HAL
open science

Alternating Direction Graph Matching

Đ.Khuê Lê-Huu, Nikos Paragios

To cite this version:

Đ.Khuê Lê-Huu, Nikos Paragios. Alternating Direction Graph Matching. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Jul 2017, Honolulu, United States. hal-01580824v1

HAL Id: hal-01580824

https://inria.hal.science/hal-01580824v1

Submitted on 2 Sep 2017 (v1), last revised 23 Feb 2018 (v2)

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.

Alternating Direction Graph Matching

D. Khuê Lê-Huu Nikos Paragios
CentraleSupélec, Université Paris-Saclay, France
{khue.le, nikos.paragios}@centralesupelec.fr

Abstract

In this paper, we introduce a graph matching method that can account for constraints of arbitrary order, with arbitrary potential functions. Unlike previous decomposition approaches that rely on the graph structures, we introduce a decomposition of the matching constraints. Graph matching is then reformulated as a non-convex non-separable optimization problem that can be split into smaller and much-easier-to-solve subproblems, by means of the alternating direction method of multipliers. The proposed framework is modular, scalable, and can be instantiated into different variants. Two instantiations are studied exploring pairwise and higher-order constraints. Experimental results on widely adopted benchmarks involving synthetic and real examples demonstrate that the proposed solutions outperform existing pairwise graph matching methods, and competitive with the state of the art in higher-order settings.

1. Introduction

The task of finding correspondences between two sets of visual features has a wide range of applications in computer vision and pattern recognition. This problem can be effectively solved using graph matching [28], and as a consequence, these methods have been successfully applied to various vision tasks, such as stereo matching [14], object recognition and categorization [1, 12], shape matching [1], surface registration [31], *etc*.

The general idea of solving feature correspondences via graph matching is to associate each set of features an attributed graph, where node attributes describe local characteristics, while edge (or hyper-edge) attributes describe structural relationships. The matching task seeks to minimize an energy (objective) function composed of unary, pairwise, and potentially higher-order terms. These terms are called the *potentials* of the energy function. In pairwise settings, graph matching can be seen as a quadratic assignment problem (QAP) in general form, known as Lawler's QAP [19]. Since QAP is known to be NP-complete [5, 27], graph matching is also NP-complete [13] and only approximate solutions can be found in polynomial time.

Graph matching has been an active research topic in the computer vision field for the past decades. In the recent literature, [13] proposed a graduated assignment algorithm to iteratively solve a series of convex approximations to the matching problem. In [22], a spectral matching based on the rank-1 approximation of the *affinity matrix* (composed of the potentials) was introduced, which was later improved in [9] by incorporating affine constraints towards a tighter relaxation. In [23], an integer projected fixed point algorithm that solves a sequence of first-order Taylor approximations using Hungarian method [18] was proposed, while in [28] the dual of the matching problem was considered to obtain a lower-bound on the energy, via dual decomposition. In [6], a random walk variant was used to address graph matching while [32] factorized the affinity matrix into smaller matrices, allowing a convex-concave relaxation that can be solved in a path-following fashion. Their inspiration was the path-following approach [29] exploiting a more restricted graph matching formulation, known as Koopmans-Beckmann's QAP [17]. Lately, [7] proposed a max-pooling strategy within the graph matching framework that is very robust to outliers.

Recently, researchers have proposed higher-order graph matching models to better incorporate structural similarities and achieve more accurate results [11, 30]. For solving such high-order models, [30] viewed the matching problem as a probabilistic model that is solved using an iterative successive projection algorithm. The extension of pairwise methods to deal with higher-order potentials was also considered like for example in [11] through a tensor matching (extended from [22]), or in [31] through a third-order dual decomposition (originating from [28]), or in [21] through a high-order reweighted random walk matching (extension of [6]). Recently, [26] developed a block coordinate ascent algorithm for solving third-order graph matching. They lifted the third-order problem to a fourth-order one which, after a convexification step, is solved by a sequence of linear or quadratic assignment problems. Despite the impressive performance, this method has two limitations: (a) it cannot be applied to graph matching of arbitrary order other than third and fourth, and (b) it cannot deal with graph matching

where occlusion is allowed on both sides, nor with many-to-many matching.

In this paper, a novel class of algorithms is introduced for solving graph matching involving constraints with arbitrary order and arbitrary potentials. These algorithms rely on a decomposition framework using the alternating direction method of multipliers.

The remainder of this paper is organized as follows. Section 2 provides the mathematical foundations of our approach while in Section 3 the general decomposition strategy is proposed along with two instantiations of this framework to a pairwise and higher-order approach. Section 4 presents in-depth experimental validation and comparisons with competing methods. The last section concludes the paper and presents the perspectives.

2. Mathematical background and notation

Let us first provide the elementary notation as well as the basic mathematical foundations of our approach. In the first subsection we will give a brief review of tensor, which will help us to compactly formulate the graph matching problem, as will be shown in the subsequent subsection.

2.1. Tensor

A real-valued D^{th}-order tensor $\mathcal{F}$ is a multidimensional array belonging to $\mathbb{R}^{n_1 \times n_2 \times \cdots \times n_D}$ (where $n_1, n_2, \ldots, n_D$ are positive integers). We denote the elements of $\mathcal{F}$ by $\mathcal{F}_{i_1 i_2 \ldots i_D}$ where $1 \leq i_d \leq n_d$ for $d = 1, 2, \ldots, D$. Each dimension of a tensor is called a *mode*.

We call the *multilinear form* associated to a tensor $\mathcal{F}$ the function $F : \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \times \cdots \times \mathbb{R}^{n_D} \to \mathbb{R}$ defined by

$$F(\mathbf{x}_1, \ldots, \mathbf{x}_D) = \sum_{i_1=1}^{n_1} \cdots \sum_{i_D=1}^{n_D} \mathcal{F}_{i_1 i_2 \ldots i_D} x^1_{i_1} x^2_{i_2} \cdots x^D_{i_D} \tag{1}$$

where $\mathbf{x}_d = (x^d_1, x^d_2, \ldots, x^d_{n_d}) \in \mathbb{R}^{n_d}$ for $d = 1, 2, \ldots, D$.

A tensor can be multiplied by a vector at a specific mode. Let $\mathbf{v} = (v_1, v_2, \ldots, v_{n_d})$ be an n_d dimensional vector. The *mode-d product* of $\mathcal{F}$ and $\mathbf{v}$, denoted by $\mathcal{F} \otimes_d \mathbf{v}$, is a $(D-1)^{\text{th}}$-order tensor $\mathcal{G}$ of dimensions $n_1 \times \cdots \times n_{d-1} \times n_{d+1} \times \cdots \times n_D$ defined by

$$\mathcal{G}_{i_1 \ldots i_{d-1} i_{d+1} \ldots i_D} = \sum_{i_d=1}^{n_d} \mathcal{F}_{i_1 \ldots i_d \ldots i_D} v_{i_d}. \tag{2}$$

With this definition, it is straightforward to see that the multilinear form (1) can be re-written as

$$F(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_D) = \mathcal{F} \otimes_1 \mathbf{x}_1 \otimes_2 \mathbf{x}_2 \cdots \otimes_D \mathbf{x}_D. \tag{3}$$

Let us consider for convenience the notation $\bigotimes_{d=a}^{b}$ to denote a sequence of products from mode a to mode b:

$$\mathcal{F} \bigotimes_{d=a}^{b} \mathbf{x}_d = \mathcal{F} \otimes_a \mathbf{x}_a \otimes_{a+1} \mathbf{x}_{a+1} \cdots \otimes_b \mathbf{x}_b. \tag{4}$$

By convention, $\mathcal{F} \bigotimes_{d=a}^{b} \mathbf{x}_d = \mathcal{F}$ if $a > b$.

In this work, we are interested in tensors having the same dimension at every mode, *i.e.* $n_1 = n_2 = \ldots = n_D = n$. In the sequel, all tensors are supposed to have this property.

2.2. Graph and hypergraph matching

A matching configuration between two graphs $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{V}_2, \mathcal{E}_2)$ can be represented by a *assignment matrix* $\mathbf{X} \in \{0,1\}^{n_1 \times n_2}$ where $n_1 = |\mathcal{V}_1|, n_2 = |\mathcal{V}_2|$. An element $x_{(i_1,i_2)}$ of $\mathbf{X}$ equals 1 if the node $i_1 \in \mathcal{V}_1$ is matched to the node $i_2 \in \mathcal{V}_2$, and equals 0 otherwise.
Standard graph matching imposes the one-to-(at most)-one constraints, *i.e.* the sum of any row or any column of $\mathbf{X}$ must be ≤ 1. If the elements of $\mathbf{X}$ are binary, then $\mathbf{X}$ obeys the *hard matching constraints*. When $\mathbf{X}$ is relaxed to take real values in $[0,1]$, $\mathbf{X}$ obeys the *soft matching constraints*.

In this paper we use the following notations: $\mathrm{vec}(\mathbf{V})$ denotes the column-wise vectorized replica of a matrix $\mathbf{V}$; $\mathrm{mat}(\mathbf{v})$ is the $n_1 \times n_2$ reshaped matrix of an n-dimensional vector $\mathbf{v}$, where $n = n_1 n_2$; $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ the assignment matrix and $\mathbf{x} = \mathrm{vec}(\mathbf{X}) \in \mathbb{R}^n$ the *assignment vector*; $\mathcal{M}^*$ (respectively $\mathcal{M}$) is the set of $n_1 \times n_2$ matrices that obey the hard (respectively, the soft) matching constraints.

Energy function. Let $x_i = x_{(i_1,i_2)}$ be an element of $\mathbf{X}$ representing the matching of two nodes i_1 and i_2. Suppose that matching these nodes requires a potential $\mathcal{F}^1_i \in \mathbb{R}$. Similarly, let $\mathcal{F}^2_{ij}$ denote the potential for matching two edges (i_1, j_1), (i_2, j_2), and $\mathcal{F}^3_{ijk}$ for matching two (third-order) hyper-edges (i_1, j_1, k_1), (i_2, j_2, k_2), and so on. Graph matching can be expressed as minimizing

$$\sum_i \mathcal{F}^1_i x_i + \sum_{ij} \mathcal{F}^2_{ij} x_i x_j + \sum_{ijk} \mathcal{F}^3_{ijk} x_i x_j x_k + \cdots \tag{5}$$

The above function can be re-written more compactly using tensors. Indeed, let us consider for example the third-order potentials. Since $\mathcal{F}^3_{ijk}$ has three indices, $(\mathcal{F}^3_{ijk})_{1 \leq i,j,k \leq n}$ can be seen as a third-order tensor belonging to $\mathbb{R}^{n \times n \times n}$ and its multilinear form (*c.f.* (1)) is the function

$$F^3(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{k=1}^{n} \mathcal{F}^3_{ijk} x_i y_j z_k \tag{6}$$

defined for $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^n$. Clearly, the third-order terms in (5) can be re-written as $F^3(\mathbf{x}, \mathbf{x}, \mathbf{x})$. More generally, D^{th}-order potentials can be represented by a Dth-order tensor $\mathcal{F}^D$ and their corresponding terms in the objective function can be re-written as $F^D(\mathbf{x}, \mathbf{x}, \ldots, \mathbf{x})$, resulting in the following reformulation of graph matching.

Problem 1 (D^{th}-order graph matching) *Minimize*

$$F^1(\mathbf{x}) + F^2(\mathbf{x}, \mathbf{x}) + \cdots + F^D(\mathbf{x}, \mathbf{x}, \ldots, \mathbf{x}) \tag{7}$$

subject to $\mathbf{x} \in \mathcal{M}^*$*, where* F^d *(*$d = 1, \ldots, D$*) is the multilinear form of a tensor* $\mathcal{F}^d$ *representing the* d^{th}*-order potentials.*

In the next section, we propose a method to solve the continuous relaxation of this problem, *i.e.* minimizing (7) subject to $\mathbf{x} \in \mathcal{M}$ (soft matching) instead of $\mathbf{x} \in \mathcal{M}^*$ (hard matching). The returned continuous solution is discretized using the usual Hungarian method [18].

3. Alternating direction graph matching

3.1. Overview of ADMM

We briefly describe the (multi-block) alternating direction method of multipliers (ADMM) for solving the following optimization problem:

$$\begin{aligned} \text{Minimize} \quad & \phi(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p) \\ \text{subject to} \quad & \mathbf{A}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 + \cdots + \mathbf{A}_p\mathbf{x}_p = \mathbf{b}, \\ & \mathbf{x}_i \in \mathcal{X}_i \subseteq \mathbb{R}^{n_i} \quad \forall 1 \leq i \leq p, \end{aligned} \tag{8}$$

where $\mathcal{X}_i$ are closed convex sets, $\mathbf{A}_i \in \mathbb{R}^{m \times n_i}\ \forall i, \mathbf{b} \in \mathbb{R}^m$.

The augmented Lagrangian of the above problem is

$$\begin{aligned} L_\rho(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p, \mathbf{y}) &= \phi(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p) \\ &+ \mathbf{y}^\top \left(\sum_{i=1}^p \mathbf{A}_i\mathbf{x}_i - \mathbf{b} \right) + \frac{\rho}{2} \left\| \sum_{i=1}^p \mathbf{A}_i\mathbf{x}_i - \mathbf{b} \right\|_2^2, \end{aligned} \tag{9}$$

where $\mathbf{y}$ is called the *Lagrangian multiplier vector* and $\rho > 0$ is called the *penalty parameter*.

In the sequel, let $\mathbf{x}_{[a,b]}$ denote $(\mathbf{x}_a, \mathbf{x}_{a+1}, \ldots, \mathbf{x}_b)$ (by convention, if $a > b$ then $\mathbf{x}_{[a,b]}$ is ignored). Standard ADMM solves problem (8) by iterating:

1. For $i = 1, 2, \ldots, p$, update $\mathbf{x}_i$:

$$\mathbf{x}_i^{k+1} = \operatorname*{argmin}_{\mathbf{x} \in \mathcal{X}_i} L_\rho(\mathbf{x}_{[1,i-1]}^{k+1}, \mathbf{x}, \mathbf{x}_{[i+1,p]}^k, \mathbf{y}^k). \tag{10}$$

2. Update $\mathbf{y}$:

$$\mathbf{y}^{k+1} = \mathbf{y}^k + \rho \left(\sum_{i=1}^p \mathbf{A}_i\mathbf{x}_i^{k+1} - \mathbf{b}^{k+1} \right). \tag{11}$$

The algorithm converges if the following *residual* converges to 0 as $k \to \infty$:

$$r^k = \left\| \sum_{i=1}^p \mathbf{A}_i\mathbf{x}_i^k - \mathbf{b}^k \right\|_2^2 + \sum_{i=1}^p \left\| \mathbf{A}_i\mathbf{x}_i^k - \mathbf{A}_i\mathbf{x}_i^{k-1} \right\|_2^2. \tag{12}$$

We will discuss the convergence of ADMM in Section 3.4.

3.2. Graph matching decomposition framework

Decomposition is a general approach to solving a problem by breaking it up into smaller ones that can be efficiently addressed separately, and then reassembling the results towards a globally consistent solution of the original non-decomposed problem [2, 4, 10]. Clearly, the above ADMM is such a method because it decomposes the large problem (8) into smaller problems (10).

In computer vision, decomposition methods such as Dual Decomposition (DD) and ADMM have been applied to optimizing discrete Markov random fields (MRFs) [15, 16, 20, 25] and to solving graph matching [28]. The main idea is to decompose the original complex graph into simpler subgraphs and then reassembling the solutions on these subgraphs using different mechanisms. While in MRF inference, this concept has been proven to be flexible and powerful, that is far from being the case in graph matching, due to the hardness of the matching constraints. Indeed, to deal with these constraints, [28] for example adopted a strategy that creates subproblems that are also smaller graph matching problems, which are computationally highly challenging. Moreover, subgradient method has been used to impose consensus, which is known to have slow rate of convergence [2]. One can conclude that DD is a very slow method and works for a limited set of energy models often associated with small sizes and low to medium geometric connectivities [28].

In our framework, we do not rely on the structure of the graphs but instead, on the nature of the variables. In fact, the idea is to decompose the assignment vector $\mathbf{x}$ (by means of Lagrangian relaxation) into different variables where each variable obeys weaker constraints (that are easier to handle). For example, instead of dealing with the assignment vector $\mathbf{x} \in \mathcal{M}$, we can represent it by two vectors $\mathbf{x}_1$ and $\mathbf{x}_2$, where the sum of each row of $\mathrm{mat}(\mathbf{x}_1)$ is ≤ 1 and the sum of each column of $\mathrm{mat}(\mathbf{x}_2)$ is ≤ 1, and we constrain these two vectors to be equal. More generally, we can decompose $\mathbf{x}$ into as many vectors as we want, and in any manner, the only condition is that the set of constraints imposed on these vectors must be equivalent to $\mathbf{x}_1 = \mathbf{x}_2 = \cdots = \mathbf{x}_p \in \mathcal{M}$ where p is the number of vectors. As for the objective function (7), there is also an infinite number of ways to re-write it under the new variables $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p$. The only condition is that the re-written objective function must be equal to the original one when $\mathbf{x}_1 = \mathbf{x}_2 = \cdots = \mathbf{x}_p = \mathbf{x}$. For example, if $p = D$ then one can re-write (7) as

$$F^1(\mathbf{x}_1) + F^2(\mathbf{x}_1, \mathbf{x}_2) + \cdots + F^D(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_D). \tag{13}$$

Each combination of (a) such a variable decomposition and (b) such a way of re-writing the objective function will yield a different Lagrangian relaxation and thus, produce a different algorithm. Since there are virtually infinite of such combinations, the number of algorithms one can design from

them is also unlimited, not to mention the different choices of the reassembly mechanism, such as subgradient methods [2, 4], cutting plane methods [2], ADMM [3], or others. We call the class of algorithms that base on ADMM Alternating Direction Graph Matching (ADGM) algorithms. A major advantage of ADMM over the other mechanisms is that its subproblems involve only one block of variables, regardless of the form the objective function.

As an illustration of ADGM, we present below a particular example. Nevertheless, this example is still general enough to include an infinite number of special cases.

Problem 2 (Decomposed graph matching) *Minimize*

$$F^1(\mathbf{x}_1) + F^2(\mathbf{x}_1, \mathbf{x}_2) + \cdots + F^D(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_D) \quad (14)$$

subject to

$$\mathbf{A}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 + \cdots + \mathbf{A}_D\mathbf{x}_D = \mathbf{0}, \quad (15)$$

$$\mathbf{x}_d \in \mathcal{M}_d \quad \forall\, 1 \le d \le D, \quad (16)$$

where $(\mathbf{A}_d)_{1\le d\le D}$ *are* $m \times n$ *matrices, defined in such a way that* (15) *is equivalent to* $\mathbf{x}_1 = \mathbf{x}_2 = \cdots = \mathbf{x}_D$*, and* $(\mathcal{M}_d)_{1\le d\le D}$ *are closed convex subsets of* $\mathbb{R}^n$ *satisfying*

$$\mathcal{M}_1 \cap \mathcal{M}_2 \cap \cdots \cap \mathcal{M}_D = \mathcal{M}. \quad (17)$$

It is easily seen that the above problem is equivalent to the continuous relaxation of Problem 1. Clearly, this problem is a special case of the standard form (8). Thus, ADMM can be applied to it in a straightforward manner. The augmented Lagrangian of Problem 2 is

$$L_\rho(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_D, \mathbf{y}) = \sum_{d=1}^{D} F^d(\mathbf{x}_1, \ldots, \mathbf{x}_d) + \mathbf{y}^\top \left(\sum_{d=1}^{D} \mathbf{A}_d\mathbf{x}_d\right) + \frac{\rho}{2}\left\|\sum_{d=1}^{D} \mathbf{A}_d\mathbf{x}_d\right\|_2^2. \quad (18)$$

The $\mathbf{y}$ update step (11) and the computation of the residual (12) is trivial. Let us focus on the $\mathbf{x}$ update step (10):

$$\mathbf{x}_d^{k+1} = \operatorname*{argmin}_{\mathbf{x}\in\mathcal{M}_d} L_\rho(\mathbf{x}_{[1,d-1]}^{k+1}, \mathbf{x}, \mathbf{x}_{[d+1,D]}^{k}, \mathbf{y}^k). \quad (19)$$

Denote

$$\mathbf{s}_d^k = \sum_{i=1}^{d-1} \mathbf{A}_i\mathbf{x}_i^{k+1} + \sum_{j=d+1}^{D} \mathbf{A}_j\mathbf{x}_j^k, \quad (20)$$

$$\mathbf{p}_d^k = \sum_{i=d}^{D} \mathcal{F}^i \bigotimes_{j=1}^{d-1} \mathbf{x}_j^{k+1} \bigotimes_{l=d+1}^{i} \mathbf{x}_l^k. \quad \text{(see (4))} \quad (21)$$

It can be seen that (details given in the supplement)

$$\sum_{i=d}^{D} F^i(\mathbf{x}_{[1,d-1]}^{k+1}, \mathbf{x}, \mathbf{x}_{[d+1,i]}^{k}) = (\mathbf{p}_k)^\top\mathbf{x}. \quad (22)$$

Thus, let cst be a constant independent of $\mathbf{x}$, we have:

$$L_\rho(\mathbf{x}_{[1,d-1]}^{k+1}, \mathbf{x}, \mathbf{x}_{[d+1,D]}^{k}, \mathbf{y}^k) = (\mathbf{p}_d^k)^\top\mathbf{x} + (\mathbf{y}^k)^\top(\mathbf{A}_d\mathbf{x} + \mathbf{s}_d^k) + \frac{\rho}{2}\left\|\mathbf{A}_d\mathbf{x} + \mathbf{s}_d^k\right\|_2^2 + \mathrm{cst}, \quad (23)$$

and the subproblems (19) are reduced to minimizing the following quadratic functions over $\mathcal{M}_d$ ($d = 1, 2, \ldots, D$):

$$\frac{1}{2}\mathbf{x}^\top\mathbf{A}_d^\top\mathbf{A}_d\mathbf{x} + \left(\mathbf{A}_d^\top\mathbf{s}_d^k + \frac{1}{\rho}(\mathbf{A}_d^\top\mathbf{y}^k + \mathbf{p}_d^k)\right)^\top\mathbf{x}. \quad (24)$$

In summary, an ADGM algorithm has three main steps: 1) choose $(\mathbf{A}_d)_{1\le d\le D}$ and $(\mathcal{M}_d)_{1\le d\le D}$ satisfying the conditions stated in Problem 2, 2) update $\mathbf{x}_d^{k+1}$ by minimizing (24) over $\mathcal{M}_d$, and 3) update $\mathbf{y}^{k+1}$ using (11) (and repeat 2), 3) until convergence).

3.3. Two simple ADGM algorithms

Let us follow the above three steps with two examples.

Step 1: Choose $(\mathbf{A}_d)_{1\le d\le D}$ and $(\mathcal{M}_d)_{1\le d\le D}$. First, $(\mathcal{M}_d)_{1\le d\le D}$ take values in one of the following two sets:

$$\mathcal{M}_\mathrm{r} = \{\mathbf{x} : \text{ sum of each row of } \mathrm{mat}(\mathbf{x}) \text{ is } \le 1\}, \quad (25)$$

$$\mathcal{M}_\mathrm{c} = \{\mathbf{x} : \text{ sum of each column of } \mathrm{mat}(\mathbf{x}) \text{ is } \le 1\}, \quad (26)$$

such that *both* $\mathcal{M}_\mathrm{r}$ *and* $\mathcal{M}_\mathrm{c}$ *are taken at least once*. If no occlusion is allowed in $\mathcal{G}_1$ (respectively $\mathcal{G}_2$), then the term "≤ 1" is replaced by "$= 1$" for $\mathcal{M}_\mathrm{r}$ (respectively $\mathcal{M}_\mathrm{c}$). If many-to-many matching is allowed, then these inequality constraints are removed. In either case, $\mathcal{M}_\mathrm{r}$ and $\mathcal{M}_\mathrm{c}$ are closed and convex. Clearly, since $\mathcal{M}_\mathrm{r} \cap \mathcal{M}_\mathrm{c} = \mathcal{M}$, condition (17) is satisfied.

Second, to impose $\mathbf{x}_1 = \mathbf{x}_2 = \cdots = \mathbf{x}_D$, we can for example choose $(\mathbf{A}_d)_{1\le d\le D}$ such that

$$\mathbf{x}_1 = \mathbf{x}_2, \quad \mathbf{x}_1 = \mathbf{x}_3, \ldots, \quad \mathbf{x}_1 = \mathbf{x}_D \quad (27)$$

or alternatively

$$\mathbf{x}_1 = \mathbf{x}_2, \quad \mathbf{x}_2 = \mathbf{x}_3, \ldots, \quad \mathbf{x}_{D-1} = \mathbf{x}_D. \quad (28)$$

It is easily seen that the above two sets of constraints can be both expressed under the general form (15). Each choice leads to a different algorithm. Let ADGM1 denote the one obtained from (27) and ADGM2 obtained from (28).

Step 2: Update $\mathbf{x}_d^{k+1}$. Plugging (27) and (28) into (15), the subproblems (24) are reduced to (details given in the supplement)

$$\mathbf{x}_d^{k+1} = \operatorname*{argmin}_{\mathbf{x}\in\mathcal{M}_d} \left\{\frac{1}{2}\|\mathbf{x}\|_2^2 - \mathbf{c}_d^\top\mathbf{x}\right\}, \quad (29)$$

where $(\mathbf{c}_d)_{1\le d\le D}$ are defined as follows, for ADGM1:

$$\mathbf{c}_1 = \frac{1}{D-1}\left(\sum_{d=2}^{D}\mathbf{x}_d^k - \frac{1}{\rho}\sum_{d=2}^{D}\mathbf{y}_d^k - \frac{1}{\rho}\sum_{d=1}^{D}\mathcal{F}^d\bigotimes_{i=2}^{d}\mathbf{x}_i^k\right), \tag{30}$$

$$\mathbf{c}_d = \mathbf{x}_1^{k+1} + \frac{1}{\rho}\mathbf{y}_d^k - \frac{1}{\rho}\left(\sum_{i=d}^{D}\mathcal{F}^i\bigotimes_{i=1}^{d-1}\mathbf{x}_i^{k+1}\bigotimes_{j=d+1}^{i}\mathbf{x}_j^k\right) \tag{31}$$

for $2 \le d \le D$, and for ADGM2:

$$\mathbf{c}_1 = \mathbf{x}_2^k - \frac{1}{\rho}\mathbf{y}_2^k - \frac{1}{\rho}\sum_{d=1}^{D}\mathcal{F}^d\bigotimes_{i=2}^{d}\mathbf{x}_i^k, \tag{32}$$

$$\mathbf{c}_D = \mathbf{x}_{D-1}^{k+1} + \frac{1}{\rho}\mathbf{y}_D^k - \frac{1}{\rho}\mathcal{F}^D\bigotimes_{i=1}^{D-1}\mathbf{x}_i^{k+1}, \tag{33}$$

$$\mathbf{c}_d = \frac{1}{2}(\mathbf{x}_{d-1}^{k+1} + \mathbf{x}_{d+1}^k) + \frac{1}{2\rho}(\mathbf{y}_d^k - \mathbf{y}_{d+1}^k) - \frac{1}{2\rho}\sum_{i=d}^{D}\mathcal{F}^i\bigotimes_{i=1}^{d-1}\mathbf{x}_i^{k+1}\bigotimes_{j=d+1}^{i}\mathbf{x}_j^k \tag{34}$$

for $2 \le d \le D-1$.

Step 3: Update $\mathbf{y}^{k+1}$. From (27) and (28), it is seen that this step is reduced to $\mathbf{y}_d^{k+1} = \mathbf{y}_d^k + \rho(\mathbf{x}_1^{k+1} - \mathbf{x}_d^{k+1})$ for ADGM1 and $\mathbf{y}_d^{k+1} = \mathbf{y}_d^k + \rho(\mathbf{x}_{d-1}^{k+1} - \mathbf{x}_d^{k+1})$ for ADGM2.

Remark. When $D = 2$ the two algorithms are identical.

Note that (29) means $\mathbf{x}_d^{k+1}$ is the projection of $\mathbf{c}_d$ onto $\mathcal{M}_d$. Since $(\mathcal{M}_d)_{1\le d\le D}$ obey only row-wise or column-wise constraints, the projection becomes row-wise or column-wise and can be solved based on the projection onto a simplex [8]. We show how to do that and give sketches of the above algorithms in the supplement.

3.4. Convergent ADGM

Note that the objective function in Problem 2 is neither separable nor convex in general. Convergence of ADMM for this type of functions is unknown. Indeed, our ADGM algorithms do not always converge, especially for small values of the penalty parameter ρ. When ρ is large, they are likely to converge. However, we also notice that small ρ often (but not always) achieves better objective values. Motivated by these observations, we propose the following adaptive scheme that we find to work very well in practice: starting from a small initial value ρ_0, the algorithm runs for T_1 iterations to stabilize, after that, if no improvement of the residual r^k is made every T_2 iterations, then we increase ρ by a factor β and continue. The intuition behind this scheme is simple: we hope to reach a good objective value with a small ρ, but if this leads to slow (or no) convergence, then we increase ρ to put more penalty on the consensus of the variables and that would result in faster convergence. Using this scheme, we observe that our algorithms always converge in practice. In the experiments, we set $T_1 = 300, T_2 = 50, \beta = 2$ and $\rho_0 = \frac{n}{1000}$.

4. Experiments

We adopt the adaptive scheme in Section 3.4 to two ADGM algorithms presented in Section 3.3, and denote them respectively ADGM1 and ADGM2. In pairwise settings, however, since these two algorithms are identical, we denote them simply ADGM. We compare ADGM and ADGM1/ADGM2 to the following state of the art methods:

Pairwise: Spectral Matching with Affine Constraint (SMAC) [9], Integer Projected Fixed Point (IPFP) [23], Reweighted Random Walk Matching (RRWM) [6], Dual Decomposition with Branch and Bound (DD) [28] and Max-Pooling Matching (MPM) [7]. We should note that DD is only used in the experiments using the same energy models presented in [28]. For the other experiments, DD is excluded due to the prohibitive execution time. Also, as suggested in [23], we use the solution returned by Spectral Matching (SM) [22] as initialization for IPFP.

Higher-order: Probabilistic Graph Matching (PGM) [30], Tensor Matching (TM) [11], Reweighted Random Walk Hypergraph Matching (RRWHM) [21] and Block Coordinate Ascent Graph Matching (BCAGM) [26]. For BCAGM, we use MPM [7] as subroutine because it was shown in [26] (and again by our experiments) that this variant of BCAGM (denoted by "BCAGM+MP" in [26]) outperforms the other variants thanks to the effectiveness of MPM. Since there is no ambiguity, in the sequel we denote this variant "BCAGM" for short.

We should note that, while we formulated the graph matching as a *minimization* problem, most of the above listed methods are *maximization* solvers and many models/objective functions in previous work were designed to be maximized. For ease of comparison, ADGM is also converted to a maximization solver (by letting it minimize the additive inverse of the objective function), and the results reported in this section are for the maximization settings (*i.e.* higher objective values are better). In the experiments, we also use some pairwise minimization models (such as the one from [28]), which we convert to maximization problems as follows: after building the affinity matrix $\mathbf{M}$ from the (minimization) potentials, the new (maximization) affinity matrix is computed by $\max(\mathbf{M}) - \mathbf{M}$ where $\max(\mathbf{M})$ denotes the greatest element of $\mathbf{M}$. Note that one cannot simply take $-\mathbf{M}$ because some of the methods only work for non-negative potentials.

And last, due to space constraints, we leave the reported running time for each algorithm in the supplement (except for the very first experiment where this can be presented compactly). In short, ADGM is faster than SMAC [9]

	Methods	Error (%)	Global opt. (%)	Time (s)
House	MPM	42.32	0	0.02
	RRWM	90.51	0	0.01
	IPFP	87.30	0	0.02
	SMAC	81.11	0	0.18
	DD	**0**	**100**	14.20
	ADGM	**0**	**100**	0.03
Hotel	MPM	21.49	44.80	0.02
	RRWM	85.05	0	0.01
	IPFP	85.37	0	0.02
	SMAC	71.33	0	0.18
	DD	**0.19**	**100**	13.57
	ADGM	**0.19**	**100**	0.02

Table 1: Results on House and Hotel sequences using the **pairwise model A**, described in Section 4.1 and previously proposed in [28].

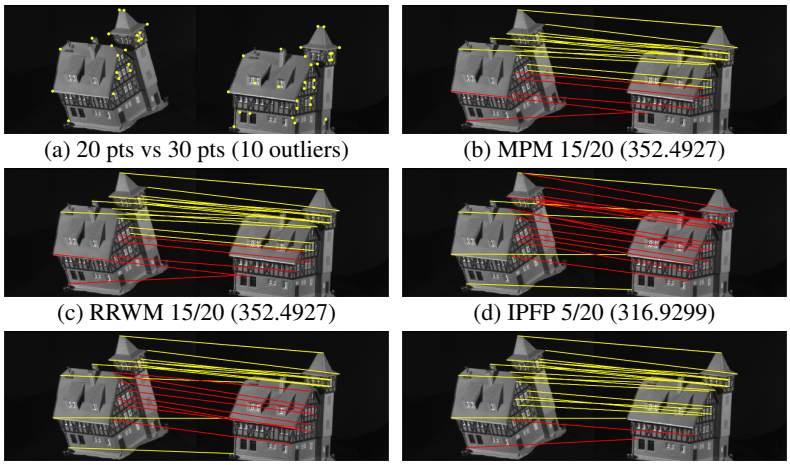

(a) 20 pts vs 30 pts (10 outliers) (b) MPM 15/20 (352.4927) (c) RRWM 15/20 (352.4927) (d) IPFP 5/20 (316.9299) (e) SMAC 12/20 (315.0426) (f) ADGM **18/20** (**353.3569**)

Figure 1: House matching using the **pairwise model B** described in Section 4.1. Ground-truth value is 343.1515. (Best viewed in color.)

(in pairwise settings) and ADGM1/ADGM2 are faster than TM [11] (in higher-order settings) while being slower than the other methods.

4.1. House and Hotel

The CMU House and Hotel sequences[1] have been widely used in previous work to evaluate graph matching algorithms. It consists of 111 frames of a synthetic house and 101 frames of a synthetic hotel. Each frame in these sequences is manually labeled with 30 feature points.

Pairwise model A. In this experiment we match all possible pairs of images in each sequence, with all 30 points (*i.e.* no outlier). A Delaunay triangulation is performed for these 30 points to obtain the graph edges. The unary terms are the distances between the Shape Context descriptors [1]. The pairwise terms when matching (i_1, j_1) to (i_2, j_2) are

$$\mathcal{F}^2_{ij} = \eta \exp\left(\delta^2/\sigma_l^2\right) + (1-\eta)\exp\left(\alpha^2/\sigma_a^2\right) - 1 \quad (35)$$

where η, σ_l, σ_a are some weight and scaling constants and δ, α are computed from $d_1 = \|\overrightarrow{i_1 j_1}\|$ and $d_2 = \|\overrightarrow{i_2 j_2}\|$ as

$$\delta = \frac{|d_1 - d_2|}{d_1 + d_2}, \quad \alpha = \arccos\left(\frac{\overrightarrow{i_1 j_1}}{d_1} \cdot \frac{\overrightarrow{i_2 j_2}}{d_2}\right). \quad (36)$$

This experiment is reproduced from [28] using their energy model files[2]. It should be noted that in [28], the above unary potentials are subtracted by a large number to prevent occlusion. We refer the reader to [28] for further details. For ease of comparison with the results reported in [28], here we also report the performance of each algorithm in terms of overall percentage of mismatches and frequency of reaching the global optimum. Results are given in Table 1. One can observe that DD and ADGM always reached the global optima, but ADGM did it hundreds times faster. Even the recent methods RRWM and MPM performed poorly on this model (only MPM produced acceptable results). Also, we notice a dramatic decrease in performance of SMAC and IPFP compared to the results reported in [28]. We should note that the above potentials, containing both positive and negative values, are defined for a *minimization* problem. It was unclear how those *maximization* solvers were used in [28]. For the reader to be able to reproduce the results, we make our software publicly available.

Pairwise model B. In this experiment, we match all possible pairs of the sequence with the baseline (*i.e.* the separation between frames, *e.g.* the baseline between frame 5 and frame 105 is 100) varying from 10 to 100 by intervals of 10. For each pair, we match 10, 20 and 30 points in the first image to 30 points in the second image. We set the unary terms to 0 and compute the pairwise terms as

$$\mathcal{F}^2_{ij} = \exp\left(-\left|\|\overrightarrow{i_1 j_1}\| - \|\overrightarrow{i_2 j_2}\|\right| / \sigma^2\right), \quad (37)$$

where $\sigma^2 = 2500$. It should be noted that the above pairwise terms are computed for every pair (i_1, j_1) and (i_2, j_2), *i.e.* the graphs are fully connected. This experiment has been performed on the House sequence in previous work, including [6] and [26]. Here we consider the Hotel sequence as well. We report the average objective ratio (which is the ratio of the obtained objective value over the ground-truth value) and the average accuracy for each algorithm in Figure 2. Due to space constraints, we only show the results for the harder cases where occlusion is allowed, and leave the other results in the supplement. As one can observe,

[1] http://vasc.ri.cmu.edu/idb/html/motion/index.html

[2] http://www.cs.dartmouth.edu/~lorenzo/Papers/tkr_pami13_data.zip.

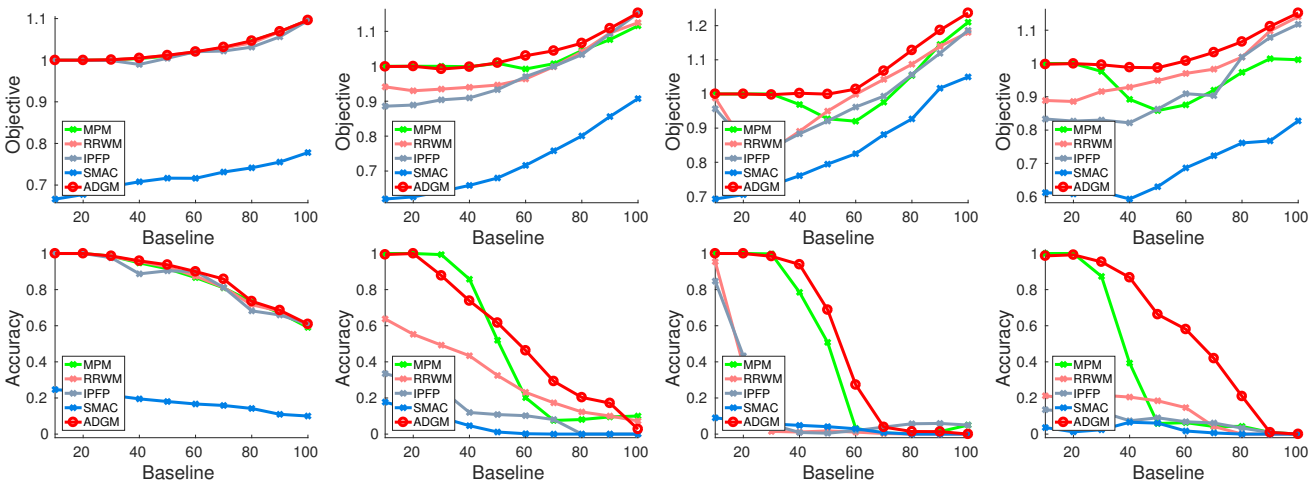

(a) House: 20 pts vs 30 pts (b) House: 10 pts vs 30 pts (c) Hotel: 20 pts vs 30 pts (d) Hotel: 10 pts vs 30 pts

Figure 2: Results on House and Hotel sequences using the **pairwise model B** described in Section 4.1.

ADGM produced the highest objective values in almost all the tests.

Third-order model. This experiment has the same settings as the previous one, but uses a third-order model proposed in [11]. We set the unary and pairwise terms to 0 and compute the potentials when matching two triples of points (i_1, j_1, k_1) and (i_2, j_2, k_2) as

$$\mathcal{F}^3_{ijk} = \exp\left(-\|f_{i_1j_1k_1} - f_{i_2j_2k_2}\|_2^2/\gamma\right), \qquad (38)$$

where f_{ijk} is a feature vector composed of the angles of the triangle (i, j, k), and γ is the mean of all squared distances. We report the results for House sequence in Figure 3 and provide the other results in the supplement. One can observe that ADGM1 and ADGM2 achieved quite similar performance, both were competitive with BCAGM [26] while outperformed all the other methods.

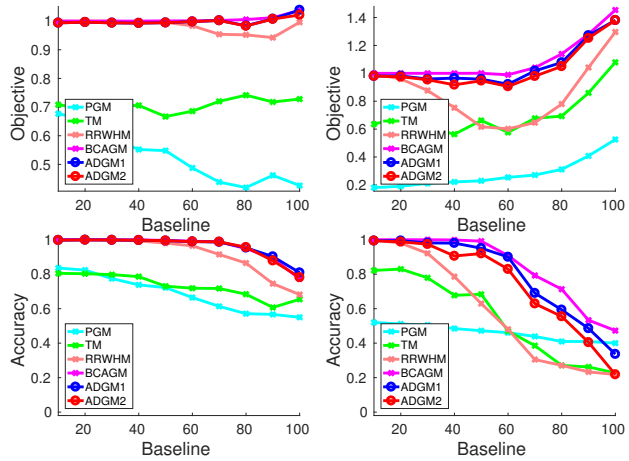

(a) 20 pts vs 30 pts (b) 10 pts vs 30 pts

Figure 3: Results on House sequence using the **third-order model** described in Section 4.1.

4.2. Cars and Motorbikes

The Cars and Motorbikes dataset was introduced in [24] and has been used in previous work for evaluating graph matching algorithms. It consists of 30 pairs of car images and 20 pairs of motorbike images with different shapes, view-points and scales. Each pair contains both inliers (chosen manually) and outliers (chosen randomly).

Pairwise model C. In this experiment, we keep all inliers in both images and randomly add outliers to the second image. The number of outliers varies from 0 to 40 by intervals of 5. We tried the **pairwise model B** described in Section 4.1 but obtained unsatisfactory matching results (showed in supplementary material). Inspired by the model in [28], we propose below a new model that is very simple yet very suited for real-world images. We set the unary terms to 0 and compute the pairwise terms as

$$\mathcal{F}^2_{ij} = \eta\delta + (1-\eta)\frac{1-\cos\alpha}{2}, \qquad (39)$$

where $\eta \in [0, 1]$ is a weight constant and δ, α are computed from $d_1 = \|\overrightarrow{i_1j_1}\|$ and $d_2 = \|\overrightarrow{i_2j_2}\|$ as

$$\delta = \frac{|d_1 - d_2|}{d_1 + d_2}, \quad \cos\alpha = \frac{\overrightarrow{i_1j_1}}{d_1} \cdot \frac{\overrightarrow{i_2j_2}}{d_2}. \qquad (40)$$

Intuitively, $\mathcal{F}^2_{ij}$ computes the geometric agreement between $\overrightarrow{i_1j_1}$ and $\overrightarrow{i_2j_2}$, in terms of both length and direction. The above potentials measure the *dissimilarity* between the edges, as thus the corresponding graph matching

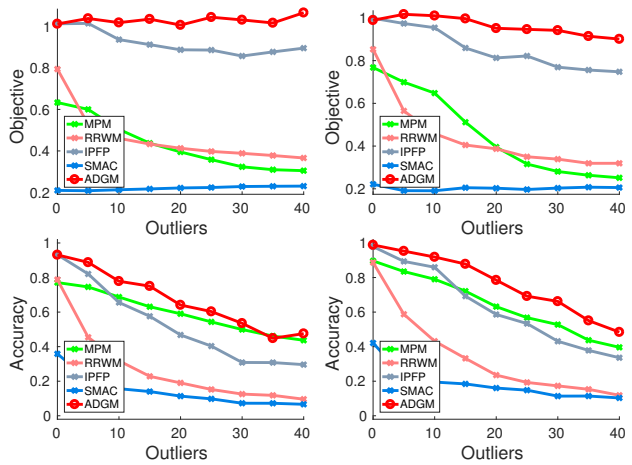

(a) Cars (b) Motorbikes

Figure 4: Results on Cars and Motorbikes using the **pairwise model C** described in Section 4.2.

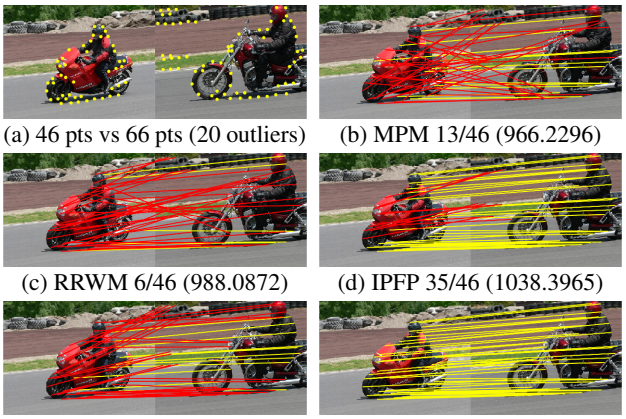

(a) 46 pts vs 66 pts (20 outliers) (b) MPM 13/46 (966.2296)

(c) RRWM 6/46 (988.0872) (d) IPFP 35/46 (1038.3965)

(e) SMAC 11/46 (1028.7961) (f) ADGM **46/46** (**1043.0687**)

Figure 5: Motorbike matching using the **pairwise model C** described in Section 4.2. (Best viewed in color.)

problem is a *minimization* one. Pairwise potentials based on both length and angle were previously proposed in [24, 28] and [32]. However, ours are the simplest to compute. In this experiment, $\eta = 0.5$.

We match every image pair and report the average in terms of objective value and matching accuracy for each method in Figure 4. One can observe that ADGM completely outperformed all the other methods.

Third-order model. This experiment has the same settings as the previous one, except that it uses a third-order model (the same as in House and Hotel experiment) and the number of outliers varies from 0 to 16 (by intervals of 2). Results are reported in Figure 6 and a matching example is given in Figure 7. ADGM did quite well on this dataset. On Cars, both ADGM1 and ADGM2 achieved better objective values than BCAGM in $7/9$ cases. On Motorbikes, ADGM1 beat BCAGM in $5/9$ cases and had equivalent performance in $1/9$ cases; ADGM2 beat BCAGM in $8/9$ cases.

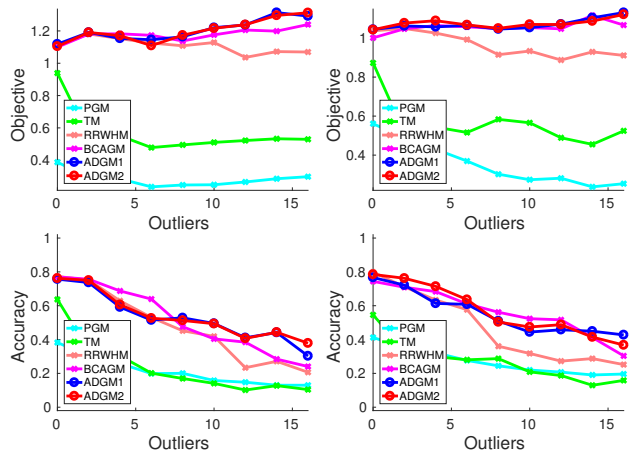

(a) Cars (b) Motorbikes

Figure 6: Results on Cars and Motorbikes using the **third-order model** described in Section 4.2.

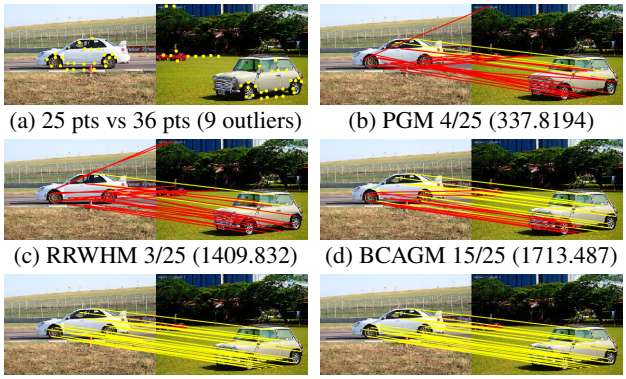

(a) 25 pts vs 36 pts (9 outliers) (b) PGM 4/25 (337.8194)

(c) RRWHM 3/25 (1409.832) (d) BCAGM 15/25 (1713.487)

(e) ADGM1 **25/25** (**2161.5354**) (f) ADGM2 **25/25** (**2161.5354**)

Figure 7: Car matching using the **third-order model** described in Section 4.2. (Best viewed in color.)

5. Conclusion and future work

We have presented ADGM, a novel class of algorithms for solving graph matching. Two examples of ADGM were implemented and evaluated. The results demonstrate that they outperform existing pairwise methods and competitive with the state of the art higher-order methods. In future work, we plan to adopt a more principled adaptive scheme to the penalty parameter, and to study the performance of different variants of ADGM. A software implementation of our algorithms are available for download on our website.

Acknowledgements. We thank the anonymous reviewers for their insightful comments and suggestions.

References

[1] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE transactions on pattern analysis and machine intelligence*, 24(4):509–522, 2002.

[2] D. P. Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.

[4] S. Boyd, L. Xiao, A. Mutapcic, and J. Mattingley. Notes on decomposition methods. *Notes for EE364B, Stanford University*, pages 1–36, 2007.

[5] R. E. Burkard, E. Cela, P. M. Pardalos, and L. S. Pitsoulis. The quadratic assignment problem. In *Handbook of combinatorial optimization*, pages 1713–1809. Springer, 1998.

[6] M. Cho, J. Lee, and K. M. Lee. Reweighted random walks for graph matching. In *Computer Vision–ECCV 2010*, pages 492–505. Springer, 2010.

[7] M. Cho, J. Sun, O. Duchenne, and J. Ponce. Finding matches in a haystack: A max-pooling strategy for graph matching in the presence of outliers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2083–2090, 2014.

[8] L. Condat. Fast projection onto the simplex and the ℓ_1 ball. *Mathematical Programming*, pages 1–11, 2014.

[9] T. Cour, P. Srinivasan, and J. Shi. Balanced graph matching. *Advances in Neural Information Processing Systems*, 19:313, 2007.

[10] G. B. Dantzig and P. Wolfe. Decomposition principle for linear programs. *Operations research*, 8(1):101–111, 1960.

[11] O. Duchenne, F. Bach, I.-S. Kweon, and J. Ponce. A tensor-based algorithm for high-order graph matching. *IEEE transactions on pattern analysis and machine intelligence*, 33(12):2383–2395, 2011.

[12] O. Duchenne, A. Joulin, and J. Ponce. A graph-matching kernel for object categorization. In *2011 International Conference on Computer Vision*, pages 1792–1799. IEEE, 2011.

[13] S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 18(4):377–388, 1996.

[14] V. Kolmogorov and R. Zabih. Computing visual correspondence with occlusions using graph cuts. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 508–515. IEEE, 2001.

[15] N. Komodakis and N. Paragios. Beyond pairwise energies: Efficient optimization for higher-order mrfs. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2985–2992. IEEE, 2009.

[16] N. Komodakis, N. Paragios, and G. Tziritas. Mrf energy minimization and beyond via dual decomposition. *IEEE transactions on pattern analysis and machine intelligence*, 33(3):531–552, 2011.

[17] T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica: journal of the Econometric Society*, pages 53–76, 1957.

[18] H. W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

[19] E. L. Lawler. The quadratic assignment problem. *Management science*, 9(4):586–599, 1963.

[20] D. K. Lê-Huu. Dual decomposition with accelerated first-order scheme for discrete markov random field optimization. Technical report, CentraleSupélec, Université Paris-Saclay, 2014.

[21] J. Lee, M. Cho, and K. M. Lee. Hyper-graph matching via reweighted random walks. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1633–1640. IEEE, 2011.

[22] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 2, pages 1482–1489. IEEE, 2005.

[23] M. Leordeanu, M. Hebert, and R. Sukthankar. An integer projected fixed point method for graph matching and map inference. In *Advances in neural information processing systems*, pages 1114–1122, 2009.

[24] M. Leordeanu, R. Sukthankar, and M. Hebert. Unsupervised learning for graph matching. *International journal of computer vision*, 96(1):28–45, 2012.

[25] A. F. Martins, M. A. Figueiredo, P. M. Aguiar, N. A. Smith, and E. P. Xing. Ad 3: Alternating directions dual decomposition for map inference in graphical models. *The Journal of Machine Learning Research*, 16(1):495–545, 2015.

[26] Q. Nguyen, A. Gautier, and M. Hein. A flexible tensor block coordinate ascent scheme for hypergraph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5270–5278, 2015.

[27] S. Sahni and T. Gonzalez. P-complete approximation problems. *Journal of the ACM (JACM)*, 23(3):555–565, 1976.

[28] L. Torresani, V. Kolmogorov, and C. Rother. A dual decomposition approach to feature correspondence. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(2):259–271, 2013.

[29] M. Zaslavskiy, F. Bach, and J.-P. Vert. A path following algorithm for the graph matching problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2227–2242, 2009.

[30] R. Zass and A. Shashua. Probabilistic graph and hypergraph matching. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[31] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 382–389. IEEE, 2010.

[32] F. Zhou and F. De la Torre. Factorized graph matching. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 127–134. IEEE, 2012.