROCEEDINGS*, volume 5875 of *Lecture Notes in Computer Science*, pages 1101–1109, 2009.

[16] R. Gupta, S. Sarawagi, and A. A. Diwan. Collective Inference for Extraction MRFs Coupled with Symmetric Clique Potentials. *JOURNAL OF MACHINE LEARNING RESEARCH*, 11: 3097–3135, NOV 2010.

[17] G. Hamarneh. Multi-label MRF Optimization via a Least Squares s-t Cut. In , *ADVANCES IN VISUAL COMPUTING, PT 1, PROCEEDINGS*, volume 5875 of *Lecture Notes in Computer Science*, pages 1055–1066, 2009.

[18] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Matin, F. V. Gleeson, S. M. Brady, and J. A. Schnabel. MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration. *MEDICAL IMAGE ANALYSIS*, 16(7, SI):1423–1435, OCT 2012.

[19] X. Hu and P. Mordohai. A Quantitative Evaluation of Confidence Measures for Stereo Vision. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 34(11): 2121–2133, NOV 2012.

[20] O. Jamriska, D. Sykora, and A. Hornung. Cache-efficient Graph Cuts on Structured Grids. In *2012 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 3673–3680, 2012.

[21] H. Jiang, S. X. Yu, and D. R. Martin. Linear Scale and Rotation Invariant Matching. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 33(7):1339–1355, JUL 2011.

[22] R. Klowsky, A. Kuijper, and M. Goesele. Modulation transfer function of patch-based stereo systems. In *2012 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 1386–1393, 2012.

[23] M. Koch, A. G. Schwing, D. Comaniciu, and M. Pollefeys. FULLY AUTOMATIC SEGMENTATION OF WRIST BONES FOR ARTHRITIS PATIENTS. In *2011 8TH IEEE INTERNATIONAL SYMPOSIUM ON BIOMEDICAL IMAGING: FROM NANO TO MACRO*, IEEE International Symposium on Biomedical Imaging, pages 636–640, 2011.

[24] N. Komodakis. Towards More Efficient and Effective LP-Based Algorithms for MRF Optimization. In , *COMPUTER VISION-ECCV 2010, PT II*, volume 6312 of *Lecture Notes in Computer Science*, pages 520–534, 2010.

[25] D. Kwon, K. J. Lee, I. D. Yun, and S. U. Lee. Nonrigid Image Registration Using Dynamic Higher-Order MRF Model. In , *COMPUTER VISION - ECCV 2008, PT I, PROCEEDINGS*, volume 5302 of *LECTURE NOTES IN COMPUTER SCIENCE*, pages 373–386, 2008.

[26] D. Kwon, I. D. Yun, and S. U. Lee. Rolled Fingerprint Construction Using MRF-Based Nonrigid Image Registration. *IEEE TRANSACTIONS ON IMAGE PROCESSING*, 19(12):3255–3270, DEC 2010.

[27] K. J. Lee, D. Kwon, I. D. Yun, and S. U. Lee. Optical Flow Estimation with Adaptive Convolution Kernel Prior on Discrete Framework. In *2010 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 2504–2511, 2010.

[28] V. Lempitsky, S. Roth, and C. Rother. FusionFlow: Discrete-continuous optimization for optical flow estimation. In *2008 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, VOLS 1-12*, PROCEEDINGS - IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, pages 3184–3191, 2008.

[29] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. T. Freeman. SIFT Flow: Dense Correspondence across Different Scenes. In , *COMPUTER VISION - ECCV 2008, PT III, PROCEEDINGS*, volume 5304 of *Lecture Notes in Computer Science*, pages 28–42, 2008.

[30] C. Liu, J. Yuen, and A. Torralba. Nonparametric Scene Parsing: Label Transfer via Dense Scene Alignment. In *CVPR: 2009 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, VOLS 1-4*, IEEE Conference on Computer Vision and Pattern Recognition, pages 1972–1979, 2009.

[31] C. Liu, J. Yuen, and A. Torralba. SIFT Flow: Dense Correspondence across Scenes and Its Applications. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 33(5):978–994, MAY 2011.

[32] Z. T. Lu, Q. J. Feng, S. J. Zhou, L. F. He, and W. F. Chen. Medical image elastic registration based on discontinuity adaptive Markov random field model. *IMAGING SCIENCE JOURNAL*, 58(4):193–201, AUG 2010.

[33] D. Mahapatra and Y. Sun. Nonrigid Registration of Dynamic Renal MR Images Using a Saliency Based MRF Model. In , *MEDICAL IMAGE COMPUTING AND COMPUTER-ASSISTED INTERVENTION - MICCAI 2008, PT I, PROCEEDINGS*, volume 5241 of *Lecture Notes in Computer Science*, pages 771–779, 2008.

[34] D. Mahapatra and Y. Sun. An MRF framework for joint registration and segmentation of natural and perfusion images. In *2010 IEEE INTERNATIONAL CONFERENCE ON IMAGE PROCESSING*, pages 1709–1712, 2010.

[35] D. Mahapatra and Y. Sun. Joint Registration and Segmentation of Dynamic Cardiac Perfusion Images Using MRFs. In , *MEDICAL IMAGE COMPUTING AND COMPUTER-ASSISTED INTERVENTION - MICCAI 2010, PT I*, volume 6361 of *Lecture Notes in Computer Science*, pages 493–501, 2010.

[36] D. Mahapatra and Y. Sun. MRF-Based Intensity Invariant Elastic Registration of Cardiac Perfusion Images Using Saliency Information. *IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING*, 58(4):991–1000, APR 2011.

[37] D. Mahapatra and Y. Sun. Integrating Segmentation Information for Improved MRF-Based Elastic Image Registration. *IEEE TRANSACTIONS ON IMAGE PROCESSING*, 21(1):170–183, JAN 2012.

[38] H. Mirzaalian and G. Hamarneh. Vessel Scale-Selection using MRF Optimization. In *2010 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 3273–3279, 2010.

[39] M. C. Robini and I. E. Magnin. Optimization by Stochastic Continuation. *SIAM JOURNAL ON IMAGING SCIENCES*, 3(4):1096–1121, 2010.

[40] S. Roth, V. Lempitsky, and C. Rother. Discrete-Continuous Optimization for Optical Flow Estimation. In , *STATISTICAL AND GEOMETRICAL APPROACHES TO VISUAL MOTION ANALYSIS*, volume 5604 of *Lecture Notes in Computer Science*, pages 1–22, 2009.

[41] D. Schlesinger. General Search Algorithms for Energy Minimization Problems. In , *ENERGY MINIMIZATION METHODS IN COMPUTER VISION AND PATTERN RECOGNITION, PROCEEDINGS*, volume 5681 of *Lecture Notes in Computer Science*, pages 84–97, 2009.

[42] Z. Stone, T. Zickler, and T. Darrell. Toward Large-Scale Face Recognition Using Social Network Context. *PROCEEDINGS OF THE IEEE*, 98(8, SI):1408–1415, AUG 2010.

[43] D. Sun, E. B. Sudderth, and M. J. Black. Layered Segmentation and Optical Flow Estimation Over Time. In *2012 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 1768–1775, 2012.

[44] L. Tang, G. Hamarneh, and R. Abugharbieh. Reliability-Driven, Spatially-Adaptive Regularization for Deformable Registration. In , *BIOMEDICAL IMAGE REGISTRATION*, volume 6204 of *Lecture Notes in Computer Science*, pages 173–185, 2010.

[45] Y. Tian and S. G. Narasimhan. Globally Optimal Estimation of Nonrigid Image Distortion. *INTERNATIONAL JOURNAL OF COMPUTER VISION*, 98(3):279–302, JUL 2012.

[46] E. Tola, C. Strecha, and P. Fua. Efficient large-scale multi-view stereo for ultra high-resolution image sets. *MACHINE VISION AND APPLICATIONS*, 23(5):903–920, SEP 2012.

[47] H.-H. Vu, P. Labatut, J.-P. Pons, and R. Keriven. High Accuracy and Visibility-Consistent Dense Multiview Stereo. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 34(5):889–901, MAY 2012.

[48] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Intrinsic Dense 3D Surface Tracking. In *2011 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR)*, IEEE Conference on Computer Vision and Pattern Recognition, pages 1225–1232, 2011.

[49] Q. Zhang, X. Song, X. Shao, R. Shibasaki, and H. Zhao. Unsupervised skeleton extraction and motion capture from 3D deformable matching. *NEUROCOMPUTING*, 100(SI):170–182, JAN 16 2013.

[50] H. Zhou, Y. Yuan, Y. Zhang, and C. Shi. Non-rigid object tracking in complex scenes. *PATTERN RECOGNITION LETTERS*, 30(2):98–102, JAN 15 2009.

Resumé in Czech

Velkým úspěchem počítačového vidění v posledních letech je objev efektivních deterministických algoritmů na minimalizaci částečně separabilních funkcí diskrétních proměnných (což je známo také pod názvy 'minimalizace energie', max-plus značkování či problém splňování vážených podmínek). Optimalizční problém v této obecné formě se ukázal užitečným téměř ve všech oblastech. V počítačovém vidění na tento problém vede úloha max-aposteriorní inference v markovských náhodných polích a podmíněných náhodných polích, kterými se modeluje mnoho úloh od husté stereo-korespondence a segmentace obrazů až po detekci a rozpoznávání objektů v obrazech.

Příspěvek této práce lze rozdělit do dvou částí. První příspěvek je sjednocení metod pro výpočet části optimálního řešení. Je-li dána instance úlohy, tyto metody najdou hodnoty podmnožiny proměnných, které jsou částí všech (nebo alespoň některých) optimálních řešení. Ukazujeme, že několik široce užívaných avšak dříve nesouvisejících metod lze získat z jediných počátečních podmínek. Díky tomu lze vidět, že tyto metody mají některé vlastnosti společné, např. zachovávají řešení standardní LP-relaxace. Tyto nové sjednocující postačující podmínky v práci charakterizujeme a navrhujeme systematické metody na hledání části optimálního řešení vzhledem k těmto podmínkám. Konkrétně předkládáme novou netriviální podtřídu úloh, pro které nalezení části optimálního řešení vede na sekvenci úloh na minimální s-t řez v grafu.

Druhý příspěvek se týká úlohy minimálního s-t řezu v grafu. Mnoho metod na minimalizaci energie je založeno na redukci původního problému (příp. jeho restrikce či relaxované formy) na úlohu minimálního s-t řezu. Protože velikost výsledných instancí této úlohy je často obrovská (např. ve 3-D segmentaci), jsou třeba efektivní algoritmy na řešení velkých řídkých úloh minimálního s-t řezu. V práci navrhujeme nový distribuovaný algoritmus na tuto úlohu. Sekvenční verze tohoto algoritmu umožňuje řešit velké instance úlohy na jediném počítači s omezenou operační pamětí a externí (diskovou) pamětí. Paralelní verze algoritmu umožňuje urychlit výpočet s pomocí více procesorů či řešit problém paralelně na několika počítačích posílajících si zprávy po síti. Dokazujeme, že horní meze na počet diskových operací resp. poslaných zpráv pro náš algoritmus jsou lepší než pro známé algoritmy. V experimentech s naší efektivní implementaci algoritmu ukazujeme, že dosahuje srovnatelných výsledků jako známé algoritmy, ovšem při významně nižším počtu drahých diskových resp. síťových operací.